



HAL
open science

Décrypter la société des plateformes : Organisations, marchés et réseaux dans l'économie numérique

Paola Tubaro

► **To cite this version:**

Paola Tubaro. Décrypter la société des plateformes : Organisations, marchés et réseaux dans l'économie numérique. Sociology. Institut d'Etudes Politiques de Paris, 2019. tel-04547405

HAL Id: tel-04547405

<https://hal.science/tel-04547405>

Submitted on 15 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Institut d'études politiques de Paris

Paola Tubaro

Décrypter la société des plateformes :
Organisations, marchés et réseaux dans l'économie numérique

Dossier préparé en vue de l'obtention de l'habilitation à diriger des
recherches en Sociologie

présenté le 11 décembre 2019

Volume 2

Jury :

M. Gilles Bastin, Professeur des universités, Institut d'Etudes Politiques de Grenoble (rapporteur)

M. Rodolphe Durand, Professeur (*Joly Family Professor of Purposeful Leadership*), HEC Paris

M. Emmanuel Lazega, Professeur des universités, Institut d'Etudes Politiques de Paris (garant et rapporteur)

Mme Béatrice Milard, Professeure des universités, Université de Toulouse Jean Jaurès (rapporteure)

M. José Luís Molina González, Professeur (*Catedrático de Universidad*), Universitat Autònoma de Barcelona

M. T.A.B. (Tom) Snijders, Professeur (*Hoogleraar*), Rijksuniversiteit Groningen

Table des matières

Introduction	xi
Paola Tubaro	
1 Contexte et objectifs	xi
2 La plateforme entre marché et entreprise	xiii
2.1 La plateforme n'est pas un marché, mais son faiseur	xiv
2.2 La plateforme comme organisation	xvii
3 De la société organisationnelle à la société de plateformes?	xviii
3.1 L'emprise des plateformes sur la société.....	xix
3.2 Deux niveaux d'action	xx
4 Économie de partage et collaborative : la plateforme autrement? ..	xxi
4.1 La brève illusion du partage	xxii
4.2 Décentralisation et coopération	xxiv
5 Plateformes, travail et production de données	xxv
5.1 Formaliser l'informel pour sa mise en données	xxv
5.2 Un marché du travail de la donnée	xxvii
6 Contenus du mémoire original d'HDR.....	xxix
7 Annexe	xxxii
7.1 SharingNetworks (SN)	xxxii
7.2 Digital Platform Labor (DiPLab).....	xxxiv
7.3 Recent ethical challenges in social network analysis (RECSNA)	xxxvi
Références	xxxvii

Disembedded or deeply embedded? Toward a multi-level network analysis approach of the online platform economy	1
Paola Tubaro	
1 Introduction : embeddedness in markets, firms and platforms	1
2 Conceptions of embeddedness : from Polanyi to Granovetter and beyond	4
3 A case study : Platforms for online labor	5
3.1 Digital platforms and embeddedness.....	7

3.2	Online labor platforms	8
4	An inventory of platforms	9
5	Platform business models and embeddedness	10
5.1	Platform as markets : intermediaries	11
5.2	Platforms as firms : multiple layers	13
5.3	Platform-market, platform-firm and embeddedness	15
6	Conclusions	18
7	Acknowledgements	21
	References	21
	Festival attendance as a network of attention: A case study in the sharing economy	27
	Paola Tubaro	
1	Introduction	27
2	A “Sharing Economy” event	30
2.1	The OuiShare Fest as a mini-economy of attention	31
2.2	Networks of speakers and attendees	32
3	Whom do speakers pay attention to?	33
3.1	Thematic affinity	34
3.2	Competition	34
3.3	Status	36
3.4	Reciprocity	37
4	Empirical setting and data	37
4.1	Dependent variable	38
4.2	Independent variables	39
5	An Exponential Random Graph Model of attention choices	42
6	Results	45
6.1	Estimates	45
6.2	Interpretation	46
7	Conclusions	49
8	Annex	50
8.1	Oui Share Fest session themes, 2013-17	50
8.2	Logistic regression model of posting attendance choices	52
8.3	ERGM goodness-of-fit diagnostics	52
9	Acknowledgments	53
	References	54
	Whose results are these anyway? Reciprocity and the ethics of “giving back” after social network research	59
	Paola Tubaro	
1	Introduction: communicating research results as an ethical problem	59
2	Bridging literatures	62
3	Social networks research and case-based reasoning	64
4	Giving personal network results to individuals	65
5	Giving whole network results to communities	68

6	Beyond networked communities: sharing results with the media and wider audiences	71
7	Discussion: dual vulnerability and reciprocation	72
8	Conclusions	74
9	Acknowledgements	76
	References	76
	The trainer, the verifier and the imitator: Three ways in which human platform workers support automation	81
	Paola Tubaro, Antonio A. Casilli, and Marion Coville	
1	Introduction	82
2	Micro-work as an instance of digital platform labor	83
3	The labor-for-data needs of machine learners	85
3.1	How data fuel machine-learning algorithms	85
3.2	A structural demand of micro-work for AI data preparation?	87
4	Insights from fieldwork	87
5	Micro-work for AI preparation	89
5.1	Platforms' offer of AI preparation services	89
5.2	Micro-workers' experience of AI preparation	91
6	Micro-work for "AI impersonation"	92
7	Micro-work for AI verification	94
8	Discussion and conclusions	95
9	Acknowledgments	98
	References	99
	Counting "micro-workers": Methodological and societal challenges around new forms of labor	101
	Paola Tubaro, Clément Le Ludec, and Antonio A. Casilli	
1	Introduction	102
2	The challenges of measuring platform labor (especially micro-work).	104
2.1	Why surveys and administrative data fail to detect micro-work	104
2.2	Why data reported by micro-work companies cannot be taken at face value	106
2.3	The limitations of alternative approaches to estimating micro-work	107
3	Three approaches to estimating the number of micro-workers in a specific country	108
3.1	What micro-work platforms operate in France?	109
3.2	Analyzing public data available from platforms	110
3.3	Analyzing experimental data with the "capture-recapture" model	111
3.4	Analyzing web traffic data to measure the audience of micro-work platforms	114

3.5	Combining methods to obtain a range of values	116
3.6	Accounting for users who micro-work on several platforms	117
4	How to interpret results as indicators of micro-workers' level of activity?	118
4.1	Very active micro-workers	119
4.2	Regular micro-workers	119
4.3	Occasional micro-workers	120
5	Discussion: strengths and limitations of present methods	121
5.1	Methodological limitations	121
5.2	Why our methods are nevertheless reliable	122
5.3	Concluding remarks	123
6	Acknowledgments	124
	References	124
	Who are the micro-workers behind AI?	127
	Paola Tubaro	
1	Introduction	127
2	Who are the micro-workers?	128
2.1	Low income vs industrial countries	129
2.2	Not only millennials	129
2.3	Not only men	130
2.4	Highly educated workers, in work, with family	132
2.5	A lower-income population that needs extra earnings	132
2.6	How do micro-workers fare?	133
3	What do micro-workers do?	134
3.1	What are micro-working tasks?	134
3.2	Why tasks are challenging after all	135
4	How microwork feeds AI	136
4.1	AI preparation	137
4.2	AI verification	137
4.3	AI impersonation	138
5	Conclusions	138
6	Acknowledgements	139
	References	139

Introduction

Paola Tubaro

1 Contexte et objectifs

À la pointe de l'innovation socio-économique contemporaine, les **plateformes numériques** comme Google, Facebook, Uber et Airbnb sont exemplaires des gains de productivité que rendent possibles les technologies de l'information et de la communication (Benavent, 2016; Gawer, 2009). Les plateformes ouvrent la voie à de nouvelles pratiques de production, d'échange et de consommation, en s'appuyant notamment sur des dispositifs novateurs de mise en relations de groupes d'acteurs divers. Elles investissent un nombre important de secteurs d'activité, allant de la production de biens immatériels et culturels au tourisme (voyages, hébergement), au transport urbain, à l'alimentation et aux services à la personne. Elles sont porteuses d'une remise en cause – saluée par certains, crainte par d'autres – d'acquis sociaux en termes notamment de conditions de travail, protection des consommateurs, structures de la concurrence, frontière public-privé, liberté de parole et accès à l'information (Amar & Viossat, 2016; Berg et al., 2018; Codagnone et al., 2016; INRS, 2018). Les débats sur le « capitalisme des plateformes » (Abdelnour & Bernard, 2018; Srnicek, 2016) s'étendent à la « plateformisation de l'État » (Brown et al., 2017; O'Reilly, 2011) autant qu'à des modèles alternatifs comme le « coopérativisme des plateformes » (Scholz & Schneider, 2017)).

À la base, tous les acteurs s'accordent pour voir dans les plateformes des **dispositifs de coordination facilitant la rencontre, la mise en relation et l'interaction d'acteurs divers** (Brousseau & Penard, 2007). Pour autant, le champ ne s'est pas encore stabilisé et même le mot « plateforme » reste ambigu en raison de ses origines plurielles puisant dans l'architecture, l'informatique (de Reuver et al., 2018), l'économie (Evans, 2011), la politique (Gillespie, 2010). Les typologies sont aussi nombreuses qu'éphémères (Codagnone et al., 2016; Howcroft & Bergvall-

Kåreborn, 2019; Schmidt, 2017). Ce manque de clarté freine, voire même fourvoie, toute tentative de concevoir des politiques publiques adaptées.

Pour y voir plus clair, je propose de conceptualiser l'avènement des plateformes numériques comme une transformation qui actualise et ravive trois tensions majeures qui ont traversé les sciences sociales tout au long de leur histoire : entre **marché et entreprise** comme modes de coordination de l'action économique, entre **profit et don** comme logiques de régulation des échanges, et entre **formel et informel** comme modalités de structuration de l'action collective.

Cet exercice exige de se situer au carrefour des disciplines, et dans une certaine mesure, d'en parcourir l'histoire pour révéler toutes les facettes des concepts et des théories qui sont aujourd'hui mobilisés, mais qui portent avec eux l'héritage de leurs origines et des contextes dans lesquels ils ont été déployés. Je montre que la tension entre marché et entreprise peut être résolue en articulant les théories de la « société organisationnelle » en sociologie économique et en sociologie des organisations (Lazega, 2015, 2016) avec l'économie des coûts de transactions (Williamson, 1975) et, dans une alliance inédite, avec les modèles économiques des « marchés multifaces » (ou « multi-versants ») (Evans & Schmalensee, 2016; Tirole, 2016). L'approche sociologique ramène les plateformes à des structures organisationnelles susceptibles de générer des asymétries de pouvoir, les approches économiques éclairent les arbitrages des acteurs et la distribution des gains et des pertes qui en résulte. Cette perspective fait penser que la transition de la société organisationnelle à l'actuelle **société des plateformes** se réalise dans la continuité.

La tension profit/don se comprend également dans la continuité, dans la mesure où les plateformes reproposent, dans un contexte technologiquement rénové, le dilemme qui a traversé l'histoire de toute l'économie sociale et solidaire, tiraillée entre l'exigence de contribuer au changement social et celle d'assurer sa solidité financière – la dernière finissant souvent par peser plus que la première. Ainsi, il n'est pas surprenant de découvrir que les discours d'innovation sociale qui entouraient l'essor des plateformes dites « de partage » ou « collaboratives » (Benkler, 2006; Botsman & Rogers, 2010; Filippova, 2016) au début, restent aujourd'hui en retrait. Moins médiatisées que leurs homologues qui poursuivent le profit, des expériences socialement plus innovantes en sont venues à occuper des positions de niche, devenant souvent l'apanage de collectivités locales et d'organisations à but non lucratif.

À son tour, la tension formel/informel s'inscrit dans le sillage de la sociologie de la quantification (Desrosières, 2008) en transformant un objet de réflexion méthodologique, la **donnée**, en objet substantif lui-même –une ressource autour de laquelle s'organisent les rapports sociaux sur les plateformes et se définissent des relations de pouvoir dont découlent des inégalités. La numérisation (et par là, on le verra, la formalisation) de pans de plus en plus étendus de la société peut propulser des avancements importants en sciences sociales, mais ces données restent le plus souvent l'apanage de grands groupes privés (Lazer et al., 2009) réticents à les partager. Les possibilités de changements méthodologiques sont alors limitées par des droits de propriété qui définissent un usage commercial des données, et qui engendrent par ailleurs des conséquences importantes en termes autant sociétaux (menaces pesant

sur la vie privée) qu'économiques (émergence d'un marché globalisé du travail de la donnée).

Cette réflexion est conceptuelle, interrogeant les théories et se confrontant aux auteur.e.s pour bâtir un cadre général d'analyse. La conclusion s'ouvre sur des pistes de recherche surtout empirique visant à en mettre à l'épreuve les aspects principaux.

2 La plateforme entre marché et entreprise

En tant que **dispositifs de coordination**, les plateformes ont été comparées aux deux alternatives classiques du marché et de l'entreprise (Benavent, 2016; Casilli, 2019; Sundararajan, 2016). Le premier opère par l'ajustement des prix, *ex post*, reposant sur l'idée que les réactions des acteurs aux incitations du marché suffisent à produire un ordre sans que celui-ci n'ait été intentionnellement poursuivi. La seconde est une coordination par une forme de plan, *ex ante*, que la littérature attribue à la fonction de direction des ressources qui revient au chef d'entreprise (Coase, 1937) ou à la hiérarchie (Williamson, 1975, 1981), mais qui pourrait découler de toute instance de décision collective (Ouchi, 1979)¹. Dans les systèmes économiques contemporains, les deux formes peuvent coexister, le plan *dans* les entreprises et les prix *entre* les entreprises.

Ce sont les **coûts de transaction** qui déterminent le choix entre marché et entreprise au cas par cas. Cette expression désigne l'ensemble des coûts, souvent de nature informationnelle, nécessaires pour mener à bien un échange entre deux agents économiques (par exemple des coûts de prospection, d'information, de vérification, de négociation, de suivi etc.). Et c'est par sa maîtrise de l'information que la plateforme permettrait de dépasser cette dualité, en enclenchant une baisse extraordinaire des coûts de transaction. En tant qu'infrastructure qui fait rencontrer des groupes d'acteurs distincts, elle a un accès privilégié aux données qu'ils produisent, aux « traces » de leurs actions qu'elle enregistre (section 5.1). Ces données, qui alimentent des algorithmes de plus en plus puissants, sont sa ressource productive première (Srnicsek, 2016), qui diminue les coûts informationnels et facilite les transactions.

Or il peut y avoir des coûts de transaction autant dans l'échange marchand qu'au sein de l'entreprise, où ils sont parfois qualifiés de coûts de la bureaucratie. Dans quelle mesure, donc, la capacité de la plateforme à infléchir ces coûts la rapproche-t-elle du marché, et dans quelle mesure de l'entreprise ? L'intérêt de cette question n'est pas purement théorique, car une *plateforme-comme-marché* serait un simple intermédiaire qui se limite à faciliter la mise en relation entre parties contractantes, et ne serait pas responsables de leurs actions et des résultats qu'elles engendraient. C'est la ligne de défense choisie par plusieurs plateformes lors de procédures judiciaires qui les ont mises en cause. Si Uber était un simple intermédiaire, par

¹ Dans ce qui suit, et sauf mention contraire, j'utilise le mot "entreprise" comme le fait Coase (1937), pour désigner toute structure de prise de décision *ex ante*, qu'elle soit ou non hiérarchique au sens strict du terme. Je reviendrai sur ce point plus loin.

exemple, toute demande de requalification en salariés de ses chauffeurs serait difficile à soutenir. Il en serait autrement si la plateforme était vue davantage comme une entreprise.

Pour répondre à cette question, il faut revenir sur le type spécifique d'actions économiques dans lesquelles la plateforme intervient.

2.1 La plateforme n'est pas un marché, mais son faiseur

L'idée d'un rapprochement avec le marché est répandue dans les débats publics, où les plateformes sont souvent présentées comme des places de marché (*market-places*) sur le modèle d'Amazon ou eBay. Il est dit parfois qu'elles « créent » des marchés là où ils n'existaient pas, ou les font passer à l'échelle lorsqu'ils étaient auparavant petits (par exemple, le marché de la livraison de repas à domicile) ou informels (par exemple, les services d'aide à domicile). D'aucuns ont même pu parler d'une tendance à la « désintermédiation ». Sur quels fondements ces idées reposent-elles ?

2.1.1 Le rapprochement hâtif entre plateforme et marché

En effet, les plateformes se servent du mécanisme de coordination par le prix que les économistes attribuent aux forces impersonnelles et supra-individuelles des marchés. L'algorithme de Uber capable de suivre en temps réel l'évolution de l'offre et de la demande de trajets urbains, quartier par quartier, et d'ajuster sa grille tarifaire afin de compenser d'éventuels déséquilibres (« tarification dynamique » ou « *surge pricing* »)², semble à première vue l'incarnation même du marché d'un bien théorisé par les économistes depuis l'œuvre fondatrice d'Alfred Marshall (Tricou, 2008). Uber paraît même exemplaire de l'efficacité d'une coordination par les prix – passant par la mise en place d'incitations monétaires pour éliminer d'éventuels excès d'offre ou de demande (Cachon et al., 2017; Castillo et al., 2017; Hall et al., 2015; Zha et al., 2017) – opposé à d'autres formes de régulation comme la tarification fixée d'avance imposée aux taxis, incapable d'équilibrer offre et demande (Cramer & Krueger, 2016).

Le succès d'une interprétation de la plateforme comme marché correspond à une perspective enracinée dans l'histoire de la discipline économique. Le programme de recherche dit de « Arrow-Debreu », socle de la micro-économie néoclassique (Arrow & Hahn, 1971; Debreu, 1959), était axé sur des modèles théoriques de sys-

² La tarification dynamique ajuste les prix en fonction de l'état de la demande et de l'offre instantanées. En cas de pic de demande des passagers dans une zone donnée, l'algorithme d'Uber augmente les prix dans cette zone pour y attirer des chauffeurs avec la promesse d'une meilleure rémunération ; cette hausse des prix va aussi décourager certains passagers qui renonceront au service, le trouvant trop cher. Ces deux ajustements simultanés visent à retrouver l'équilibre entre offre et demande.

tèmes de marchés décentralisés purs, où les entreprises n'étaient que des participants au même titre que les individus. L'idée de Alchian & Demsetz (1972) que l'entreprise ne serait qu'une variante du marché, car son fonctionnement repose au fond sur des *contrats*, évacue toute différence majeure entre (par exemple) le rapport employeur-employé et le rapport vendeur-acheteur. Suivant Segrestin & Hatchuel (2012, p. 24), on peut entrevoir dans ces visions réductrices de l'entreprise une persistance de pratiques remontant au début du XIXe siècle, quand il n'y avait pas de relation durable d'emploi et quand les ouvriers « louaient » leurs services au moyen de contrats de louage ponctuels.

De manière plus subtile, Gillespie (2010) remarque que les étymologies du terme « plateforme » évoquent un espace d'opportunités, un lieu ouvert accueillant les apports des uns et des autres de manière égalitaire – ce qui n'est pas sans rappeler les espoirs d'émancipation dont l'essor du « marché » était historiquement porteur par rapport à des formes pré-modernes d'organisation économique (Fontaine, 2014; Peart & Levy, 2005).

L'élément idéologique est donc très important, et ne se limite pas à des arguments reflétant les intérêts matériels et immédiats des gérants des plateformes, car il s'inscrit dans une histoire longue et résonne avec la culture économique, centrée sur l'imaginaire du marché et de ses vertus plus ou moins avérées. Mais dans le cas spécifique des plateformes, de quel marché parle-t-on ? Pour répondre à cette question, il faut se tourner vers les théories récentes des marchés dits « multi-faces » ou « multi-versants ».

2.1.2 Plateforme et marchés multi-faces

La théorie économique des marchés bi- ou multi-faces (*two-sided* ou *multi-sided markets*) (Eisenmann et al., 2006; Evans & Schmalensee, 2016; Hagiu & Wright, 2015; Rochet & Tirole, 2003; Rysman, 2009; Tirole, 2016) s'intéresse à la mise en relation de catégories distinctes d'acteurs par un intermédiaire, comme justement une plateforme. Par exemple, Uber fait interagir deux catégories, les passagers et les chauffeurs, alors que Deliveroo en coordonne trois, les mangeurs, les restaurateurs et les livreurs. Dans tous les cas, il existe des externalités dites « de réseau » par lesquelles un plus grand nombre d'utilisateurs est une incitation à participer pour d'autres. Ces externalités peuvent être intra-groupe (s'inscrire sur Facebook pour communiquer avec les nombreux autres inscrits), ou inter-groupes (conduire pour Uber s'il y a un grand nombre de passagers, ou inversement, se faire conduire par Uber s'il y a un grand nombre de chauffeurs). Ces externalités sont un levier pour la plateforme qui en capture la valeur : sa base d'utilisateurs représente l'une de ses principales richesses, à côté de son capital informatique et logiciel. Pour s'assurer un maximum d'utilisateurs, la plateforme peut être amenée à exploiter d'éventuelles différences d'élasticité de la demande des divers groupes, en subventionnant celui qui est le plus sensible aux prix : par exemple aujourd'hui, un particulier utilise gratuitement les fonctionnalités de base de Facebook alors que les annonceurs publicitaires paient.

Cette approche explique les dynamiques *winner-take-all* qui ont vite transformé certaines plateformes en des acteurs incontournables à l'échelle multinationale (Google, Facebook...), avec des chiffres d'affaires en progression toujours plus forte. L'échelle et la concentration sont des issues fréquentes, quoique non inévitables (Brousseau & Penard, 2007) des externalités de réseau dans un marché multi-faces. Elles expliquent aussi la « gratuité » de services comme le moteur de recherche de Google, qui ne relève pas d'une forme de bienveillance de la part de ses producteurs mais résulte d'un **calcul de tarification optimale**, le prix nul pour la face la plus sensible du marché (les usagers particuliers) encourageant l'adhésion, pour rendre la plateforme plus attractive pour l'autre face (les annonceurs publicitaires), à qui le service est tarifé.

En quoi cette perspective théorique renforce, confirme ou met à l'épreuve le rapprochement entre plateforme et marché évoqué plus haut (sous-section 2.1.1) ?

2.1.3 La plateforme comme faiseur de marchés

L'essor de la théorie des marchés multi-faces s'inscrit dans la dynamique de déclin graduel du paradigme Arrow-Debreu, accompagné d'un intérêt croissant pour les enjeux de l'information (Steiner, 2017). La différence est de taille : en effet les théories néoclassiques, à la Arrow-Debreu, supposaient une coordination impersonnelle par les seuls prix qu'aucun acteur ne saurait entièrement contrôler, dans des conditions de concurrence « pure et parfaite », alors que ce qui ressort ici est l'importance d'un acteur, la plateforme, dans la mise en place, l'exécution et le suivi des transactions entre les autres parties. La théorie des marchés multi-faces redécouvre ainsi les intermédiaires et les organisateurs – les « **faiseurs de marchés** » pour ainsi dire, et finit par (implicitement) corroborer une vision auparavant minoritaire et hétérodoxe en économie, mais répandue en sociologie économique, selon laquelle la coordination marchande n'est pas entièrement émergente et *ex post*, mais est au moins partiellement coconstruite *ex ante* par les institutions qui régulent, organisent, et surveillent les échanges.

Ces institutions ne sont pas seulement des règles de fond, comme l'attribution des droits de propriété, mais de cadres construits, organisés, peuplés d'acteurs humains qui collaborent pour rendre possibles les transactions. Dans les deux dernières décennies, les économistes ont pris la mesure du rôle décisif d'acteurs comme par exemple les organisateurs d'enchères (Kirman, 2001), ou les chambres de compensation dans les problèmes d'appariement (Roth, 2002) ; les sociologues sont allés plus loin en étudiant, entre autre, les salons commerciaux (Garcia-Parpet, 1986), les normes de qualité (Karpik, 2000), les professionnels de l'intermédiation (Cochoy & Dubuisson-Quellier, 2000). Fligstein (2002) définit le marché par la coexistence de différents ensembles d'institutions incluant les droits de propriété, les structures de gouvernance des firmes, les règles de l'échange et les formes du contrôle. Ces institutions sont incarnées par des acteurs, le plus souvent des collectifs organisés (publics ou privés), qui les mettent en œuvre et en assurent le suivi.

C'est ce rôle que reprend aujourd'hui la plateforme, se positionnant au milieu des parties prenantes pour encourager, contrôler et réguler leurs échanges. Elle n'est pas un marché *stricto sensu* mais son faiseur, l'acteur qui en établit et en applique les règles, qui recrute des participants à l'échange, apparie les contreparties dans les transactions, gère les flux d'information. Il semble alors qu'il faut voir la plateforme davantage comme une structure organisée – *en amont* du système de contrats marchands et non pas au même niveau que celui-ci. Pour tester la solidité de cette idée, il faut la confronter aux théories de l'organisation – laissant donc de côté le champ de l'économie pour se rapprocher de la sociologie.

2.2 La plateforme comme organisation

Si la plateforme est un faiseur de marchés, elle peut être vue comme une **organisation** – en entendant par ce terme non pas un mode de coordination alternatif au marché comme chez Coase (1937), mais les règles qu'il faut dans tous les cas se donner pour circonscrire les possibilités d'action et orienter les comportements des acteurs. En ce sens, autant l'entreprise que le marché sont des espaces organisés, que l'on peut situer sur un continuum au lieu de les opposer, comme proposait de le faire l'école française d'analyse stratégique des organisations (Crozier & Friedberg, 2014; Reynaud, 1988, 1997). Ainsi, malgré les hésitations terminologiques que l'on trouve dans la littérature, l'organisation n'est pas nécessairement hiérarchique : elle peut être régie par des proportions variées de bureaucratie (Weber, 1922; Crozier, 1971) et de collégialité (Lazega, 2001), de règles formelles et informelles, d'accords explicites et implicites. Elle s'étend à couvrir **toute forme de structuration et de régulation (même faible) des contextes à l'intérieur desquels l'action se développe** (Crozier & Friedberg, 2014; Friedberg, 1997). Selon Lazega & Mounier (2002), les marchés peuvent être vus comme des organisations collégiales puisqu'il n'y a pas de hiérarchie formelle entre entrepreneurs, et que les processus relationnels observables à leur échelle, qui les aident à gérer leurs interdépendances, sont objectivables et modélisables. Dans cette même perspective, les marchés multi-faces coordonnés par une plateformes sont des organisations mixtes, mi-collégiales et mi-hiérarchiques.

Une définition très générale voit dans l'organisation un « outil ayant une vie propre » (Selznick, 1949, p. 10), qui ressemble, tout en les dépassant, les motivations et les intérêts de ses membres qui peuvent tous, d'une manière ou d'une autre, y trouver leur compte. Elle implique un **projet commun qui s'inscrit dans la durée, aboutit à des formes d'action collective, et met en place une gestion coordonnée (*ex ante*) de ses ressources, que ce soit par des relations d'autorité ou de coopération entre pairs**. Cette gestion interne est souvent assise sur un capital important de connaissances partagées, par exemple techniques. La plateforme remplit ces conditions en mettant en place des formes de gestion *ex ante* des ressources humaines et matérielles dont elle dispose, s'appuyant d'ailleurs sur un patrimoine conséquent notamment en termes de ressources informatiques (outils de

calcul, propriété intellectuelle de logiciels). C'est en ce sens qu'elle se rapproche d'une organisation – dont la fonction est d'organiser le marché.

Ces premiers résultats sont déjà lourds de conséquences, car ils suggèrent qu'au lieu d'ouvrir la voie à des processus de décentralisation comme certains ont pu le croire, les plateformes peuvent avoir l'effet inverse, centralisant l'information. Par exemple, Uber collecte l'information sur l'offre et la demande de déplacements pour faire opérer son algorithme de tarification dynamique – en contrôlant le processus de formation des prix et en les imposant à ses chauffeurs et passagers, au lieu de les laisser établir librement comme le ferait le marché selon l'orthodoxie économique. Pour ce faire, et encore une fois contrairement à l'idéal-type du marché économe en information (Hayek, 1945), la plateforme en est gourmande : elle se nourrit de données numériques (« *big data* »), traces d'actions et interactions humaines médiatisées par ordinateur, grâce auxquelles elle réalise ses fonctions de coordination et en monétise les résultats (Bastin & Tubaro, 2018).

C'est ainsi que les plateformes peuvent se mettre en position à pouvoir contrôler les systèmes de marchés pour lesquels elles servent d'intermédiaires. La théorie économique des marchés « multi-faces », déjà évoquée, met en avant les « externalités de réseaux » imputables à la simple présence d'utilisateurs sur les différents versants (sous-section 2.1.2). Lorsque ces effets produisent un passage à l'échelle très rapide, ils peuvent permettre d'atteindre des positions de monopole ou quasi-monopole, forme de pouvoir économique par excellence.

3 De la société organisationnelle à la société de plateformes ?

Si le marché n'est pas la meilleure description de la plateforme, il n'est pas non plus emblématique des systèmes socio-économiques qui ont précédé l'avènement des technologies numériques. Selon Simon (1991) « *the economies of modern industrialized society can more appropriately be labeled organizational economies than market economies* » (p. 42). De façon similaire, Coleman (1974, 1982), Perrow (1991), Presthus (1962) et Stinchcombe (2001) décrivaient la société qui les entourait comme une « *organizational society* ». Il s'agit de voir les organisations comme les forces structurantes de nos économies et sociétés – qui s'appréhendent donc au niveau méso (Lazega & Mounier, 2002; Lazega, 2003). Le contraste d'avec les analyses des marchés à la Arrow-Debreu (sous-section 2.1.1), qui espéraient pouvoir faire du mécanisme supposé impersonnel des prix le seul lien entre le niveau micro des individus (et des entreprises réduites elles-mêmes à des individus) et le niveau macro du système, est flagrant.

Si l'on admet aujourd'hui que nous sommes en présence d'une « société des plateformes », s'agirait-il d'un prolongement ou d'une reconfiguration des lignes directrices de la société organisationnelle ? Cette question revient à se demander dans quelle mesure la plateforme du XXI^e siècle étend son contrôle à des pans de plus en plus étendus de la société, au même degré que l'organisation du XX^e siècle.

Pour répondre, il faut comparer l'activité et le fonctionnement des plateformes aux différentes interprétations qui ont été proposées de la société organisationnelle.

3.1 *L'emprise des plateformes sur la société*

Perrow (1991), qui analysait la société organisationnelle surtout en relation aux grandes entreprises privées, montrait que celles-ci créent des formes de **dépendance économique**, s'assurant la disponibilité des individus pour leurs besoins de main d'œuvre. Si les plateformes mettent à mal le modèle du salariat qui dominait au moment où Perrow écrivait, elles s'inscrivent dans une continuité d'objectifs et de pratiques par rapport à certaines caractéristiques et pratiques de l'entreprise organisée traditionnelle. Dans un contexte économique fortement appauvri par la crise financière de 2008, les opportunités d'emploi salariat se réduisant dans les secteurs traditionnels, le travail pour les plateformes (conduire pour Uber, livrer des repas pour Deliveroo, transcrire des textes pour Amazon Mechanical Turk), pourtant précaire et dépourvu de formes de protection sociale, devient incontournable pour une partie de la population (Casilli, 2019; Casilli, Tubaro, Le Ludec, Coville, Besenval, Mouhtare & Wahal, 2019; Graham & Anwar, 2019; Gray & Suri, 2019; Lehdonvirta, Kässi, Hjorth, Barnard & Graham, 2019).

Perrow (1991) notait également l'**externalisation des coûts sociaux de l'activité organisée**, comme la pollution issue des processus productifs. La plateforme n'est pas étrangère à cette tendance. Son appétit pour les données (j'y reviendrai) engendre des coûts environnementaux colossaux (notamment avec les applications de l'intelligence artificielle, très gourmandes en ressources naturelles comme le montrent par exemple Crawford & Joler (2018) et Strubell et al. (2019)) et des risques en termes d'insuffisante protection des informations personnelles et de la vie privée des usagers, voire même des dérapages (Tubaro et al., 2014; Tubaro, 2018b).

Un dernier aspect mis en avant par Perrow (1991) est le contrôle, réalisé non seulement par voie hiérarchique mais surtout par des moyens automatisés et peu intrusifs, passant par la formalisation, la standardisation et la mesurabilité des activités réalisées. Les plateformes d'aujourd'hui poussent à l'extrême ces formes de **contrôle, au moyen d'algorithmes basés sur une exploitation de plus en plus sophistiquée des données** (Zuboff, 2019), capables d'imposer une discipline de manière autant discrète qu'efficace, pouvant aller jusqu'à « désactiver » des usagers sur la base de mesures de performance (Rosenblat & Stark, 2016; Rosenblat, 2018).

D'autres visions de la société organisationnelle comme celle de Tilly (1998) soulignent comment différents acteurs et groupes d'acteurs peuvent s'avérer gagnants ou perdants par leur capacité d'interagir avec l'organisation, dans un processus **d'accapement d'opportunités** (*opportunity hoarding*). Une version de cette idée s'applique aux plateformes qui, selon Schor (2018), font en sorte que des membres des classes moyennes en trajectoire sociale descendante s'approprient des opportunités de travail rémunéré auparavant réservées aux classes populaires,

entrant en concurrence directe avec celles-ci et parvenant dans bien de cas à les exclure. La responsabilité de plateformes comme Uber ou Deliveroo est d'avoir déstigmatisé des activités relativement peu qualifiées (conduire des voitures, livrer des marchandises), notamment en les associant à une culture dynamique et technophile. Les plateformes auraient ainsi creusé les inégalités, notamment au sein des couches les moins aisées de la population (Schor, 2017).

Dans tous les cas, les plateformes ont des effets potentiellement forts sur la société qui les entoure, et d'où elles puisent les ressources, notamment humaines, qui les font vivre. Elles prolongent ainsi des effets déjà avérés des organisations, tout en modifiant la manière dont ils se donnent à voir.

3.2 Deux niveaux d'action

Une conception alternative de la société organisationnelle, inspiré du principe de la « dualité » de Breiger (1974), est celle qui est proposée par Lazega (2015, 2016, 2017), partant du principe que l'action collective se déroule à plusieurs niveaux interdépendants, autant à celui de l'individu qu'à celui de l'organisation. La société organisationnelle apparaît alors comme **une société de classes où le pouvoir et la répartition des ressources sont liés au contrôle des appareils organisés**. Dans ce contexte, les individus façonnent leurs opportunités, stratégies et trajectoires à travers leurs capacités à intégrer les organisations et s'en servir. Chaque niveau (individu et organisation) a sa propre temporalité, et le résultat final se joue souvent en reportant le coût de leur « **synchronisation** » sur le niveau le plus faible.

Comment étendre ces idées aux plateformes ? Il est possible de voir les coûts de synchronisation comme une variante des coûts de transaction (Coase, 1937), dans le contexte des activités de coordination exercées par les plateformes, au milieu de différentes parties prenantes qui constituent des faces séparées du marché. Ces coûts émanent de la difficulté à prendre en compte simultanément les attentes, les besoins et les contraintes de toutes ces parties : par exemple, Deliveroo doit synchroniser l'activité de ses mangeurs, restaurateurs et livreurs pour que chaque repas arrive à destination dans le temps statué. Ces coûts que les technologies n'arrivent pas à comprimer peuvent affaiblir le modèle d'affaires de la plateforme : la réponse est alors de les **reporter sur les acteurs de la face de marché la plus dépendante économiquement**. L'exemple de la livraison instantanée montre que ce sont les travailleurs qui en font les frais (Aguilera et al., 2018). Ne pouvant guère modifier la temporalité des mangeurs qui commandent à midi et le soir, ayant peu de levier sur les restaurateurs qui doivent gérer leurs salles en plus des services à emporter, la plateforme pousse les livreurs à s'ajuster, pour qu'ils soient plus présents dans les zones et aux horaires où la demande est la plus élevée – quelles que soient leurs préférences et contraintes personnelles.

Un autre exemple est celui des plateformes de travail en ligne, où la réalisation de prestations ne nécessite pas la coprésence physique des parties intéressées. Ces plateformes sous-traitent des prestations allant du *design* et de la programmation in-

formatique jusqu'à de petites tâches de préparation de données (sous-section 5.2) à des exécutants qui les effectuent en ligne. La temporalité de la plateforme est essentiellement alignée sur celle des clients, qui se trouvent généralement dans des pays du Nord du monde (Graham et al., 2017) et qui publient des tâches durant leurs horaires de bureau. Les travailleurs n'ont pas la même temporalité, ne serait-ce qu'en raison du décalage horaire car beaucoup d'entre eux se trouvent dans des pays du Sud. Un excès d'offre massif entretenu par les politiques des plateformes (Graham & Anwar, 2019) les oblige à s'adapter, sans tenir compte de leurs circonstances de vie, préférences ou contraintes personnelles, les poussant jusqu'à travailler la nuit pour se synchroniser avec les horaires des clients (Lehdonvirta, 2018).

Ces coûts importent parce qu'ils sont **déclencheurs des processus de création (ou renforcement) d'inégalités** (sous-section 3.1) : c'est en imposant à un seul groupe de les prendre en charge, que la plateforme ouvre la voie à des formes de dépendance économique et d'accaparement d'opportunités ; c'est pour les gérer qu'elle en externalise une partie et/ou met en place des systèmes de contrôle. Par cet ensemble complexe de mécanismes socio-économiques, la plateforme peut devenir un système producteur d'inégalités et d'exclusion.

J'ai tenu à souligner la nature d'organisation de la plateforme du XXI^e siècle, et à l'inscrire dans le prolongement des transformations sociétales forgées par les organisations tout au long du XX^e siècle, comme une étape nécessaire pour dé-idéologiser les débats qui la rapprochaient trop vite du marché, et pour clarifier son rôle dans le système économique. Il s'agissait aussi de donner plus de cohérence à des résultats empiriques provenant d'approches disparates, et parfois de champs disciplinaires divers et déconnectés. Le résultat est un cadre riche qui met au jour la force des effets de l'avènement des plateformes sur l'ensemble de nos économies et sociétés. Les éléments de continuité qui se dégagent sont une aide pour la recherche tout autant que pour l'action, suggérant la possibilité d'étendre aux plateformes certains des concepts, outils, méthodes, et préconisations de politique publique qui ont déjà fait leurs épreuves dans l'analyse des organisations. Bien sûr, cette extension ne pourra pas se faire sans adaptation, et un enjeu pour la recherche à venir est de définir très précisément comment celle-ci peut se configurer en tenant compte des spécificités du monde des plateformes (comme la disparition du salariat).

4 Économie de partage et collaborative : la plateforme autrement ?

La conceptualisation de la société des plateformes ne saurait se passer d'une discussion des promesses de changement dont elle a été porteuse. Sa toute première diffusion vers la fin des années 2000 se démarquait par son lien étroit avec des espoirs d'innovation non seulement économique et technologique, mais aussi sociale, écologique et culturelle. On parlait moins d'économie des plateformes que d'une « économie du partage » (*sharing economy*), un terme plus chargé moralement même si

son sens et son périmètre étaient loin d'être consensuels. Forme d'action collective émergente, la plateforme semblait initialement adresser, dans un contexte technologiquement transformé, le problème récurrent de la conciliation entre logiques de profit et logiques de don³, qui s'était été déjà posé maintes fois sans jamais être résolu (par exemple au regard du secteur coopératif, de l'économie sociale et solidaire, de la micro-finance, de l'investissement socialement responsable).

Ces promesses à l'apparence inédites se situaient dans la continuité d'un imaginaire du numérique qui, dès les années 1990, avait cru en une « économie du don *high-tech* » (Barbrook, 2000 (1999)), une forme de « communauté » (Wellman & Gulia, 1999) en ligne cohésive et régie par un sentiment d'identité fédérateur (Rheingold, 2000), avec des comportements pro-sociaux (Kollok, 1999). Encore dans les années 2000, le vocabulaire du web faisait ressortir la réciprocité, la gratuité, le don (Alter, 2009; Caillé, 2000) tout autant que l'« amitié » (Boyd, 2006; Casilli, 2010; Doueïhi, 2011), avec d'ailleurs des exemples à succès de projets à la structure communautaire et coopérative comme Wikipedia.

L'économie du partage des années 2000 est héritière de ces traditions, même s'il n'y en a pas une seule définition, s'agissant d'une juxtaposition d'objectifs, de positions politiques, de visions de la société. Il est utile de les parcourir brièvement, avant d'analyser leur mise en œuvre.

4.1 La brève illusion du partage

L'idée d'une économie du partage a pris différentes formes, visant d'abord à une rénovation de nos **modes de consommation**, fondée sur la mise en commun d'actifs excédentaires, matériels et immatériels, comme dans le célèbre exemple de la perceuse, que cinquante millions d'américains posséderaient mais qui ne serait utilisée, en moyenne, qu'une dizaine de minutes pendant toute sa durée de vie (Botsman & Rogers, 2010). La partager allait alors éliminer le gaspillage, permettre de faire des économies, et contribuer à un changement des normes sociales en dévalorisant la propriété par rapport au simple accès (Rifkin, 2001). Dans la sphère de la **production**, étaient mis en avant des modèles en pair-à-pair, avec partage des connaissances (Bauwens & Sussan, 2005) et réappropriation des savoirs traditionnels à côté des technologies numériques. Les communautés de développeurs de logiciels « *open source* » étaient un exemple souvent évoqué des gains de productivité se produisant lorsque chacun contribue, sans que personne n'ait l'exclusive propriété du résultat.

Ces espoirs ne se sont pas matérialisés. Il est désormais clair qu'il ne s'agit guère de partage ni de collaboration, mais d'une offre de services commerciaux, portée par la croissance parfois spectaculaire de grandes plateformes multi-nationales à

³ J'utilise le concept de don au sens large pour dénoter toute logique d'échange qui s'écarte de la pure poursuite du profit. Je peux ainsi couvrir le large éventail de formes que cette notion a pris au sein de l'économie des plateformes, tout en permettant de la mettre en rapport avec une littérature plus générale et à une histoire dont les origines précèdent de loin celles des technologies numériques.

but lucratif comme Airbnb et Uber. La consommation contrôlée et respectueuse de l'environnement que prônaient Botsman & Rogers (2010) a largement laissé la place à ce que d'aucuns appellent l'économie de la « fonctionnalité », où l'intérêt pour l'efficacité d'un service prévaut sur ses retombées sociales, culturelles ou environnementales. Des plateformes comme Uber et Airbnb qui suscitaient des enthousiasmes au début par leur promesse de contribuer au lien social en valorisant la rencontre et l'hospitalité, se sont trouvées au centre d'âpres polémiques pour leur traitement des travailleurs, leur désinvolture fiscale, leurs externalités négatives sur le trafic urbain, la disponibilité de logements et les prix de l'immobilier (Ravenelle, 2019; Rosenblat, 2018). Le débat s'est déplacé sur le travail, précarisé et dépouillé de ses prérogatives en termes de droits et d'accès à la protection sociale (Casilli, 2019; Gray & Suri, 2019; Roberts, 2019; Scholz, 2016)). Certains des promoteurs des débuts ont pris leurs distances par rapport aux formes les plus ouvertement commerciales (Bauwens, 2015; Bauwens & Kostakis, 2017), et un représentant de premier plan de OuiShare a déclaré le collaboratif « fini » (De Grave, 2016).

La consommation engagée (achat de produits du commerce équitable, de l'agriculture biologique, ou de producteurs locaux) n'a pas pour autant complètement disparu, et certains usagers continuent de penser qu'elle fait partie des solutions, non du problème (Fitzmaurice et al., 2018). Elle est pourtant restée en retrait, gérée principalement au niveau local, paradoxalement par des organisations pas ou peu numérisées : par exemple dans le champ de l'alimentation, des coopératives (comme La Louve à Paris), des AMAP (Associations pour le Maintien de l'Agriculture Paysanne), des initiatives de jardins partagés comme les « Incroyables comestibles » (Bouré, 2017). La production en pair-à-pair résiste dans des **niches** comme les *makerspaces*, les *repair cafés* et les *fablabs* (Berrebi-Hoffmann et al., 2018). D'autres initiatives voient la participation active des collectivités locales, notamment les « *Sharing cities* », les villes du partage, qui permettent à l'échelle locale d'expérimenter de nouvelles solutions dans la mobilité, le logement, la prise en charge des enfants ou des personnes âgées dans le voisinage, sous l'exemple de villes pionnières comme Séoul, Barcelone et Amsterdam (McLaren & Agyeman, 2015).

Loin d'être exhaustive, cette liste suggère que toute la volonté de renouvellement social des premiers porteurs des principes du partage ne s'est pas effacée. Le champ s'est restructuré sans totalement expulser sa dimension la plus idéaliste, mais en confiant sa gestion à des organisations autres que les grandes plateformes désormais familières à tout un chacun. Pour résister, elles en sont venues à occuper des positions marginales. La société des plateformes n'arrive pas mieux que la société organisationnelle à concilier objectif de profit d'une part, et objectifs sociaux, culturels et environnementaux de l'autre : ils coexistent, le premier au centre, les seconds en périphérie.

4.2 *Décentralisation et coopération*

Notons avant de conclure une spécificité de la diffusion de ces mouvances en France, où la terminologie du partage a été toujours un peu en retrait par rapport à la désignation « économie collaborative » ou « société collaborative », en partie sous l'influence du groupe OuiShare. Selon les dires de ses fondateurs, cette interprétation faisait référence à « des systèmes qui ne sont fondés ni sur la hiérarchie, ni sur la compétition », éloignés d'une économie « qui n'est après tout que l'agrégat de structures organisées verticalement entrant en compétition entre elles pour l'accès aux ressources » (De Grave, 2016). Dans le livre-manifeste du groupe, la « société collaborative » est définie comme « la fin des hiérarchies » (Filippova, 2016). Ces prises de position qui se réfèrent indirectement, quoique un peu naïvement, à la tension entreprise/marché (section 2), révèlent tout de même une grande ambition d'ouvrir la voie à des modèles égalitaires de création et de répartition de la valeur.

La réflexion sur la coordination économique et la recherche de formes plus poussées de décentralisation et horizontalité continuent aujourd'hui. Ces questions sont au cœur des expérimentations non seulement de OuiShare, mais aussi d'autres petites organisations engagées comme le collectif néo-zélandais Enspiral, qui est allé très loin dans le développement de technologies pour la démocratie directe dans la décision collective. Se développe en même temps le mouvement du coopératisme des plateformes (Scholz, 2016; Scholz & Schneider, 2017), qui sans vouloir éliminer le modèle de la plateforme dans son ensemble, met en cause son orientation au service du profit, dans la tradition du mouvement coopératif qui se posait en alternative à l'entreprise tout en coexistant avec elle.

Encore une fois, résistent des initiatives et actions qui ont pu se positionner sur des niches. La prise en compte des coûts de synchronisation (section 3) aide à comprendre les raisons de cette dynamique de marginalisation des versions plus idéalisées de l'économie des plateformes. C'est la participation des individus à des organisations, plateformes ou autres entités collectives, porteuses de valeurs, qui engendre une dynamique d'interactions résultant en une distribution de ressources et de pouvoir. Les coûts de synchronisation ont été moins élevés pour les plateformes (et leurs gérants) quand elles ont épousé des logiques de profit, par rapport à celles qui ont milité pour des logiques et des valeurs alternatives. La temporalité plus longue qu'exigerait un changement plus profond des valeurs, a défavorisé ces dernières à partir du moment où elles se sont vues concurrencées par des rivales acceptant d'adopter des logiques commerciales.

L'avenir des objectifs d'innovation sociale outre que technologique et économique, restés à la périphérie, est ouvert. La lecture proposée ici amène à voir leur évolution potentielle comme un **terrain de lutte sur les utilisations possibles de la technologie** : les valeurs qui s'affrontent et leurs chances d'affirmation éventuelles ne peuvent être comprises qu'en lien avec les acteurs (individuels et collectifs) qui en sont porteurs, leurs liens avec les plateformes les plus centrales et l'évolution des modèles d'affaires de celles-ci. L'enjeu pour la recherche est de saisir le positionnement des acteurs dans le système des valeurs et leurs interdépendances, afin

de faire apparaître les processus dynamiques qui soutiennent certaines valeurs et en délaissent d'autres, produisant différents scénarios pour l'avenir.

5 Plateformes, travail et production de données

L'économie des plateformes nécessite des données, tout comme l'économie industrielle du XXe siècle exigeait de l'énergie. Les données constituent l'atout informationnel qui permet aux plateformes de réduire les coûts de transaction et de s'imposer comme intermédiaires entre différentes parties prenantes (section 1). Elles en constituent donc la ressource économique primaire. On voit émerger autour des données une reconfiguration des systèmes de production et d'accumulation tout autant que des relations sociales. Pour appréhender ces transformations, il convient alors de reprendre au sujet des données numériques le questionnement que Desrosières (2005, p. 6) avait appliqué aux statistiques et à la quantification du social, considérant sans jamais les séparer « leurs apports de connaissance et les circuits sociaux de leur mise en forme et de leurs usages ». Les données sont à la fois partie de la boîte à outils des chercheur.e.s, et leur **objet d'étude substantif** dans la mesure où elles transforment la société qui les entoure.

Kitchin (2014) nous rappelle que loin du sens littéral du mot, qui fait référence à ce qui nous aurait été donné, qui se trouverait (pour ainsi dire) dans la nature et ne devrait qu'être saisi et utilisé, **la donnée est produite**. Comme une enquête statistique est issue du travail de ses concepteurs, enquêteurs et codeurs, soutenus par des investissements institutionnels parfois conséquents, les données numériques n'existeraient pas sans une importante activité humaine et sans la direction ou coordination centralisée qui la met en œuvre et l'accompagne (Bastin & Tubaro, 2018, p. 390). Il faut alors regarder de plus près les conditions de production et d'usage des données, en considérant d'abord la codification, la formalisation de toute action ou interaction médiatisée par ordinateur, débouchant l'émergence d'un véritable marché du travail en ligne lié à la donnée.

5.1 Formaliser l'informel pour sa mise en données

L'avènement des technologies numériques a ravivé une réflexion déjà ancienne qui opposait le formel et l'informel comme deux modalités différentes de l'action, individuelle et collective. Les premières recherches sur la socialisation en ligne mettaient l'accent sur une participation à l'apparence large, libre et ouverte, sur la spontanéité des interactions et des discussions, sur l'anonymat ou le pseudonymat. On a pu parler d'expert-amateur (Flichy, 2010), d'influenceur, de micro-célébrité (Senft, 2008) au sujet des personnes dont le statut officiel ne faisait pas entrevoir le succès qu'ils allaient obtenir en ligne. Selon Lance Bennett & Segerberg (2013), l'« action connective » spontanée rendue possible par les technologies numériques dépasserait

désormais la capacité des organisations traditionnelles (l'« action collective » classique). Ces aspects ne pouvaient que séduire les premiers promoteurs de la société qu'ils appelaient collaborative (section 4), intéressés à la flexibilité, au changement et à la mobilité. L'opposition du collectif OuiShare à la « hiérarchie » (sous-section 4.2) peut s'interpréter aussi comme un rejet de tout ce qui est officiel, explicite, écrit, statué dans des règles et par là, rigide et susceptible de freiner le dynamisme d'une organisation.

La recherche sur les organisations avait déjà résolu cette tension dans des contextes plus classiques. Depuis l'école des « Relations Humaines » (Mayo, 1949), on connaissait l'existence d'une organisation **informelle** (officieuse, non reconnue) avec des règles implicites, une connaissance partagée « tacite », une culture et même des rôles sociaux qui peuvent avoir émergé spontanément par l'interaction des acteurs. De fait, les structures formelles et informelles définissent simultanément et concurremment les règles d'action et le pouvoir réel des membres (Crozier & Friedberg, 2014; Reynaud, 1988). Ainsi, la prise en compte des structures formelles ne peut pas perdre de vue ce qui relève de formes organisationnelles plus floues ou en voie de constitution.

Les actions et interactions médiatisées par ordinateur ne sont pas pour autant de même nature que les actions informelles en présentiel et, malgré leur spontanéité apparente, laissent réapparaître des éléments de formalisme. Toute action en ligne – aussi petite soit-elle, y compris un simple clic ou un « j'aime » – laisse une trace numérique (Merzeau, 2009) qui alimente des bases de données : elle est codifiée, enregistrée, mémorisée et mise en relation avec d'autres traces/données. Malgré des interfaces faciles d'usage et même ludiques, l'utilisateur est de plus en plus directement en prise avec les systèmes de données, dont il est consommateur tout en les alimentant par ses propres actions (Liu, 2004). Aussi, ces données sont agrégées dans des profils personnels de plus en plus riches et détaillés – d'autant plus que l'anonymat et le pseudonymat de l'internet des origines se sont largement perdus (Boyd, 2012). Les usagers le savent et y font de plus en plus attention. Les pages sur lesquelles ils ont cliqué, les contenus qu'ils ont « aimé », influenceront l'offre de contenus et de pages qui leur seront proposés à l'avenir. Leur présence web sera d'autant plus conforme aux règles et aux attendus des fournisseurs de services qu'ils en anticipent les conséquences, dans une logique de mise en visibilité de soi très contrôlée (Casilli, 2010). On ne saurait retrouver ici l'invisibilité traditionnellement attribuée aux processus sociaux informels.

Plus les actions obéissent à des codes et des règles formelles simples, plus elles peuvent être intégrées aisément dans des bases de données. Ajouter quelqu'un sur Facebook ou LinkedIn exige une interaction codifiée, qui n'a guère d'équivalent dans la communication en face-à-face, avec une requête explicite à laquelle le destinataire doit répondre positivement. Le choix d'établir un lien est enregistré dans les profils des deux personnes concernées, de sorte qu'aucune divergence d'interprétation n'est plus possible. Il s'agit d'un acte quasiment public, visible et traçable (Casilli, 2010, p. 272), sans rapport avec les contenus de la relation, et encore moins avec les sentiments ou la proximité émotionnelle.

En ce sens, les « réseaux sociaux » en ligne ne représentent guère les mêmes relations que l'amitié ou la collaboration au sens traditionnel de ces termes – même s'il est désormais admis que la sociabilité en ligne s'articule avec la sociabilité humaine au sens plus général (Wellman et al., 2001; Wellman, 2001). Les relations qu'enregistrent les données numériques reflètent les principes et les priorités des plateformes, et conservent la forme et la structure que celles-ci imposent. Les plateformes et leur conception technique établissent (par exemple) ce que veut dire être en contact, ou si la réciprocité est nécessaire pour valider un lien. Ces considérations invitent à la prudence lors de toute tentative de généraliser la connaissance des comportements numériques à d'autres univers sociaux (ou inversement).

En somme, **la numérisation de pans entiers de la vie économique et sociale entraîne ce retour du formalisme, essentiel pour la mise en données de l'expérience des usagers.** Le quotidien devient calculable, et mis au service de l'économie de la donnée dans laquelle prospère la plateforme (Alaimo & Kallinikos, 2017). Les relations en ligne s'insèrent dans un ensemble de médiations techniques et de contextes et modes d'usage qui leur sont propres. Elles sont l'émanation des orientations politiques, ontologiques et épistémologiques ainsi que des intérêts économiques de la plateforme.

5.2 Un marché du travail de la donnée

La codification et la formalisation de toute action ou interaction médiatisée par ordinateur, sont les conditions de production de la donnée numérique. Il s'agit parfois d'instruire et de diriger les usagers pour que des données exploitables puissent être extraites de leurs pratiques. Par exemple, TripAdvisor incite les vacanciers à formaliser leur expérience en notant un hôtel sur une échelle donnée, et se sert de cette note pour ajuster la position de l'hôtel dans son classement, et/ou pour décider des publicités à lui associer. von Ahn et al. (2008) ont révélé que le bien connu « re-Captcha » que le système d'identification de Google utilisait pour distinguer un usager humain d'un bot, servait également à classer et interpréter des données visuelles lorsque les algorithmes de reconnaissance d'images n'arrivaient pas à reconnaître (par exemple) des noms de rues dans des photographies de trafic urbain.

Dans d'autres cas, la plateforme paie des travailleurs pour leur faire produire ou enrichir les données dont elle a besoin. Il s'agit par exemple de se faire aider à catégoriser les photos des logements que les hôtes mettent sur leur profil Airbnb, lorsque les légendes rajoutées par ceux-ci sont insuffisantes ou imprécises (Yao et al., 2018). Il faut alors que des fournisseurs extérieurs classent ces images correctement et y rajoutent les étiquettes nécessaires, par exemple pour distinguer des pièces (cuisine, salle à manger etc.). La commande est segmentée en un grand nombre de petites tâches rapides et fortement standardisées, pour qu'elles soient exécutées par des prestataires payés à la pièce à l'aide de plateformes de « micro-travail » comme Amazon Mechanical Turk, Appen, Clickworker, Lionbridge et Microworkers. Chaque prestataire va classer et étiqueter un petit nombre de ces images, mais

collectivement, ils peuvent compléter l'ensemble de la commande en un temps très court. Le processus de formalisation atteint son comble avec la fragmentation et la standardisation de ces tâches, qui en permettent la réalisation rapide, efficace et peu coûteuse.

Il s'agit là d'un véritable processus de production à grande échelle, comparable à ce que les organisations déployaient à l'ère industrielle, mettant au travail des masses pour satisfaire les besoins en données de ces plateformes. Leur importance est cruciale, car ces annotations servent à préparer les données pour qu'elles puissent être traitées par des méthodes d'apprentissage statistique (*machine learning*), à la base des technologies actuelles de l'intelligence artificielle. Même si les entreprises essaient d'économiser en automatisant une partie du travail de production de données, une composante humaine est toujours présente : on parle en informatique de modèles avec « **l'humain dans la boucle** » (*human-in-the-loop*) (Schmidt, 2019; Tubaro & Casilli, 2019).

La production des données qui alimentent l'activité des plateformes est **le résultat d'un système complexe de rapports socio-économiques fortement asymétriques** (Ekbia & Nardi, 2017). La faible rémunération des travailleurs, leur manque de contrôle sur les conditions de travail, la difficulté à faire remonter des soucis (Casilli, Tubaro, Le Ludec, Coville, Besenval, Mouhtare & Wahal, 2019; Gray & Suri, 2019; Roberts, 2019), sont des problèmes particulièrement aigus surtout lorsque l'exécution des tâches est délocalisée dans des pays pauvres ou émergents, chez des travailleurs quasi-captifs qui manquent le plus souvent de sources alternatives de revenu (Graham et al., 2017; Graham & Shaw, 2017; Graham & Anwar, 2019). D'autres enjeux sont de nature éthique et politique, notamment en termes des violations des règles de protection de la vie privée lorsque les données traitées par ces travailleurs extérieurs constituent des informations personnelles voire sensibles (Casilli, Tubaro, Le Ludec, Coville, Besenval, Mouhtare & Wahal, 2019).

On peut sourire à l'idée que les technologies numériques, initialement saluées comme des outils d'interaction spontanée, s'avèrent productrices de formalismes qui chiffrent et enregistrent des aspects de plus en plus étendus du quotidien. En même temps, les promoteurs de la société « collaborative », qui se battent contre le formalisme et les règles, empruntent ces mêmes technologies pour faire avancer leurs projets d'organisation décentralisée. Il peut également paraître curieux que dans le langage courant, la locution « réseaux sociaux » en soit venue à désigner quasi-exclusivement les services de réseautage en ligne, malgré leurs spécificités par rapport à d'autres types de relations humaines, et malgré leur dépendance des modèles d'affaires des plateformes. Il est également paradoxal que l'essor actuel de l'intelligence artificielle s'appuie sur les « petites mains » d'humains qui, derrière leurs écrans d'ordinateur ou de *smartphone*, réalisent des tâches de la données sans lesquelles les puissants algorithmes de *machine learning* seraient défaillants. Ce sont les contradictions de la nouvelle économie de la donnée, qui pourtant fournit aux plateformes les ressources informationnelles par lesquelles elles baissent les coûts de transaction et s'imposent dans le système économique actuel.

Les voies qui s'ouvrent à la recherche future sont nombreuses, beaucoup de questions se posent autour de la production de données et des marchés mondialisés où elle puise ses inputs productifs : le marché du micro-travail, bien sûr, mais aussi les marchés de l'énergie et des matières premières comme les métaux rares (Crawford & Joler, 2018), nécessaires à construire et maintenir l'infrastructure *hardware* et *software* des plateformes.

6 Contenus du mémoire original d'HDR

La vision qui vient d'être proposée prend les plateformes numériques comme point d'entrée pour revenir sur des grandes questions qui ont traversé l'histoire des sciences sociales : la variété des formes de coordination entre les deux pôles du marché et de l'entreprise, la tension entre logiques d'échange marchandes et non marchandes, la place des structures informelles par rapport au formel. J'ai visé à construire un cadre analytique cohérent, prenant appui sur le paradigme de l'organisation, vue comme structure d'interaction où se mettent en œuvre des processus relationnels qui aident les acteurs à gérer leurs interdépendances, et véhiculent des comportements complexes, entre le marché et le don. En même temps, il a fallu prendre en compte les grandes masses de données (*big data*), dans leurs aspects économiques (extraction de valeur commerciale, constitution d'un marché globalisé de travail de la donnée) et sociaux (inégalités, brouillage de la frontière public-privé). Il s'agit d'un élément spécifique au fonctionnement de l'économie numérique, qui ne manque d'affecter le modèle d'affaires des plateformes tout autant que les retombées sociales.

L'objet complexe que constituent les plateformes numériques exige un développement unique d'articulation de la théorie, de la méthode et de terrains nouveaux. Résultant d'une réflexion au croisement entre la sociologie économique et la sociologie des organisations d'un côté, et l'économie des marchés multi-faces et la théorie des coûts de transaction de l'autre, sans oublier les apports de l'informatique et des sciences de gestion, les conclusions provisoires obtenues ci-dessus invitent à maintenir une perspective trans-disciplinaire. Aussi, investir les terrains nouveaux du numérique exige d'articuler les méthodes. L'objet même de l'investigation invite à l'utilisation des bases de données nativement numériques, qui constituent un ingrédient essentiel du fonctionnement des plateformes et qui documentent bien l'importance de l'information dans leur modèle économique. La capacité de recueillir et d'analyser des quantités massives de données reste en retrait dans les sciences sociales par rapport à d'autres domaines comme la biologie et la physique, car ces données sont souvent la propriété privée des plateformes (et parfois, d'administrations publiques tout autant puissantes) et ne sont pas toujours accessibles à la recherche publique. À l'appel pour une *computational social science* ouverte et accessible (Lazer et al., 2009), qui tarde à se concrétiser, doit alors se rajouter un effort systématique de continuer à mobiliser les méthodes issues de la tradition de recherche sociologique, notamment les enquêtes, qu'elles soient quantitatives ou qualitatives

(Bastin & Tubaro, 2018). Pour que la recherche publique ne soit pas anéantie par sa pauvreté en données relativement aux grandes plateformes, elle doit savoir extraire un maximum de connaissances des sources qui lui sont toujours accessibles.

Dans ce contexte, l'analyse des réseaux sociaux peut s'avérer d'un grand secours. Comme dans la sociologie des organisations, elle peut aider à faire apparaître les interdépendances entre acteurs à tous les niveaux, et à en dégager les effets sur les décisions, l'accès aux ressources, et la performance. Avec l'analyse des réseaux, les mêmes métriques et visualisations peuvent aider à représenter les structures sociales formelles comme l'organigramme, les relations informelles comme l'amitié, et les relations médiatisées par ordinateur, offrant par là une possibilité de comparaison. Le sens de ces interdépendances n'est toutefois pas le même dans ces différents contextes, et il faut savoir faire des distinctions – quitte à renoncer à la représentation réticulaire lorsqu'elle risque de nous fourvoyer.

Les considérations qui précèdent ouvrent la voie à des analyses, conceptuelles ainsi qu'empiriques, qui en examinent de manière plus détaillée les présupposés et les implications. La suite de ce mémoire original d'HDR consiste en un recueil de six articles inédits, en langue anglaise, qui mettent en œuvre ces idées en les confrontant avec des données.

1. L'article « **Disembedded or deeply embedded? Towards a multi-level network analysis of the platform economy** » reprend la question de l'identité de la plateforme entre les deux pôles du marché et de l'entreprise (section 2), à la lumière de l'un des concepts clé de la « nouvelle sociologie économique », à savoir l'encastrement (*embeddedness*). En m'appuyant sur cette littérature, je propose la notion de « deep embeddedness » pour décrire l'encastrement dans les structures économiques complexes que la plateforme crée autour d'elle. Cet article, dont je présente ici une toute première version, est essentiellement conceptuel, tout en utilisant quelques données empiriques issues du projet *Digital Platform Labor*, DiPLab (présenté en Annexe).
2. L'article « **Festival attendance as a network of attention : A case study in the sharing economy** » développe l'idée de mettre en relation les valeurs de l'économie « de partage » ou « collaborative » avec les acteurs (individuels et collectifs) qui en sont porteurs, et les relations entre eux, pour saisir les évolutions possibles de ce que j'ai appelé un terrain de lutte (sous-section 4.2). La dynamique que je mets en exergue est plus générale et peut investir d'autres terrains d'interactions comme les événements professionnels qui réunissent périodiquement les représentants d'un secteur ou champ économique. Cet article, qui reprend et re-élabore des idées présentées lors de la conférence *Complex Networks* de 2018 (Tubaro, 2018a), est issu du projet *Sharing Networks* (présenté en Annexe).
3. L'article « **Whose results are these anyway? Reciprocity and the ethics of “giving back” after social network research** » est issu d'une réflexion sur la manière dont l'avènement des plateformes, et plus généralement du numérique (les spécificités des réseaux en ligne, sous-section 5.1, les menaces sur la vie privée, sous-section 3.1, les marchés du travail des données, sous-section 5.2), affecte la responsabilité des chercheur.e.s, en l'occurrence lors de la restitution

des résultats. Issu du projet *Recent ethical challenges in social-network analysis* (présenté en Annexe), cet article est le plus avancé du recueil, étant en révision pour un numéro spécial, que je co-coordonne, de la revue *Social Networks*. La version jointe est la révision de l'article soumise à la revue en juillet 2019, après une demande de modifications mineures.

4. L'article « **The trainer, the verifier and the imitator : Three ways in which human platform workers support automation** » approfondit la question du travail de la donnée (sous-section 5.2) pour montrer qu'il ne s'agit pas d'un phénomène temporaire mais structurel, lié à la nature de la recherche actuelle en intelligence artificielle, basée sur l'apprentissage statistique (*machine learning*). Il est co-écrit avec A.A. Casilli et M. Coville, et issu du projet *Digital Platform Labor*, DiPLab (Annexe). Une version préliminaire a été présentée en mai 2019 à la journée de recherche du centre S&O à HEC Paris.
5. L'article « **Counting “micro-workers” : Methodological and societal challenges around new forms of labor** » développe davantage les questions du micro-travail (sous-section 5.2) et de la donnée (section 5) pour discuter comment le travail des plateformes est particulièrement rebelle aux statistiques et demande à revoir nos approches à la quantification. Quelque peu paradoxalement, l'essor des données numériques va de pair avec de plus grandes difficultés à mettre en chiffre des dimensions pourtant essentielles de nos sociétés, notamment la population active sur les plateformes. Co-écrit avec C. Le Ludec et A.A. Casilli, cet article est issu du projet *Digital Platform Labor*, DiPLab (Annexe). Une version préliminaire a été mise en ligne en tant que document de travail au printemps 2019 (Le Ludec, Tubaro & Casilli, 2019a,b), nous permettant entre temps d'avoir des retours autant des collègues que d'acteurs de la politique publique, après avoir été beaucoup relayée par la presse. La version jointe a été profondément révisée pour en tenir compte.
6. Le dernier texte, intitulé « **Who are the micro-workers behind AI?** » est un chapitre pour un livre collectif qui sera intitulé « *Automating the Crowd. Real people behind artificial intelligence* », coordonné par K. Whitaker (Alan Turing Institute), A. Alexander (Université de Cambridge), A. Borda (Université de Melbourne) et R. Milne (Wellcome Genome Campus). S'agissant d'un livre grand public, le texte est plus court et plus descriptif que les autres. Il dépeint la population qui micro-travaille en France (sous-section 5.2), et explique en langage non technique pourquoi leur activité est importante dans l'actuelle économie des plateformes. Il reprend quelques-uns des messages essentiels du rapport conclusif du projet *Digital Platform Labor*, DiPLab (Casilli, Tubaro, Le Ludec, Coville, Besenval, Mouhtare & Wahal, 2019).

7 Annexe

Cet annexe résume brièvement les enjeux et les techniques mis en œuvre dans les projets de recherche auxquels sont liés les articles présentés dans ce mémoire original.

7.1 *Sharing Networks (SN)*

L'étude *Sharing Networks* reconstitue les réseaux sociaux dans le champ de l'économie collaborative dans son ensemble, à partir de la fréquentation d'un grand festival international, le OuiShare Fest, qui a eu lieu à Paris chaque année entre 2013 et 2017, avec une édition locale à Barcelone (aussi annuelle, 2015-17). Ce terrain que j'ai baptisé *Sharing Networks*, a été réalisé entre 2016 et 2017, et l'analyse des données est en cours.

Organisé par le *think-tank* OuiShare (Filippova, 2016), le OuiShare Fest se proposait de rassembler des acteurs divers, tous porteurs d'une vision de l'usage des technologies dans la société – que ce soit pour promouvoir la décentralisation grâce à la *blockchain*, les nouvelles formes de travail indépendant à l'heure du retrait du salariat, la connaissance libre et les *open data*, ou la démarche zéro-déchets. Chaque année, y ont participé plusieurs centaines de personnes provenant de différents pays. En cartographiant les relations entre participants, je visais à établir un état des lieux des valeurs mis en avant dans ce milieu, tout en les associant aux personnes et aux collectifs qui en sont les porteurs, et à suivre les débats et les interactions entre eux.

J'ai eu l'accord de OuiShare pour réaliser une **enquête par questionnaire lors des deux dernières éditions de l'événement parisien (mai 2016 et juillet 2017), et de sa variante locale, le OuiShare Fest Barcelona de novembre 2017**. Dans les trois cas, j'ai distribué un questionnaire aux participant.e.s avec trois « générateurs de noms », à savoir, des questions couramment utilisées dans la recherche sur les réseaux sociaux pour amener les enquêtés à nommer leurs contacts. Je leur ai demandé qui, parmi les autres participants à l'événement, ils/elles connaissaient déjà avant ; avec qui, parmi ces mêmes participants, ils/elles venaient de faire connaissance ; et qui, parmi les personnes déjà nommées, ils/elles souhaitaient recontacter dans l'avenir proche. Avec des qualifications pour chacun des contacts nommés (genre, pays d'origine, profession etc.), je peux établir l'évolution des réseaux personnels des participants avant, pendant et après l'occasion de rencontre qu'est le festival. En 2016, j'ai aussi interviewé en profondeur un sous-ensemble des répondants pour mieux connaître leurs motivations, leurs activités dans les plateformes et la place qu'elles ont dans leur environnement social.

J'ai eu aussi accès, par les organisateurs du OuiShare Fest, aux données nativement numériques issues de l'application logicielle (« Sched ») qu'ils proposaient aux participants pour choisir leur emploi du temps et pour se mettre en relation avec les personnes qu'ils rencontraient au festival. Par rapport à mes générateurs de noms, ces données renseignent sur un autre type de relations, celles qui sont pas-

sées par la voie numérique que représente l'application. Elles fournissent aussi des informations précieuses sur les intérêts thématiques des enquêtés.

Ce projet a été une importante expérience de collaboration avec des collègues, et d'encadrement d'étudiant.e.s et jeunes chercheur.e.s. Lors des éditions parisiennes du OuiShare Fest de 2016 et 2017, la collecte de données a été réalisée en collaboration avec A.A. Casilli (Telecom Paris) et L. Mounier (ancienne IR CNRS), bénéficiant en 2016 d'un financement de 5 000 euros par le Fonds Action Carnot. Pour l'édition espagnole de 2017, j'ai sollicité la collaboration de J.L. Molina (Universitat Autònoma de Barcelona). Des étudiant.e.s m'ont aidée pour la passation du questionnaire (N. Casati, G. Compain, T. Helsens, A. Wang à Paris, et K. Genna, A. Kussy et M. Lobato à Barcelone), et des stagiaires m'ont assistée dans le travail de codage et de mise en forme des données après collecte (S. Benmarhnia et S. Sinz, niveau M2; C. Oster et E. Wong, niveau L3). J'ai aussi encadré une étudiante de l'ENSAE, A. Papp, qui a réalisé son mémoire de M2 sur ces données.

L'analyse des données issues de ce projet n'est pas terminée, mais des résultats préliminaires ont été obtenus. C'est dans des événements comme le OuiShare Fest que **le milieu de l'économie collaborative, réuni temporairement en un lieu géographique précis, définit son identité** (Brailly et al., 2016), sa vision de l'avenir et ses valeurs (Anand & Watson, 2004). Cela ne se fait pas sans des tensions entre acteurs animés d'intérêts et d'objectifs divergents, et une double logique apparaît – d'une part, un alignement vers les proches, défenseurs des mêmes visions, et de l'autre, un conformisme qui conduit plutôt à se ranger du côté des acteurs au statut le plus élevé (Tubaro, 2018a). Le discours ambiant qui valorise l'horizontalité et les rapports "entre pairs" mitige les effets des facteurs qui tendraient à faire apparaître une hiérarchie informelle, et résulte en des différences de statut relativement faibles.

Le mémoire de A. Papp s'intéresse à la place des femmes dans ce milieu, à la fois dominé par une forte culture technologique généralement peu féminisée, et par un ensemble d'initiatives et d'objectifs sociétaux qui relèvent de la tradition de l'économie sociale et solidaire – celle-ci, beaucoup plus habituée à la présence féminine. Les données montrent que des inégalités préalables persistent, les femmes arrivant à l'événement avec moins de contacts (masculins autant que féminins) que les hommes; mais **le festival joue un rôle d'intégration pour les femmes**, plus présentes dans les réseaux informels des participants après qu'avant l'événement, et plus présentes parmi les contacts jugés durables. Plus nombreuses que les hommes à faire partie de la composante principale du réseau, elles arrivent à faire avancer des sujets qui leur tiennent à cœur, réussissant à contrecarrer d'autres tendances qui les auraient davantage marginalisées (Papp, 2019).

D'autres résultats, non encore publiés, montrent que le côté social de l'économie collaborative n'a pas disparu dans le temps, comme beaucoup de critiques le suggèrent, mais s'est déplacé vers d'autres acteurs et sujets. Un foisonnement d'initiatives, d'organisations et de personnes continuent l'expérimentation de formes alternatives de production et de consommation à la marge du système. Il ne s'agit pas seulement de petites organisations engagées, mais aussi, de façon croissante, de grandes entreprises traditionnelles, d'administrations publiques et surtout de collectivités locales (McLaren & Agyeman, 2015). Des initiatives à la faveur du travail,

du développement local et de l'environnement promettent de manière plus crédible, quoique moins visible par les politiques et le grand public, ce que la culture hautement technologique des débuts n'a pas réussi à livrer.

7.2 *Digital Platform Labor (DiPLab)*

Le micro-travail consiste en la réalisation de petites tâches courtes et standardisées, la plupart desquelles sont réalisées exclusivement en ligne (comme la reconnaissance d'objets dans des images, la traduction de phrases, la transcription d'enregistrements sonores), que des plateformes spécialisées allouent à des prestataires anonymes, en contrepartie d'une rémunération généralement très faible. La plateforme de micro-travail la plus connue est Amazon Mechanical Turk, mais il en existe beaucoup d'autres, autant à l'échelle internationale (comme l'australienne Appen, l'allemande Clickworker et les américaines Lionbridge et Microworkers) que nationale (Foule Factory en France). Le micro-travail est parfois considéré comme la forme la plus invisible, la plus « émietlée » du travail des plateformes, l'évolution jusqu'ici la plus extrême de la fragmentation de la chaîne de valeur. C'est un prisme qui permet de porter un regard plus général sur les transformations actuelles du travail et leurs aboutissements possibles.

Depuis l'été 2017 j'ai entrepris, en collaboration avec A.A. Casilli (Telecom Paris) le montage et la réalisation d'un projet de recherche visant à documenter la réalité du micro-travail en France et dans le monde francophone – jusqu'ici très peu connue dans la mesure où la littérature naissante sur le micro-travail s'est concentrée essentiellement sur les États-Unis et tout au plus sur les pays anglophones. Nous avons mené une enquête par questionnaire (mai - juillet 2018) auprès des micro-travailleurs de Foule Factory, la principale plateforme française, et réalisé des entretiens en profondeur avec des micro-travailleurs, des responsables de plateformes, des clients et d'autres parties prenantes (septembre 2017 - novembre 2018). Nous avons aussi procédé à une collecte de données numériques, en obtenant un jeu de données de la plateforme *Net Business Rating*, qui offre aux micro-travailleurs la possibilité d'évaluer les plateformes de micro-travail ainsi que les clients, et de se donner mutuellement des conseils. L'objectif est de faire ressortir les conditions de travail, les contextes de vie, les motivations de participation, et la position des micro-travailleurs dans la structure sociale.

Ce projet a bénéficié d'une articulation de subventions du syndicat Force Ouvrière, par une agence d'objectif de l'IRES, Institut de Recherches Économiques et Sociales (32 000 euros), de France Stratégie, un service du Premier Ministre (12 000 euros) et de la Maison des Sciences de l'Homme Paris-Saclay (24 000 euros), en plus d'une aide logistique de la plateforme française Foule Factory et du financement d'un stage de 4 mois par l'Inria. L'équipe du projet inclut des jeunes chercheur.e.s que j'ai co-encadré.e.s, M. Besenval (stagiaire M2), M. Coville (post-doctorante), C. Le Ludec (ingénieur d'études), T. Mouhtare (assistante

de recherche), N. Sapkota (stagiaire L3), M. Venkat (stagiaire M1), et E. Wahal (ingénieure d'études).

Ce projet a tenté une estimation du nombre de micro-travailleurs en France – qui échappent à la statistique publique et pour lesquels on manque donc de données fiables. S'agissant le plus souvent d'une activité secondaire, il a fallu distinguer leurs niveaux d'activité : **il y aurait environ 15 000 « très actifs », 50 000 « réguliers » et 260 000 « occasionnels »**. Il s'agit d'ordres de grandeur, obtenus en superposant des méthodes différentes et dont le but principal est d'orienter l'action, tout en posant une première pierre pour des recherches futures (Casilli, Tubaro, Le Ludec, Coville, Besenval, Mouhtare & Wahal, 2019).

Qui micro-travaille en France ? Ce ne sont pas seulement des jeunes passionnés de technologies, **la majorité de nos enquêtés ayant entre 25 et 44 ans** : le micro-travail est un complément de revenu pour des personnes en âge de travailler, dont la plupart ont terminé des études même longues. Ce n'est pas non plus un phénomène uniquement urbain, les personnes actives étant distribuées sur toute la France en fonction (à peu près) de la densité de population. Et il ne s'agit pas d'une activité principalement masculine, comme des enquêtes précédentes, réalisées dans d'autres contextes géographiques, auraient pu le suggérer (Berg et al., 2018). Nous constatons une **légère majorité de femmes**, dont l'investissement dans le micro-travail, assez important dans certains cas, suggère un glissement vers une « triple-journée » : à l'emploi principal, souvent à temps partiel, s'ajoutent les tâches ménagères et l'activité sur les plateformes. Surtout, nous observons **un taux d'inactivité très élevé par rapport à la population générale, et un nombre important de micro-travailleurs et micro-travailleuses vivant en dessous du seuil de pauvreté**, défini comme la moitié du revenu médian : quoique peu rémunérée, l'activité sur les plateformes est une manière de pallier au manque de pouvoir d'achat (Casilli, Tubaro, Le Ludec, Coville, Besenval, Mouhtare & Wahal, 2019).

Nous montrons aussi que le micro-travail joue un rôle clé dans l'essor de l'intelligence artificielle. Il contribue à **l'entraînement de l'intelligence artificielle** : la préparation, à la main, de données visant à entraîner les algorithmes de *machine learning*. Le micro-travail est nécessaire pour étiqueter des images, transcrire des mots, interpréter des bouts de conversation orale enregistrée par des assistants vocaux. Une autre fonction est la **vérification de l'intelligence artificielle** : il s'agit de s'assurer que des moteurs de recherche donnent les résultats espérés, ou encore que des logiciels reconnaissent correctement des textes ou des sons. La dernière fonction que nous avons identifiée est **l'imitation de l'intelligence artificielle** : l'intervention humaine pour pallier aux limites des algorithmes, pour les remplacer quand ils sont défaillants (Casilli, Tubaro, Le Ludec, Coville, Besenval, Mouhtare & Wahal, 2019; Tubaro & Casilli, 2019).

En résumé, le micro-travail se définit de plus en plus comme un nouveau moyen de remédier à une précarité économique. Dans le même temps, en raison d'une demande très fluctuante de la part des entreprises, il n'apporte pas de stabilité professionnelle, et ses impacts de long terme sur les trajectoires professionnelles des personnes concernées restent à établir. L'enjeu est donc de réguler cette nouvelle

force de travail et de renforcer sa protection sociale actuellement quasi-inexistante, tout en reconnaissant son apport à l'innovation technologique.

7.3 *Recent ethical challenges in social network analysis (RECSNA)*

Le mouvement général de **recul de la tolérance vis-à-vis des atteintes à la vie privée** des individus depuis quelques années, la **capture des données par des plateformes à des finalités commerciales**, et le flou légal sur nombre de questions liées aux données, posent de nouveaux défis éthiques et déontologiques aux chercheurs.e.s. Ces enjeux sont **particulièrement aigus en analyse des réseaux sociaux** car même des efforts récents pour établir une base légale au *scraping* de données ont surtout visé la fouille de textes ou la recherche de contenus textuels, régies par le droit d'auteur et le principe de la liberté d'expression. Les données relationnelles qui intéressent les sociologues des réseaux sociaux ne sont pas du texte et n'entrent donc pas dans ce cadre, de sorte que les conditions de leur usage restent souvent indéfinies.

Ces freins s'ajoutent à des difficultés déjà bien connues en analyse des réseaux sociaux, comme l'impossibilité d'anonymat dans la collecte des données, difficulté à recueillir le consentement des *alter* nommés par un *ego* dans une enquête avec générateurs de noms (Breiger, 2005). Il y a aussi des défis plus récents résultant de la diffusion des réseaux socio-numériques et de la **sur-valorisation de la connectivité**, qui peuvent gêner les enquêtés, susciter des réponses biaisées, ou rendre la restitution délicate dans la mesure où les résultats d'une recherche peuvent être mal interprétés.

Au titre de mon rôle de coordinatrice du groupe BSA-SNAG, j'ai souhaité lancer une réflexion à ce sujet. Avec A. D'Angelo (Université de Nottingham), L. Ryan (Université de Sheffield) et A.A. Casilli (Telecom Paris), j'ai organisé un atelier sur « *Recent Ethical Challenges in Social Network Analysis* » (RECSNA), qui s'est tenu à l'Institut d'Études Avancées de Paris en décembre 2017, avec un financement de la Maison des Sciences de l'Homme Paris-Saclay. En plus de la BSA, se sont associés à l'organisation de l'événement le groupe ENDL et le Réseau Thématique « Réseaux Sociaux » de l'Association Française de Sociologie (coordonné à l'époque par G. Favre, Université de Toulouse Jean Jaurès, et E. Penalva-Icher, Université Paris-Dauphine).

L'intérêt qu'a suscité cette thématique nous a ensuite amené à monter un projet de numéro spécial sur ce même sujet de la revue *Social Networks*, la plus importante dans le domaine de l'analyse des réseaux sociaux. Des 17 pré-propositions initialement reçues, 1 a été déjà acceptée et 8 sont en révision au moment où j'écris ces lignes. Il y aura aussi une longue introduction par les coordinateurs du numéro et une post-face du Prof. R. Breiger, coordinateur d'un précédent numéro de la même revue sur l'éthique de la recherche sur les réseaux sociaux (2005). La date de la publication définitive du numéro n'est pas encore établie (2020 ou 2021).

Références

- Abdelnour, S. & S. Bernard. 2018, «Vers un capitalisme de plateforme ? Mobiliser le travail, contourner les régulations», *La nouvelle revue du travail*, vol. 13.
- Aguilera, A., L. Dablanc & A. Rallet. 2018, «L'envers et l'endroit des plateformes de livraison instantanée : Enquête sur les livreurs micro-entrepreneurs à Paris», *Réseaux*, vol. 212, n° 6, p. 23–49.
- von Ahn, L., B. Maurer, C. McMillen, D. Abraham & M. Blum. 2008, «reCAPTCHA : Human-based character recognition via web security measures», *Science*, vol. 321, n° 5895, p. 1465–1468.
- Alaimo, C. & J. Kallinikos. 2017, «Computing the everyday : Social media as data platforms», *The Information Society*, vol. 33, n° 4, p. 175–191.
- Alchian, A. & H. Demsetz. 1972, «Production, information costs, and economic organization», *The American Economic Review*, vol. 62, n° 5, p. 777–795.
- Alter, N. 2009, *Donner et prendre. La coopération en entreprise*, La Découverte.
- Amar, N. & L. C. Viosat. 2016, «Les plateformes collaboratives, l'emploi et la protection sociale», Rapport IGAS n. 2015-121 R.
- Anand, N. & M. Watson. 2004, «Tournament rituals in the evolution of fields : the case of the Grammy Awards», *Academy of Management Journal*, vol. 47, n° 1, p. 59–80.
- Arrow, K. & F. Hahn. 1971, *General Competitive Analysis*, North Holland.
- Barbrook, R. 2000 (1999), «L'économie du don high-tech», dans *Libres enfants du savoir numérique. Une anthologie du libre*, édité par O. Blondeau & F. Latrive, Eclat, p. 143–157.
- Bastin, G. & P. Tubaro. 2018, «Le moment *big data* des sciences sociales», *Revue française de sociologie*, vol. 59, n° 3, p. 375–394.
- Bauwens, M. 2015, *Sauver le monde : Vers une économie post-capitaliste avec le peer-to-peer*, Les liens qui libèrent.
- Bauwens, M. & V. Kostakis. 2017, *Pour une véritable économie collaborative : Vers une société des communs*, Editions Charles Léopold Meyer.
- Bauwens, M. & R. Sussan. 2005, «Le peer to peer : nouvelle formation sociale, nouveau modèle civilisationnel», *Revue du MAUSS*, vol. 26, n° 2, p. 193–210.
- Benavent, C. 2016, *Plateformes*, FYP Editions.
- Benkler, Y. 2006, *The Wealth of Networks*, Yale University Press.
- Berg, J., M. Furrer, E. Harmon, U. Rani & M. Silberman. 2018, «Digital labour platforms and the future of work : Towards decent work in the online world», Rapport de l'OIT.
- Berrebi-Hoffmann, I., M.-C. Bureau & M. Lallement. 2018, *Makers. Enquête sur les laboratoires du changement social*, Seuil.
- Botsman, R. & R. Rogers. 2010, *What's Mine Is Yours : The Rise of Collaborative Consumption*, HarperBusiness.
- Bouré, M. 2017, «Réappropriation des systèmes alimentaires par les citoyens : une logique de communs urbains», *Netcom*, vol. 31, n° 1-2, p. 175–192.

- Boyd, D. 2006, «Friends, Friendsters, and MySpace top 8 : Writing community into being on social network sites», *First Monday*, vol. 11, n° 2. URL www.firstmonday.org/.
- Boyd, D. 2012, «The politics of "real names"», *Communications of the ACM*, vol. 55, n° 8, p. 29–31.
- Brailly, J., G. Favre, J. Chatellet & E. Lazega. 2016, «Embeddedness as a multilevel problem : A case study in economic sociology», *Social Networks*, vol. 44, p. 319–333.
- Breiger, R. 2005, «Introduction to special issue : ethical dilemmas in social network research», *Social Networks*, vol. 27, n° 2, p. 89 – 93.
- Breiger, R. L. 1974, «The duality of persons and groups», *Social Forces*, vol. 53, n° 2, p. 181–190.
- Brousseau, E. & T. Penard. 2007, «The economics of digital business models : A framework for analyzing the economics of platforms», *Review of Network Economics*, vol. 6, n° 2, p. 81–114.
- Brown, A., J. Fishenden, M. Thompson & W. Venters. 2017, «Appraising the impact and role of platform models and Government as a Platform (GaaP) in UK Government public service reform : Towards a Platform Assessment framework (PAF)», *Government Information Quarterly*, vol. 34, n° 2, p. 167–182.
- Cachon, G., K. Daniels & R. Lobel. 2017, «The role of surge pricing on a service platform with self-scheduling capacity», *Manufacturing Service Operations Management*, vol. 19, n° 3, p. 368–384.
- Caillé, A. 2000, *Anthropologie du don. Le tiers paradigme*, Desclée de Brouwer.
- Casilli, A. 2010, *Les liaisons numériques. Vers une nouvelle sociabilité ?*, Seuil.
- Casilli, A. 2019, *En attendant les robots. Enquête sur le travail du clic*, Seuil.
- Casilli, A., P. Tubaro, C. Le Ludec, M. Coville, M. Besenval, T. Mouhtare & E. Wahal. 2019, «Le micro-travail en France. Derrière l'automatisation, de nouvelles précarités au travail ?», Rapport final du projet Digital Platform Labor (DiPLab).
- Castillo, J., D. Knoepfle & G. Weyl. 2017, «Surge pricing solves the wild goose chase», dans *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, ACM, New York, NY, USA, p. 241–242.
- Coase, R. H. 1937, «The nature of the firm», *Economica*, vol. 4, n° 16, p. 386–405.
- Cochoy, F. & S. Dubuisson-Quellier. 2000, «Les professionnels du marché : vers une sociologie du travail marchand», *Sociologie du travail*, vol. 42, n° 3, p. 359–368.
- Codagnone, C., F. Abadie & F. Biagi. 2016, «The future of work in the "sharing economy". Market efficiency and equitable opportunities or unfair precarisation ?», Rapport pour la Commission Européenne.
- Coleman, J. 1974, *Power and the Organisation of Society*, Norton.
- Coleman, J. 1982, *The asymmetric society*, Syracuse University Press.
- Cramer, J. & A. Krueger. 2016, «Disruptive change in the taxi business : The case of Uber», *American Economic Review*, vol. 106, n° 5, p. 177–82.
- Crawford, K. & V. Joler. 2018, «Anatomy of an AI system. The Amazon Echo as an anatomical map of human labor, data and planetary resources», Anatomy of AI projet, url : <https://anatomyof.ai/>.
- Crozier, M. 1971, *Le phénomène bureaucratique*, Seuil.

- Crozier, M. & E. Friedberg. 2014, *L'acteur et le système : Les contraintes de l'action collective*, Seuil, 1ère édition 1977.
- De Grave, A. 2016, «L'économie collaborative, c'est fini», *OuiShare Magazine*.
- Debreu, G. 1959, *Theory of Value. An Axiomatic Analysis of Economic Equilibrium*, Yale University Press.
- Desrosières, A. 2005, «Décrire l'État ou explorer la société : les deux sources de la statistique publique», *Genèses*, vol. 58, n° 1, p. 4–27.
- Desrosières, A. 2008, *Gouverner par les nombres : L'Argument statistique II*, Presses des Mines.
- Doueihi, M. 2011, *Pour un humanisme numérique*, Seuil.
- Eisenmann, T. R., G. Parker & M. van Alstyne. 2006, «Strategies for two-sided markets», *Harvard Business Review*, vol. 84, n° 10.
- Ekbia, H. & B. Nardi. 2017, *Heteromation, and Other Stories of Computing and Capitalism*, MIT Press.
- Evans, D. S. 2011, *Platform Economics : Essays on Multi-Sided Businesses*, Competition Policy International.
- Evans, D. S. & R. Schmalensee. 2016, *Matchmakers : the New Economics of Multi-Sided Platforms*, Harvard Business Review Press.
- Filippova, D. 2016, *Société collaborative. La fin des hiérarchies*, Editions Rue de l'Echiquier.
- Fitzmaurice, C. J., I. Ladegaard, W. Attwood-Charles, M. Cansoy, L. B. Carfagna, J. B. Schor & R. Wengronowitz. 2018, «Domesticating the market : moral exchange and the sharing economy (online first)», *Socio-Economic Review*.
- Flichy, P. 2010, *Le Sacre de l'amateur. Sociologie des passions ordinaires à l'ère numérique*, Seuil.
- Fligstein, N. 2002, *The Architecture of Markets : An Economic Sociology of Twenty-First-Century Capitalist Societies*, Princeton University Press.
- Fontaine, L. 2014, *Le marché*, Gallimard.
- Friedberg, E. 1997, *Le pouvoir et la règle*, Seuil, 1ère édition 1993.
- Garcia-Parpet, M. 1986, «La construction sociale d'un marché parfait : Le marché au cadran de Fontaines-en-Sologne», *Actes de la recherche en sciences sociales*, vol. 65, n° 1, p. 2–13.
- Gawer, A. 2009, *Platforms, Markets and Innovation*, Edward Elgar.
- Gillespie, T. 2010, «The politics of "platforms"», *New Media & Society*, vol. 12, n° 3, p. 347–364.
- Graham, M. & M. Anwar. 2019, «The global gig economy : Towards a planetary labour market?», *First Monday*, vol. 24, n° 4.
- Graham, M., I. Hjorth & V. Lehdonvirta. 2017, «Digital labour and development : impacts of global digital labour platforms and the gig economy on worker livelihoods», *Transfer : European Review of Labour and Research*, vol. 23, n° 2, p. 135–162.
- Graham, M. & J. Shaw. 2017, *Towards a Fairer Gig Economy*, Meatspace Press.
- Gray, M. & S. Suri. 2019, *Ghost Work : How to Stop Silicon Valley from Building a New Global Underclass*, Houghton Mifflin Harcourt.

- Hagiu, A. & J. Wright. 2015, «Multi-sided platforms», *International Journal of Industrial Organization*, vol. 43, p. 162–174.
- Hall, J., C. Kendrick & C. Nosko. 2015, «The effects of Uber’s surge pricing : A case study», Working paper, University of Chicago Booth school of business.
- Hayek, F. 1945, «The use of knowledge in society», *American Economic Review*, vol. 35, n° 4, p. 519–530.
- Howcroft, D. & B. Bergvall-Kåreborn. 2019, «A typology of crowdwork platforms», *Work, Employment and Society*, vol. 33, n° 1, p. 21–38.
- INRS. 2018, «Plateformisation 2027. conséquences de l’ubérisation en santé et sécurité au travail», Rapport INRS.
- Karpik, L. 2000, «Le guide rouge Michelin», *Sociologie du Travail*, vol. 42, n° 3, p. 369–389.
- Kirman, A. 2001, «Market organization and individual behavior : Evidence from fish markets», dans *Networks and Markets*, édité par J. Rauch & A. Casella, Russell Sage Foundation, p. 155–196.
- Kitchin, R. 2014, *The Data Revolution : Big Data, Open Data, Data Infrastructures and Their Consequences*, Sage.
- Kollok, P. 1999, «The economics of online cooperation : gifts and public goods in cyberspace», dans *Communities in Cyberspace*, édité par M. Smith & P. Kollok, Routledge, p. 220–239.
- Lance Bennett, W. & A. Segerberg. 2013, *The Logic of Connective Action. Digital Media and the Personalization of Contentious Politics*, Cambridge University Press.
- Lazega, E. 2001, *The Collegial Phenomenon : The Social Mechanisms of Cooperation Among Peers in a Corporate Law Partnership*, Oxford University Press.
- Lazega, E. 2003, «Rationalité, discipline sociale et structure», *Revue française de sociologie*, vol. 44, n° 2, p. 305–329.
- Lazega, E. 2015, «Mobilités, turnover relationnel et coûts de synchronisation : Comprendre l’action collective par ses infrastructures relationnelles dynamiques et multi-niveaux», *L’Année sociologique*, vol. 65, n° 2, p. 391–424.
- Lazega, E. 2016, «Synchronization costs in the organizational society : Intermediary relational infrastructures in the dynamics of multilevel networks», dans *Multilevel Network Analysis for the Social Sciences : Theory, Methods and Applications*, édité par E. Lazega & T. Snijders, Springer, p. 47–77.
- Lazega, E. 2017, «Organized mobility and relational turnover as context for social mechanisms : A dynamic invariant at the heart of stability from movement», dans *Knowledge and Networks*, édité par J. Glückler, E. Lazega & I. Hammer, Springer.
- Lazega, E. & L. Mounier. 2002, «Interdependent entrepreneurs and the social discipline of their cooperation : The research program of structural economic sociology for a society of organizations», dans *Conventions and Structures in Economic Organization : Markets, Networks, and Hierarchies*, édité par O. Favereau & E. Lazega, Edward Elgar, p. 147–199.
- Lazer, D., A. Pentland, L. Adamic, S. Aral, A. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, T. Jebara, G. King, M. Macy, D. Roy

- & M. Van Alstyne. 2009, «Computational social science», *Science*, vol. 323, n° 5915, p. 721–723.
- Le Ludec, C., P. Tubaro & A. A. Casilli. 2019a, «Combien de personnes micro-travaillent en France ? Estimer l'ampleur d'une nouvelle forme de travail», Document de travail i3 19-SES-02.
- Le Ludec, C., P. Tubaro & A. A. Casilli. 2019b, «How many people microwork in France ? Estimating the size of a new labor force», arxiv 1901.03889.
- Lehdonvirta, V. 2018, «Flexibility in the gig economy : managing time on three online piecework platforms», *New Technology, Work and Employment*, vol. 33, n° 1, p. 13–29.
- Lehdonvirta, V., O. Kässi, I. Hjorth, H. Barnard & M. Graham. 2019, «The global platform economy : A new offshoring institution enabling emerging-economy microproviders», *Journal of Management*, vol. 45, n° 2, p. 567–599.
- Liu, A. 2004, «Transcendental data : Toward a cultural history and aesthetics of the new encoded discourse», *Critical Inquiry*, vol. 31, p. 49–84.
- Mayo, E. 1949, *The Social Problems of an Industrial Civilization*, Routledge.
- McLaren, D. & J. Agyeman. 2015, *Sharing Cities. A Case for Truly Smart and Sustainable Cities*, MIT Press.
- Merzeau, L. 2009, «Du signe à la trace : l'information sur mesure», *Hermès*, vol. 53, p. 21–29.
- O'Reilly, T. 2011, «Government as a platform», *Innovations : Technology, Governance, Globalization*, vol. 6, n° 1, p. 13–40.
- Ouchi, W. 1979, «Compte rendu de *Markets and Hierarchies : Analysis and Antitrust Implications*», *Administrative Science Quarterly*, vol. 22, n° 3, p. 540–544.
- Papp, A. 2019, *L'influence du genre sur la négociation d'une définition de l'économie collaborative*, mémoire de master, ENSAE.
- Peart, S. & D. Levy. 2005, *The "Vanity Of The Philosopher" : From Equality To Hierarchy In Postclassical Economics*, The University of Michigan Press.
- Perrow, C. 1991, «A society of organizations», *Theory and society*, vol. 20, p. 725–762.
- Presthus, R. 1962, *The Organizational Society*, Knopf.
- Ravenelle, A. 2019, *Hustle and Gig : Struggling and Surviving in the Sharing Economy*, University of California Press.
- de Reuver, M., C. Sørensen & R. C. Basole. 2018, «The digital platform : a research agenda», *Journal of Information Technology*, vol. 33, n° 2, p. 124–135.
- Reynaud, J. 1988, «Les régulations dans les organisations : régulation de contrôle et régulation autonome», *Revue française de sociologie*, vol. 29, n° 1, p. 5–18.
- Reynaud, J. 1997, *Les règles du jeu. L'action collective et la régulation sociale*, A. Colin, 1ère édition 1989.
- Rheingold, H. 2000, *The Virtual Community : Homesteading on the Electronic Frontier*, MIT Press.
- Rifkin, J. 2001, *The Age of Access : The New Culture of Hypercapitalism, Where all of Life is a Paid-For Experience*, Tarcher.
- Roberts, S. 2019, *Behind the Screen : Content Moderation in the Shadows of Social Media*, Yale University Press.

- Rochet, J.-C. & J. Tirole. 2003, «Platform competition in two-sided markets», *Journal of the European Economic Association*, vol. 1, n° 4, p. 990–1029.
- Rosenblat, A. 2018, *Uberland : How Algorithms Are Rewriting the Rules of Work*, University of California Press.
- Rosenblat, A. & L. Stark. 2016, «Algorithmic labor and information asymmetries : A case study of Uber’s drivers», *International Journal of Communication*, vol. 10, n° 0.
- Roth, A. E. 2002, «The economist as engineer : Game theory, experimentation, and computation as tools for design economics», *Econometrica*, vol. 70, n° 4, p. 1341–1378.
- Rysman, M. 2009, «The economics of two-sided markets», *Journal of Economic Perspectives*, vol. 23, n° 3, p. 125–43.
- Schmidt, C. 2017, «Digital labour markets in the platform economy : Mapping the political challenges of crowd work and gig work», Report, Friedrich-Ebert-Stiftung.
- Schmidt, F. 2019, «Crowdproduktion von Trainingsdaten : Zur Rolle von Online-Arbeit beim Trainieren autonomer Fahrzeuge», Report, Hans-Böckler-Stiftung.
- Scholz, T. 2016, *Platform Cooperativism : Challenging the Sharing Economy*, Rosa Luxemburg Stiftung.
- Scholz, T. & N. Schneider. 2017, *Ours to Hack and to Own : The Rise of Platform Cooperativism : A New Vision for the Future of Work and a Fairer Internet*, OR Books.
- Schor, J. 2017, «Does the sharing economy increase inequality within the eighty percent? Findings from a qualitative study of platform providers», *Cambridge Journal of Regions, Economy and Society*, vol. 10, n° 2, p. 263–279.
- Schor, J. 2018, «Dependence and precarity in the platform economy», Keynote at Reshaping Work 2018 Conference, Amsterdam, 25-26 October 2018.
- Segrestin, B. & A. Hatchuel. 2012, *Refonder l’entreprise*, Seuil.
- Selznick, P. 1949, *TVA and the Grass Roots : A Study in the Sociology of Formal Organization*, University of California Press.
- Senft, T. 2008, *Camgirls : Celebrity and community in the age of social networks*, Peter Lang.
- Simon, H. 1991, «Organizations and markets», *Journal of Economic Perspectives*, vol. 5, n° 2, p. 25–44.
- Srnicek, N. 2016, *Platform Capitalism*, Polity Press.
- Steiner, P. 2017, «Economy as matching», Hal working paper.
- Stinchcombe, A. 2001, *When Formality Works*, The University of Chicago Press.
- Strubell, E., A. Ganesh & A. McCallum. 2019, «Energy and policy considerations for deep learning in NLP», *CoRR*, vol. abs/1906.02243.
- Sundararajan, A. 2016, *The Sharing Economy : The End of Employment and the Rise of Crowd-based Capitalism*, MIT Press.
- Tilly, C. 1998, *Durable inequality*, University of California Press.
- Tirole, J. 2016, *Économie du bien commun*, PUF.
- Tricou, F. 2008, *La loi de l’offre et de la demande : une enquête sur le libéralisme économique*, Presses Universitaires du Septentrion.

- Tubaro, P. 2018a, «Emergent relational structures at a “sharing economy” festival», dans *Proceedings of the 7th International Conference on Complex Networks and Their Applications*, vol. 2, édité par L. Aiello, C. Cherifi, H. Cherifi, R. Lambiotte, P. Lió & L. Rocha, p. 559–571.
- Tubaro, P. 2018b, «La vie privée, un bien commun ?», *Regards croisés sur l'économie*, vol. 23, n° 2, p. 129–137.
- Tubaro, P. & A. Casilli. 2019, «Micro-work, artificial intelligence and the automotive industry», *Journal of Industrial and Business Economics*, vol. 46, n° 3, p. 333–345.
- Tubaro, P., A. Casilli & Y. Sarabi. 2014, *Against the Hypothesis of the End of Privacy : An Agent-Based Modelling Approach to Social Media*, Springer.
- Weber, M. 1922, *Wirtschaft und Gesellschaft : Grundriss der verstehenden Soziologie*, Mohr.
- Wellman, B. 2001, «Computer networks as social networks», *Science*, vol. 293, n° 5537, p. 2031–2034.
- Wellman, B. & M. Gulia. 1999, «Net surfers don't ride alone : Virtual communities as communities», dans *Networks in the Global Village*, édité par B. Wellman, Westview Press, p. 167–194.
- Wellman, B., A. Quan-Haase, J. Witte & K. Hampton. 2001, «Does the internet increase, decrease, or supplement social capital? Social networks, participation and community commitment», *American Behavioral Scientist*, vol. 45, n° 3, p. 436–455.
- Williamson, O. 1975, *Markets and Hierarchies : Analysis and Antitrust Implications*, Free Press.
- Williamson, O. 1981, «The economics of organization : The transaction cost approach», *American Journal of Sociology*, vol. 87, n° 3, p. 548–577.
- Yao, S., Q. Zhu & P. Siclait. 2018, «Categorizing listing photos at Airbnb», *Medium*.
- Zha, L., Y. Yin & Y. Du. 2017, «Surge pricing and labor supply in the ride-sourcing market», *Transportation Research Procedia*, vol. 23, p. 2–21.
- Zuboff, S. 2019, *The Age of Surveillance Capitalism : The Fight for a Human Future at the New Frontier of Power*, Profile Books.

Disembedded or deeply embedded? Toward a multi-level network analysis approach of the online platform economy

Paola Tubaro

Abstract

This paper extends the economic-sociological concept of embeddedness to encompass not only social networks of, for example, friendship or kinship ties, but also economic networks of ownership and control relationships. Applying these ideas to a case study of digital platform labor pinpoints two possible scenarios. When platforms take the role of market intermediaries, economic ties are thin and workers are left to their own devices, in a form of “disembeddedness”. When platforms are parts of intricate inter-firm outsourcing structures, however, economic ties envelop workers in a “deep embeddedness” which involves both stronger constraints and higher rewards. The discovery of deep embeddedness brings to light the far-reaching social implications of platform business models that were previously under-researched. More generally, it enriches the notion of embeddedness and reinforces its capacity to account for actors’ position in today’s globalized, highly specialized production system.

Key words: Digital labor, embeddedness, outsourcing, multi-level network analysis, platform economy, transaction costs

1 Introduction: embeddedness in markets, firms and platforms

Embeddedness, or the idea that economic action is grounded in social structures, has been a founding block of economic sociology for the past few decades. In a path-breaking article, Granovetter (1985) contended that people do not act as atoms in an impersonal market, nor do they passively internalize outside norms: rather, their behaviors result from the history of their relationships and from their positions

Laboratoire de Recherche en Informatique, Centre National de la Recherche Scientifique, Paris Saclay, e-mail: paola.tubaro@lri.fr

with respect to other relationships. Granovetter explicitly suggested to operationalize embeddedness through *networks* of social relationships, although he did not use graph-theoretic formalisms in his article. Soon afterwards, advances in theory (for example, Burt (1992)) and analytical techniques (Wasserman and Faust, 1994) firmly established social network analysis at the heart of embeddedness scholarship. Since then, a rich empirical literature has shown that, for example, firms form partnerships based on their networks of prior alliances (Gulati and Gargiulo, 1999), that long-term personal relationships help companies to face the highs and lows of the market (Uzzi, 1996, 1997), and that friendship among managers dampens the worst effects of price competition (Ingram and Roberts, 2000).

This – still growing – current of research has deployed the construct of the network in manifold variants and combinations. But overall, network ties are of two main types: when internal to the boundaries of the firm, they serve to explain how this firm works – who gains knowledge or power internally, or how consensus forms around a decision. When they cross firm boundaries, ties are channels for market relationships – whether competitive or oligopolistic, global or local, trusted/repeated or “arm’s length” (Uzzi, 1997). Put differently, economic sociology implicitly inherits a vision *à la* Ronald Coase (1937) that identifies the firm and the market as two alternative, albeit possibly coexisting, devices for economic coordination.

Rooted in transaction cost economics (Coase, 1937; Williamson, 1975, 1981) the firm-market duality is admittedly not an indigenous concept to sociology. Rather than framing problems as referring to firms *vis-à-vis* markets, economic sociologists tend to use the language of organizations more generally: for example, they characterize networks as intra- or inter-organizational. Sociology and transaction cost theory converge in recognizing that economic activity requires some degree of organization to avoid chaos: there must be a system for matching actors, spreading information, arranging resources, dividing up the work. Even the market is not the impersonal, automated and cost-free mechanism that neoclassical economics imagined, but rather needs rules, intermediaries and support services (Favereau and Lazega, 2002). Against this common ground, there is still value in explicitly separating markets from firms, because they elicit different types of costs from actors, with distinct effects on their actions. In markets, there can be search costs, bargaining/decision costs, and policing/enforcement costs, while in firms, there are costs of planning, of assigning activities to individuals, of grouping activities into divisions, of controlling results and of delegating authority. Carruthers and Uzzi (2000) claim that within sociology, the study of the firm has started earlier, and progressed faster, than the study of the market.

This distinction becomes all the more relevant today, as the rise of digital platforms (Gawer, 2009; Srnicek, 2016) has been saluted by many as a third way in which coordination may occur, beyond the firm and the market (Sundararajan, 2016). While this characterization is subject to discussion, it nevertheless challenges the practices of economic sociologists and their underlying conceptualizations: can the notion of embeddedness and its common operationalization through social network analysis be easily transposed to the world of the digital platform? If not, what is that does not fit, and what aspects of our thinking should change?

These questions are obviously of practical importance for scholars of the platform economy, who may need to adapt existing theoretical constructs to the specificity of their object. They are also of broader theoretical importance insofar as they re-open the question of whether the notion of embeddedness, as shaped by Granovetter's influential article, really succeeds in integrating the economy and society. Critics did not much question the capacity of today's embeddedness research to shed light on the firm, but were dissatisfied with its representation of the market. Calnitsky (2014) and Krippner (2002) claim that economic sociology still follows neoclassical economics in presupposing the existence of a core asocial market which uncomfortably sits surrounded by society. If we further stretch this concept to describe the platform, will we uncover other implications of it that have remained dormant so far?

As a first step toward answering these questions, I propose a view of platforms that benefits simultaneously from recent theoretical developments in economics and in economic sociology. I use insight from transaction cost theory (and, to a lesser extent, two-sided market models) within economics, and from the "multi-level analysis of networks" applied to economic sociology. For tractability and to ensure comparison with recent literature (Wood et al., 2019b), I narrow down my analysis to the case of online labor, where digital platforms allocate tasks and jobs to providers who execute them remotely from their laptops or smartphones.

In this setting, I show that the common interpretation of Granovetter's embeddedness in terms of *social* networks of individuals – peers, friends, relatives – is limiting. Drawing on the multi-level approach proposed by Brailly et al. (2016), I extend the concept of embeddedness to also encompass *economic* networks, which include combinations of firm-to-firm and individual-to-firm networks. Their juxtaposition brings to light the linkages between platforms' business models, their structure of commercial ties, and workers' individual outcomes. Two configurations can be observed on this basis: when platforms take the role of mere market intermediaries that connect workers and clients, economic ties are weak regardless of social ties, and workers are left to their own devices, in a form of *disembeddedness*. When platforms take on management roles akin to firms, economic ties take highly complex forms with multiple layers of contractors, brokers and suppliers that frame the activity of the individual. These layers constitute a structure that is evocatively reminiscent of complex neural networks in so-called "deep learning", and I thus propose the notion of "*deep embeddedness*" to describe them.

The discovery of deep embeddedness has three main implications for research in economic sociology. Substantively, it contributes to the empirical literature on the platform economy by describing the business models of some online labor platforms that have remained largely under-researched so far. Epistemologically, it demonstrates the value of dialogue between sociology, social network analysis and parts of economics, under an organizational umbrella. Theoretically, it indicates that the relationships that matter are not only personal ones such as friendship or kinship, but also economic ones – relations that describe who owns which hardware and software tools, who recruits workers, who provides quality controls. Augmented to include more diverse sets of relationships, Granovetter's original intuition is capable of syn-

thesizing both the social relations that originate outside any economic action, and the economic relations that derive from actors' position in today's highly specialized production system. It is therefore more powerful than its critics contend, and can respond to concerns that economic sociology focuses too much on personal ties while paying insufficient attention to productive structures. An integrated understanding of Granovetter's proposition, enriched with insight from recent network-theory research, has value for analyzing contemporary economic transformations.

2 Conceptions of embeddedness: from Polanyi to Granovetter and beyond

The concept of embeddedness is not an easy one to define. Long before Granovetter (1985), its origins are attributed to the highly influential work of Karl Polanyi (1944) where it took a broader, more comprehensive meaning. According to Beckert (2009), it encompassed the limits that institutions and the moral fabric of society impose on economic transactions, while also alluding to the normative task of stabilizing the organization of society through the institutional regulation of the economy. However, Polanyi's formulation is confusing, not least because he used the term only twice in his book, although the concept underpins his overall view. Peck (2013) identifies a sharp contradiction, whereby Polanyi argued in some places that all economic activity is always "embedded and enmeshed in institutions economic and noneconomic" (Polanyi, 1957), but proposed a more radical view elsewhere, with focus on the disembedding of the "fictitious commodities" of labor, land and money and the importance of societal-level legal, normative and cultural constraints that limit commodification. Peck (2013) refers to the former as "soft Polanyi" and to the latter as "hard Polanyi". In this perspective, the work of Granovetter (1985) is a development of the soft Polanyi that leaves aside the hard version¹.

Granovetter's formulation has not escaped criticism, and a common argument pertains to the social networks approach that has quickly come to dominate much of economic sociology. Structure, it has been said, becomes the only explanation of economic outcomes but at the expense of forgetting the social content underlying it and ultimately failing to explain its own origins – that is, how it emerges, or how networks are shaped the way they are (Krippner, 2002; Beckert, 2003). Motivated by concerns that loosely remind those of the defenders of a "hard Polanyi", other critics found that emphasis on networks of personal ties abstracts away the social relations of capitalism – those defined around ownership of means of production – and therefore misses the modes in which individuals participate in socioeconomic life (Calnitsky, 2014).

To be sure, research on social networks has made significant advances in the past fifteen years that dampen many of these criticisms. A "theory of networks" endeavor-

¹ Historically, Granovetter seems to have developed his ideas and terminology independently (Granovetter, 2004, p. 113).

ors to explain network structures, as a complement to already existing “network theory” that uses networks to explain behaviors (Borgatti and Halgin, 2011). There is growing interest in the content and meaning of ties, notably through qualitative and mixed-method network analyses (Crossley, 2010; Dominguez and Hollstein, 2014). A research program in “neo-structural sociology” combines structure, culture, and agency (individual and collective) in new ways to study contemporary economies (Lazega, 2018).

Specifically on embeddedness, a step forward is the re-framing of the concept as a multi-level problem. Recognizing that most extant research separates the analysis of networks between firms and analysis of networks between individuals, Brailly et al. (2016) propose to study them jointly, each constituting a level in a multi-level analysis (Snijders and Bosker, 1999). In practice, they consider at the same time the networks of inter-individual ties $I - I$, the network of inter-firm ties $F - F$, and the affiliation network (*à la* Breiger (1974)) that links individuals to firms $I - F$. The first two, within-level networks may be characterized as “horizontal”, and the third one, across levels, as “vertical”, as illustrated schematically in figure 1. Brailly et al. (2016) show that while each level has its own specific processes, they are partly nested and a full understanding must take into account their inter-dependencies. Published in a social network analysis journal, these ideas still wait to be fully recognized within sociology. They suggest a way in which networks reconcile the social and the economic, which correspond to different levels: the horizontal networks between individuals $I - I$ are typically social (friendship, advice, information exchange) while those between firms $F - F$ are economic (contracts, alliances) and the vertical networks that tie individuals to firms $I - F$ are also economic (mostly employment relationships, though sometimes they may take other forms such as appointments to board directorships). But then, can a multi-level approach that combines all of these respond to (at least some of) the criticisms addressed to Granovetter’s concept of embeddedness? Does it suffice to take on board the insights that readers found in Polanyi, but not in the Granovetterian tradition?

This brief overview summarizes aspects of the debates surrounding embeddedness that are of relevance to better understand how digital platforms are changing the current socio-economic landscape. While network structures are crucial to operationalize embeddedness, most discussions revolve around the nature and meaning of these networks (whom they connect, what relationships they represent), as well as their structures and types (within and across levels, or horizontally and vertically). If these ideas are (more or less) well-positioned to improve the capacity of the embeddedness concept to enlighten its original objects of study, the firm and the market, how do they perform with regard to platforms?

3 A case study: Platforms for online labor

In essence, platforms are coordination devices based on use of digital technologies. The engineering literature defines them as extensible hardware or software func-

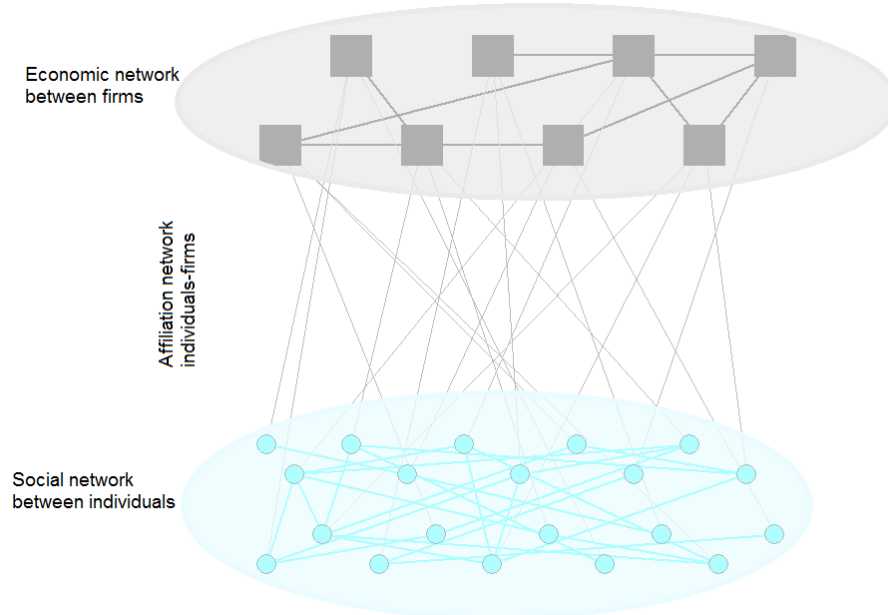


Fig. 1 A schematic graph representation of multi-level embeddedness. Nodes are actors (cyan circles = individuals, black squares = firms) and lines are ties between actors (black = inter-firm ties $F - F$, cyan = inter-individual ties $I - I$, gray = individual-firm ties $I - F$). Source: author's elaboration, adapted from Brailly et al. (2016).

tionalities to which third-party modules can be added, while economics and management stress the organizational processes and standards with which they facilitate participation and involvement of relevant third parties (Gawer, 2014; de Reuver et al., 2018). The latter view has been highly influential throughout the social and economic sciences, with the idea of platforms as intermediaries between distinct groups of economic actors (Srnicsek, 2016) borrowed from “two-sided” (or “multi-sided”) market economics (Eisenmann et al., 2006; Rochet and Tirole, 2003; Rysman, 2009). Two-sided markets are characterized by so-called “network externalities”, meaning that a large number of users on one side (for example, Uber drivers) triggers participation on the other side (passengers). The platform captures the value of these externalities, which can occur within and/or across sides: hence, its user base is one of its main assets. To increase its size, the platform may choose to subsidize the more price-sensitive side (for example, users of the Google search engine who access the service for free) at the expense of the other (advertisers, who bid for “eyeballs” in dedicated auctions).

3.1 *Digital platforms and embeddedness*

As technology-enabled coordination devices, internet platforms interrogate extant understandings of embeddedness in two main ways. First, they challenge the established distinction between firm and market as devices for economic coordination (Sundararajan, 2016, pp. 78-79). Are they just another instance of firms or of markets, a hybrid, or an entirely new third way? Is their alleged capacity to reduce transaction costs (Lobel, 2018) enough to supersede the firm in favor of the price mechanism (or *vice-versa*)? It is difficult to answer these questions, owing to lack of shared definitions and widespread use of the everyday-language meaning of the terms “firm” and “market” – not least because different disciplines take part in the debate, each with its own vocabulary.

Perhaps in no other area is ambiguity more apparent than in the inter-mediation of *labor* – where even greater theoretical rigor might bring little improvement, if we believe with Coase (1937, p. 392) that labor services are most relevant to determine the arbitrage between firm and market, all other production inputs being far easier to accommodate into existing theories. Labor platforms are those that bring together clients and workers to buy and sell “gigs”. Because these workers are recruited on demand and their commitment ends as soon as their output is delivered, they are legally not employees. But because they rarely have the freedom to shape and manage their work, they are not fully independent providers either. Labor platforms blur the economic firm-market duality as much as the legal notion of an employer-employee relationship (De Stefano, 2016; De Stefano and Aloisi, 2018).

Second, platforms transform relational mechanisms such as reciprocity and trust, that the embeddedness literature thought to support economic action. When they reached center-stage in the early 2010s, and were surrounded by enthusiasms around the promises of “collaborative” consumption (Botsman and Rogers, 2010) and a “sharing” economy (Lessig, 2008), platforms seemed to elevate these mechanisms from informal to formal (and valuable). But it soon turned out that platforms replace these social mechanisms with impersonal systems of reputation-monitoring (Wood et al., 2019a) which impose discipline and are akin to surveillance. Again, this tension is particularly apparent in labor platforms, whose algorithmic infrastructure records workers’ activities and transforms them into publicly-visible performance scores (Rosenblat, 2018). These systems arouse strategic reactions on both sides and call for caution before interpreting any observations.

On this basis, I contend that labor platforms are an ideal arena to study the present transformations of embeddedness. But even so, they constitute a very broad object of study encompassing heterogeneous cases. In what follows, I review extant categorizations to identify particularly relevant sub-groups.

3.2 *Online labor platforms*

Existing efforts to draw typologies of labor platforms (Codagnone et al., 2016; Howcroft and Bergvall-Kåreborn, 2019; Schmidt, 2017), share two common criteria even though not always expressed in the same terminology. One is the extent to which work is performed offline while inter-mediation occurs online, in comparison to all-online labor platforms. The former category encompasses, among others, platforms for food delivery (for example, Deliveroo), urban transport (Uber) and cleaning (Helpling), while the latter includes platforms for non-manual work such as design, computer-programming, accounting, transcriptions, translations, and data entry: examples are Upwork, Freelancer.com and Fiverr. The other main criterion distinguishes platforms that allocate jobs in full, as whole projects to individual professional freelancers (design of the logo of a company or translation of a complete document, for example through the platform Upwork) or small parts of it to myriad non-specialists (“micro-tasks” such as identification of objects in an image, or translation of only one or few sentences, as on Amazon Mechanical Turk).

Barely visible in the media and less concerned by legal fights so far, all-online platforms are nevertheless particularly interesting to study. With loose geographical boundaries, workers are exposed to wide competition that drives remunerations down and sometimes exposes them to unconventional working hours and possibly overtime – for example, to meet demand from clients in different time zones (Graham and Anwar, 2019). Thus, the conditions for embeddedness may be more difficult to meet. This is why Wood et al. (2019b) choose precisely this setting to undertake the first-ever study of embeddedness in the platform economy – with particular focus on global platforms for full freelancing projects. They find instances of embeddedness in the sense of Peck (2013)’s “soft Polanyi” (section 2) which is also similar to the view of Granovetter (1985): many of the workers they observe rely on their peer networks of, for example, friends, kin or even co-workers to get tasks done. To some extent, they restore the personal relationships that platforms’ formal reputation systems obfuscate (subsection 5.3). However, workers appear disembedded in the “hard Polanyi” (Peck, 2013) sense that Granovetter allegedly left aside, and that sees embeddedness as absence of commodification: without labor regulations and rights, they are exposed to fluctuations in the labor market, and have limited access to healthcare and social protection.

It is tempting to try and generalize these ideas to other areas – notably micro-work, the other main sub-category within all-online platform labor. Here, one would also observe absence of labor rights and exclusion from welfare. In terms of peer networking, however, results from the literature are mixed. A team at Microsoft Research mapped the communication ties of Amazon Mechanical Turk workers and found that they constitute a structured network (Yin et al., 2016). In a related ethnographic study across four different platforms (Gray et al., 2016), they found that there is collaboration, sometimes spontaneously re-created by workers regardless of platforms’ unsupportive policies. But before interpreting these results in terms of micro-workers’ network embeddedness, two caveats are in order. First, in the network maps of Yin et al. (2016), only 13.4% of workers are connected, so that there

is a long tail of disembeddedness as per Granovetter’s original definition. Second, the network is enabled by forums, a specificity of Amazon Mechanical Turk that may have no equivalent elsewhere. Initiatives to raise awareness and improve working conditions have coalesced around this platform, older and better-known than its competitors.

These studies place emphasis on the social networks of peers and focus exclusively on individuals ($I - I$, left panel of figure 2). They remain silent on the economic ties $I - F$ that link individuals to firms (or platforms or other economic organizations that shape the production process), represented as gray lines in the central panel of figure 2. Neither do they make explicit any economic linkages between firms (or platforms, etc.), represented as black thick lines in the central panel of figure 2 ($F - F$). Without these economic dimensions, network analysis cannot link social ties such as mutual help and support to the economic positioning of individuals in the productive structure unfolding around the platform. It is necessary to simultaneously address the two, as schematically represented in the right panel of figure 2, which is also consistent with the multi-level structure of figure 1. To move forward in this direction, I now explore different platform business models and probe them against the above-outlined multi-level view of embeddedness.

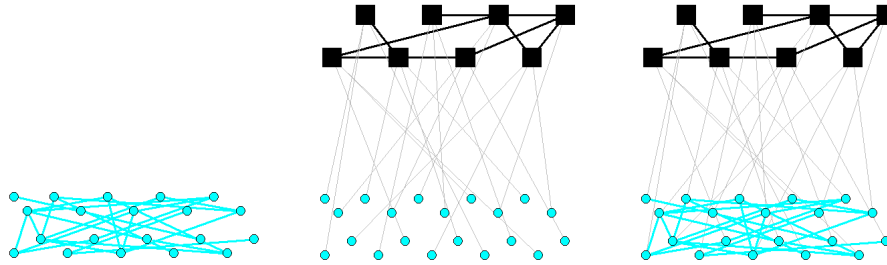


Fig. 2 Different dimensions of embeddedness. Nodes are actors (cyan circles = individuals, black squares = firms) and lines are ties between actors (black = inter-firm ties $F - F$, cyan = inter-individual ties $I - I$, gray = individual-firm ties $I - F$). Left: social ($I - I$) network. Center: economic ($I - F$ and $F - F$) networks. Right: multilevel structure of social and economic networks. Source: author’s elaboration.

4 An inventory of platforms

To provide empirical support for these ideas, I draw on a comprehensive study of platform labor in France, DiPLab (“Digital Platform Labor”), undertaken in 2017-19 (Casilli et al., 2019). France offers an original perspective relative to extant literature focused on the English-speaking world, while still enabling cross-national compa-

rability and allowing generalizability of (many) findings beyond its borders. Its demand for remotely-provided data services is high because it is investing heavily in the development of digital technologies and in particular, of artificial intelligence. Supply is sizable too, as the country features a mix of international platforms (that operate mostly, if not exclusively, in English) and local ones (mostly in French), offering a range of options to workers. While in principle, one may expect online labor to be undertaken anywhere, competition is not always global because some tasks require language proficiency or local knowledge. It is for this reason that even if a lot of the online workforce is located in low-income countries in the Global South (Casilli, 2017; Graham et al., 2017), niches remain in the richer North.

The dataset that I use here consists of a detailed inventory of online labor platforms and mobile applications in use in the country, with information on their company structures, business models, history, location, financial situation, marketing strategies and human resources. A total of 28 platforms have been identified, combining information from different sources: business databases such as Crunchbase and *societe.com*, a local repertoire; log data from *NetBusinessRating*, a dedicated discussion forum where workers share their reviews of platforms; and a set of exploratory interviews that asked workers to name the platforms they used and those they were at least aware of. Once compiled and double-checked for completeness, this list was enriched with details for each of the platforms, taken from their websites, press releases, media features, and in some cases personal communications with founders or managers. The majority of these platforms offer online-only micro-working services, though some of them do so secondarily, their main activity concerning micro-tasks that requires physical presence in some place, like taking pictures of products in shops (subsection 3.2). Other platforms included in the inventory offer inter-mediation of full-fledged freelancing services such as design and computer programming. The broad scope of the inventory allows comparisons across cases and types (notably micro-work and freelancing), and helps to distinguish widespread practices from any idiosyncrasies.

As a complement to this data, I use a set of 92 in-depth interviews with online workers, clients, platform staff and other stakeholders, mostly though not exclusively based in France. For the purposes of this paper, interviews are solely meant to provide concrete examples of configurations that may seem otherwise abstract or unfamiliar, and to support my interpretation of results.

5 Platform business models and embeddedness

Observation of the selected platforms clearly indicates that they adopt two main types of business models. I now illustrate them by expanding from some first results published in the DiPLab project report (Casilli et al., 2019). I then extend the discussion to the question of how these elements inform the analysis of embeddedness framed above.

5.1 *Platform as markets: intermediaries*

Some of the observed platforms function just as classical textbook cases of two-sided market intermediaries (section 3). The prototypical example of this mode of functioning is Amazon Mechanical Turk. The platform organizes interactions between two user groups, workers on the one side and clients (called “requesters”) on the other, and provides the technical infrastructure upon which their activities occur. This model mimics a lean, decentralized marketplace where the intermediary is simply a facilitator, and decisions are left to clients (how much to pay for a task, how many workers should complete it, how long it should be available on the site) and workers (what tasks to do).

For the system to run as smoothly as possible, the platform offers multiple built-in functionalities so that it can host the whole process internally, end-to-end (figure 3, top panel). Once logged in, workers can browse and choose from available tasks. They then accept the conditions imposed for these tasks (if different from the general terms of use of the platform), execute them, submit them and claim payment. On the side of the client, the platform makes available a range of tools and templates to design tasks, post them on the site, review and accept (or reject) workers’ submissions, and make payments at the end. On Mechanical Turk, all these functions are accessible through an API (Application Programming Interface) that obviates the need for any other form of personal contact between clients and workers. The latter remain anonymous and their personal characteristics are invisible to the former². From the viewpoint of the requesting client, the transaction is technically un-distinguishable from any remote procedure call, completely de-personalizing the labor relationship.

After the successful launch of this market-oriented model by Amazon in the mid-2000s, many other micro-tasking platforms such as the American Microworkers and the French Foule Factory have adopted variants of it. Most freelancing platforms like those observed by Wood et al. (2019b) function in about the same way, though with some adaptation. Instead of letting workers get tasks on a first-come, first-served basis as on Mechanical Turk, freelancing platforms let clients select workers based on their characteristics, track record and possibly interest in their request. This means that there is no anonymity, and that contact between client and worker is not only allowed but encouraged. Remunerations are not fixed by the client and equal for all participating workers, but negotiated on a case-by-case basis, often based on a quote provided by the worker (and sometimes, they are the outcome of competitive bidding in an auction). Interestingly, my data show that pricing policies differ markedly: freelancing platforms subsidize clients and charge commission fees to workers, while micro-working platforms subsidize workers and charge only clients

² Specific functionalities allow clients to request that workers meet some pre-selected requirements (skills or socio-demographic characteristics), but it is the platform that operates the selection through automated worker screening and qualifications.

(Casilli et al., 2019)³. Despite these differences, most freelancing platforms can be said to function as market-like structures.

If digital platforms are often praised for their capacity to reduce transaction costs (section 5.3), marketplace-like structures for labor inter-mediation fail to achieve this. The time that workers spend searching for tasks, and working on tasks that are rejected or ultimately not submitted (often due to unclear instructions from requesters) is a cost that, on Amazon Mechanical Turk, results in a large gap between the \$11 per hour that the average requester pays, and the median hourly wage of only \$2 that workers earn (Hara et al., 2018). This is also true of most global online freelancing platforms, where workers spend considerable time searching for requests, applying, bidding unsuccessfully and waiting for work (Wood et al., 2019b). The massive over-supply that characterizes online labor platforms both in freelancing and micro-work (Graham and Anwar, 2019; Schmidt, 2019) also contributes to raising search costs for workers who need to set up alert notifications, connect several times a day, or keep their screens continuously on in order not to miss opportunities.

Clients also face high transaction costs – in the form of policing and enforcement costs. Clients struggle to monitor anonymous external contributors, and platforms' reputation systems only provide signals of potential worker quality, without guaranteeing engagement or performance. Thus, requesters have set up multiple solutions to assure quality, some of them expensive: having several workers do the same task and choosing the majority solution, or posting follow-up tasks that ask workers to check the outputs of others (Vaughan, 2018). Clients' demand for quality is increasing today, as data security becomes a major concern for large companies that use proprietary data-sets and that need to protect their cutting-edge R&D activities from the gaze of competitors. Thus, they become wary of generalist platforms where anyone can see the posted tasks and, worse, unidentified contributors can access the underlying data.

Another reason why high transaction costs may arise for clients is the growing demand, at least by the most technologically advanced among them, for specialist software infrastructures with state-of-the-art customized features. In particular, there are rising expectations regarding annotation and labeling tools such as those needed to prepare training data-sets for artificial intelligence algorithms (Schmidt, 2019; Tubaro and Casilli, 2019). Transaction cost theory (Williamson, 1981) suggests how provision of such specific assets is unlikely to occur on the open market, requiring instead tighter, long-term engagement between the involved parties. It is a major challenge to Amazon's model.

³ According to two-sided market theory (Evans and Schmalensee, 2016), this means that freelancing workers are less price-sensitive than their clients, while the opposite is true of micro-workers. See section 3.

5.2 *Platforms as firms: multiple layers*

To reduce transaction costs, an alternative model has recently emerged, in which the workforce is managed through vendor relationships on behalf of (mostly large) corporate customers. In this case, the organization of the workflow is managed like in a firm, which centralizes all transactions and does not need the price system to ensure coordination. Champions of this model are three large and well-known actors of the global digital data business, notably Appen, Lionbridge and Pactera. This does not mean that they employ workers for a salary as traditional companies do, as their costs would raise too much. Rather, they arrange tightly-knit outsourcing networks of providers and intermediaries coordinated through a set of contracts and common infrastructures. Centralized direction distinguishes them from platform-as-market models, while multiplicity of actors demarcates them from classical firms.

As a result, the whole work process is sliced into several “layers” with different interlocutors in each of them (figure 3, bottom panel). First, potential new workers see advertisements of online earning opportunities on standard jobbing websites, such as Monster.com. They apply and then go through a selection process, which may include an interview, and is managed by the vendor that posted the ads – for example Appen – or one of its subsidiaries. Selected workers are then invited to connect to a dedicated platform, or another external vendor, to establish a contract. In contrast to simple, market-like platforms where labor is systematically paid by piece-rate, here contracts often specify a regular amount of daily or weekly time that the worker commits to, for a generically defined type of tasks rather than a single, specific one (although the worker may at times be offered less work than agreed, or no work at all, depending on clients’ demand), and a hourly remuneration. Finally, contracted-out workers are directed to the platform that provides the technical infrastructure to execute tasks. This is often a closed, proprietary platform belonging to the final client, a monopsonist: for example UHRS (Universal Human Relevance System) for Microsoft, RaterHub for Google and TryRating for Apple. The functioning of these technical platforms is highly confidential, with no access allowed from the outside and often, not even a public web page. Its management is often entrusted to another sub-contractor, notably if there is bespoke technical infrastructure that requires maintenance, or if there is a need to supervise a pool of workers operating in the same premises. There is evidence that providers in India, Indonesia, the Philippines, and parts of Africa set up offices where they manage teams of local workers (Graham et al., 2017; Graham and Anwar, 2019; Roberts, 2019; Schmidt, 2019). Such arrangements are particularly interesting for large-sized client companies owing to the relatively low cost of labor and the high technical expertise in some of these countries (Murali, 2019). Evidence so far is that such arrangements are less common in Europe and North America, but can be observed in specific cases, for example, when data need to be accessed securely in a dedicated onsite facility.

One interview with a female micro-worker whom I will call J., illustrates concretely the multi-layer structure that the worker is integrated into, and its global outreach. A resident of South-Eastern France, J. was contracted by a Chinese data company to do transcriptions on behalf of an American client, the producer of a

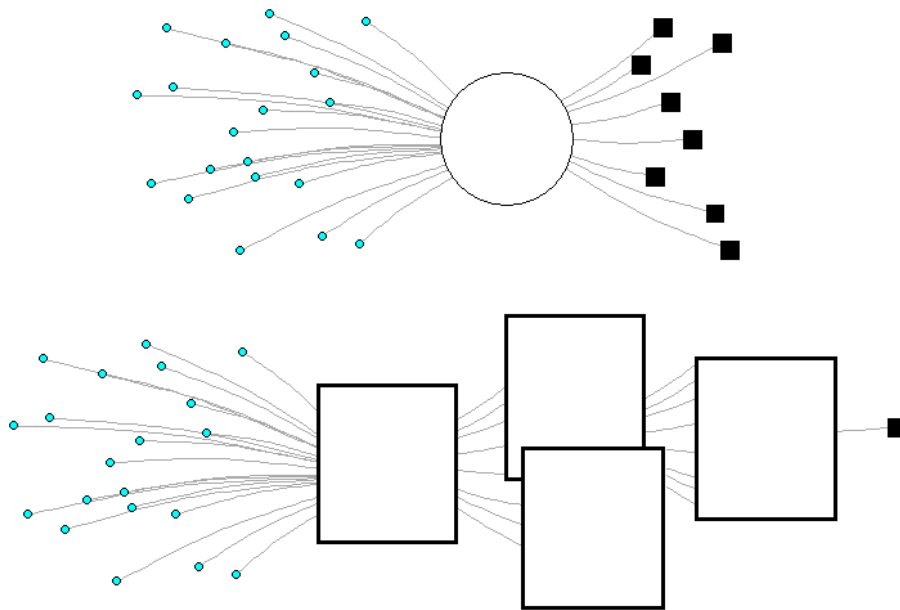


Fig. 3 Two models of platform inter-mediation for micro-work. Top panel: intermediaries in market-like two-sided structures of interaction. Bottom panel: firm-like multi-layer structures. Cyan circles = workers, black squares = final clients, white large circle = marketplace platform, white rectangles = firm-like platforms, their contractors and subsidiaries. Source: author's elaboration.

vocal assistant. J. had to listen to recordings of the interactions of French users with this vocal assistant, check them against the automated transcriptions that the assistant had done, and correct any errors. She was aware of micro-workers in neighboring countries who were doing the same tasks in their own languages, and they were all under the supervision of a subcontractor in Italy who managed day-to-day issues and practicalities. She received US tax forms for the income she earned.

Under these conditions, transaction costs go down. Workers do not have to spend time searching for tasks, because they receive assignments every day or every week, their only choice being whether to accept them or not. Because the system is tightly managed and deals with large orders, workers usually enjoy better remunerations and more predictable working hours than with market-like platforms. A., a 40-year-old mother of two who lives in Southern France, has been doing 20-30 hours a week of tasks on an American platform for the last 3 years. Throughout this period, she arranged never to take holidays in order not to miss assignments. Although she considers this activity temporary and plans to change job when her children grow up, she is relatively satisfied of her pay of 15 euros per hour – which is above the

French minimum wage. She also signed up with a marketplace platform but does not find it interesting, and uses it rarely.

On their side, clients are reassured about quality and therefore, incur lower policing and enforcing costs. These platforms that Schmidt (2019) qualifies as “second generation” in comparison to earlier market-like models *à la* Amazon Mechanical Turk, know their workers and are able to set up schemes to assure quality like firms do: careful selection of who does what tasks, attribution of monitoring responsibilities to the most motivated workers, creation of qualified teams to do more complex tasks together. Further, these structures offer greater confidentiality and security in that clients’ data and technical tools (for example, specialist annotation and labeling software used to prepare training data-sets for artificial intelligence algorithms) are only accessed by selected workers and are under the direct control of the client or one of its suppliers, as on UHRS.

5.3 Platform-market, platform-firm and embeddedness

How do the preceding remarks relate to the concept of multi-level embeddedness outlined in section 2? Recall that the idea is to reconstitute the inter-individual $I - I$ and the inter-firm $F - F$ networks of ties, and the affiliation network that links individuals to firms $I - F$. Let us temporarily leave aside the inter-individual network $I - I$ on which the literature has already provided evidence as discussed in subsection 3.2. The above discussion of business models can help reconstitute the other two which, as mentioned, consist of economic relationships.

In the platform-as-market model, there is no $F - F$ network of ties between firms. Clients act independently and autonomously of one another, and their only interlocutor is the platform. If there is a network, it is star-shaped, with the platform at the center. Similarly, affiliation ties $I - F$ are evanescent. Workers are not employees and their engagement with the platform (and/or its clients), construed as a spot contract between independent businesses, deprives them of the human contact and sense-making that normally accompany work. Excess supply of labor (Graham and Anwar, 2019) makes them easily replaceable, and in the case of micro-work, they are anonymous in the eyes of their clients as mentioned above. Further, interviews suggest that micro-workers do not always know the identity of the client or the purposes of the tasks (even though platforms encourage requesters to disclose this information). For example, D. described a videogame-like task in which she had to move characters back or forth on her screen depending on their appearance (North African or French). She ignored who was sponsoring this task and what its purposes were: developing a tool to monitor prejudice or to spread anti-immigrant ideas? Under these conditions, decisions at the client (F) level affect conditions at the worker (I) level, but there can be no inter-level feedback. In sum, the platform-as-market model comes with disembeddedness in a multi-level sense – whether or not workers are embedded in the narrower sense of having a network of peers $I - I$ who may support their activity.

Instead, the platform-as-firm model generates dense economic networks. The inter-firm $F - F$ network consists of the contractual agreements between the final client (for example, Microsoft and its proprietary service UHRS), its vendors (for example, Pacter), their subcontractors (in charge of, say, managing the workforce on a daily basis or of maintaining the technical infrastructure) and any additional service provider (such as a jobbing website). The affiliation network $I - F$ still includes no proper employment relationships, but it ties workers to those in charge of the direction of their activity: to the company or platform which contracted them out, to the final client who (hopefully) files tax documents for them, to the subcontractors who provide technical tools or day-to-day management. The client can know who the workers are and *vice-versa*, particularly when orders are large or repeated.

There is certainly here a form of embeddedness, reminiscent of the multi-level conception described by Brailly et al. (2016), whereby individuals are integrated into complex contractual structures that directly or indirectly orient their productive activity. Its specificity is the multi-layer configuration that distinguishes its outer appearance from more classical cases with just firms and individuals. This form may be dubbed “deep” embeddedness, following an intuition of Casilli et al. (2019) who talk about “deep labor” platforms to highlight the apparent similarity between their multiple-layer structure and the complex neural networks that constitute the building blocks of the mathematical field of *deep learning*, where “layers” are containers that receive input, transform it with sophisticated non-linear functions and then pass the results to the next layer. Beyond the surface, the reference to deep learning is motivated by the prominent place of the latter in the development of today’s artificial intelligence research, and its intensive use of online labor, particularly micro-work, for its needs.

But there’s more to the “deep” adjective in the case of the multi-layer structures described above. Depth signals opaqueness of the overall construction to the outside viewer, especially if it is part of a corporate net that stretches globally. The worker sees some of the layers (and typically, can name clients as discussed in subsection 5.2) but not necessarily the whole structure, and essential parts of it remain black boxes. For example, A. said she had no idea how the platform chose tasks for her to do. The reason why she never refused a task – at the expense of being exposed to disturbing material when she had to moderate violent or pornographic content – is that she feared a refusal might lead the platform to give her less work, but she had no way to check whether her apprehension was genuine.

Additionally, the “deep” metaphor points to the hidden role of labor, whose contribution to datafied production processes is essential but often left in the shadow (Gray and Suri, 2019). Corporate communication highlights the role of technology, not human contribution, especially in the artificial intelligence industry (Casilli, 2019; Casilli et al., 2019; Schmidt, 2019; Tubaro and Casilli, 2019). In this sense, laborers are deeply embedded in these multi-layer structures because their contribution is elicited in ways that are little visible from the outside – neither the final client such as UHRS, nor its main vendor such as Appen, hire them as salaried employees – but still integrate them in the system and to some extent, control their activities.

Both workers and their client(s) can be said to be deeply embedded. The client renounces the user-friendly interfaces of market platforms such as Mechanical Turk, and their offer of standardized technical tools, to develop its own infrastructure – generally lacking fancy communication solutions and sometimes even devoid of support services, as these are delegated to vendors and their subcontractors. The ties that bind a client to these vendors usually involve some mutual dependency due to asset specificity, that is, investment of the vendor in the development of tools and services that respond very precisely to the client's needs, but would be difficult to re-deploy elsewhere in case of contract termination (Williamson, 1975, 1981).

Let us now re-introduce the question of the $I-I$ social networks among individual workers, previously left aside. Such networks may exist under deep embeddedness. Some platforms invest in the creation of networking tools for their workers via online forums and chats, in the hope of increasing their engagement (Schmidt, 2019, p. 37). However in another sense, the siloed architecture of these platforms restricts interactions among workers: an example is UHRS, where different vendors manage the workforce and the few existing discussion boards are internal to each of them and therefore asynchronous (Gray et al., 2016). In practice, in such cases inter-individual ties $I-I$ exist (almost) only between workers who share affiliation ties $I-F$ to the same firm F_i . Additionally, participation in these networks does not necessarily concern all workers: A. claims not to be in touch with anyone except (occasionally) platform support services, and not to know any other micro-worker.

Is deep embeddedness good or bad for workers? In exchange for lower search costs, it removes some of the flexibility and autonomy that disembedded, market-like platforms such as Mechanical Turk allow. Indeed workers have to accept a certain discipline (the amount of work they engage themselves to do regularly) and in some cases, supervision by others (whenever a subcontractor is in charge of workforce management). Discipline has some advantages: Gray et al. (2016) suggest that the internal organizational structures of each single vendor or subcontractor may enable mentoring and skill-building in ways that would not be possible on Mechanical Turk.

But there is a profound power asymmetry insofar as the client side retains the flexibility that workers lose – notably in terms of the possibility to dispose of the workforce if and when needed. An interview illustrates what this means. F. is a 25-year-old man from Northern France who used to work for an international deeply-embedded platform structure when he was a student. He had engaged himself to do one hour a day for seven days a week, and in this way, he earned about 400 euros a month – a welcome supplement to his meager student budget. Yet the platform discontinued his contract after six months, without any explanations.

Brailly et al. (2016) and Lazega (2016) show that in a multi-level structure, different levels have different temporalities, and require actors to adjust. Here, platforms move faster, and this asynchrony comes to workers' disadvantage.

6 Conclusions

I have shown that the early platform-as-market model embodied in Amazon Mechanical Turk *disembeds* workers, exposing them to high costs of participation to the platform (transaction costs) and, one might add, potential psycho-social risks due to lack of purpose and difficulty to give meaning to their work (for example when clients disclose too little information). The emerging multi-layer model adopted by, among others, Pacter and Appen *deeply embeds* workers, paying them more but also reducing their flexibility and subjecting them to stricter discipline and sometimes outright control, while also restricting their opportunities for peer interactions at least in some cases. With minor tweaks, the same can be said of clients. Market-like platforms offer them user-friendly interfaces and a range of off-the-shelf technical tools, but impose high enforcement costs and discourage development of specific assets: their contracts can be seen as arm's length ties in the sense of Uzzi (1997), therefore disembedded. Multi-layer platforms enable customization of technical tools and more controlled management of the workforce, but require engagement in tighter, longer-term contracts that are more difficult to terminate or change: they are more similar, in this respect, to Uzzi's embedded ties.

Epistemologically, the study undertaken showcases the benefits of blending traditions and disciplines. Transaction costs economics has helped to identify the two main business models of online labor platforms and the circumstances under which either of them is chosen, sometimes beyond the predictions of two-sided market theory. Social network analysis combined with economic sociology has provided the multi-level structure as a framework to analyze these business models and derive implications from them. Organization theory broadly interpreted links these different insights, interpreting platforms as coordination devices and raising the question of the modes and costs through which they function.

Implications for the sociological study of digital platforms are of import. To begin with, these results add an extra dimension to extant categorizations of labor platforms, summarized in subsection 3.2. Most freelancing and some micro-work platforms adopt the marketplace model, though they implement it differently, notably in terms of their pricing structures; while only some micro-work platforms embrace the multi-layer model as of yet. Remunerations and working conditions differ markedly. Specifically, deep embeddedness limits some of the worst consequences of the platformization of labor: it re-introduces some degree of information and sense-making for the worker, a certain regularity of engagement, and at least minimal structures of support. Therefore, a revised typology that also introduces a dividing line based on platforms' business model, in addition to the degree of qualification or the duration of the tasks they offer, will more accurately describe the diverse situations that workers face.

Further, this research exposes the limitations of the platform-as-market model, as industry leaders increasingly adopt the alternative multi-layer structures. It is too early to discern the effects of this ongoing evolution, but provisional hypotheses can be formulated. One is that asymmetries across clients may drive the future unfolding of these tendencies. The multi-layer model may become more prominent if it is

the preferred choice of larger or more powerful clients, but if smaller players cannot afford the asset-specific investments it requires, they will have to be content with increasingly less attractive market-like platforms, resulting in inequalities in access to labor-powered data resources. The other is that, against a popular idea, the platform economy may not mean the end of centralization – and possibly of more traditional models of labor supervision. This insight goes against a defense line often adopted by platforms in court cases that opposed them to workers. If they are mere intermediaries between independent businesses, as they typically claim, they have hardly any responsibilities toward their workers, but if they are firm-like entities, they do. The policy implications of the platform-as-firm model may strengthen a view of the *platform as an employer*, thereby touching upon current lively disputes around the legal status of platform workers (De Stefano, 2016; Prassl and Risak, 2016).

Related to the above, existing conceptualizations of platforms reveal their limitations. Two-sided market economics (section 3), fits well with the platform-as-market model, but less so with the multi-layer model. An organizational lens along the lines of Gawer (2014) is more comprehensive, interpreting organization broadly as a generic coordination device that may take different forms, each entailing specific costs and incentives. However, Gawer's definition of platforms as *meta-organizations*, which presupposes that members are "legally autonomous and not linked through employment relationships" (p. 1240), sweeps under the carpet the controversies surrounding the status of platform workers. A more suitable organizational perspective would stress how all online labor platforms are systems of coordinating activities *in view of supporting outsourcing*. Both the market-like and the multi-layer models aim to remove the need for employment relationships, allowing the client to access labor on demand: they only differ in the way they achieve this, one through a spot market system and the other through a set of (longer-term) subcontracting ties to other organizations. Both operate on a planetary scale and serve a flourishing just-in-time economy, in contrast to the local labor market settings of most embeddedness research – including the work of Polanyi and Granovetter themselves. These new geographies of online labor (Graham et al., 2019) matter because they shift the balance of power in favor of capital (clients, platforms, and their subsidiaries and contractors) and against labor – insofar as capital benefits from its greater mobility compared to labor (Graham and Anwar, 2019). Online labor thus appears as a site of power struggles – much like the large corporation of the twentieth century – rather than a place where trade occurs between equals.

More generally, the findings illustrated in this article advance research on embeddedness within economic sociology. Critics asserted that focus on just the social network of friends and kin is too narrow to really reconcile the social and economic dimensions (section 2). I have submitted that the multi-level network analysis approach of Brailly et al. (2016) offers a solution by combining the social networks of individuals $I - I$, the economic networks of firms $F - F$, and the (likewise economic) affiliation network between individuals and firms $I - F$. Application of these schemes to platforms suggests that economic networks may be even more important than social networks as they reveal important asymmetries of power that affect action. In the multi-level networks that represent market-like platforms, there are no

inter-level feedback loops, as individuals remain separate from firms for all practical purposes; while in those that represent multi-layer platforms, firms influence individuals but not the other way round. Analysis of economic network structures suffices to obtain these results, abstracting from any personal relationships that workers (or clients) may have with their peers. Even beyond the specific context of the platform, this suggests that the social network of peers and friends does not suffice to establish whether labor is (or is not) embedded, without also commenting on the economic network of organizations to which workers are tied in various ways.

Future research will have to establish the conditions under which, metaphorically, the number of layers matters more than the number of friends — that is, the economic structure of outsourcing ties affects workers' conditions and outcomes more than their social networks of peers. In online labor, this occurs to the extent that peer networks of workers rarely exist. Siloed platform designs often limit their opportunities to interact with other workers (Gray et al., 2016), and sometimes, platform management tightly controls the few interaction spaces they have, such as web-based discussion forums (Casilli et al., 2019). They do not even have the comfort of sharing their experience with their family, friends, or other people in their broader social environment (sub-section 5.3). Because over-supply of labor and globalization lead to price competition, workers have limited opportunities to organize collectively, and “mostly interact as competitors rather than collaborators” (Graham and Anwar, 2019). Under these conditions, the interactions that matter are those that occur in the economic layers that embed the worker. It would be interesting to see, then, if initiatives to support the creation of independent, non-platform-controlled networking spaces for workers (such as online discussion forums) mitigates or even reverses this effect.

The discovery of deep embeddedness in the case of multi-layer platforms is evidence of the merits of the multi-level view, which completes and refines the conclusions derived by Wood et al. (2019b) based on Peck (2013)'s distinction between hard and soft Polanyi. In the case of marketplace-like platforms (section 5.3), I have made the same claim of disembodiedness that they insist upon. The difference is that they impute this result to commodification (in the hard-Polanyi sense), while I base my argument on the thinness of affiliation ties between individuals and firms ($I - F$). In this respect, the proposed multi-level view is entirely consistent with a broad view of embeddedness and can smoothly integrate the hard and soft dimensions in one common scheme. In sum, far from weakening the grandiose view that pervades Polanyi's work, the network perspective inaugurated with the seminal article of Granovetter (1985) does have potential to advance this program of research. Future studies of embeddedness – regardless of their specific setting – should combine the Polanyian and Granovetterian perspectives, provided the latter is re-interpreted in multi-level perspective.

Of course, this paper has limitations, and the main one is the small empirical basis on which the analysis is based. It does not allow extending conclusions to platforms that do not offer labor mediation services (for example, social media such as Facebook or Twitter), and with its focus on France, it may have missed some forms of platform-enabled outsourcing organizations. For example in the Global South,

there may be more mingling of platform labor with traditional business-process-outsourcing services. Nevertheless, I believe the insight provided here contributes to the (growing, but still scant) literature on platform labor, has value as an example of fruitful cross-disciplinary research, and offers a fresh perspective on embeddedness.

7 Acknowledgements

The data used in this paper are from the DiPLab (“Digital Platform Labor”) research project, co-funded by Maison des Sciences de l’Homme Paris-Saclay; Force Ouvrière, a workers’ union, through IRES (Social and Economic Research Institute); and France Stratégie, a service of the French Prime Minister. Foule Factory - Wirk offered logistical and financial support for fielding the survey, and Inria provided complementary funding. For more information: <http://diplab.eu/>

I am grateful to Antonio A. Casilli, Anoush Margaryan and Florian Schmidt for inspiring discussions that helped shape the arguments developed in this paper. A preliminary version of this paper was presented in a thematic session on “Networks of Economic Life: Relational Processes Within Economic Organizations and Markets” at the INSNA Sunbelt Conference 2019. I thank organizers and participants for helpful feedback. Any remaining errors or inconsistencies are exclusively mine.

References

- Beckert, J. (2003). Economic sociology and embeddedness: How shall we conceptualize economic action? *Journal of Economic Issues*, 37(3):769–787.
- Beckert, J. (2009). The great transformation of embeddedness: Karl Polanyi and the new economic sociology. In Hann, C. and Hart, K., editors, *Market and Society: The Great Transformation Today*, page 38–55. Cambridge University Press.
- Borgatti, S. P. and Halgin, D. S. (2011). On network theory. *Organization Science*, 22(5):1168–1181.
- Botsman, R. and Rogers, R. (2010). *What’s Mine Is Yours: The Rise of Collaborative Consumption*. HarperBusiness.
- Brailly, J., Favre, G., Chatellet, J., and Lazega, E. (2016). Embeddedness as a multilevel problem: A case study in economic sociology. *Social Networks*, 44:319–333.
- Breiger, R. L. (1974). The duality of persons and groups. *Social Forces*, 53(2):181–190.
- Burt, R. (1992). *Structural Holes : The Social Structure of Competition*. Harvard University Press.

- Calnitsky, D. (2014). Economic sociology as disequilibrium economics: A contribution to the critique of the new economic sociology. *The Sociological Review*, 62(3):565–592.
- Carruthers, B. and Uzzi, B. (2000). Economic sociology in the new millennium. *Contemporary Sociology*, 29(3):486–494.
- Casilli, A. (2017). Global digital culture/digital labor studies go global: Toward a digital decolonial turn. *International Journal of Communication*, 11(0).
- Casilli, A. (2019). *En attendant les robots. Enquête sur le travail du clic*. Seuil.
- Casilli, A., Tubaro, P., Le Ludec, C., Coville, M., Besenval, M., Mouhtare, T., and Wahal, E. (2019). Le micro-travail en France. Derrière l'automatisation, de nouvelles précarités au travail ? Final report of project Digital Platform Labor (DiPLab).
- Coase, R. H. (1937). The nature of the firm. *Economica*, 4(16):386–405.
- Codagnone, C., Abadie, F., and Biagi, F. (2016). The future of work in the "sharing economy". Market efficiency and equitable opportunities or unfair precarisation? Report for the European Commission.
- Crossley, N. (2010). The social world of the network. Combining qualitative and quantitative elements in social network analysis. *Sociologica, Italian journal of sociology on line*, (1/2010).
- de Reuver, M., Sørensen, C., and Basole, R. C. (2018). The digital platform: a research agenda. *Journal of Information Technology*, 33(2):124–135.
- De Stefano, V. (2016). The rise of the "just-in-time workforce": On-demand work, crowdwork, and labor protection in the "gig economy". *Comparative Labor Law & Policy Journal*, 37:471–504.
- De Stefano, V. and Aloisi, A. (2018). European legal framework for "digital labour platforms". Report jrc112243, european commission, luxembourg.
- Dominguez, S. and Hollstein, B. (2014). *Mixed Methods Social Networks Research: Design and Applications*. Cambridge University Press.
- Eisenmann, T. R., Parker, G., and van Alstyne, M. (2006). Strategies for two-sided markets. *Harvard Business Review*, 84(10).
- Evans, D. S. and Schmalensee, R. (2016). *Matchmakers: the New Economics of Multi-Sided Platforms*. Harvard Business Review Press.
- Favereau, O. and Lazega, E. (2002). *Conventions and Structures in Economic Organization: Markets, Networks, and Hierarchies*. Edward Elgar.
- Gawer, A. (2009). *Platforms, Markets and Innovation*. Edward Elgar.
- Gawer, A. (2014). Bridging differing perspectives on technological platforms: Toward an integrative framework. *Research Policy*, 43(7):1239–1249.
- Graham, M. and Anwar, M. (2019). The global gig economy: Towards a planetary labour market? *First Monday*, 24(4).
- Graham, M., Hjorth, I., and Lehdonvirta, V. (2017). Digital labour and development: impacts of global digital labour platforms and the gig economy on worker livelihoods. *Transfer: European Review of Labour and Research*, 23(2):135–162.
- Graham, M., Ojanperä, S., and Dittus, M. (2019). Internet geographies: Data shadows and digital divisions of labour. In Graham, M. and Dutton, W., editors, *So-*

- ciety and the Internet. How Networks of Information and Communication are Changing Our Lives*, pages 58–79.
- Granovetter, M. (1985). Economic action and social structure: the problem of embeddedness. *American Journal of Sociology*, 91(3):481–510.
- Granovetter, M. e. (2004). Polanyi symposium: A conversation on embeddedness. *Socio-Economic Review*, 2(1):09–135.
- Gray, M. and Suri, S. (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.
- Gray, M., Suri, S., Ali, S., and Kulkarni, D. (2016). The crowd is a collaborative network. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, pages 134–147.
- Gulati, R. and Gargiulo, M. (1999). Where do interorganizational networks come from? *American Journal of Sociology*, 104(5):1439–1493.
- Hara, K., Adams, A., Milland, K., Savage, S., Callison-Burch, C., and Bigham, J. (2018). A data-driven analysis of workers’ earnings on Amazon Mechanical Turk. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI ’18, pages 449:1–449:14, New York, NY, USA. ACM.
- Howcroft, D. and Bergvall-Kåreborn, B. (2019). A typology of crowdwork platforms. *Work, Employment and Society*, 33(1):21–38.
- Ingram, P. and Roberts, P. W. (2000). Friendships among competitors in the Sydney hotel industry. *American Journal of Sociology*, 106(2):387–423.
- Krippner, G. R. (2002). The elusive market: Embeddedness and the paradigm of economic sociology. *Theory and Society*, 30(6):775–810.
- Lazega, E. (2016). Synchronization costs in the organizational society: Intermediary relational infrastructures in the dynamics of multilevel networks. In Lazega, E. and Snijders, T., editors, *Multilevel Network Analysis for the Social Sciences: Theory, Methods and Applications*, pages 47–77. Springer.
- Lazega, E. (2018). Networks and institutionalization: A neo-structural approach [EUSN 2017 Keynote address]. *Connections*, 37:7–22.
- Lessig, L. (2008). *Remix: Making Art and Commerce Thrive in the Hybrid Economy*. Penguin.
- Lobel, O. (2018). Coase and the platform economy. In Davidson, N., Finck, M., and Infranca, J., editors, *The Cambridge Handbook of the Law of the Sharing Economy*, pages 67–77. Cambridge University Press.
- Murali, A. (2019). How India’s data labellers are powering the global AI race. *Factor Daily*.
- Peck, J. (2013). Disembedding Polanyi: Exploring Polanyian economic geographies. *Environment and Planning A: Economy and Space*, 45(7):1536–1544.
- Polanyi, K. (1957). The economy as instituted process. In Polanyi, K., Arensberg, C., and Pearson, H., editors, *Trade and Market in the Early Empires: Economies in History and Theory*. Free Press and Falcon’s Wing Press.
- Polanyi, K. (2001 (1944)). *The Great Transformation: The Political and Economic Origins of Our Time*. Beacon Press.

- Prassl, J. and Risak, M. (2016). Uber, Taskrabbit, and co.: Platforms as employers? Rethinking the legal analysis of crowdwork. *Comparative Labor Law & Policy Journal*, 37(3).
- Roberts, S. (2019). *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press.
- Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029.
- Rosenblat, A. (2018). *Uberland: How Algorithms Are Rewriting the Rules of Work*. University of California Press.
- Rysman, M. (2009). The economics of two-sided markets. *Journal of Economic Perspectives*, 23(3):125–43.
- Schmidt, C. (2017). Digital labour markets in the platform economy: Mapping the political challenges of crowd work and gig work. Report, Friedrich-Ebert-Stiftung.
- Schmidt, F. (2019). Crowdproduktion von Trainingsdaten: Zur Rolle von Online-Arbeit beim Trainieren autonomer Fahrzeuge. Report, Hans-Böckler-Stiftung.
- Snijders, T. and Bosker, R. (2011 (1999)). *Multilevel Analysis: An introduction to basic and advanced multilevel modeling*. Sage.
- Srnicek, N. (2016). *Platform Capitalism*. Polity Press.
- Sundararajan, A. (2016). *The Sharing Economy: The End of Employment and the Rise of Crowd-based Capitalism*. MIT Press.
- Tubaro, P. and Casilli, A. (2019). Micro-work, artificial intelligence and the automotive industry. *Journal of Industrial and Business Economics*, 46(3):333–345.
- Uzzi, B. (1996). The sources and consequences of embeddedness for the economic performance of organizations: The network effect. *American Sociological Review*, 61(4):674–698.
- Uzzi, B. (1997). Social structure and competition in interfirm networks: The paradox of embeddedness. *Administrative Science Quarterly*, 42(1):35–67.
- Vaughan, J. (2018). Making better use of the crowd: How crowdsourcing can advance machine learning research. *Journal of Machine Learning Research*, 18(193):1–46.
- Wasserman, S. and Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge University Press.
- Williamson, O. (1975). *Markets and Hierarchies : Analysis and Antitrust Implications*. Free Press.
- Williamson, O. (1981). The economics of organization: The transaction cost approach. *American Journal of Sociology*, 87(3):548–577.
- Wood, A., Graham, M., Lehdonvirta, V., and Hjorth, I. (2019a). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1):56–75.
- Wood, A., Graham, M., Lehdonvirta, V., and Hjorth, I. (2019b). Networked but commodified: The (dis)embeddedness of digital labour in the gig economy. *Sociology*, page (online first).

Yin, M., Gray, M., Suri, S., and Vaughan, J. (2016). The communication network within the crowd. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1293–1303.

Festival attendance as a network of attention: A case study in the sharing economy

Paola Tubaro

Abstract

I show that one mechanism through which business events contribute to configuring the shared values of a community is attention – the way in which people allocate their cognitive effort and time. Using a festival of the “sharing economy” as a case study, I construct the set of choices of whom to pay attention to as a network, and I explain its structure with exponential random graph models (ERGM). I find two critical processes underpinning it: a quest for non-redundancy, or novelty, which may go as far as overriding any status effects, and a strategic use of popularity and reciprocity, as ways through which, depending on commonalities with others and on situations of competition, individuals endeavor to adjust their position in the pecking order. These compound that processes result from combinations of simpler mechanisms, can explain how attention-paying choices may engender complex patterns of domination through which different actors and values may gain, lose or maintain legitimacy.

Key words: Events, economy of attention, social networks, ERGM

1 Introduction

A growing literature has emphasized how events such as music awards (Anand and Watson, 2004), literary prizes (Anand and Jones, 2008), book fairs (Moeran, 2011) and film festivals (Smits, 2018) call to mind the metaphor of “tournaments of value” that Appadurai (1986, p. 21) coined to describe the Melanesian *kula* ring :

Laboratoire de Recherche en Informatique, Centre National de la Recherche Scientifique, Paris Saclay, e-mail: paola.tubaro@lri.fr

Complex periodic events that are removed in some culturally-defined way from the routine of everyday economic life. Participation in them is likely to be both a privilege for those in power and an instrument of status contests between them [...] What is at issue in such tournaments is not just status, rank, fame, or reputation of actors, but the disposition of the central tokens of value in the society in question.

What all these events have in common is that they bring together and render visible “imagined communities” (Anderson, 1983) through periodically recurring (albeit temporary) physical co-location. They contribute to giving cognitive legitimacy to an area of activity, making it real and consequential to participants. Their powerful symbolic structure often “takes the form of a public spectacle” (Anand and Watson, 2004, p. 60) with carnival-like elements (Moeran and Pedersen, 2011, p. 5). Events shape collective identity, helping participants define what is shared among them, what should be considered as worthy, where the boundaries with the external world lie. Depending on the circumstances, events may strengthen the dominant values (if any) or challenge them: Delacour and Leca (2011) refer to these two cases, respectively, as “tournaments of value” in the singular and “tournaments of values” in the plural.

The literature on creative industries engages most explicitly with the trope of tournaments of value(s) (or, as Anand and Watson (2004) and Anand and Jones (2008) put it, “tournament rituals”). Management scholars prefer to call them “field-configuring events” (Lampel and Meyer, 2008) and social anthropologists “large-scale professional gatherings” (Leivestad and Nyqvist, 2017), referring to a wider range of settings from policy (Hardy and Maguire, 2010; Schüssler et al., 2014) to high-tech industries (Aspers and Darr, 2011). Despite terminological differences, there is large consensus on the role of events in defining shared values, broadly interpreted as the criteria by which people judge what is legitimate (Suchman, 1995, p. 574). According to Moeran and Pedersen (2011), the idea of tournaments of value(s) covers about the same ground as that of field-configuring events. Either way, focus is on institutionalization of a shared culture and formation of the behavioral scripts that allow actors to be seen by others as legitimate (Aspers and Darr, 2011). Following Simmel’s idea that values are judgements made by people, Moeran and Pedersen (2011, p. 10) sees events as the sites where values are given: not only economic values that result from negotiations, but quality and legitimacy judgements made collectively by members of the concerned community.

While the literature has convincingly shown *that* tournaments do inflect values, it has only scratched the surface of *how* they do so – how tournaments work in practice. If many scholars – from Anand and Watson (2004) to Moeran (2011) – emphasize the importance of networking as a way through which participants collectively negotiate, define and revise values, rare are those who apply formal social network analysis methods to unpack the underlying social mechanisms, like reciprocity or clustering. Brailly et al. (2016) and Favre and Brailly (2015) do so without recurring to the metaphor of tournaments, but come close to it by leveraging Selznick (1957)’s notion of “precarious values”. A value is precarious when it is at risk of losing its carriers and representatives – that is, the active support of relevant individuals or groups that help preserve it against competing values. Thus, mapping

networks of people who carry different values helps detect patterns of domination that eventually impose some values and marginalize others.

To advance this line of research, it is important to acknowledge the effects of differences in the material infrastructure and technical organization of events (Urry, 2003). Informal networking matters most in trade fairs (Brailly et al., 2016; Brailly, 2016; Favre and Brailly, 2015; Moeran, 2011; Panitz and Glückler, 2017), that is, business events where sellers hold stands that buyers visit and most interactions are one-to-one, private meetings that participants arrange autonomously and often informally – although event organizers sometimes provide software applications or other services to facilitate contacts. Negotiations started at the event may afterward lead to signature of contracts between the companies of the buyer and the seller (Brailly et al., 2016) – whereby an initially informal interaction becomes formal. The nature of interactions differs in events shaped as *conferences*, which consist of multiple, sometimes parallel sessions (whether they be talks, panels, workshops etc.). In these settings, participants exchange ideas rather than goods, and interactions occur mainly through *attendance*. It is by attending (for example) a panel that participants become aware of the ideas of the speakers, weigh them against their previously held knowledge and beliefs, and have the opportunity to ask questions. Interactions can be schematized as one (session) to many (members of the audience), and publicly visible to all attendees. While the conference organizers' choice of topics, speakers, schedule and venues constrains participants' options, there is also an important element of decentralization as participants autonomously choose what to attend from a given event program.

What shapes attendance decisions in conference-like settings? Oddly, this question has received hardly any scholarly attention so far, despite the diffusion of the conference format – sometimes as a standalone event (Forrest, 2017), sometimes on the side of trade fairs (Favre and Brailly, 2015). The organization of festivals often resembles that of conferences too (Arcodia and Whitford, 2006). This paper is a first step toward addressing this question, shedding new light on the mechanisms that drive participants' choices.

I frame attendance as the *outcome of individual decisions of what (and more specifically, whom) to pay attention to* in a conference setting. I rely on Goldhaber (1997)'s definition of the conference as a small-sized, closed, temporary economy of attention, as he presented it as part of his keynote at a, well, conference:

While you are here, your main concern is how you pay attention and where you pay it [...] Even between sessions, the exchange of attention is what mostly tends to occupy people at a conference. Of course, there are material considerations [...] But they tend to be secondary issues, taken for granted, and not occupying much attention. We are living a temporary attention economy in miniature right at this moment.

This perspective enriches our understanding of events as tournaments of value(s) with ideas on the economy of attention, supporting the claim that the choice of what / whom to pay attention to is one of the mechanisms through which tournaments take place. Although the processes I pinpoint in what follows are derived from a single case study – that of the OuiShare Fest, a festival in the so-called “sharing economy” – they provide insights of potentially wider applicability. I argue that

paying attention to sessions or speakers that defend precarious values *à la* Selznick (1957) is a way to support them; turning one’s attention to alternative speakers challenges or dismisses them. Thus, attention relationships are one way in which tournaments unfurl, so to speak. I model these relationships as a network to highlight the way they breed interdependencies across actors and sustain alliances, conflicts, and the emergence of hierarchies. My findings invite to see them as channels through which different actors and values gain, lose or maintain legitimacy, and contribute to engendering patterns of domination that may have far-reaching effects – possibly beyond the specific event attended.

2 A “Sharing Economy” event

The OuiShare Fest is an international event that a non-profit organization of the same name organized every year in Paris between 2013 and 2017, with a local edition in Barcelona over 2015-17. It was structured as a conference, with multiple sessions running in parallel (except a few plenaries such as Opening and Closure) over three days. OuiShare presented it as follows:

“OuiShare Fest is an interdisciplinary festival that gathers creative leaders, entrepreneurs, movement builders, purpose-driven organizations and communities from across sectors and countries who want to drive systemic and meaningful change. Though diverse, attendees of OuiShare Fest have many things in common: a willingness to experiment, step out of their comfort zone, collaborate and connect with other humans in a profound way”.

OuiShare and its Fest were created in the wake of enthusiasm for what came to be called the “sharing economy” in the early 2010s (Schor and Cansoy, 2020). The rapid popularity of the “sharing” (or alternatively “collaborative”, see Botsman and Rogers (2010)) label stemmed from its entrenched promise of societal change in addition to digitization and technological innovation. OuiShare and other associations such as Shareable in the USA, and ShareNL in the Netherlands, explicitly leveraged these aspirations to federate stakeholders’ actions and to promote (some version of) this agenda. Surely, the socially engaged flavor of sharing economy discourses sparked controversy, and some saw it as a cover for mere commercial transactions (Slee, 2015), with forms of discrimination (Schor and Attwood-Charles, 2017) and degrading working conditions (Rosenblat and Stark, 2016; Ravenelle, 2019). Nonetheless, many users still see the sharing economy as a means to develop a morally acceptable exchange system, a sense of community, and meaningful working lives (Fitzmaurice et al., 2018).

In what follows, I motivate my choice of the OuiShare Fest as an ideal setting to study how attention choices shape values in a nascent area of socio-economic activity that strives for legitimacy. I then explain how I operationalize the notion of attention-paying.

2.1 The OuiShare Fest as a mini-economy of attention

The OuiShare Fest can be seen as a small-sized, temporary economy of attention as suggested in section 1. Participants did not (and did not intend to) constitute a market or an industry, did not take the role of buyers or sellers, and did not trade: the event was about ideas, not goods. Its core was a rich program from which participants chose what to attend. Every year, the program included manifold session formats, ranging from classical keynotes and panels to interactive workshops and more innovative fishbowls, games, shows, artistic performances and even “unconferences” and “unpanels”. Some sessions took place in traditionally-arranged conference rooms, others in more extravagant venues such as a boat and a tent. Program preparation took several months prior to the event, and resulted both from the organizers’ own choices and from submissions to an open call for session proposals. Once all speakers (also including chairs, moderators, facilitators, performers etc.) and sessions had been scheduled, the organizers aggregated them into thematic tracks, each distinguished by a different color, so that participants could more easily navigate the program and make their attendance choices.

The Annex reports the list of themes for each year of the Paris edition. These themes define the scope of the event as understood by the organizers, and as implicitly agreed to by all those who chose to participate. Their designation and the labels attached to it are an attempt to categorize the sharing economy at a given point in time, assigning priorities and differentially giving legitimacy to specific areas of activity and expertise. Of note, themes denote value cleavages. For example, personal development is an area beloved by many, but frowned upon by others; similarly, decentralization principles arouse the enthusiasm of some and the skepticism of others. Thematic interests are thus indicative of the standing of each speaker in the tournament that the event represents.

That thematic tracks changed over time is in itself a sign of the evolution of legitimacy criteria within this community. Some differences are minor and depend on overall design (for example, number of themes included each year) or terminology (the 2016 “Education & Personal Development” covering about the same contents as “Presence & consciousness” in 2017). Other differences are substantive. “Collaborative consumption”, the term launched by Botsman and Rogers (2010), was prominent at start and disappeared afterward. The city became more and more important, becoming the overarching theme of the 2017 event (“Cities of the world, unite!”). Blockchain and decentralization rose between 2015 and 2016, and vanished in 2017. At each edition, OuiShare took into account the lessons learned from the previous year’s experience to update its offer – that is, in addition to exogenous influences, the outcomes of past tournaments contributed to shaping each new tournament.

2.2 *Networks of speakers and attendees*

The attractiveness of the OuiShare Fest program was embodied in the choice of speakers¹. Formally, participants chose sessions to attend, but ultimately, it was to speakers that their choices were directed. Sessions' titles and contents mirrored the knowledge and interests of those who contributed to them. Speakers were from multiple backgrounds (such as start-up founders, technology specialists, business consultants, policy-makers, academics, charity champions, and artists) and had expertise in diverse areas (the themes discussed above, subsection 2.1). They carried with them different views of the sharing economy, political orientations, cultural norms and values, which they endeavored to convey to participants.

In 2016, speakers constituted about 17% of all participants – against 12% team members, 5% press and 66% non-speaking attendees. As in other events, theirs was a privileged position reserved to a minority (Favre and Brailly, 2015) and made visible to all through a differentiated name tag. Among other things, speakers were entitled to free entrance throughout the duration of the event. They enjoyed high visibility: they all had a profile on the Fest's web page, their talks were advertised to prospective participants through mailing lists and newsletters, they were live-streamed and live-tweeted, and videos were posted online after the event. They were cherished for bringing prestige and epistemic authority: a well-known speaker could not only attract attendees but also raise the standing and respectability of the event.

On this basis, I operationalize choices to pay attention as human relationships between participants and speakers, which can be represented and analyzed with the formalism of graph theory, as is commonly done in social network analysis. This approach is useful to measure and visualize the interdependencies that link actors at the OuiShare Fest, and builds on a rich literature to derive insight into the forces that shape these interdependencies (Amati et al., 2018; Lazega and Snijders, 2016; Lusher et al., 2013). Specifically, the left panel of Figure 1 represents attendee-to-speaker (A2S) relationships as a graph whose nodes are individuals and ties indicate attention choices. The A2S network has “two-mode” structure in that nodes can be either senders (A, attendees) or receivers (S, speakers) of attention ties, but not both. The graph is valued (attendee A_i may choose to attend one or more sessions of speaker S_j) and directed (mutual attendance may not occur).

Speakers were themselves also participants and, outside the session(s) assigned to them for their talks (or exhibitions, performances, workshops etc.), they could choose to spend their free time at the event by attending sessions of other speakers. Besides spreading their own message, they could listen to the messages of others. It becomes then particularly interesting to ask *who, from among speakers, paid attention to whom*. Put differently, paying-attention choices can also take the form of

¹ In what follows, I use this word to cover all roles of those responsible for at least part of the contents of a session, whether they be keynotes, panelists, round table moderators, workshop facilitators, exhibitors, or performers. Similarly, I call “attendees” all those who were not speakers, whether they registered as attendees, press or team (the latter group included OuiShare members and other volunteers who helped the organizers). “Participants” are all those who took part, both speakers and attendees.

speaker-to-speaker ($S2S$) relationships. The central panel of Figure 1 is the $S2S$ attendance graph, which is valued, directed, and one-mode (any node may be sender and/or receiver of an attendance tie).

Juxtaposed, these two graphs form a richer, “multi-level” network structure $A2S2S$ driven by two sources of agency, whereby speakers’ choices are taken as separate from attendees’ choices (right panel in Figure 1) Lazega and Snijders (2016).

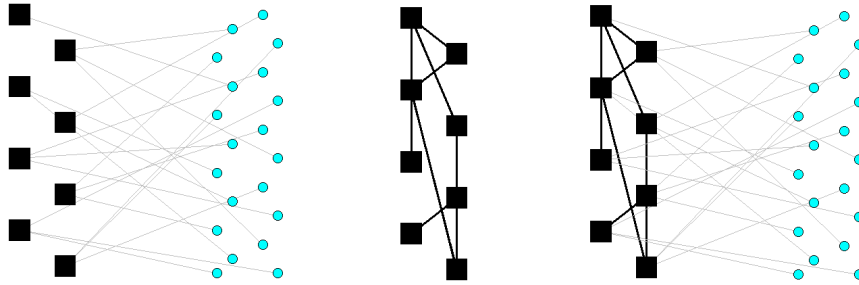


Fig. 1 Left: attendee-to-speaker two-mode network $A2S$, where black squares are speakers (S), cyan circles are non-speaking attendees (A), and light gray lines are attention-paying choices (of A toward S). Center: speaker-to-speaker one-mode network $S2S$, where black squares are speakers, black lines are attention-paying choices (of S toward another S , excluding self-loops). Right: multi-level combination of both $A2S2S$.

$S2S$ relationships are particularly interesting because as noted above, speakers constituted the core and soul of the OuiShare Fest program and were a selling point that organizers used to attract registrations. By paying attention to one another, they could mold the take-home messages of the event, profoundly affecting participants’ understandings and their sense of being part of a common enterprise. Because the way they chose to direct their attention could have far-reaching consequences on the survival of the values they embodied, I make them the focus of the remainder of this paper.

3 Whom do speakers pay attention to?

While attention networks are uncharted territory, the organizational network analysis literature suggests four social mechanisms that can provide the foundation for my subsequent theorizing.

3.1 *Thematic affinity*

People commonly direct their attention to the themes they care for. OUIShare Fest speakers are themselves specialists of some area, topic or theme (say, the blockchain), and may be interested in those other speakers who share the same interests. Similarities facilitate mutual recognition of competences. In the broader network analysis literature, commonality of interests and of areas of expertise is known to trigger connections of some sort – an instance of the principle of “homophily” (McPherson et al., 2001).

In a setting like the OUIShare Fest, thematic affinity can drive choices for two reasons. Paying attention to a speaker who is an expert in one’s own area is a way to be comforted in one’s stances, to see them indirectly confirmed (Centola et al., 2007). It can also be part of an effort to affirm the importance of one’s own area of interest and to legitimate its position in a collective undertaking. This is especially true in uncertain situations and in discursive struggles where inclusion of specific areas or categories is ambiguous or contested (Anand and Watson, 2004).

Brought to the extreme, this criterion would sustain the expectation of a fragmented network of speaker-to-speaker (*S2S*) attention ties consisting of separate cliques defined by areas of interest – with, say, all the blockchain experts attending one another’s sessions but not those of the smart city experts, and vice versa (top left panel, figure 2). Reality is likely more nuanced – if only because there can be speakers whose expertise spans different areas and activities.

3.2 *Competition*

Are speakers competitors and when they are, do they pay attention to each other, or rather avoid each other? By “competitors”, I do not mean speakers who are scheduled at the same time in different venues, because they do not necessarily represent mutually exclusive alternatives in the eyes of participants. If their offer differs, they will appeal to distinct audiences. Further, participants have the opportunity to attend one of them live and watch a video-recording of the other afterward.

To define competing speakers, I rather use the social network analysis notion of *positional similarity*. What is relevant here is the *A2S2S* multi-level participant-to-speaker network (right panel in figure 1), considering all those who pay attention to a given speaker, not just those who are themselves speakers. Two speakers occupy similar – in the limit, identical – network positions (top right panel, figure 2) if they attract essentially the same attendees (whether they are themselves speakers or not) and therefore, display similar patterns of ties in the Fest attention network. Similarity induced by structure is what the literature calls “equivalence” (Lorrain and White, 1971), a notion that goes beyond the mere dyadic level in that two speakers might occupy similar or identical positions without having a tie to each other. The reason why actors with similar ties to others are likely competitors is that they depend on the same resources and face the same opportunities – or re-

sources/opportunities controlled by the same third parties (Burt, 1987; Mizruchi, 1993).

If speakers refuse to pay attention to their competitors, then there will be no ties between them in the *S2S* network whenever they receive ties from the same others.

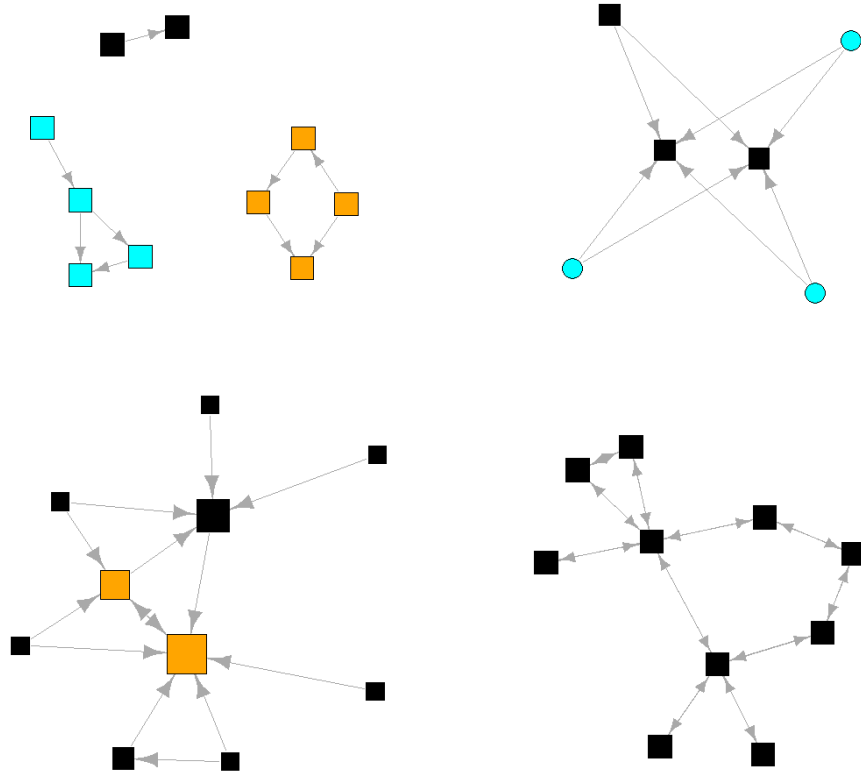


Fig. 2 Schematic representation of relational structures resulting from four different social mechanisms. Squares and circles = actors (nodes), lines = choices (ties). **Top left:** a clustered speaker-to-speaker (*S2S*) network in which choices follow thematic similarities, where node colors represent shared themes. **Top right:** competition resulting from positional similarity in the *A2S2S* network. Two *S* nodes represented as black squares and located at the center, receive choices from the same other nodes, including another node of type *S* represented as a black square, and three of type *A* represented as cyan circles. **Bottom left:** a centralized *S2S* network as may result from status-driven choices, with uneven distribution of received choices. Size of nodes is proportional to number of choices received (in-degree). The orange color represents an observable quality: two of the three most chosen nodes have it, while the third receives choices based on popularity effects only. **Bottom right:** an *S2S* network where all choices are directly reciprocated.

3.3 Status

Regardless of themes, topics and competitive positions, people may just follow perceived patterns of domination and direct their attention to speakers with high status. In other words, they may not defend specific “precarious” values (Selznick, 1957), but just align themselves to the values supported by those that they see at the top of the pecking order. If such motivations are intuitively clear, how do people know the pecking order? Indeed, status is difficult to measure owing to its multi-faceted definition. It can be exogenous and derive from some observable quality of individuals: for example, a known professional achievement or an award in a competition. But status can also be endogenous and result from self-reinforcing processes, whereby initially small quality differences between individuals get amplified over time (Manzo and Baldassarri, 2015). This happens when individual qualities are not observable, and actors attribute status to someone based on cues – that may or may not be accurate. In particular, *popularity* is often taken as an indication of better quality, insofar as absence of precise information may drive people to use the choices of others as a guide towards making their own choices. This tendency results in homogeneous, though not necessarily optimal, behavior (Raafat et al., 2009).

Popularity itself may be a reflection of status through feedback loops. The complexity of status stems from its being both an antecedent of network ties – as high-status individuals are more likely to be selected by others – and a consequence of them – insofar as better-connected individuals are ascribed higher status (Torlò and Lomi, 2017). The very relationship that this paper focuses on – paying attention to someone – can be seen as a source of status. A speaker who pays attention to another speaker confers to the latter some degree of honor, respect or prestige. Status can be seen as the accumulation of such acts (Podolny and Phillips, 1996). Like other deference-conferring gestures, attention-paying establishes an ordering or ranking of individuals, and within a bounded setting such as an event, it can be represented as a network (Sauder et al., 2012). This is one reason why status rankings may reinforce themselves and diverge from distribution of observable qualities.

If speakers are more likely to pay attention to other speakers with high status (however construed and operationalized), one should expect to observe an *S2S* network where almost all nodes are directly or indirectly connected to all other nodes, in a centralized structure in which some “stars” receive many choices and the great majority receive very few (bottom left panel, figure 2). There are, however, limits to the pecking order that may result from status-based choices, due to the symbolic “cost” of attending another speaker’s session – the cost of recognizing the others’ status, implicitly admitting that there is something to learn from them (Blau, 1964). This cost may act as a disincentive.

3.4 Reciprocity

Do speakers tend to reciprocate the choice of other speakers who pay attention to them? Reciprocity is a known powerful driver of human behavior and an antidote to status hierarchies. According to Gould (2002, p. 1151), “Someone who pays less attention to you than you pay to her implicitly asserts that she is superior to you in status. If you do not respond by withdrawing your attention, you have implicitly agreed”. He submits that people are less willing to demonstrate esteem toward those who do not return the favor and conversely prefer to receive such demonstrations without reciprocating. When it prevails, the desire to have one’s gestures reciprocated curbs the self-reinforcing character of status attributions.

Reciprocity may sometimes be replaced by group-based mechanisms (such as clustering) that sustain one-sided ties without sharpening status hierarchies. Put differently, a social group may provide a forum for interaction that mitigates status even when ties are not reciprocated (Block, 2015).

In addition to direct reciprocity, there may be a form of indirect reciprocity, which according to Nowak and Sigmund (2005), comes in two flavors. Upstream indirect reciprocity occurs when a speaker B who has received attention from speaker A, feels motivated to demonstrate esteem to someone too – and chooses to do so toward speaker C. Downstream indirect reciprocity is based on reputation and is observed when speaker A has paid attention to B and therefore receives attention from C. With indirect reciprocity, speakers who pay attention to many others, also receive many choices: this is another mechanism that may mitigate status hierarchies.

If speakers (directly) reciprocate attention choices, then the resulting relational structure will tend to be symmetric and can be represented as a graph where most ties are mutual (bottom right panel of figure 2). Indirect reciprocity will give rise to a network structure with high correlation between in-degrees (number of choices received) and out-degrees (number of choices directed to others).

4 Empirical setting and data

I explore these questions with data from the fourth OuiShare Fest, which took place on 18-20 May 2016 in Paris. I use back-office data from the event’s software application, “Sched”, and the register of participants kept by the organizers. Sched is an online and mobile tool to present oneself to others, connect, and choose a personalized program: each participant had the opportunity to create a profile (with name, photo, organization, job title, self-description), import contacts from social media such as Twitter and LinkedIn (from among those who were also attending the event), and choose a schedule (from among up to seven simultaneous parallel sessions going on at any one time during the event). The latter function was particularly attractive as no printed version of the program was available to participants, in line with the “Zero-Waste” policy that OuiShare promotes. Sched is designed to

suit the needs of a population that makes intense use of, and is comfortable with, digital technologies.

App data are an ideal source of information to address people’s attention-paying choices as discussed above (section 2). By taking the time to choose their agenda on Sched, participants devoted attention to the program and expressed interest in a selection of sessions and their speakers. Of note, their choices were visible to all online, being displayed on the Sched personal page of the choosing participant, on the page of the chosen session and the pages of all its speakers (also including moderators, facilitators etc.). They are publicly-visible endorsements, akin to the “Like” function of some social media platforms.

Sched data are traces of an attention exercise that participants did online prior to the event, based solely on their interests. They do not mirror practical factors such as time clashes: the application let users choose sessions that took place at the same time in different rooms (session-hopping was allowed), and there was no check-in at the entrance of sessions to monitor actual behaviors.

Of the 1830 registered participants in 2016 (of whom 1505 actually attended the event at least for one day), about two fifths have completed Sched pages. Importantly, this includes all the 258 speakers, whose profiles had been pre-filled by the organizers in advance, noting the sessions in which they were involved as speakers, panelists or moderators. 88 of these speakers chose a program of sessions to attend beyond the one(s) in which they were scheduled to be involved. 475 non-speakers did so. A logistic regression (Annex) shows that controlling for time spent at the event, participation in previous editions, place of residence, and gender, speakers were less likely than others to choose a schedule on Sched. As suggested in subsection 3.3, paying attention to others has a “cost” that not all are prepared to pay, whilst such cost is lower (or perhaps non-existent) for non-speaking participants, whose main reason to attend the event is the program. Only press representatives were less likely than speakers to use Sched – presumably because they had narrow reporting goals rather than an overarching interest in the Fest.

4.1 Dependent variable

The dependent variable Y is the $S2S$ network of speakers paying attention to other speakers. It is a dyadic variable, whose values are defined for each pair of the 258 OuiShare Fest 2016 speakers. Disregarding self-attendance, it corresponds to $258 \times 257 = 66306$ dyads. Each value Y_{ij} (with $i, j = 1, \dots, 258, i \neq j$) is equal to the number of times speaker i chose to pay attention to speaker j by choosing one of their sessions on Sched. Y_{ij} will differ from Y_{ji} whenever the choice of i to pay attention to j is not reciprocated.

The $S2S$ network is almost fully connected, with a main component of 255 nodes and 3 isolates – speakers who neither send out nor receive choices. Figure 3 illustrates differences between speakers depending on the number of attention choices they receive from, and direct to, other speakers, also known as their in-degrees (left

panel) and out-degrees (right panel). Both distributions are right skewed, consistently with most empirical social networks: positive skewness indicates higher variation in relational activities than would be produced by a random process, and points to differences in the outgoing attitudes of individuals (in the case of out-degrees) and differences in status, prominence or prestige in the case of in-degrees (Ahuja et al., 2009; Pallotti et al., 2013).

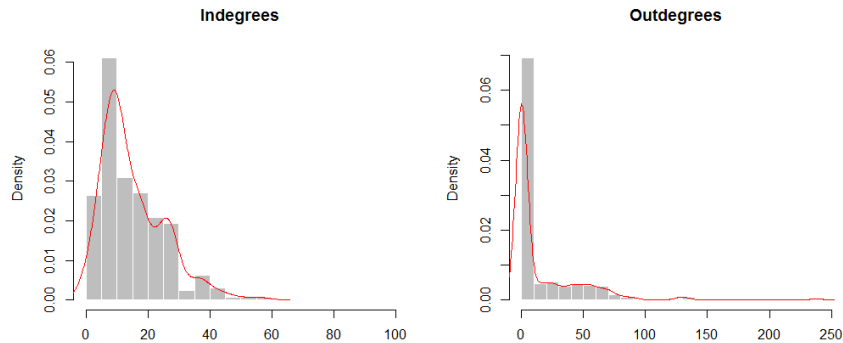


Fig. 3 Degree distributions. In-degrees (left panel) and out-degrees (right panel) in the *S2S* network of 258 speakers. Histograms (gray bars) and density plots (red curved lines).

In what follows, for analytical convenience and without significant loss of information, the values of Y are binarized (that is, each $Y_{ij} > 0$ is set to 1).

4.2 Independent variables

Independent variables are also defined at dyadic level, with one value for each of the 66306 pairs of speakers. They measure the different social mechanisms that may affect the network structure Y .

4.2.1 Thematic affinity

To remain as close as possible to the experience of Sched users, I map speakers to themes, or areas of expertise, using the categorization proposed by the event organizers, who had assigned each session to either a general category (for sessions such as Fest Opening and Closure), a satellite event, or one of six following themes: “Building Enterprises for the Digital Age”, “Digital Institutions and The City”, “Education and Personal Development”, “Power and Capital in the 21st Century Orga-

nization”, “The (Present) Future of Work”, “Understanding Decentralization and The Blockchain” (subsection 2.1).

Figure 4 represents the extent to which speakers who are competent in each of the six themes (rows) pay attention to speakers in theirs and other themes (columns)². It appears clearly that most choices come from speakers involved in sessions on “Education and Personal Development”, “Power and Capital in the 21st Century Organization” and to a lesser extent “Building Enterprises for the Digital Age” and “Understanding Decentralization and The Blockchain”, while the themes that receive most choices are “Power and Capital in the 21st Century Organization”, “Digital Institutions and The City”, and to a lesser extent “Building Enterprises for the Digital Age” and “Understanding Decentralization and The Blockchain”. On the one hand, these choices differ from those of the 475 non-speaking attendees who have a Sched schedule, whose preferred theme is “Digital Institutions and The City” (22.9%), followed by “Building Enterprises for the Digital Age” (20.7%) and “Power and Capital in the 21st Century Organization” (17.6%). Thus, speakers’ choices have specificities that are worth exploring. On the other hand, same-theme ties are overall rarer than cross-theme ones. This is enough to rule out the possibility that the network is split into separate thematic components.

Formally, this is a dyadic, binary, symmetric variable TA (table 1) whose value $TA_{ij} = TA_{ji}$ is 1 if speakers i and j (with $i \neq j$) have at least one theme in common, 0 otherwise.

4.2.2 Positional similarity

Attention choices not only result from exogenous factors such as thematic affinity, but are also partly self-sustained through an endogenous tendency to reproduce existing network structures. As discussed in sub-section 3.2, I interpret competition between speakers as the extent to which they attract attention from the same participants, whether they be themselves speakers or simple attendees (positional similarity). This requires examining the attendance choices of all OuiShare Fest participants who used Sched to build a personalized online program, in the expectation that they will inflect speakers’ choices of whom, among other speakers, to pay attention to. Technically, this means calculating an indicator that reflects similarity in incoming ties within the $A2S2S$ network structure (right panel of figure 1), as a potential determinant of $S2S$ (central panel of figure 1).

A commonly used measure of positional similarity is the Jaccard index which, for any two speakers, is defined as the number of common participants (intersection),

² Thematic expertise is attributed to speakers based of the theme(s) of the session(s) in which they speak. Each speaker can thus be expert in one or more themes. I follow Everett and Borgatti (2017) in representing these “overlapping categories” as a person-by-category affiliation matrix $B = (b_{ij})$, of dimensions 258×6 . Matrix $C = AB$, where A is the 258×258 square matrix of speakers’ attention-paying choices, is a speaker-by-theme matrix that measures speakers’ access to different themes not through their own expertise, but through their attention-paying choices. The heatmap of figure 4 represents the product matrix $B^T C = B^T AB$.

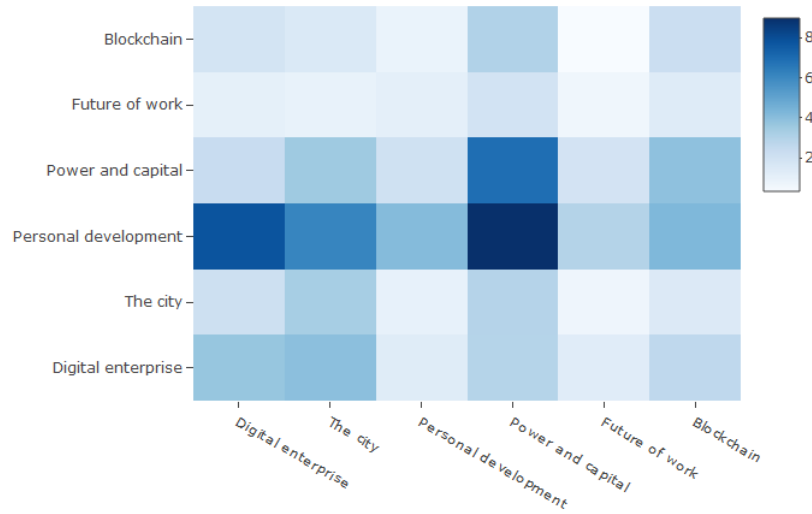


Fig. 4 Heatmap representing the extent to which speakers with expertise in each of the six themes of OuisShare Fest 2016 (rows) pay attention to speakers in theirs and other themes (columns). Darker colors represent higher numbers. Interpretation: the highest number of choices are from speakers in the “Education and Personal Development” theme toward speakers in the “Power and Capital in the 21st Century Organization” theme.

divided by the cardinality of the union (sum of the participants of each of the two speakers, both separately and jointly). Scaling by the size of the union rules out the possibility that similarity scores are affected by the tendency of more prominent speakers to receive more choices than lesser-known speakers (Borgatti, 2012). This dyadic variable $J_{ij} = J_{ji}$ can be computed for each pair of speakers, is symmetric and varies between 0 and 1 (table 1).

4.2.3 Status

What matters here is the status of the receiver, rather than the sender, of attendance choices – the potential object of others’ attention. Hence for each of the 66306 pairs of speakers $[i, j]$, focus is on the status of j (see table 1). As discussed in subsection 3.3, some dimensions of status are exogenous, and I retrieve them from speakers’ Sched profiles and from registration information. Variable St_j^1 is binary and indicates whether speaker j was the founder of a sharing economy start-up, think-tank, charity or other type of organization. Founders are entrepreneurs who invest their talent and resources to advance the sharing economy, and are highly respected in this community that values initiative, innovation and dynamism. Variable St_j^2 is also

binary and indicates whether speaker j had already participated in the 2015 and/or 2014 OuiShare Fest, to denote experience in, and familiarity with, this milieu.

As discussed above, some status cues that can be used to infer quality when it cannot be objectively measured, depend on the network structure itself. Again, I need to include status indicators that reflect the $A2S2S$ network structure, under the expectation that they affect $S2S$. One of these indicators is the number of attention choices that each speaker received through Sched – a proxy for popularity (subsection 3.3). To retrieve this information, I sum over the attendance choices of all participants who had a Sched profile (variable St_j^3 , taken on $A2S2S$). A speaker could be chosen more than once by the same participant.

The other indicator that can be computed from the $A2S2S$ network structure is the number of followers that speakers had on social media, notably Twitter and LinkedIn, insofar as it was imported into, and made visible through, Sched. Social media are relevant because Fest participants are heavy technology users, interacting online as much as face-to-face. OuiShare has always actively used Twitter and LinkedIn to communicate, creating a large online network around itself, and all its events were live-tweeted. Being part of the conversation signals online notoriety and legitimacy within this community. Technically, I count the number of times each speaker appears as an online contact in the Sched pages of OuiShare Fest participants, summing over all those (speakers and non-speakers in the $A2S2S$ structure) who had linked their Sched profiles to their social media accounts (variable St_j^4).

4.2.4 Control variables

Other potentially relevant factors are participation in the previous edition of the event, 2015, which may facilitate integration of participants (variable $T1$), place of residence (distinguishing, for simplicity, France where 50% of participants and about 40% of speakers came from, and other countries: variable Fr) and gender (variable Ge). They are originally monadic and taken from the organizers' register. To adapt them to the dyadic structure of the data, I compute an exact match indicator for each variable, to identify pairs of speakers in the same category – for example, both from France (table 1).

5 An Exponential Random Graph Model of attention choices

Attendance choices are driven both by exogenous and endogenous factors. I have already introduced the number of choices received by a speaker as a measure of status, and the Jaccard index of positional similarity – two measures that I take on the multi-level relational structure $A2S2S$, and that I use to explain the speaker-to-speaker ($S2S$) relational structure. The presence of such systematic dependencies between network ties requires explicit modeling and cannot be left to standard sta-

Table 1 Dyadic variables included in empirical model specification: definitions and descriptive statistics. Those in italics are endogenous to the multi-level A2S2S network.

Variable	Construct	Min	Max	Mean	St.Dev
Speaker-to-speaker attendance	Dependent variable	0	1	0.059	0.235
TA: Common themes	Thematic affinity	0	1	0.234	0.538
St1: Founder of organization	Status (receiver)	0	1	0.31	0.463
St2: Experience		0	2	0.527	0.753
<i>St3: Number choices</i>	<i>Status (receiver)</i>	<i>4</i>	<i>500</i>	<i>105.7</i>	<i>72.407</i>
<i>St4: Number online followers</i>		<i>0</i>	<i>63</i>	<i>4.77</i>	<i>10.642</i>
<i>Ja: Jaccard index</i>	<i>Positional similarity</i>	<i>0</i>	<i>1</i>	<i>0.145</i>	<i>0.119</i>
T1: Participation previous year		0	1	0.314	0.465
Fr: Residence in France	Control (match)	0	1	0.376	0.485
Ge: Gender		0	1	0.616	0.487

tistical inference approaches which assume independence of observations (Pallotti et al., 2013).

Recently developed Exponential Random Graph Models (ERGM) allow specifying the dependence among relations that such micro-mechanisms entail, in addition to modeling the effects of exogenous covariates (Robins et al., 2009; Snijders et al., 2006). These models address a fundamental issue that arises with network data: the presence of dependencies across observations as one node may be involved in multiple relationships with other nodes. They do not limit themselves to “controlling for” such dependencies, but go as far as representing them directly, uncovering specific local (sub-graph) configurations of ties implied by dependence mechanisms, and testing their statistical significance (Amati et al., 2018).




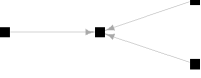
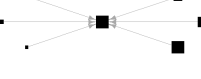
With ERGM, one may go as far as exploring effects that are fully endogenous to the S2S network structure (table 2). Reciprocity, a potentially important determinant of choices (sub-section 3.4), is one such effect: only speakers can reciprocate an attention choice by other speakers. Direct reciprocity is the tendency of the network structure to be symmetric, and concerns 7% of ties in this network. Indirect reciprocity is the tendency of outgoing and incoming ties to co-occur (technically, a mixed-two-paths configuration, which can represent both upstream and downstream indirect reciprocity), which entails a tendency of in-degrees and out-degrees to be correlated.

Some status-related effects are also fully endogenous to the S2S network (sub-section 3.3). They capture differences in the propensity of individual speakers to be attended by other speakers – independently of the other choices they might receive from the broader population of participants. Such differences may create status disparities regardless of other individual qualities. To account for this possibility, I

include so-called in-degree effects, which measure popularity interpreted as the extent to which attention choices are more likely to be directed at speakers who already have incoming ties, all other things equal.

Specifically, I use two-stars, three-stars and geometrically-weighted in-degree. The meaning of two- and three-star effects is intuitive: the more edges on a node, the more two- (or three-) stars an additional edge will create (table 2). The in-degree effect is an higher-order effect, suitable for cases in which the in-degree distribution is uneven. It estimates the change in tie likelihood given the in-degree of the nodes involved, and it is “geometrically weighted”, that is, with marginally decreasing weighting as degree increases. The weights avoid degeneracy, a common problem with ERGM estimations³.

Table 2 Structural dependencies to be modeled, representing endogenous feedback effects to the S2S network.

Variable	Construct	Configuration
DR: Reciprocity	Direct reciprocity	
IR: Mixed-two-paths	Indirect reciprocity	
2I: 2-instar		
3I: 3-instar	In-degree effects	
gwI: geometrically weighted in-degree		

Despite requiring a major departure from the way data structures are commonly conceptualized in the social sciences, ERGM provide the same approach to inference, whereby an outcome variable is predicted by several independent variables. Parameter values help identify a probability distribution for all graphs of the same

³ Degeneracy occurs when the ERGM estimation process produces graphs that are complete, or empty, or with all edges concentrated in a small region of the graph, or otherwise irrelevant (Handcock, 2003; Snijders et al., 2006).

size, and estimates aim at finding the parameter values that best match the observed network structure. This is obtained via Markov Chain Monte Carlo Maximum Likelihood (MCMCML) or other simulation-based techniques; interested readers may find more statistical details in Lusher et al. (2013).

6 Results

Table 3 reports ERGM estimates of a model in which the dependent variable $Y = S2S$ is estimated on the basis of the exogenous covariates and endogenous network dependencies outlined above. Among endogenous configurations, I do not include transitive closure, widely used in the literature but not salient in this context (meaning that two speakers who pay attention to the same third speaker are also likely to choose each other). Neither do I model out-degree effects: rather, I constrain the estimation process to preserve observed out-degrees (so that only networks whose out-degrees are the same as those in the data have non-zero probability). Indeed, speakers' overall reluctance to make online attendance choices (Annex) suggests that the specificity of the few Sched-using speakers is not an incidental occurrence but a constitutive element of the relational structure under study. The model is estimated with R package 'ergm' in the Statnet suite (Handcock et al., 2003; Hunter et al., 2008).

6.1 Estimates

Let us first comment on quality of estimations. The MCMC diagnostics (not reported here) indicate good convergence: the MCMC sample statistics vary randomly around the observed values at each step and the difference between the observed and simulated values of the sample statistics have a roughly bell-shaped, zero-mean distribution indicating stationarity.

Goodness-of-fit can be assessed following the approach of Hunter et al. (2008) who suggested that the network statistics measured in the original data should be reproduced reasonably well on networks simulated from the fitted model (Li and Carriere, 2013). The observed statistics can be plotted against the simulated statistics to assess whether ERGM fits the data well. The Annex includes plots of the statistics commonly used to measure goodness-of-fit, showing that the model represents key patterns in the data rather well, though it slightly amplifies the skewness of the in-degree distribution.

Table 3 ERGM Maximum Likelihood estimates. Dependent variable: network of attention choices $S2S$, independent variables: covariates and endogenous structural dependency statistics outlined in Tables 1 (upper section) and 2 (lower section).

Variable	Estimate	Std. err.	p-value	Signif. ^a	Odds-ratio
Thematic affinity					
Mutual ($TA_{ij,ji}$)	0.9118	0.1963	< 1e-04	***	2.4888
Upper (TA_{ij})	0.2904	0.0734	< 1e-04	***	1.3369
Lower (TA_{ji})	0.4507	0.075	< 1e-04	***	1.5694
Positional similarity					
Mutual ($JA_{ij,ji}$)	-0.2875	0.0472	< 1e-04	***	0.7502
Upper (JA_{ij})	0.038	0.2657	0.8861		1.0388
Lower (JA_{ji})	-0.1027	0.2756	0.7094		0.9024
Status					
Receiver founder	0.0606	0.0375	0.106		1.0624
Receiver experience	-0.057	0.0286	0.0463	*	0.9446
Receiver No. attendees	0.0098	0.0009	< 1e-04	***	1.01
Receiver No. followers	-0.0033	0.0022	0.1334		0.9967
Controls					
Co-presence past year	-0.1128	0.0431	0.0089	**	0.8934
Both from France	0.2248	0.0413	< 1e-04	***	1.2521
Same gender	0.0437	0.0417	0.2938		1.0447
Structural dependencies					
Reciprocity					
Direct reciprocity	0.127	0.174	0.4655		1.1354
Indirect reciprocity	0.0001	0.0008	0.873		1.0001
Indegree effects					
2-instar	0.0804	0.0091	< 1e-04	***	1.0838
3-instar	-0.0034	0.0004	< 1e-04	***	0.9966
Gwidegree ^b	-1.6607	0.3274	< 1e-04	***	0.19

^a Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

^b Geometrically-weighted in-degree (section 5).

N = 66306, AIC: -2133, BIC: -1969

6.2 Interpretation

Interpretation of specifications and outcomes closely resembles that of a logistic regression: generally, a large positive (negative) parameter indicates that the corresponding configuration is observed in the network under study more (less) frequently than what would be expected by chance alone, conditional on the presence of configurations associated with other effects in the data.

6.2.1 Thematic affinity

Thematic affinity does not go as far as splitting the community into separate islands, but it has a positive and significant effect: speakers are more likely to give and receive choices from others with similar interests, and this is true whether their choice is reciprocated (“mutual” effect) or not (ij and ji combinations respectively). Their behavior is highly cooperative provided there is a shared interest: they give back when they receive, but are also happy to give without a return. As suggested in sub-section 3.1, this may be because choosing like-minded others may reinforce one’s thematic interests and contribute to legitimizing them. In this sense, Sched attendance choices are a way for speakers to promote their ideas beyond their own talks, by publicly displaying interest in those similar to them.

6.2.2 Competition

Competition, interpreted as positional similarity (sub-section 3.2) has a negative effect only on mutual (reciprocal) ties. Put differently, if speakers i and j attract largely the same attendees, then they are less likely to choose each other, but their competitive position does not significantly affect the likelihood of a non-reciprocated tie between them (ij or ji). Speakers actively counter reciprocity when they compete for attention of the same third parties – hence, for the same resources. Gould (2002) theorizes that people are torn between the desire to receive demonstrations of esteem from others without returning the favor – because this raises their status – and the desire to see their own demonstrations of attention to others reciprocated – because this avoids lowering their status (sub-section 3.4). The finding reported in table 3 suggests that the former tendency prevails under competitive conditions.

6.2.3 Status

Regarding the observable dimensions of status (sub-section 3.3), being founder of an organization does not matter, while surprisingly, experienced speakers are slightly less likely to receive attention from other speakers. This result is best understood in conjunction with the effect of one control variable, co-participation the year before, which is negative and significant. Together, they mean that all other things equal, a returning speaker is less likely than a newcomer to attend the talk of another returning speaker. This attitude is in line with OuiShare’s explicit orientation towards “a willingness to experiment, step out of one’s comfort zone” (section 2). Interestingly, preference for novelty seems to counter some dimensions of status – those, at least, that can be associated to durable involvement in the community. A trade-off appears between novelty/openness and continuity/fixed anchors.

The model includes several proxies for the unobservable dimensions of status, notably popularity (sub-section 3.3). They are operationalized in the A2S2S network by counting number of online followers and number of attendees: the former

has no effect while the latter is positive and significant. That this finding resonates with the literature on herding behavior does not mean that the choice to attend the most attended should be interpreted as mindless imitation. In an international event that hosts multiple views and standpoints, a large number of attendees may denote a *diverse* audience. Thus, following their choices may signal (again) a quest for novelty, a sense of curiosity for what others like. If the preference for novelty counters some dimensions of status, as noted above, it is consistent with others.

Other popularity indicators are in-degree effects, all significant: two-in-star is positive, meaning that ties are more likely to be directed at nodes that already have an incoming tie, while three-in-star and geometrically weighted in-degree distribution, a higher-order effect, are negative, meaning that after accounting for all other model terms, nodes with large numbers of ties are less likely to receive more ties. Taken together, these results indicate that a speaker does not want to be the first to pay attention to another speaker, but does not follow mere popularity either. Distinction from peers is valued, but nobody would go as far as taking the risk of setting entirely new trends. That speakers value distinction while being reluctant to overly differentiate themselves suggests a form of conformity as often observed in middle-status individuals Phillips and Zuckerman (2001).

6.2.4 Reciprocity

Neither direct reciprocity nor indirect reciprocity (sub-section 3.4) are significant. Paying attention to someone else does not increase one's odds of receiving attention. Further, using Sched to choose one's own agenda does not affect the chances of receiving attention choices by others, so there is no tendency toward correlation of in- and out-degrees. These tendencies are better understood in light of the results already outlined: if in general, reciprocated and unreciprocated ties are equally likely to be observed, speakers are more likely to form mutual ties when there is the comfort of thematic affinity between them (sub-section 6.2.1), while they are less likely to do so when there is the discomfort of competing for the same attendees (sub-section 6.2.2). Put differently, collaborative contexts favor reciprocity, while competitive conditions hinder it.

The effects of reciprocity (or lack of it) are important because they inevitably spill over onto status (sub-section 6.2.3). Reciprocity within same-theme network niches may mitigate status effects, but lack of an overall tendency to reciprocity in the network (and active resistance to it among competitors) tend to reinforce them. Overall, the interplay of these opposing social mechanisms results in mild status hierarchies. This is consistent with a setting such as the sharing economy whose definition, scope and boundaries were still moving when the research took place, and there were no strict social norms to regulate how participants should pay attention to, and recognize the value of, one another.

Overall, this setting is one in which actors have multiple forms of status that do not systematically strengthen one another. Some interactions that generate, maintain

or exacerbate status hierarchies, coexist with others that mitigate them. Newcomers are welcome, while struggles to gain prominence occur: actors compete for attention and make choices instrumentally to discover novelties and to promote their themes of interest.

7 Conclusions

The motivating question of this paper was, what shapes people's choice to pay attention to others in conference-like settings. The choice of whom to pay attention to is one of the mechanisms through which values (broadly interpreted as themes, orientations and areas of activity) gain or lose legitimacy within a community. The speakers who give and receive attention in a conference carry with them values that are essentially precarious, whose survival depends on the support they receive (Selznick, 1957), especially in a nascent area of activity such as the sharing economy in 2016.

The above analysis suggests two main answers. The first is a quest for novelty: people pay attention to those who offer new experiences, who get them out of their comfort zone as *OuiShare* put it. The desire for novelty may or may not lead to the choice of high-status others: in particular when status involves some form of rigidity, whether real or perceived, novelty prevails (sub-section 6.2.3). At first sight, this result seems at odds with the literature which, focusing mostly on trade fairs (section 1), finds that ties are more likely to be formed after a history of co-participation (Brailly et al., 2016). But the structure of interactions differs. In trade fairs, researchers map networks of informal contacts through which participants prepare future commercial transactions: connecting with previous acquaintances is a way for them to reduce uncertainty and facilitate negotiations. Instead in conferences, where the key interaction is attendance at sessions and the goal is to gain knowledge or insights rather than contracts, new speakers are more likely to be non-redundant. This is why novelty is appreciated. Research in social network analysis has made clear that when the stakes are of informational or cognitive nature, redundancy is an issue, while access to non-redundant ideas brings advantages (Burt, 1992, 2004). Thus, one way in which speakers' attention-paying choices create value in an event like the *OuiShare Fest* is through their orientation toward novel, non-redundant content.

The second answer is a strategic use of popularity and reciprocity. Participants fine-tune their use of popularity cues to affirm their standing while limiting the risks of status loss that might arise if they were the first to pay attention to someone else, especially under conditions of uncertainty (sub-section 6.2.3). Similarly, and perhaps ironically in a festival formed around the premise of sharing, people use reciprocity (sub-section 6.2.4) as a flexible instrument to manage their position in the social structure, depending on whether their goal is to secure support for their thematic niche (sub-section 6.2.1) or to stay afloat in situations of competition. It would be too hasty to conclude that these results challenge received wisdom that

popularity is always (more or less) self-reinforcing, and that dislike of symmetry is uncommon. Rather, the insight they provide is that popularity and reciprocity themselves are too general principles to directly explain behaviors. They subtly combine with other relevant social mechanisms (here, thematic homophily and competition) and play out differently depending on the resulting mix of factors. In the specific case of the OuiShare Fest, limited and selective use of reciprocity contributes to generating a status hierarchy, which creates value by providing signals of quality when other information is lacking. Conversely, circumspect reliance on popularity keeps this hierarchy mild, and creates value by avoiding a rigidity that would hamper creativity and innovation. Together, they manage status in this community so that its above-discussed drive for novelty does not entirely exclude some degree of fixed anchoring, and its homophily does not split the community.

Future research will establish how general these empirical findings are. For example, it could be that competitive events such as awards display sharper status differences, based on different combinations of reciprocity, popularity, or homophily. A specific limitation of my study stems from the design of the digital application used to make attendance choices. Perhaps, modified software features might trigger more (or less) intense use of Sched: for example, one might ask whether a more user-friendly application could lower the cost of making attendance choices.

What is general beyond my setting is the importance of considering the joint effects of multiple relational mechanisms (reciprocity, popularity, homophily and so on), each interfering with the others in complex ways. Within the narrower literature on tournaments of values (also called field-configuring events or large-scale professional gatherings as discussed earlier), my results suggest the need of distinguishing between types of interactions, as cognitive-oriented attention choices and (for example) trade-oriented informal networking do not rely on the same social processes. Given the wide diffusion of attention-based events such as conferences and festivals, and their potential impact on the communities that they gather, the ideas discussed here may be of broader interest than the specific setting considered. Attempts to replicate the results reported above and extend them to different settings will stimulate further reflection on the determinants and consequences of attention-paying choices.

8 Annex

8.1 Oui Share Fest session themes, 2013-17

Table 4 OuiShare Fest session themes, with number of sessions per theme, 2013-17. Non-thematic sessions (such as Opening and Closure) and scheduled leisure and breaks (coffee and lunch breaks, conclusive “love” party etc.) not included.

2013		2014		2015		2016		2017	
Themes	N. sessions	Themes	N. sessions	Themes	N. sessions	Themes	N. sessions	Themes	N. sessions
Business & Design	7	Co-Create	21	Capital	7	Building Enterprises for the Digital Age	23	APIs for new urban services	1
Collaborative Consumption	17	Connect	17	Cities	5	Digital Institutions & The City	25	Fighting climate change	10
Economics & Society	17	Debate	28	Culture	10	Education & Personal Development	30	Future of work	9
Finance & Money	8	Inspire	43	Decentralization	9	Power and Capital in the 21st Century Organization	14	Living together	49
Makers & Open Manufacturing	13	Play	12	Experience & Other	39	The (Present) Future of Work	14	New technologies for impact	25
Open Knowledge	8			Makers	9	Understanding Decentralization & The Blockchain	15	Presence & consciousness	12
UnConference	12			Open Knowledge	8			Social justice	12
				Organizations	23			Taking power back	18
				Politics	9				
		Sharing economy	18						
		Social impact	7						
		Sustainability	9						
		Work	9						
Total	82		121		162		121		136

8.2 Logistic regression model of posting attendance choices

Table 5 reports the results of a logistic regression model of posting attendance choices on Sched. The dependent variable takes values 1 = Yes, 0 = No. N = all persons registered to OuiShare Fest 2016. Covariates are Presence (number of days spent at the event, between 0 and 3); Experience (number of Fest participations in previous two years, between 0 and 2); residence in France (1 = Yes, 0 = No); Gender (1 = Male, 0 = Female); and type of participant (Attendee as reference category; other types: Attendee 1 day, Press, Speaker, Team).

Table 5 Results of a logistic regression model of posting attendance choices on Sched^a.

	Estimate	Std. Error	z-value	Pr(> z)	Signif. ^b	Odds-ratios
Presence	0.8514	0.0667	12.759	< 0.0000	***	2.343
Experience	-0.008	0.1009	-0.079	0.9368		0.992
France	-0.4283	0.1224	-3.499	0.0005	***	0.6516
Gender	-0.0295	0.1123	-0.262	0.793		0.971
Attendee	<i>ref.</i>					
Attendee 1 day	-0.2798	0.1946	-1.438	0.1505		0.7559
Press	-0.6036	0.2859	-2.112	0.0347	*	0.5468
Speaker	-0.5516	0.1648	-3.348	0.0008	***	0.5761
Team	0.0855	0.1705	0.501	0.6162		1.0892
(Intercept)	-1.85	0.1892	-9.78	< 0.0000	***	0.1572

^a N = 1830, AIC: 1942.1

^b Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

8.3 ERGM goodness-of-fit diagnostics

In figure 5, the observed network is compared to networks simulated from the fitted ERGM in terms of standard summary statistics: in-degree (number of incoming ties of a node), edge-wise shared partners (nodes connected to both sender and receiver of a tie), and minimum geodesic distance, or path length (minimum number of connected ties, by which two nodes are linked). The black solid lines are the observed statistics in the actual network and the box plots reflect networks simulated from the model. Fit is good if they overlap as much as possible.

Figure 6 plots the model's observed statistics as quantiles of the simulated sample. In a good fit, the observed statistics are near the sample median (0.5).

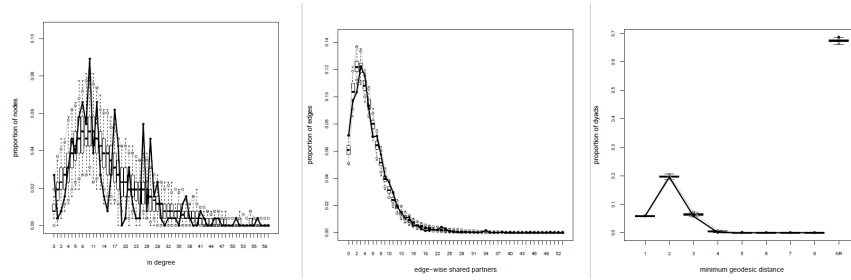


Fig. 5 The diagnostics plots of Hunter et al. (2008).

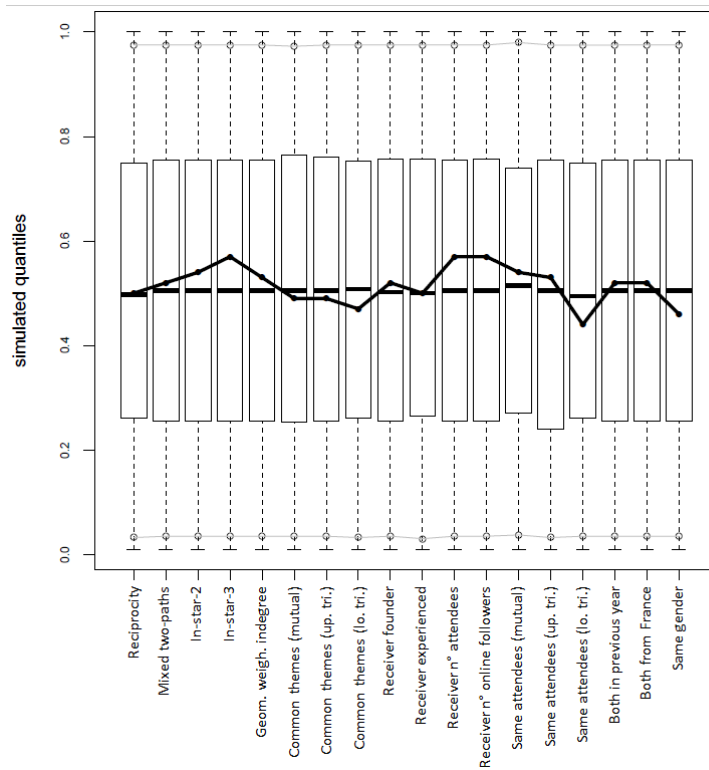


Fig. 6 Observed statistics as quantiles of the simulated sample.

9 Acknowledgments

This paper is part of the research project “Sharing Networks”. I gratefully acknowledge logistical support from the OuisShare Fest 2016 organization team, in particular Francesca Pick (Chair), Khushboo Balwani (Communications officer), Camille

Chapuis (Project manager), and additionally OuiShare founders Arthur de Grave, Diana Filippova and Benjamin Tincq. I thank Sarah Benmarhnia for research assistance in preparation of the data, and Lise Mounier for advice on a draft of this paper. Preliminary results were presented at the DYMLNA workshop at Sciences Po, Paris (March 2017), a seminar of the RITM research group at Université Paris Sud, Sceaux (May 2017), the XXXVII INSNA Sunbelt conference in Beijing (June 2017), the 3rd EUSN conference in Mainz (September 2017), the 2018 Research Day of the Society & Organizations (S & O) research group at HEC Paris, and the 2018 International Conference on Complex Networks and their Applications, Cambridge, UK (Tubaro, 2018). I am grateful to all participants and reviewers for very helpful and constructive feedback, and I am particularly indebted to Julien Brailly, Emilio J. Castilla, Guillaume Favre, Emmanuel Lazega, Alessandro Lomi and Tom A.B. Snijders for their insightful suggestions. Any remaining errors or inconsistencies are exclusively mine.

References

- Ahuja, G., Polidoro Jr., F., and Mitchell, W. (2009). Structural homophily or social asymmetry? The formation of alliances by poorly embedded firms. *Strategic Management Journal*, 30(9):941–958.
- Amati, V., Lomi, A., and Mira, A. (2018). Social network modeling. *Annual Review of Statistics and Its Application*, 5(1):343–369.
- Anand, N. and Jones, B. C. (2008). Tournament rituals, category dynamics, and field configuration: The case of the Booker Prize. *Journal of Management Studies*, 45(6):1036–1060.
- Anand, N. and Watson, M. (2004). Tournament rituals in the evolution of fields: the case of the Grammy Awards. *Academy of Management Journal*, 47(1):59–80.
- Anderson, B. (2006 [1983]). *Imagined Communities: Reflections on the Origin and Spread of Nationalism*. Verso.
- Appadurai, A. (1986). Introduction: Commodities and the politics of value. In Appadurai, A., editor, *The social life of things*, pages 3–63. Cambridge University Press.
- Arcodia, C. and Whitford, M. (2006). Festival attendance and the development of social capital. *Journal of Convention & Event Tourism*, 8(2):1–18.
- Aspers, P. and Darr, A. (2011). Trade shows and the creation of market and industry. *The Sociological Review*, 59(4):758–778.
- Blau, P. (1964). *Exchange and Power in Social Life*. Wiley.
- Block, P. (2015). Reciprocity, transitivity, and the mysterious three-cycle. *Social Networks*, 40:163–173.
- Borgatti, S. (2012). Two-mode concepts in social network analysis. In Myers, R., editor, *Computational Complexity: Theory, Techniques, and Applications*, pages 2912–2924. Springer.

- Botsman, R. and Rogers, R. (2010). *What's Mine Is Yours: The Rise of Collaborative Consumption*. HarperBusiness.
- Brailly, J. (2016). Dynamics of networks in trade fairs—a multilevel relational approach to the cooperation among competitors. *Journal of Economic Geography*, 16(6):1279–1301.
- Brailly, J., Favre, G., Chatellet, J., and Lazega, E. (2016). Embeddedness as a multilevel problem: A case study in economic sociology. *Social Networks*, 44:319–333.
- Burt, R. (1992). *Structural Holes : The Social Structure of Competition*. Harvard University Press.
- Burt, R. S. (1987). Social contagion and innovation: Cohesion versus structural equivalence. *American Journal of Sociology*, 92(6):1287–1335.
- Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology*, 110(2):349–399.
- Centola, D., González-Avella, J. C., Eguíluz, V. M., and Miguel, M. S. (2007). Homophily, cultural drift, and the co-evolution of cultural groups. *Journal of Conflict Resolution*, 51(6):905–929.
- Delacour, H. and Leca, B. (2011). A salon's life. Field configuring event, power and resistance in a creative field. In Moeran, B. and Pedersen, J., editors, *Negotiating Values in the Creative Industries: Fairs, Festivals and Competitive Events*, pages 36–58. Cambridge University Press.
- Everett, M. and Borgatti, S. (2017). Dealing with overlapping categorical attribute data. Paper presented at the XXXVII Sunbelt Conference of the International Network for Social Network Analysis (INSNA) 30 May–4 June 2017 in Beijing, China.
- Favre, G. and Brailly, J. (2015). Salons et définition de normes marchandes : Le cas de la distribution de programmes de télévision en Afrique sub-saharienne. *L'Année sociologique*, 65(2):425–456.
- Fitzmaurice, C. J., Ladegaard, I., Attwood-Charles, W., Cansoy, M., Carfagna, L. B., Schor, J. B., and Wengronowitz, R. (2018). Domesticating the market: moral exchange and the sharing economy (online first). *Socio-Economic Review*.
- Forrest, M. (2017). Establishing the complexity of obesity: The conference as a site of understanding obesity as a medical condition. In Høyer Leivestad, H. and Nyqvist, A., editors, *Ethnographies of Conferences and Trade Fairs: Shaping Industries, Creating Professionals*, pages 87–105. Springer.
- Goldhaber, M. H. (1997). The attention economy and the Net. *First Monday*, 2(4).
- Gould, R. V. (2002). The origins of status hierarchies: A formal theory and empirical test. *American Journal of Sociology*, 107(5):1143–1178.
- Handcock, M. (2003). Assessing degeneracy in statistical models of social networks. Mimeo.
- Handcock, M. S., Hunter, D. R., Butts, C. T., Goodreau, S. M., and Morris, M. (2003). *statnet: Software tools for the Statistical Modeling of Network Data*. Seattle, WA.

- Hardy, C. and Maguire, S. (2010). Discourse, field-configuring events, and change in organizations and institutional fields: Narratives of DDT and the Stockholm convention. *The Academy of Management Journal*, 53(6):1365–1392.
- Hunter, D., Handcock, M., Butts, C., Goodreau, S., and Morris, M. (2008). ergm: A package to fit, simulate and diagnose exponential-family models for networks. *Journal of statistical software*, 24(3):1–29.
- Lampel, J. and Meyer, A. D. (2008). Guest editors' introduction. *Journal of Management Studies*, 45(6):1025–1035.
- Lazega, E. and Snijders, T. (2016). *Multilevel Network Analysis for the Social Sciences: Theory, Methods and Applications*. Springer.
- Leivestad, H. and Nyqvist, A. (2017). *Ethnographies of Conferences and Trade Fairs: Shaping Industries, Creating Professionals*. Palgrave Macmillan.
- Li, Y. and Carriere, K. (2013). Assessing goodness of fit of exponential random graph models. *International Journal of Statistics and Probability*, 2(4):64–74.
- Lorrain, F. and White, H. C. (1971). Structural equivalence of individuals in social networks. *The Journal of Mathematical Sociology*, 1(1):49–80.
- Lusher, D., Koskinen, J., and Robins, G. (2013). *Exponential Random Graph Models for Social Networks: Theory, Methods, and Applications*. Cambridge University Press.
- Manzo, G. and Baldassarri, D. (2015). Heuristics, interactions, and status hierarchies: An agent-based model of deference exchange. *Sociological Methods & Research*, 44(2):329–387.
- McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444.
- Mizruchi, M. S. (1993). Cohesion, equivalence, and similarity of behavior: a theoretical and empirical assessment. *Social Networks*, 15(3):275–307.
- Moeran, B. (2011). The book fair as a tournament of values. In Moeran, B. and Pedersen, J., editors, *Negotiating Values in the Creative Industries: Fairs, Festivals and Competitive Events*, pages 119–144. Cambridge University Press.
- Moeran, B. and Pedersen, J. (2011). Introduction. In Moeran, B. and Pedersen, J., editors, *Negotiating Values in the Creative Industries: Fairs, Festivals and Competitive Events*, pages 1–35. Cambridge University Press.
- Nowak, M. and Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437:1291–1298.
- Pallotti, F., Lomi, A., and Mascia, D. (2013). From network ties to network structures: Exponential Random Graph Models of interorganizational relations. *Quality and Quantity*, 47(3):1665–1685.
- Panitz, R. and Glückler, J. (2017). Rewiring global networks at local events: congresses in the stock photo trade. *Global Networks*, 17(1):147–168.
- Phillips, D. J. and Zuckerman, E. W. (2001). Middle-status conformity: Theoretical restatement and empirical demonstration in two markets. *American Journal of Sociology*, 107(2):379–429.
- Podolny, J. and Phillips, D. (1996). The dynamics of organizational status. *Industrial and Corporate Change*, 5(2):453–471.

- Raafat, R. M., Chater, N., and Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, 13(10):420–428.
- Ravenelle, A. (2019). *Hustle and Gig: Struggling and Surviving in the Sharing Economy*. University of California Press.
- Robins, G., Pattison, P., and Wang, P. (2009). Closure, connectivity and degree distributions: Exponential random graph (p^*) models for directed social networks. *Social Networks*, 31(2):105–117.
- Rosenblat, A. and Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber’s drivers. *International Journal of Communication*, 10(0).
- Sauder, M., Lynn, F., and Podolny, J. M. (2012). Status: Insights from organizational sociology. *Annual Review of Sociology*, 38(1):267–283.
- Schor, J. and Cansoy, M. (forthcoming 2020). The sharing economy. In Wherry, F. and Woodward, I., editors, *The Oxford Handbook of Consumption*. Oxford University Press.
- Schor, J. B. and Attwood-Charles, W. (2017). The “sharing” economy: labor, inequality, and social connection on for-profit platforms. *Sociology Compass*, 11(8):e12493.
- Schüssler, E., RÜling, C.-C., and Wittneben, B. B. F. (2014). On melting summits: The limitations of field-configuring events as catalysts of change in transnational climate policy. *Academy of Management Journal*, 57(1):140–171.
- Selznick, P. (1957). *Leadership in Administration: A Sociological Interpretation*. University of California Press.
- Slee, T. (2015). *What’s Yours is Mine: against the sharing economy*. Scribe Publications.
- Smits, R. (2018). Tournaments of values at the European Film Market. *Arts and the Market*, 8(1):64–79.
- Snijders, T. A. B., Pattison, P. E., Robins, G. L., and Handcock, M. S. (2006). New specifications for exponential random graph models. *Sociological Methodology*, 36(1):99–153.
- Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *The Academy of Management Review*, 20(3):571–610.
- Torlò, V. J. and Lomi, A. (2017). The network dynamics of status: Assimilation and selection. *Social Forces*, 96(1):389–422.
- Tubaro, P. (2018). Emergent relational structures at a “sharing economy” festival. In Aiello, L., Cherifi, C., Cherifi, H., Lambiotte, R., Lió, P., and Rocha, L., editors, *Proceedings of the 7th International Conference on Complex Networks and Their Applications*, volume 2, pages 559–571.
- Urry, J. (2003). Social networks, travel and talk. *The British Journal of Sociology*, 54(2):155–175.

Whose results are these anyway? Reciprocity and the ethics of “giving back” after social network research

Paola Tubaro

Abstract

Bridging the social networks, field methods and ethics literatures, I make the case that the process of reporting research findings is an ethical issue, and recommends elevating it in the research design. I draw on a reflective account of three research experiences with settings in, respectively, online health communities, economic organizations, and the mainstream media. I proceed in steps, discussing release of personal network results to individual participants, of whole network results to the researched community, and finally of general results to wider audiences, under a unifying idea that a reciprocity obligation underlies the reporting process. I claim that communication should follow an iterative rather than a linear approach to reach all relevant stakeholders, thereby mitigating the vulnerabilities that arise from research.

Key words: Research ethics, reporting results to study participants, reciprocity, communication outside academia, social network research, network visualization.

1 Introduction: communicating research results as an ethical problem

What ethical issues may arise in the process of communicating the results of a social network study? I faced this question when I was researching the online and face-to-face social networks of persons living with anorexia nervosa and other eating disorders. The study looked at the linkages between their health condition and their social environment, showing that relationships offer support complementary

Laboratoire de Recherche en Informatique, Centre National de la Recherche Scientifique, Paris Saclay, e-mail: paola.tubaro@lri.fr

to treatment, and that web-based groups could be particularly helpful (Tubaro and Mounier, 2014). One day I received an invitation to speak at a popular TV show, whose success relied on a sensationalist approach – juxtaposing people at opposite ends of the body weight spectrum and chastising their unhealthy eating habits. The producers had heard about “pro-ana” websites and expected me to warn viewers about the potential risks of online networking and social influence allegedly triggering disordered eating behaviors. Accepting would certainly earn wide visibility for my study, but would also indirectly endorse the show’s problematic concept – while I knew from my research that judgment and stigma could be far more harmful for sufferers than any contact with (and even influence from) peers. After some hesitation, I eventually declined the invitation, explaining my reasons in a long letter to the show’s producers¹.

This is just an instance of a more general, but largely overlooked, research ethics problem: where and in what contexts should we, as researchers, draw on our research results to provide feedback to members of the social group we investigated, and more generally to the public? Standard ethical guidelines provide relatively scarce guidance on such issues. Surely, there are general principles that recommend dissemination (Iphofen (2011, chap. 12), and several research institutions, publishers, and learned societies have created organizations and adopted codes of conduct (such as COPE, Committee on Publication Ethics) that cover matters such as authorship, originality, acknowledgement of sources, declaration of conflicting interests. Nevertheless, most of these initiatives concern scientific publication, while communication with stakeholders outside the world of research is seldom discussed. Ethical approval procedures usually do require researchers to take into account these stakeholders, but they are undertaken before a study starts and focus on the early stages: key concerns are anonymization, informed consent, and any distress that might occur during data collection, but it is more difficult to anticipate reporting and communication issues that might arise long afterward. Neither is it practical to go back to ethics review boards later for one-off advice: their procedures are too lengthy for media requests that typically come at very short notice, or for unanticipated stakeholders’ reactions that require rapid responses. Further, board members may have limited knowledge of the specific setting and circumstances. In the end, the researcher is to make a decision alone, often in no time.

Reporting social network research outputs is especially tricky insofar as, in recent times, ubiquitous online social networking services and the blurring of boundaries between “public” and “private” spheres (Tubaro et al., 2014b) make people’s relationships widely visible. In the wake of internet-mediated communication at both local and global levels, people have progressively learned to recognize their connectiveness and to attach values and status to it (Christakis and Fowler, 2011; Rainie and Wellman, 2012). The potential benefits of an extended range of social contacts are now widely understood, and are among the motivations for using social platforms and applications, so much so that people sometimes overestimate the extent and

¹ I made public an anonymized version of the letter: <https://web.archive.org/web/20131217023109/http://www.anamia.fr/reponse-a-un-journaliste-tele-ayant-sollicite-une-interview-sur-les-troubles-alimentaires-et-les-sites-pro-ana/>

reach of their online social capital (Dunbar, 2016). Thus, any research results that make people aware of the relative richness (or poverty) of their social networks are likely to have an effect on their personal life satisfaction and well-being (D'Angelo and Ryan, 2019).

Another reason why in present times, reporting research outcomes is a particularly sensitive part of social network research is that stakeholders who hold setting-specific knowledge may well re-identify study participants and misuse results. In their influential contribution to the *Social Networks* special issue on ethical dilemmas, Borgatti and Molina (2005) made precisely this point, in the case of organizational network research that managers may use to reward or penalize employees. Today, increased visibility of relationships through digital technologies extends this challenge beyond workplaces. Using the above example, websites on eating disorders might expose to unwanted viewers (including potential future employers) not only their authors, but also their readers.

A related challenge stems from the common confusion between social networks as our research community construes them – as structures of relationships, interactions or interdependencies between individuals and groups – and social networks popularly interpreted as digital tools for social networking (or social media). This is a common source of misunderstandings, especially for those of us who study non-technology-mediated networks, and calls for very careful word choice at every step. Kadushin (2005, p. 139)'s point that “the social network field may have become a victim of its own successes” is still, and perhaps increasingly, true: we benefit from improved visibility but struggle more to get our message through.

This paper aims to advance social network scholars' endeavor to comprehend the conditions for responsible return of research results. It draws on a reflective account of three cases from my social network research activity. All three concern communication to laypersons: thus, scientific presentations and peer-reviewed publications fall outside the scope of my analysis. I share my field experiences as a step towards a collective effort to co-construct an appropriate ethical framework, in the belief that researchers can (and should) take an active role.

To achieve this, I set out to bridge different literatures, in addition to research ethics strictly interpreted. Particularly in field methods, there is a large and mature tradition of reflecting on, and developing best practices for, returning results to the community where the researcher originally collected the data, and on which those results might exert an effect. At the crossroads of these different lines of thought, I frame return of results as a reciprocity obligation – a way of “giving back” to society in exchange of the resources (be they informational, financial or other) that it made available to research. Accordingly, return of results addresses not only informants in the field, but also a wider range of stakeholders, including funders and promoters of research projects, partners, policymakers and the general public. Instead of a linear progression in which reporting is just the final stage of research after data collection and analysis, the process I envision is a feedback loop between stakeholders and researchers that forms an integral part of research design since inception.

I argue that social network research is an ideal setting to develop these ideas. While other domains, especially biomedical research, sometimes notoriously shaped

the key principles of ethics for all disciplines, the social sciences may benefit from fresh perspectives that take into account both their substantive and methodological specificities. Social network research is particularly well suited to serve as a terrain of experimentation as, unlike clinical trials, its focus is not on individuals taken independently of one another but on interpersonal relationships that link individuals and groups. It thus introduces fundamental dependencies across people's behaviors and is capable of providing unique insights into society seen as a cohesive whole, rather than as a sum of separate entities.

In addition, social network research helps answer a fundamental question that the literature often leaves open: what is the "community" to which researchers should return results. If it is not limited to the narrow circle of study participants, where should one draw the boundaries? Network studies suggest that a community may be understood as the set of social relationships that are meaningful to a focal individual or group (Wellman et al., 1988), and prove that mapping such relationships very precisely and assessing how actors make sense of them, can be very helpful. For example in my above-mentioned study of the networks of persons with eating disorders, doctors and other health professionals appeared as alters together with friends and family members, and were therefore targets of my communication as part of participants' "personal communities" (Chua et al., 2011).

2 Bridging literatures

While frequently ignored in established research ethics guidelines, presentation of results has begun to receive a greater deal of attention in recent times, notably in the literature on clinical trials. Returning results to patient participants becomes a greater priority as part of the ongoing shift in healthcare toward participant engagement and "patient-centered care" (World Health Organisation, 2007). At the very least, this means communicating aggregate research results to general audiences, in summary form and using language that is understandable and usable by non-specialists. Some regulators have already adopted policies and processes to this end: both the US Food and Drug Administration, and the EU Clinical Trials regulation, require publication of summaries of results for laypersons.

There is also an emergent, though less consensual need to give clinical study participants access to their individual research results, that is, the outcomes of any tests (such as blood tests, scans etc.) performed on them individually. The hope to get such results is a primary motivation for many patients to participate in a clinical study at all, especially in a context in which digital self-tracking devices and the services of commercial companies such as 23andMe have accustomed people to easily accessing, handling and visualizing their own health data (Wong et al., 2018). Ethically, return of individual results is a way to ensure participants receive the greatest possible benefit from their contribution to the study. Research bodies such as MRCT in the United States (2017) are issuing guidelines to support release of individual outcomes, although there are disagreements over the optimal timing,

the extent to which results can place a burden on patients, and the need to withhold findings in some circumstances. One contested issue is report to the family, including after the death of the research participant, notably in the case of genomic research (Wolf et al., 2015).

Calls for open science and open access reinforce the expectation of accessibility of research results to even wider communities – not limited to research participants. Recent debates on open access rely upon the moral framing that requires the outputs of publicly funded research to be freely available to those who paid for it – the taxpayers. This is an argument about obligations of reciprocity (Parker, 2013), whereby researchers receive (financial) resources from the community, and give knowledge to it in exchange. In this perspective, dissemination of knowledge is a key responsibility of researchers and grounds their social legitimacy. The public is implicitly supposed to participate in the process, although this current of thought focuses mostly on scientific publications as such (with issues such as replicability), rather than on outreach to lay audiences (Cahill and Torre, 2007).

These approaches have the merit of replacing presentation of results at the center-stage of research ethics, shifting from a focus on data collection and data management to a more comprehensive view that covers all phases of a scientific study. Yet report of research results is not an unprecedented outcome of patient-centered care, as the idea that researchers have a duty to engage in conversations with wider audiences predates the current movement toward open science. Traditional approaches to the social sciences, not always framed as ethics, have long taken strong normative stances in this matter. Pierre Bourdieu (2015) believes researchers have a duty to expose the power relationships that shape society, while Michael Burawoy (2005a,b)'s "public sociology" advocates dialogue between researchers and diverse audiences beyond academia. However, these arguments are often cast in too general terms to help answer specific, practical questions: what is the relevant community? What results can be made available? What is the appropriate format?

More to the point, ethnography and field methods have long reflected on "giving results back" to the community where researchers originally collected the data. Field study design encompasses not only entry into, but also exit from, the target community, so that sharing of results is naturally constitutive of it. While firmly grounded in methodology, this literature raises issues that have a strong bearing on ethics. For example, what is the place (if any) of the researcher's autobiographical experience and emotions in the field (Ellis, 1995)? Do participants have a say on narration of results concerning them? What information can intermediaries, community leaders, and any other non-participant stakeholders expect to receive? How should the researcher address any criticism or disagreement? Especially within participatory and advocacy research (MacKenzie et al., 2015), it is increasingly recognized that researchers have an ethical obligation to share results with the community under study, in order for it to reap a maximum of benefits from participation. Indeed the potential effects of research on stakeholders may vary from shedding light on a problem they are facing, to giving them more voice in the policy process, or to facilitating work towards a solution (Reyes-García and Sunderlin, 2011).

Schurmans et al. (2014) broaden the concept of reporting results, to address not only informants in the field but all who may be concerned, including funders and promoters of research, partners, and fellow researchers. One may also add policy-makers, for their role in steering both research and social action, as well as associations and other civil society bodies, according to the teachings of public sociology (Burawoy, 2005b; Jeffries, 2011). This approach requires re-framing report of findings not as the last stage of a linear process that starts with data collection and terminates with scientific publication, but as a cyclical one with multiple feedback loops. Continuing dialogue with a range of relevant audiences enables researchers to validate their work, builds trust with participants and other stakeholders, and may lead to maintaining relationships over time, possibly through multiple projects (Chavis et al., 1983; MacKenzie et al., 2015). Some social science traditions like action research already operate in such a cyclical mode. In this long-run perspective, social network research may benefit from its capacity to identify the relevant community of stakeholders, as discussed above, to establish what results to return to each of its different components at what time.

These approaches draw on different perspectives and types of disciplinary expertise, but resonate with, and complement, each other. Both the field methods and open science currents build on a principle of reciprocity. Bringing them together enlarges the range of stakeholders to consider: ideally, all those relationships that shape the social environment in which the research is undertaken, and on which it may have an impact. At the crossroads of ethical, scientific and political issues, such a comprehensive view accounts for the multiple commitments that researchers have to diverse actors and the difficulties that arise in the effort of reconciling diverging agendas (Tubaro, 2017).

3 Social networks research and case-based reasoning

So far, the social networks literature has not addressed these issues explicitly. The 2005 Social Networks special issue on ethical dilemmas touches on them only indirectly and partially, in terms for example of potential misuses of results and of commitments to problematic sponsors such as the military. Yet social network research exemplifies at best some of the ethical issues that the above discussed literatures highlight. For example, there are specific concerns regarding report of individual results to participants, insofar as revealing the set of ties surrounding a person may jeopardize confidentiality and more generally trigger consequences. Social network research also raises specific problems as a consequence of the power and growing popularity of visualizations, which constitute an effective tool to communicate results but may be deceptive or distressing under some conditions. Further, the common confusion between social networks and online social media adds to the difficulty of communicating results clearly and unambiguously.

Conversely, social network research provides excellent material to work toward, and test the efficacy of, possible solutions that draw on the integrated approach pro-

posed above. Return of individual (ego) networks is already practiced and, with appropriate ethical safeguards, may pioneer extensions of this practice to the social sciences at large. Similarly, work on visualizations may improve researchers' capacity to communicate to audiences beyond academia – and again, social networks may be a test for solutions of general applicability. Also as mentioned above, network research can help us identify more precisely the community of stakeholders – who have an interest in the research or on whom the research may have an effect – whom researchers should interact with.

In what follows, I explore these ideas with the help of three example cases from my own research experience: one about use of visual sociograms to return individual results, another about whole network representations of the researched community/organization, the third expanding on the TV show dilemma introduced earlier. A case study approach is most suitable to a situation characterized by limited preliminary evidence and a research context that is not yet clearly specified, in an emerging domain of investigation (Yin, 2009). I do not limit my inquiry to the exploratory level, though, and move forward to using the cases at hand to extract more general arguments, the conclusions of which can be actionable. Thus I make an instrumental use of my cases, seeing them as concrete instances of the broader issue of interest, into which they can provide insight (Baxter and Jack, 2008; Stake, 1995). Cases require in-depth study of particular empirical contexts in which they are embedded, in order for the researcher to distinguish the general from the specific, to propose interpretations, and to devise new ideas (Ragin, 1992).

As Crossley and Edwards (2016) maintain, social network research is essentially a case-study method, where focus is exclusively upon a given, well-demarcated social group, especially in whole-network designs. Contrary to classical statistical inference, there is no assumption that the case or cases under consideration are representative, in a statistical sense, of any wider population. This is one reason why social network research has caught on in studies of specific, historically or geographically limited worlds such as organizations and social movements. In principle, the case study requires engagement with all possible sources of data, methods and analytic strategies that may shed light on the phenomenon of interest. In practice, it is often associated with qualitative research, except in social network studies that draw on a notoriously strong quantitative tradition. Recent approaches lead to greater convergence, whereby qualitative case study specialists use more quantitative methods (Mills et al., 2010) and network researchers increasingly draw upon multiple data sources, triangulate observations and build richer pictures of their cases (Crossley, 2010; Dominguez and Hollstein, 2014; Ryan and D'Angelo, 2018).

4 Giving personal network results to individuals

Social network research makes visible patterns of relationships that would otherwise escape human gaze. Therefore, giving individual-level results back to participants provides them with unique, otherwise inaccessible information about themselves.

Should we, then, systematically provide study participants with maps of their own networks – representing their “personal communities” in the case of personal network research, or extractions of their ego-centered subgraphs in the case of whole network research?

This practice is already spreading with the help of network visualizations, which are being integrated into different stages of research, from data gathering to presentation of results (D’Angelo et al., 2016). Researchers who have shown participants visualizations of their ego networks note that they shifted their role from being observed to becoming observers themselves (Molina et al., 2014) and felt the need to explain and justify what they saw (Bellotti, 2016). But if people understand visuals, they may take some type of action as a consequence. This is especially likely in a society that values connectedness (Christakis and Fowler, 2011), and in which individuals are aware of the benefits that might accrue from their relationships (Rainie and Wellman, 2012), so that they will have a sense of themselves as more or less successful networkers. Wide availability of social media visualizations until about 2013 has contributed to increasing people’s appetite for this type of information, although most digital platforms have recently restricted access to these data (Hogan, 2018). The potential existence of such effects involves a duty for the researcher to anticipate them and as far as possible, to prevent any drawbacks.

A case from my research illustrates some of the ensuing dilemmas. As part of the above-mentioned study of the social networks of eating disorder sufferers, I used a graphical concentric-circles sociogram structure embedded in a web survey to elicit personal networks (Tubaro et al., 2014a; Tubaro and Mounier, 2014; Tubaro et al., 2016). Participants had to fill in two such sociograms, for their face-to-face and for their computer-mediated ties. The top panels of figure 1 are two examples of these sociograms, representing a face-to-face personal network (left), and an online personal network (right). From these images, drawn by participants themselves, were extracted edgelist and tables with alters’ attributes, for use in analysis. My team and I had initially planned to offer respondents an optimized visualization of their full (face-to-face and online) personal networks at the end of the survey, to reward them and acknowledge their participation. The bottom panel of figure 1 is the layout of one early option we considered (and eventually dismissed), representing online and face-to-face networks as concentric circles as before, but placing them together in the same image as perpendicular plans intersecting in a three-dimensional space (horizontal: face-to-face, vertical: online).

But there were risks. If participants liked the visualization, they would export it and perhaps post it onto their personal pages and social media profiles. If the visualization included the names of their contacts (alters), anonymity would be jeopardized; but without names (as is the case in this figure), the visuals might appear confusing and little meaningful, especially with sizable networks and/or major overlaps between online and face-to-face contacts. Another risk would arise if our tool confronted participants to their isolation and loneliness – utterly visible in a graphical representation. At the time (2010), there was no discussion of such a risk in the social networks literature, and the rare previous experiences of sociogram use positively acknowledged participants’ comments on how “interesting their personal

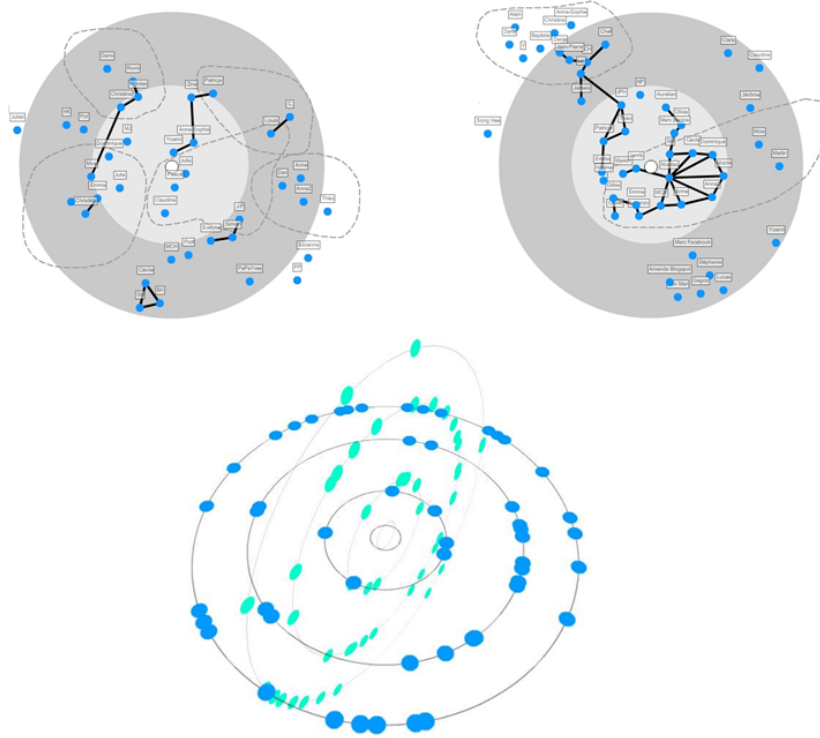


Fig. 1 Graphical, concentric-circle sociograms for data collection (top) and planned return to participants (bottom) in a web survey. Top left panel: a face-to-face personal network, right panel: an online personal network, as could be drawn by a participant. The bottom panel is an unfinished (and anonymized) project to aggregate these same data into one image that participants could export at the end of the survey. Horizontal plan, blue: face-to-face network; vertical plan, green: online network. All the graphs reproduced in this figure are from technical tests and do not represent real persons.

networks look” (Hogan et al., 2007, p. 137). Yet ours was a self-administered web survey where no interviewer’s mediation was available to offer reassurance to participants if needed. We were particularly concerned because we were surveying persons living with eating disorders, whose experiences of illness and hospitalization are known to limit their inter-personal interactions and to disrupt their social functioning (Levine, 2012; Patel et al., 2016). We eventually dropped the plan to give individual results back to participants and never developed the tool depicted in the bottom panel of figure 1.

Later research provides hard evidence of the issues we intuitively anticipated – even outside the health field and in interviewer-led studies. Ryan et al. (2016) interviewed highly-skilled French professionals in London, many of whom appeared quite self-conscious about the visual representations of their personal networks.

Those who thought of themselves as successful, confident networkers experienced a sense of unease when presented with, for example, lack of diversity in their relationships – as a woman who had described her network as “cosmopolitan”, but realized that her visual included almost only co-nationals. Their attitude betrayed a form of vulnerability – that of migrants whose relatively high economic status did not ensure full integration.

In sum, results may reveal (possibly untold) vulnerabilities and trigger reactions that are not observed in more classical participant-researcher interactions. For all their power and appeal, visualizations can be double-edged swords. This is not to say that personal network results should never be shared with study participants. Rather, the lesson learned is that researchers should try and anticipate the conditions that might in some cases lead to distress, and adapt their use of visualizations accordingly.

5 Giving whole network results to communities

Let us now turn to a second case study, in the context of a research about networking patterns at an international festival of the so-called “sharing economy”, the annual OuiShare Fest, which took place in Paris from 2013 to 2017. At the 2016 edition, I distributed a questionnaire to participants, with name generators that asked them about their relationships pre-dating the event, the ties formed during its course, and those expected to give rise to further contacts or collaborations in the near future. Almost one year after that, and about one month before the following edition of the festival, I approached the organizers with some preliminary results. I did so to give them confidence that the study was producing relevant outcomes, to renegotiate access to the event and re-run the survey again, and to offer them insights – which I expected to matter to them as professional networking was a watchword of the OuiShare experience. Another motivation was the interest of study participants, which I clearly perceived during fieldwork. Asked in 2016 if she would like to see my (aggregate) results, one of them had replied enthusiastically:

I think it will be really interesting. The fact that somebody is interested in studying what's happening, provides the feeling that what's happening is really happening, and that it's big enough to be studied [...] So I think if you do that, it will be really huge (resp. 26, 2016).

Indeed OuiShare members were thrilled with the Gephi visuals that I produced from my data. One graph showed that relationships between participants that pre-existed the 2016 OuiShare Fest formed clusters largely disconnected from one another, with many isolates; another revealed that through the new ties formed during the event, most people had become part of a main, densely-knitted connected component.

OuiShare leaders immediately produced an animated gif from my graphs and shared it widely on Twitter. Two of them went as far as printing my visualizations on T-shirts that they wore at the following (2017) edition of the event, labeling them

as “before” and “after”. Interestingly, they advertised themselves to attendees as “human tinders” – the match-makers who, like the well-known app, connect people “in real life” (figure 2).

This particular episode happened not to harm anyone, and it was even rather funny; but it alerted me about potential issues. One brings me back to visualization and its power. It was indeed just those two graphs that caught the attention of OuiShare, not the other visuals I provided, nor the explanations I gave on the side. Those two graphs alone were so successful, partly because I deliberately chose a layout that made very apparent the outcome that OuiShare members were so keen to see – the effectiveness of their event in connecting people. The effect was not spurious, but other graphical choices, such as colors or layouts, could have made it much less visually striking. Although there was no deception here, graph visualizations may lead to selective appropriation by stakeholders, and researchers need to be very careful about their use to communicate results.

Another reason for reflecting on this case is that the key users of my results were not study participants individually, but the organizers of the event where the research took place: among the latter, few actually responded to my questionnaire, while many of them appeared in the network as alters (that is, they were nominated by other respondents). This confirms the tenet of the field methods literature that reporting results to the researched community does not simply mean getting back to participants narrowly defined, and that a more holistic perspective is needed. Network analysis can help refine and extend this approach through its capacity to map who matters to whom, so as to account for the full range of relevant stakeholders, whether they are study participants themselves or not. The fact that my graphs first circulated through Twitter, a social networking service, ensured targeted diffusion to stakeholders within the OuiShare community much more effectively than any publication in a generalist outlet.

When results are shared with a whole social setting rather than with single participants, additional issues may arise depending on the power structure and influence channels within that setting. OuiShare includes prominent community leaders, and their flagship event has played an important role in shaping the “sharing economy”, an emerging and still loosely defined field of activity. They asked me to what extent the results I reported were attributable to their members: were they instrumental in linking participants at the event? From the point of view of their internal management, this would be useful information: “connector” is an official role in their organization and they expect members to be effective networkers. But from a researcher’s standpoint, this was a difficult decision to make. Even if OuiShare prides itself to be a flat, collegial organization of (mainly) freelancers and independent professionals interacting via high-tech digital devices, fear of introducing or exacerbating power relationships prevented me from deanonymizing network nodes to identify the best connectors. Because invisible, informal dominance structures might still be in place behind any formal shape (Krackhardt and Hanson, 1993), with potential consequences for individual members, the same issues might arise as those highlighted by Borgatti and Molina (2003, 2005) in the context of more traditional, hierarchical organizations. As Howard Becker (1978, p. 330) first noted, full

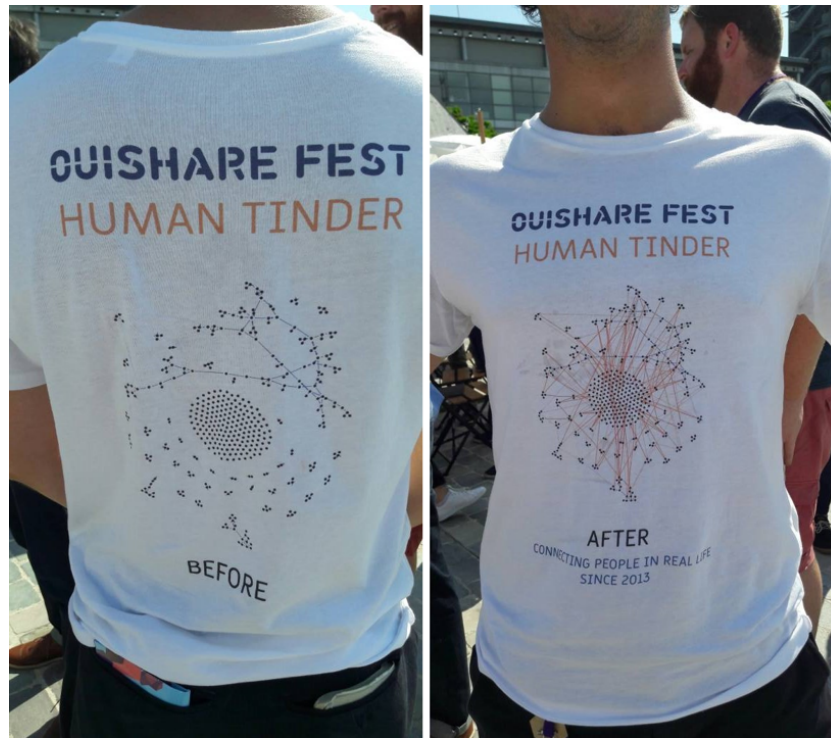


Fig. 2 T-shirts printed by OuisShare members in 2017, with my graphs of networks of ties between participants at the previous edition of the event. The “before” graph refers to ties that pre-existed the event; the “after” one to ties formed during it. I drew the graphs in Gephi using questionnaire data I collected in 2016. The labels, the “Human Tinder” title, and the “Connecting people in real life since 2013” caption were added by OuisShare.

consensus does not exist in any community or organization, and the social scientist may end up pleasing one faction while disserving another, so that “the problem is not to avoid harming people but deciding which people to harm”. Even anonymized results would not be neutral, producing potential tensions if, for example, they revealed that it was outsiders rather than insiders who mostly contributed to linking participants.

In a sense, there was also a risk for the researcher. I gave my findings to this community at an early stage to build trust, receive feedback and negotiate a return to the field. But the unexpectedly wide publicity could have backfired, in that any unhappy members could have obstructed my second wave of data collection. In addition, because my authorship was not acknowledged on the T-shirts (although it was in the animated gif that circulated via Twitter), I feared my re-use of such material in future publications would need to be very carefully thought out. In essence, researchers are to be included among those potentially affected, that an ethical approach should take into account – as will be discussed more extensively later. Extending reporting

efforts beyond individual participants is thus a challenge. Visuals can serve as effective and easy-to-understand tools, but they may be distressing or deceptive. When are difficulties most likely to arise? Extant literature (Borgatti and Molina, 2003, 2005) already emphasized the risks that arise in organized settings, and my case appears as an extension of these ideas to looser organizational forms. Even mild power or influence relationships produce specific opportunities and threats that researchers should anticipate as far as possible. My case study also resonates with recent evidence that people attach values and status to connectedness (Christakis and Fowler, 2011; Rainie and Wellman, 2012) – and are therefore keen to demonstrate their role as actors of this connectedness, possibly using research results to this purpose.

The key is to ensure reciprocity criteria are met, at least to some extent: the fact that they were in the OuiShare Fest case (the organizers gave me access to the field, I gave results to them), kept potential risks under control and gave me a second access to the field, and even further logistical support, one year later (after which I also returned results to them as I had done before). But in other cases, such conditions are not met and return of results is more problematic, as discussed in the next section.

6 Beyond networked communities: sharing results with the media and wider audiences

What about communicating general results to publics that are larger than the specific social setting in which a study was run (such as the OuiShare community)? Research institutions increasingly encourage outreach to wider non-academic audiences who are indirect financial contributors (notably through public funding schemes) and/or potential beneficiaries (for example, families, neighborhoods or communities of interest). Ethical issues arise as sensitive decisions must be made, such as choice of sufficiently accurate wording when technical language cannot be used. So far, only a few institutions and learned societies have published ethical guidelines for responsible press communications (for example the Society for Neuroscience²).

As before, I draw on my own research experience to stimulate a discussion. I already mentioned a case in which I declined to speak at a TV show to avoid risks of distortion of my findings on online networks and eating disorders. This is a far-too-common trade-off between two ethical imperatives, one that urges researchers to reach out to the public, and one that invites to caution in light of the potential negative consequences of any misrepresentation of results, notably when they have health implications. Other challenges related to this case speak expressly to the social network research community. The producers of the show implicitly endorsed the idea that online relationships are somehow less valuable than face-to-face ones. As network scholars, we know this is a more complex matter. A long-lasting de-

² See <https://www.sfn.org/Membership/Professional-Conduct/Guidelines-for-Responsible-Conduct-Regarding-Scientific-Communication>, accessed on 27 July 2019.

bate has opposed those who stress the “internet paradox” (Kraut et al., 1998) of a communication technology that drives people apart and undermines human propensity towards face-to-face conversations (Turkle, 2012), and those who emphasize seamless integration of the two (Wellman and Gulia, 1999; Wellman and Haythornthwaite, 2002). Even young people are not all savvy with technologies, and down-to-earth inequalities persist online (Hargittai, 2010).

What’s more, the show producers took for granted that social influence through peers on the web triggers imitation of unhealthy eating behaviors and is therefore undesirable. Beyond its neglect of influences from non-computer-mediated relationships (Valente, 2010), this view puts the blame on personal networks, without considering the responsibilities of mainstream media that can reach out to much larger audiences. My own research showed that personal networks mediate the reception of media contents, sometimes attenuating their effects, sometimes amplifying them, with different strength and significance depending on individuals’ body size (Pallotti et al., 2018). This very insight from my social networks research suggested that, with little time to elaborate and no control on the social environments of the show’s publics, it would be wiser not to communicate widely on such sensitive matters.

Alternatively, I could have chosen to attend the show and explain myself; but even if I had been given sufficient time to make my point, my approach could be seen as lecturing others, as if I were in a position of higher standing, rather than just engaging in a mutually beneficial conversation.

This brings me to the reciprocity argument I introduced above. Broadly defined, dialogue with stakeholders is a way for researchers to give back to their social environment in exchange of what they received from it – not just the financial resources they ultimately obtain from taxpayer money, and not just responses from participants to a study, but any form of information and social support. The problem is that the conditions for reciprocity must be in place for it to happen. In the case of this TV show, I could not ensure appropriate “giving back”, as the risk of distortions and sensationalism was too high. Besides, reciprocity was made difficult by the lack of common language and shared understandings of the matter at stake.

7 Discussion: dual vulnerability and reciprocation

Building on the thought of Emmanuel Levinas, philosopher Chardel (2013), who was a partner in my eating disorders study, describes the mirroring of instability and risks in respondents (or other stakeholders) and in researchers as dual vulnerability. First, vulnerability concerns study participants: in our research on eating disorder networks, it was mainly due to their health condition, exposure to social stigma, and relatively young age. The second vulnerability is that of researchers who take human, scientific and ultimately legal responsibility when they handle sensitive information which participants (especially if in situations of illness or distress) entrust them with. Chardel’s notion of researcher vulnerability stems specifically from French data protection law, which had stringent requirements even before entry into

force of the GDPR in May 2018, all the more so when data contained sensitive information (e.g. on health).

Yet this notion can be easily transposed to regulations in place in other countries, where researchers are subject to Review Boards, with further constraints coming from the legal and risk management services of their institutions (Molina and Borgatti, 2019). Chardel's view has the advantage of taking into account various sources of researcher vulnerability, deriving from the broad range of ethical and even scientific risks that arise from interactions with multiple stakeholders. These extra sources of vulnerability are looser and less recognizable, but their effects are no less forceful.

Some aspects of my case studies can be understood in these terms. The first (return of personal networks to individuals) suggests that not only participants, but also researchers are vulnerable if presentation of results fails to protect anonymity, because their status and reputation in the profession may suffer. Perhaps in a subtler way, both parties are vulnerable if researchers are unable to reassure participants when findings generate discomfort. Researchers may lack the necessary skills, training, or information to provide support, especially in difficult cases that would require specialized clinical or psychological expertise. Deciding whether to try and give support is in itself a challenge, insofar as social science research is not usually meant to have therapeutic value.

The second case study (return of whole networks to communities) hints at dual vulnerability arising from potential misuse of research results, especially in organized settings. Local leaders may leverage these results for commercial purposes, and even use them to promote their own advantage, possibly to the detriment of weaker members. The researched community is vulnerable to the potentially destabilizing effects of such actions. Researchers are vulnerable too, insofar as they may have unwittingly contributed to developing tools for purposes that do not fit with their values.

The third case (communication of generic results to the public at large) is probably the one where the vulnerability of all parties is at its highest. Misrepresentation through highly-visible popular media may unduly stigmatize the researched population or community while also de-legitimizing researchers if they fail to strike the right balance between the complexity of the problem under study and the need to use simplified language. This is difficult to manage as researchers typically lack control over the final media product (content of the article or report) and its framing (insertion in a broader story, choice of title and accompanying images). Further, researchers are rarely prepared to handle exposure to public debates (although universities and learned societies increasingly offer training and support for press communications).

In all three cases, a reciprocation process may mitigate the emerging dual vulnerability. I argued earlier that a reciprocity obligation underlies the reporting process notably in the context of social network research. Now, it can be added that reciprocation emerges in the mirroring of reduced vulnerability as stakeholders develop trust in researchers (and vice versa) when communication to relevant audiences proceeds safely. The cases presented above show that when the conditions for reci-

procuity are met, researchers are fully able to “give back” to the communities that supported their work, while keeping vulnerability low for all parties.

8 Conclusions

Researchers’ presence in the field is far from neutral and naturally involves some form of feedback to those affected. In what precedes, I have argued that the act of reporting results is an important part of the ethical reflection that should permeate a research project at all stages. Some preliminary plan to share results with participants should be made at the very beginning, even though not all circumstances can be anticipated, and some adjustments will need to be made along the way.

A thorny issue is *who* to return results to. Under some conditions, research participants may expect to receive their own individual results – notably the structure and composition of their personal network, possibly though not necessarily in visual form. Using the case of a study of users of health-related websites, I highlighted some circumstances in which this is not suitable. Visuals offer an intuitive, effective means for the social network researcher to communicate, but may be misleading or unhelpful in some cases.

Whole network results may be of interest to the researched organizations, groups or communities. Using the case of an event whose attendees I surveyed, I highlighted several ethical risks such as appropriation (and possibly, misuse) of results by stakeholders in a position of power or leadership. These influential members may sometimes raise awareness of a research in a given community, encourage participation, and provide domain-expert feedback on preliminary results. It is thus essential that researchers identify such leaders at an early stage and inform and involve them in all phases of the fieldwork. In cases in which leadership is emergent rather than determined by formal structures, social network research may detect individuals in key positions and help understand which relationships to leverage.

Generalized research results (such as the fact that people’s reception of media contents is mediated by their social contacts) are also relevant for the public at large, and researchers are increasingly encouraged to communicate to the generalist press. While this serves laudable transparency and accountability purposes, the risk of misrepresentation of results is often high. As discussed above, it is wiser in some cases to renounce the visibility that media presence offers. This is by no means a call to shy away from the press: journalists can be allies who help find a larger audience for research and may be instrumental to pass the message especially when it can have policy impact. Ideally, researchers should be able to include journalists in their social networks, establishing long-term relationships that help create the conditions for “giving back” as discussed above.

These considerations prompt a reflection on *when* to return results. The Society for Neuroscience, one of the few to have released guidelines for ethical communication with the press, recommends waiting until a finding is peer-reviewed and accepted for publication before announcing it to the public at large. However

in the social sciences, this should not rule out the possibility of earlier release to researched and/or affected communities, as long as the preliminary or descriptive nature of findings is clear. I construe report of findings as a cyclical process involving dynamic, mutual interchange between researchers and relevant stakeholders. I shared my preliminary OuiShare Fest findings with the organizers before showing them to academic audiences. By so doing, I both offered a reward to OuiShare and benefited my research, earning the possibility to go back to the field for a new wave of data collection, and receiving feedback that helped me refine my interpretation. Later, as peer reviews increased my confidence in the results, I communicated more broadly through the OuiShare radio and newsletter. In short, communication with field actors both preceded and followed dissemination to fellow researchers, and each step involved larger circles. While this is just one case and different sequences may be envisioned in other settings, the common challenge is to identify relevant stakeholders and decide what to share with each of them at what time.

Overall, these considerations suggest that there is insight to gain by sharing researchers' first-hand experience and the lessons learned from it. While ethical guidelines may include some general principles of how to communicate with the press, they cannot anticipate all possible issues, and some solutions are difficult to transfer beyond disciplinary boundaries. Learning from one another and working together to create the conditions for appropriate "giving back" is a form of "indirect" reciprocity that defines a more promising way forward.

We could set the example as social network scholars, while also engaging in a collective effort to rethink the way we have practiced our research and branded ourselves since the advent of ubiquitous digital technologies. Some of the difficulties I have discussed derive from an ambiguity in the way outsiders see our work. To be sure, as a group we have benefited from the increasing popularity of social media, as it has turned out very soon that we have the right analytical tools and concepts to analyze them. Our very methods – our points-and-lines graphs, our metrics, and our vocabulary – are more widely understood today than in the pre-Facebook era. And yet, we suffer from the unprecedented imbalance of economic and political power between us – a niche scientific community – and the tech giants that own online social networking sites. It is perhaps time to engage in a more upfront discussion of the matter, to position ourselves more clearly in regard to our objects of study and our methods. It is not only a matter of identity or legitimacy toward our institutions and our funders: we also face consequences in terms of clarity of our message to the public, and capacity to illuminate broad societal issues.

Such considerations do not represent an end-point to this discussion. Our views on how to responsibly communicate our research results are bound to evolve, especially as researchers develop solutions that mitigate the forms of dual vulnerability to which they and their stakeholders are subject. It is important not to leave this to top-down regulators: we need to creatively and collectively contribute to devising an appropriate ethical framework.

9 Acknowledgements

This paper is under consideration for publication in the journal *Social Networks*, as part of a special issue on “Recent Ethical Challenges in Social Network Analysis”. This is the revised version, submitted to the journal after a request for minor modifications in July 2019. A preliminary draft was presented at the workshop RECSNA2017 that took place in Paris on 5-6 December 2017, organized by BSA-SNAG (British Sociological Association - Social Network Analysis Group) with support from Maison des Sciences de l’Homme Paris-Saclay and Institut d’Etudes Avancées, and in collaboration with RT 26 (*Réseaux sociaux*) of *Association française de sociologie*. I am grateful to the whole RECSNA team and attendees, and I am especially indebted to Bruce Cronin, Louise Ryan and two anonymous reviewers of *Social Networks* for their helpful feedback. Any remaining errors or inconsistencies are exclusively mine.

References

- Baxter, P. and Jack, S. (2008). Qualitative case study methodology: Study design and implementation for novice researchers. *The Qualitative Report*, 13(4):544–559.
- Becker, H. (1978). Problems in the publication of field studies (1st ed. 1964). In Bynner, J. and Stribley, K., editors, *Research Design: the Logic of Social Inquiry*, pages 325–340. Routledge.
- Bellotti, E. (2016). Qualitative methods and visualizations in the study of friendship networks. *Sociological Research Online*, 21(2):2.
- Borgatti, S. P. and Molina, J. L. (2003). Ethical and strategic issues in organizational social network analysis. *The Journal of Applied Behavioral Science*, 39(3):337–349.
- Borgatti, S. P. and Molina, J. L. (2005). Toward ethical guidelines for network research in organizations. *Social Networks*, 27(2):107–117.
- Bourdieu, P. (2015). *Questions de sociologie (1st ed. 1980)*. Les Editions de Minuit.
- Burawoy, M. (2005a). 2004 American Sociological Association Presidential address: For public sociology. *The British Journal of Sociology*, 56(2):259–294.
- Burawoy, M. (2005b). The critical turn to public sociology. *Critical Sociology*, 31(3):313–326.
- Cahill, C. and Torre, M. (2007). Beyond the journal article: representations, audience, and the presentation of Participatory Action Research. In Kindon, S., Pain, R., and Kesby, M., editors, *Participatory Action Research Approaches and Methods*, page 196–205. Routledge.
- Chardel, P. (2013). Réseaux numériques et pratiques de soi : le cas des communautés ‘ana-mia’. Réinterroger l’articulation du droit et de l’éthique. *Ethique publique*, 15(2).

- Chavis, D., Stucky, P., and Wandersman, A. (1983). Returning basic research to the community: A relationship between scientist and citizen. *The American Psychologist*, 38(4):424–434.
- Christakis, N. and Fowler, J. (2011). *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*. Back Bay Books.
- Chua, V., Madej, J., and Wellman, B. (2011). Personal communities: the world according to me. In Scott, J. and Carrington, P., editors, *The Sage Handbook of Social Network Analysis*, pages 101–115. Sage.
- Crossley, N. (2010). The social world of the network: Combining qualitative and quantitative elements in social network analysis. *Sociologica*, 1.
- Crossley, N. and Edwards, G. (2016). Cases, mechanisms and the real: The theory and methodology of mixed-method social network analysis. *Sociological Research Online*, 21(2):13.
- D'Angelo, A. and Ryan, L. (2019). The presentation of the networked self: Ethics and epistemology in social network analysis (online first). *Social Networks*.
- D'Angelo, A., Ryan, L., and Tubaro, P. (2016). Special section on “Visualization in mixed-methods research on social networks”. *Sociological Research Online*, 21(2).
- Dominguez, S. and Hollstein, B. (2014). *Mixed Methods Social Networks Research: Design and Applications*. Cambridge University Press.
- Dunbar, R. (2016). Do online social media cut through the constraints that limit the size of offline social networks? *Royal Society Open Science*, 3(1):150292.
- Ellis, C. (1995). Emotional and ethical quagmires in returning to the field. *Journal of Contemporary Ethnography*, 24(1):68–98.
- Hargittai, E. (2010). Digital na(t)ives? variation in internet skills and uses among members of the “net generation”. *Sociological Inquiry*, 80(1):92–113.
- Hogan, B. (2018). Social media giveth, social media taketh away: Facebook, friendships, and apis. *International Journal of Communication*, 12:592–611.
- Hogan, B., Carrasco, J. A., and Wellman, B. (2007). Visualizing personal networks: Working with participant-aided sociograms. *Field Methods*, 19(2):116–144.
- Iphofen, R. (2011). *Ethical Decision-Making in Social Research*. Palgrave Macmillan.
- Jeffries, V. (2011). *Handbook of Public Sociology*. Rowman & Littlefield Publishers.
- Kadushin, C. (2005). Who benefits from network analysis: ethics of social network research. *Social Networks*, 27(2):139–153.
- Krackhardt, D. and Hanson, J. (1993). Informal networks: the company behind the chart. *Harvard Business Review*, 71(4):104–111.
- Kraut, R., Patterson, M., Lundmark, V., Kiesler, S., Mukopadhyay, T., and Scherlis, W. (1998). Internet paradox. a social technology that reduces social involvement and psychological well-being? *The American Psychologist*, 53(9):1017–1031.
- Levine, M. (2012). Loneliness and eating disorders. *The Journal of Psychology*, 146(1-2):243–257.

- MacKenzie, C. A., Christensen, J., and Turner, S. (2015). Advocating beyond the academy: Dilemmas of communicating relevant research results. *Qualitative Research*, 15(1):105–121.
- Mills, A., Durepos, G., and Wiebe, E. (2010). *Encyclopedia of Case Study Research*. Sage.
- Molina, J. and Borgatti, S. (2019). Moral bureaucracies and social network research (forthcoming).
- Molina, J., Maya-Jariego, I., and McCarty, C. (2014). Giving meaning to social networks: methodology for conducting and analyzing interviews based on personal network visualizations. In Dominguez, S. and Hollstein, B., editors, *Mixed-Methods Social Networks Research*, pages 305–335. Cambridge University Press.
- Pallotti, F., Tubaro, P., Casilli, A., and Valente, T. (2018). “You see yourself like in a mirror”: The effects of internet-mediated personal networks on body image and eating disorders. *Health Communication*, 33(9):1166–1176.
- Parker, M. (2013). The ethics of open access publishing. *BMC Medical Ethics*, 14(16).
- Patel, K., Tchanturia, K., and Harrison, A. (2016). An exploration of social functioning in young people with eating disorders: A qualitative study. *PLOS ONE*, 11(7):1–23.
- Ragin, C. (1992). Cases of “What is a case?”. In Ragin, C. and Becker, H., editors, *What is a Case? Exploring the Foundations of Social Inquiry*, pages 1–17. Cambridge University Press.
- Rainie, L. and Wellman, B. (2012). *Networked: The New Social Operating System*. MIT Press.
- Reyes-García, V. and Sunderlin, W. (2011). Why do field research? In Angelsen, A., Larsen, H., Lund, J., Smith-Hall, C., and Wunder, S., editors, *Measuring Livelihoods and Environmental Dependence: Methods for Research and Fieldwork*. Earthscan.
- Ryan, L. and D’Angelo, A. (2018). Changing times: Migrants’ social network analysis and the challenges of longitudinal research. *Social Networks*, 53:148–158.
- Ryan, L., Mulholland, J., and Agoston, A. (2016). Talking ties: Reflecting on network visualisation and qualitative interviewing. *Sociological Research Online*, 19(2):16.
- Schurmans, M., Charmillot, M., and Dayer, C. (2014). Introduction du dossier “la restitution des savoirs”. *SociologieS*.
- Stake, R. (1995). *The Art of Case Study Research*. Sage.
- Tubaro, P. (2017). Les tensions entre sociologie et politique à l’aune d’une tentative de législation des sites web sur les troubles alimentaires. *SociologieS*.
- Tubaro, P., Casilli, A., and Mounier, L. (2014a). Eliciting personal network data in web surveys through participant-generated sociograms. *Field Methods*, 26(2):107–125.
- Tubaro, P., Casilli, A., and Sarabi, Y. (2014b). *Against the Hypothesis of the End of Privacy: An Agent-Based Modelling Approach to Social Media*. Springer.
- Tubaro, P. and Mounier, L. (2014). Sociability and support in online eating disorder communities: Evidence from personal networks. *Network Science*, 2(1):1–25.

- Tubaro, P., Ryan, L., and D'Angelo, A. (2016). The visual sociogram in qualitative and mixed-methods research. *Sociological Research Online*, 21(2).
- Turkle, S. (2012). *Alone Together: Why We Expect More from Technology and Less from Each Other*. Basic Books.
- Valente, T. (2010). *Social Networks and Health*. Oxford University Press.
- Wellman, B., Carrington, P., and Hall, A. (1988). Networks as personal communities. In Wellman, B. and Berkowitz, S., editors, *Social Structures: A Network Analysis*, pages 130–184. Cambridge University Press.
- Wellman, B. and Gulia, M. (1999). Net surfers don't ride alone: Virtual communities as communities. In Wellman, B., editor, *Networks in the Global Village*, pages 167–194. Westview Press.
- Wellman, B. and Haythornthwaite, C. (2002). *The Internet in Everyday Life*. Wiley-Blackwell.
- Wolf, S., Branum, R., Koenig, B., Petersen, G., Berry, S., Beskow, L., Daly, M., Fernandez, C., Green, R., LeRoy, B., Lindor, N., O'Rourke, P., Radecki Breitkopf, C., Rothstein, M., Van Ness, B., and Wilfond, B. (2015). Returning a research participant's genomic results to relatives: analysis and recommendations. *The Journal of Law, Medicine & Ethics*, 43(3):440–463.
- Wong, C. A., Hernandez, A. F., and Califf, R. M. (2018). Return of Research Results to Study Participants: Uncharted and Untested. *Journal of the American Medical Association (JAMA)*, 320(5):435–436.
- World Health Organisation (2007). People-centred health care: A policy framework. Report, World Health Organization, Geneva.
- Yin, R. (2009). *Case Study Research: Design and Methods (4th ed.)*. Sage.

The trainer, the verifier and the imitator: Three ways in which human platform workers support automation

Paola Tubaro, Antonio A. Casilli, and Marion Coville

Abstract

This paper aims to shed light on the role of digital platform labor in the development of today’s artificial intelligence (AI), predicated on data-intensive machine learning algorithms. Focus is on the specific ways in which outsourcing of data tasks to myriad “micro-workers”, recruited and managed through specialized platforms, powers virtual assistants, self-driving vehicles and connected objects. Using qualitative data from multiple sources, we show that micro-work performs a variety of functions, between three poles that we label, respectively, “AI preparation”, “AI verification”, and “AI impersonation”. Because of its wide scope of application, it is a structural component of automation and an essential part of contemporary AI production processes – not an ephemeral form of support that may vanish once the technology reaches maturity stage. Through the lens of micro-work, we prefigure the policy implications of a future in which automation does not replace human workforce but implies its marginalization and precarization.

Key words: Digital platform labor, micro-work, datafied production processes, artificial intelligence, machine learning

Paola Tubaro
Laboratoire de Recherche en Informatique, Centre National de la Recherche Scientifique, Paris
Saclay, e-mail: paola.tubaro@lri.fr

Antonio Casilli
Telecom Paris, e-mail: antonio.casilli@telecom-paris.fr

Marion Coville
IAE Poitiers, e-mail: mcoville@poitiers.iae-france.fr

1 Introduction

Recent spectacular progress in research on artificial intelligence (AI) has revamped concerns that date back to the early nineteenth century, when the idea that machines may supersede human labor first spread among scholars, policymakers and workers. A well-publicized prospective literature emphasizes potential job losses due to automation (Brynjolfsson and McAfee, 2014; Chui et al., 2016; Frey and Osborne, 2017). While some scientific research corroborates these predictions (Acemoglu and Restrepo, 2018), other studies highlight historical dynamics of complementarity rather than substitution between human labor and machinery (Autor, 2015; Bessen, 2017), with more complex outcomes such as polarization between high- and low-skilled workers (Autor and Dorn, 2013).

However divergent these analyses may be, they share a quasi-exclusive focus on the expected spillovers of AI to other economic sectors, inferred from long-run industry trends and from known effects of previous waves of labor-saving mechanization. Less commonly discussed is the place of labor in the production of AI itself. Are there any distinctive effects of AI as it is conceived and implemented at the present moment?

The present paper builds on the assumption that the specificity of today's AI production processes, and their unique effects on labor, reside in their role in a "datafied" economy (Mayer-Schönberger and Cukier, 2013). Contemporary AI solutions are predicated on machine learning algorithms with a voracious appetite for data, despite a history of diverse approaches and visions (Domingos, 2017). Viewed as part of a broader drive to accumulate and leverage data resources (Kitchin, 2014), today's AI production reveals its need for human help not only to design cutting-edge algorithms (highly-qualified engineers and computer scientists) but also at a much more basic level, to produce, enrich, and curate data. This is the role of "micro-workers" (Irani, 2015a): barely visible and poorly compensated contributors, who operate remotely online from their computer or smartphone to flag inappropriate web content, label images, transcribe or translate bits of text, read aloud and record short sentences. These are activities that humans can do quickly and easily (whence the "micro" adjective), yet more efficiently than computers. They are fragments of large-scale data projects that support machine-learning algorithms.

Against popular discourses, the very existence of micro-work suggests that today's AI industry is in fact labor-intensive, although under less-than-ideal working conditions – with low pay, no employment contract, and absence of social protection (Berg et al., 2018). Even whilst automation is still in the making and has not yet been deployed at large scale, its demand for micro-tasks is already transforming the daily practices, experiences and career trajectories of thousands of workers worldwide (Gray and Suri, 2019). In the language of (Ekbj and Nardi, 2017), this is an instance of "heteromation", a neologism that stresses how, against a myth of automation capable of liberating people from the need to toil, labor demand is still high but humans operate on the margin of machines and computerized systems. Gray and Suri (2017) call this mixed configuration of machinery and human activ-

ity the “paradox of automation’s last mile”, the incessant creation of residual human tasks as a result of technological progress.

This (still scant) literature provides only the global picture, though, without looking deeper into the specific functions of micro-work in today’s AI industry. At what stage(s) of the production chain are humans needed? That is, what micro-tasks support what processes? Which of these tasks, if any, are temporary solutions to fill workflow gaps that will be probably resolved in future? Which ones, instead, fulfill structural needs and are therefore likely to be permanently needed? Are there any tasks or processes where human intervention is more embarrassing for AI producers – and which, therefore, are more likely to obfuscate its role? Answering these questions is important to unpack the nature of the linkages between artificial and human intelligence, and the re-organizations of labor that data-powered technologies induce. A better understanding of these processes is also important to inform policy action: if demand for micro-work is a transitory phenomenon, short-term measures will suffice, but if not, a more profound re-thinking of labor protection will be needed.

In this paper, we review a wide range of micro-working activities and show that they perform not just one, but a continuum of crucial functions, between three poles that we label, respectively, “AI preparation”, “AI verification”, and “AI impersonation”. We conclude that because of its wide and diverse scope of application, micro-work is a structural component of automation and an essential part of data-intensive AI production processes – not an ephemeral form of support that may vanish once the technology reaches maturity stage. In short, it is here to stay.

Our findings invite platforms and regulators alike to take concrete steps in regard to the working conditions, remunerations and career prospects of the people who toil behind the successes and promises of present-day AI. This requires a major effort to raise awareness and change mindsets, insofar as micro-work has remained largely out-of-sight so far – not explicitly considered even in otherwise laudable attempts to develop ethical principles for AI (Jobin et al., 2019).

2 Micro-work as an instance of digital platform labor

Let us start with some background on micro-work. Like other forms of digital labor (Casilli, 2019), it is an outcome of the emergence of platforms as devices to coordinate economic activity between service providers and clients – both construed as independent businesses that endeavor to make a one-off deal, rather than parts of a long-term employer-employee relationship. Platforms support large-scale outsourcing by enabling client companies to access workforce on demand, at a fraction of the cost of salaried staff, and usually with quicker turnaround times. They advertise themselves to clients as AI-service vendors, and to workers as providers of online earnings opportunities.

The most famous micro-work platform is Amazon Mechanical Turk, originally an internal service that the company developed to remove duplicates from its cata-

logue. It turned out that distributing tasks to a broad set of people working in parallel, was much more effective than any attempt to solve the problem algorithmically. Toward the mid-2000s, Amazon opened its service to external users, and is now an intermediary in a two-sided market (Rochet and Tirole, 2003) of workers and “requesters”. Many more platforms, such as Microworkers at international level, and Foule Factory¹ in France, adopt this very business model today, with minor variants. Few micro-work platforms serve a single monopsonist, usually a major tech company: TryRating for Apple, RaterHub for Google, and UHRS (Universal Human Relevance System) for Microsoft. There are also mixed models: the German Clickworker offers both a marketplace like Amazon and a managed service for more demanding clients, as part of which it recruits workers for, and administers access to, UHRS.

Additionally, micro-work platforms differ in size and scope. Some are tiny start-ups, others have grown to multi-nationals, such as Appen, a publicly-traded company headquartered in Sydney which has lately acquired other major players, notably the Americans Leapforce in 2017 and Figure Eight (formerly called Crowd-Flower) in 2019. Finally, some platforms such as Mechanical Turk and Foule Factory cater to a diverse range of corporate needs in addition to AI, others such as the French IsAHit propose a narrower set of services including an explicit AI offer, and a growing number specialize exclusively in AI services (Schmidt, 2019). The latter often involve alliances between a company that specializes in sales to AI companies, and another that manages micro-workers: for example Figure Eight and Clixsense, Mighty AI (acquired by Uber in June 2019) and Spare5, Wirk and Foule Factory.

In the typology proposed by Schmidt (2017), and re-elaborated by Berg et al. (2018), platform micro-work is an instance of “cloud-work”, performed remotely online. It differs from the other main variety of cloud-work, web-based freelancing, which concerns creative work (such as graphic design and software development), involves qualified professionals, and entrusts them with whole, relatively long projects rather than single short tasks. Examples of freelancing platforms include Upwork and Freelancer.com. Another way to see the difference between micro-work and freelancing is that the former is dispersed to an undefined set of anonymous, replaceable contributors via the platform, rather than assigned in full to a selected, identified contractor. For this reason, micro-work is sometimes referred to as “crowd-work” or “crowdsourcing”, contraction of “crowd” and “outsourcing” – a term that, however, wipes away the fragmented nature of the tasks and the small remunerations they attract (as low as one or two cents).

Both forms of cloud-work differ from “gig” labor where platform-mediated services are performed offline, even though coordination occurs online: for example, driving for Uber, delivering food for Glovo, cleaning for Helpling. Yet one form of gig-work, which Schmidt (2017, p. 7) calls “local micro-tasking”, is similar to micro-work insofar as it consists of small tasks (such as taking pictures of products in shops) given to an unspecific set of providers. Local micro-tasking platforms with

¹ “Foule” means crowd in French.

operations in varied European countries are Roamler and BeMyEye, while Clickworker offers both online and local micro-tasking.

Some of the tasks that micro-workers do are similar to the unpaid contributions of many users of digital technologies who unwittingly feed into tech companies' databases with their everyday routines. The reCAPTCHA algorithm that Google deployed to tell apart people from bots famously had the accessory purpose of leveraging human computation to recognize words and characters that software could not read (von Ahn et al., 2008). The difference is that micro-work platforms pay – and even if rates are low, they are enough to affect users' motivations and clients' expectations.

The nascent literature on platform labor has often conflated micro-work with other forms of platform-based digital labor, especially freelancing. Indeed some (online-only) micro-work platforms occasionally make available more qualified tasks, such as translations or text-writing, and conversely some freelancing platforms happen to publish simpler, lower-compensated tasks. This approach has been useful in a recent valuable effort to assess the size and growth rate of the whole online (or cloud-based) global labor market (Kässi and Lehdonvirta, 2018). However, this obscures the linkages with AI production. There is now a need to decouple micro-work from other forms of platform labor, which serve a range of economic and societal needs – from urban transportation to companies' branding – which do not necessarily tie in to automation. How do paid micro-tasks precisely affect the data economy behind automation, and how does their role in this particular production chain distinguish them from platform work more generally? For greater clarity and focus, we address this question in reference to online-only micro-work, leaving aside local micro-tasking.

3 The labor-for-data needs of machine learners

AI producers are companies, start-ups and research labs that use machine learning to develop applications ranging from chatbots and hands-free vocal assistants, to automated medical image analysis, self-driving vehicles and drones. Let us first review the basic functioning of machine learning, and derive preliminary conjectures on where and when it may need human intervention.

3.1 How data fuel machine-learning algorithms

At the crossroads of informatics and statistics, machine learning “teaches” computers to find solutions from data, without each step having to be explicitly programmed (Alpaydin, 2014, 2016). Its quality can get progressively better over time, depending not only on the algorithm but also on the data given to it. For example, development

of a vocal assistant requires huge audio datasets with examples of potential user requests (like “turn on the light in the kitchen”, “call mum”, etc.).

So-called “supervised” machine learning algorithms, the most widely used in both research and industry to date, need not only high-quantity, but also high-quality data, that is, complete with annotations. Supervised machine learning aims to infer a function that maps an input to an output based on exemplary input-output pairs (“training” dataset). The learned function must be able to assess new cases in “test” datasets. For example, to teach a computer to distinguish between images of dogs and other animals, one would need a training dataset that associates each image (input) to an annotation, such as a tag that says whether the image shows a “dog” or “other” (output); after having been exposed to many tagged images, the algorithm will be able to classify new, untagged images and determine whether they represent dogs. The more accurate the tags in the training dataset, the more the solution can be fine-tuned and generalized to a wide range of real-world cases.

Despite its apparent simplicity, the image recognition classification algorithm just outlined has far-reaching implications, for example in the development of self-driving cars, which need to recognize objects such as a dog crossing the street, before they can make decisions. A state-of-the-art example of the supervised family is “deep” learning, which analyzes data through a layered structure of algorithms inspired by the neural network of the human brain, leading to more effective learning.

Less demanding in terms of data quality is “unsupervised” machine learning, where data have no labels, and the algorithm is left to find its underlying structure based on common patterns. Two main types of algorithms can be distinguished, dimensionality reduction which consists in mapping a multidimensional dataset into more interpretable two-dimensional structures, and clustering, which groups observations into coherent classes. Today’s unsupervised learning brings to the next level some quantitative techniques traditionally used in social science, namely factorial analysis and hierarchical clustering (Boelaert and Ollion, 2018). Like these older tools, it is often used when the objective is unclear, or for exploratory analysis. It is unlikely to wipe out the supervised variant, because there are still many tasks it cannot do. Besides, interpretation of results can be problematic due to lack of objective standards to judge algorithmic performance.

A third family is reinforcement learning, a formalized version of human trial-and-error which uses mapping between input and output like supervised learning, but unlabeled data like unsupervised learning. It includes a feedback loop that gives the algorithm positive and negative signals, so that it adjusts accordingly. Because of its massive data needs, reinforcement learning is often applied to domains such as games, where simulated data are relatively easy to source.

3.2 A structural demand of micro-work for AI data preparation?

The above summary suggests that AI companies depend heavily on data resources, including not only raw data but also annotations that add extra meaning by associating each data point, such as an image, with relevant attribute tags. To account for these dual aspects, we propose to distinguish data generation and data annotation² as two separate sub-processes in the production chain of AI, even though they have much in common. Both are part of the preliminary phase, the first stage of the AI production chain – which we label “AI preparation”. The problem is that both are challenges for AI producers, despite a widespread rhetoric of “data deluge” (Anderson, 2008): the right data are not always available or accessible, and when they are, they often lack suitable annotations, and need some interventions before they can be used. On this basis, we can formulate a preliminary expectation, to guide our empirical analysis. It is that micro-work caters precisely to these unmet data needs: it contributes to AI preparation, both in terms of data generation and data annotation. Micro-work is an input to AI in the current data economy.

Our other expectation is that micro-work is a structural rather than a temporary input to AI production. We are wary of the views of some technology enthusiasts who believe that data generation and annotation tasks will ultimately be fully automated. The “heteromation” paradigm (Ekbia and Nardi, 2017) implies instead that some essential tasks will always be directed to humans as indispensable though hidden providers, any solution resulting from a division of labor between the two. Our above review of machine learning techniques corroborates this line of thought, suggesting that their huge data needs, which grow with every new challenge, will keep demand for micro-work high in the foreseeable future.

4 Insights from fieldwork

Our expectations are not hypotheses to be tested *stricto sensu*. We take them simply as the starting points with which we enter the research setting, the prior assumptions that derive from the still limited social research on AI and platform labor, and from contextual knowledge of the basic linkages between AI, machine learning, and data. Being reflexively aware of these initial hunches is a guide toward comparing them to stakeholders’ discourses and their actual experience, but does not engage us to stick to them. Rather, we use empirical evidence to enrich and substantiate these preliminary expectations, to complexify them, and if necessary to revise them, in an iterative, emergent process.

The data we use are from DiPLab (“Digital Platform Labor”), a mixed-methods study of tech companies and of the day-to-day routines of platform workers that took

² For the purposes of this paper, it is unnecessary to distinguish between training, validation and test data as is often done in machine learning: all data need to be generated and, in the case of supervised learning, annotated.

place in 2017-18 in France (Casilli et al., 2019). A country with traditionally high technological and scientific development, France is currently the second European country by number of AI start-ups (MMC Ventures, 2019), after significant public and private investments (French Government (Ministry for the Economy, Ministry of Education and Ministry of Digital Technologies), 2017; Villani et al., 2018). In addition to its inherent interest, focus on France enables to extend our gaze beyond the high-profile platforms, particularly Amazon Mechanical Turk, which have been overrepresented in the literature to date, despite being little used outside the United States. With its numerous, competing players, some of which operate only within national boundaries, France is exemplary of a trend toward diversification and specialization of micro-work platforms, attracting an ever-wide range of users.

We combine, on the one hand, insight from AI producers and micro-work platforms, and on the other hand, the views of people who perform micro-tasks online. We triangulate information obtained from these different stakeholders in order to reach greater consistency and completeness, to cross-check findings and to corroborate them. We do so because each of these stakeholders has a different perspective and, taken in isolation, would provide only a partial and incomplete view. Leaders and staff of AI companies have the best understanding of technology and its needs, but as we will show, they are often unwilling or unable to disclose their use of human workers. In turn, micro-work platforms that mediate between AI companies and workers, are best positioned to know the structure of the market, but their need to attract clients and investors may bias their communication strategies. Finally, workers can share the unique, concrete experience of doing micro-tasks, but they are not always aware of their purposes and final uses.

To uncover the viewpoint of AI producers and of the platforms/vendors that supply data laborers to them, we use primarily an inventory of micro-work platforms and related AI data service vendors. Although focus is on France (Wirk - Foule Factory, IsAHit), the inventory also includes information on international platforms whose scope of activity is global and includes France (like Appen, Clickworker, Lionbridge, Mechanical Turk, Microworkers), and for comparison purposes, it adds a few AI start-ups with more limited penetration in France (like Mighty AI). We compiled this inventory using desk research. We explored industry reports, newspaper articles, and most importantly the websites, press kits, and other communication tools of the platforms and companies concerned. To validate our findings from this material, and to gain more insight into less publicized aspects, we use in-depth interviews with 3 French clients and platform owners.

To account for micro-workers' perspective, we rely on a survey that we distributed in 2018 as a paid task on Foule Factory, collecting 908 unique, complete responses. For the purposes of the present paper, we use only one particularly relevant open-ended question, which asked micro-workers to describe in their own words the last task done online. We coded responses independently and cross-checked our categories for greater reliability. Here, we analyze results qualitatively, to identify common patterns; a quantitative description can be found in Casilli et al. (2019). We also use in-depth interviews with micro-workers. We invited a sub-sample of 72 questionnaire respondents to a follow-up interview of 30 - 60 minutes. We did

another set of interviews (60 - 120 minutes) with 14 French micro-workers active on varied international platforms such as Appen, Clickworker and Lionbridge, and with 3 African micro-workers who do tasks for French requesters through IsAHit. All interviews have been audio-recorded and accompanied by a written report by the interviewer(s).

5 Micro-work for AI preparation

Let us start by looking at the role of micro-work platforms in the provision of data generation and annotation services for AI companies. Appen says openly that effectively harnessing the power of machine learning requires human skills³. Lionbridge AI sells “Machine intelligence, powered by humans”⁴. We first look at these value propositions through platforms’ communication materials, before turning our attention to the views of the underlying workforce.

5.1 Platforms’ offer of AI preparation services

An example of data generation that most platforms advertise to their clients is audio utterance collection, important to train voice-controlled devices. Platforms can leverage their contributor base to gather this data with a variety of vocal timbres, regional accents, uses of slang, and contexts (such as background noise). Platforms that operate at global level can replicate the data collection in different languages – Appen boasts over 180, Lionbridge 300. Platforms that operate at national level also have advantages: a vocal assistant to be sold in France, for example, must be trained in the country to learn French accents, the names of French cities and personalities, and other local specificities. The market seems small, but the process can scale: a producer of vocal assistant software that we interviewed, has built an application allowing users to customize the assistant to their needs. It integrates a “data generation” functionality through which users can request bespoke datasets: the company manages the order by passing it to a standard platform such as Mechanical Turk or Foule Factory, monitoring execution and ensuring delivery⁵.

Platforms present data annotation as their core offer to clients. With sound or text data, they propose services such as categorization of topics in a conversation, determination of emotions behind a statement, classification of intents, and identification

³ Appen website, consulted on 15 March 2019, URL: <https://appen.com/why-human-annotated-data-key/>

⁴ Lionbridge website, consulted on 14 March 2019, URL: <https://www.lionbridge.com/artificial-intelligence/>

⁵ In passing, we noted that this is a challenge for start-up engineers who, trained in technology, often feel unprepared to take on the community management or HR roles they would need to engage with human workers.

of parts of speech. With images and videos, the offer includes assignment of images to categories, detection of objects within images with dedicated tools such as bounding boxes (rectangles around the objects of interest), cuboids (3-D bounding boxes) or polygons (precise drawings around objects of interest, possibly of irregular shape), addition of in-image tags to each object, and labeling of anatomical or structural points of interest (like eyes in faces) with so-called “landmark annotation”.

Technology moves fast, and computers are now pretty good at tasks that seemed unsurmountable even just a few years ago, such as (to use the above example) telling apart a dog from another animal. Human capacity is now in demand to recognize details and nuances, indispensable to increase the precision of software for sensitive applications such as autonomous vehicles and medical image analysis. A state-of-the-art technique is semantic segmentation, much more precise than those mentioned above because it involves painstakingly separating every pixel of an image into the parts that a computer vision algorithm will have to recognize. On Lionbridge’s blog, a machine learning specialist speculates that pixel-accurate annotation is becoming the new norm, while rougher tools such as bounding boxes may eventually disappear⁶.

Such accuracy would be impossible if workers had to draw shapes with the functionalities of standard software. Micro-work platforms such as Appen and Lionbridge compete fiercely to develop cutting-edge tools, themselves based on machine learning, that increase the precision, scale and speed of human data annotators. Some AI start-ups have also joined the race, usually focusing on technological development and using one or more standard micro-work platforms to access human contributors. One solution consists in having workers manually label a sub-set of data, and then letting an algorithm learn those annotations and applying them to the rest of the dataset. Another relies on an automated tool that roughly pre-annotates objects (for example, by forming lines around cars in a traffic image), so that the worker only needs to adjust the details. Figure Eight’s “active learning” distributes labor between humans and machines:

Computers can automate a portion but not all of the data, thus requiring a human-in-the-loop workflow. In this environment, computers can complete the high confidence rows and humans the lower confidence⁷.

To summarize, technological progress has not eliminated the need for micro-tasking, but transformed it, integrating humans and computers more tightly. These evolutions accompany the growth of the business of AI preparation: the industry think-tank Cognilytica estimates the worldwide market for what we call data generation and annotation at over \$500M in 2018, expecting it to rise to \$1.2B by 2023. As part of this trend, the market for third-party data solutions attained \$150M in 2018, and will exceed \$1B by 2023 (Cognilytica, 2019). According to Lukas Biewald,

⁶ Lionbridge website, consulted on 17 August 2019, URL: <https://lionbridge.ai/articles/why-pixel-precision-is-the-future-of-image-annotation-an-interview-with-ai-researcher-vahan-petrosyan/>

⁷ Figure Eight website, consulted on 14 March 2019, URL: <https://success.figure-eight.com/hc/en-us/articles/202703295-Getting-Started-Ideal-Jobs-for-Crowdsourcing>

founder of Figure Eight, the recent rise of deep learning has boosted demand, because its complex algorithmic structures require much larger (labeled) data sets than other machine learning techniques:

Deep learning has been fantastic [...] We began noticing deep learning when we started having customers who would ask for tens of millions of data rows right off the bat.⁸

5.2 *Micro-workers' experience of AI preparation*

The concrete experience of micro-workers broadly confirms preliminary insight from platforms' communication. Our online survey provides evidence of data generation in the form of voice recordings: many participants reported having read aloud a few short sentences in French and audio-recorded them. Variants of this task include requests to record, say, five ways to ask a virtual assistant about the weather. Some micro-workers understood that this was "to help design intelligent virtual assistants controlling connected objects" (L.⁹). This task requires large numbers of participants to ensure sufficient variety, and it is unlikely to change much over time, in that linguistic skills and local knowledge cannot be easily replaced or outsourced to offshore providers.

Data annotation tasks are also common. Workers are often asked to classify objects such as DVD titles, photos and "virtual avatars" (R.). Sometimes, they have to associate images and names of commercial products as found in multiple online marketplaces – clearly to teach computers to recognize essential similarities (the product is always the same) despite dissimilar contexts (the websites differ). CV anonymization, reported by almost one fifth of respondents, was understood to be about "removing all distinctive marks that could be discriminatory" (C.). Of note, workers also had to tag the spaces in the document where names, birth dates and addresses were placed originally – arguably to help some recruitment algorithm to understand the structure of a CV. This is another task that requires local knowledge insofar as job application standards vary across countries.

In terms of image annotation, some respondents mentioned a task they called "motocross" where they had to identify different roads and tracks in photographs and to indicate the nature of the ground (pebbles, road, sand etc.). Some thought it was for a video game, others for a census of racetracks. This is because, as we soon realized, requesters vary widely in the extent to which they provide detailed information on their tasks, and on the purposes they serve, leaving workers often confused. A more dramatic example of the consequences of erratic information from clients is a task that asked micro-workers to tag vegetables (tomatoes, carrots, etc.) in pictures of salads. M. (30 years old, married, resident of a mid-sized city, full-time teacher and micro-worker in her spare time) found this task "silly" but simple

⁸ Interview of L. Biewald with B. Lorica, O'Reilly Data Podcast, 4 May 2017, URL: <https://www.oreilly.com/ideas/data-preparation-in-the-age-of-deep-learning>

⁹ Names of respondents have been changed to initials for confidentiality.

and adequately paid for the limited effort it required. She grasped that it served to develop some software application for nutrition. But D. (25 years old, unmarried, living in a rural area, unemployed) could not make sense of it:

They tell you: draw a circle around a tomato. We don't know why. I think everyone knows what a tomato is, I hope [...]. Then I think to myself: if it's there, it must be useful to someone, for something, but... Why, I don't know.

A type of task that did not surface in the questionnaire, but was mentioned in interviews, consists in categorizing videos, text and images circulating on the web. Flagging violence and pornography serves to train algorithms for content moderation. After the attacks of 2015-16, A. checked “monstrous” terrorist content for several weeks, 30 hours a week, as clients “were panicking”. Exposure to this content can be distressing (Roberts, 2019), although A. assures that she has found ways not to be personally affected. Only a small part of content moderation can be automated: any new types of data (here, videos) first require micro-workers to train future automated solutions.

In sum, micro-workers' experience confirms their important contribution to data generation and annotation for AI, although they are not always aware of it, and suggests that this role is neither temporally nor spatially concentrated.

6 Micro-work for “AI impersonation”

Our prior expectations about the linkages between micro-work and AI did not factor in scandals, yet examples abound. In 2019, investment firm MMC Ventures reviewed over 2,800 purported AI start-ups across Europe, and found evidence of AI consistent with their value proposition in about 60% of them. The newspapers that covered the story were eager to stress that, well, a whopping 40% of these start-ups do not do AI (see for example Ram (2019)).

The year before, we had heard similarly outraged voices – not from micro-workers, who often lack awareness of the ultimate goals that their activity is serving (subsection 5.2) and are therefore ill-positioned to judge whether an alleged AI is genuine. We interviewed K., a Parisian entrepreneur and start-up founder who blamed his competitors for their claim to do AI while, instead, they outsource all work to humans recruited through platforms overseas. He went as far as to claim that “Madagascar is the leader in French artificial intelligence”. Even more upset was S., a student who did an internship in an AI start-up that offers personalized luxury travel recommendations to the better-off. His company's communication strategy emphasized automation, with an alleged recommender system based on users' preferences extracted from social media. But behind the scenes, it outsourced all its processes to micro-providers in Madagascar. It did no machine learning, and the intern could not gain the high-tech skills he had dreamt of.

Why do start-ups cheat? One hint comes from the prices that data vendors post online (only a few of them do, while the majority propose only personalized quotes

to individual clients). Top-quality data are expensive: semantic segmentation costs a few dollars per image, compared to bounding boxes that are priced less than a dime, and simple categorizations that are available for one or two cents. Costs go further up if accuracy of results is sought, for example by having each data point annotated by multiple workers. To this, start-ups should add the costs of powerful hardware and the salaries of highly qualified computer scientists. Under pressure to perform, they may find it cheaper to fragment the work into micro-tasks and sub-contract them to low-paid workers through platforms.

If so, there can be another role for micro-work in addition to AI preparation, and we propose to label it “AI impersonation”. It happens when humans, so to speak, steal computers’ jobs. After all, this is the very idea behind Amazon Mechanical Turk, the platform that first popularized micro-work. Its name is that of a fake chess-playing machine built in the late eighteenth century and dressed in seemingly oriental clothes, but in fact operated by a human player hidden inside. That Amazon dubbed its creation “artificial artificial intelligence” is also indicative of its intent: filling the gap of what artificial intelligence is expected but unable to do (Irani, 2015b).

Seen in this way, impersonation is not just about fraud, and indeed Amazon has always been upfront about it. It is an interpretation of the “human-in-the-loop” that makes workers hardly distinguishable from algorithms. Amazon’s original idea was to allow programmers to seamlessly integrate the two into their processes, whereby managing a task for “Turkers” would be similar to sending a remote request for an algorithm to execute. Today, the typical system has some sort of self-monitoring capability, which allows it to recognize situations in which it has doubts, or in which it cannot autonomously bring an activity to completion: in these cases, it hands control over to a human operator. This is the approach followed, among others, by Google Duplex, a conversational assistant that makes restaurant reservations, where up to 25% of calls are made by humans as of May 2019 (Chen and Metz, 2019). An apparent deception, it is nevertheless a way to gradually train the assistant¹⁰. In some cases, impersonation may involve qualified labor rather than micro-work. The creators of Julie Desk, a French start-up producing an email-based scheduling assistant, initially did the job by hand, in place of an algorithm that had yet to be coded:

We worked as assistants ourselves for a period of 8 months, and manually answered all the requests we received! It allowed us to understand what the recurrent patterns in the meeting scheduling process were and then, with the help of data scientists, we coded them to give birth to Julie¹¹.

Micro-work platforms do provide human labor force to meet these needs, but tend not to explicitly advertise these roles: they arguably negotiate impersonating tasks individually with clients as part of their managed service packages. Humans

¹⁰ Google AI Blog, accessed 15 March 2019, URL: <https://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html>

¹¹ Julie Desk blog, accessed on 15 March 2019, URL: <https://www.juliedesk.com/blog/artificial-intelligence-empathy/>

are always in the loop, but they are even less visible here, than when they do “preparation” tasks. To take these aspects into account, it is important to include impersonation in our framework: it is not just a temporary strategy to keep afloat an insufficiently funded start-up, but part of the “heteromated” system that increases demand for human workers with every new problem to be solved, while keeping them in marginal and unrecognized roles.

7 Micro-work for AI verification

Even after inclusion of AI impersonation together with AI preparation, it appears from our interviews there are more services that micro-workers provide to AI. In spring 2019, public outcry followed the revelation in the news that human workers listen to users’ conversations with smart assistants (see for example Hern (2019)). The year before, we had interviewed J., a transcriber who worked for six months to improve the quality of the French version of one of these virtual assistants, produced by a tech multinational and already in commerce. Her job was to check that the virtual assistant correctly understood what its users said. She listened to audio recordings (usually short tracks, averaging between 3 and 15 seconds), then compared her understanding to the automated transcription produced by the virtual assistant. If the transcription was inaccurate, she had to correct it: any misunderstanding, conjugation or spelling mistakes had to be highlighted. Another part of her work consisted in adding tags to the transcribed text indicating any sounds or events that could explain the virtual assistant’s performance – why some sentences were well understood, some not. J. knew that fellow transcribers were doing these tasks in other European countries and languages, all following the same guidelines.

Because the role of J. was undisclosed to users, her case reminds of impersonation, but the difference is that she was not replacing a failing algorithm: the one she checked for quality was up and running. She realized that the results of her work would help engineers and computer scientists to ensure the virtual assistant would not make the same mistakes in future. In this sense, the case of J. also reminds of what we called AI preparation, with the difference that she was intervening “post-robot”: she produced training data from the amended outputs of an already-trained algorithm. Therefore, we propose to call this case, similar yet not identical to the other two, “AI verification”.

We found other examples of AI verification in our fieldwork. A., the micro-worker who moderated violent content (section 5.2), also did relevance scoring. This type of tasks consists in assessing the extent to which the outputs of search engines or conversational agents are relevant to a user’s request, and sending feedback to help computer scientists improve their model. Another post-robot task consists in checking the results of optical character recognition (OCR) software. For example, a firm that aims to digitize its invoices may scan the original documents, use OCR to convert the resulting images into character codes, and get human help to look at the outcomes, fill the gaps, and make corrections if necessary. We interviewed three

African micro-workers who did such tasks for French clients through the Paris-based platform IsAHit. In both cases, humans intervene well after the preparatory phase, when the AI solution has already been trained, tested and brought to market.

Micro-work platforms do advertise services such as relevance scoring and transcription checking to their AI-producing clients, but do not group them together in a separate category corresponding to our AI verification. Partly, this is due to the same reasons that keep them quiet about impersonation: any discovery that a supposedly automated solution is at least partly hand-made, may be seen as a form of deception. Additionally, output checks performed by humans sometimes involve privacy leaks (to the discomfort of our interviewee J.) that may damage the reputation of the company or platform. Even more than in the other cases, micro-workers' contribution is surrounded by silence.

We make space for AI verification in our analysis, because of its potentially wide applicability. Checks of the accuracy and quality of algorithmic solutions will always be needed, regardless of whether supervised, unsupervised or reinforcement learning is used. By the same token, verification is not a temporary need but a recurrent one. As the sales of AI-based tools increase and concern a more diverse range of users, there will be a growing need to ensure that outputs meet expectations.

8 Discussion and conclusions

To map the linkages between AI and micro-work in our datafied economies, we started from the expectations that micro-work contributes to the preliminary, input phase of the AI production process, and that its contribution is structural rather than temporary. The former is in line with the communication strategies of platforms, and their insistence on the value of data produced and annotated with a "human touch"; the latter is at odds with the opinions of tech enthusiasts who anticipate full automation of data generation and annotation, but resonates with industry reports that the global market for human-powered data services for AI is growing.

We reflectively thought through our expectations in light of empirical evidence from desk research, responses to an online questionnaire and in-depth interviews. We compared and contrasted the voices of all stakeholders to probe and refine our ideas. This material corroborated our initial assumption of an important role of micro-work in AI preparation. But we also noted some anomalies that led us to broaden the set of roles that micro-work may play in AI production. Industry actors brought us to identify AI impersonation, which occurs whenever humans outperform computers, so that it is advantageous to use them instead of (parts of) algorithms. In turn, micro-workers' accounts of interventions to check the outputs of an automated system – that is, at the end of the production chain – prompted us to single out AI verification.

The result of this analysis is a typology, summarized in Figure 1. The process of AI production starts with preparation (left panel), which includes both data generation and annotation. This may concern image, text, sound, video or other types of

data, and it is largely outsourced to online micro-workers. The data that they produce or enrich feed an algorithm that learns a model (central panel) which in turn, returns an output with some degree of certainty. For example, the output of an image classification algorithm can be “it is 90% likely to be a dog, 10% likely to be another animal”. If impersonation occurs at all, it is at this stage. Humans replace part of the algorithm (when they step in to complete a task that, say, Google Duplex struggles to achieve) or all of it (when they entirely simulate an algorithm that has not yet been coded, as in the early days of Julie Desk). AI verification (right panel) is the process through which outputs are sent to micro-workers to be checked for accuracy and if necessary, corrected.

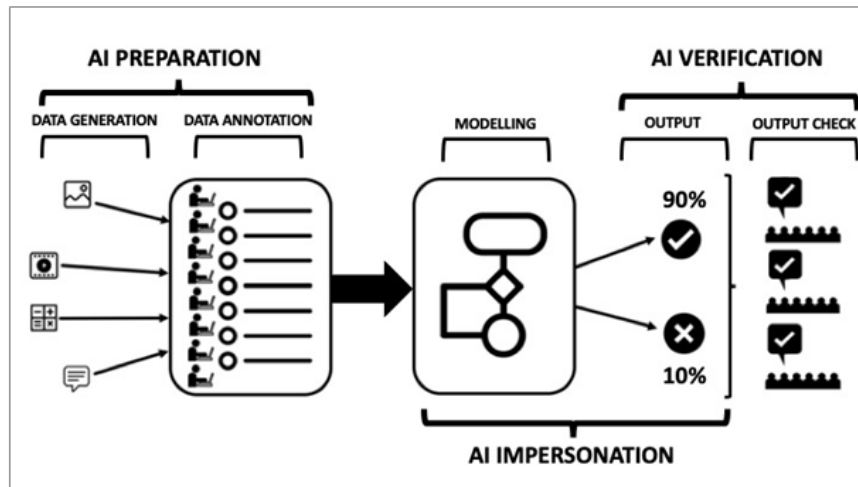


Fig. 1 The three main functions of micro-work in the development of data-intensive, machine-learning based AI solutions of today. Source: authors’ elaboration based on Casilli et al. (2019).

To be sure, there are possible overlaps between our three cases of AI preparation, AI verification and AI impersonation. On the one hand, both impersonation and verification may be first steps toward developing datasets that can be subsequently used for preparation purposes. On the other hand, the boundaries between verification and impersonation become fuzzy when humans intervene to correct errors in real time. There is in fact a continuum of functions for micro-work in AI, many real-world cases being positioned in-between the three main types that we have singled out.

There might even be cases that do not fit with any of these types. In France, we did not find so-called click-farms, where workers have to “like” (or dislike, share etc.) the webpages of brands, products, celebrities, sports teams or politicians. Usually even less paid than standard micro-work, and less often disclosed than the cases of impersonation mentioned here, these tasks are commonly outsourced to providers in low-income countries. Because they artificially inflate the indicators of quality and popularity used by (among others) search engines, they have lately been blamed

for contributing to the diffusion of fake news. In our perspective, they may be said to perform some kind of “AI disruption”. Future research may be able to assess their scope and consequences.

Refining our initial idea to add more roles of micro-work, at different stages of the AI production process, helps us find a comprehensive answer to the question of the extent to which micro-work is a temporary or structural component of AI. If we had focused exclusively on AI preparation, as in the discourse of most micro-work platforms, we might have thought that the need to accumulate labor-powered data is specific to the current times, in which AI is growing fast but has not reached maturity, and the less data-demanding unsupervised learning has not made enough progress. But in our data economy, this is unlikely to happen. Data availability will never reach a steady state: most use cases for machine learning require ongoing acquisition of new sources to continuously adjust to changing conditions, resulting in a steadily-growing need for humans to produce data for more accurate, more precise, and more profitable results. The discovery of AI verification strengthens this idea, in that some of the data used to re-train an existing algorithm and adapt it to new circumstances, come from the quality checks routinely done by humans. Taking into account AI verification, not just preparation, also contributes to dismiss the related idea that progress in unsupervised learning might eliminate the need for humans: even then, verification tasks will have to be performed.

Similarly, impersonation should be understood in light of the other two types. A single, perhaps highly mediatized case may suggest that micro-work is a transitory phenomenon, to disappear as companies accumulate the necessary data, skills, and computational capacity. But our typology supports the idea that impersonation is systemic and will always be present to some degree, because it ensures the necessary connection between AI preparation and AI verification, supplementing algorithms when they fail. Impersonation also demonstrates that the durability of demand for micro-work depends not only on technological, but also on economic factors. As long as there are humans who can perform tasks more cheaply than AI, perhaps (but not necessarily) because they reside in countries where the cost of labor is low, it will be advantageous to substitute them for machines. Overall, our typology hints that full automation is not to be expected any time soon, and that micro-work will continue to play an important role in keeping the industry going.

That micro-tasking is not a transitory phenomenon, and can occupy multiple roles in the AI production process, may make platforms pride themselves on countering the gloomy predictions of AI-induced job losses, by creating earning opportunities that would not exist without the technology. But difficult questions must be asked about the conditions under which micro-work is performed. Although a detailed analysis of working conditions is outside the scope of this paper, the interviews we used hint that executing un-challenging tasks such as labelling images for an unknown purpose can be destabilizing; that involuntarily accessing personal data of other people, or witnessing grossly deceptive forms of impersonation, brings ethical dilemmas; and that exposure to violent web content may generate distress (Casilli et al., 2019). Because platform labor occurs in a legal grey area (Prassl, 2018), no systematic ways to alleviate these problems exist.

We repeatedly noted the silence surrounding micro-work. Platforms tell clients that human contribution has value, but not who these humans are and in what conditions they work. As a result, clients know little about micro-workers – just as the latter are often unaware of the purposes of their tasks – and may find it difficult to interact with them as mentioned above. Ironically, the full extent of human intervention is unclear even to key industry actors. The incentive to obscure the role of human contributors is highest when the credibility of full automation promises is at stake. As a general tendency and beyond one-off revelations, this contributes to keeping micro-work far from the gaze of the general public and from the agenda of policy-makers. Out of the reach of institutional regulations, subject only to the forces of a market plagued by an excess supply of workers (Graham and Anwar, 2019), remunerations remain structurally low.

Micro-work is the already-visible tip of an iceberg of transformations brought about by data-driven automation. AI is not the end of human labor, but risks depriving it of the quality, meaning and social status that it had acquired over time. There is a need for ambitious, long-term policies that frame the further development of AI by taking into account the concrete conditions of its production, in light of ongoing debates on digital platform labor and its shortcomings – from low remuneration and precariousness to lack of social security (Graham and Shaw, 2017). Put differently, credible commitment to socially responsible AI requires the definition of labor standards in the processes that underpin it. More transparency is needed, toward workers as well as the general public, to ensure the full extent of human participation is understood and recognized for what it is worth.

9 Acknowledgments

This paper presents results of the “Digital Platform Labor” (DiPLab) project, co-funded by Maison des Sciences de l’Homme Paris-Saclay; Force Ouvrière, a workers’ union, as part of a grant from Institut de recherches économiques et sociales (IRES); and France Stratégie, a service of the French Prime Minister’s office. The platform FouleFactory/Wirk offered logistical support, and Inria provided complementary funding. For more information: <http://diplab.eu/>

The authors would like to thank current and former DiPLab team members, notably Maxime Besenval, Odile Chagny, Clément Le Ludec, Touhfat Mouhtare, Lise Mounier, Manisha Venkat and Elinor Wahal.

Preliminary versions of this paper were presented at the Reshaping Work Conference in Amsterdam, October 2018, and at the annual Research Day of the Society and Organizations (S&O) Center of HEC Paris, May 2019. Many thanks to discussants, organizers and participants for constructive feedback. We are especially grateful to Mitali Banerjee, Magali Delmas, Vili Lehdonvirta, and Anoush Margaryan for helpful suggestions. Any remaining errors or inconsistencies are exclusively ours.

References

- Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.
- Alpaydin, E. (2014). *Introduction to Machine Learning (3rd edition)*. MIT Press.
- Alpaydin, E. (2016). *Machine Learning: the New AI*. MIT Press.
- Anderson, C. (2008). The end of theory: The data deluge makes the scientific method obsolete. *Wired*.
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30.
- Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, 103(5):1553–97.
- Berg, J., Furrer, M., Harmon, E., Rani, U., and Silberman, M. (2018). Digital labour platforms and the future of work: Towards decent work in the online world. ILO Report.
- Bessen, J. (2017). Automation and jobs: When technology boosts employment. Boston Univ. School of Law, Law and economics research paper, no. 17-09.
- Boelaert, J. and Ollion, E. (2018). The great regression: Machine learning, econometrics, and the future of quantitative social sciences. *Revue française de sociologie*, 59(3):475–506.
- Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. Norton.
- Casilli, A. (2019). *En attendant les robots: Enquête sur le travail du clic*. Seuil.
- Casilli, A., Tubaro, P., Le Ludec, C., Coville, M., Besenval, M., Mouhtare, T., and Wahal, E. (2019). Le micro-travail en France. Derrière l’automatisation, de nouvelles précarités au travail ? Report of the Digital Platform Labor (DiPLab) project.
- Chen, B. and Metz, C. (2019). Google’s duplex uses a.i. to mimic humans (sometimes). *The New York Times*.
- Chui, M., Mayika, J., and Miremadi, M. (2016). Four fundamentals of workplace automation. McKinsey q. 2016 (1).
- Cognilytica (2019). Data engineering, preparation, and labeling for ai 2019. Report n. cgr-de100.
- Domingos, P. (2017). *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. Penguin.
- Ekbia, H. and Nardi, B. (2017). *Heteromation, and Other Stories of Computing and Capitalism*. MIT Press.
- French Government (Ministry for the Economy, Ministry of Education and Ministry of Digital Technologies) (2017). France intelligence artificielle – rapport de synthèse. Report.
- Frey, C. and Osborne, M. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114(C):254–280.

- Graham, M. and Anwar, M. (2019). The global gig economy: Towards a planetary labour market? *First Monday*, 24(4).
- Graham, M. and Shaw, J. (2017). *Towards a Fairer Gig Economy*. Meatspace Press.
- Gray, M. and Suri, S. (2017). The humans working behind the AI curtain. *Harvard Business Review*, (9):2–5.
- Gray, M. and Suri, S. (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.
- Hern, A. (2019). Apple contractors 'regularly hear confidential details' on Siri recordings. *The Guardian*.
- Irani, L. (2015a). The cultural work of microwork. *New Media & Society*, 17(5):720–739.
- Irani, L. (2015b). Difference and dependence among digital workers: The case of Amazon Mechanical Turk. *South Atlantic Quarterly*, 114.
- Jobin, A., Ienca, M., and Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1:389–399.
- Kitchin, R. (2014). *The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences*. Sage.
- Kässi, O. and Lehtonvirta, V. (2018). Online Labour Index: Measuring the online gig economy for policy and research. *Technological Forecasting and Social Change*, 137:241–248.
- Mayer-Schönberger, V. and Cukier, K. (2013). *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. Houghton Mifflin Harcourt.
- MMC Ventures (2019). The state of AI 2019: Divergence. Report.
- Prassl, J. (2018). *Humans as a Service. The Promise and Perils of Work in the Gig Economy*. Oxford University Press.
- Ram, A. (2019). Europe's AI start-ups often do not use AI, study finds. *Financial Times*.
- Roberts, S. (2019). *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press.
- Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029.
- Schmidt, C. (2017). Digital labour markets in the platform economy: Mapping the political challenges of crowd work and gig work. Report, Friedrich-Ebert-Stiftung.
- Schmidt, F. (2019). Crowdproduktion von Trainingsdaten: Zur Rolle von Online-Arbeit beim Trainieren autonomer Fahrzeuge. Report, Hans-Böckler-Stiftung.
- Villani, C., Schoenauer, M., Bonnet, Y., Berthet, C., Cornut, A., Levin, F., and Rondenpierre, B. (2018). Donner un sens à l'intelligence artificielle : pour une stratégie nationale et européenne. Report, France Premier Ministre.
- von Ahn, L., Maurer, B., McMillen, C., Abraham, D., and Blum, M. (2008). reCAPTCHA: Human-based character recognition via web security measures. *Science*, 321(5895):1465–1468.

Counting “micro-workers”: Methodological and societal challenges around new forms of labor

Paola Tubaro, Clément Le Ludec, and Antonio A. Casilli

Abstract

Standard statistics struggle to apprehend the workforce that performs tasks and gigs on digital platforms. Particularly difficult to measure are "micro-work" platforms that allocate fragmented data tasks to remote online providers, for remunerations as low as a few cents, with hardly any form of contractual or social protection. As a first step toward filling this gap, we develop an innovative mix of qualitative and quantitative tools (online surveys, in-depth interviews, capture-recapture techniques, web traffic analytics) to count micro-workers in an industrialized country, France. We identify three distinct worker populations, corresponding to different levels of engagement. A group of about 15,000 "very active" individuals, most of whom use these platforms at least once a week; a second featuring over 50,000 "regular" workers, overall more selective and present at least once a month; and a third circle of approximately 260,000 "occasional" workers, more heterogeneous and who often alternate inactivity and various levels of work practice. These figures urge a re-thinking of some key principles of labor statistics, and invite a broader discussion on the place of labor in our society.

Key words: Micro-work, digital platforms, labor statistics

Paola Tubaro
Laboratoire de Recherche en Informatique, Centre National de la Recherche Scientifique, Paris
Saclay, e-mail: paola.tubaro@lri.fr

Clément Le Ludec
MSH Paris Saclay

Antonio Casilli
Telecom Paris, e-mail: antonio.casilli@telecom-paris.fr

1 Introduction

The rise of “platform work” in the wake of the 2008 financial crisis undermines the accuracy of the instruments classically used to observe employment, unemployment and inactivity, such as the Labor Force Surveys that national statistical institutes conduct in European countries. Based on international classifications that allow comparisons across time and space, these surveys have traditionally well served the need of measuring population activity levels and of describing jobs. Yet they struggle to apprehend platform-based “gigs”, “rides” and “tasks” that challenge established definitions of salaried employment.

Lack of consensual definitions is a major obstacle to measuring this phenomenon. The surveys conducted so far, whether by official statistical agencies (for example, by adding a module or a few questions to an existing survey) or smaller-sized research organizations (for example Huws et al. (2017)), diverge in scope (which activities to include? which platforms, beyond the most high-profile ones?) and periodization (include all workers who used the platforms in the last week, as per Labor Force Surveys traditional protocols? or extend to the last month or even year? or perhaps include anyone who used platforms once in life?). There are also sampling problems: in large surveys fielded to the general population, endemic non-response (sometimes due to unclear questions) leads to underestimations (Eichhorst et al., 2017), while small, more focused surveys are unrepresentative and lead to over-estimations. Depending on their design, extant studies obtain very divergent estimates ranging from 0.4% to 22% of the size of Western economies (Frenken, 2018), although many of them are in the range of about 0.5% to 2% (O’Farrel and Montagnier, 2019).

Particularly difficult is the measurement of platform work that is performed entirely online. Some estimates, such as the one featured in ter Weel et al. (2018), even exclude it from scope. Unlike “geographically sticky” (Ojanperä et al., 2017) car rides and bike delivery services, such activities are often executed remotely, at home behind a computer or on a mobile phone while commuting. They even lack visual clues such as the pieces of equipment and corporate logos that gig workers wear, punctuating urban traffic. A recent study at the Oxford Internet Institute has taken an important step toward measuring online work, advancing that its global volume has increased by a whopping 25% between 2016 and 2018 (Kässi and Lehdonvirta, 2018).

However, this otherwise remarkable result does not distinguish between the two main configurations of online work. On the one hand, there are freelancing platforms that cater to companies’ need for on-demand talents such as graphic designers, software developers, and strategy consultants. These are selected, sometimes very well-paid, professionals. On the other hand, “micro-work” platforms entrust much smaller and simpler tasks to indistinct masses of largely anonymous, and often little qualified, providers¹. Micro-work is the most elusive form of work in the

¹ We prefer the term “micro-work” to the alternative “crowdsourcing”, first proposed by the *Wired* magazine in 2006 and resulting from contraction of the words “crowd” and “outsourcing”. Crowd-

platform economy. It encompasses tasks such as identifying or labeling objects on images, transcribing invoices, translating bits of text, moderating content (such as videos), sorting or ranking photographs, recording sentences spoken aloud. Standardized and repetitive, these tasks usually take only a few minutes, are typically paid on a piecework basis, and attract compensations that can be as low as a few cents.

Online micro-work is a by-product of today's "datafied" economy (Mayer-Schönberger and Cukier, 2013) because it supports the digital transformation of companies – for example, transcription of scanned invoices may serve the creation of a paperless archive. Disassembling bigger "projects" into tiny "tasks", and re-assembling workers' outputs at the end, comes with a promise of speed and cost-effectiveness. By mobilizing many micro-workers simultaneously for a few hours rather than a single subcontractor for a longer time, client companies conjugate process scalability and workforce flexibility. More prominently, micro-work feeds into the scientific and industrial development of artificial intelligence (Gray and Suri, 2017; Irani, 2016). Labels that identify objects in natural settings can be used to train machine learning models. For instance, labelled images of cars in road traffic are necessary to operate autonomous vehicles (Schmidt, 2019; Tubaro and Casilli, 2019), conversation transcripts help improve automatic speech recognition of smart speakers, and sets of predefined questions and answers can be used to develop virtual assistants (Casilli et al., 2019).

Like other forms of online platform labor, the size of this workforce is unknown. Platforms themselves claim that their user-base is large: if Amazon Mechanical Turk, the most widely-known micro-working service, boasted 500,000 "Turkers" as early as 2014², the Chinese giant Witmart alone is said to exceed 12 millions³. However as we will see, researchers have challenged these figures and come to much more conservative estimates using a range of methods (section 2.3). Also, these figures refer to single platforms without taking into account the variety of players that co-exist, both globally and in specific geographies. Thus, they are of limited help in orienting local policy actions.

As a first step toward filling this gap, this paper leverages a range of methods to estimate the number of people micro-working in one single country, namely France. This case is inherently interesting because France is a pioneer in information technologies and a heavy investor in artificial intelligence research. Nevertheless, the country is little documented in the (still scant) literature on micro-work, and its residents are relatively less present in international platforms such as Amazon Mechanical Turk. Partly for language reasons, France has its own micro-work platforms, the

sourcing includes both micro-tasking and some parts of freelancing, notably with platform designs that rely on spec work and tournament job where workers bid competitively against one another to win a client – so that the "crowd" does the job, but only one gets paid.

² "500,000 Workers from 190 countries", cited in M. Harris, "Amazon's Mechanical Turk workers protest: 'I am a human being, not an algorithm'", *The Guardian*, December 3, 2014.

³ *Les Echos*, "L'actu tech en Asie : Female founders veut percer le plafond de verre", June 19, 2015.

most popular of which is Foule Factory⁴, whose webpage displays 50,000 contributors.

To estimate the size of the French micro-working population, we draw on publicly available sources combined with the results of DiPLab (“Digital Platform Labor”), a comprehensive study of micro-work in the French-speaking world⁵. Using three methods to collect and analyze, respectively, public data sources, experimental data and web traffic data, we identify three distinct pools of micro-workers and we link them to diverse levels of engagement. They comprise approximately 15,000 “very active”, 50,000 “regular” and 260,000 “occasional” online platform micro-workers respectively.

To the best of our knowledge, ours is the first attempt to estimate the number of micro-workers in a specific country across micro-work platforms, considering a range of different uses and levels of engagement. It demonstrates the need to go beyond the most high-profile cases and to take into account local conditions in addition to global trends. In this sense, our results may inspire researchers interested in re-applying our methodology to other countries.

2 The challenges of measuring platform labor (especially micro-work)

Micro-work does not fit neatly into the categories ordinarily used to measure the labor force. Some of the difficulties are common to other forms of platform labor, others are specific to it. In what follows, we explore them and discuss how different tools (from general and ad-hoc surveys to publicly-available company data, and non-standard methods) may or may not help. We come to conclusions as to the best method-mix for our case.

2.1 *Why surveys and administrative data fail to detect micro-work*

Official statistical surveys are usually designed to report the main occupation of respondents, while the flexibility afforded by online platforms allows arranging work as a secondary activity, to earn a supplement of income. This is one reason why early studies such as that of Katz and Krueger (2019), which still followed standard statistics in tracing only people who perform platform work as a main job, obtained low estimates (0.5% of the adult population). A related issue is the possibility to perform platform work irregularly – an aspect that, again, Labor Force Surveys miss

⁴ “Foule” is the French for “Crowd”.

⁵ As part of DiPLab, we built an inventory of micro-work platforms and mobile applications being used in France, collected online questionnaire data, and interviewed in-depth a selection of micro-workers, clients, platform managers and other stakeholders. The data collection was conducted in 2017-18 (see Section 3.3 for more details).

insofar as they assume that those who answer “yes” to the question of whether they did any paid work in a given reference period, did so with some regularity.

Measuring platform work cannot be content with such binary questions and additionally requires assessing regularity, intensity and significance (Urzi-Brancati et al., 2019). If these issues are common to many forms of platform labor, micro-work amplifies them. The “unbundling” (Pesole et al., 2018) of large data jobs into tiny tasks makes them more likely to be part of a secondary or occasional activity rather than a main, stable occupation, at least in Western countries. Because these tasks are mostly simple, short and sometimes even “gamified”, and their purposes are not always disclosed, workers sometimes fail to even recognize them as labor. As a result, more than other platform-based activities, micro-work may go under-reported in general surveys.

While the language and structure of official statistical agencies slowly adapts to the fast-moving world of the platform economy, several groups of academic researchers, unions and think-tanks have attempted to launch smaller, dedicated surveys at country or cross-country levels (O’Farrel and Montagnier, 2019). The majority have attempted to estimate the number of platform workers in general, without specifically targeting micro-workers. However, a new wave of a European comparative survey (Huws et al., 2017) run by the University of Hertfordshire in collaboration with other research bodies has released more detailed results for Spain (Fundación Felipe González and University of Hertfordshire, 2019), where allegedly 17% of the adult population do platform work at least once a week, and 40% of them, amounting to over 3 million people, do online clerical work, small tasks and “click-work” – that is, a set of activities that is broader than our definition of micro-work, but includes most of it. Other studies enable indirect estimations. The COLLEEM survey, fielded in 14 European countries in 2017, estimates that slightly less than 10% of the 16-74 years old have ever done any platform work (Pesole et al., 2018), with differences across countries (12.5% in Spain, 7.6% in France). Of them, 24.7% are said to have performed micro-tasks at least sporadically (Urzi-Brancati et al., 2019). Applying these ratios to Spain would give a number of micro-workers just below one million, while for France, this would amount to roughly 900,000.

While these surveys are valuable tools to situate micro-work in the broader economic context, differences in how they operationalize the notions of secondary and irregular activity produce widely divergent results, and invite to caution in interpretation. The numbers for Spain in the two studies just cited differ markedly, and even the lower figure may be an overestimate insofar as the COLLEEM survey produced systematically higher values than a comparable one, the 2018 Eurobarometer (O’Farrel and Montagnier, 2019, p. 14)⁶.

In other settings, administrative data have proven to be a valuable resource that compensates for the limitations of surveys. Here, administrative data may also open

⁶ The Flash Eurobarometer 467 “The use of the collaborative economy” of 2018 estimates that 11% of the adult French population have ever offered services through platforms. 5% of them (about 300,000 people) have performed online services such as IT and accounting, while 7% (about 420,000) have offered “other” online services: neither category corresponds precisely to micro-work, but both cover parts of it.

new avenues if for example the fiscal authorities require platforms to report users' earnings for tax collection purposes (O'Farrel and Montagnier, 2019). However, a consequence of irregularity is that some workers earn very little, so that their platform incomes fall below the given tax thresholds. Thus, the extent to which their presence will be reported at all, will depend on the details of the specific legal requirements that countries might impose.

2.2 Why data reported by micro-work companies cannot be taken at face value

In the absence of official statistical data, many studies so far have simply chosen to rely on the information that platforms themselves release publicly, when available. On micro-tasking specifically, a widely cited 2015 World Bank report summed the number of registered users on the main international platforms then known (essentially Amazon Mechanical Turk and Crowdflower⁷), coming to an estimate of 5.8 million micro-workers worldwide (Kuek et al., 2015). In France, user registration data from platforms are featured in an influential report of the French General Inspectorate of Social Affairs (Amar and Viossat, 2016).

However, this method is problematic for at least three reasons. First, the number of sign-ups says nothing of people's actual levels of engagement, and it is therefore an incomplete information (unless the platform actively enforces a systematic deregistration policy for inactive users). Second, this method over-emphasizes high-profile players such as Mechanical Turk while under-researching the large number of smaller local and international platforms, which release incomplete or no information at all about their user basis. To the extent that these less visible platforms attract different users (or function differently under some other dimension), it may produce unreliable results. Third, this approach neglects the practice of registering with multiple platforms. Extensively analyzed in the economics of two-sided markets ("multi-homing", see Athey et al. (2018)) but curiously little explored in research on platform labor, multi-activity can partly be seen as another consequence of the accessory and/or irregular nature of micro-work. Insofar as workers sign up to multiple services, the sheer number of registered users can be leveraged only to analyze one single platform, but it will count some workers more than once if it sums registrations across platforms.

A specificity of micro-work, notably its remote nature and its geographical "unstickiness", adds an additional difficulty in that even platforms that provide global figures do not always give the breakdown of the number of registered workers by country. The researcher, then, must rely on educated guesses. Existing evidence on (mainly) freelancing online labor platforms suggests that most of the demand is located in North America and Europe, and most of the supply in South and

⁷ The platform Crowdflower re-branded as Figure Eight in 2018 and was acquired by another platform, Appen, in 2019.

South-East Asia (Graham et al., 2017). Thus, one might expect similar patterns to characterize micro-work whereby in the case of France, supply might largely come from French-speaking African countries following historical "coloniality" patterns (Casilli, 2017).

More importantly perhaps, the weakening of geographical boundaries exposes online micro-workers to broad competition, possibly producing over-supply of labor as has already been documented in the case of freelancing (Graham and Anwar, 2019). Wide swings in demand, dependent on the rapidly changing needs of the global artificial intelligence industry, exacerbate this tendency (Schmidt, 2019). This likely results in a very skewed distribution of levels of engagement, with many registered micro-workers whose activity is actually very low because they do not manage to align with task availability fluctuations.

Overall, data publicly provided by platforms, notably the number of their registered users, are more promising than general population surveys, but cannot be taken at face value. Corrections and adjustments are in order, and their interpretation calls for further thought.

2.3 The limitations of alternative approaches to estimating micro-work

To move beyond the limitations of data publicly provided by platforms, some researchers endeavored to interrogate them privately – assuming that platform owners have gut sense of the market in which they operate. Following this approach, Information Systems researchers at the University of Kassel looked at the case of Germany. They first mapped all platforms established in Germany and then surveyed them about their views of themselves and the whole sector. Notice that they did not focus on micro-workers but on “crowdworkers”, a notion that as noted above (footnote 1), encompasses both micro-workers and some types of freelancers. One question invited respondents to estimate the gross number of German crowdworkers throughout platforms. The average of all answers gives a rough figure of over one million, of which one quarter are allegedly “active”, though this notion is not precisely defined (Mrass and Peters, 2017). While interesting, this result is based on a low number of responses (17) with high variance.

Other studies explicitly probe publicly-announced company data, specifically the 500,000 announced by Amazon Mechanical Turk, the oldest and by far the best known platform. A key motivation of many of these studies is to assure the quality of the results of scientific surveys, questionnaires and experiments for which research teams in disciplines such as psychology, marketing, and linguistics recruit “Turkers” (Fort et al., 2011). This involves ensuring access to a sufficiently large pool of respondents (Keith et al., 2017). In this perspective, Stewart et al. (2015) borrowed the bio-ecological “capture-recapture” model (see section 3.3) to show the existence of a pool of 7,300 potential study participants on Mechanical Turk at any given time. In the same year, activist Kristy Milland conducted a six-week sur-

vey and counted about 30,000 workers⁸. Difallah et al. (2018) improved Stewart's model, allowing for a longer observation time and taking into account the heterogeneous propensity of individuals to accept tasks. They showed that about 100,000 people work on Jeff Bezos' platform, or one fifth of the advertised population. Finally, using data from TurkPrimer, an independent company that helps researchers design and implement studies on Mechanical Turk, a recent study finds 250,000 Turkers worldwide (Robinson et al., 2019).

Notice the shift of focus in these studies, which do not target a country across platforms, but a platform across countries. However, about 75% of Amazon Mechanical Turk workers are Americans (Difallah et al., 2018), and some of these estimates paid particular attention to them, partly to serve the needs of scientific studies that targeted US participants. Therefore, these findings say little on other platforms and on other countries – especially as Mechanical Turk is so unique by fame and longevity, that its characteristics may not transfer easily to other settings.

The immediate lesson to be learned from this brief review is that no method is perfect, but a combination of tools may ensure that they compensate for each other's limitations. Leaving aside general population surveys due to their high cost and uncertain results, in what follows we start by mapping the platforms that operate in one specific country, France (section 3.1). Then (section 3), we retrieve the numbers that these platforms provide publicly, and endeavor to correct them to consider different levels of workers' engagement as well as multi-activity. This will require deployment of a wider range of methods, including some of those previously discussed.

3 Three approaches to estimating the number of micro-workers in a specific country

In light of the above discussions, we now deploy a range of approaches to estimate the size of the French micro-working population. We start by identifying the micro-tasking platforms that recruit French workers (subsection 3.1). Then we follow the literature in referring to the figures publicly provided by these platforms (subsection 3.2), and then applying a “capture-recapture” model to a single platform (subsection 3.3). We then present an alternative method based on website audience tracking (subsection 3.4). We combine our different methods to improve accuracy and obtain a meaningful range of estimates (subsection 3.5) and finally, we refine our results by taking into account micro-workers performing tasks on several platforms (subsection 3.6).

⁸ The study is not published but the author has made her calculations available: Kristy Milland, 150717 Preliminary results - Mapping study, July 17, 2015, URL: <https://docs.google.com/spreadsheets/d/1T3yP\Jo4qELrwsE2NAPNs07L1AWmpAEr9vnhreGJl-K0/edit#gid=1993074859> (accessed 18 December 2018).

3.1 What micro-work platforms operate in France?

To map the micro-work market in France, we cannot limit our investigation to Amazon Mechanical Turk, where demand from French companies outweighs the supply of French micro-workers. The platform has long discouraged registrations outside the USA and, to a lesser extent, India. Recent research (Berg et al., 2018; Forde et al., 2017) covering multiple countries including France, has added to Amazon Mechanical Turk also Crowdfunder (see footnote 7), Microworkers, Clickworker⁹ and, in the case of Berg et al. (2018), Prolific. We include their selection of platforms except Prolific, a site specialized in surveys that exists only in English, and which does not seem to have users in France (see Berg et al. (2018)). We also retain ClixSense, which provides access to Crowdfunder.

In addition to these international platforms, we include FouleFactory, a platform that exclusively recruits its micro-workers in France and since its creation in 2014, has established itself as a key player in the market. It sells services to artificial intelligence and other hi-tech companies through its sister start-up Wirk, launched in 2018. We also include Ferpection, a French platform that recruits internationally. FouleFactory and Ferpection are potentially more attractive for French speakers than other platforms like Microworkers, where English language is required even to sign up. Finally, we add Appen, an international platform open to recruitment in France, which lists tasks of the same nature as the others, but has a tighter and more closed business model that enables it to manage large orders for clients, and to offer more regular assignments to workers. As noted earlier (footnote 7), Appen owns the former Crowdfunder but has left it to operate independently throughout 2019.

This list is not exhaustive and there are other micro-work websites and applications in France. We exclude micro-working platforms that seem to be very rarely used in the country, such as CrowdGuru and Spare5, and French platforms that are in-between micro-tasking and freelancing, such as Malt and 5euros.com, as the part of micro-work in their overall offer is small. We also exclude income-generating mobile applications such as BeMyEye and Roamlar, which do offer micro-tasks, but of a particular nature that requires physical presence in a given place (for example, taking pictures of products in supermarkets), so that they are not online-only. Finally, we must unfortunately exclude platforms that are in scope, but for which we could not find sufficient data, notably Lionbridge and Pacteria.

It is useful to characterize this selection of platforms according to their access policies (open or limited) and the presence of at least one welcome page in the local language (Table 1). With regard to the first criterion, notice that FouleFactory has closed registration of new members but reopens it intermittently in case of surge in the demand of micro-work by client firms. As for the second criterion, the presence of a French page seems to encourage access by new users.

⁹ In addition to connecting its own registered micro-workers and clients, Clickworker also provides an entry point for UHRS (Universal Human Relevance System), Microsoft's proprietary micro-work platform, where it operates as a monopsony.

Table 1 Characteristics of micro-work platforms included in our study

Platform	Open sign-up	Local language page
Amazon Mechanical Turk	No	No
Microworkers	Yes	No
Clickworker	Yes	Yes
ClixSense	Yes	No
FouleFactory	No	Yes
Ferpection	Yes	Yes
Appen	No	Yes

Source: authors' elaboration based on platforms' websites.

3.2 Analyzing public data available from platforms

As discussed above (subsection 2.2), figures published by micro-work platforms can be a useful starting point, although their interpretation calls for caution. Table 2 presents the figures provided by our selection of platforms. Clearly, international platforms have a much larger user base than national ones. Other differences depend on platforms' policies: while anyone can sign up to Clickworker or ClixSense in just a couple of simple steps, Appen screens potential new micro-workers after they go through a full-fledged application process, and FouleFactory has closed its registrations, created a waiting list, and put in place a deregistration policy for inactive users. Platforms also differ in their degree of precision: some, such as Microworkers, give very detailed figures and update them regularly, while others, such as Amazon, just provide approximations.

Table 2 mirrors platforms' practice to communicate the size of their micro-workforce without breaking it down by country. The literature and platforms' narratives provide more insight. On Amazon Mechanical Turk, French workers represent only a small fraction of the user base (Difallah et al., 2018). On another international platform, Microworkers, France is not among the 10 most represented countries, and remains part of the group of those where less than 20% of total transactions take place (Hirth et al., 2011). On the German platform Clickworker, France and other European countries (excluding Germany) account for a total of 25% of the workforce¹⁰. No information is available for ClixSense and Appen, and French platforms present a variety of situations: while FouleFactory is only accessible to residents of France, Ferpection has users signing up from France and from the United Kingdom, the United States, Ireland and other countries¹¹.

¹⁰ See <https://www.clickworker.com/about-us/clickworker-crowd/> accessed on 3/12/2018

¹¹ See <http://help.ferpection.com/l/fr/les-utilisateurs/combien-y-a-t-il-d-utilisateurs-dans-la-communaute> accessed on 3/12/2018

Table 2 Number of people signed-up to micro-working sites, worldwide.

Platform	Registered users
Amazon Mechanical Turk	500,000
Microworkers	1,215,829
Clickworker	1,200,000
ClixSense	7,000,000
FouleFactory	50,000
Ferpection	50,000
Appen	1,000,000
Total	11,015,829

Source: websites of platforms included in analysis (accessed in August 2017 for Amazon Mechanical Turk; September 2018 for all other platforms).

3.3 Analyzing experimental data with the “capture-recapture” model

Since Stewart et al. (2015) and Difallah et al. (2018) challenged Amazon’s publicly available data by implementing a “capture-recapture” approach to Mechanical Turk, this method commonly used in bio-ecology and epidemiology has proven to be a valuable tool to estimate the number of micro-workers on a platform. Having posted a micro-task for a relatively long period of time, and allowing repeated participation, the technique consists in “capturing” participants a first time, identifying them as having already done the task, then seeing how many of them get “recaptured” a second time.

We were able to replicate this approach on the French leading micro-working platform FouleFactory. As part of our DiPLab study (see footnote 1), we distributed a questionnaire in the form of a paid task, administered in two steps over two months: the first phase obtained 505 responses and the second 492 responses¹². While purely logistical issues motivated this two-step collection, we saw it as an opportunity to make an estimate of the number of persons micro-working for FouleFactory. Indeed, in the second collection, we detected 89 returning respondents - thus enabling us to apply the capture-recapture model.

The capture-recapture model is based on two assumptions: first, the population is closed (that is, nobody entered or left the platform during the study) and second, all subjects have the same chances of being captured. The former was satisfied in our case, as FouleFactory’s closed-registration policy and the short time interval between the two waves limited the number of those entering or exiting the platform. The so-called Lincoln-Petersen formula applied to the data thus collected, with N

¹² Calculations after cleaning and filtering of data.

the total number to be estimated, n_1 the number of users captured in the first wave, n_2 the number in the second wave and m the number of recaptured participants, gives :

$$N = \frac{n_1 n_2}{m}$$

Based on the number of individuals observed, the results of the calculation give a population of 2792 microworkers (Table 3). This figure is only 5.6% of the platform's public declarations, but it is consistent with the domain-specific knowledge of its managers who, answering our questions about the ideal duration of our data collection, were confident that "over two months, [we] can mobilize up to 3000 people on this type of task".

However, this is likely an underestimation. The second assumption of the model (equal chance of capture) does not apply if users display an uneven level of involvement in micro-work on FouleFactory. When facing the same issue on Amazon Mechanical Turk, Difallah et al. (2018, p. 141) introduce a latent variable (a) that they interpret as propensity to participate: for example, it could reflect the fact that some people perform all the tasks available on the platform, regardless of their appeal, difficulty or remuneration, while others choose tasks more selectively. The initial formula is modified as follows¹³:

$$N^* = \frac{n_1 n_2}{m} \left(1 + \frac{\text{Var}(a)}{E(a)^2}\right) = N \left(1 + \frac{\text{Var}(a)}{E(a)^2}\right)$$

Difallah et al. (2018) estimated the distribution of a based on 28 observations over two years. In our case, their parameter setting¹⁴ would give $N^* = 34,166$ (Table 3), a value that still falls short of the public declarations of FouleFactory (50,000 as indicated above), but approximates it if we interpret it as an order of magnitude.

However, there is no guarantee that our population shares the same characteristics of the people who work on Mechanical Turk. We therefore propose an alternative estimate of N^* using a proxy for a , based on one of the questions asked in our survey, about the number of micro-tasks carried out in the previous month. The answers to this question, summarized in Figure 1, show a high diversity of participation - inconsistent with the assumption of equal chances of capture¹⁵. Using the above formula to calculate N^* , we obtain a population of 6531 micro-workers on FouleFactory¹⁶.

¹³ $N \approx N^*$ when the variance of a is close to the mean, but $N < N^*$ otherwise.

¹⁴ Difallah et al (2018, p. 142) consider that the propensity to participate follows a beta probability distribution, and that therefore the chance to recapture the same subject n times follows a beta-binomial law with parameters $\alpha = 0.29$ and $\beta = 20.9$. With these values, and a sample of S single observations, the population N^* is equal to $\frac{S}{1-f(0|n,\alpha,\beta)}$.

¹⁵ As our variable was originally a categorical (ordered) one, we have first transformed it into a numerical variable (taking the value of the minimum of the interval in each case: 0, 1, 3, 6, 10) in order to infer a distribution from it, thus enabling calculation of its mean and variance. Alternative transformations of this variable into a numerical one give qualitatively similar results.

¹⁶ Variations of this calculation, estimating distributions for each possible value in each interval, produce results of a similar order of magnitude (between 5,800 and 8,000).

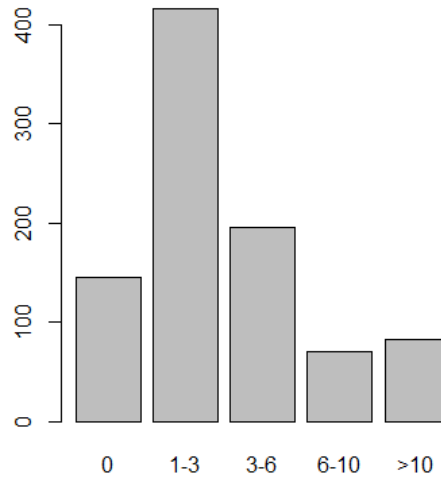


Fig. 1 Number of micro-tasks carried out on FouleFactory in the month preceding the survey, n = 908. Source: Authors’ elaboration with DiPLab survey data.

Table 3 summarizes these results. The plain capture-recapture model (Lincoln-Petersen formula) produces an underestimation and its correction (parameters of Difallah et al. (2018)) approximates the number of signed-up users reported by the platform (Table 2). The third estimate, correcting the Lincoln-Petersen figure with a proxy from our DiPLab survey, gives an intermediate figure.

Table 3 Estimate of the size of the FouleFactory population using the capture-recapture method (basic formula and corrections).

Method	Result	% of declared (50,000)
Lincoln-Petersen formula	2,792	5.6%
Correction with Difallah et al. (2018) parameters	34,166	68.3%
Correction with a proxy from DiPLab survey	6,531	13 %

Source: authors’ elaboration with DiPLab survey data.

The latter estimate is more difficult to interpret based on the information provided so far. We now need to find clues to make sense of it. Also, we need to extend our

analysis beyond the case of one single platform and go back to the full list of the micro-work intermediaries that we have identified as relevant for France (section 3.1).

3.4 Analyzing web traffic data to measure the audience of micro-work platforms

We submit that user panels to measure the audience of these platforms, a method borrowed from the study of media but never used in research on platform labor so far, can provide another perspective on the population of micro-workers. For France, we used the service SimilarWeb.com¹⁷, which makes it possible to estimate the number of unique monthly visitors to micro-working sites worldwide (Table 4, left column).

A potential difficulty that might arise with this type of sources is the double counting that might occur if, for example, the same person connected via two different devices or from two different IP addresses. Although we have no access to SimilarWeb's algorithm, the company ensures that this is controlled for in their calculations. In addition, our own DiPLab survey suggests that most people always connect from the same device and place, typically at home: therefore, this problem is likely negligible.

In addition, given that most platforms have separate web interfaces for client companies and for micro-workers, we can focus exclusively on the latter, namely on the pages dedicated to carrying out the tasks (for example, Fouleurs.com for FouleFactory). We must also exclude simple visitors (also including journalists or researchers like ourselves) from our analysis. To check for this, we look at the average duration of visits, which as per Table 4, is relatively long (central column). This allows us to exclude the presence of simple visitors or internet users who are just searching information, and to argue that these sections of the platforms are almost exclusively accessed by micro-workers. The only exception is Amazon Mechanical Turk, where information from SimilarWeb.com is not entirely reliable as the large number of observed unique visitors even exceeds the platform's own estimate of 500,000 (left column), and must therefore include surfers who are not themselves micro-workers. As discussed above, Mechanical Turk should be seen as a unique rather than a representative case.

Let us now comment especially the right column of Table 4, which provides estimates of the average number of French unique monthly visitors per platform. Starting from FouleFactory, we observe 6,958 unique micro-workers accessing it on average every month – which is very close to the figure of 6,531 that results from

¹⁷ The data used by SimilarWeb.com come from two main sources: partnerships that the site has established with Internet service providers and data that the site retrieves with an add-on on users' internet browsers – which provides users with metrics of website traffic in exchange for their browsing data. We have chosen SimilarWeb.com instead of the French market leader Médiamétrie, because the latter sends too few visits to our platforms of interest.

Table 4 Average number of unique visitors on micro-work platforms (left), average duration of visits (center) and average number of monthly unique visitors from France (right).

Platform	Average monthly unique visitors	Average visit duration	Average monthly unique French visitors
Worker.mturk.com (Amazon Mechanical Turk)	588,976	34:19	1,250
Microworkers	174,808	18:22	1,835
Workplace.click-worker.com (Clickworker)	242,579	7:40	14,700
ClixSense	1,083,000	11:05	20,250
Fouleurs.com (FouleFactory)	7,647	26:41	6,958
Ferpection	28,064	7:10	8,116
Appen	260,699	5:04	9,645
		Total	62,754

Source: Authors' elaboration based on information from SimilarWeb.com, accessed in September 2018. Figures are based on websites' audience tracking over the period July-September 2018.

our corrected capture-recapture calculation as detailed in subsection 3.3. This order of magnitude seems to correspond well to the size of the active population on this platform over a month (as could be measured in the summer of 2018). It should also be noted that almost all of the visits originate from France, in accordance with the platform's policy. The residual cases of people signing in from outside France essentially concern French micro-workers temporarily connecting from abroad — a trend confirmed by responses to our questionnaire. Concerning the other micro-work platforms offering access interfaces in French, ClixSense, Clickworker, Appen and Ferpection are more visited than FouleFactory. This is both due to their open registration policy and to their lifespan (at least in the case of ClixSense, created as early as 2009). Some also benefit from their role as gateways to other micro-work platforms: as noted above, ClixSense gives access to Figure Eight (formerly Crowdfunder, see footnote 7), while Clickworker provides access to UHRS (footnote 9).

To estimate the number of French micro-workers on Amazon Mechanical Turk, where as discussed, SimilarWeb provides only very limited help, we use instead data

from a survey conducted by the team of Panos Ipeirotis at New York University¹⁸. Extrapolating from the size of the group of French “Turkers” that they observed, i.e. 0.25% of the total, and assuming that their sample is representative, we estimate that the number of French residents operating on Mechanical Turk is around 1,250.

3.5 Combining methods to obtain a range of values

The calculation methods we have implemented so far provide differing estimates of the population of micro-workers in France: a high estimate is the number of signed-up users advertised by the platforms themselves (11,015,829 individuals worldwide), a low estimate is the plain capture-recapture model (2,792 individuals in a single platform, FouleFactory) and an intermediate estimate is the audience-based measure (62,754 individuals across platforms all over France).

Each of these approaches has limitations. At least in its basic, uncorrected version, the capture-recapture model underestimates the target population, while the sum of platforms’ public declarations overestimates it. The use of audience measurements brings a helpful alternative with intermediate results, but it is based on proprietary techniques whose actual implementation is difficult to audit for external researchers. At this stage, we have no means to choose one single metric from among those three, all the more so as they derive from very different measurement approaches. However, we can use these three results to cross-pollinate each other and help us derive more consistent measures.

Notice, first, that in the case of FouleFactory alone, the two corrections to the capture-recapture model that we have proposed to compensate for the downward bias of the basic Lincoln-Petersen formula (34,166 and 6,531), are close to, respectively, the publicly-released number of registered workers provided by the platform itself (50,000) and audience figures (6,958). These results implicitly validate all these metrics, although they are all lower (to different degrees) than the reference 50,000.

More importantly, audience figures enable an assessment of the number of registered micro-workers who are located in France. To compute their number, we need to extrapolate their frequency of connection to the platforms (calculated from Table 4) as a percentage of the total number of persons signed-up worldwide (Table 2)¹⁹. Table 5 shows the results of these calculations.

¹⁸ Data from the survey “Analyzing MTurk demographics” are available at https://github.com/ipeirotis/mturk_demographics and serve as the basis for the article by Difallah et al. (2018).

¹⁹ Given the national scope of FouleFactory’s activity, we consider all its registrations to be French. For Mechanical Turk we refer to the figures from P. Ipeirotis’s survey mentioned in footnote 18. For the other platforms, we take into account the number of visitors from France (Table 4, last column) divided by the total number of visitors (Table 4, first column), then apply the resulting percentage to the number of registrations (Table 2). For example, Ferpection has 8,116 unique visitors from France, out of a total of 28,064 visitors, or 28.92%; multiplying by its publicly declared number of registered users, equal to 50,000, we obtain 14,460.

Table 5 Estimated number of people registered on micro-work platforms and connecting from France.

Platform	Number French
Amazon Mechanical Turk	1,250
Microworkers	12,766
Clickworker	72,720
ClixSense	130,900
FouleFactory	50,000
Ferpection	14,460
Appen	37,000
Total	319,096

Source: Authors' elaboration.

On this basis, we can go a step further and attempt to generalize the capture-recapture model initially used just on one platform, by applying its ratios (Table 3, last column) to the total number of French micro-workers we have just estimated (319,096). We thus obtain a range, with a low estimate of 17,869 (5.6% of the declared total for France), a high estimate of 217,943 (with a “corrected” rate of 68.3%) and an intermediate estimate of 41,482 (with the alternative correction of 13%).

3.6 Accounting for users who micro-work on several platforms

We cannot rely on these three estimates until we have made an additional correction to take into account workers performing micro-tasks on several platforms. As discussed above, disregarding this type of multi-activity amounts to over-estimating the number of workers by counting the same persons multiple times. How many platforms do people use at the same time? To answer this question, let us turn to our DiPLab survey, in which participants were asked to indicate their use of different micro-work platforms from a list of options. 16.6% of them use at least two platforms from among the 7 included in our sample, which amounts to say that on average, each person uses 1.27 micro-tasking platforms.

Assuming our sample is representative, and extrapolating its behavior to the entire French micro-working population, amounts to applying this rate to the measures previously obtained (Table 6).

Little is known of micro-workers' multi-activity in other contexts but, comparing our survey to studies of other online practices such as the buying and selling of goods and services (where for example, Oxera (2015) reports an average of 2.2 platforms per user), it appears that our rate of multi-activity is low. This is somewhat

Table 6 Estimate of the population of micro-workers refined to account for use of multiple platforms.

Range	Measure	Result (reminder)	Result (minus multi-activity)
Low	Uncorrected capture-recapture	17,869	14,903
Intermediate	Audience measure	62,754	52,337
High	Number of declared users (France)	319,096	266,126

Source: Authors' elaboration.

surprising as the globally competitive context of micro-work, together with systematic excess supply of labor (subsection 2.2) would lead to expect higher rates. The DiPLab survey suggests that practical impediments play a role. Some respondents hint that entry costs of maintaining active accounts on several platforms are relatively high, both in terms of incompatibility of payment systems (for example, FouleFactory uses Mango Pay while Microworkers uses Paypal, Skrill and Payeroneer) and in terms of limited portability of “qualifications” (unpaid tests needed to access some types of micro-tasks, for example tasks that require knowledge of a specific language). Additionally, the English language used on many international platforms puts off some French micro-workers, even though some of their tasks are in French.

4 How to interpret results as indicators of of micro-workers' level of activity?

We now turn to the meaning of these different measures of the size of the French micro-working population. Building on the qualitative evidence collected via 92 semi-structured interviews conducted within the DiPLab project, we associate our figures to characterizations of users according to their level of activity. We distinguish a group of “very active” (subsection 4.1), a group of “regular” (subsection 4.2) and one of “occasional” (subsection 4.3) micro-workers.

4.1 Very active micro-workers

Our first, lower estimate of 14,903 micro-workers is based on the uncorrected capture-recapture model, first applied to FouleFactory, subsequently extrapolated to the total number of French registrations over the seven platforms of interest, and finally corrected to remove duplicates (subsection 3.3). As discussed, the Lincoln-Petersen formula on which this figure is based, is downward biased when the chances of capture (here, levels of engagement with the platform) are unequal, and we can add now that it tends to over-sample users with a high level of activity. Indeed, the number of FouleFactory users who were captured and then also re-captured is remarkably higher than we would have predicted (Stewart et al., 2015; Difallah et al., 2018).

These remarks suggest that this figure should not be simply dismissed as an under-estimate, but provides instead a rough guess of the number of what we might call “very active” micro-workers. While we do not have elements to characterize them in much detail, we can say at least that they connect to micro-working platforms at least weekly – as 90% respondents to our sample do. We also have examples from our qualitative fieldwork of what high activity may mean in practice. In addition to her primary occupation as a nurse, one respondent performs “simple and quick tasks” every day during her breaks. Another, expressing her “fear of missing out tasks”, even goes so far as leaving FouleFactory continuously open on her computer.

4.2 Regular micro-workers

The estimate we obtain through audience measurements from the seven micro-work platforms considered is much larger, with 52,337 users (average monthly number of unique visitors from France, calculated over three months, summed across platforms (section 3.4) and corrected downward to account for workers who switch across platforms (section 3.6)). Unlike the above, this method is more likely to over-sample users who connect at least monthly – indeed any more frequent activity won’t change the total, and conversely more sporadic connections are less likely to be counted at all. Thus, we can take this measure as an estimate of the number of micro-workers with an intermediate level of activity, and we label it “regular”. It likely involves using platforms for a secondary source of income.

For a more concrete sense of what it means in practice, let us resort again to our DiPLab interviews. An editorial manager in a communication agency has been micro-working for three years. She signed up during a period in which her “lifestyle required an additional income”, resulting in “a compulsive need to earn more money”. Since then, she has been routinely returning to micro-work, though not with the same intensity as before, and doing essentially questionnaire-based micro-tasks that she sees as an extension of her education. This case suggests that “regular” micro-workers often have a selective attitude towards tasks, and just won’t

perform any job that they see online. Notice, also, that users may move from very active to regular (and presumably, the other way round) at different times.

4.3 Occasional micro-workers

Our last estimate features 266,126 persons and corresponds to registered users as declared by platforms themselves, after removing non-French users and duplicates. Because these are just users who signed up with one or more micro-work platforms, in principle they may have any level of activity. But overall, people with low levels of activity will be much more numerous within this group. Thus, we can expect the great majority among them to be “occasional” micro-workers, who perform less than one task per month. Some of them may do tasks very rarely, and there might even be cases of users who never did any tasks at all. Using a corrected measure of the number of registered users is the only way to observe the long tail of this population that connects less frequently, is more difficult to reach through platform-based dedicated surveys (like the one we launched), and may escape web traffic panels.

Our interviews suggest that broadly speaking, two sub-populations can be distinguished within the “occasional”. One includes users who signed up out of curiosity and quickly abandoned micro-work, possibly because of the fluctuating availability and unsuitability of tasks – two points often emphasized by our study participants. They never de-registered, but their use remains sporadic. Another includes individuals who were “very active” or “regular” at some point in time, have currently reduced their involvement with these platforms, but may get back to them again in the future depending on their circumstances.

One respondent, holding a day-job as administrative secretary, illustrates this situation well. He signed up to FouleFactory in 2016, but “quickly dropped out”. A few months later, he received an email indicating that the platform would delete his account due to inactivity. He then decided to resume his micro-working activity in a relatively sustained way. He now claims he earns about 80 euros a month from micro-tasks, almost four times more than the average in our survey (Casilli et al., 2019). Another micro-worker, currently seeking employment, also alternates between periods of inactivity and of intensive performance. In her working life, she has always experienced alternating phases of unemployment and employment. She says she stops micro-work as soon as she finds a job because she “no longer has time”, while her phases of unemployment correspond to periods of intensive micro-work because she prefers to “do this rather than nothing”.

5 Discussion: strengths and limitations of present methods

Our three estimates respond to the need of assessing the size of a workforce that escapes the definitions and measurements of official statistics. Rather than opting for the relatively safe and simple option of focusing on a single platform, we look at one country as the basic unit at which any policies can be devised, but where multiple platforms operate simultaneously. We fully take into account the specificities of micro-work, notably diversity in levels of engagement, with many people practicing it as a secondary or infrequent activity; absence of geographical anchor, which implies co-existence of both national and international platforms; and, in an attempt to bridge the digital labor and the two-sided-markets literatures, multiple use of different platforms. Our choice of methods borrows from some of the most advanced literature on micro-work platforms (the capture-recapture model) and from neighboring fields (web traffic measures) to bring more precision to publicly available data from companies. Our mix of methods enables us to do comparisons and quality checks, to obtain more reliable results.

5.1 Methodological limitations

The methods we have deployed were a journey into uncharted territory, and required us to make some strong assumptions and simplifications. Lack of data restricted the set of platforms we could study. The capture-recapture method was applied to only one platform, FouleFactory, and we extrapolated its basic result (the uncorrected Lincoln-Petersen formula) to all others, assuming they would exhibit the same behavior. What's more, our survey included just two waves of capture and recapture, allowing only for approximate corrections of the downward bias resulting from the equal-capture-chances assumption. Regarding audience measures, constructed by private companies for business purposes, there was very limited scope for us to control their quality. Finally to assess the prevalence of micro-workers operating on several platforms, we had to assume our FouleFactory survey sample was representative of the French micro-working population.

For all these reasons, we believe that *our figures should be interpreted as orders of magnitude rather than exact amounts*: we will retain about 15,000 “very active”, 50,000 “regular” and 260,000 “occasional” micro-workers. Because these figures have been obtained through different methods, they cannot meaningfully be added; neither can they be construed as subsets of one another, the “very active” constituting an inner circle and the “occasional” an outer circle, with the “regular” in-between. These labels only indicate the levels of engagement that each of these estimates over-represents in comparison to the others, but we cannot exclude overlaps.

5.2 *Why our methods are nevertheless reliable*

Despite these limitations, these figures provide helpful indicators for research and policy. Moreover, the reliability of our methods lies in their consistency with existing evidence. The 15,000 can be seen as a sort of lower bound on the number of French micro-workers: there are at least as many people who perform this activity. At a very basic level, it is evidence that there are micro-workers in France at all: even if cost optimization may lead companies to search for cheaper online labor overseas as discussed in section 2.2, and even if many large data orders are perhaps more likely to transit through English-language international platforms, there is still one fraction of micro-work that is done in the country, and people who are available to do it. We have interpreted this estimate as inclusive of mostly very active people, which means the volume of economic activity they generate is roughly comparable to that of French Uber drivers, whose number was estimated at about 14,000 in an aforementioned IGAS report (section 3.2, Amar and Viossat (2016)), and of Deliveroo couriers, who are 9,300 couriers in France if we trust company figures (Deliveroo, 2018). If these more high-profile cases have attracted a lot of academic and policy attention, micro-work then deserves the same.

The 50,000 is an intermediate level, probably the most representative of the widespread trend to do platform work as a secondary or accessory activity, requiring regularity but not necessarily high frequency. The volume of economic activity that these workers generate is presumably small, but not negligible. Use of platforms for extra income rather than a main job is not unique to micro-work, and has already been observed in surveys of platform labor in general (Huws et al., 2017; Pesole et al., 2018; Urzì-Brancati et al., 2019; O'Farrel and Montagnier, 2019). Regarding micro-work specifically, this usage may be specific to Europe and North America, with providers in low-income and emerging countries choosing instead to micro-work full-time (or almost so) to meet growing demand by large digital companies (Schmidt, 2019). In countries such as France, the need for an extra source of income that the existence of these workers brings to light, calls for a better understanding of the conditions under which it emerges: what are the unmet needs that push people to seek additional resources through micro-work?

The 260,000 corresponds to the number of French users registered on one or more of the seven micro-tasking platforms considered. This might seem a surprisingly large number, especially for a type of platform work that has attracted limited attention from the general public so far. However, it is the same order of magnitude as the estimate of about 300,000 that might be derived from the 2018 Eurobarometer, and lower than the 900,000 of the European COLLEEM survey, as reported in subsection 2.1 (although these comparisons are only lightly indicative because definitions vary across these studies). Also, its full-time equivalent may be very small, because not all of these 260,000 do a lot of tasks²⁰. The real question is what this

²⁰ A back-of-the-envelope calculation suggests that, with micro-workers earning an average of 20 euros per month (a result from our DiPlab survey, see subsection 4.3), and a monthly net minimum wage set at around 1,230 euros in 2018, then 260,000 occasional micro-workers would correspond

count of 260,000 itself represents – before comparing it to data obtained with different methods, or converting it into categories whose fit with the specific features of platform economy may be questionable.

We contend that this very figure has strong implications for both theory and policy. As the freelancers studied by Graham and Anwar (2019), these micro-workers constitute a large pool of manpower that clients may tap into whenever their data wants become more pressing. To meet the (sometimes large) peaks in companies' demand, platforms make this human workforce available to them in amounts that are huge and often even excessive. Endemic to all forms of online labor, *this massive excess supply should attract policymakers' attention* because it shifts the power balance toward platforms (and clients), pushing remunerations down to the detriment of workers.

5.3 Concluding remarks

A final thought before concluding brings us back to labor statistics and the question of what it can learn from a study like ours. We have used unconventional methods, approximations and admittedly strong assumptions to come to results that are mere orders of magnitude. So far, such methodological bricolage was commonly found in research on hidden and difficult-to-reach populations – for example persons living with HIV or other stigmatized health conditions, sex workers or illegal drug users. In contrast, scholars of work and employment could realistically expect to rely on the clarity and precision of instruments such as Labor Force Surveys and Censuses.

These certainties no longer hold: in the data economy that thrives through platforms, one part of the world of work escapes established definitions and misses the gaze of official statistics, gradually sliding into invisibility. To push the metaphor a bit further, our object of study ends up being an analytical equivalent of a disenfranchised population. Scholars, then, have to lower their expectations of accuracy, and roll their sleeves up to devise creative ways of observing the less-and-less observable.

There are challenges ahead for labor research, which will have to open itself to a wider and more diverse range of methodological options. More importantly, the traditional bases for labor policies are profoundly affected. If counting workers becomes a challenging trial-and-error process, with potentially disputed results, then it becomes difficult to, for example, determine eligibility requirements for retirement, develop programs to support disadvantaged categories of workers, monitor the application of workplace safety programs, identify opportunities for local development, or flag cases of discrimination. Behind statistical measures, it is the more fundamental question of the place of labor in our society that needs to be addressed.

to about 4,200 full-time workers paid at minimum wage. These calculations are only suggestive because the distribution of monthly earnings from micro-work is highly skewed.

6 Acknowledgments

This paper is based on results of the “Digital Platform Labor” (DiPLab) project, co-funded by Maison des Sciences de l’Homme Paris-Saclay; Force Ouvrière, a workers’ union, as part of a grant from Institut de recherches économiques et sociales (IRES); and France Stratégie, a service of the French Prime Minister’s office. The platform FouleFactory/Wirk offered logistical support, and Inria provided complementary funding. For more information: <http://diplab.eu/>

The authors would like to thank current and former DiPLab team members, notably Maxime Besenval, Odile Chagny, Marion Coville, Touhfat Mouhtare, Lise Mounier, Manisha Venkat and Elinor Wahal, for their inputs to the project in general, and this paper in particular.

A preliminary version of this paper was posted on arXiv (n:1901.03889 [econ.GN]) and in French, as a working paper in the i3 series²¹. We thank Mark Graham, Mary Gray, Lilli Irani, Emmanuel Lazega, Vili Lehdonvirta, Antoine Naboulet, Etienne Ollion, Elise Penalva-Icher, Maurice Ronai and Alex Wood for useful and constructive feedback. Any remaining errors or inconsistencies are exclusively ours.

References

- Amar, N. and Viossat, L.-C. (2016). Les plateformes collaboratives, l’emploi et la protection sociale. Report IGAS n. 2015-121R, IGAS.
- Athey, S., Calvano, E., and Gans, J. (2018). The impact of consumer multi-homing on advertising markets and media competition. *Management Science*, 64(4):1574–1590.
- Berg, J., Furrer, M., Harmon, E., Rani, U., and Silberman, M. (2018). Digital labour platforms and the future of work: Towards decent work in the online world. ILO Report.
- Casilli, A. (2017). Digital labor studies go global: Toward a digital decolonial turn. *International Journal of Communication*, 11(0).
- Casilli, A., Tubaro, P., Le Ludec, C., Coville, M., Besenval, M., Mouhtare, T., and Wahal, E. (2019). Le micro-travail en France. Derrière l’automatisation, de nouvelles précarités au travail ? Report of the Digital Platform Labor (DiPLab) project.
- Deliveroo (2018). Deliveroo lance marketplace+. *Deliveroo France Newsroom*.
- Difallah, D., Filatova, E., and Ipeirotis, P. (2018). Demographics and dynamics of Mechanical Turk workers. In *Proceedings of WSDM 2018: the Eleventh ACM International Conference on Web Search and Data Mining*, pages 135–143. ACM.

²¹ URL: <http://i3.cnrs.fr/wp-content/uploads/2019/02/WPi3-19-SES-02-LeLudec-Tubaro-Casilli.pdf>

- Eichhorst, W., Hinte, W., Rinne, U., and Tobsch, V. (2017). How big is the gig? assessing the preliminary evidence on the effects of digitalization on the labor market. *management revue - Socio-Economic Studies*, 28(3):298–318.
- Forde, C., Stuart, M., Joyce, S., Oliver, L., Valizade, D., Alberti, G., Hardy, K., Trappmann, V., Umney, C., and Carson, C. (2017). The social protection of workers in the collaborative economy. Report for the committee for employment and social affairs, European Parliament.
- Fort, K., Adda, G., and Cohen, K. (2011). Last words: Amazon Mechanical Turk: Gold mine or coal mine? *Computational Linguistics*, 37(2):413–420.
- Frenken, K. (2018). Unpacking the gig economy – are we asking the right questions? Presentation at the Reshaping Work 2018 Conference, Amsterdam, 25-26 October.
- Fundación Felipe González and University of Hertfordshire (2019). Huella Digital: La plataformización del trabajo en Europa. Ficha informativa de España (report of the Digital Footprint comparative project).
- Graham, M. and Anwar, M. (2019). The global gig economy: Towards a planetary labour market? *First Monday*, 24(4).
- Graham, M., Hjorth, I., and Lehdonvirta, V. (2017). Digital labour and development: impacts of global digital labour platforms and the gig economy on worker livelihoods. *Transfer: European Review of Labour and Research*, 23(2):135–162.
- Gray, M. and Suri, S. (2017). The humans working behind the AI curtain. *Harvard Business Review*, (9):2–5.
- Hirth, M., Hoßfeld, T., and Tran-Gia, P. (2011). Anatomy of a crowdsourcing platform - using the example of Microworkers.com. In *2011 Fifth International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing*, pages 322–329.
- Huws, U., Spencer, N., Syrdal, D., and Holts, K. (2017). Work in the european gig economy: Research results from the UK, Sweden, Germany, Austria, The Netherlands, Switzerland and Italy. Commissioned report, Foundation for European Progressive Studies.
- Irani, L. (2016). The labor that makes AI "magic". AINow Summit, White House and NYU.
- Kässi, O. and Lehdonvirta, V. (2018). Online Labour Index: Measuring the online gig economy for policy and research. *Technological Forecasting and Social Change*, 137:241–248.
- Katz, L. and Krueger, A. (2019). The rise and nature of alternative work arrangements in the united states, 1995–2015. *ILR Review*, 72(2):382–416.
- Keith, M., Tay, L., and Harms, P. (2017). Systems perspective of Amazon Mechanical Turk for organizational research: review and recommendations. *Frontiers in Psychology*, (8):1359.
- Kuek, S., Paradi-Guilford, C., Fayomi, T., Imaizumi, S., and Ipeirotis, P. (2015). The global opportunity in online outsourcing. Report, World Bank.
- Mayer-Schönberger, V. and Cukier, K. (2013). *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. Houghton Mifflin Harcourt.

- Mrass, V. and Peters, C. (2017). Crowdworking-Plattformen in Deutschland. Leimeister, J.M. (ed.), Working paper series, nr. 16. Kassel, Germany.
- O'Farrel, R. and Montagnier, P. (2019). Measuring platform mediated workers. OECD Digital Economy Papers, n. 282.
- Ojanperä, S., Graham, M., Straumann, R., De Sabbata, S., and Zook, M. (2017). Engagement in the knowledge economy: Regional patterns of content creation with a focus on sub-saharan africa. *Information Technologies and International Development*, (13):33–51.
- Oxera (2015). Benefits of online platforms. Report prepared for google.
- Pesole, A., Urzì-Brancati, C., Fernández-Macías, E., and González-Vázquez, I. (2018). Platform workers in Europe: Evidence from the COLLEEM survey. Report EUR 29275 EN, Publications Office of the European Union, JRC112157.
- Robinson, J., Rosenzweig, C., Moss, A., and Litman, L. (2019). Tapped out or barely tapped? recommendations for how to harness the vast and largely unused potential of the mechanical turk participant pool. PsyArXiv preprint, doi: [org/10.31234/osf.io/jq589](https://doi.org/10.31234/osf.io/jq589).
- Schmidt, F. (2019). Crowdproduktion von Trainingsdaten: Zur Rolle von Online-Arbeit beim Trainieren autonomer Fahrzeuge. Report, Hans-Böckler-Stiftung.
- Stewart, N., Ungemach, C., Harris, A., Bartels, D., Newell, B., Paolacci, G., and Chandler, J. (2015). The average laboratory samples a population of 7,300 Amazon Mechanical Turk workers. *Judgment and Decision Making*, 10(5):479–491.
- ter Weel, B., van der Werff, S., Bennaars, H., Scholte, R., Fijnje, J., Westerveld, M., and Mertens, T. (2018). De opkomst en groei van de kluseconomie in Nederland. SEO Report n. 2018-30.
- Tubaro, P. and Casilli, A. (2019). Micro-work, artificial intelligence and the automotive industry. *Journal of Industrial and Business Economics*, 46(3):333–345.
- Urzì-Brancati, C., Pesole, A., and Fernández-Macías, E. (2019). Digital labour platforms in Europe: Numbers, profiles, and employment status of platform workers. Report EUR 29810 EN, Publications Office of the European Union, JRC117330.

Who are the micro-workers behind AI?

Paola Tubaro

Abstract

Micro-workers undertake menial online tasks to generate and annotate data for machine learning, to correct and verify the outputs of artificial intelligence systems, and to replace algorithms when they fail. This chapter examines what is micro-work and introduces some of the real people who carry out micro-tasks.

Key words: Micro-work, platform economy, digital labor, artificial intelligence, working conditions.

1 Introduction

Real men and women contribute to producing artificial intelligence (AI) solutions through “micro-work” (Irani, 2016; Schmidt, 2019; Tubaro and Casilli, 2019). Data-related jobs are fragmented into myriad small tasks that can be performed remotely online, and specialized digital platforms such as Amazon Mechanical Turk, Microworkers and Clickworker allocate them to anonymous providers who work behind their computer or smartphone screens. Such tasks may consist in labeling images, for example by drawing a box around, and/or adding a tag to, each car in a picture of street traffic. Other tasks are textual, such as transcription of license plates and shopping receipts, or translation of short sentences. Others are about content moderation, for example flagging derogatory words in a web page (Roberts, 2019). Quick and repetitive, these tasks are typically paid on a piecework basis, with rates as low as a few cents. Their silent but essential contribution to the functioning of smart equipment, computer vision systems and voice assistants, is what Gray and Suri (2019, 2017) call the “paradox of automation’s last mile”: smart technologies

Laboratoire de Recherche en Informatique, Centre National de la Recherche Scientifique, Paris Saclay, e-mail: paola.tubaro@lri.fr

incessantly create a need for so-called “human-in-the-loop” contributions, so that the final outcomes result from a mix of machinery and people.

Micro-work is an instance of digital platform labour, whereby workers are not employees but sub-contractors usually paid by piece-rate, subject to algorithmic control performed by the platform (Wood et al., 2019). Because it is mostly performed online, it differs from geographically sticky (Graham et al., 2017) “gig” work such as urban transportation through Uber and food delivery through Deliveroo. Neither is it the same as *online freelancing*, which is also performed online but concerns more creative work such as design and software development, involves qualified professionals, and entrusts them with full projects rather than single tasks (Schmidt, 2017). Micro-work pushes to the extreme the “datafication” processes that underlie all these activities.

Whose work is it? Who are the people behind today’s AI, what motivates them to engage in this activity, and what are their working conditions? In what follows, I show that they are often well-educated and professionally active but unable to earn enough, sometimes owing to care duties. The unregulated, unprotected setting of digital platform work is little suited to improve their chances, especially for the most vulnerable among them. Yet, they are indispensable to AI production. They prepare data for the machine learning algorithms that sustain today’s AI solutions, check the outputs of AI systems for accuracy and correct them if necessary, and even impersonate AI when algorithms fail. Their contribution is needed at all stages of AI development and commercialization and therefore, it is not a temporary stopgap measure but a systemic input. Thus, credible commitment to responsible, ethical AI in our economy and society requires considering micro-working conditions.

To support these claims, I combine insights from the nascent international literature on micro-work with empirical results of a mixed-methods research on micro-work in France, DiPLab (“Digital Platform Labor”). The study comprised a questionnaire distributed through a local micro-work platform, Foule Factory¹, and in-depth interviews with micro-workers, platform owners, business clients and other stakeholders (Casilli et al., 2019). As a highly industrialized country and a pioneer in information technologies, France is exemplary of more general trends, while language and some other specificities allow comparing and contrasting it to other known cases studied in the literature which, so far, has privileged English-speaking settings.

2 Who are the micro-workers?

Let us first discover the identity of the people who engage in micro-work. Where are they located? Are they men or women, young or old, more or less educated? What else do they do in life, and why do they micro-work?

¹ “Foule” means “crowd” in French.

2.1 *Low income vs industrial countries*

The first attempts to describe the micro-working population date back to 2010, just a few years after the international launch of Amazon's pioneer platform Mechanical Turk (Ipeirotis, 2010; Ross et al., 2010). Its demographics were less skewed at the time, with about 50% US and 40% India residents, while these percentages are now 75% US and about 16% India, the rest consisting of negligible proportions from third countries (Difallah et al., 2018). However, this does not mean that there are no micro-workers elsewhere. Hirth et al. (2011) analyzed the platform Microworkers and found it much more geographically diverse than Amazon Mechanical Turk, with 60% of the population coming from Bangladesh, India and the US together. More recently, Posch et al. (2018) surveyed the platform Figure Eight (previously called Crowdflower and now part of Appen, an international giant in data services for machine learning and artificial intelligence) to find, again, significant geographical diversity. A study commissioned by the European Parliament found European micro-workers on four international platforms (Mechanical Turk, Clickworker, Figure Eight and Microworkers) (Forde et al., 2017).

The first conclusion, then, is that micro-work is a global phenomenon, although different platforms may have different country compositions. For this reason, it is important not to limit our analysis to just international platforms, but to take into account local ones too: in the setting of this study, France, one of them is Foule Factory, a French-language platform that serves the needs of local clients through local workers.

2.2 *Not only millennials*

In an early report on worldwide online labor (including both qualified freelancing and micro-work), Kuek et al. (2015) show that most workers are "millennials" (born in the 1980s and 1990s), especially in emerging countries such as India. However, later studies suggest that the youngest among them, including students, are not the majority. If 60% of Forde et al. (2017)'s respondents (across four platforms and throughout Europe, see sub-section 2.1) were below 40 years of age, a long-term demographic survey of Amazon Mechanical Turk finds that 65.5% of the sample are 25-44 years old (Difallah et al., 2018)².

The DiPlab study comes to similar results: all workers are 18 years or older, due to the conditions imposed by the Foule Factory platform that exclude under-age users, and only one person in the sample is over 75. Most respondents (63.4%) are 25-44 years old – in line with the just-cited result of Difallah et al. (2018). This age range is clearly over-represented relative to the general French population, as can be observed by comparing this result to figures by INSEE, the National Statistical

² For comparison purposes, I have independently re-analyzed the data used in this study, available online at https://github.com/ipeirotis/mturk_demographics

Institute, which indicate that this age range amounts to only 35% of the general French population of age 18-74³.

	DiPLab	France
18-24	12.1	11.7
25-34	33.4	17
35-44	30	18
45-54	15.2	19.5
55-64	6.3	18.2
65-74	3	15.6

Fig. 1 Age distribution of DiPLab survey participants, compared to the general French population (in percentage). Interpretation: 33.4% of DiPLab respondents are 25-34 years old. Source: adapted from Casilli et al. (2019).

Micro-work, then, is not a student job but an activity for people whose age is usually associated to having completed even long studies, having a job and a (potentially young) family. Let us then look at other variables to better understand their situation.

2.3 *Not only men*

Are micro-workers more women or men? On Amazon Mechanical Turk, Difallah et al. (2018) observe about equal participation of men and women, although percentages differ greatly across countries : in particular, women are about 55% of US “Turkers”, but are much less numerous than men in India. Throughout Europe, Forde et al. (2017) count 60% men. A recent ILO research report that surveyed five international platforms (Mechanical Turk, Clickworker, Figure Eight, Microworkers and Prolific Academic) finds a more extreme gender imbalance, with only one third of workers in industrial countries being women, and barely one fifth in low-income and emerging countries (Berg et al., 2018).

Yet women are slightly over-represented in the DiPLab survey of French micro-workers, comprising 56.1% of the Foule Factory population – close to the percentage of American women active on Mechanical Turk reported above. Ferpection, another French platform and one of the very few that publicly provide figures about their workforce, says it has 57% women.

³ INSEE, Demographic state of France 2018. URL: <https://www.insee.fr/fr/statistiques/1892086?sommaire=1912926>.

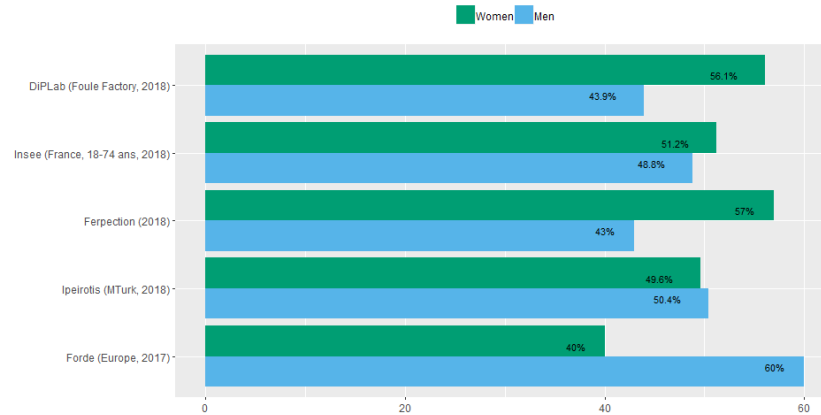


Fig. 2 Gender distribution of DiPLab survey participants (platform Foule Factory, 2018), compared to: 1/ the general French population aged 18-74, observed by INSEE (2018); 2/ proportions displayed by the French platform Ferpection in 2018; 3/ data collected by the team of Difallah et al. (2018) on Amazon Mechanical Turk across countries; 4/proportions found in the survey of Forde et al. (2017). Interpretation: 56.1% of participants in the 2018 DiPLab survey are women. Source: adapted from Casilli et al. (2019).

It is not true, then, that micro-work is a predominantly male activity. Rather, women and men choose different platforms. There is a majority of women on Ferpection and Foule Factory, two local, French-language platforms; but among the Mechanical Turk workers from France surveyed by Difallah et al. (2018), only about 25% are women. This depends largely on platforms' history and policies: Mechanical Turk restricted access of non-US residents around 2013, so its French users are early adopters, more likely to be men as with most digital technologies (Borgonovi et al., 2018). Conversely, sign-ups to Foule Factory boomed after a 2015 TV report that presented it as an earning opportunity for all. In passing, it is for this reason that it is important not to limit analysis to the most high-profile international platforms, as most of the literature has done to date, missing important parts of the population.

Perhaps more importantly, country of residence changes the role of gender with respect to micro-work. Why do women micro-work – when they do? Ipeirotis (2010) consider that more male workers from India participate on Amazon Mechanical Turk because it is their primary source of income, while more women from the US participate because they consider Mechanical Turk a secondary source of income. Is this a more general difference between emerging and industrial countries? Using the case of France to provide a clue, what follows looks at the extent to which DiPLab respondents use micro-work as a complementary activity, in relation to their family and employment status.

2.4 Highly educated workers, in work, with family

Kuek et al. (2015) was one of the first studies to show that online workers, including micro-workers, are more likely to hold a University degree than the general population – a finding that almost all subsequent literature confirms. DiPLab is not an exception: 66.5% of the surveyed micro-workers have a bachelor's degree or higher, in comparison to just 27.8% of the general French population. It is certainly ironic that highly qualified people do tasks such as tagging objects in images or transcribing sentences.

As expected (sub-section 2.3), micro-work is more often practiced as a complement to a full- or part-time job, especially in industrial countries. Only 52% of the workers surveyed worldwide by Berg et al. (2018) use micro-work as a secondary source of income, against 68% of Forde et al. (2017)'s European micro-workers and just over 60% of French DiPLab respondents (which, for just over half of them, is a clerical occupation). Of the 40% French DiPLab micro-workers who have a long-term employment contract, more than two thirds work full-time. As in the general population, women are more numerous to work part-time (32%, against 11% of men). Women are more numerous (about 60%) than men (slightly less than half) to depend at least partly on their spouse or partner for their income.

Yet even in industrial countries, micro-work is sometimes practiced as a remedy to inactivity. 36% of the French micro-workers surveyed by DiPLab are inactive, that is, they have no (main or secondary) professional activity aside from their micro-tasks, but only 15% of them are retired. A small number (mostly men) are students and the others are in diverse situations, but always with a majority of women. For these people, micro-work is an essential source of income.

Regarding family status, DiPLab data show that a larger number of women have children compared to men (55% vs 39%). Women also devote 6 to 12 hours per week more than men to domestic chores.

Such imbalances help to answer the question of why women micro-work (sub-section 2.3): mostly in working age but more numerous to work part-time (or to be inactive) than men, more frequently depending on their spouse for income, more often responsible for family care, and spending more time at home, they use all their free time and all their breaks between work and home activities to do online micro-tasks. Micro-work provides some welcome additional earnings, but also adds to the burden of formal employment and care work – eating into their leisure time.

2.5 A lower-income population that needs extra earnings

With these situations, are micro-workers poorer than the general population? To answer this question, reference can be made to the French Observatory of Inequalities which defines the lower classes as the 30% of the population with the lowest incomes. Among the surveyed DiPLab micro-workers, the lower classes are largely over-represented, with 51% meeting this criterion.

An alternative criterion is the poverty threshold, computed as 50% of the country’s median income⁴. 22% of the micro-workers under study live under this threshold, compared with only 8% of the general French population.

Against this background, online micro-tasks are an attempt to cope. The DiPlab survey asked respondents three reasons why they micro-work. The great majority say they need money – either as their first, second or third answer (Figure 3). In interviews, some said they actively look for better-paid tasks and feel frustrated if these turn out to be less rewarding than expected.

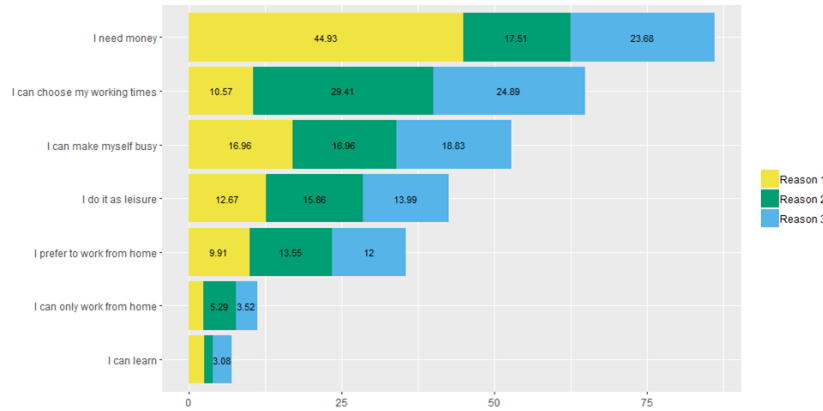


Fig. 3 Reasons for micro-work. Each DiPLab survey participant was invited to choose three of them, and rank them in order of importance. Interpretation: 44.93% choose “I need money” as the first reason why they micro-work. Source: adapted from Casilli et al. (2019).

If micro-work is a necessary – albeit often complementary – income source, does it provide enough for workers to support themselves? Let us now look at their earnings from this activity.

2.6 How do micro-workers fare?

Earnings from micro-work are low. On Amazon Mechanical Turk, the median hourly remuneration is only about US\$ 2 (Hara et al., 2018). The figures of Berg et al. (2018), calculated across five different international platforms, are similar (US\$ 2.16/hour). One reason why remunerations are so low is that a lot of time is spent in unpaid activities, such as searching from the list of available tasks and passing qualification tests (for example, language proficiency tests) to have access

⁴ This level is currently estimated at 855 euros per month for a single person. Source: INSEE and French Observatory of Inequalities.

to a wider range of tasks. Hourly rates are higher when only paid work is taken into account.

But hourly rates do not account for variations in the number of hours worked – which are huge because micro-work is often a complement rather than a main earning activity (sub-section EducationWork). Monthly, rather than hourly, earnings are better indicators in this respect: DiPLab respondents earn 22 euros a month on average, although the distribution is very skewed with a small number of people who make up to 2,000 euros a month, sometimes through multi-activity over different platforms (as do over 80% of DiPLab respondents).

In sum, people in different life situations micro-work. They live both in industrial and in low-income countries, they are both men and women, both with and without a primary professional activity beyond micro-tasking. What is common to many of them is a high educational level, somewhat at odds with the typically repetitive nature of micro-tasks. It is also the need for an extra source of income, either because they experience financial difficulties or because care duties prevent them from seeking better-paid employment opportunities. However, most people earn very little from micro-work, and their engagement is highly variable. Why? To answer this question, it is now time to look more closely at what micro-workers do.

3 What do micro-workers do?

Let us first review the activities that micro-workers do, to better understand the challenges that they face, before looking at the industrial needs they help to meet.

3.1 What are micro-working tasks?

Among other things, the DiPLab survey asked micro-workers to briefly describe the last task they had performed online. Although a single question at a specific point in time cannot provide an exhaustive description of all possible tasks, it can give a concrete idea of what the daily activities of a micro-worker consist in. Let us first review the most frequently mentioned answers, before trying to make sense of them.

A common task consists in responding to online surveys, to the satisfaction of many micro-workers who find them clearly presented and relatively well-paid, although others dislike their length. Another task reported by one in five DiPLab respondents involve voice: reading aloud a few short sentences in French and audio-recording them. Variants of this task include requests to record five different ways to, say, ask what the weather is like. Other recurring tasks are about recognition of objects in images: one of them asked micro-workers to draw bounding boxes around vegetables (tomatoes, carrots, etc.) in pictures of salads. There are also document management tasks such as CV anonymization, reported by one fourth of respon-

dents. According to Claire⁵, this task was about “removing all distinctive marks that could be discriminatory”, but it also required tagging the spaces in the document where names, birth dates, addresses and other personal or sensitive information were placed originally. Other examples of document management tasks involve classifying objects such as DVD titles, sorting photos, or recognizing whether commercial products found in different online marketplaces are actually the same.

It is useful to extend our gaze beyond the questionnaire, to consider additional types of tasks that did not surface there, but were mentioned in in-depth interviews. One is about moderating videos, text and images circulating on the web, in order to remove violence, pornography, drugs, and other inappropriate material. Aline said that, after the attacks of 2015-16, she checked “monstrous” terrorist content for several weeks, full-time, as clients “were panicking” for fear that gore images would flood their websites.

Another task that was mentioned in interviews, consists in checking the accuracy of speech recognition by voice assistants. Juliette spent several months comparing automated transcriptions and audio recordings of users’ conversations with their virtual assistant, check them for accuracy, and if needed, rectify them. Other interviewees checked the accuracy of automated transcriptions of scanned documents such as commercial receipts and invoices, adding corrections as appropriate.

Other interviews paint a more sinister picture. Kader, an entrepreneur, derided his competitors who allegedly made Madagascar “the French AI leader”. Why? Because, he argued, they claim they do AI while in reality, they just recruit cheap micro-workers in low-income, French-speaking African countries to do the work manually. Some micro-workers steal robots’ jobs, so to speak – for example when they have to infer clients’ tastes from social media and recommend some products to them.

3.2 Why tasks are challenging after all

However sketchy, this list of tasks helps us understand why micro-work is not necessarily straightforward – and why some workers engage with it very little. Routine tasks such as recording one’s voice or looking for tomatoes in pictures are relatively light to perform. Although they are often perceived as “silly”, as Anne said, some workers are just happy with the little pay they get for them, given the limited effort required. However, some tasks require specific personal characteristics or skills that not everyone has: for example, audio-recording tasks require native speakers.

More to the point, some micro-workers find these tasks perplexing, partly because requesters do not always clearly explain what their goal is. Diane could not make sense of why she was told to “draw a circle around a tomato” because “everyone knows what a tomato is, I hope”, and insisted that “I don’t know why”. Lack of understanding of the purposes of one’s own activity is known to entail psycho-

⁵ Interviewees’ names have been changed for confidentiality.

social risks, against which micro-workers – who are formally not employees but independent contractors – are not protected.

Other psycho-social risks arise with content moderation tasks. Exposure to violence and other sensitive content can be distressing (Roberts, 2019), although our above-cited interviewee assures that she has found ways not to be personally affected. She admitted that she manages to resist because of fear that she might be given fewer tasks if she refused –while she badly needs the money.

Handling other people’s personal data is also a source of psycho-social risks. Juliette, who listened to users’ conversations with their virtual assistants (sub-section 3.1), felt uncomfortable at the thought that she might be unwittingly violating their privacy. She wondered whether they realized that human beings would listen to their “sexual delusions”, as she put it.

Another issue is isolation: micro-workers do tasks at home behind their screens and are rarely in contact with their peers, except a minority who join dedicated forums (less than 60% according to Berg et al. (2018)). In many cases, they also find it difficult to get in touch with platforms and clients. Another source of stress is fierce competition for tasks, due to high volatility of clients’ demand: several interviewees said they keep their micro-working platforms always open or have set up alerts in order not to miss tasks when they are posted. This means unusual working times, including at night, for workers in low-income countries in Asia and Africa who do tasks for requesters mostly located in Europe and North America.

Against this background, it is easier to see why not all micro-workers engage fully with the activity, why some of them work (and earn) much more than others, and why micro-work is not a solution to the financial problems of all those who practice it. That said, what purposes does their activity serve?

4 How microwork feeds AI

It is now time to look at the other side of the market – the clients who post tasks on micro-working platforms and collect the outputs of workers’ effort. What do they look for, and why do they ask people to – to re-use the same example as above – draw circles around tomatoes in pictures of salads?

Leaving aside surveys which serve a range of diverse purposes (from marketing to scientific research), all the other tasks mentioned above have linkages to AI production, although this is not always immediately clear to workers.

Today’s AI production relies heavily on machine learning, a technique in-between computer science and statistics that “teaches” computers to find solutions from data, without each step having to be explicitly programmed (Alpaydin, 2014, 2016). “Learning” means that the quality of outputs can get progressively better over time, depending not only on the algorithm but also on the data given to it. Put differently, data are a necessary input just as algorithms. Micro-work prepares these data and supports the algorithms that handle it, as I show in what follows.

4.1 AI preparation

A first function of micro-work is what may be called “AI preparation” (Casilli et al., 2019) and comprises both generation of new data for machine learning, and “annotation” of already-existing data for better quality. For an example of data generation, take the above-mentioned task to read aloud short sentences and audio-recording them (sub-section 3.1). The requesters of such tasks are producers of virtual assistants who need to assemble large sets of diverse examples of voices, all pronouncing the same words, so that its vocal recognition algorithm can learn that they all mean the same despite differences in regional accents, tone, background noise, etc. It also needs examples of different ways to ask the same thing – such as, what the weather is like.

In turn, data annotation is needed for “supervised” machine learning algorithms, the best performing and most widely used ones so far. For example, to teach a computer to distinguish tomatoes from carrots, many images of these vegetables will not be enough, because the machine will not know what is what. Thus, the images should include tags that indicate, for each of them, whether it shows a tomato or a carrot, and ideally indications of where in the image the tomato or carrot is. After exposure to many images annotated in this way (so-called “training” datasets), the algorithm will be able to classify new, untagged images (“test” dataset) and determine whether they represent tomatoes, carrots or both. This is why micro-workers like Diane were asked to draw circles around tomatoes (sub-section 3.2). Their input will ultimately serve to build, for example, smart nutritional advisers that count calories based on the contents of a plate.

Many of the tasks described in sub-section 3.1 are for data annotation too. CV anonymization, with addition of tags to identify the location of sensitive information, was arguably meant to help some recruitment algorithm to understand the structure of a CV. Sorting documents such as videos or DVDs performs the same function, and so does content moderation – aiming to teach a computer to recognize violent, pornographic or otherwise sensitive content.

AI preparation can be partly automated, to the extent that algorithms can “learn” the data structure or its annotations from a first set of data prepared by micro-workers. However, there will always be a need for human input whenever new data are to be prepared. For example, selling a virtual assistant in a new country market will require a set of examples in that country’s language, dialect(s) and conditions.

4.2 AI verification

If AI preparation occurs at input level, to generate or enrich training datasets, another usage of micro-work is at output level, to review the results. This is AI verification. For example, checking the accuracy of speech recognition by virtual assistants is an example of verification, and so are checks of the quality of automated transcriptions of scanned invoices (sub-section 3.1). AI verification is more widely

applicable than preparation. First, it is needed with any type of machine learning algorithms, not only with the supervised variant. Second, it is not only required when AI solutions are being developed, but also, and more importantly, when they are already in use: for example, every time a company commercializes a virtual assistant in a new country, it needs new, local workers to check speech recognition results in the languages and dialects of that country.

4.3 AI impersonation

AI impersonation occurs when micro-work does not support the processes of data production or algorithmic quality assurance, but replaces them when they fail. This happens when humans outperform computers, either in terms of efficacy or cost. The very idea that prompted Amazon to launch Mechanical Turk in 2006, was to integrate humans directly into software programming when they were more efficient than automating. According to (Irani, 2015, p. 225), this platform was “born out of the failures of artificial intelligence to meet the needs of internet companies”. AI impersonation may seem unfair or even unlawful when use of low-paid humans instead of algorithms is undisclosed – as in the case of the AI start-ups that outsource work to providers in low-income countries rather than developing real AI, as one interviewee Kader accused them to do (sub-section 3.1). This is sometimes just a quick-and-dirty remedy to pressure from investors, but it is sometimes due to genuine reasons: “faking” AI may be a way to understand the production process of some service in the first place, before actually developing a suitable algorithm to automate it.

These considerations hint that the need for micro-work is unlikely to be a temporary one. As more industrial sectors integrate AI-based solutions, the need for micro-workers to prepare datasets for AI and to perform quality checks on algorithmic outputs will remain high. As long as wealth and income disparities across countries maintain repositories of cheap human labor in the developing world, AI impersonation will remain a cost-effective option.

5 Conclusions

Micro-work offers an opportunity for extra earnings to women with childcare duties who do not have a full-time job and do tasks primarily from home. It also attracts the inactive and more generally, lower-income (albeit highly-educated) persons. However, these extra earnings are often tiny or irregular, largely due to highly fluctuating demand from companies. The extent to which micro-work offers an opportunity for personal and/or professional development is doubtful, to the extent that tasks are mostly unqualified, that some micro-workers cannot give meaning to them, and that incentives to engage more intensely in it are scarce due to competition for tasks,

uneven working hours, and social isolation. While the study presented here focuses on one single country, France, its results are likely transferable to other settings.

And yet, micro-work is an indispensable ingredient of AI development. Consequently, current reflections about ethical AI should fully integrate the issue of labor standards in the backstage of automation. We should not look at AI as a pure product of techno-scientific progress: it also depends on economic cost considerations that lead companies to add humans to robots (or even to replace robots with humans) whenever convenient, and on socio-economic inequalities inherited from the past, whereby micro-taskers from low-income countries sometimes appear as a preferable alternative to highly-expensive engineers in the Western world. The consequences of AI on society are not independent of these hidden providers: for example, personal data leakages are all the more serious as independent sub-contractors potentially located anywhere have access to them, as recent scandals around voice-activated assistants demonstrate (Day et al., 2019).

Before even thinking of regulating this new labour force and strengthening its almost-inexistent social protection, it is necessary to at least recognize explicitly the contribution of workers to technological innovation, and the important role of this new form of work in the production and deployment of AI solutions.

6 Acknowledgements

This paper presents results of the “Digital Platform Labor” (DiPLab) project, co-funded by Maison des Sciences de l’Homme Paris-Saclay; Force Ouvrière, a workers’ union, as part of a grant from Institut de recherches économiques et sociales (IRES); and France Stratégie, a service of the French Prime Minister’s office. The platform FouleFactory/Wirk offered logistical support, and Inria provided complementary funding. For more information: <http://diplab.eu/>

I would like to thank current and former DiPLab team members, notably Maxime Besenval, Antonio A. Casilli, Odile Chagny, Marion Coville, Clément Le Ludec, Touhfat Mouhtare, Lise Mounier, Manisha Venkat and Elinor Wahal.

A preliminary version was presented at the “Automating the Crowd” workshop that took place at the Alan Turing Institute, London, on 22 January 2019, and I am grateful to organizers and participants for helpful feedback. Any remaining errors or inconsistencies are exclusively mine.

References

- Alpaydin, E. (2014). *Introduction to Machine Learning (3rd edition)*. MIT Press.
Alpaydin, E. (2016). *Machine Learning: the New AI*. MIT Press.

- Berg, J., Furrer, M., Harmon, E., Rani, U., and Silberman, M. (2018). Digital labour platforms and the future of work: Towards decent work in the online world. ILO Report.
- Borgonovi, F., Centurelli, R., Dernis, H., Grundke, R., Horvát, P., Jamet, S., Keese, M., Liebender, A., Marcolin, L., Rosenfeld, D., and Squicciarini, M. (2018). Bridging the digital gender divide: Include, upskill, innnovate. OECD report.
- Casilli, A., Tubaro, P., Le Ludec, C., Coville, M., Besenval, M., Mouhtare, T., and Wahal, E. (2019). Le micro-travail en France. Derrière l'automatisation, de nouvelles précarités au travail ? Report of the Digital Platform Labor (DiPLab) project.
- Day, M., Turner, G., and Drozdiak, N. (2019). Amazon workers are listening to what you tell alexa. *Bloomberg*.
- Difallah, D., Filatova, E., and Ipeirotis, P. (2018). Demographics and dynamics of Mechanical Turk workers. In *Proceedings of WSDM 2018: the Eleventh ACM International Conference on Web Search and Data Mining*, pages 135–143. ACM.
- Forde, C., Stuart, M., Joyce, S., Oliver, L., Valizade, D., Alberti, G., Hardy, K., Trappmann, V., Umney, C., and Carson, C. (2017). The social protection of workers in the collaborative economy. Report for the committee for employment and social affairs, European Parliament.
- Graham, M., Hjorth, I., and Lehdonvirta, V. (2017). Digital labour and development: impacts of global digital labour platforms and the gig economy on worker livelihoods. *Transfer: European Review of Labour and Research*, 23(2):135–162.
- Gray, M. and Suri, S. (2017). The humans working behind the AI curtain. *Harvard Business Review*, (9):2–5.
- Gray, M. and Suri, S. (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.
- Hara, K., Adams, A., Milland, K., Savage, S., Callison-Burch, C., and Bigham, J. (2018). A data-driven analysis of workers' earnings on Amazon Mechanical Turk. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 449:1–449:14, New York, NY, USA. ACM.
- Hirth, M., Hoßfeld, T., and Tran-Gia, P. (2011). Anatomy of a crowdsourcing platform - using the example of Microworkers.com. In *2011 Fifth International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing*, pages 322–329.
- Ipeirotis, P. (2010). Analyzing the Amazon Mechanical Turk marketplace. *XRDS*, 17(2):16–21.
- Irani, L. (2015). Difference and dependence among digital workers: The case of Amazon Mechanical Turk. *South Atlantic Quarterly*, 114.
- Irani, L. (2016). The labor that makes AI "magic". AINow Summit, White House and NYU.
- Kuek, S., Paradi-Guilford, C., Fayomi, T., Imaizumi, S., and Ipeirotis, P. (2015). The global opportunity in online outsourcing. Report, World Bank.
- Posch, L., Bleier, A., Flöck, F., and Strohmaier, M. (2018). Characterizing the global crowd workforce: A cross-country comparison of crowdworker demographics. *CoRR*, abs/1812.05948.

- Roberts, S. (2019). *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press.
- Ross, J., Irani, L., Silberman, M., Zaldivar, A., and Tomlinson, B. (2010). Who are the crowdworkers?: Shifting demographics in Mechanical Turk. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '10, pages 2863–2872, New York, NY, USA. ACM.
- Schmidt, C. (2017). Digital labour markets in the platform economy: Mapping the political challenges of crowd work and gig work. Report, Friedrich-Ebert-Stiftung.
- Schmidt, F. (2019). Crowdproduktion von Trainingsdaten: Zur Rolle von Online-Arbeit beim Trainieren autonomer Fahrzeuge. Report, Hans-Böckler-Stiftung.
- Tubaro, P. and Casilli, A. (2019). Micro-work, artificial intelligence and the automotive industry. *Journal of Industrial and Business Economics*, 46(3):333–345.
- Wood, A., Graham, M., Lehtonvirta, V., and Hjorth, I. (2019). Networked but commodified: The (dis)embeddedness of digital labour in the gig economy. *Sociology*, page (online first).