



HAL
open science

Why should I comply with taxes if others don't?: an experimental study testing informational effects

Nathalie Etchart-vincent, Marisa Ratto, Emmanuelle Taugourdeau

► To cite this version:

Nathalie Etchart-vincent, Marisa Ratto, Emmanuelle Taugourdeau. Why should I comply with taxes if others don't?: an experimental study testing informational effects. 2024. hal-04635966

HAL Id: hal-04635966

<https://hal.science/hal-04635966v1>

Preprint submitted on 5 Jul 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.

Why should I comply with taxes if others don't?: an experimental study testing informational effects

Nathalie Etchart-Vincent*, Marisa Ratto[†] and Emmanuelle Taugourdeau[§]

July 3, 2024

Abstract

Our experimental study investigates the impact of information about others' tax behaviour on the subjects' subsequent tax decisions. A novel framework allows us to test the taste for social conformity and behavioural convergence hypothesis. Two kinds of individual information are introduced, namely information about the income reported on average, within the whole subject's group and within a subgroup, made of either peers or non peers and chosen by the subject. Our main results are fourfold. First, we replicate usual results as regards the influence of tax morale, probability of audit and redistribution on tax compliance. Second, our data show that many subjects are more interested in non-peers' than peers' tax behaviour. Third, with regard to our main point, our data display a huge variety of behavioural responses at the individual level. Roughly 50% of the subjects, most of whom are full tax compliers, are insensitive to others' tax behaviour, thereby exhibiting strong intrinsic preferences towards taxes. At the same time, our data provide strong evidence of behavioural convergence towards others' average behaviour, and a taste for social anti-conformity is also found for a minority of subjects. Finally, the kind of information appears to matter, and we find some asymmetry in upward and downward behavioural variations.

Keywords: Tax compliance, Information, Tax morale, Peer effects, Social norms, Behavioral contagion, Social conformity, Artefactual field experiment, Lab experiment

JEL Codes: C91, D83, H26

*CNRS, Centre d'Economie de la Sorbonne, Université Paris 1 panthéon-Sorbonne, CNRS-IRD; MSE 106-112 Bd de l'Hôpital, 75013 PARIS, nathalie.etchart-vincent@univ-paris1.fr

[†]Université Paris-Dauphine - PSL, LEDa, UMR CNRS 8805- UMR IRD 260, 75016 Paris, France, maria_luisa.ratto@dauphine.psl.eu

[‡]Corresponding author. CNRS, CREST, 5 avenue le Chatelier 91120 Palaiseau, emmanuelle.taugourdeau@ensae.fr

[§]Financial support from iCODE (Institute for Control and Decision of the Idex Paris Saclay) is gratefully acknowledged. The experiment was run in the Laboratoire Experimental of Ecole Polytechnique, and we are grateful to Sri Srikandi for the programming, as well as for his invaluable practical assistance. We also wish to thank the participants to the numerous conferences and seminars where the paper was presented (PET, LAGV, IIPF, SOPE workshop, among others), and we are especially grateful to Yves Le Yaouanq for his helpful comments on a previous version of the paper. Finally, we would like to thank Mathilde Bouchardon for her early involvement in the research project during her Master's dissertation. She left academia and allowed us to pursue the project on our own.

1 Introduction

In the recent years, a number of huge tax evasion cases, such as the series of Offshore Leaks between 2013 and 2021, have been made public and highly-publicized. A non negligible consequence of people learning about these scandals has been to boost social discontent, enhance tax and redistribution concerns within the population (Ait Bihi Ouali, 2020) and even erode trust in the political and social system. However, in a moment when public debt soars, public budgets are severely constrained, and the lost revenue from tax evasion reaches record levels, another major challenge for public authorities is to avoid that taxpayers react to observing others cheating by questioning the rationale of their own tax compliance and start to evade taxes in their turn. Indeed, learning about others evading their taxes (be it through the media, during a family dinner or at the coffee machine at work) is likely to weaken, and possibly overcome, the social and moral barriers that initially prevented the taxpayer from evading their taxes, while fostering their motives to evade taxes.¹ Boosted by network effects (Di Gioacchino and Fichera, 2020; Traxler, 2010), the contagion of tax evasion may then occur, following the same general dynamics as the contagion of dishonesty in general (Innes and Mitra, 2013; Keizer et al., 2008).² More generally, in a world where trying to promote tax compliance while preventing (the contagion of) tax evasion is such an important concern for public policy, a major question to be addressed is whether, how and to what extent, learning about others' tax behaviour may affect the individual's subsequent tax decisions.

To date, the yet huge body of lab and field experimental evidence remains inconclusive, with high between-and within-study heterogeneity (Bazart and Bonein, 2014). A particularly salient finding is the wide range of individual behavioural reactions to social information about others' tax behaviour. However, we suspect that the diversity of reactions might be even larger than previously documented: being based on the provision to each and every subject of standardized social information, most previous experimental designs could not capture the twofold fact that people often have the opportunity to choose on whom to get information, and this information of interest may not be the same as their neighbour's. Based on an innovative and flexible experimental design, the purpose of our experimental study is precisely to enrich the informational part of the design to further investigate whether and how the subjects react to learning about others' tax behaviour.

Our methodological contribution to existing experimental research is twofold. First, we extend the characterization of 'others' (the people on whose tax behaviour the subjects are given information) to non-peers. Second, we refine the characterization of the 'provision' of social information, by giving the subjects the possibility to choose whether to get information or not, and which information to get. The starting point of our approach is that a given individual is unlikely to be interested in the tax behaviour of any kind of others (Akin, 2019). They will rather have some specific others of interest,

¹Social disapproval towards tax evaders and dishonesty contagion are likely to be stronger when, because of their outstanding position, these tax evaders are especially expected to set an example of good citizenship, as is the case with politicians (Hammar et al., 2009) or the rich and famous (Martinangeli and Windsteiger, 2024; Rockenbach et al., 2021).

²The public disclosure of non-anonymous information about the taxpayers' tax behaviour has been implemented in a number of countries as a tool to deter tax evasion and promote tax compliance, based on social image motives (such as shame and pride) and direct incentives (such as the fear of punishment) (see Bø et al., 2015; Blaufus et al., 2015; Hasegawa et al., 2013; Reck et al., 2022; Slemrod et al., 2022). In our study, information about others' behaviour remains anonymous, and our aim is to investigate whether and how learning about others' tax behaviour may induce behavioural contagion of either tax compliance or tax evasion among the taxpayers.

or “referent others” (Gächter et al., 2012), which means that providing them with information about the tax behaviour of some ‘others’ is likely to be pointless if these others are not relevant to them. So, our first challenge with regard to the experimental design is to enlarge the range of admitted “referent others” to better respond to the diversity of individual preferences.

Typically, two main kinds of referent others have been used in the experimental tax literature so far. Of the first kind are the individual’s fellow citizens. They are considered as such based on the idea that citizenship induces a strong sense of belonging to the whole society/tax community, being therefore the primary factor of social identification and social conformity in tax decision making (Gino and Pierce, 2009). The influence, on the subjects’ subsequent tax behaviour, of social information about the (average) tax behaviour among their fellow citizens (i.e. within the whole tax community) has been investigated in virtually all existing experimental studies.³ We choose to do the same here, both for its own sake and for comparison purposes. In the following, the corresponding kind of information will be referred to as ‘whole group information’.

The second kind of referent others used in a number of previous experimental tax settings are peers, loosely defined as close/similar others. The notion of peer refers to that kind of social identification induced by social proximity (Bicchieri et al., 2023) and homophily (McPherson et al., 2001).⁴ Even though both the fellow citizens and the peers have been shown to be others of (some) interest for most people, they may not be the most influential ones, and to some individuals they may even be of no interest at all. For instance, as shown in Reck et al. (2022) and Martinangeli and Windsteiger (2024), a taxpayer might also/instead be interested in the tax behaviour of some others who are far and/or different from them (e.g., politicians, the rich and famous), and thus can be referred to as ‘non-peers’. Based on this observation, we choose to extend our investigation to the non-peer kind of referent others. But this is not enough.

Not only is a given individual unlikely to be interested in the behaviour of any kind of others, but it is most unlikely that all individuals will be interested in the same others.⁵ Moreover, taxpayers in the real world may be willing to search for information by themselves, and actually find access to the most relevant-to-them piece of information (Akin, 2019).⁶ So that we can investigate the subjects’ informational strategies and their subsequent tax decisions, each subject in our experimental design is given the opportunity to choose their own reference group, in this case a subgroup of their experimental group, among three subgroups built each on the basis of their relative level of tax morale (Luttmer and Singhal, 2014) as compared to theirs. By construction, our strategy thus enables us to provide each interested subject with their endogenously determined, thus completely personalized, reference group. Depending on their choice, this chosen reference group will be made up of either peers (when

³In the lab, the whole society is usually proxied by the subject’s experimental group. Less common proxies are an experimental group in a past session (Alm, 2019), or a fictitious past experimental group (Trivedi et al., 2003).

⁴The related concept of neighbours and neighbourhood can be found in the literature too. It may either represent the geographic notion (as in Del Carpio, 2022) or be another (metaphorical) word for peers (Di Gioacchino and Fichera, 2020; Méder et al., 2012).

⁵This does not preclude some referent others from being shared by a high number of individuals, as is precisely the case at the whole society level.

⁶In this respect, the public tax disclosure system in Norway offers a wonderful life-size playground, by providing the taxpayers with a comprehensive set of information about their fellow citizens, and the possibility to choose among that enormous amount of information the very pieces they are interested in (if any). Stimulating data and analyses about the searchers and targets of the searches in 2014-2015 are provided in Reck et al. (2022).

the subject chooses the subgroup in which the level of tax morale is identical to theirs) or non-peers (when they choose a subgroup in which the level of tax morale is either higher or lower than theirs). In the following, the corresponding kind of information will be referred to as ‘chosen reference group information’. A last option allows the subjects not to get any information.

Even though the way we implement our informational strategy is novel, two recent experimental studies by Akin (2019) and Garcia et al. (2020) show similarities with ours. Akin (2019)’s study is concerned with generic dishonest behaviour rather than tax behaviour. Yet, as in our study, the subjects were given the possibility to choose which social information to get about others’ tax behaviour, and this self-selection process was a pivotal feature of the study. Now, a key difference with our framework is in the structure of informational settings: the subjects in Akin (2019) had to decide upstream whether to get information about the average or maximum score of dishonesty among others. In our within-subject design, all the subjects get average information, and all of them are also given the possibility to get specific information by choosing their preferred one among four informational options. The second close to ours lab experimental study is Garcia et al. (2020)’s. As in our study, both whole group information and reference group information were introduced. However, contrary to ours, their study was based on a between-subject experimental design, the source of information was different across the two informational settings (official vs. unofficial information), and the subjects were allocated either to the whole group (official) information treatment, or to the reference group (unofficial) information treatment. A further difference with our experimental design is that the reference groups were not chosen by the subjects but determined randomly, and reference group information was drawn from informal conversations that took place in each reference group.

Based on our experimental design, an artefactual field study is conducted with 240 subjects.⁷ The main experimental task consists in making income reporting decisions in three different informational settings: a no-information baseline, a whole group/same informational setting, and a chosen reference group/personalized informational setting. For our purpose, tax morale is elicited at the individual level at the outset of the experiment. Furthermore, two tax environments are introduced on a between-subject basis, namely a basic one with exogenous probability of audit and no redistribution, and a more realistic one with endogenous probability of audit and a redistribution mechanism, thereby inducing strategic interdependence between the taxpayers.⁸ Finally, attitude towards risk is elicited at the individual level and basic socio-demographic information is collected for the sake of experimental control.

Our main results are fourfold. First, we replicate usual results as regards the influence of tax morale, probability of audit and redistribution on tax compliance. Second, interesting patterns of attitude towards information emerge from the data, with only a small minority of subjects being not interested in information at all, and a majority of subjects being interested in non-peers’ rather than peers’ tax behaviour. Third, with regard to our main point, the data show that roughly 50% of the subjects keep on behaving the same way throughout the experiment, regardless the way other subjects behave,

⁷The artefactual field experiment (Harrison and List, 2004) departs from a pure lab experiment in that it is conducted with a more diverse subject-pool.

⁸More basically, the redistribution of collected taxes in the form of a public good obviously is a non-negligible factor for tax compliance (see Alm et al., 1992).

thereby showing strong intrinsic preferences towards taxes. Most of them are fully tax compliers, but there also is a minority of irreducible tax evaders. Even so, our data provide clear evidence of behavioural convergence towards others' average tax behaviour (thereby bringing clear support to the taste for social conformity hypothesis) while displaying a huge amount of between-subject behavioural heterogeneity. In particular, we find non negligible evidence of social anti-conformity-oriented behaviour. Finally, the kind of information appears to matter, with whole group information being more influential on tax behaviour than reference group information, and we find some upward and downward asymmetry in behavioural responses to information.

The remainder of the paper is organized as follows. An overview of the related literature is provided in Section 2. Section 3 presents the theoretical framework and predictions. Section 4 outlines the experimental design, derives the four hypotheses that will be tested, and describes the experimental procedures. The results are presented in Section 5, while the last section discusses the results and concludes.

2 Related literature

Since Allingham and Sandmo (1972)'s seminal model, the determinants of individual tax compliance have been extensively investigated. Based on a portfolio approach, this model captures the influence of basic economic factors (tax rate, probability of audit, level of the fine) on the taxpayer's decision (under risk) to evade taxes or not (see for instance Alm, 2012 for a survey). However, this purely economic approach has been strongly challenged by the observation that tax compliance usually prevails despite actual low levels of deterrence. Even though widely used third-party reporting systems actually have a deterring effect on tax evasion (see Kleven et al., 2011; Slemrod, 2007), there still are many opportunities to evade taxes, but not all are actually seized. Based on the insights drawn from behavioural economics, the social and psychological nature of the taxpayer has been acknowledged, and a number of non-economic factors to tax compliance have been identified and studied, such as the taxpayer's moral values (e.g., ethics, sympathy, or fairness; see for instance Alm and Torgler, 2011; Bordignon, 1993; Christian and Alm, 2014; Jimenez and Iyer, 2016; Kirchler et al., 2008), their moral feelings (e.g., the desire to feel proud or relieved, and to avoid guilt, remorse or shame; see for instance Elster, 1989; Erard et al., 1994; Kirchler et al., 2007; Wenzel, 2005b), and the influence of social norms.⁹ Since our study is essentially concerned with the influence of social information on behaviour, our focus here is rather on social norms (Cialdini and Trost, 1998; Cialdini and Goldstein, 2004), and more specifically, on the norms at play before information is provided, as well as on those activated by the provision of social information.

The main and most compelling social, actually societal, norm towards compliance is called tax morale.¹⁰ Tax morale can be viewed as an umbrella notion that subsumes all the individual's non-pecuniary motives to tax compliance (Luttmer and Singhal, 2014), including their moral feelings and values. It is thus usually loosely defined as the taxpayer's intrinsic motivation, or internalized will-

⁹Recent surveys can be found in Alm et al. (2012), Alm (2019), Hashimzade et al. (2013), Pickhardt and Prinz (2014).

¹⁰Tax morale is close to a number of related notions that highlight its primacy, such as tax ethics (Alm and Torgler, 2011; Pommerehne and Frey, 1992; Torgler and Murphy, 2004), civic duty (that is, the feeling of having a responsibility towards society; Orviska and Hudson, 2003), and even patriotism (Konrad and Qari, 2012).

ingness, to pay taxes (Braithwaite and Ahmed, 2005; Frey and Torgler, 2007) and contribute to the society (Cummings et al., 2009; Torgler, 2007). It is thus likely to considerably vary across taxpayers, depending on the way they metabolize internal and external pressure to tax compliance, which itself may be affected by a number of factors including socio-demographics (see Doerrenberg and Peichl, 2013 for instance). In our study, tax morale has a special status: elicited for each subject at the very outset of the experiment, it was then used during the experiment as a criterion for defining each subject’s chosen reference group.¹¹

The second kind of social norms of interest to us are those embedded in, and conveyed by, others’ behaviour. Learning about others’ behaviour is thus likely to activate the individual’s propensity to group identification and social conformity and induce them to follow (i.e., conform to) these norms by adjusting their behaviour towards others’. The need for social conformity has been shown to increase with social proximity and the strength of group identification (Gino and Pierce, 2009; Gioia, 2017; Wenzel, 2004), as best illustrated by so-called peer effects (Alm et al., 2017; Charroin et al., 2022; Dimant, 2019 and the many references therein). Therefore, learning about others’ behaviour is likely not to have the same behavioural impact on the individual depending on who these others are to them. Moreover, when the individual is made aware that others complied more than they did, both their taste for social conformity and tax morale contribute to induce them to adjust their behaviour towards others’. Conversely, when the individual realizes that they complied more than others, their tax morale and need for social conformity may come into conflict, and the resulting impact on their tax behaviour is undetermined.¹²

A number of other mechanisms have been identified, that may contribute to behavioural convergence towards others’ behaviour. The most studied case is when the individual realizes that they complied with taxes while people around evaded their taxes. First, the “everybody does it” phenomenon (Bardach, 1989) offers a comfortable self justification, as well as guilt relief, to the complying taxpayer who would be tempted to start evading taxes (or evading more) in their turn (see for instance Keizer et al., 2008 and Lefebvre et al., 2015 for empirical evidence of Wilson and Kelling, 1982’s “broken window theory”). Second, most people have a strong need for reciprocity and fairness, which has been shown to translate into conditional cooperation (Bazart and Bonein, 2014; Bordignon, 1993; Doerrenberg and Peichl, 2022; Fortin et al., 2007; Keser and Van Winden, 2000; Traxler, 2010). Uncovering that others evaded taxes, thus did not contribute their part to the common effort, the compliant individual may experience bad feelings of bitterness, resentment, “external frustration” (Méder et al., 2012), and a feeling of injustice, which may lead them to start evading taxes in their turn (Bazart and Bonein, 2014). Frustration of having been the “fall guy” will be even more intense if the tax system includes a public good/redistribution, from which the tax evader will benefit without having contributed to it (Bazart and Bonein, 2014; Bicchieri et al., 2023; Frey and Torgler, 2007; Traxler, 2010).¹³ An important point to make here is that bad norms tend to be more sticky than good ones (Dimant, 2019

¹¹In data analysis, tax morale was also classically introduced as an explanatory variable when analysing the determinants of tax compliance, and as a control variable when studying the impact of social information on tax behaviour.

¹²For instance, behavioural adjustment towards others’ behaviour might be restricted to low-tax morale individuals (Charroin et al., 2022).

¹³Note that, in the opposite case (i.e., when the individual evaded taxes and learns about others’ complying behaviour), there also may be bad feelings of guilt and shame, combining with basic tax morale to prompt the individual to adjust their behaviour towards others’, thus to comply more.

and the many references therein), thereby making bad behaviour more contagious than good one. In that respect, downward behavioural adjustments towards others' less honest/compliant behaviour have been shown to be of greater magnitude than upward behavioural adjustments towards others' more honest/compliant behaviour (see Dimant, 2019; Innes and Mitra, 2013; Isler and Gächter, 2022; Lefebvre et al., 2015; Trivedi et al., 2003 for instance).

In spite of the good reasons for behavioural convergence to occur, experimental evidence remains inconclusive and somewhat contradictory. Whether in the lab or in the field, both supportive (Alm et al., 2015, 2009, 2017; Bazart and Bonein, 2014; Charroin et al., 2022; Del Carpio, 2022; Diekmann et al., 2015; Fosgaard et al., 2013; Gino et al., 2009; Hallsworth et al., 2017; Innes and Mitra, 2013; Jamison et al., 2021; Lefebvre et al., 2015) and ambiguous or not significant results (Blumenthal et al., 2001; Castro and Scartascini, 2015; Dwenger et al., 2016; Fellner et al., 2013; Jamison et al., 2021; Rauhut, 2013; Wenzel, 2004) have been brought to light. However, the reasons behind that mixed evidence are to be found in the huge between- and within-study heterogeneity that actually obscures the overall picture.¹⁴ In the last twenty years, the influence of social information on tax behaviour has been shown to strongly vary depending on both the experimental stimuli and the subjects' individual characteristics, such as their baseline level of tax compliance (Bazart and Bonein, 2014; Castro and Scartascini, 2015; Del Carpio, 2022; Dwenger et al., 2016; Garcia et al., 2020), prior beliefs (Fellner et al., 2013), sensitivity to social norms (Kimbrough and Vostroknutov, 2016) and personal factors (Hofmann et al., 2017; Trivedi et al., 2003).

Even though behavioural convergence appears to be the most intuitive reaction to social information about others' behaviour, alternative reactions are still possible. First, people may exhibit behavioural status quo, meaning that they will not change their behaviour after learning about others' behaviour. There are several reasons why this may happen. Maybe they just do not care about this piece of information they did not choose to get. Maybe they are interested in that information but only out of (pure) curiosity, while knowing that they will stick to their initial (complying/evading) tax behaviour anyway (which may happen if the target is not a really significant other to them; see Reck et al., 2022 for a nice example).¹⁵ Maybe they consider that the informational content does not call for a change in behaviour. Anyway, regardless their very motive, these people are not affected by the social norm embedded in (these) others' behaviour, and they do not exhibit any taste for social (anti-)conformity. Second, instead of adjusting their behaviour towards others' behaviour, some people may be willing to adjust their behaviour away from others' behaviour (Fortin et al., 2007; Garcia et al., 2020), thereby signalling that they wish to deviate from the social norm embedded in others' tax behaviour. There are several reasons why it may happen.¹⁶ In particular, the individual may have a taste for social anti-conformity (see Dvorak et al., 2024 and the many references therein; see also Fortin et al., 2007;

¹⁴Typically, inconclusive overall results may result from the offsetting of two, or more, opposite effects (see Castro and Scartascini, 2015; Rauhut, 2013; Wenzel, 2004).

¹⁵Some highly compliant subjects may look strongly confident in their willpower and conviction to resist temptation to misbehave (see Charroin et al., 2022).

¹⁶The consequences of wrong prior beliefs updating should also be mentioned: in case of initial wrong beliefs about others' attitude towards taxes, divergent behavioural adjustment may then be the unexpected, but regular, result of a rational belief updating process after learning about others' actual behaviour (see Diekmann et al., 2015; Rauhut, 2013; Wenzel, 2005a for nice analyses).

Ushchev and Zenou, 2020). Indeed, even though the taste for social conformity appears to be both a dominant neurally-based human trait (Morgan and Laland, 2012) and a socially dominant trait, it is not universal. A minority of individuals actually exhibit a taste for social anti-conformity, based on either a desire for rebellion, the rejection of the rules of the society as a matter of principle, or a desire to assess one’s individuality and distinctive identity and carve out one’s own behavioural path (Gino and Pierce, 2009; Garcia et al., 2020; Wenzel, 2004). Then, it is as if learning about others behaving differently would strengthen and harden the individual’s initial behavioural orientation, inducing them to commit even more to that behaviour and run counter to the majority’s behaviour. Observing others evading (resp. complying with) taxes all around, the anti-conformist tax complier (resp. evader) will become even more (resp. less) compliant, thereby widening the gap between others’ behaviour and theirs and inducing behavioural divergence.¹⁷

As just shown, the range of possible behavioural reactions to social information is rather wide. Our experimental design is precisely built to allow the expression of that behavioural diversity, and our data will be analysed with that all-around heterogeneity in mind.

3 Conceptual framework

We now present the theoretical background of our work, the rationale for which is to substantiate our experimental design and provide a guidance for upcoming working hypotheses. To that end, an ad hoc micro-founded model was developed to describe the taxpayer’s preferences towards income reporting, in the presence of information about referent others’ income reporting behaviour, thus in line with our experimental purpose.

As was done in a number of recent papers that also aimed to combine “rational choice with social influence mechanisms” (Noguera et al., 2014) and/or moral motives (see for instance Christian and Alm, 2014; Fortin et al., 2007), the taxpayer’s preferences are described using a several-component separable model. In our case, three components are considered.

The first one is a standard Allingham-Sandmo-Yitzhaki’s expected-utility set-up (Allingham and Sandmo, 1972; Yitzhaki, 1974), which captures the influence of basic economic factors on the (portfolio-like) income reporting decision (see for instance Alm, 2012 for a survey).¹⁸ In our model as in our experiment, we authorize the audit probability to be endogenous to the taxpayer’s income reporting decisions. The reason for this is twofold. First, the idea is to make the tax system more realistic since audit probabilities are actually endogenous in the real world (Torgler, 2002). Second, by creating strategic interdependence between the taxpayers, endogenous audit probability is an extra (social norm-free) channel through which social information about others’ tax behaviour may impact the taxpayer’s behaviour.

The second component of the model introduces a basic redistribution mechanism/public good provi-

¹⁷Note that, in a realistic tax setting with redistribution and endogenous audit probability, the strategic motive to free ride also is boosted when others comply, which gives the committed evader an additional reason to keep on evading their taxes while (because!) the majority complies. In Garcia et al. (2020), tax evaders precisely seem to be more committed to tax evasion than tax compliers to tax compliance.

¹⁸Even though our main focus is on the influence of social norms on tax behaviour, the basic financial components of the tax decision setting (tax rate, probability of audit and fine) obviously remain deciding factors in the income reporting decision (see Alm, 2019; Alm and Malézieux, 2021; Dularif et al., 2019; Slemrod, 2019).

sion mechanism, as was done in Blaufus et al. (2015), Christian and Alm (2014), Cowell and Gordon (1988), Hammar et al. (2009), and Méder et al. (2012). Such a redistribution mechanism offers three advantages. First, as the endogenous audit probability, it allows greater realism by capturing the basic economic nature of taxes in the real world (which is to finance public expenses/goods and allow for redistribution among taxpayers). Second, by inducing strategic interdependence between the taxpayers (as the endogenous audit probability), it provides an extra (social norm-free) channel through which social information about others' tax behaviour may impact the taxpayer's behaviour. Third, a specific advantage of redistribution is that it may create additional fairness concerns, as well as free riding temptations, when the taxpayer learns about others' tax behaviour.

Finally, the third key component of our model characterizes the psychic cost of tax compliance. This cost is understood as the sum of all non-monetary costs associated with the process of tax compliance, from reporting income to paying taxes (see Lopes and Martins, 2013; Nar, 2023, and the references therein).¹⁹ They first consist of basic time and psychological costs (such as the amount of time involved, the trouble experienced when looking up and reading the instructions, the anxiety and stress of making mistakes, but also the jealousy, bitterness and resentment towards other taxpayers, as well as the feeling of injustice towards the tax system; see Lopes and Martins, 2013), but they also include opportunity costs, that is all the costs that could have been avoided if one had chosen to evade taxes instead of complying with taxes (Nar, 2023).

These costs obviously depend on the taxpayer's current income reporting behaviour: the higher the reported income, the higher the cost. We further assume that they may also be affected by the previous income reporting behaviour of other taxpayers, and more specifically by the gap between the income previously reported by the taxpayer and that previously reported by other taxpayers.²⁰ Called 'reported income gap' in the following, this gap is integrated into the model in the form of the difference, at the previous period, between the income reported by the individual and the average income reported within their reference group. Finally, we assume that compliance costs depend on two key personal characteristics of the taxpayer, namely their level of tax morale and their taste for social (anti-)conformity. Since it is meant to capture the individual's intrinsic attitude towards taxes, thus a rather stable individual feature (at least in the short run; see Doerrenberg and Peichl, 2013; Halla, 2012), tax morale is modeled as an exogenous parameter (as was done in Méder et al., 2012, but contrary to what was done in Besley et al., 2023; Di Gioacchino and Fichera, 2020; Traxler, 2010). Compliance costs are assumed to decrease as tax morale increases. Similarly, since the taste for social (anti-)conformity parameter is meant to capture the individual's intrinsic attitude towards conforming, or not, to (the social norm embedded in) others' behaviour (Fortin et al., 2007; Ushchev and Zenou, 2020), it is modeled as an exogenous parameter, which can be either positive in case of a taste for social conformity or negative in case of a taste for social anti-conformity. Relaxing the usual restriction put on (the sign of) that parameter allows our model to account for more diverse behaviour:

¹⁹Note that this approach is dual to the more usual one based on (tax evasion) "psychic costs" (Benjamini and Maital, 1985; Gordon, 1989), which capture the many non-economic costs borne when evading taxes (Frey and Torgler, 2007; Myles and Naylor, 1996; Traxler, 2010), and especially those related to experiencing bad feelings such as guilt, remorse and shame (as well as a social stigma when tax evading behaviour is publicized).

²⁰In the standard "psychic cost" (of tax evasion) literature, a number of theoretical papers also introduce the income reporting behaviour of other taxpayers in the taxpayer's utility function, and do so using different modeling strategies (see for instance Fortin et al., 2007; Gordon, 1989; Méder et al., 2012; Myles and Naylor, 1996).

a given reported income gap will not affect the taxpayer's compliance costs (and their final tax behaviour) the same way depending on their having a taste for either social conformity or anti-conformity.

In our setting, the taxpayer's utility function is thus given by:

$$U(x) = E[u(x_i); p] + \frac{g}{n}[(n-1)t\bar{X}_{-i} + tx_i] - F(x_i; \Delta_{i,k}, \theta, \alpha_i) \quad (1)$$

where:

$E[u(x_i; p)]$ denotes taxpayer i 's expected utility as in Yitzhaki (1974), and is given by:

$$E[u(x_i; p)] = (1-p)u(R - tx_i) + pu(R - tx_i - t(1+\pi)(R - x_i))$$

with p is the audit probability, R the income of the individual, x_i their reported income, t the tax rate and π the fine rate on evaded taxes.

As regards the audit probability p , two possibilities are contemplated: p is either held fixed as in Yitzhaki (1974), or made endogenous. In the latter case, we assume that $p = p(x_i, \bar{X}_{-i})$ is a decreasing function of the reported income x_i ($p'_{x_i} \leq 0$), as well as an increasing function of \bar{X}_{-i} ($p'_{\bar{X}_{-i}} \geq 0$).²¹ We further assume that $p''_{x_i} = 0$, $p''_{\bar{X}_{-i}} = 0$ and $p''_{x_i \bar{X}_{-i}} = 0$.

$\frac{g}{n}[(n-1)t\bar{X}_{-i} + tx_i]$ represents the part of the public good from which taxpayer i benefits. In eq.(1), \bar{X}_{-i} denotes the average income reported by the $(n-1)$ taxpayers with whom the individual contributes to the public good through the payment of taxes, x_i their own contribution (i.e. their reported income) and g an efficiency factor that leverages the amount of collected taxes to constitute the "public fund", as in Christian and Alm (2014). This public fund is redistributed equally among the n taxpayers.

$F(x_i; \Delta_{i,k}, \theta_i, \alpha_i)$ is a cost function that reflects the compliance costs previously defined. F is a function of variables x_i (individual i 's reported income) and $\Delta_{i,k}$, the reported income gap, which is defined as the difference between the income individual i reported in a previous period (denoted \hat{x}_i) and the income reported on average in individual i 's reference group k in a previous period (denoted \bar{X}_k). So the reported income gap can be written as:

$$\Delta_{i,k} = (\hat{x}_i - \bar{X}_k) \quad (2)$$

The compliance cost also depends on individual parameters $\theta_i \in [0, 1]$ (which captures taxpayer i 's tax morale), and α_i (which captures taxpayer i 's taste for social (anti-)conformity with group k and can be either positive if individual i has a taste for social conformity and wishes to adjust their behaviour towards group k 's behaviour, or negative if they rather have a taste for social anti-conformity and

²¹This assumption is grounded on the screening models of audit (Reinganum and Wilde, 1985; Scotchmer, 1987; Macho-Stadler and Pérez-Castrillo, 2002). In these models, tax enforcement authorities pre-commit to an audit strategy which depends on the reported income. The optimal income reporting decision of the taxpayer implies that the optimal audit rule declines with reported income.

wish to adjust their behaviour away from group k 's behaviour).

We assume that the compliance cost is convex with x_i : $F'_{x_i} \geq 0$ and $F''_{x_i} \geq 0$. Moreover, we assume that the marginal cost of income reporting decreases as the level of tax morale θ_i increases: $F''_{x_i\theta_i} \leq 0$. A taxpayer who is fully tax moral ($\theta_i = 1$) bears a smaller cost for paying taxes than a taxpayer whose tax morale is lower ($\theta_i < 1$).

Now, we further assume that the marginal cost of income reporting depends on the reported income gap $\Delta_{i,k}$. If individual i has a taste for social conformity ($\alpha_i > 0$), the marginal cost of income reporting increases in $\Delta_{i,k}$ and individual i has an incentive to conform to group k 's (average) behaviour. Therefore, if the reported income gap is positive (resp. negative), the conformity-oriented individual will be induced to decrease (resp. increase) their reported income. Conversely, if individual i has a taste for social anti-conformity ($\alpha_i < 0$), the marginal cost of income reporting decreases in $\Delta_{i,k}$ and individual i prefers to depart from group k 's (average) behaviour. Therefore, if the reported income gap is positive (resp. negative), the anti-conformity oriented individual will be induced to increase (resp. decrease) their reported income.

Then, the sign of the cross derivative depends on the sign of the taste for social (anti-)conformity of individual i (α_i):

$$\begin{aligned} F''_{x_i, \Delta_{i,k}} &> 0 \quad \text{if} \quad \alpha_i > 0 \\ F''_{x_i, \Delta_{i,k}} &< 0 \quad \text{if} \quad \alpha_i < 0 \end{aligned}$$

We now turn to the predictions that can be derived from the comparative statics of our theoretical framework (See Appendix 7.1). Predictions P1, P2 and P3 relate to the influence of standard economic factors, namely probability of audit, public good/redistribution and tax morale respectively, on economic reporting behaviour. Prediction P4 is specific to our theoretical framework, and it is the cornerstone of our work. The correspondence between our theoretical predictions, denoted Pj , and our upcoming experimental hypotheses, denoted Hj , is made through the subscript j .

P 1. *When the audit probability is fixed ($p = \bar{p}$), the amount of tax evasion decreases as the audit probability increases $\left(\frac{\partial x_i}{\partial p} > 0\right)$;*

When the audit probability is not fixed, an increase in the average reported income within the taxpayer's reference group increases their audit probability; consequently, an increase in the average reported income within the reference group decreases the taxpayer's amount of tax evasion: $\left(\frac{\partial x_i}{\partial \bar{X}_{-i}} > 0\right)$

P 2. *The amount of tax evasion decreases as the return from the public good contributions increases $\left(\frac{\partial x_i}{\partial g} > 0\right)$*

Note that this effect decreases as the number of taxpayers in the community increases, to finally disappear when the community is extremely large (free riding effect) (See Equation 6 in Appendix).

P 3. *The amount of tax evasion decreases as the level of tax morale increases $\left(\frac{\partial x_i}{\partial \theta_i} > 0\right)$*

The following prediction is specifically related to our model:

P 4. *The effect of the reported income gap (i.e. the gap between the reported income of taxpayer i and the average reported income within their reference group k) on taxpayer i 's amount of tax evasion depends on their taste for social (anti-)conformity (i.e. on the sign of α_i).*

- *if $\alpha_i = 0$: an increase in the reported income gap has no impact on taxpayer i 's amount of tax evasion: they do not care about others' tax behaviour (P 4.a).*
- *if $\alpha_i > 0$ i.e. taxpayer i likes to conform to others' tax behaviour (taste for social conformity):*
 - *if $x_i > \bar{X}_k$, an increase in the reported income gap ($|\Delta|_{i,k}$ increases) increases taxpayer i 's amount of tax evasion ($\frac{\partial x_i}{\partial \Delta_{i,k}} < 0$). This case can be referred to as a “fall guy” effect (P 4.b).*
 - *if $x_i < \bar{X}_k$, an increase in the reported income gap decreases taxpayer i 's amount of tax evasion ($\frac{\partial x_i}{\partial |\Delta|_{i,k}} > 0$). This can be referred to as a “guilty feeling” effect (P 4.c).²² Since the taxpayer's marginal cost of income reporting is an increasing function of the reported income gap, having previously reported a lower income than others leads the taxpayer to reduce their amount of tax evasion.*
- *if $\alpha_i < 0$ i.e. taxpayer i does not like to conform to others' tax behaviour (taste for social anti-conformity):*
 - *if $x_i > \bar{X}_k$, an increase in the reported income gap decreases taxpayer i 's amount of tax evasion ($\frac{\partial x_i}{\partial \Delta_{i,k}} > 0$). This case can be referred to as a “self sacrifice” effect (P 4.d).*
 - *if $x_i < \bar{X}_k$, an increase in the reported income gap increases taxpayer i 's amount of tax evasion ($\frac{\partial x_i}{\partial |\Delta|_{i,k}} < 0$). This case can be referred to as a “pure selfishness” effect (P 4.e).²³ Since the taxpayer's marginal cost of income reporting is a decreasing function of the reported income gap, having previously reported a lower amount than others leads the taxpayer to increase their amount of tax evasion.*

4 The Experimental design and tested hypotheses

The purpose of our experiment is to investigate whether and how information about others' income reporting behaviour is likely to affect the individuals' level of tax compliance. For that purpose, two types of information were introduced in the experimental design. The first one was information about the average income reported by the other members of the subject's experimental group (meant to embody the subject's tax community). In the following, this kind of information will be referred to as “whole group information” (WGI). It was provided to all the subjects. The second one is information about the average reported income within a specific (sub-)group of the subject's experimental group.

²²It is easier to understand the mechanism with the absolute value of Δ . However, even with the regular value, we can see that an increase in the reported income gap ($\Delta_{i,k} < 0$ becomes more negative) lowers the taxpayer's amount of tax evasion and ($\frac{\partial x_i}{\partial \Delta_{i,k}} < 0$).

²³Again, it is easier to understand the mechanism with the absolute value of Δ . However, even with the regular value we would have ($\frac{\partial x_i}{\partial \Delta_{i,k}} > 0$).

Selected by the subject among three subgroups, built each on the basis of their relative level of tax morale as compared to theirs, the chosen subgroup was thus meant to embody the subject's reference group. From now on, this kind of information will be referred to as "chosen reference group information" (CRGI). More specifically, each subject was asked whether they wished to get information about the average income reported within the subgroup of subjects (from the same experimental session) whose level of tax morale was either equal to, lower than, or higher than theirs. Since they were also offered the possibility not to get any of these three pieces of information, they were actually proposed a four-(informational) option menu.

Even though, for obvious control and feasibility reasons, the criterion on which the subgroups were based (that is, tax morale) was not chosen by the subjects, we think that the subgroup chosen by a subject to get behavioural information about can be fairly considered as their credible "reference group" in the context of the experiment. The reason for this is twofold. First, the subjects had complete freedom either to choose a subgroup if they had some interest in one of the three informational options, or to decline information if they were not interested in any. The fact that the subjects were offered a way out increases the likelihood that the decision to choose a subgroup be really motivated (at least by curiosity), which enhances the credibility of the chosen subgroup as a reference group. Second, both the tax morale criterion and the way it was used to build the reference groups (based on the relative level of tax morale of others as compared to one's) were chosen for their affinity with actual attitudes in real life. Typically, it is rather easy in France to learn about others' (often black-and-white) opinions on State intervention and general, and taxes in particular, and quite natural to infer expectations about the tax behaviour of those who expressed these opinions: people who hold positive (resp. negative) opinions about taxes are expected to comply with their taxes (resp. evade their taxes if they get the opportunity to).²⁴ There may also be some curiosity as to whether actual tax behaviour is consistent with stated level of tax morality. But this could be more than pure curiosity, because (as mentioned in the introduction), most people are driven by (conscious or not) social comparison and fairness motives, which may stimulate their interest in the tax behaviour of those taxpayers whose tax morality is equal to, or different from, theirs.²⁵ Finally, given the prevalence of social comparisons in tax matters, we think the way we built our reference groups was relevant and meaningful to the subjects.²⁶

4.1 The experimental design: tax environments and decision settings

The experiment involved three parts, as well as a post-experimental questionnaire. The experimental procedures are described in details in Appendix 7.2. The third part of the study was introduced to investigate a framing effect; it will not be presented here. The post-experimental questionnaire was essentially meant to collect basic socio-demographic data. So, we just present the first and second

²⁴In France, politics-related conversations are rather frequent, be it with family and friends, colleagues or even strangers sometimes, and opinions about taxes can be expressed quite openly and give rise to lively discussions/debates.

²⁵Typically, one may be willing to check whether those who state similar-to-theirs opinions also behave similarly, and even eager to see whether those who portray themselves as people of integrity, with highest standards of tax morality, behave in accordance with these standards.

²⁶The accuracy of our strategy was actually discussed and tested informally with several people around us before the design was finalized.

parts of the experimental design.

The first part was devoted to eliciting the subjects' individual risk attitude and level of tax morale. Both had to be elicited at the very outset of the experiment, before any reference to the tax nature of the main task (tax morale also had to be elicited before reference groups were introduced). The procedure that was used to elicit risk attitude and tax morale at the individual level will be described in Subsection 4.3.

The second part of the experiment, which is the core of the study, was based on several tax decisions made in a given tax environment and decision setting. Two tax environments T1 and T2 were introduced, as well as the same four decision settings A, B, C, D in each tax environment. The two kinds of information were provided in each decision setting, but the way/order it was implemented differed across decision settings. The design of the three-round tax game in Part 2 (the sequence of which was identical in tax settings T1 and T2) is summarized in Table 1.

Decision settings A and B (resp. C and D) involved 3 rounds (resp. 2 rounds), meaning that the individuals had to make 3 (resp. 2) income reporting decisions. In all the decision settings, the first income reporting decision was made with no additional information (Round 1), while the subsequent one(s) was (were) made after being given a piece of social information (Round 2 and Round 3). This strategy provided us with a baseline (Round 1) from which to investigate the influence of social information on further income reporting decisions. Moreover, the behavioural data collected in Round 1 were used as informational input in subsequent rounds, meaning that, both in Rounds 2 and 3, information about others' income reporting behaviour was computed from the decisions they made in Round 1.²⁷ Now, in decision settings A and B, WGI and CRGI were given in separate rounds, and the subject had to report their income after each piece of information (thus in each Round 2 and 3). In decision settings C and D, WGI and CRGI were given at the same time, and a unique income reporting decision was made after that (in Round 2). The rationale for introducing A and B on the one side, and C and D on the other side, was to see whether final income reporting decisions (i.e., the decisions made in the last round, after both kinds of information had been received) would differ depending on whether an intermediate decision was made after each piece of information (as in A and B) or not (as in C and D). Finally, the order in which WGI and CRGI were provided in Rounds 2 and 3 (resp. in Round 2) was reversed between A and B (resp. C and D) to control for information delivery order effects.

Now, as regards the two tax environments (hereafter labelled T1 and T2), the first one (T1) was a basic tax environment with exogenous probability of audit, and no public good/redistribution of collected taxes. In such a tax environment, the taxpayers' decisions are independent from each other, so that social information about others' cannot influence the individual's tax decisions other than through a pure social norm effect. By contrast, the second tax environment (T2) was made more realistic by including both an endogenous probability of audit and the redistribution of collected taxes (or the funding of a public good), which are actual components of real tax systems. These two additional

²⁷In Rounds 2 and 3, all the subjects (including each subject's others) were likely to make their income reporting decisions based on others' previous income reporting behaviour. Due to the endogeneity of the decision process in our framework, providing the subjects with information about others' Round 2 tax behaviour would then have induced strong path dependency from round to round. Conversely, the subjects' income reporting decisions in Round 1 were independent from each other, thus provided a sound exogenous ground to investigate the effects of social information.

Table 1: The design of the tax game in Part 2

Decision setting Round i	A	B	C	D
Round 1	No info			
	Income reporting decision (1)	Income reporting decision (1)	Income reporting decision (1)	Income reporting decision (1)
Round 2	Info: WGI	Before info: choice of the reference group Info (or not): CRGI	Info: WGI + Before info: choice of the reference group Info (or not): CRGI	Before info: choice of the reference group Info (or not): CRGI + Info: WGI
	Income reporting decision (2)	Income reporting decision (2)	Income reporting decision (2)	Income reporting decision (2)
Round 3	Before info: choice of the reference group Info (or not): CRGI	Info: WGI		
	Income reporting decision (3)	Income reporting decision (3)		

Note: WGI: Whole group information; CRGI: Chosen reference group information

components make the taxpayers' decisions interdependent. Therefore, social information about others' tax behaviour in T2 may affect the individual's subsequent tax decisions twice: directly through social norm effects, as in T1, but also indirectly through its effects on endogenous audit probability (which structurally depends on others' tax behaviour and should be updated after learning about it) and public good funding (since information about others' average tax behaviour also provides information about the actual level of public good funding/redistribution).²⁸

4.2 The main hypotheses to be tested experimentally

The specific features of our experimental design allowed us to test most previous theoretical predictions (see Section 3, *supra*) as well as some additional hypotheses.

First, note that, in our experimental design, endogenous audit probability and public good funding were introduced together, while we only drew separate theoretical predictions about the influence

²⁸Since endogenous probability and public good funding were introduced together in T2, their respective influence on tax compliance could obviously not be disentangled. But this was not a concern, since our strategy was not to quantify these effects separately, but to build a realistic, thus comprehensive, tax environment (T2) that could be compared to a non realistic basic setting (T1).

on tax compliance of endogenous audit probability (P1) on the one side, and of public good funding (P2) on the other side. However, the direction of predictions P1 and P2 suggest that the combined effect of endogenous audit probability and redistribution might be positive. So, our conjecture is that having both an endogenous probability of audit and a redistribution mechanism (as is the case in tax environment T2) might enhance tax compliance, as compared to a situation where the probability of audit is exogenous and the taxpayers do not benefit from the collected taxes (as is the case in our baseline tax setting T1). From the preceding, we thus draw our first hypothesis H1:

H 1. *The level of tax compliance should be higher in T2 than in T1.*

The impact of risk attitude could not be directly derived in the theoretical model. However, since tax evasion is a risk taking behaviour, we draw our second hypothesis H2:

H 2. *The higher the level of risk aversion, the higher the level of tax compliance.*

Now, as regards the influence of tax morale on income reporting behaviour, theoretical prediction P3 suggests that tax morale works in favour of tax compliance. From which our third hypothesis H3:

H 3. *The higher the level of tax morale, the higher the level of tax compliance.*

Now we come to the cornerstone of our work, which was to investigate whether, and in what direction, social information about other taxpayers' tax behaviour may influence one's tax behaviour. As was established based on our theoretical model (see Prediction P4 in Section 3, supra), testing whether social information impacts the taxpayer's income reporting behaviour boils down to testing whether the reported income gap (i.e., the difference, denoted $\Delta_{i,k}$, between the income previously reported by the taxpayer and the income previously reported by others) influences the taxpayer's current income reporting decision (x_i). This in turn is equivalent to testing whether the taxpayer's taste for social (anti-)conformity parameter (α_i) is significantly different from 0 or not. Moreover, in case α_i actually turns out to be significantly different from 0, the direction in which the taxpayer's reported income will be impacted is given by the sign of α_i , which indicates whether the taxpayer has a taste for either social conformity or anti-conformity.

As suggested by the comparative statics of our theoretical model, the sign of α_i in our experimental design can be easily determined, based on the observation of the sign of $\frac{\partial x_i}{\partial \Delta_{i,k}}$. So, investigating experimentally the correlation between $\Delta_{i,k}$ and x_i should give us some insights into how the taste for social (anti-)conformity actually enters the subjects' utility function. From a behavioural perspective, the question we aimed to answer with our experimental design is whether, learning about others' tax behaviour, the subjects would adjust their own tax behaviour towards that behaviour, or away from it, thereby exhibiting either behavioural convergence to, or divergence from, others' behaviour. From a social norms standpoint, the question rather becomes whether, being faced with the social norm embedded in others' tax behaviour, the subjects would tend to comply with, or deviate from, that norm. Three cases should therefore be distinguished, depending on whether the individual cares or not about the social norm embedded in others' tax behaviour and, if they do, in which direction (towards either social conformity or anti-conformity). When the individual does not care about the social norm, α_i is null and there is no reason for their tax behaviour to change after getting social

information. When they are sensitive to the social norm (α_i is non null), the direction of their change in behaviour (towards either convergence or divergence) after getting some piece of information will depend on the sign of α_i .

At this stage, we do not expect any differences in behavioural reactions to social information depending on the type of information received (WGI vs. CRGI). Besides, we expect no change in behaviour among those subjects who were not interested in getting CRGI. From the preceding, we thus draw our fourth hypothesis **H 4**:

H 4. *In the whole group/chosen reference group information setting, the individual's income reporting behaviour is expected to be affected by social information about others' income reporting behaviour. The direction of the individual's behavioural adjustment will depend on their taste for social (anti-)conformity.*

H 4a. *In case the subject does not care about the social norm embedded in others' income reporting behaviour, social information should not affect their income reporting behaviour.*

H 4b. *In case of a taste for social conformity, behavioural adjustment should go downward (resp. upward) for the subjects whose reported income was initially higher (resp. lower) than average.*

The adjustment downward (resp. upward) refers to as the “fall guy” effect (resp. “guilty feeling” effect) previously defined in prediction **P 4b** (resp. **P 4c**).

H 4c. *In case of a taste for social anti-conformity, behavioural adjustment should go upward (resp. downward) for the subjects whose reported income was initially higher (resp. lower) than average.*

The adjustment upward (resp. downward) refers to as a “strong conviction” effect (resp. “pure selfishness” effect) previously defined in prediction **P 4d** (resp. **P 4e**).

H 4d. *Regardless of their attitude towards social (anti-)conformity, no change in income reporting behaviour should occur in the subjects who refused to get reference group information.*

4.3 The experimental procedures

The experiment was programmed using the z-tree software (Fischbacher, 2007). It was conducted in February and March 2016 at the Experimental Economics Laboratory at Ecole Polytechnique with 240 participants, who were divided into twelve 20-subject sessions. They all chose to participate on a voluntary basis, from the Laboratory subject database.²⁹ All the participants-to-be registered for one of the proposed slots, with no information about the content of the experiment. Since our design required 20-subject sessions, we used overbooking to make sure of it. Supernumerary participants who showed up did not take part in the session but they received a 5 euro flat-fee. Each session was characterized by the combination of two features, namely the underlying tax environment (either T1 or T2) and the decision setting (A, B, C, or D). In each tax environment T1 and T2, two (resp. one) sessions were conducted for each decision setting A and B (resp. C and D). So, six sessions were run in each tax environment. Over the whole study, four sessions were conducted within decision setting A

²⁹2/3 of the subject pool had already taken part in an experiment before ours.

(resp. B), and two within decision setting C (resp. D). Each session (with its three parts and post-hoc questionnaire) lasted approximately one hour.

In each experimental session, each of the 20 subjects was welcomed in front of the experimental room and invited to sign a registration form before entering the room. Subjects were then invited to take a seat at a given, selected at random, computer station. They were well separated and not allowed to communicate during the session. Anonymity was preserved at every stage of the experiment, to avoid any social pressure that would have been a confounding factor. Even though the wording was kept as neutral as possible, the experimental design itself was meant to help the subjects of a given session see their experimental group as a real tax community.

On their computer table, the subjects found a consent form which they had to fill in first. Then they were given time to read a general instruction-sheet, before we let them listen to a recorded message, with the same instructions. After that, they could call on us (privately, to avoid social image concerns) for some clarification if necessary. General instructions were silent about the aim of the experiment, as well as about the further tasks to be completed: the subjects were only informed that the experiment would involve several parts and they would have to make economic decisions. The subjects were informed that they could earn some money above the 5 euro show-up fee, depending on their decisions during the experiment, and that their total gain could go up to 20 euros (resp. 23.50 euros) in T1 (resp. T2). No further detail was given at this stage. They were invited to turn off their mobile phone and to double-click on their computer to start the experiment.

The first part of the experiment started with a lottery game meant to assess each subject's risk attitude based on the "Holt-Laury" procedure (Holt and Laury, 2002). Lottery amounts were chosen to be of the same order of magnitude as those manipulated in the tax part of the experiment, in which all the subjects started with a 1000 (tokens) income. Further details about the procedure are provided in Appendix 7.2. After that, each subject's level of tax morale was elicited in a very standard way, based on a single survey question drawn and adapted from the 2008 European Values Survey (EVS).³⁰ At this early stage, the priority was to prevent the subjects from realizing that our interest was in their tax morality. Indeed, it was both important to minimize self-image issues when stating their level of tax morality, and to avoid undesirable memory and self-consistency effects in their subsequent tax decisions. To that end, we used a muddy-the-waters strategy, concealing the tax morale-related question among a series of eight pairs of opposing statements on different political economy matters (such as environment, liberalization, or competition), also adapted from the 2008 EVS. For each pair of eight statements, the subject was asked to state their own position on a 7-point scale (from 0 to 7) between the two extreme statements.

As regards tax morale, a basic index was thus obtained for each subject, based on their stated score to the following (fourth) question, with 0 indicating the absence of tax morale and 7 indicating full tax morale:

Some people think that not reporting one's whole income to the tax authority is fully legitimate when given the opportunity to (score: 0). Others think that accurately reporting one's income to the tax authority is a civic or moral duty (score 7).

³⁰The relevance of that widely-used basic measure is nicely discussed in Guerra and Harrington, 2018

Table 2: The parameters and their numerical value

Notation	Meaning	Numerical value
R	Income	1000 tokens
π	Fine rate	1
g	Efficiency factor of the public good	1
p	Audit probability	0.3 in T1 $p \in [0.1, 0.5]$ in T2 ³¹

Where do you consider yourself in the score ranking?

After the risk attitude and tax morale elicitation tasks were completed, a new screen indicated “Part 2”. Once all the subjects had attained this screen, they were told verbally that they were now going to make several income reporting decisions. A new series of documents were distributed to provide them with instructions, which they were asked to read carefully. The first document provided them with general instructions about the tax game and the payoff they would receive at the end of the session. A second folder contained a thorough description of the tax system as well as some examples, meant to help the subjects understand the mechanisms at work and the consequences of their income reporting decisions. We had to provide the subjects with a number of details and examples to ensure that the procedures were well-understood (especially in T2, in which the probability of audit was endogenous). A third document presented the details of the performance-based payment procedure (only Part 2 was actually incentivized), as well as the exchange rate, here equal to 15/1000, since all the amounts involved in the tax games were in experimental money (tokens). In T2, the subjects were additionally informed that their probability of audit would depend on both their own tax behaviour (i.e., their reported income) and others’ average tax behaviour (i.e., the average reported income among the other members of the subject’s experimental group). They were provided with a fourth instruction sheet, on which a table reported the value of the audit probability as a function of these two variables (see Figure 1 in Appendix 7.2).

The main stimuli values are given in Table 2. The notations here refer to those introduced in the theoretical model. For the sake of comparability, as well as to avoid additional fairness concerns that would be a confounding factor here, all the subjects were given the same income R ($= 1000$ tokens) and informed of this feature. The fine rate π ($= 1$) was also identical across sessions. However, some parameters were specific to T2 sessions, namely the endogenous (varying) probability of audit p and the efficiency factor of the public good g . In T2, the audit probability p took values within interval $[0.1, 0.5]$, centered at $p = 0.3$, which was the exogenous (fixed) audit probability in T1. The efficiency factor was given the simple value $g = 1$, and the amount of taxes collected in the first place was re-distributed equally among the 20 taxpayers/subjects. The subjects were informed that the taxes and penalties further collected after audits would not be included in the public fund (as in the theoretical model).

The subjects were given comfortable time (10 minutes) to read the documents and call on us (privately) for some clarification if necessary, to be sure that both the tax system and payment rule were

well-understood. The subjects were informed that, at the end of the whole experiment, one of their Part 2 income reporting decisions would be selected at random to be played for real. We also emphasized the importance of making their tax decisions as they would do in real life: they were invited to really try to put themselves in the shoes of a taxpayer who has to fill out their tax return. After that, the subjects were asked to read carefully and follow the instructions on their computer, and to start making two trials before turning to real decisions. Finally, even though an explicit tax environment was introduced, we used neutral wording to avoid unwanted social norm effects. Typically, in French, the word which is commonly used in the medias to refer to the tax authority is “fisc”. But this word often has a pejorative flavour. So we rather chose to say “administration fiscale”, which is more neutral.

In each of their 2 or 3 experimental rounds, the subjects were invited to assess how much of their 1000-token income they were willing to report to the tax authority. In Round 1, the subjects made a simple income reporting decision, with no further information than that given in instruction sheets. In Rounds 2 and 3, information about the tax behaviour in Round 1 of either all, or a chosen subgroup of other taxpayers of the same session was provided/proposed, following the design presented in Subsection 4.1 and the timing of information delivery detailed in Table 1 in the same subsection. Each session ended with a questionnaire including standard socio-demographic items (e.g., gender, age, occupation, yearly income). The subjects were also asked whether they had already participated to an experimental study in economics, and whether they had already filled in a tax return in their life.

After that, the payment procedure was automatically implemented on each computer. At that point, each subject uncovered which of their tax decisions had been selected at random to be played for real, and whether they had been audited. The final earnings of each subject thus depended on their previous income reporting behaviour in the selected tax decision, as well as on the realization, or not, of the risk of audit based on its probability. Remember that in T2, the probability of audit was endogenous: for a given subject in a given tax decision, the actual probability of them being audited was impacted by both their income reporting behaviour and others’ average income reporting behaviour in that decision. It was calculated by the computer program based on the audit probability table shown in Figure 1 (Appendix 7.2). The final gain of each subject then appeared on the screen of their computer. The next step for them was to go to a separate room, one by one, to get their payment in cash. The average earning (including the show-up fee) was 18.46 euros.

The main characteristics of our subject pool are given in Table 9 (Appendix 7.2). Given the purpose of our study, it is important to note that 1) a huge majority of our subjects (87%) regularly earned a living, so they had already had the opportunity to fill an income tax return for real, 2) nearly half of our subjects were no longer students, they were much older than students and had long experience filling in income tax returns and paying their taxes. So, while obviously not representative of the French population, our subject pool was quite more diverse than a standard lab experimental one and looked fairly appropriate for our purpose.

5 The experimental results

As regards tax morale, our data show that for almost half of the subjects (45%), the answer to the tax-related question indicated full tax morale (score=7). A very low or null tax morale (score ≤ 3) was found for only 10% of the subject pool. Even though the possibility of self-image concerns among the subjects was not completely excluded, the experimental design was rather meant to minimize it. Moreover, complete anonymity precluded any social image and reputational concerns.

In Subsection 5.1, we consider the income reporting behaviour in Round 1 (the no-information baseline) and its determinants. The impact of social information on tax compliance will be examined in Subsections 5.2 ad 5.3.

5.1 Tax compliance in Round 1 and its determinants

Even though the subjects' tax decisions in Round 1 were primarily collected to provide a baseline as well as informational input for the sequel, they are worth investigating for themselves. First, the comparison in Round 1 of income reporting behaviour between tax environments T1 and T2 allows us to see whether subjects' interdependence in T2 affects their tax behaviour and in which direction. Second, our 240 subjects making the same income reporting decisions in the same "pure" experimental conditions provided us with a nice set of data, which allowed us to investigate the influence, in our study, of some basic determinants of tax compliance (namely here: tax morale, risk attitude and strategic interdependence between subjects), while controlling for a number of socio-demographic and experimental variables.

Decision setting and session heterogeneity

Before coming to the results, we take a detour to address the question of across-sessions and across-decision settings subject pool heterogeneity. Throughout Round 1, the experimental design was strictly the same in all experimental sessions of a given tax environment (being either T1 or T2). Therefore, conditionally on a random allocation of the participants to experimental groups, the distribution of income reporting behaviour in Round 1 should also be statistically the same across sessions run in T1 (resp. T2). However, random allocation to experimental groups could not be performed for obvious practical reasons: typically, participants could choose and sign up for any open slot, depending on their availability and preferences. Obviously, time preferences and constraints do not only differ between individuals, but also across categories of people: they are unlikely to be the same for students, administrative staff and researchers. So, we tried to reach a fairly equal representation of all the categories of participants in each experimental group by designing a time schedule that could suit them all. First, all the sessions were conducted on Tuesdays and Thursdays, either during the lunch break or at the end of the afternoon, which were the right times of the day for most people. Moreover, for each decision setting and in each tax environment, half the sessions were held during the lunch break, while the other half were scheduled at the end of the afternoon: even though heterogeneity could occur across experimental sessions depending on the moment of the day they were held, it should disappear at the decision setting level with the pooling of lunch time and afternoon data.

Basic non-parametric Kolmogorov-Smirnov tests run on the reported income distributions in Round

1 show that we partly attained our goal, with no significant differences in behaviour, either between decision settings C and D, or between decision settings A and B. However, while there is no significant difference in the distribution of tax behaviour between (pooled) C/D and B, a significant difference occurs between C/D and A (p -value = 0,008), with a much lower level of tax compliance in C/D than in A. Similar results were obtained using t -tests. Moreover, as shown in Table 3, the standard deviation of the reported income variable in Round 1 appears to be much larger in decision settings C and D than in A and B, which suggests that C and D groups were strongly polarized, including both highly tax compliant and highly tax evading subjects.

These results show that, despite our efforts, the composition of experimental groups was not really balanced. The main salient feature is the over-representation of Polytechniciens (the students of Ecole Polytechnique) in C and D experimental sessions (See Table 9 in Appendix 7.2). This feature is not innocuous in that Polytechniciens are a rather specific population with specific socio-demographic features, very high scientific skills, and also, for some of them, a specific mindset that may induce them to consider economic tasks in the lab as pure mathematical/financial games. In our experiment, they actually were more likely than other kinds of subjects to behave as in the standard portfolio-like standard tax model, thus to engage in tax evasion.³² Subject pool heterogeneity across decision settings is obviously a concern, in that it may introduce potentially confounding factors and biases in data analysis and interpretation. For this reason, caution will be the watchword in upcoming qualitative and statistical assessments, and a number of controls were actually introduced in econometric regressions to overcome this issue.

Differences in behaviour between tax environments T1 and T2

Now we come to the influence of the tax environment on tax compliance. Table 3 presents the descriptive characteristics of income reporting behaviour in Round 1 in each tax environment T1 and T2, depending on the decision setting. First, the reported income in T2 was 21% higher than in T1 on average. A t -test shows that this difference is highly significant (p -value=0.0013). Quite expectedly, this result indicates that the behavioural norm conveyed by the T2 tax environment is more conducive to tax compliance than that conveyed by T1 (all other things being equal). Moreover, the standard deviation was 11 % lower in T2 than in T1. This result supports the notion of higher behavioural homogeneity in T2 than in T1, which suggests that the strength of the T2 behavioural norm to induce a given behaviour was also stronger. However, as could be expected, there are some non-negligible variations across decisions settings. Non-parametric Mann-Whitney tests show significantly higher (at the 5 % level at least) reported income in T2 than in T1 in decision settings A, B, and D, while no difference was found in decision setting C, where the mean reported income was even lower in T2 than in T1 (p -value = 0.5). Here again, the standard deviation was higher in decision settings C and D, thereby indicating polarized behaviour within the corresponding experimental groups.

A second salient feature of our data set is the much higher number/proportion of “fully tax moral”

³²Quite interestingly, the proportion of fully tax moral subjects was not lower in C and D than in A and B (it was actually the same). Most Polytechniciens in our experiment actually reported a very high level of tax morale. However, for a number of them, being highly tax moral was not an obstacle to committing fraud in the context of the experiment. This disconnection between attitude to taxes and tax behaviour was not found in any other category of subjects.

Table 3: Income reporting behaviour in Round 1

Decision setting	T1			T2			T1 + T2		
	Obs	Mean x_i (Std. dev.)	Full TM (%)	Obs	Mean x_i (Std. dev.)	Full TM (%)	Obs	Mean x_i (Std. dev.)	Full TM (%)
A	40	722.00 (385.96)	19 (47.5%)	40	860.5 (183.90)	22 (55%)	80	791.25 (308.37)	41 (51.25%)
B	40	661.25 (330.79)	10 (25%)	40	776.25 (260.42)	21 (52.5%)	80	718.75 (301.41)	31 (38.75%)
C	20	645.00 (414.82)	20 (40%)	40	635.00 (302.22)	8 (40%)	40	640.00 (358.27)	16 (40%)
D	20	443.5 (442.96)	9 (45%)	20	767.5 (204.73)	11 (55%)	40	605.5 (378.06)	20 (50%)
(A+B+C+D)	120	642.5 (390.36)	46 (38.3%)	120	779.33 (345.62)	62 (51.7%)	240	710.92 (332.58)	108 (45%)

Note: Full TM: Fully tax moral subjects

subjects (score=7) in T2 than in T1 (62 vs. 46; 52% vs. 38%). Since tax morale was elicited at the outset of the experiment, before any reference to the tax environment, this difference in the rate of fully tax moral people across tax settings can only be due to the fact that, while having obviously no information about the T1 or T2 nature of the upcoming sessions, fully tax moral subjects actually signed-up more often for (planned-to-be) T2 session slots. The resulting higher level of tax morale in T2 unfortunately introduces a confounding factor in the descriptive analysis: (part of) the extra level of tax compliance observed in T2 might be explained by the subjects' higher level of tax morale, rather than by the very features of T2 as compared to T1. To examine this point further, we restrict our analysis to the fully tax moral subjects to compare their level of tax compliance in T1 and T2. The data in Table 10 (Appendix 7.3) show that the level of tax compliance within the fully tax moral subject pool remains higher in T2 than in T1, thereby suggesting that a genuine tax environment effect was actually at work. Regressions to come in next paragraph will help quantify the size of this effect, while controlling for potentially confounding factors such as tax morale. However, the current descriptive analysis already suggests that a tax environment that combines an endogenous probability of audit and a redistribution mechanism (as is actually the case in the real world) may contribute to enhance tax compliance as compared to a basic tax environment with no redistribution and exogenous audit probability, thereby bringing some preliminary support to Hypothesis **H1**.

The determinants of tax compliance

Now we turn to the econometric analysis of the determinants of tax compliance, based on the income reporting decisions made in Round 1 by the 240 subjects. The reported income in Round 1 was used as the explained variable, while tax morale, attitude towards risk and tax environment were introduced as explanatory variables, and socio-demographic features (sex, age, status and income) and decision setting as controls.³³ The detailed list of variables and their modalities is provided in Table 14 in

³³Tax environment and Sex were treated as dummies, while Tax morale, Risk Aversion, Age, Status and Income were

Appendix 7.4. Since the tests for heteroskedasticity were significant, the regressions were conducted with generalized least-squares (GLS) on various model specifications.

The results of the main regressions are presented in Table 4. Both the explained variable (the reported income in Round 1) and the coefficients are expressed in units of tokens (remember the income was the same for all the subjects and equal to 1000 tokens). The baseline is the behaviour, in T1 and Decision Setting A, of a male student from Ecole Polytechnique, aged between 18 and 24, who exhibits full tax morale (score=7) and risk neutrality. A number of results can be drawn from the regressions, with respect to both explanatory variables and controls.

First, in the baseline, around 2/3 of the income was reported, meaning that more than 30% of the income was evaded. Second, all other things being equal, being assigned to T2 (with redistribution and endogenous audit probability) increased the reported income by more than 110 tokens as compared to T1. Quite expectedly, the tax environment appears to have a highly significant (at the 0.01 level) and huge quantitative impact on the income reporting decision. This new result brings further support to **H1**, thereby confirming the key role played by redistribution and endogenous probability of audit in promoting overall tax compliance (at least in the lab but maybe also in the real world where both features are actually present). To put it differently, our results suggest that making taxpayers' tax decisions and outcomes properly interdependent, and making the taxpayers highly aware of that interdependence, may contribute to sustain overall tax compliance.

Third, tax morale appears to be a significant variable for tax compliance: the lower the subjects' level of tax morale, the stronger (up to 160 tokens) the decrease in their reported income as compared to the full tax morale baseline. By contrast, risk attitude did not significantly affect income reporting as compared to the risk neutrality baseline, except marginally (at the 0.1 level) for low risk averse subjects (RA3).³⁴ Therefore, **H3** appears to be strongly supported by our data, while **H2** is not.

Now, as regards the controls, a contrasting picture emerges as regards the influence of the decision setting on income reporting behaviour. Being assigned to any other decision setting than A decreases the level of tax compliance as compared to that in A. The difference does not reach significance for Decision setting B, while in C the level of tax compliance is significantly and markedly lower than in A (by around 115 tokens) and in D the difference is even more significant and larger (around 165 tokens, which corresponds to a 25% drop in tax compliance as compared to A). Obviously, the decision setting should not affect tax behaviour at this stage, since the experimental design was strictly identical across decisions settings. The fact that it actually does confirms the failure of our subjects' allocation process and underlines the need to circumvent the problem when it comes to analysing Round 2 and Round 3 decisions in Subsections 5.2 and 5.3.

Finally, as regards socio-demographic controls, the influence of age and sex on tax compliance appears to be in line with both our expectations and existing empirical results (see Alm and Malézieux, 2021 for instance). First, the data show that women complied significantly more than men, with an extra level of reported income up to 120 tokens. Second, the older the subjects, the higher their reported income: as compared to the 18-24 baseline, the increase in reported income is significant, massive, and increases with age. The reported income may be up to 222 tokens higher in the +55 group than in the 18-24

treated as categorical variables.

³⁴Risk attitude is usually found to affect tax compliance (see for instance Alm and Malézieux, 2021).

Table 4: The determinants of income reported behaviour in Round 1

	Reported income in Round 1			
	(1)	(2)	(3)	(4)
Constant	610.4*** (63.29)	688.1*** (50.93)	682.4*** (51.26)	708.2*** (54.74)
T2	120.4*** (38.59)	118.4*** (38.76)	122.8*** (38.72)	110.8*** (39.03)
TAXMORAL2	-123.6*** (41.34)	-131.3*** (41.09)	-128.9*** (40.90)	-126.1*** (40.71)
TAXMORAL3	-151.7** (65.99)	-161.4** (65.71)	-159.0** (65.60)	-143.5** (65.19)
RA2	123.5 (76.00)			
RA3	93.37* (51.85)			
RA4	91.16 (65.20)			
RA5	69.48 (60.86)			
DS B	-57.37 (46.98)	-68.05 (46.84)	-76.26 (47.34)	-57.83 (46.47)
DS C	-123.9** (57.44)	-118.0** (57.16)	-114.9** (56.87)	-115.7** (56.66)
DS D	-173.6*** (56.43)	-171.4*** (56.88)	-167.6*** (56.52)	-165.7*** (56.46)
WOMEN	96.78** (43.90)	102.0** (42.86)	110.9** (44.95)	120.0*** (43.42)
AGE2534	117.2** (48.40)	106.3** (47.89)	135.8 (87.53)	140.1** (62.36)
AGE3554	143.7** (58.47)	132.5** (56.98)	208.4** (98.19)	143.7** (71.61)
AGE55	151.7** (70.69)	147.4** (67.47)	222.3** (104.4)	138.3* (81.05)
STATUS2			54.13 (88.16)	
STATUS3			-80.56 (104.7)	
STATUS4			-76.10 (97.55)	
INCOME2				-102.3* (57.78)
INCOME3				13.53 (77.97)
INCOME4				-28.88 (87.26)
Observations	240	240	240	240

Notes: The table presents GLS estimations. Variables and their modalities are detailed in Appendix 7.4. Different combinations of control variables are used in the regressions presented in Columns (1) to (4). Standard errors are in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

baseline, which corresponds to a huge increase of more than one-third in the level of tax compliance. In contrast, neither Income nor Status appears to significantly impact the income reporting decision.³⁵ The likely reason for this is that income and status are largely correlated with each other, and each of them is also strongly correlated with age in our subject pool. In upcoming analyses involving Round 2 and 3 income reporting decisions, only sex and age will be kept as socio-demographic controls.

5.2 Tax compliance in Rounds 2 and 3: a descriptive analysis

In this subsection, we conduct a descriptive analysis of the subjects' income reporting behaviour after getting information about others' tax behaviour, with a special focus on attitude towards information as well as on the behavioural response to information and social (anti-)conformity hypothesis.

Overall picture of the data

To begin with, we consider income reporting behaviour in each round, as well as across rounds, differentiating the results with respect to decision setting and tax environment.

Basic descriptive data are provided in Table 5. The mean and standard deviation of the reported income variable in each round are reported in Column (1), differentiating for decision and tax setting. The number (and percentage) of subjects who changed their income reporting behaviour from round to round is given in column (2), while column (3) and (4) report the number of subjects whose reported income respectively increased and decreased, as well as the mean variation in the reported income. Three remarks should help read the table. First, in Round 2, the subjects from Decision settings C and D had already received both WGI and CRGI pieces of information, while subjects in A and B had only received one (WGI and CRGI respectively), so comparisons of income reporting behaviour between A/B and C/D in Round 2 should be made with caution. Second, the C and D subjects were no longer in the subject pool in round 3, which precludes any direct comparison of aggregate data from Round 2 to Round 3. In particular, the apparent increase in reported income observed between Round 2 and 3 on aggregate data is a pure artefact, due to the withdrawing of C and D subjects, whose specific tax behaviour pulled down the overall level of tax compliance in Rounds 1 and 2. Third, since the subject pool in tax settings T1 and T2 changed between Rounds 2 and 3, no direct comparison of T1 (resp. T2) data across the three rounds could be conducted. To circumvent this problem, the subject pool in each tax setting was restricted to Decision settings A and B (denoted *DS A&B* in Table 5), and between-rounds comparisons were carried out over this restricted data base.³⁶

Now, at first sight, a general result that can be drawn from the data in column (1) is the apparently null or quasi null average effect of information on income reporting behaviour, be it at the aggregate level, or at the level of each decision setting or tax setting. However, a further look at columns (2) to (4) shows that this neutral result actually hides a wide variety of responses at the individual level. First, regardless of the data aggregation level, the data in column (2) show that a majority of subjects did not change their behaviour in response to information about others' tax behaviour (be it WGI,

³⁵An exception is income category 9700-26700 (euros per year) which induces a decrease in reported income by 100 tokens as compared to the very low income baseline.

³⁶In Round 3, only A and B subjects were still present in the subject pool, so the data for T1 and T2 were by construction reduced to A and B subjects.

Table 5: Income reporting behaviour in Rounds 1, 2 and 3

	Obs	Reported income	Variation in reported income among reactive subjects		
		x_i Mean (Std. dev.) (1)	Δx_i Nb of subjects (% of subjects) (2)	Δx_i^+ Nb of subjects (mean Δx_i^+) (3)	Δx_i^- Nb of subjects (mean Δx_i^-) (4)
Round 1	240	710.92 (332.58)			
<i>T1</i>	120	642.5 (390.36)			
-- > <i>DS A&B</i>	80	691.625 (358.46)			
<i>T2</i>	120	779.33 (345.62)			
-- > <i>DS A&B</i>	80	818.375 (227.98)			
<i>DS A</i>	80	791.25 (308.37)			
<i>DS B</i>	80	718.75 (301.41)			
<i>DS C</i>	40	640 (358.27)			
<i>DS D</i>	40	605.5 (378.06)			
Round 2	240	712.45 (324.95)	109 (45.42%)	58 (+171.7)	51 (-188)
<i>T1</i>	120	637.4 (381.57)	48 (40%)	23 (+207.78)	25 (-215.64)
-- > <i>DS A&B</i>	80	696.84 (353.48)	32 (40%)	17 (+219.88)	15 (-221.4)
<i>T2</i>	120	787.5 (234.76)	61 (50.83%)	35 (+148)	26 (-161.54)
-- > <i>DS A&B</i>	80	831 (228.24)	34 (42.5%)	23 (+137.39)	11 (-195.45)
<i>DS A</i>	80	815.21 (301.51)	28 (35%)	19 (+166.74)	9 (-139)
<i>DS B</i>	80	712.62 (299.86)	38 (47.5%)	21 (+177.62)	17 (-248.23)
<i>DS C</i>	40	640.75 (340.00)	20 (50%)	8 (+197.5)	12 (-129.2)
<i>DS D</i>	40	609.51 (342.34)	24 (57.5%)	11 (+143.73)	13 (-197.69)
Round 3	160	776.32 (282.28)	64 (40%)	28 (+254.1)	36 (-142.5)
<i>T1</i> -- > <i>DS A&B</i>	80	696.46 (336.94)	30 (37.5%)	11 (+242.06)	19 (-159.47)
<i>T2</i> -- > <i>DS A&B</i>	80	856.19 (184.2)	34 (42.5%)	17 (+242.06)	17 (-123.53)
<i>DS A</i>	80	818.15 (277.67)	25 (31.25%)	11 (+215)	14 (-152.14)
<i>DS B</i>	80	734.5 (282.36)	39 (48.75%)	17 (+279.41)	22 (-163.36)

Note: Δx_i is the variation in reported income between rounds (in tokens) and mean Δx_i^+ (Δx_i^-) corresponds to the increase (decrease) in reported income.

CRGI, or both). From now on, these subjects will be called non reactive, while the subjects who changed their reported income between successive rounds after learning about others' tax behaviour will be called reactive. A second feature in columns (3) and (4) is that upward and downward changes in behaviour among reactive subjects were rather massive, while almost compensating each other on aggregate and thereby giving the false overall picture that nothing happened at all. A first likely reason for this behavioural heterogeneity is that, by design, each subject received individually-determined, thus specific, pieces of information: the diversity of informational stimuli called for a diversity of behavioural reactions. The second likely reason for the heterogeneity of behavioural reactions is the large diversity of the subjects in terms of tax morale, risk attitude and socio-demographics, thanks to which a given stimulus may entail a variety of (possibly opposite) behavioural responses. To get a better picture of what is really happening, we now turn to the descriptive analysis of income reporting behaviour at the level of decision setting and tax environment.

Tax behaviour in Rounds 2 and 3 depending on the decision setting

To begin with, we focus on within- and between-decision setting income reporting decisions. As regards Decision settings A and B, the data in Table 5, column (1), show that, from Round 1 to Round 2, reported income did not change on average in B, while increasing by 3% in A. From Round 2 to Round 3, the situation is symmetrical to the previous one: average reported income did not change in A, while increasing by 3% in B. Put together, these results might suggest that the kind of information matters and WGI has more impact on tax compliance than CRGI. This conjecture will be put to econometric test in next Subsection 5.3. However, here again, we need to go beyond aggregate results. First, less than half of the subjects actually changed their behaviour from round to round (and even only 35% in A from Round 1 to Round 2). Second, among reactive subjects, behavioural responses are diverse, with no clear pattern either by round, decision setting or tax environment.³⁷

Now, as regards Decision settings C and D, in which both kinds of information WGI and CRGI were provided together in Round 2, the average reported income remained quasi constant between Round 1 and (last) Round 2.³⁸ Going beyond aggregate data, we observe that only half of the subjects changed their reported income after getting information (20 out of 40 in C and 24 out of 40 in D), and among these reactive subjects, huge behavioural heterogeneity prevails: in Decision setting C, the subjects who decreased their reported income were in a majority, but the absolute value of that decrease (by -129 tokens on average) was lower than the value of the increase in reported income (by $+197$ tokens on average) observed among other subjects. The situation was nearly opposite in

³⁷For instance, between Rounds 1 and 2, twice as many subjects increased their reported income as decreased it (19 against 9) in A, and the average increase was 25 tokens higher than the absolute value of the average decrease. By contrast, nearly as many subjects decreased their reported income as increased it in B, but the absolute value of the average decrease was 75 tokens higher than the average increase. Things look rather different between Rounds 2 and 3: in A, 27% more subjects decreased their reported income as increased it (14 against 11), but the average increase was 63 tokens higher than the absolute value of the average decrease. Similarly, 30% more subjects decreased their reported income than increased it in B, but the absolute value of the average increase was 116 tokens higher than the average decrease.

³⁸Note that tax behaviour in C and D was characterized by a low average level of tax compliance in Round 1. Since informational stimuli were based on Round 1 tax behaviour within the groups, learning that others around evaded taxes too may have contributed to a dynamics that could barely escape a kind of "tax evasion trap".

Decision setting D (with average variations of about +144 against about -198).

Another variable of interest here is the evolution of the standard deviation of the reported income distribution throughout the rounds. Typically, a decrease in that standard deviation over rounds indicates that the gap between the subjects' reported income is narrowing, thus that their income reporting decisions are somewhat converging. In our experimental framework, a decrease in standard deviation may be further interpreted as a clue that the social norm conveyed by social information about others' tax behaviour has some effectiveness in driving tax behaviour. However, our data in column (1) of Table 5 show no spectacular drop in the standard deviation value across rounds, in either decision setting. A 10% (resp. 6%) decrease occurred in A (resp. B) between Rounds 1 and 3, but most of it took place between Rounds 2 and 3. In Decision settings C and (even more) D, the standard deviation was initially much higher than in A and B, probably due to the bi-modal polarization of behaviour in these experimental groups; its evolution from Round 1 to last Round 2 is characterized by a 9.5% (resp. 5%) decrease in D (resp. C), thereby suggesting that social information had a slight influence to overcome polarization and induce a tenuous behavioural convergence process.

Tax behaviour in Rounds 2 and 3 depending on the tax setting

Now, we turn to the descriptive analysis of within- and between- tax environment income reporting decisions. Several observations deserve our attention. First, as expected, the absolute level of tax compliance remained much lower in T1 than in T2 throughout the rounds (see column (1) in Table 5): the combined introduction of redistribution and endogenous probability of audit seemed powerful enough to increase the level of tax compliance in T2 by more than 20% as compared to T1. Second, the mean reported income did not significantly change from Round 1 to Round 2, either in T1 or in T2, with less than half of reactive subjects, and opposite variations in reported income, of similar magnitude, among them. To go a step further and examine the change in reported income from Round 1 to Round 3, we need to restrict our analysis to the subjects from Decision settings A and B. What we observe is that the reported income remained remarkably constant through rounds in T1, while constantly increasing in T2 (though not spectacularly). In other words, the gap between the level of tax compliance in T1 and that in T2 increased through rounds (by 26%, from 127 to 160 tokens). This result supports the notion that not only the baseline level of tax compliance in Round 1, but also the impact of social information on the variations in reported income in Rounds 2 and 3, are higher in T2 than in T1. This result thus brings some empirical support to our expectation that the amount of interdependence between subjects (introduced in T2 through both an endogenous probability of audit and a redistribution mechanism) might leverage the influence that social information about others' tax behaviour is likely to have on the subjects' subsequent tax decisions.

Finally, as regards the standard deviation of the reported income distribution in T1 and T2, several observations can be drawn from the (restricted to A and B) data. First, the standard deviation is at least 35% lower in T2 than in T1. This large gap may be due to differences in the behavioural norm conveyed by each tax environment.³⁹ Second, the standard deviation decreased from Round 1

³⁹Typically, the amount of between-subjects interdependence in T2 does not only increase the subjects' financial incentive to comply with taxes, but it also involves an implicit, both descriptive and injunctive, behavioural norm towards tax compliance, which thus has some effectiveness to guide and frame the majority's behaviour. Conversely, basic tax setting T1 looks like a pure mathematical game and does not involve any implicit social guide to behaviour. A

to Round 3 in both T1 and T2, but to a much larger extent in T2 (by 19% against 6% in T1). Social information about others' average tax behaviour induced behaviour to partly converge, both in T1 and T2, but it seems that, here again, the amount of interdependence in T2 leveraged the influence of information, thereby accelerating behavioural convergence. The possibility of a positive interaction between the specific features of tax setting T2 and the influence of social information is rather supported at this stage by the widening gap (from 36% in Round 1 to 45% in Round 3) in the standard deviation between T1 and T2.⁴⁰

Subjects' attitude towards chosen reference group information

We now turn our attention to the way the subjects dealt with chosen reference group information (CRGI). Two kinds of behaviour are worth investigating: first, the deliberate decision not to get any (free) information, and second, the decision to choose one piece of social information over another, which may give us useful insights into the notion of 'others of interest' in our experiment, and especially into the alleged preference for peers. The detailed features of the subjects' informational choices are given in Table 11 and Table 12 in Appendix 7.3, which respectively focus on the individual characteristics of the subjects depending on their informational choices (CRGI), and on the informational choice of the subjects depending on their individual characteristics.

The first noteworthy observation is that a huge majority of the subjects seized the opportunity to get (free) information: only roughly 10% of them (25 out of 240) chose not to get any information at all. This small no-info group is somewhat different from the whole subject-pool, being both older and wealthier, with a strong over-representation of both women and inconsistent (in the risk task) subjects, and an under-representation of fully tax moral subjects. Interestingly, the proportion of Polytechniciens in this group was only 20%, against 40% in the whole subject-pool. Also, more than half (56%) of these no-info subjects belonged to the highly tax moral group, while it only represented 44.58% of the subject pool.⁴¹

Though small, the no-info subject sample allowed us to test our hypothesis H4d that individuals who are not interested in getting information about others' tax behaviour should not change their behaviour when making their next income reporting decision. A Wilcoxon test was conducted to compare the income reporting decision made by the subjects after choosing not to get information and that made in the previous round. Only the 23 subjects from decision settings A and B could be used (in C and D, WGI and CRGI were provided together, so their respective impact could not be isolated). Moreover, 2 subjects were discarded because they made erratic decisions (with a variation in reported income of more than 900 tokens). While keeping in mind the small size of that subject sample, the data suggest that there was no significant change in subsequent tax behaviour among those who were not interested in getting information, thereby bringing some cautious support to hypothesis **H4d**.⁴²

The second noticeable result is that only 1/3 (33.75%) of the subjects asked for information about

number of subjects will simply solve the mathematical problem, but others may call upon their own behavioural norm in order to reach a decision. From which a wide variety of income reporting decisions.

⁴⁰Econometric checks are in order to control for possibly confounding factors before drawing any conclusion.

⁴¹As further shown in Table 12, 13% of highly tax moral subjects were not interested in information, against only 8% for both fully and slightly tax moral subjects.

⁴²When keeping the 23 subjects, the difference becomes marginally significant at the 10% level.

the tax behaviour of subjects who had the same level of tax morale as themselves, that is, their peers in the context of the experiment. This result suggests that, when given the possibility to choose their others of interest, it may be that only a minority of people are primarily interested in the behaviour of similar others. This result somewhat qualifies the basic belief that social interest should primarily stem from proximity and similarity. Moreover, in our study, the peer-oriented group was somewhat different from the whole subject-pool: it included a higher proportion of women and fully tax moral subjects, it was older and wealthier, with a higher proportion of administrative staff and a lower proportion of professors/researchers (but roughly the same proportion of Polytechniciens).

The third salient result here is that 56% of the subjects actually chose to get information about different others (non-peers), and among them, more than 71% (about 40% of the whole subject pool) chose to get information about the behaviour of less tax moral others.⁴³ As regards the socio-demographic pattern behind the decision to choose the lower tax morale reference group, the only difference between this group of subjects and the whole subject pool is the over-representation of men and the under-representation of administrative staff and subjects over 55. Interestingly, fully tax morale subjects were not over-represented in that subgroup: even though they only had two really open informational options (since they knew that the higher tax morale reference group would be empty in their case), it seems that they did not make a transfer onto the lower-tax morale option.⁴⁴

Several additional observations can be drawn from the subjects' informational choices depending on their individual characteristics (provided in Table 12 in Appendix 7.3). As regards the no-info option (chosen by 10.4% of the subjects), the most over-represented categories are subjects over 55, non-Polytechnicien students, researchers and related, as well as the subjects who gave inconsistent answers in the risk task. Conversely, the wealthiest subjects and Polytechniciens are strongly under-represented. As regards the decision to get information about their peers (that is, the reference group with the same level of tax morale, chosen by 33.75% of the whole subject pool), the most over-represented categories are the oldest and wealthiest subjects as well as, to a lower extent, the fully tax moral subjects, the administrative staff and the subjects who gave inconsistent answers in the risk task. Conversely, risk neutral subjects, researchers and related and, to a lower extent, high tax morale subjects were strongly under-represented. In contrast, as regards the decision (made by 16.25% of the whole subject pool) to choose the higher tax morale reference group, the strongly over-represented category were the low tax morale subjects, while fully tax morale and strongly risk averse subjects were heavily under-represented. Finally, the choice of the lower tax morale informational option (made

⁴³Two general motives can be put forward to explain why people may be interested in the behaviour of less tax moral subjects. First, since low tax morale subjects can be expected to evade taxes more than average, learning that they actually do may provide the individual with a comfortable self-justification for cheating in their turn. Second, a basic curiosity motive may be at work: one may simply be curious about the actual behaviour of those who have been identified with lower tax morale standards.

⁴⁴As compared to other categories of subjects, fully tax morale subjects could not really choose the higher tax morale option (even though some did). However, they massively preferred the same-tax morale option to the lower-tax morale one. In this respect, it is worth comparing the informational strategy of highly tax moral subjects with that of fully tax moral subjects. Indeed, the groups have roughly the same size, but highly tax moral subjects could choose between the three informational options while fully tax morale subjects could only choose between the lower tax morale and same tax morale options. Interestingly, the proportion/number of lower tax morale choices was roughly the same in both groups. While the remaining part of fully tax moral subjects chose the same tax morale option, highly tax moral subjects roughly evenly split between the same tax morale and higher tax morale options.

by 39.58% of the whole subject pool) was heavily over-represented among risk neutral subjects as well as among researchers and related, and strongly under-represented among the oldest subjects, low tax morale subjects and those who gave inconsistent answers in the risk task.

Non reactive subjects and behavioural persistence

A population of particular interest to us are the 108 (out of 240, thus 45%) non reactive subjects who constantly reported the same income throughout the 2 or 3 rounds, thereby showing no sensitivity to others' tax behaviour and no taste for social (anti-)conformity as we defined it.

Within this population, two subgroups deserve some specific attention: the super compliers, who constantly reported 1000 tokens, and the super evaders, who constantly reported less than 500 tokens.

Sixty percent of the non reactive subjects (65 out of 108), thus more than a quarter (27.1%) of the whole subject pool, were actually super compliers, who reported their entire income either twice or thrice depending on their decision setting. This remarkable behaviour strongly suggests that these subjects made their income reporting decisions based on their own behavioural rule, thereby paying no attention to the way others around behaved, and being not influenced by the descriptive social norm embedded in others' average tax behaviour.⁴⁵ An interesting point is that these subjects were still interested in getting information about others' tax behaviour (only 6 out of 65 did not ask for CRGI), which tends to support the notion of a pure curiosity motive (see Reck et al., 2022 for instance). Among the 59 subjects who asked for information, 31 were interested in the behaviour of their peers (i.e., the subgroup of subjects with the same level of tax morale), and 28 in the behaviour of non peers, which brings support to our starting idea that others of interest are not necessarily peers as usually assumed.⁴⁶ The detailed characteristics of super compliers are given in Table 13 in Appendix 7.3. As compared to the whole subject-pool, these subjects were more often fully tax moral (68% subjects against 45% in the whole subject pool) and strongly risk averse. Were also over-represented the 25-34 and the women. Conversely, the 18-24 and high tax morale individuals were strongly under-represented.

On the other hand, 26 (thus 24%) of the 108 non reactive subjects are super evaders who reported less than 500 tokens either twice or thrice depending on their decision setting.⁴⁷ Even though the population of super evaders was rather small here, their attitude towards information exhibits some

⁴⁵This strong behavioural rule is likely based on the societal injunctive norm, which makes full tax compliance the appropriate normative behaviour, based on civic duty and related notions.

⁴⁶Quite expectedly given their high own level of tax morale, most of these subjects (24 out of 28) asked for information about the subjects with a lower level of tax morale (the subgroup with higher tax morale would likely have been empty).

⁴⁷In Decision settings C and D, the notion of super evaders might be weakened by the fact that they only made 2 income reporting decisions (which is not much to study behavioural persistence). Moreover, they belonged to experimental groups in which the level of tax evasion was higher than average: their own persistently low level of tax compliance could be due to the single fact that they received information about people being as little compliant as they were, thereby having no reason to change their behaviour. Persistent tax evading behaviour among C/D subjects might thus be a pure experimental artefact. So, our claim that this behaviour resulted from a genuine preference for tax evasion needs to be substantiated. To that end, we isolated the subjects who constantly reported less than 500 tokens in C and D, and compared the level of reported income in both their whole and reference groups (that is, the pieces of information they received) and their first reported income. The data show that, for all of them, the level of tax compliance in their group was much superior to their first reported income. Had they been sensitive to information about others' tax behaviour, they would have reported a higher income in Round 2, which they did not. This is not because information did not call for a behavioural change that these subjects did not change their behaviour, but because they deliberately chose to do so and persistently evade taxes.

noticeable features. First, as super compliers, they seemed to exhibit a pure curiosity motive: 22 out of 26 subjects did ask for information while not taking it into account in their subsequent decisions. Second, our data show that super evaders were much more interested in the behaviour of non-peers' behaviour than in that of peers: only 5 subjects out of 22 asked for information about the subgroup of subjects with the same level of tax morale.

The characteristics of super evaders (displayed in Table 13 in Appendix 7.3) deserve some specific attention. As compared to the whole subject-pool, super evaders were much more often males (84.6% against 58.3% in the whole subject pool), very young people (76.9% belonged to the youngest 18-24 category, against 45.4% in the whole subject pool) and Polytechniciens (being 73.1% against 40% in the whole subject pool).⁴⁸ Were also heavily over-represented strong risk averse subjects (being 30.8% against 13.8% in the whole subject pool) and high tax morale subjects (being 65.4% against 44.6% in the whole subject pool), while fully tax moral subjects were strongly under-represented (being 23.1% against 45% in the whole subject pool). However, it is worth noticing that the proportion of high-at-least tax morale subjects among super evaders (88.4%) is roughly equal to the proportion of high-at-least tax morale subjects in the whole subject pool. The fact that, in our study, the level of tax morale appears not to be decisive at all in the decision to super evade is largely due to the high proportion of Polytechniciens in this group: while debriefing informally with us after the payment stage, several of them admitted that they had processed income reporting decisions as pure gambles, based on the very mathematical features of the decision problem, thereby not really taking the tax-related nature of the task into account.⁴⁹

Finally, a salient difference arises between the behavioural path of super compliers and that of super evaders. Among the 85 subjects who fully complied in the first round, 65 subjects (thus more than 75% of them) still did so in the last round (be it the second or third one). By contrast, among the 53 subjects who reported less than 500 tokens in Round 1, only 26 subjects (50%) were still super tax evaders in last round. So, in our experiment, super tax compliance seemed to be much more persistent than super tax evasion. Our interpretation of this difference involves the (different) nature of the underlying behavioural rule in each case. On the one hand, super tax compliance is likely sustained by a strong normative behavioural rule, in relation with the societal injunctive norm towards compliance. Super compliant behaviour is a cue that the injunctive rule is strong enough to overcome the descriptive social norm (conveyed by information about other subjects' tax behaviour) in case they contradict each other. On the other hand, strong tax evasion is likely based on a purely computational behavioural rule. It seems that, after learning about others around being more compliant than they are, it gets hard not to give in to the (both descriptive and injunctive) pressure to conform to others' average tax behaviour and comply with taxes.

Reactive subjects and the taste for social (anti-)conformity

The data provided in Table 6 show that 45% of the subjects (109 out of 240) were reactive in Round

⁴⁸A majority of super evaders were in either Decision setting C or D (14 out of 26, while there were twice fewer subjects in C/D than in A/B), and a huge majority of them were in T1 (22 out of 26).

⁴⁹In the field, the combination of actual cheating behaviour and high (self-assessed) tax morale is much more likely to emerge among people who are very concerned about their social image and claim to care about being honest while actually misbehaving. In our experiment, anonymity prevents social image concerns.

2. In Round 3 (which existed only for A and B subjects), this proportion slightly decreased, with 64 reactive subjects out of 160 (40%).⁵⁰ Nevertheless, the reactive subjects were not necessarily the same across rounds: among the 64 subjects (from decision settings A and B) who reported a different income in Round 3 than in Round 2, 23 (36%) had not changed their income reporting behaviour from Round 1 to Round 2.⁵¹ Conversely, among the 41 subjects (from Decisions settings A and B) who changed their reported income between Rounds 1 and 2, 25 exhibited a behavioural status quo between Rounds 2 and 3.

We now turn to the question of whether the reactive (to social information) subjects rather exhibited a taste for social conformity or for anti-conformity. From a behavioural standpoint, the taste for conformity leads the individual to adjust their tax behaviour towards others' tax behaviour. The taste for social anti-conformity is defined the opposite way, resulting in a behavioural adjustment away from others' tax behaviour. Social conformity (resp. anti-conformity) is thus expected to induce behavioural convergence (resp. divergence) among people.⁵² Typically, a subject is said to have a taste for social conformity if they report a lower (resp. higher) income after being informed that they previously reported a higher (resp. lower) income than others did. The first part of the proposition is reversed to define the taste for anti-conformity.

The question of whether our subjects behaved in an either conformist or anti-conformist way is roughly answered in Table 6 (left and right sides, in bold), where $INFO_t$ stands for the piece of information received at the outset of Round t , x_{it-1} the reported income in Round $t-1$, and $x_{it-1} - INFO_t$ the reported income gap in Round t .

Table 6: Taste for social (anti-)conformity among reactive subjects in Rounds 2 and 3

Round 2	N=109	$x_{i2} > x_{i1}$	$x_{i2} < x_{i1}$	Round 3	N=64	$x_{i3} > x_{i2}$	$x_{i3} < x_{i2}$
$x_{i1} > INFO_2$	55	13	42	$x_{i2} > INFO_3$	39	10	29
$x_{i1} \leq INFO_2$	54	45	9	$x_{i2} \leq INFO_3$	25	18	7

Note: x_{it} is reported income in Round t , $INFO_t$ is information received at the beginning of Round t .

Unsurprisingly in view of our previous results, most changes in behaviour between Round 1 and Round 2 (87 out of 109, 79.8%) and between Round 2 and Round 3 (47 out of 64, 73.4%) reflect a clear (albeit somewhat less marked in Round 3) pattern of behavioural adjustment towards others' tax behaviour. In both rounds, the proportion of the subjects whose decisions were consistent with the behavioural convergence pattern varies between 72% and 83% depending on the round and sign of their reported income gap.⁵³ For instance, when learning at the outset of Round 2 that they in Round 1 had reported a higher income than others' (i.e., $x_{i1} > INFO_2$), 42 subjects out of 55 (76.4%)

⁵⁰The last two figures were already given, but from a different perspective, in Table 5, column (2), supra.

⁵¹For half of these 23 subjects, others' reported income was very close to theirs, so they had no reason to change their behaviour in reaction to social information. Some other subjects could not obtain their requested information because their chosen reference group was actually empty, and they did not change their behaviour after that either.

⁵²In our theoretical framework, the taste for (anti-)conformity is technically captured through a non-zero parameter α_i and the direction of this taste is given by the sign of α_i . In our experimental study, this sign can actually be determined empirically, based on the direction of the subject's behavioural reaction (i.e., the sign of $x_{it} - x_{it-1}$) after learning about others' tax behaviour (i.e. the sign of the reported income gap $x_{it-1} - INFO_t$).

⁵³Independence between the sign of the reported income gap and subsequent behaviour is rejected at the 1% level, both in Rounds 2 and 3.

responded by lowering their own reported income in Round 2 ($x_{i2} < x_{i1}$). This response corresponds to the “fall guy” effect depicted in Section 3 and Subsection 4.2: when learning that others were less virtuous than they were, the taxpayer tends to feel like the fall guy, which provides them with a good reason to evade in their turn. Conversely, among the 54 subjects who initially reported a lower income than their fellow taxpayers ($x_{i1} < INFO_2$), 45 (83.3%) decided to report a higher income in the next round ($x_{i2} > x_{i1}$), in compliance with the “guilty feeling” effect: when learning that others were more virtuous than they were, the taxpayer tends to feel guilty, which prompts them to behave more honestly thereafter. In Round 3, the behavioural pattern is qualitatively roughly the same, even though the numbers are somewhat different.⁵⁴ Our results thus bring strong support to **H4b** and they also are in line with most previous findings about social conformity, behavioural contagion and related topics (see the many references in Section 2).

Yet, it is worth noting that, in both rounds, a significant minority (between 16.7% and 28%) of the subjects chose to react to others’ average tax behaviour by adjusting their behaviour away from that behaviour, becoming even more compliant (resp. evading even more) after learning about others being less (resp. more) compliant than they were. These two behavioural patterns are compatible with the taste for social anti-conformity hypothesis and refer to the “strong conviction” and “pure selfishness” effects respectively (described in Section 3 and Subsection 4.2). So, even though the taste for social conformity appears to prevail among our subjects, our data also bring some non negligible support to the social anti-conformity hypothesis **H4c**.

5.3 The influence of social information on reported income: an econometric analysis

To further investigate the influence of social information on income reporting behaviour, we now come to the econometric analysis of the data. We investigate the influence of social information on the variations in reported income from two perspectives, depending on either the type of information (regardless the round in which it was provided) or the round in which information was provided (regardless the type of information). Unless otherwise specified, the baseline in the regressions is, as before, the behaviour, in T1 and Decision setting A, of a risk neutral, fully tax moral, young male student from Ecole Polytechnique. Again we used a GLS method to cope with heteroskedastic residuals. All the regressions include Decision setting, Tax environment and the major individual variables (Sex, Age, and Tax morale, and Risk attitude) as controls. For the sake of clarity, individual variables are not reported in the presented tables.⁵⁵

⁵⁴Both the group size and the distribution of the subjects with respect to the sign of their reported income gap are different across the rounds, partly due to the withdrawing of C and D subjects from the subject-pool at the end of Round 2, which mechanically reduced the size of the subject pool and increased the overall level of tax compliance among the remaining subjects. While in Round 2 there were as many subjects with a positive reported income gap as with a negative one, the proportion of subjects with a positive reported income gap sharply increased in Round 3.)

⁵⁵Full regressions show no impact of sex and age as compared to the baseline. However, depending on the model specification, being strongly/weakly risk seeking, as well as being highly/slightly tax moral could significantly reduce the impact of information on the variation in reported income as compared to the risk neutral fully tax morale baseline.

The influence of information depending on its type

Let us start by examining the influence of social information on income reporting behaviour depending on its type (either WGI or CRGI). This analysis is restricted to the data from Decision settings A and B.⁵⁶ Econometric results are provided in Table 7.

In our model, the logic of the subject's income reporting decision in Round t is based on their behavioural reaction (captured through the variation in their reported income $\Delta x_{i,t}$) to the observation of a gap (called the reported income gap and denoted $\Delta INFO_{i,k,t}$) between others' reported income and theirs. To properly analyse the data, the technical question arises of how the subjects actually dealt with information \bar{X} , and which of their previous reported income they chose in practice as the benchmark against which to assess the reported income gap and the variation in reported income. We obviously are not able to answer this empirical question. However, two plausible alternative decision processes can be identified, depending on whether that benchmark would be the subject's first income reporting decision (thus their Round 1 reported income) or the income reporting decision they made just before the current one (thus, their previous round reported income). Called R1-DP in the following, the Round 1-based decision process is all the more plausible since information with respect to others' tax behaviour (denoted \bar{X}) was about their behaviour in Round 1. Following this logic, the subjects would make their Round t income reporting decision while assessing both the variation in reported income $\Delta x_{i,t}$ and the reported income gap $\Delta INFO_{i,k,t}$ with respect to the income $x_{i,1}$ they reported in Round 1. The fact that, in each round 2 and 3, the current decision would be made based on the decisions the subject, and others, made in Round 1 implies that Round 2 and Round 3 income reporting decisions would be made independently. In a way, our experimental design was meant to favour that kind of information and decision processing. However, the Previous Round-based decision process, called PR-DP in the following, is also quite plausible, because it may be intuitive to consider successive income reporting decisions as a sequence of interdependent decisions. Following that approach, the subjects would make their Round t income reporting decision while assessing both the variation in reported income and the reported income gap in Round 2 (resp. Round 3) with respect to their previous, thus Round 1 (resp. Round 2), reported income.

Although conceptually different, the two strategies obviously coincide in Round 2, where Round 1 and the previous round are the same. Therefore, only in Round 3 is the distinction between the two influential, which implies that only the data from Decision settings A and B can be used to discriminate between the two strategies.

To account for both decision processes, two sets of regressions were conducted, based on the two following specifications of the variation in reported income. The R1-DP (resp. PR-DP) is described in Eq. (3) (resp. Eq. (4)). For the sake of clarity, all technical details are provided in Appendix 7.5:

$$\mathbf{R1-DP} : \Delta x_{i,t} = x_{i,t} - x_{i,1} = \beta_0 + \beta_1 \Delta INFO_{i,k,t} + \beta_2 Y_i + \varepsilon_i \quad (3)$$

with $\Delta INFO_{i,k,t} = x_{i,1} - \bar{X}_k$.

⁵⁶C and D data could not be used here because both pieces of information WGI and CRGI were provided together and a unique income reporting decision was made after that, thereby making it impossible to isolate the impact of each type of information on reporting income behaviour.

$$\text{PR-DP} : \Delta x_{i,t} = x_{i,t} - x_{i,t-1} = \beta_0 + \beta_1 \Delta INFO_{i,k,t} + \beta_2 Y_i + \varepsilon_i \quad (4)$$

with $\Delta INFO_{i,k,t} = x_{i,t-1} - \bar{X}_k$.

Note that the sign of β_1 in the econometric model refers to the sign of α_i in our theoretical model and measures whether, on average, subjects had a taste for either social conformity or anti-conformity. Now, in both Eq. (3) and (4), the vector of control variables Y_i includes the usual variables as well as, specifically here, a variable Chosen Reference Group (introduced in CRGI regressions only) to control for the informational options chosen by the subjects depending on the relative level of tax morale of others. The modalities of this variable (namely: Same TM, Lower TM and Higher TM) are described in table 14 in 7.4.

Table 7: Post-information variations in reported income in Decision settings A and B depending on the information type (WGI/CRGI): full sample vs. subsample of reactive subjects

Type of information	R1-DP				PR-DP			
	Variations in reported income/Round 1				Variations in reported income/previous round			
	$\Delta x_{i,t} = \{x_{i,t} \bar{X}_k\} - x_{i,1}$				$\Delta x_{i,t} = \{x_{i,t} \bar{X}_k\} - x_{i,t-1}$			
	WGI ($k = whole$)		CRGI ($k = ref$)		WGI ($k = whole$)		CRGI ($k = ref$)	
Subject pool A+B	full	reactive	full	reactive	full	reactive	full	reactive
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta INFO_{whole}$	-0.274*** (0.0457)	-0.654*** (0.0840)			-0.205*** (0.0409)	-0.562*** (0.0927)		
$\Delta INFO_{ref}$			-0.185*** (0.0428)	-0.514*** (0.0979)			-0.151*** (0.0448)	-0.426*** (0.121)
Higher TM			-64.58 (40.13)	-10.22 (76.12)			-43.26 (40.54)	36.41 (119.2)
Lower TM			-32.78 (28.11)	9.188 (65.03)			-21.81 (28.11)	89.02 (80.23)
T2	18.72 (26.32)	-24.51 (49.43)	28.56 (25.17)	-10.94 (48.18)	21.30 (23.05)	6.973 (50.09)	16.84 (25.33)	-9.345 (57.41)
DS A			47.01* (24.28)	125.7*** (47.75)			23.88 (24.34)	70.71 (65.85)
DS B	-13.89 (26.29)	-4.640 (46.24)			-3.668 (23.00)	9.436 (47.11)		
Observations	160	76	127	55	160	67	127	47

Note: The table presents GLS estimations. All the regressions further include the following control variables: Sex, Age, Tax morale, and Risk attitude. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The results obtained based on R1-DP are presented in Columns (1) to (4) of Table 7, while those obtained with PR-DP are given in columns (5) to (8). In each set of results, the first (resp. last) two columns pertain to the WGI (resp. CRGI) type of information. Because WGI (resp. CRGI) was provided in Decision setting A (resp. B) first, regressions pertaining to WGI (resp. CRGI) were conducted using A (resp. B) as the baseline to enable us to capture the pure WGI (resp. CRGI) effect. Finally, each regression was conducted both over the whole A and B subject-pool and over the

subsample of reactive subjects (whose tax behaviour changed across rounds).⁵⁷ The varying size of the subject pool across columns is also indicated at the bottom of Table 7.⁵⁸

In columns (1) to (4), the regressions conducted based on R1-DP show a highly significant effect of both kinds of social information on the change in reported income relative to Round 1. The negative sign of both $\Delta INFO_{whole}$ and $\Delta INFO_{ref}$ coefficients indicates that, on average, the subjects increased (decreased) their level of tax compliance (as compared to Round 1) after learning that they had reported a lower (higher) income than average in the (whole/reference) group, reflecting a behavioural convergence pattern of behaviour. Moreover, the absolute value of the coefficients associated to variables $\Delta INFO$ reflects the intensity of the convergence. When the analysis is conducted over the whole A and B subject pool (which included about 50% of non reactive subjects), the values of the behavioural adjustment towards others' average tax behaviour remain moderate: 18.5% for CRGI and 27.4% for WGI). These values can be viewed as reasonable proxies of the average behavioural response within the standard population (which actually includes a number of non reactive subjects). When the analysis is restricted to the group of reactive subjects, the intensity of behavioural convergence reaches remarkably high values (51.4% for CRGI and 65.4% for WGI), which obviously cannot be extrapolated to any population. Anyway, this series of results brings clear and strong additional support to Hypothesis **H4b**.

Besides, thanks to the reversal of the decision setting baseline between WGI and CRGI regressions, the comparison between $\Delta INFO_{whole}$ and $\Delta INFO_{ref}$ coefficients provides direct information about the respective influence of WGI and CRGI when provided first (in Round 2). Since the (absolute) value of information coefficients is much higher in the WGI case than in the CRGI case, it can be drawn from our data that WGI had a much stronger impact towards behavioural convergence than CRGI: a positive difference of 100 tokens between the subjects' reported income and the income reported on average in their group implied a decrease in the subjects' reported income by 27 tokens in the WGI case, against only 18.5 tokens in the CRGI case. Moreover, when the subject sample is further restricted to reactive subjects, the intensity of behavioural reactions quite logically increases, but the flavour of the results remains the same: WGI appears to have more impact towards behavioural convergence than CRGI, with a decrease in reported income by nearly 65 tokens in the former case, against 51.4 tokens in the latter.⁵⁹

The results of the regressions conducted using PR-DP are shown in columns (5) to (8). Quite expectedly, the quantitative impact of reported income gaps on the variation in reported income (captured through the $\Delta INFO$ coefficients) is somewhat smaller in the PR-DP case.⁶⁰ However, the results

⁵⁷Remember that at least 50% of the subjects did not change their behaviour throughout the experiment. Focusing over the subgroup of reactive subjects is thus a way to get a more concentrate picture of the data.

⁵⁸Note that the samples are reduced due to the fact that some subjects did not get any CRGI, be it because they refused to get information or because they could not get information since their reference group of interest was empty. For instance, the subject pool in column (3) reduces to 198 subjects (instead of 240) because 25 subjects did not request any information and 17 had an empty chosen reference group.

⁵⁹As regards the controls, neither the tax environment nor the reference group chosen appears to have significant impact, while Decision setting A appears to have a significant impact (as compared to the baseline B) in both CRGI regressions.

⁶⁰In Round 2, R1-DP and PR-DP coincide. In Round 3, the variations computed with PR-DP are smaller than those computed with R1-DP since they involve a one-round gap instead of a two-round gap, from which less intense effects.

are qualitatively similar to those obtained with R1-DP. In particular, WGI again appears to have had more behavioural influence towards convergence than CRGI and, again, the influence of both WGI and CRGI is higher when assessed over the reactive subsample.⁶¹ The fact that the findings are robust to the change in the decision benchmark is rather reassuring since we do not know which theoretical decision process is closest to the subjects' empirical one.

The influence of information depending on the round

In this subsection we focus on the impact of information depending on the round at the outset of which information was provided, regardless of the type of information. The results of the regressions are presented in Table 8.

The regressions in columns (1) to (4) focus on behavioural changes between Rounds 1 and 2, while the regressions in columns (5) to (7) focus on behavioural changes between Rounds 2 and 3. In all the regressions, the general dependent variable is the variation in reported income between two successive rounds $t - 1$ and t , denoted as usual $\Delta x_{i,t,t-1} = x_{i,t} - x_{i,t-1}$. In the regressions presented in columns (2) and (5), the basic explanatory variable is the reported income gap $\Delta INFO_t = x_{i,t-1} - \bar{X}_k$, where \bar{X}_k denotes the piece of information about the average behaviour in group k (which may be either the whole or the reference group) received at the beginning of round t . The regressions in columns (1) to (4) of Table 8 thus involve $\Delta x_{i,2,1}$ and $\Delta INFO_2$, while the regressions in columns (5) to (7) involve $\Delta x_{i,3,2}$ and $\Delta INFO_3$. Since Round 3 only exists for Decision settings A and B, most of the analysis is restricted to Decision settings A and B for the sake of comparison. The only regression that also includes C and D subjects is presented in column (1), while the same regression restricted to Decision settings A and B is shown in column (2).

Obviously, the reported income gap may be either positive or negative, depending on whether the subject reported a higher or a lower income than others in the previous period. Within the behavioural convergence pattern, the question arises whether, for a given absolute value of the reported income gap, the magnitude of the behavioural adjustment towards others' average tax behaviour will be the same. A number of converging experimental findings (see Dimant (2019) for instance) suggest that it may not be the case, with downward behavioural adjustments (towards less tax compliance) being of greater magnitude than upward behavioural adjustments (towards more tax compliance). So, to further analyse the reactions to information within our subject pool, we split the variable $\Delta INFO_t$ into variables $\Delta INFO_t > 0$ and $\Delta INFO_t < 0$ depending on the sign of the reported income gap.⁶² We then used the new variables as explanatory variables in separate regressions over the corresponding restricted groups of subjects. The results are presented in columns (3)-(4) and (6)-(7).

In all the regressions, the baseline is the usual one (including Decision setting A). As before, tax environment, decision setting and all individual variables were controlled for (only the first two are reported in the table). The varying size of the subject pool across columns is indicated at the bottom of Table 8.

The results obtained over the whole subject pool in Round 2, displayed in column (1), show that

⁶¹As shown in columns (7) and (8), the only qualitative difference with R1-DP is that belonging to Decision setting A instead of B does no longer affect the variation in reported income with respect to previous round.

⁶²Variable $\Delta INFO_t > 0$ (resp. $\Delta INFO_t < 0$) is thus simply the positive (resp. negative) part of $\Delta INFO_t$

social information about others' average tax behaviour ($\Delta INFO_2$) had a highly significant effect on reported income. Moreover, since the coefficient of $\Delta INFO_2$ is negative, the results are consistent with the taste for social conformity hypothesis **H4b** and the behavioural convergence pattern: in the baseline (thus in Decision setting A and tax environment T1), being more compliant than others by 100 tokens in Round 1 resulted in a 17.6 tokens decrease in the reported income from Round 1 to Round 2.⁶³

Table 8: Post-information variations in reported income depending on the round

	Full sample	Decision settings A and B					
	$\Delta x_{i,21}$ $\Delta INFO_2$ (1)	$\Delta x_{i,21}$			$\Delta x_{i,32}$		
		$\Delta INFO_2$ (2)	$\Delta INFO_2 > 0$ (3)	$\Delta INFO_2 < 0$ (4)	$\Delta INFO_3$ (5)	$\Delta INFO_3 > 0$ (6)	$\Delta INFO_3 < 0$ (7)
$\Delta INFO_2$	-0.176*** (0.03)	-0.147*** (0.04)					
$\Delta INFO_2 > 0$			-0.211* (0.12)				
$\Delta INFO_2 < 0$				0.0278 (0.10)			
$\Delta INFO_3$					-0.220*** (0.05)		
$\Delta INFO_3 > 0$						-0.063 (0.09)	
$\Delta INFO_3 < 0$							-0.274** (0.13)
T2	13.84 (18.95)	11.57 (24.06)	-13.76 (31.71)	-14.77 (50.43)	22.96 (26.09)	15.85 (23.22)	55.55 (66.09)
DS B	-43.64* (23.43)	-49.47** (23.87)	-50.67* (30.22)	-52.54 (45.31)	10.50 (26.06)	-37.53 (23.09)	91.55 (64.05)
DS C	-12.55 (27.06)						
DS D	-48.38* (26.50)						
Observations	223	143	85	57	144	93	51

Note: The table presents GLS estimations. All the regressions further include the following control variables: Sex, Age, Tax morale, and Risk attitude. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. One subject disappeared from the subject pool between column (2) and (3)+(4) because they had a null $\Delta INFO_2$.

When splitting the A and B sample with respect to the sign of $\Delta INFO_2$, the results shown in columns (3) and (4) are very contrasting. Among the subjects who reported a higher income than others in Round 1 ($\Delta INFO_2 > 0$), the impact of social information on tax behaviour is negative again (thus compatible with the taste for social conformity hypothesis), marginally significant but quantitatively large, with an equal to 100 reported income gap ($\Delta INFO_2 > 0$) implying a behavioural adjustment downward by 21.1 tokens. On the other hand, among the subjects who reported a lower

⁶³As regards the controls, being assigned to T2 instead of T1, and C instead of A, did not make any difference. In contrast, belonging to Decision setting B (resp. D) instead of A marginally reduced the variation in reported income between Round 1 and Round 2 by 43.64 (resp. 48.38) tokens as compared to A.

income than others in Round 1 ($\Delta INFO_2 < 0$), information about others' income reporting behaviour had no significant impact on the subsequent subjects' tax behaviour. Our data thus support the notion of an asymmetrical behavioural response to a similar-in-magnitude reported income gap, with a larger adjustment downward than upward, as suggested by the stickiness of bad norms theory and found in several previous studies.

Now, in Round 3, the results displayed in column (5) show that social information about others' average tax behaviour ($\Delta INFO_3$) had a highly significant effect on reported income. The coefficient is negative, thus compatible with the taste for social conformity hypothesis, and larger in absolute value than the coefficient in Round 2, which implies a stronger behavioural convergence process in Round 3: in the baseline (thus in Decision setting A and tax environment T1), being more compliant than others by 100 tokens in Round 2 decreased reported income by 22 tokens from Round 2 to Round 3.⁶⁴ Moreover, when splitting the A and B subject pool with respect to the sign of $\Delta INFO_3$, the results shown in columns (6) and (7) are very contrasting again, but opposite to those obtained in Round 2. Among the subjects who reported a higher income than others in Round 2 ($\Delta INFO_3 > 0$), information about others' average tax behaviour had no impact on subsequent tax behaviour. In contrast, for the subjects who reported a lower income than others in Round 2 ($\Delta INFO_3 < 0$), the coefficient of that variable is significant and equal to -0.274, meaning that an equal to -100 reported income gap implied a significant and large increase by 27.4 tokens in the subjects' reported income. In Round 3, this is among the subjects who were less compliant than others in the first place that the magnitude of the behavioural response was highest. This result (of the upward adjustment being larger than the downward one) looks rather surprising with respect to the literature, but it is quite consistent with our descriptive observations in Table 5.

6 Conclusion

Based on an innovative and flexible experimental design, the purpose of our experimental study was to investigate whether and how the subjects would react to learning about others' tax behaviour. For this purpose, the informational component of the experimental design was enriched in two directions: first, the characterization of 'others' (on whose tax behaviour the subjects were given information) was extended to both peers and non-peers; second, the characterization of the 'provision' of social information was refined, by giving the subjects the possibility to choose whether to get information or not, and which information to get. Two kinds of information were introduced in our experimental design: information about the average tax behaviour within the whole subject's group (meant to proxy the tax community), and information about the average tax behaviour within a subgroup, made of either peers or non peers and chosen by the subject (meant to proxy the subject's reference group). First and foremost, our paper contributes to the literature about the ability of social information about other taxpayers' tax behaviour to affect a subject's tax compliance behaviour. Doing so, our paper also contributes to the literature about the mediating influence of social norms and social conformity in the tax field.

⁶⁴Being assigned to either Tax setting T2 or Decision setting B does not make any difference.

The main results are threefold. First, interesting patterns of attitude towards information emerge from the data, with only a small minority of subjects being not interested in information at all, and a majority of subjects being interested in non-peers' rather than peers' tax behaviour. Second, with regard to our main point, the data display a huge amount of behavioural heterogeneity at the individual level. Roughly 50% of the subjects kept on behaving the same way throughout the experiment, regardless the way other subjects behave, thereby showing strong intrinsic preferences towards taxes. Most of them were fully tax compliers, but there also was a minority of irreducible tax evaders. Even so, our data provide (both on individual and aggregate data) strong evidence of behavioural convergence towards others' average tax behaviour, thereby bringing clear support to the taste for social conformity hypothesis. At the same time, we get non negligible evidence of social anti-conformity-oriented behaviour. Finally, two additional findings are worth mentioning. On the one hand, the kind of information was found to matter, with whole group information being more influential on tax behaviour than reference group information. On the other hand, we find some upward and downward asymmetry in behavioural responses to information.

From a public policy perspective, the behavioural heterogeneity we observed within our subject pool is a reminder that the manipulation of social norms through the dissemination of social information is not a very safe strategy to use on a large scale (Mascagni, 2018), and precautions should be taken to avoid backfiring effects Castro and Scartascini (2015); Garcia et al. (2020); Wenzel (2004). Typically, the fact that around 50% of the population might be sensitive to others' tax behaviour when making their own tax decisions (whether they have a taste for social conformity or anti-conformity) may not be good news, at a time when so many tax evasion cases are being disclosed (and actually need to be disclosed). In this respect, the growing power of social media and networks nowadays is not innocuous, in that the uncontrolled use of information may contribute to initiate and circulate undesirable tax-related social norms, as well as they may create much uncertainty about the dynamics of tax compliance at the aggregate level and undermine the efforts of governments to promote tax compliance (Martinangeli and Windsteiger, 2024; Schultz et al., 2007; Traxler, 2010).

We acknowledge two limitations of our study and take this opportunity to suggest a few avenues for future research. First, even though the composition of our subject-pool was more diverse than it usually is in lab experiments, it was obviously not a representative sample of the population. Typically, even though a large majority of our subjects were actual taxpayers, most of them were either public servants or wage earners, who in real life are all submitted to third-party income reporting (which actually reduces their tax evading possibilities, without eliminating them altogether). Besides, a strong minority of the subjects were Polytechniciens, whose individual characteristics, motives and behavioural patterns are quite atypical, as was shown in the paper. However, the fact that we are able to replicate several previous findings of the experimental literature makes us confident that the findings specific to our study are somewhat relevant too. Nevertheless, the usual caveats apply: our results should obviously not be generalized to the field without caution, and the robustness of our new results deserves some further investigation. In particular, enlarging the subject-sample to other social

categories, and especially to self-employed people whose attitude towards taxes may be somewhat different, is a track for future research. Another possible limitation of our work is that the subjects could not completely choose their reference group, since they could not choose the criterion on which the reference groups were shaped, in this case tax morale. It was obviously not possible, for the sake of both tractability and control, to let each subject choose their own reference group's criterion. Yet, the robustness and external validity of our results is an open question. Further lab and field studies are needed, that could use alternative criteria (such as gender, wealth or status) to build endogenous reference groups, thereby contributing to enlarge and complement our understanding of whether and how social information may affect people's tax decisions.

To conclude, we would like to highlight a result that we think deserves some further empirical investigation. The flexibility of our experimental design allowed us to find that the subjects were as much interested in the behaviour of non-peers as in the behaviour of peers. A huge amount of literature has been devoted to the study of peer effects. Maybe it would also be worth investigating the reasons why individuals, when they are given the possibility to choose on whom to get information, might sometimes be more interested in the behaviour of people who are different and far from them, than in that of people who are similar and close to them. The twofold question of who these influential non-peers actually are, and whether and how they may influence the taxpayers' tax behaviour, is likely of high relevance for public policy.

7 Appendix

7.1 Proof of the theoretical predictions

The individual is assumed to maximize the following utility function with respect to x_i :

$$U(x) = (1 - p)u(I_{NA}) + pu(I_A) + \frac{g}{n}[(n - 1)t\bar{X}_{-i} + tx_i] - F(x_i; \Delta_{i,k}, \theta_i, \alpha_i)$$

with:

$$I_{NA} = R - tx_i$$

where I_{NA} stands for "income when non audited", and

$$I_A = R - tx_i - t(1 + \pi)(R - x_i)$$

where I_A stands for "income when audited".

The first order condition writes:

$$-t(1 - p(x_i, \bar{X}_{-i}))u'_{I_{NA}} + p'_{x_i}(u(I_A) - u(I_{NA})) + p(x_i, \bar{X}_{-i})t\pi u'_{I_A} + \frac{g}{n}t - F'_{x_i} = 0$$

We denote by H the function of the endogenous variable x_i and the set of parameters $(\bar{X}_k, t, \pi, \theta_i, \alpha_{k,i}, g, n)$ such that:

$$H(x_i; \bar{X}_k, t, \pi, \theta_i, \alpha_i, g, n) = -t(1 - p(x_i, \bar{X}_{-i}))u'_{I_{NA}} + p'_{x_i}(u(I_A) - u(I_{NA})) + p(x_i, \bar{X}_{-i})t\pi u'_{I_A} + \frac{g}{n}t - F'_{x_i}$$

Since we assumed that $p''_{x_i} = 0$, $p'_{x_i} < 0$, $u'' < 0$ and $F''_{x_i} > 0$, the second order condition gives:

$$H'_{x_i} = 2p'_{x_i}(u'_{I_A}\pi t + u'_{I_{NA}}t) + t^2(1 - p(x_i, \bar{X}_{-i}))u''_{I_{NA}} + p(x_i, \bar{X}_{-i})t^2\pi u''_{I_A} - F''_{x_i} < 0$$

The implicit function theorem is then used to derive some comparative statics. Since the denominator is strictly negative ($H'_{x_i} < 0$), the sign of each ratio is the opposite of the sign of the numerator. Equations (4), (5), (6), (7) and (8) correspond respectively to Predictions P1 with not-fixed probability, P2, P3, P4 and P1 with fixed probability :

$$\frac{dx_i}{d\bar{X}_{-i}} = \frac{-p'_{\bar{X}_{-i}}(1 + \pi)t}{H'_{x_i}} > 0 \quad (5)$$

$$\frac{dx_i}{d\theta_i} = \frac{F''_{x_i}\theta_i}{H'_{x_i}} > 0 \quad (6)$$

$$\frac{dx_i}{dg} = -\frac{t/n}{H'_{x_i}} > 0 \quad (7)$$

$$\frac{dx_i}{d\Delta_{i,k}} = \frac{F''_{\Delta_{i,k,t-1}}}{H'_{x_i}} < 0 \iff \alpha_i > 0 \quad (8)$$

For a fixed and given probability of audit, we can rewrite:

$$\hat{H}(x_i, \bar{X}_k, t, \pi, \theta_i, \alpha_i, g, n, p) = -t(1-p)u'_{INA} + pt\pi u'_{IA} + \frac{g}{n}t - F'_{x_i}$$

and

$$\hat{H}'_{x_i} = t^2(1-p)u''_{INA} + pt^2\pi u''_{IA} - F''_{x_i}$$

which is negative for any $F''_{x_i} > 0$. For $F''_{x_i} < 0$, the marginal utility effect must outweigh the marginal cost effect to ensure $\hat{H}'_{x_i} < 0$. This condition reduces to:

$$\pi > \frac{F''_{x_i} - t^2(1-p)u''_{INA}}{pt^2u''_{IA}}$$

from which we deduce:

$$\frac{dx_i}{dp} = -\frac{tu'_{INA} + t\pi u'_{IA}}{\hat{H}'_{x_i}} > 0 \quad (9)$$

7.2 The experimental procedures

The Holt-Laury procedure

In the Holt-Laury procedure, the subjects were faced with a 11-line table (numbered from 0 to 10) involving a safe lottery, called Option A ($1000, p; 800, (1 - p)$) and a riskier one, called Option B ($1925, p; 50, (1 - p)$), with p increasing (from line to line) from 0% to 100% with a constant step of 10%. On the first line ($p = 0\%$), Option A stochastically dominates Option B since $800 > 50$. On the last line ($(1 - p) = 0\%$), Option B dominates Option A since $1925 > 1000$. So, a basically rational subject should switch from Option A to Option B between the second and tenth line. The switching point actually depends on the subject's level of risk aversion: the more risk averse they are, the sooner they switch, the higher the switching point in the table. So, this point provides us with a basic index of risk attitude comprised between 1 and 9, with risk neutrality at 5 and risk seeking (resp. aversion) for inferior (resp. superior) to 5 values.

The determination of the probability of audit in T2

The subjects were provided with the following table, and explained how that table worked. The idea was to let them understand that the probability of them being controlled was a decreasing function of their own reported income and an increasing function of others' reported income.

Le revenu moyen déclaré par chaque autre membre du groupe →	0	50	100	150	200	250	300	350	400	450	500	550	600	650	700	750	800	850	900	950	1000	
Mon revenu déclaré ↓																						
0	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
50	43%	48%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
100	35%	40%	45%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
150	28%	33%	38%	43%	48%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
200	20%	25%	30%	35%	40%	45%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
250	13%	18%	23%	28%	33%	38%	43%	48%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
300	10%	10%	15%	20%	25%	30%	35%	40%	45%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
350	10%	10%	10%	13%	18%	23%	28%	33%	38%	43%	48%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
400	10%	10%	10%	10%	10%	15%	20%	25%	30%	35%	40%	45%	50%	50%	50%	50%	50%	50%	50%	50%	50%	50%
450	10%	10%	10%	10%	10%	10%	13%	18%	23%	28%	33%	38%	43%	48%	50%	50%	50%	50%	50%	50%	50%	50%
500	10%	10%	10%	10%	10%	10%	10%	10%	15%	20%	25%	30%	35%	40%	45%	50%	50%	50%	50%	50%	50%	50%
550	10%	10%	10%	10%	10%	10%	10%	10%	10%	13%	18%	23%	28%	33%	38%	43%	48%	50%	50%	50%	50%	50%
600	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	15%	20%	25%	30%	35%	40%	45%	50%	50%	50%	50%
650	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	13%	18%	23%	28%	33%	38%	43%	48%	50%	50%
700	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	15%	20%	25%	30%	35%	40%	45%	50%
750	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	13%	18%	23%	28%	33%	38%	50%
800	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	15%	20%	25%	30%	50%
850	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	13%	18%	23%
900	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	15%
950	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%
1000	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%	10%

Figure 1: Probability of audit table

The socio-demographic characteristics of the subject pool

The descriptive statistics for the subject pool are summarized in Table 9, Column (1). The subjects' main characteristics are tightly related to the composition of the population from which they were drawn. Most of them were students (53%), among whom a number of Master graduates and PhD students at Ecole Polytechnique (13%) and a majority of "Polytechniciens" (40% of the subject pool). "Polytechniciens" are not standard students: first, they earn a salary and already pay taxes from the age of 20; second, a huge majority of them are males.⁶⁵ Administrative staff, the main part of which is female, came second (34% of the whole subject pool). Finally, a number of permanent researchers, both male and female, took part in the experiment (13%).

This mix of features results in the following socio-demographic picture, with a majority of men (58% vs. 42%), rather young subjects (45% of the subjects were under 24, and only 29% were over 35), and 30% of the subjects earning a low/very low income (only 27% earned more than 26700 euros per year). In Columns (2) and (3), the subject pool is split based on the decision setting.

⁶⁵Ecole Polytechnique is the best and most famous engineering Grande Ecole (graduate school) in France. Moreover, it is a public military school, so Polytechniciens are military public servants, who get paid by the French State during their 5-year graduate program.

Table 9: The subject pool: descriptive statistics

		(1)	(2)	(3)
		Whole subject pool	Decision settings A/B	Decision settings C/D
		N = 240	N = 160	N = 80
Variable	Category	% of subjects	% of subjects	% of subjects
Sex	Male	58.33%	55.625%	63.75%
	Female	41.67%	45.375%	36.25%
Age	18 – 24	45.42%	41.875%	52.5%
	25 – 34	25.42%	26.875%	22.5%
	35 – 54	17.91%	18.75%	16.25%
	≥ 55	11.25%	12.5%	8.75%
Tax morale	Full tax morale (7)	45%	45%	45%
	High tax morale (4-6)	44.58%	44.375%	45%
	Low tax morale (0-3)	10.42%	10.625%	10%
Risk attitude	Risk seeking (score ≤ 4)	8.33%	6.25%	12.5%
	Risk neutrality (4 < score ≤ 6)	25.00%	26.875%	21.25%
	Weak risk aversion (6 < score ≤ 8)	31.25%	28.125%	37.5%
	Strong risk aversion (8 < score ≤ 10)	13.75%	13.125%	15%
	Inconsistent answers	21.67%	25.625%	13.75%
Income (euros per year)	Income ≤ 9700	31.67%	30.%	35%
	9700 < income ≤ 26700	41.67%	42.5%	40%
	26700 < income ≤ 71000	18.75%	18.125%	20%
	Income > 71000 or No answer	7.91%	9.375%	5%
Status	Student of Ecole Polytechnique	40.00%	35.625%	48.75%
	Graduate student/PhD Student/Post-Doc	12.92%	15.625%	7.5%
	Professor/ Researcher/Engineer	13.33%	11.875%	16.25%
	Administrative staff/Technician/Other	33.75%	36.875%	27.5%
Experience	Not the first experiment	62.08%	60%	66.25%
	First experiment participation	37.92%	40%	33.75%

7.3 Specific groups of subjects

Table 10: Reported income among fully tax moral subjects in each round

		Nb. of fully TM subj.	Mean reported income (in tokens)
Round 1	<i>Whole subgroup</i>	108	810.18
	<i>T1</i>	46	780.43
	<i>T2</i>	62	832.25
Round 2	<i>Whole subgroup</i>	108	811.94
	<i>T1</i>	46	771.95
	<i>T2</i>	62	841.62
Round 3	<i>Whole subgroup</i>	72	893.66
	<i>T1</i>	29	847.24
	<i>T2</i>	43	879.07

Nb. of fully TM subj.: Number of fully tax moral subjects

Table 11: The subjects' individual characteristics depending on their informational choice (CRGI)

	Kind of CRGI	(1)	(2)	(3)	(4)
		No Info	Same TM	Higher TM	Lower TM
	Size of the subgroup (/240) (% of the whole subject pool)	N = 25 (10,42%)	N = 81 (33,75%)	N = 39 (16,25%)	N = 95 (39,58%)
Variable	Category	% of subj.	% of subj.	% of subj.	% of subj.
Sex	Male	48%	52%	64%	64%
	Female	52%	48%	36%	36%
Age	18 – 24	36%	43%	51%	47%
	25 – 34	28%	21%	26%	29%
	35 – 54	20%	17%	13%	20%
	≥ 55	16%	19%	10%	4%
Tax morale	Full tax morale (7)	36%	60%	18%	45%
	High tax morale (4-6)	56%	30%	62%	47%
	Low tax morale (0-3)	8%	10%	20%	8%
Risk attitude	Risk seeking (score ≤ 4)	8%	4%	10%	12%
	Risk neutrality (4 < score ≤ 6)	20%	25%	26%	26%
	Weak risk aversion (6 < score ≤ 8)	20%	28%	33%	36%
	Strong risk aversion (8 < score ≤ 10)	16%	16%	51%	15%
	Inconsistent answers	36%	27%	26%	11%
Income (euros per year)	income ≤ 9700	24%	33%	33%	32%
	9700 < Income ≤ 26700	52%	33%	49%	43%
	26700 < income ≤ 71000	20%	21%	13%	19%
	Income > 71000 or No answer	4%	13%	5%	6%
Status	Student of Ecole Polytechnique	20%	38%	49%	43%
	Graduate student/PhD Student/Post-Doc	24%	11%	13%	12%
	Professor/Researcher/Engineer	20%	8%	10%	18%
	Administrative staff/Technician/Other	36%	43%	29%	27%

TM: tax morale
 % of subj.: % of subjects

Table 12: The subjects' informational choices (CRGI) depending on their individual characteristics

	Kind of CRGI	(1)	(2)	(3)	(4)
	No Info	Same TM	Higher TM	Lower TM	
	Size of the subgroup (/240) (% of the whole subject pool)	N = 25 (10.42%)	N = 81 (33.75%)	N = 39 (16.25%)	N = 95 (39.58%)
Variable	Category	% of subj.	% of subj.	% of subj.	% of subj.
Sex	Male	8.5%	30%	18%	43.5%
	Female	13%	39%	14%	34%
Age	18 – 24	8%	32%	19%	41%
	25 – 34	11%	28%	17%	44%
	35 – 54	12%	32%	12%	44%
	≥ 55	15%	55%	15%	15%
Tax morale	Full tax morale (7)	8%	45%	7%	40%
	High tax morale (4-6)	13%	22.5%	22.5%	42%
	Low tax morale (0-3)	8%	32%	32%	28%
Risk attitude	Risk seeking (score ≤ 4)	8%	33%	17%	42%
	Risk neutrality (4 < score ≤ 6)	10%	15%	20%	55%
	Weak risk aversion (6 < score ≤ 8)	7%	31%	17%	45%
	Strong risk aversion (8 < score ≤ 10)	12%	39%	6%	43%
	Inconsistent answers	18%	42%	19%	21%
Income (euros per year)	Income ≤ 9700	8%	36%	17%	39%
	9700 < income ≤ 26700	13%	27%	19%	41%
	26700 < income ≤ 71000	11%	38%	11%	40%
	Income > 71000 or no answer	5%	53%	10%	32%
Status	Student of Ecole Polytechnique	5%	32%	20%	43%
	Graduate student/PhD Student/Post-Doc	19%	29%	16%	36%
	Professor/Researcher/Engineer	16%	19%	12%	53%
	Administrative staff/Technician/Other	11%	43%	14%	32%

TM: tax morale

Socio-dem. var.: Socio-demographic variables

% of subj.: % of subjects

Table 13: The individual characteristics of super tax compliers and super tax evaders

	Subjects' type Reported income across rounds	Super tax compliers $x_{i1} = x_{i2} (=x_{i3}) = 1000$	Super tax evaders $x_{i1,2,(3)} \leq 500$
	Size of the group	N = 65	N = 39
Variable	Category	Nb of subjects (%)	Nb of subjects (%)
Sex	Male	31 (47.7%)	32 (82.1%)
	Female	34 (52.3%)	7 (17.9%)
Age	18 – 24	22 (33.85%)	28 (71.8%)
	25 – 34	24 (36.92%)	8 (20.5%)
	35 – 54	12 (18.46%)	2 (5.1%)
	≥ 55	7 (10.77%)	1 (2.6%)
Tax morality	Full tax morale (7)	44 (67.69%)	8 (20.5%)
	High tax morale (4-6)	18 (27.69%)	27 (69.2%)
	Low tax morale (0-3)	3 (4.61%)	4 (10.3%)
Risk aversion	Risk seeking (score ≤ 4)	3 (4.61%)	2 (5.1%)
	Risk neutrality ($4 < \text{score} \leq 6$)	12 (18.46%)	13 (33.3%)
	Weak risk aversion ($6 < \text{score} \leq 8$)	19 (29.23%)	13 (33.3%)
	Strong risk aversion ($8 < \text{score} \leq 10$)	17 (26.15%)	8 (20.5%)
	Inconsistent answers	14 (21.54%)	3 (7.7%)
Income (euros per year)	Income ≤ 9700	17 (26.15%)	15 (38.5%)
	$9700 < \text{income} \leq 26700$	25 (38.46%)	18 (46.1%)
	$26700 < \text{income} \leq 71000$	16 (24.61%)	3 (7.7%)
	Income > 71000 or No answer	7 (10.77%)	3 (7.7%)
Status	Student of Ecole Polytechnique	21 (32.31%)	26 (66.7%)
	Graduate student/PhD Student/Post-Doc	10 (15.38%)	2 (5.1%)
	Professor/Researcher/Engineer	9 (13.84%)	4 (10.3%)
	Administrative staff/Technician/Other	25 (38.46%)	7 (17.9%)

Nb of subjects: Number of subjects

7.4 Variables and variable categories in the regressions

Table 14: The variables and their categories

Variable	Category	Name of category in regressions
Tax setting	1 (basic tax environment)	T1
Decision setting	2 (realistic tax environment)	T2
	A	DS A
	B	DS B
	C	DS C
	D	DS D
Tax morale	Full (7)	TAXMORAL1
	High (4-6)	TAXMORAL2
	Low (0-3)	TAXMORAL3
Sex	Men	MEN
	Women	WOMEN
Age	18-24	AGE1824
	25-34	AGE2534
	35-54	AGE3554
	≥ 55	AGE55
Risk attitude	Risk neutrality	RA1
	Strongly risk seeking	RA2
	Weakly risk seeking	RA3
	Weakly risk averse	RA4
	Strongly risk averse	RA5
Status	Student of Ecole Polytechnique	STATUS1
	Graduate student/PhD Student/Post doc	STATUS2
	Professor/Researcher/Engineer	STATUS3
	Administrative staff/Technician/Other	STATUS4
Income (euros per year)	Income ≤ 9700	INCOME1
	$9700 < \text{income} \leq 26700$	INCOME2
	$26700 < \text{income} \leq 71000$	INCOME3
	Income > 71000 or No answer	INCOME4
Chosen Reference Group	Others having the same level of tax morale (as me)	Same TM
	Others having a lower level of tax morale (than me)	Lower TM
	Others having a higher level of tax morale (than me)	Higher TM

The categories used as the baseline in the regressions (unless specified otherwise in the main text) are in bold.

7.5 Econometric specification in 5.3

The reported income in Round 1 is given by:

$$x_{i,1} = \alpha_0 + \alpha Y_i + \varepsilon_{i,1} \quad (10)$$

In Round 2, the reported income also depends on the information received at the beginning of the round:

$$x_{i,2} = \gamma_0 + \beta(x_{i,1} - \bar{X}_k) + \gamma Y_i + \varepsilon_{i,2} \quad (11)$$

The difference gives:

$$\Delta x_{i,2} = x_{i,2} - x_{i,1} = \mu_0 + \mu_1 \Delta INFO_{i,k,2} + \mu_2 Y_i + \varepsilon_i \quad (12)$$

with $\Delta INFO_{i,k,2} = x_{i,1} - \bar{X}_k$, $\mu_0 = \gamma_0 - \alpha_0$, $\mu_1 = \beta$, $\mu_2 = \gamma - \alpha$, $\varepsilon_i = \varepsilon_{i,2} - \varepsilon_{i,1}$.

More generally, the difference between two rounds writes:

$$\Delta x_{i,t} = \beta_0 + \beta_1 \Delta INFO_{i,k,t} + \beta_2 Y_i + \varepsilon_i \quad (13)$$

To account for each of the two decision processes R1-DP and PR-DP (based on either Round 1 benchmark or previous round benchmark), two sets of regressions were conducted, based on the specification of the variation in reported income described in Eq. (13), where:

- $\Delta x_{i,t}$ is the change in taxpayer i 's reported income between a previous round (being either Round $t - 1$ in RP-DP or first round 1 in R1-DP) and current Round t . From which the general specification:

$$\Delta x_{i,t} = \begin{cases} x_{i,t} - x_{i,t-1} \\ \text{or} \\ x_{i,t} - x_{i,1} \end{cases}$$

- $\Delta INFO_{i,k,t}$, the reported income gap in Round t , is defined as the gap between the taxpayer's previous reported income in either Round $t - 1$ ($x_{i,t-1}$) or first round 1 ($x_{i,1}$) and group k 's reported income in Round 1 (\bar{X}_k), about which the taxpayer learns at the beginning of Round t . From which the general specification:

$$\Delta INFO_{i,k,t} = \begin{cases} x_{i,t-1} - \bar{X}_k \\ \text{or} \\ x_{i,1} - \bar{X}_k \end{cases}$$

In R1-DP, when facing WGI (denoted *whole* in subscripts), $\Delta x_{i,t}$ and $\Delta INFO_{i,whole,t}$ are computed as:

- $\Delta x_{i,t} = \{x_{i,t} | \bar{X}_{whole}\} - x_{i,1}$, with $x_{i,t} | \bar{X}_{whole} = x_{i,2}$ in Decision setting A and $x_{i,t} | \bar{X}_{whole} = x_{i,3}$ in B;
- $\Delta INFO_{i,whole,t} = x_{i,1} - \bar{X}_{whole}$, where \bar{X}_{whole} was provided at the outset of either Round 2 (in A)

or 3 (in B).

Similarly, when facing CRGI (denoted ref in subscripts), $\Delta x_{i,t}$ and $\Delta INFO_{i,ref,t}$ are computed as:

- $\Delta x_{i,t} = \{x_{i,t} | \bar{X}_{ref}\} - x_{i,1}$, with $x_{i,t} | \bar{X}_{ref} = x_{i,2}$ in Decision setting B and $x_{i,t} | \bar{X}_{ref} = x_{i,3}$ in A;
- $\Delta INFO_{i,ref,t} = x_{i,1} - \bar{X}_{ref}$, where \bar{X}_{ref} was provided at the outset of either Round 2 (in B) or 3 (in A).

In the in PR-DP case, $\Delta x_{i,t}$ and $\Delta INFO_{i,k,t}$ were computed the same way, with subscript 1 for Round 1 being replaced with subscript $t - 1$ for previous round (being either Round 1 or Round 2, depending on the information type WGI or CRGI as well as on the decision setting) in the equations above.

References

- Ait Bihi Ouali, L. (2020). Effects of signalling tax evasion on redistribution and voting preferences: Evidence from the panama papers. *Plos one*, 15(3):e0229394.
- Akin, Z. (2019). Dishonesty, social information, and sorting. *Journal of Behavioral and Experimental Economics*, 80:199–210.
- Allingham, M. G. and Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 1(3-4):323–338.
- Alm, J. (2012). Measuring, explaining, and controlling tax evasion: Lessons from theory, experiments, and field studies. *International Tax and Public Finance*, 19(1):54–77.
- Alm, J. (2019). What motivates tax compliance? *Journal of Economic Surveys*, 33(2):353–388.
- Alm, J., Bloomquist, K. M., and McKee, M. (2015). On the external validity of laboratory tax compliance experiments. *Economic Inquiry*, 53(2):1170–1186.
- Alm, J., Bloomquist, K. M., and McKee, M. (2017). When you know your neighbour pays taxes: Information, peer effects and tax compliance. *Fiscal Studies*, 38(4):587–613.
- Alm, J., Jackson, B. R., and McKee, M. (2009). Getting the word out: Enforcement information dissemination and compliance behavior. *Journal of Public Economics*, 93(3):392–402.
- Alm, J., Kirchler, E., and Muehlbacher, S. (2012). Combining psychology and economics in the analysis of compliance: From enforcement to cooperation. *Economic Analysis and Policy*, 42(2):133–151.
- Alm, J. and Malézieux, A. (2021). 40 years of tax evasion games: A meta-analysis. 24:699–750.
- Alm, J., McClelland, G. H., and Schulze, W. D. (1992). Why do people pay taxes? *Journal of Public Economics*, 48(1):21–38.
- Alm, J. and Torgler, B. (2011). Do ethics matter? tax compliance and morality. *Journal of Business Ethics*, 101(4):635–651.
- Bardach, E. (1989). Moral suasion and taxpayer compliance. *Law & Policy*, 11(1):49–69.
- Bazart, C. and Bonein, A. (2014). Reciprocal relationships in tax compliance decisions. *Journal of Economic Psychology*, 40:83–102.
- Benjamini, Y. and Maital, S. (1985). Optimal tax evasion and optimal tax evasion policy: Behavioral aspects. In *The Economics of the Shadow Economy*. pp. 245–264. In W. Gaertner and A. Wenig (eds.), Berlin:Springer-Verlag.
- Besley, T., Jensen, A., and Persson, T. (2023). Norms, enforcement, and tax evasion. *The Review of Economics and Statistics*, 105(4):998–1007.

- Bicchieri, C., Muldoon, R., and Sontuoso, A. (2023). Social norms. In *The Stanford Encyclopedia of Philosophy*. In Edward N. Zalta Uri Nodelman (eds.), Winter 2023 edition. Stanford University.
- Blaufus, K., Braune, M., Hundsdoerfer, J., and Jacob, M. (2015). Self-serving bias and tax morale. *Economics Letters*, 131:91–93.
- Blumenthal, M., Christian, C., Slemrod, J., and Smith, M. G. (2001). Do normative appeals affect tax compliance? evidence from a controlled experiment in minnesota. *National Tax Journal*, 54(1):125–138.
- Bordignon, M. (1993). A fairness approach to income tax evasion. *Journal of Public Economics*, 52(3):345–362.
- Braithwaite, V. and Ahmed, E. (2005). A threat to tax morale: The case of australian higher education policy. *Journal of Economic Psychology*, 26(4):523–540.
- Bø, E. E., Slemrod, J., and Thoresen, T. O. (2015). Taxes on the internet: Deterrence effects of public disclosure. *American Economic Journal: Economic Policy*, 7(1):36–62.
- Castro, L. and Scartascini, C. (2015). Tax compliance and enforcement in the pampas: Evidence from a field experiment. *Journal of Economic Behavior & Organization*, 116:65–82.
- Charroin, L., Fortin, B., and Villeval, M. C. (2022). Peer effects, self-selection and dishonesty. *Journal of Economic Behavior & Organization*, 200:618–637.
- Christian, R. C. and Alm, J. (2014). Empathy, sympathy, and tax compliance. *Journal of Economic Psychology*, 40:62–82.
- Cialdini, R. B. and Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, 55(1):591–621.
- Cialdini, R. B. and Trost, M. R. (1998). Social influence: Social norms, conformity and compliance. *The handbook of social psychology*, pages 151–192. McGraw-Hill (4th edition).
- Cowell, F. A. and Gordon, J. P. (1988). Unwillingness to pay: Tax evasion and public good provision. *Journal of Public Economics*, 36(3):305–321.
- Cummings, R. G., Martinez-Vazquez, J., McKee, M., and Torgler, B. (2009). Tax morale affects tax compliance: Evidence from surveys and an artefactual field experiment. *Journal of Economic Behavior & Organization*, 70(3):447–457.
- Del Carpio, L. (2022). Are the neighbors cheating: Evidence from a social norm experiment on property taxes in peru. *Unpublished Manuscript*. https://www.insead.edu/system/files/2023-07/Lucia-Del-Carpio-Are_the_neighbors_cheatingDecember2022.pdf, INSEAD.
- Di Gioacchino, D. and Fichera, D. (2020). Tax evasion and tax morale: A social network analysis. *European Journal of Political Economy*, 65:101922.

- Diekmann, A., Przepiorka, W., and Rauhut, H. (2015). Lifting the veil of ignorance: An experiment on the contagiousness of norm violations. *Rationality and Society*, 27(3):309–333.
- Dimant, E. (2019). Contagion of pro-and anti-social behavior among peers and the role of social proximity. *Journal of Economic Psychology*, 73:66–88.
- Doerrenberg, P. and Peichl, A. (2013). Progressive taxation and tax morale. *Public Choice*, 155(3/4):293–316.
- Doerrenberg, P. and Peichl, A. (2022). Tax morale and the role of social norms and reciprocity – evidence from a randomized survey experiment. *FinanzArchiv*, 78(1-2):44–86.
- Dularif, M., Sutrisno, T., Nurkholis, and Saraswati, E. (2019). Is deterrence approach effective in combating tax evasion? a meta-analysis. *Problems and Perspectives in Management*, 2:93–113.
- Dvorak, F., Fischbacher, U., and Schmelz, K. (2024). Incentives for conformity and anticonformity. <http://dx.doi.org/10.2139/ssrn.3754595>.
- Dwenger, N., Kleven, H., Rasul, I., and Rincke, J. (2016). Extrinsic and intrinsic motivations for tax compliance: Evidence from a field experiment in germany. *American Economic Journal: Economic Policy*, 8(3):203–32.
- Elster, J. (1989). Social norms and economic theory. *Journal of economic perspectives*, 3(4):99–117.
- Erard, B., Feinstein, J., et al. (1994). The role of moral sentiments and audit perceptions in tax compliance. *Public Finance, Supplement*, 49:70–89.
- Fellner, G., Sausgruber, R., and Traxler, C. (2013). Testing enforcement strategies in the field: Threat, moral appeal and social information. *Journal of the European Economic Association*, 11(3):634–660.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.
- Fortin, B., Lacroix, G., and Villeval, M.-C. (2007). Tax evasion and social interactions. *Journal of Public Economics*, 91(11-12):2089–2112.
- Fosgaard, T. R., Hansen, L. G., and Piovesan, M. (2013). Separating will from grace: An experiment on conformity and awareness in cheating. *Journal of Economic Behavior & Organization*, 93:279–284.
- Frey, B. S. and Torgler, B. (2007). Tax morale and conditional cooperation. *Journal of Comparative Economics*, 35(1):136–159.
- Gächter, S., Nosenzo, D., and Sefton, M. (2012). The impact of social comparisons on reciprocity. *The Scandinavian Journal of Economics*, 114(4):1346–1367.
- Garcia, F., Opromolla, L. D., Vezzulli, A., and Marques, R. (2020). The effects of official and unofficial information on tax compliance. *Journal of Economic Psychology*, 78:102265.

- Gino, F., Ayal, S., and Ariely, D. (2009). Contagion and differentiation in unethical behavior: The effect of one bad apple on the barrel. *Psychological Science*, 20(3):393–398.
- Gino, F. and Pierce, L. (2009). Dishonesty in the name of equity. *Psychological Science*, 20(9):1153–1160.
- Gioia, F. (2017). Peer effects on risk behaviour: The importance of group identity. *Experimental Economics*, 20:100–129.
- Gordon, J. P. (1989). Individual morality and reputation costs as deterrents to tax evasion. *European Economic Review*, 33(4):797–805.
- Guerra, A. and Harrington, B. (2018). Attitude–behavior consistency in tax compliance: A cross-national comparison. *Journal of Economic Behavior Organization*, 156:184–205.
- Halla, M. (2012). Tax morale and compliance behavior: First evidence on a causal link. *The BE Journal of Economic Analysis & Policy*, 12(1):1–27.
- Hallsworth, M., List, J. A., Metcalfe, R. D., and Vlaey, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of Public Economics*, 148:14–31.
- Hammar, H., Jagers, S. C., and Nordblom, K. (2009). Perceived tax evasion and the importance of trust. *The Journal of Socio-Economics*, 38(2):238–245.
- Harrison, G. W. and List, J. A. (2004). Field experiments. *Journal of Economic Literature*, 42(4):1009–1055.
- Hasegawa, M., Hoopes, J. L., Ishida, R., and Slemrod, J. (2013). The effect of public disclosure on reported taxable income: Evidence from individuals and corporations in japan. *National Tax Journal*, 66(3):571–607.
- Hashimzade, N., Myles, G. D., and Tran-Nam, B. (2013). Applications of behavioural economics to tax evasion. *Journal of Economic Surveys*, 27(5):941–977.
- Hofmann, E., Voracek, M., Bock, C., and Kirchler, E. (2017). Tax compliance across sociodemographic categories: Meta-analyses of survey studies in 111 countries. *Journal of Economic Psychology*, 62:63–71.
- Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *The American Economic Review*, 92(5):1644–1655.
- Innes, R. and Mitra, A. (2013). Is dishonesty contagious? *Economic Inquiry*, 51(1):722–734.
- Isler, O. and Gächter, S. (2022). Conforming with peers in honesty and cooperation. *Journal of Economic Behavior Organization*, 195:75–86.
- Jamison, J., Mazar, N., and Sen, I. (2021). Applying behavioral insights to tax compliance: Experimental evidence from latvia. *Journal of Tax Administration*, 6(2):6–29.

- Jimenez, P. and Iyer, G. S. (2016). Tax compliance in a social setting: The influence of social norms, trust in government, and perceived fairness on taxpayer compliance. *Advances in Accounting*, 34:17–26.
- Keizer, K., Lindenberg, S., and Steg, L. (2008). The spreading of disorder. *Science*, 322(5908):1681–1685.
- Keser, C. and Van Winden, F. (2000). Conditional cooperation and voluntary contributions to public goods. *Scandinavian Journal of Economics*, 102(1):23–39.
- Kimbrough, E. O. and Vostroknutov, A. (2016). Norms make preferences social. *Journal of the European Economic Association*, 14(3):608–638.
- Kirchler, E., Hoelzl, E., and Wahl, I. (2008). Enforced versus voluntary tax compliance: The “slippery slope” framework. *Journal of Economic Psychology*, 29(2):210–225.
- Kirchler, E., Muehlbacher, S., Kastlunger, B., and Wahl, I. (2007). Why pay taxes? a review of tax compliance decisions. In *Developing Alternative frameworks for explaining tax compliance, avoidance and evasion*. pp. 15–31. In J. Alm, J. Martinez, B. Torgler (eds.). London: Routledge.
- Kleven, H. J., Knudsen, M. B., Kreiner, C. T., Pedersen, S., and Saez, E. (2011). Unwilling or unable to cheat? evidence from a tax audit experiment in denmark. *Econometrica*, 79(3):651–692.
- Konrad, K. A. and Qari, S. (2012). The last refuge of a scoundrel? patriotism and tax compliance. *Economica*, 79(315):516–533.
- Lefebvre, M., Pestieau, P., Riedl, A., and Villeval, M. C. (2015). Tax evasion and social information: an experiment in belgium, france, and the netherlands. *International Tax and Public Finance*, 22:401–425.
- Lopes, C. and Martins, A. G. (2013). The psychological costs of tax compliance: Some evidence from portugal. *The Journal of Applied Business and Economics*, 14:53–61.
- Luttmer, E. F. and Singhal, M. (2014). Tax morale. *Journal of Economic Perspectives*, 28(4):149–68.
- Macho-Stadler, I. and Pérez-Castrillo, J. D. (2002). Auditing with signals. *Economica*, 69(273):1–20.
- Martinangeli, A. F. and Windsteiger, L. (2024). Inequality shapes the propagation of unethical behaviours: Cheating responses to tax evasion along the income distribution. *Journal of Economic Behavior Organization*, 220:135–181.
- Mascagni, G. (2018). From the lab to the field: A review of tax experiments. *Journal of Economic Surveys*, 32(2):273–301.
- 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.

- Méder, Z., Simonovits, A., and Vincze, J. (2012). Tax morale and tax evasion: Social preferences and bounded rationality. *Economic Analysis and Policy*, 42(2):171–188.
- Morgan, T. J. and Laland, K. N. (2012). The biological bases of conformity. *Frontiers in neuroscience*, 6:87.
- Myles, G. D. and Naylor, R. A. (1996). A model of tax evasion with group conformity and social customs. *European Journal of Political Economy*, 12(1):49–66.
- Nar, M. (2023). Costs stemming from tax systems: Tax compliance costs. *The Journal of Asian Finance, Economics and Business*, 10(2):267–280.
- Noguera, J. A., Quesada, F. J. M., Tapia, E., and Llàcer, T. (2014). Tax compliance, rational choice, and social influence: An agent-based model. *Revue française de sociologie*, 55(4):765–804.
- Orviska, M. and Hudson, J. (2003). Tax evasion, civic duty and the law abiding citizen. *European Journal of Political Economy*, 19(1):83–102.
- Pickhardt, M. and Prinz, A. (2014). Behavioral dynamics of tax evasion – a survey. *Journal of Economic Psychology*, 40:1–19.
- Pommerehne, W. W. and Frey, B. S. (1992). The effects of tax administration on tax morale. Technical report, Discussion Papers, Series II 191, University of Konstanz, Collaborative Research Centre (SFB) 178 "Internationalization of the Economy".
- Rauhut, H. (2013). Beliefs about lying and spreading of dishonesty: Undetected lies and their constructive and destructive social dynamics in dice experiments. *PloS one*, 8(11):e77878.
- Reck, D., Slemrod, J., and Vattø, T. E. (2022). Public disclosure of tax information: Compliance tool or social network? *Journal of Public Economics*, 212:104708.
- Reinganum, J. F. and Wilde, L. L. (1985). Income tax compliance in a principal-agent framework. *Journal of Public Economics*, 26(1):1–18.
- Rockenbach, B., Tonke, S., and Weiss, A. R. (2021). Self-serving behavior of the rich causes contagion effects among the poor. *Journal of Economic Behavior & Organization*, 183:289–300.
- Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., and Griskevicius, V. (2007). The constructive, destructive, and reconstructive power of social norms. *Psychological Science*, 18(5):429–434.
- Scotchmer, S. (1987). Audit classes and tax enforcement policy. *The American Economic Review*, 77(2):229–233.
- Slemrod, J. (2007). Cheating ourselves: The economics of tax evasion. *Journal of Economic Perspectives*, 21(1):25–48.
- Slemrod, J. (2019). Tax compliance and enforcement. *Journal of Economic Literature*, 57(4):904–954.

- Slemrod, J., Rehman, O. U., and Waseem, M. (2022). How do taxpayers respond to public disclosure and social recognition programs? evidence from pakistan. *The Review of Economics and Statistics*, 104(1):116–132.
- Torgler, B. (2002). Speaking to theorists and searching for facts: Tax morale and tax compliance in experiments. *Journal of Economic Surveys*, 16(5):657–683.
- Torgler, B. (2007). *Tax compliance and tax morale: A theoretical and empirical analysis*. Edward Elgar Publishing.
- Torgler, B. and Murphy, K. (2004). Tax morale in australia: What shapes it and has it changed over time. *Journal of Australian Taxation*, 7:298–335.
- Traxler, C. (2010). Social norms and conditional cooperative taxpayers. *European Journal of Political Economy*, 26(1):89–103.
- Trivedi, V. U., Shehata, M., and Lynn, B. (2003). Impact of personal and situational factors on taxpayer compliance: An experimental analysis. *Journal of Business Ethics*, 47(3):175–197.
- Ushchev, P. and Zenou, Y. (2020). Social norms in networks. *Journal of Economic Theory*, 185:104969.
- Wenzel, M. (2004). An analysis of norm processes in tax compliance. *Journal of Economic Psychology*, 25(2):213–228.
- Wenzel, M. (2005a). Misperceptions of social norms about tax compliance: From theory to intervention. *Journal of Economic Psychology*, 26(6):862–883.
- Wenzel, M. (2005b). Motivation or rationalisation? causal relations between ethics, norms and tax compliance. *Journal of Economic Psychology*, 26(4):491–508.
- Wilson, J. Q. and Kelling, G. L. (1982). Broken windows. *Atlantic monthly*, 249(3):29–38.
- Yitzhaki, S. (1974). A note on income tax evasion: A theoretical analysis. *Journal of Public Economics*, 3:201–202.