What Do Our Choices Say About Our Preferences?

Malgorzata Sulkowska, Krzysztof Grining, Marek Klonowski

What Do Our Choices Say About Our Preferences?

Krzysztof Grining[1], Marek Klonowski[2], Małgorzata Sulkowska[1,3]

[1] Department of Fundamentals of Computer Science
Wrocław University of Science and Technology
{firstname.secondname}@pwr.edu.pl

[2] Department of Theoretical Physics
Wrocław University of Science and Technology
{firstname.secondname}@pwr.edu.pl

[3] Université Côte d'Azur, CNRS, Inria, I3S, France
{firstname.secondname}@inria.fr

Abstract. Taking online decisions is a part of everyday life. Think of buying a house, parking a car or taking part in an auction. We often take those decisions publicly, which may breach our privacy - a party observing our choices may learn a lot about our preferences. In this paper we investigate the online stopping algorithms from the privacy preserving perspective, using a mathematically rigorous differential privacy notion.
In differentially private algorithms there is usually an issue of balancing the privacy and utility. In this regime, in most cases, having both optimality and high level of privacy at the same time is impossible. We propose a natural mechanism to achieve a controllable trade-off, quantified by a parameter, between the accuracy of the online algorithm and its privacy. Depending on the parameter, our mechanism can be optimal with weaker differential privacy or suboptimal, yet more privacy-preserving. We conduct a detailed accuracy and privacy analysis of our mechanism applied to the optimal algorithm for the classical secretary problem. Thereby the classical notions from two distinct areas - optimal stopping and differential privacy - meet for the first time.

Key words: privacy preserving algorithm, optimal stopping, differential privacy, secretary problem

1 Introduction

We make online decisions every day - buying a house, trading stock options or parking a car. The choices we make are mostly based on our knowledge and experience. Those decisions are most often publicly visible and it raises concern about the internal information on which the choice was based. Namely, our choices may somewhat leak that information, which we consider as sensitive. Someone could observe our choices and deduce our preferences or domain knowledge of our algorithmic trading company. In this paper we consider the security of such information in case of optimal stopping algorithms.

The optimal stopping algorithms are widely known and thoroughly researched. One of the most classical models in this area is the *secretary problem*. We have a set of n

linearly ordered candidates and a *selector* would like to choose the best candidate from that set. The caveat is that the candidates appear online one by one in some random order. I.e., the selector cannot see all of them at once and simply pick the best one. He has only one choice and has to make it online. The decision must be based on an incomplete knowledge gathered by comparing the current candidate with the previously seen ones, and it is irreversible, as the candidate cannot be hired after he or she was rejected. There is a known, optimal solution to that problem which gives the asymptotic probability of $1/e$ for choosing the best candidate (consult [1]). The problem was popularized in 1960 by Martin Gardner in *Scientific American* column under the name of *googol game* and the optimal solution was first written down by Lindley in [1]. For a historical overview of the secretary problem consult Ferguson's survey [2]. Many generalizations of this classical version were considered later on. Stadje was the first one who replaced a linear order of candidates with the partial order (poset) and aimed at choosing any maximal element [3]. An account of research considering threshold strategies for posets was given by Gnedin in [4]. Optimal strategies for a particular posets as well as universal algorithms for the whole families of posets have been featured in [5,6,7,8,9]. The problem was investigated even on much more general structures, like matroids (consult [10]). Other interesting extensions consider different payoff functions. A natural reformulation of the classical case is to aim at minimizing the expected rank of the candidate instead of at selecting just the best one. This model was introduced by Lindley in [1] but fully solved by Chow et al. in [11]. In [12] the authors present the stopping problem in the auction setting where the seller has a multidimensional objective function with only a partial order among the outcomes. Even though formulated probably in the mid-1950s, the secretary problem with its generalizations is still vivid and attracts the attention of theoretical computer science community, see the latest results in [13,14,15].

In this paper we investigate the optimal stopping algorithms from a privacy-protection perspective, which, to the best of our knowledge, has not been done so far. Assuming that the online choice we make is publicly visible, it may leak some information about our preferences. Note that if the choice was possible offline, i.e., we could first see all the candidates and then pick the best one, the leakage would be unavoidable, as the chosen candidate would simply be our best one. But when the choice has to be made in an online regime (we pick or reject the candidates on the fly, without the possibility of revisiting the rejected ones) then on one hand the outcome is not perfectly accurate but on the other our preferences are not that visible. The optimal stopping algorithms can be used for example in algorithmic trading. One might care whether the visible action on the market (e.g., closing an American option position) leaks something about an internal knowledge. Hereby we try to answer the question whether the inherent uncertainty of the online stopping algorithms is enough to sufficiently hide our preferences.

Our analysis of privacy is based on *differential privacy* notion commonly considered as the only state-of-the-art approach. Its idea was introduced by Dwork et al. in [16], however its precise formulation in the widely used form appeared in [17]. Differential privacy is mathematically rigorous and formally provable in contrary to the previous anonymity-derived privacy definitions (for a comprehensive study check [18] and references therein). Informally, the idea behind differential privacy is as follows: for two "similar" inputs, a differentially private mechanism should provide a response chosen

from a very similar distributions. In effect, judging by the output of the mechanism one cannot say if a given individual was taken into account for producing a given output, as it cannot distinguish two outputs produced from two data sets differing with only one user. This is mostly done by adding an auxiliary randomness (e.g., a carefully calibrated noise) to data. Vast majority of the fundamental papers (e.g. [19,20,21]) consider centralized model as is assumed in our contribution. In this model a trusted party (called *curator*) holds a database. He is entitled to gather and process all participants' data. Curator also releases the computed statistics to a wider (possibly untrusted) audience. Such approach to the privacy-preserving protocols can be used to give a formal guarantee for privacy resilient to any form of post-processing. Analysis of the protocols based on differential privacy is usually technically much more involved comparing with previous approaches, but they give immunity against various linkage attacks (see e.g. [22,23]).

1.1 Results

We can think of two extremes while making an online decision - one is to act according to the known, optimal algorithm and the other is to act completely randomly. Obviously, the first approach yields the best possible chance to make a correct pick but leaks internal information while the latter does not leak anything but selects a candidate only at random. Depending on the nature of the problem and the importance of information to be hidden we have to make a trade-off.

In this paper we present a natural algorithm that is such a trade-off. It has a steering parameter $p \in [0,1]$: by $p = 1$ it is optimal, and by $p = 0$ it hides preferences perfectly. We analyze its effectiveness and privacy properties. Subsequently, we apply it to the classical secretary problem. It turns out that already the optimal secretary algorithm itself ensures some privacy, but rather only in weak metrics. Our approach, basing on adding a randomized perturbation, is common in privacy mechanisms. This intuitively obvious algorithm turned out to be surprisingly difficult to analyze. This work can be seen as the first step to tackle privacy and information hiding in the optimal stopping algorithms so we focused only on the classical model.

Main results

- We introduce the definition of a *differentially private stopping time* (see Definition 5 in Section 2) with a metric defined on a set of orderings representing selector's preferences (see Definition 6 in Section 2).
- In Theorem 1 we show a fundamental lower bound for the possibility of constructing privacy preserving algorithms that are close to optimal.
- We propose a natural mechanism transforming any optimal stopping time for a linearly ordered set of candidates into an algorithm preserving accuracy at a controllable level and having better information hiding properties. Formulation of the algorithm is presented in Section 4. General results concerning its accuracy and differential privacy are stated in Section 4.1 (see Fact 5 and Theorem 3).
- We conduct a detailed analysis of our mechanism applied to the optimal stopping algorithm for the classical secretary problem. The results on its accuracy and differential privacy are stated in Section 4.2 (see Fact 6 and Theorem 4). In particular, we obtain privacy properties of the optimal algorithm for the secretary problem.

In order to obtain the privacy results for our mechanism applied to the secretary problem we had to prove a series of collateral properties of the optimal algorithm for the secretary problem (see the Appendix).

2 Formal model

2.1 Stopping time

Let $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ be the set of n candidates. Let $\mathcal{S}_n$ be the set of all permutations of elements from $\mathcal{C}$, $|\mathcal{S}_n| = n!$. We interpret a single permutation σ from $\mathcal{S}_n$ as the ordering of candidates with respect to their qualifications, i.e., $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_n) = (C_{\sigma_1}, C_{\sigma_2}, \ldots, C_{\sigma_n})$ refers to the ordering $C_{\sigma_1} > C_{\sigma_2} > \ldots > C_{\sigma_n}$ (in particular, it means that the candidate C_{σ_1} has the best and the candidate C_{σ_n} has the worst qualifications from the whole group $\mathcal{C}$). These orderings are called the *qualification orderings*.

Let $\mathcal{T}_n$ be also the set of all permutations of the elements from $\mathcal{C}$, just this time we interpret $\tau \in \mathcal{T}_n$ as the sequence giving the order in which the candidates appear in some online game ($\tau = (\tau_1, \tau_2, \ldots, \tau_n)$ means, in particular, that the candidate C_{τ_1} appeared as the first one and the candidate C_{τ_n} as the last one in our online game). These orderings are called the *time orderings*.

Throughout this paper we refer to the following model of an online stopping problem. Fix σ (choose a particular qualification ordering on the set of candidates). Note that, in fact, σ is a preference which we want to hide. Assume that τ is chosen uniformly at random from $\mathcal{T}_n$. The player knows σ but he does not know τ. Candidates from $\mathcal{C}$ appear one by one following the order given by τ. At time t, i.e., when t candidates appeared, the player observes the qualification order induced by $\{\tau_1, \tau_2, \ldots, \tau_t\}$. That is, he knows the relative ranks of the candidates seen so far but he does not know their total ranks. At each time step he has to decide whether to continue the game and reveal the next element or to stop the game meaning that he selects the element τ_t. If he decides to reveal another element, he is not allowed to come back to the previous steps of the game. His task is to maximize (or ensure relatively high) probability that the selected candidate belongs to some previously defined set (e.g. is maximal in the whole set $\mathcal{C}$).

Formally, we define a probability space $(\mathcal{T}_n, \mathcal{P}, \mathbb{P})$, where $\mathcal{P}$ is the set of all subsets of $\mathcal{T}_n$ and the probability measure is defined by $\mathbb{P}[\{\tau\}] = 1/n!$ for any $\tau \in \mathcal{T}_n$. A *stopping time* is a function $M : \mathcal{S}_n \times \mathcal{T}_n \to \{1, 2, \ldots, n\}$ such that its value, say $t = M(\sigma, \tau)$, depends only on information the player gathered up to time t, which is the qualification order induced by $\{\tau_1, \tau_2, \ldots, \tau_t\}$. The value can not depend on any future events. We give a strict formal definition of a stopping time below.

Definition 1 *Let $\mathcal{P}_1 \subseteq \mathcal{P}_2 \subseteq \ldots \subseteq \mathcal{P}_n \subseteq \mathcal{P}$ be a sequence of σ-algebras (such a sequence is called a filtration). A random variable $M : \mathcal{S}_n \times \mathcal{T}_n \to \{1, 2, \ldots, n\}$ is a stopping time with respect to a filtration $(\mathcal{P}_t)_{t=1}^n$ if, truncating M to any $\sigma \in \mathcal{S}_n$, we have that $(M|_\sigma)^{-1}(t) \in \mathcal{P}_t$ for all $t \leq n$.*

In our case the sets A in $\mathcal{P}_t$ are those with the following property. Fix $\sigma \in \mathcal{S}_n$. If $\tau = (\tau_1, \tau_2, \ldots, \tau_n) \in A$ then for every $\tilde{\tau} \in \mathcal{T}_n$ such that the orders of candidates induced by τ and $\tilde{\tau}$ are identical up to time t we have $\tilde{\tau} \in A$.

The expression $M(\sigma,\tau) = t$ means that the algorithm M stopped at time t, thus selected the candidate C_{τ_t}. Notation for a candidate returned by M while playing on the time ordering τ is $\tau_{M(\sigma,\tau)}$. We will often refer to the probability that a candidate returned by M belongs to some subset S of the set of all candidates $\mathcal{C}$, i.e. $\mathbb{P}[\tau_{M(\sigma,\tau)} \in S]$. In order to clarify notation we introduce a notion $CM(\sigma,\tau)$ for $\tau_{M(\sigma,\tau)}$.

Definition 2 *Let $\mathcal{M}$ denote the set of all stopping times. We say that M^{opt} is an optimal stopping time if*

$$M^{opt} = \operatorname{argmax}_{M\in\mathcal{M}}\mathbb{P}[\tau_M \in D],$$

where $D \subseteq \mathcal{C}$ is a set of previously defined candidates and τ is chosen uniformly at random from $\mathcal{T}_n$.

The set $\{1,2,\ldots,n\}$ will be denoted by $[n]$. We write $f(n) \sim g(n)$ whenever $\lim_{n\to\infty} f(n)/g(n) = 1$.

2.2 Secretary problem

From now on M^* always denotes the optimal stopping time for the classical *secretary problem*. In the secretary problem the player aims at maximizing the probability of selecting the candidate which is the best from the whole $\mathcal{C}$. That is, for a particular $\sigma \in \mathcal{S}_n$, the set D from Definition 2 is given by $D = \{\sigma_1\}$. A full solution to this problem (the optimal algorithm and its probability of success) is well known, consult [1] or [2]. The asymptotic results are also established. We present them below.

Definition 3 *Let $\sigma \in \mathcal{S}_n$ be the qualification ordering of the elements from $\mathcal{C}$. Let the value t_n be a so-called threshold of the algorithm. The optimal algorithm M^* for the secretary problem (the one maximizing $\mathbb{P}[\tau_M = \sigma_1]$ over $M \in \mathcal{M}$) is defined as follows. For any $\tau \in \mathcal{T}_n$ we have $M^*(\sigma,\tau) = k$ if and only if*

(1) $k > t_n - 1$ and
(2) τ_k is the maximal element in the qualification ordering induced by $\{\tau_1,\tau_2,\ldots,\tau_k\}$ and
(3) for $i \in \{t_n,\ldots,k-1\}$ the element τ_i is not maximal in the qualification ordering induced by $\{\tau_1,\ldots,\tau_i\}$.

If τ is such that the above three conditions are never altogether satisfied, then $M^*(\sigma,\tau) = n$.

Fact 1 *The threshold t_n of the optimal algorithm M^* for the classical secretary problem with n candidates is defined as the smallest integer t for which*

$$\frac{1}{t} + \frac{1}{t+1} + \ldots + \frac{1}{n-1} \leqslant 1.$$

We have $t_n \sim n/e$ (consult [1]).

Fact 2 *Fix $\sigma \in \mathcal{S}_n$. The probability that the optimal algorithm M^* selects the k^{th} best candidate is given by*

$$\mathbb{P}[CM^*(\sigma,\tau)=\sigma_k]=\frac{t_n-1}{n}\left(\sum_{i=t_n}^{n-k+1}\frac{\binom{n-k}{i-1}}{\binom{n-1}{i-1}}\frac{1}{i-1}+\frac{1}{n-1}\right),$$

where t_n is the threshold from Fact 1 (consult [24] or see Theorem 5 in the Appendix). In particular, the probability that it wins (selects the best candidate) is

$$\mathbb{P}[CM^*(\sigma,\tau)=\sigma_1]=\frac{t_n-1}{n}\sum_{i=t_n}^{n}\frac{1}{i-1}\xrightarrow{n\to\infty}1/e\approx 0.37.$$

In general, when k is a constant or a function of n such that $k(n)=o(n)$

$$\mathbb{P}[CM^*(\sigma,\tau)=\sigma_k]\sim\frac{1}{e}\sum_{s=k}^{\infty}\frac{1}{s}\left(1-\frac{1}{e}\right)^s$$

(consult [24] or see Theorem 6 in the Appendix).

2.3 Differential privacy

In this subsection we recall the definition of *differential privacy* and present the privacy model used throughout the paper. For more details about differential privacy see e.g. [18].

We assume that there exists a trusted curator who holds data of individuals in the database. A *privacy mechanism* is an algorithm, used by the curator, that takes as an input a database and produces an output (a release) using randomization. By $\mathcal{X}$ we denote the space of all possible rows in a database (each row consits of data of some individual). The privacy mechanism has a domain $\mathbb{N}^{|\mathcal{X}|}$ representing the set of databases. Thus each database is represented as an $|\mathcal{X}|$-tuple $(n_1, n_2, \ldots, n_{|\mathcal{X}|})$, where n_k is interpreted as the number of rows of kind k in this database. If $x=(n_1,n_2,\ldots,n_{|\mathcal{X}|})$ then $n_1+n_2+\ldots+n_{|\mathcal{X}|}$ is the number of rows in x. The goal is to protect data of every single individual, even if all the users except one collude with the Adversary to breach the privacy of this single, uncorrupted user.

Definition 4 (Differential Privacy, [18]) *A randomized algorithm M with the domain $\mathbb{N}^{|\mathcal{X}|}$ is (ε,δ)-differentially private, if for all $S\subseteq$ Range(M) and for all $x,y\in\mathbb{N}^{|\mathcal{X}|}$ such that $\|x-y\|_1\leqslant 1$ the following condition is satisfied:*

$$\mathbb{P}[M(x)\in S]\leqslant e^{\varepsilon}\cdot\mathbb{P}[M(y)\in S]+\delta,$$

where the probability space is over the outcomes of M and $\|\cdot\|_1$ denotes the standard l_1 norm.

The intuition behind the (ε,δ)-differential privacy is that if we choose two consecutive databases (that differ exactly on one record), then the mechanism is very likely to return

indistinguishable values. Speaking informally, it preserves privacy with high probability, thus the outcome could be distinguishable, thus not privacy-preserving, only with probability at most δ.

In this paper we consider the optimal stopping algorithms and by the database we understand a permutation of the set of n choices (e.g., candidates in the secretary problem). Rather than hiding the participation of a candidate in a stopping game, we wish to hide the preferences of the selector. Indeed it is the preference that is sensitive, not the participation itself. For example, in financial markets the set of candidates is publicly known (say, the prices of a stock). On the other hand, our reaction to that set based on the underlying assumptions and the domain knowledge, which are connected with the preferences permutation, is sensitive and needs protection.

Below we present a differential privacy definition reformulated for our purposes.

Definition 5 *Let $(\mathcal{S}_n, d)$ be a metric space. A randomized algorithm M with the domain $\mathcal{S}_n \times \mathcal{T}_n$ is an (ε, δ)-differentially private stopping time w.r.t. metric d if for all $S \subseteq$ Range(M) and for all $\sigma, \rho \in \mathcal{S}_n$ such that $d(\sigma, \rho) \leqslant 1$ we have*

$$\mathbb{P}[CM(\sigma, \tau) \in S] \leqslant e^{\varepsilon}\mathbb{P}[CM(\rho, \tau) \in S] + \delta,$$

where τ is chosen uniformly at random from $\mathcal{T}_n$.

Since we assume a uniform distribution on $\mathcal{T}_n$, all the probabilities are calculated with respect to the fact that τ is chosen uniformly at random from $\mathcal{T}_n$. From now on, for short, we sometimes write (ε, δ)-DP for an (ε, δ)-differentially private stopping time.

We introduce a definition of a parametrized metric on $\mathcal{S}_n$.

Definition 6 *For $l \in \{1, 2, \ldots, n-1\}$ let $d_l : \mathcal{S}_n \times \mathcal{S}_n \to [0, \infty)$. We say that d_l is an l-distance between the permutations $\sigma, \rho \in \mathcal{S}_n$ if*

$$d_l(\sigma, \rho) = \min\{k : \sigma = \pi_1 \circ \pi_2 \circ \ldots \circ \pi_k \circ \rho\},$$

where $\pi_i \in \mathcal{S}_n$ and each π_i is a transposition of the elements being at most l apart in the permutation $\pi_{i+1} \circ \ldots \circ \pi_k \circ \rho$.

Example 1 *Let $\sigma = (C_1, C_2, C_3, \ldots, C_{n-1}, C_n)$ and $\rho = (C_n, C_2, C_3, \ldots, C_{n-1}, C_1)$. We have $d_1(\sigma, \rho) = 2n - 3$ since we need to make at least $2n - 3$ swaps of the neighboring elements in order to obtain σ from ρ:*

$$\sigma = (C_n\ C_{n-1}) \ldots (C_n\ C_3) \circ (C_n\ C_2) \circ (C_1\ C_n) \ldots (C_1\ C_{n-2}) \circ (C_1\ C_{n-1}) \circ \rho.$$

However, $d_{n-1}(\sigma, \rho) = 1$ since we need just one swap of the elements $n - 1$ apart in order to obtain σ from ρ:

$$\sigma = (C_1\ C_n) \circ \rho.$$

Fact 3 *Let $l \in \{1, 2, \ldots, n-1\}$. It is easy to show that (S_n, d_l) is a metric space.*

Note that d_{n-1} is the strongest metric as it treats two permutations with any pair of swapped elements as neighboring. On the other hand we have d_1 (here permutations

are neighboring only if it is a pair with the neighboring elements swapped). From the privacy perspective point of view these metrics significantly differ. A privacy mechanism using d_{n-1}, intuitively, hides our preferences completely. It should e.g. output a similar outcome even for the qualification orderings σ and $\tilde{\sigma}$, where $\tilde{\sigma}$ swaps in σ the best candidate with the worst one. In general, d_l metric hides preferences for up to l distance. I.e., say that a picked candidate was in reality k-th on the full preference list. Then from the Adversary perspective he or she could have been between $k-l$ and $k+l$ on the preference list.

Note that if an algorithm is private in a stronger metrics, the Adversary cannot really guess the preferences of the selector, despite knowing the picked candidate. Intuitively, constructing such an algorithm having a high probability of success seems hard to achieve, as then the final choice should say almost nothing about the selector's preferences. Indeed, in the next section we prove that achieving a reasonable differential privacy parameters (ε and δ) and a constant probability of success in case of metrics d_l such that $l = l(n) \xrightarrow{n\to\infty} \infty$ is impossible.

3 General results

In this section we present two general results for the (ε, δ)-DP stopping times with the metric d_l. In the first one we refer to the problem of choosing the best candidate, thus for a fixed $\sigma \in \mathcal{S}_n$ the set D from Definition 2 is given by $D = \{\sigma_1\}$. This result tells that by $l = l(n) \xrightarrow{n\to\infty} \infty$ it is impossible to construct an algorithm with a constant probability of success having reasonable privacy parameters. The second result constitutes the lower bound for ε for a given optimal stopping time and δ.

Theorem 1. *Fix $\sigma \in \mathcal{S}_n$ and let $\delta \geq 0$. Consider a metric space $(\mathcal{S}_n, d_l)$ for $l = l(n) \xrightarrow{n\to\infty} \infty$. Then, for every stopping time M such that $\mathbb{P}[CM(\sigma,\tau) = \sigma_i] > 0$ for all $i \in \{1, 2, \ldots, n\}$, if M is (ε, δ)-DP w.r.t. d_l, then $\mathbb{P}[CM(\sigma,\tau) = \sigma_1] \xrightarrow{n\to\infty} 0$ or $\varepsilon \xrightarrow{n\to\infty} \infty$.*

Remark 1. Here we assume that δ is significantly smaller than $\mathbb{P}[CM(\sigma,\tau) = \sigma_1]$, i.e. that the difference between δ and $\mathbb{P}[CM(\sigma,\tau) = \sigma_1]$ tends to some positive constant with $n \to \infty$.

Proof. Assume by contradiction that $l = l(n) \xrightarrow{n\to\infty} \infty$ and the stopping algorithm M is such that $\mathbb{P}[CM(\sigma,\tau) = \sigma_1]$ is a positive constant and M is (ε, δ)-DP w.r.t. d_l with a constant ε. Consider the following qualification orderings, all being at distance 1 from σ w.r.t. d_l:

$$\begin{aligned}
\rho_2 &= (\sigma_1\ \sigma_2) \circ \sigma = (\sigma_2, \sigma_1, \sigma_3, \ldots, \sigma_n),\\
\rho_3 &= (\sigma_1\ \sigma_3) \circ \sigma = (\sigma_3, \sigma_2, \sigma_1, \sigma_4, \ldots, \sigma_n),\\
&\ldots\\
\rho_{l+1} &= (\sigma_1\ \sigma_{l+1}) \circ \sigma = (\sigma_{l+1}, \sigma_2, \ldots, \sigma_l, \sigma_1, \sigma_{l+2}, \ldots, \sigma_n).
\end{aligned}$$

Let $\mathbb{P}[CM(\sigma,\tau) = \sigma_j] = q_{j,n}$. Note that for $j = 2, 3, \ldots, l+1$

$$\mathbb{P}[CM(\rho_j,\tau) = \sigma_1] = q_{j,n}.$$

Since M is (ε, δ)-DP w.r.t. d_l the following system of l inequalities is satisfied

$$\varepsilon \geq \ln \frac{q_{1,n} - \delta}{q_{2,n}}, \quad \varepsilon \geq \ln \frac{q_{1,n} - \delta}{q_{3,n}}, \ldots, \quad \varepsilon \geq \ln \frac{q_{1,n} - \delta}{q_{l+1,n}}$$

(consider $S = \{\sigma_1\}$ in Definition 5). We infer that since $q_{1,n}$ is a constant and ε is a constant, the probabilities $q_{j,n}$ for $j = 2, 3, \ldots, l+1$ are also constant. Now, since $l = l(n) \xrightarrow{n\to\infty} \infty$ we get $q_{1,n} + q_{2,n} + \ldots + q_{l+1,n} \xrightarrow{n\to\infty} \infty$ which contradicts the fact that $\sum_{j=1}^{n} q_{j,n} = 1$.

Therefore throughout the rest of the paper we always assume that l is a constant.

Theorem 2. *Consider a metric space $(\mathcal{S}_n, d_l)$. Let $\delta \geq 0$ and let M be a stopping algorithm which is (ε, δ)-DP w.r.t. d_l. For a given $\sigma \in \mathcal{S}_n$ let*

$$\mathbb{P}[CM(\sigma, \tau) = \sigma_i] = q_{i,n} \in (0, 1), \quad i = 1, 2, \ldots, n.$$

If there exists at least one pair $i, j \in [n]$ such that $i \neq j$, $|i - j| \leqslant l$, $q_{i,n} \geqslant q_{j,n}$ and $\delta < q_{i,n} - q_{j,n}$ then

$$\varepsilon \geqslant \max_{\substack{1 \leqslant i,j \leqslant n \\ |i-j| \leqslant l}} \ln \left\{ \frac{q_{i,n} - \delta}{q_{j,n}} \right\}.$$

Otherwise $\varepsilon \geq 0$. The inequalities are tight.

Remark 2. In this paper most of the time we find ourselves in the first situation. Usually δ will be already significantly smaller than $|q_{1,n} - q_{2,n}|$.

Proof. Since M is (ε, δ)-DP w.r.t. d_l, the following inequality has to be satisfied for all $S \subseteq \mathcal{C}$ and for all $\rho \in \mathcal{S}_n$ such that $d_l(\sigma, \rho) = 1$

$$\mathbb{P}[CM(\sigma, \tau) \in S] \leqslant e^{\varepsilon} \mathbb{P}[CM(\rho, \tau) \in S] + \delta. \tag{1}$$

Equivalently, we will work with the inequality

$$e^{\varepsilon} \geqslant \frac{\mathbb{P}[CM(\sigma, \tau) \in S] - \delta}{\mathbb{P}[CM(\rho, \tau) \in S]} \tag{2}$$

assuming that $\mathbb{P}[CM(\rho, \tau) \in S] \neq 0$. (When $\mathbb{P}[CM(\rho, \tau) \in S] = 0$ the inequality (1) holds with any $\varepsilon \geqslant 0$.) Thus let us investigate what is the maximal value that the right-hand side of the inequality (2) may attain. Since $d_l(\sigma, \rho) = 1$, let us express ρ as

$$\rho = (\sigma_i\ \sigma_j) \circ \sigma = (\sigma_1, \ldots, \sigma_{i-1}, \sigma_j, \sigma_{i+1}, \ldots, \sigma_i, \ldots, \sigma_n),$$

where $i, j \in [n]$, $i \neq j$ and $|i - j| \leqslant l$. For $k \in [n] \setminus \{i, j\}$ we have

$$\mathbb{P}[CM(\rho, \tau) = \sigma_k] = \mathbb{P}[CM(\sigma, \tau) = \sigma_k] = q_{k,n}.$$

In the remaining cases $\mathbb{P}[CM(\rho, \tau) = \sigma_i] = q_{j,n}$ and $\mathbb{P}[CM(\rho, \tau) = \sigma_j] = q_{i,n}$. Note that whenever S neither contains σ_i nor σ_j, the probabilities $\mathbb{P}[CM(\sigma, \tau) \in S]$ and $\mathbb{P}[CM(\rho, \tau) \in S]$ are equal and the above inequality holds with any $\varepsilon \geqslant 0$. The

situation is analogous whenever S includes both, σ_i and σ_j. Thus let us consider S such that $\sigma_i \in S$ and $\sigma_j \notin S$ (the symmetric case with $\sigma_j \in S$ and $\sigma_i \notin S$ is analogous). We can write

$$\mathbb{P}[CM(\sigma,\tau) \in S] = q_{i,n} + q \quad \text{and} \quad \mathbb{P}[CM(\rho,\tau) \in S] = q_{j,n} + q,$$

where $q = \mathbb{P}[CM(\sigma,\tau) \in S \setminus \{\sigma_i\}] = \mathbb{P}[CM(\rho,\tau) \in S \setminus \{\sigma_i\}]$. We have

$$\frac{\mathbb{P}[CM(\sigma,\tau) \in S] - \delta}{\mathbb{P}[CM(\rho,\tau) \in S]} = \frac{q_{i,n} + q - \delta}{q_{j,n} + q} =: f(q).$$

Note that $f'(q) = \frac{q_{j,n} - q_{i,n} + \delta}{(q+q_{j,n})^2}$. Thus whenever $q_{i,n} \geqslant q_{j,n}$ and $\delta < q_{i,n} - q_{j,n}$ the function f is decreasing. In the remaining cases f is weakly increasing. Therefore if $q_{i,n} \geqslant q_{j,n}$ and $\delta < q_{i,n} - q_{j,n}$ we get

$$\frac{\mathbb{P}[CM(\sigma,\tau) \in S] - \delta}{\mathbb{P}[CM(\rho,\tau) \in S]} \leqslant \frac{q_{i,n} - \delta}{q_{j,n}}$$

which means that we maximize the expression setting $q = 0$ (i.e., setting $S = \{\sigma_i\}$). The right-hand side of the above inequality is greater than or equal to 1. In the remaining cases (when $q_{i,n} \geqslant q_{j,n}$ and $\delta \geqslant q_{i,n} - q_{j,n}$ or when $q_{j,n} > q_{i,n}$) we get

$$\frac{\mathbb{P}[CM(\sigma,\tau) \in S] - \delta}{\mathbb{P}[CM(\rho,\tau) \in S]} \leqslant \frac{1 - q_{j,n} - \delta}{1 - q_{i,n}}$$

which means that we maximize the expression setting $q = 1 - q_{i,n} - q_{j,n}$ (i.e., setting $S = \mathcal{C} \setminus \{\sigma_j\}$). Note that the right-hand side of the above inequality is smaller than or equal to 1. Thus in this case the inequality (2) holds for $\varepsilon \geq 0$.

We conclude that whenever there exists at least one pair $i, j \in [n]$ such that $i \neq j$, $|i-j| \leqslant l$, $q_{i,n} \geqslant q_{j,n}$ and $\delta < q_{i,n} - q_{j,n}$ then

$$\varepsilon \geqslant \max_{\substack{1 \leqslant i,j \leqslant n \\ |i-j| \leqslant l}} \ln\left\{\frac{q_{i,n} - \delta}{q_{j,n}}\right\}.$$

Otherwise $\varepsilon \geq 0$. Note that in the proof in both cases we have indicated S realizing the maximum. Therefore the bounds are tight. □

4 Hiding preferences

Each optimal stopping algorithm is (ε, δ)-differentially private at some level, i.e., for some values of ε and δ. These values will be often too high to meet user's expectations. What the user can do is to resign from the optimality of the algorithm (however try to keep the accuracy of the algorithm at some acceptable level) gaining a higher level of privacy. It can be achieved by a careful modification of the distribution of the outcome of the algorithm. Analyzing the definition of differential privacy one can deduce that the closer this distribution to the uniform one is, the smaller the values of ε and δ in the

definition of differential privacy may become. E.g., below in Fact 4 we explain that the algorithm with uniform outcome is $(0,0)$-differentially private regardless of the metric. Thus what the user should do in order to achieve the desired privacy level is to modify the algorithm M^{opt} such that the distribution of its outcome comes in some sense closer to the uniform distribution.

In this section we analyze a natural mechanism transforming an arbitrary optimal stopping time M^{opt} into the algorithm meeting stricter privacy requirements, yet preserving some level of accuracy. It is equipped with a parameter $p \in [0,1]$ controlling the smooth transition between optimality and $(0,0)$-DP.

Definition 7 *Let M'' and M' be two online stopping algorithms for the same stopping problem. A p-mix on M'' and M' is defined as follows. We toss a coin that comes down heads with probability p. If it comes down heads, we play according to M''. If it comes down tails, we play according to M'.*

Definition 8 *We call the algorithm $\tilde{M}$ a blind choice if for a fixed $\sigma \in \mathcal{S}_n$ and for any $\tau \in \mathcal{T}_n$ it always stops at τ_1, i.e., $C\tilde{M}(\sigma,\tau) = \tau_1$, equivalently $\tilde{M}(\sigma,\tau) = 1$.*

Fact 4 *A blind choice $\tilde{M}$ is $(0,0)$-differentially private regardless of the metric we use.*

Proof. Note that for a fixed σ and any $C \in \mathcal{C}$ we have $\mathbb{P}[C\tilde{M}(\sigma,\tau) = C] = 1/n$. Indeed, the candidate C will be selected by $\tilde{M}$ if and only if it is the first candidate in the time ordering τ. Time ordering τ is chosen uniformly at random from $\mathcal{T}_n$, thus the probability that it starts with C is $1/n$. In particular, for $S \subseteq \mathcal{C}$ we get $\mathbb{P}[C\tilde{M}(\sigma,\tau) \in S] = |S|/n$ (this probability is independent of σ). Thus for fixed $\sigma, \rho \in \mathcal{S}_n$ and for any $S \subseteq \mathcal{C}$ we always get $\mathbb{P}[C\tilde{M}(\sigma,\tau) \in S] = \mathbb{P}[C\tilde{M}(\rho,\tau) \in S] = |S|/n$ thus

$$\mathbb{P}[C\tilde{M}(\sigma,\tau) \in S] \leqslant e^0 \mathbb{P}[C\tilde{M}(\rho,\tau) \in S] + 0.$$

□

Note that if we consider a p-mix on M^{opt} and $\tilde{M}$ (i.e., a p-mix on an optimal algorithm and a blind choice) then we get a controllable by p algorithm being a trade-off between two extremes. One of them is optimality (the case when $p = 1$) and the other one is $(0,0)$-differential privacy (the case when $p = 0$). Setting higher p means relaxing the requirements for (ε,δ)-differential privacy but at the same time obtaining a larger probability of choosing the proper candidate. Setting smaller p we resign from the high accuracy of the algorithm but we gain a higher level of privacy.

Some general results on the accuracy and the privacy of a p-mix algorithm are given in Section 4.1. The detailed analysis of a classical case, i.e., of a p-mix on the optimal algorithm for the secretary problem and a blind choice, is given in Section 4.2.

4.1 Accuracy and differential privacy of a p-mix M

In this section we formulate some general results on the accuracy and the privacy of a p-mix algorithm. We start with a simple fact about the minimum level to which the accuracy of this algorithm may drop.

Fact 5 *Fix $\sigma \in \mathcal{S}_n$. Let $D \subseteq \mathcal{C}$ be the set of the desired candidates from Definition 2. Let M^{opt} be the optimal stopping time and M' any other stopping algorithm for this problem. Let $p \in [0,1]$ and M be the p-mix on M^{opt} and M'. Then*

$$\mathbb{P}[CM(\sigma,\tau) \in D] \geqslant p \cdot \mathbb{P}[CM^{opt}(\sigma,\tau) \in D].$$

Proof. Obviously, by the definition of a p-mix algorithm we get

$$\begin{aligned}\mathbb{P}[CM(\sigma,\tau) \in D] &= p \cdot \mathbb{P}[CM^{opt}(\sigma,\tau) \in D] + (1-p) \cdot \mathbb{P}[CM'(\sigma,\tau) \in D] \\ &\geqslant p \cdot \mathbb{P}[CM^{opt}(\sigma,\tau) \in D].\end{aligned}$$

□

Assume that the optimal algorithm for some stopping problem is (ε, δ) - differentially private. The following theorem explains how differential privacy improves (i.e., how parameters ε and δ drop) if we mix this optimal algorithm with a strategy whose outcome distribution is uniform.

Theorem 3. *Fix $\sigma \in S_n$. Let M^{opt} be some optimal stopping algorithm. Consider a metric space $(\mathcal{S}_n, d_l)$. Assume that M^{opt} is (ε,δ)-differentially private w.r.t. metric d_l. Let M' be the algorithm whose outcome distribution is uniform, i.e., for all $k \in [n]$ $\mathbb{P}[CM'(\sigma,\tau) = \sigma_k] = 1/n$ (e.g., it may be a blind choice $\tilde{M}$). Then the p-mix M on M^{opt} and M' is $\left(\ln\left(e^{\varepsilon} - \frac{(1-p)(e^{\varepsilon}-1)}{n}\right), p \cdot \delta\right)$-differentially private w.r.t. metric d_l.*

Proof. The algorithm M^{opt} is (ε,δ)-differentially private w.r.t. metric d_l thus for all $\rho \in \mathcal{S}_n$ such that $d_l(\sigma,\rho) \leqslant 1$ and for all $S \subseteq \mathcal{C}$ we have

$$\mathbb{P}[CM^{opt}(\sigma,\tau) \in S] \leqslant e^{\varepsilon}\mathbb{P}[CM^{opt}(\rho,\tau) \in S] + \delta.$$

Moreover, by the definition of a p-mix for every $\pi \in \mathcal{S}_n$ we get

$$\mathbb{P}[CM(\pi,\tau) \in S] = p \cdot \mathbb{P}[CM^{opt}(\pi,\tau) \in S] + (1-p) \cdot \mathbb{P}[CM'(\pi,\tau) \in S].$$

Hence for all $\rho \in \mathcal{S}_n$ such that $d_l(\sigma,\rho) \leqslant 1$ and for all $S \subseteq \mathcal{C}$ such that $S \neq \emptyset$ (if $S = \emptyset$ then the differential privacy inequality for M holds for any $\varepsilon \geq 0$ and any $\delta \geq 0$)

$$\begin{aligned}\mathbb{P}[CM(\sigma,\tau) \in S] &= p \cdot \mathbb{P}[CM^{opt}(\sigma,\tau) \in S] + (1-p) \cdot \mathbb{P}[CM'(\sigma,\tau) \in S] \\ &\leqslant p \cdot e^{\varepsilon} \cdot \mathbb{P}[CM^{opt}(\rho,\tau) \in S] + p \cdot \delta + (1-p) \cdot \mathbb{P}[CM'(\sigma,\tau) \in S] \\ &= e^{\varepsilon}(\mathbb{P}[CM(\rho,\tau) \in S] - (1-p) \cdot \mathbb{P}[CM'(\rho,\tau) \in S]) \\ &\quad + p \cdot \delta + (1-p) \cdot \mathbb{P}[CM'(\sigma,\tau) \in S] \\ &= e^{\varepsilon}\mathbb{P}[CM(\rho,\tau) \in S] + p \cdot \delta - (1-p)(e^{\varepsilon}-1)\frac{|S|}{n} \\ &\leqslant \left(e^{\varepsilon} - \frac{(1-p)(e^{\varepsilon}-1)}{n}\right)\mathbb{P}[CM(\rho,\tau) \in S] + p \cdot \delta.\end{aligned} \tag{3}$$

□

4.2 Trade-off between optimality and differential privacy in the secretary problem

This section is an analytical discussion about the optimal stopping algorithm for the classical secretary problem in the context of differential privacy. Below we present a detailed analysis of the accuracy and differential privacy of a p-mix on M^* and $\tilde{M}$, where M^* is the optimal solution for the secretary problem, and $\tilde{M}$ is the blind choice. Recall that in the secretary problem one aims at choosing only the best out of all n candidates (i.e., for a fixed $\sigma \in \mathcal{S}_n$ at selecting σ_1). Any other choice is interpreted as a loss. Let us start with a simple fact about the accuracy of the p-mix on M^* and $\tilde{M}$.

Fact 6 *Fix $\sigma \in \mathcal{S}_n$. Let M be a p-mix on M^* and $\tilde{M}$, where M^* is the optimal algorithm for the secretary problem, and $\tilde{M}$ is the blind choice. Then*

$$\mathbb{P}[CM(\sigma,\tau) = \sigma_1] = p \cdot \mathbb{P}[CM^*(\sigma,\tau) = \sigma_1] + \frac{1-p}{n} \sim \frac{p}{e}.$$

Proof. The result follows straight from the definition of a p-mix and Fact 2. □

Before we move on to analyzing differential privacy of the classical p-mix on M^* and $\tilde{M}$, let us introduce some simplifications in notation. Throughout this section $p \in [0,1]$ is always a constant and M is always a p-mix on M^* and $\tilde{M}$. (The case $p = 0$ when M is just a blind choice was already discussed thus we will often assume $p \in (0,1]$.) We also introduce a shorter notation for the probabilities that M^* or M select the k^{th} best candidate, namely, for $\sigma \in \mathcal{S}_n$

$$r_{k,n} = \mathbb{P}[CM^*(\sigma,\tau) = \sigma_k] \quad \text{and} \quad q_{k,n} = \mathbb{P}[CM(\sigma,\tau) = \sigma_k].$$

By this notation, following Definition 7, we can write

$$q_{k,n} = p \cdot r_{k,n} + \frac{1-p}{n}. \tag{4}$$

Therefore, by Theorem 6 (see the Appendix), we formulate the following corollary.

Corollary 1. *Let $\sigma \in \mathcal{S}_n$, let $p \in [0,1]$ be a constant and let M be a p-mix on M^* and $\tilde{M}$. Then for $k \geqslant 1$ being a constant or a function of n such that $k(n) = o(n)$*

$$q_{k,n} \sim \frac{p}{e} \cdot \sum_{s=k}^{\infty} \frac{1}{s} \left(1 - \frac{1}{e}\right)^s, \quad \textit{in particular} \quad q_{1,n} \sim \frac{p}{e}.$$

For transparency of notation let also

$$a_k = \sum_{s=k}^{\infty} \frac{1}{s} \left(1 - \frac{1}{e}\right)^s.$$

Note that the above sum always converges, in particular $a_1 = 1$ and $a_2 = 1/e$. By Fact 2 and Corollary 1 for $k \geqslant 1$ being a constant or a function of n such that $k(n) = o(n)$ we can simply write

$$r_{k,n} \sim \frac{1}{e} \cdot a_k \quad \text{and} \quad q_{k,n} \sim \frac{p}{e} \cdot a_k.$$

In Theorem 4 and the two following corollaries we give constraints for ε, δ and p which guarantee that a p-mix M is (ε, δ)-differentially private with respect to metric d_l. Recall that we always assume that l is a constant (consult Theorem 1). The following two technical lemmas will be helpful by proving the main theorem.

Lemma 1. *Fix $\sigma \in \mathcal{S}_n$. The sequence $\{q_{k,n}\}_{k\in[n]}$ is non-increasing in k.*

Proof. By Corollary 4 (see the Appendix) we know that the sequence $\{r_{k,n}\}_{k\in[n]}$ is non-increasing in k. By (4) we have

$$q_{k,n} - q_{k+1,n} = p \cdot (r_{k,n} - r_{k+1,n}) \geqslant 0.$$

□

Lemma 2. *Fix $\sigma \in \mathcal{S}_n$. Let $l \geqslant 1$ be a constant. Let $\delta \in [0,1]$. For any $k \in [n-1]$, if $n \geqslant 7$ then*

$$\frac{q_{1,n} - \delta}{q_{l+1,n}} \geqslant \frac{q_{k,n} - \delta}{q_{k+l,n}}.$$

Proof. First, we are going to show that for $k \geqslant 2$

$$\lim_{n\to\infty} \frac{q_{1,n}}{q_{l+1,n}} > \lim_{n\to\infty} \frac{q_{k,n}}{q_{k+l,n}}. \tag{5}$$

By (4) and by Fact 2 we have

$$\lim_{n\to\infty} \frac{q_{1,n}}{q_{l+1,n}} = \lim_{n\to\infty} \frac{r_{1,n}}{r_{l+1,n}} = \frac{1}{a_{l+1}} \quad \text{and} \quad \lim_{n\to\infty} \frac{q_{k,n}}{q_{k+l,n}} = \lim_{n\to\infty} \frac{p \cdot r_{k,n} + \frac{1-p}{n}}{p \cdot r_{k+l,n} + \frac{1-p}{n}}.$$

By Theorem 7 (see the Appendix) we know that for $k \geqslant 2$ we have $\lim_{n\to\infty} \frac{r_{k,n}}{r_{k+l,n}} < \frac{1}{a_{l+1}}$. Thus for sufficiently large n (one can verify that $n \geqslant 7$ is enough) we can write $r_{k,n} < \frac{r_{k+l,n}}{a_{l+1}}$ and thereby get (note that $a_{l+1} < 1$)

$$\frac{p \cdot r_{k,n} + \frac{1-p}{n}}{p \cdot r_{k+l,n} + \frac{1-p}{n}} < \frac{p \cdot \frac{r_{k+l,n}}{a_{l+1}} + \frac{1-p}{n}}{p \cdot r_{k+l,n} + \frac{1-p}{n}}$$

$$\xrightarrow{n\to\infty} \begin{cases} \frac{p\cdot\frac{c}{a_{l+1}}+1-p}{p\cdot c+1-p} < \frac{1}{a_{l+1}} & \text{for} \quad r_{k+l,n} = c \cdot 1/n + o(1/n), \\ 1 < \frac{1}{a_{l+1}} & \text{for} \quad r_{k+l,n} = o(1/n). \end{cases}$$

Whenever $r_{k+l,n} = \omega(1/n)$ we also have $r_{k,n} = \omega(1/n)$ (indeed, by Corollary 4 the sequence $r_{k,n}$ is non-increasing in k) and again by Theorem 7 (see the Appendix) we get

$$\lim_{n\to\infty} \frac{p \cdot r_{k,n} + \frac{1-p}{n}}{p \cdot r_{k+l,n} + \frac{1-p}{n}} = \lim_{n\to\infty} \frac{r_{k,n}}{r_{k+l,n}} < \frac{1}{a_{l+1}}.$$

We conclude that for any $k \in [n-1]$ and sufficiently large n (again $n \geqslant 7$ is enough) we get $\frac{q_{1,n}}{q_{l+1,n}} \geqslant \frac{q_{k,n}}{q_{k+l,n}}$ and, by Lemma 1,

$$\frac{q_{1,n} - \delta}{q_{l+1,n}} = \frac{q_{1,n}}{q_{l+1,n}} - \frac{\delta}{q_{l+1,n}} \geqslant \frac{q_{1,n}}{q_{l+1,n}} - \frac{\delta}{q_{k+l,n}} \geqslant \frac{q_{k,n}}{q_{k+l,n}} - \frac{\delta}{q_{k+l,n}} = \frac{q_{k,n} - \delta}{q_{k+l,n}}.$$

□

Theorem 4. *Fix $\sigma \in \mathcal{S}_n$. Consider a metric space $(\mathcal{S}_n, d_l)$ for $l \geqslant 1$ being a constant. Let $p \in (0,1]$ and $\delta \in [0,1]$. If $n \geqslant 7$ and*

$$\varepsilon \geqslant \begin{cases} \ln\left(\frac{q_{1,n}-\delta}{q_{l+1,n}}\right) \sim \ln\left(\frac{p-\delta\cdot e}{a_{l+1}\cdot p}\right) & \text{for} \quad \delta < q_{1,n} - q_{l+1,n} \\ 0 & \text{for} \quad \delta \geqslant q_{1,n} - q_{l+1,n}, \end{cases}$$

then the p-mix M is (ε,δ)-differentially private. The bounds are tight.

Proof. By Theorem 2 we know that if there exists at least one pair $i, j \in [n]$ such that $i \neq j$, $|i-j| \leqslant l$, $q_{i,n} \geqslant q_{j,n}$ and $\delta < q_{i,n} - q_{j,n}$ then

$$\varepsilon \geqslant \max_{\substack{1\leqslant i,j\leqslant n \\ |i-j|\leqslant l}} \ln\left\{\frac{q_{i,n}-\delta}{q_{j,n}}\right\}.$$

Otherwise $\varepsilon \geq 0$. Consequently, by Lemma 1 and Lemma 2 we get

$$\varepsilon \geqslant \begin{cases} \ln\left(\frac{q_{1,n}-\delta}{q_{l+1,n}}\right) & \text{for} \quad \delta < q_{1,n} - q_{l+1,n} \\ 0 & \text{for} \quad \delta \geqslant q_{1,n} - q_{l+1,n}. \end{cases}$$

Additionally, by Corollary 1 we get

$$\ln\left(\frac{q_{1,n}-\delta}{q_{l+1,n}}\right) \sim \ln\left(\frac{p-\delta\cdot e}{a_{l+1}\cdot p}\right).$$

□

Figure 1 shows the shape of the asymptotic region of pairs (ε,δ) for which the p-mix M is (ε,δ)-differentially private.

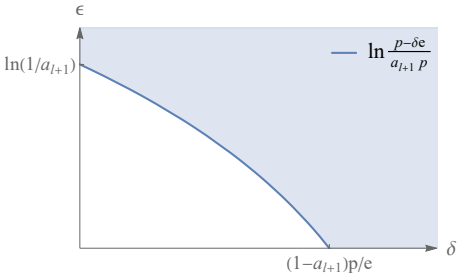

Fig. 1: Shaded area is an asymptotic region in which the p-mix M is (ε,δ)-differentially private.

Figure 2 shows how the boundaries of the asymptotic regions in which the p-mix M is (ε,δ)-differentially private change with p for a given l. Figure 3 shows how the boundaries of an asymptotic region in which the p-mix M is (ε,δ)-differentially private change with l for a given p.

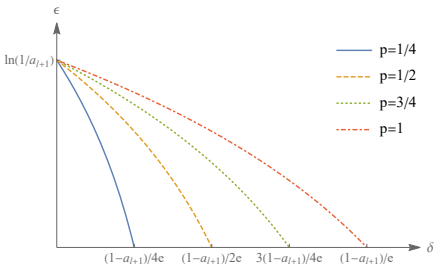

Fig. 2: Boundaries of the asymptotic regions in which the p-mix M is (ε, δ)-differentially private for a given l and various values of p.

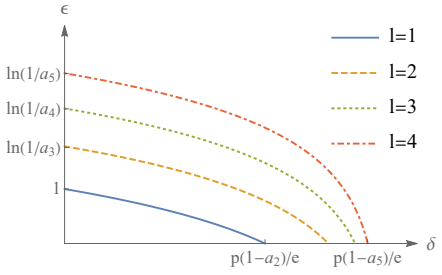

Fig. 3: Boundaries of the asymptotic regions in which the p-mix M is (ε, δ)-differentially private for a given p and various values of l.

Corollary 2. *Fix $\sigma \in \mathcal{S}_n$. Consider a metric space $(\mathcal{S}_n, d_l)$ for $l \geqslant 1$ being a constant. Let $p \in (0, 1]$ and $\varepsilon \geqslant 0$. If $n \geqslant 7$ and*

$$\delta \geqslant \begin{cases} q_{1,n} - e^{\varepsilon} q_{l+1,n} \sim \frac{p}{e}\left(1 - a_{l+1} e^{\varepsilon}\right) & \text{for} \quad \varepsilon < \ln\left(\frac{q_{1,n}}{q_{l+1,n}}\right) \\ 0 & \text{for} \quad \varepsilon \geqslant \ln\left(\frac{q_{1,n}}{q_{l+1,n}}\right), \end{cases}$$

then the algorithm M is (ε, δ)-differentially private.

Proof. Note that the inequality $\delta \geqslant q_{1,n} - e^{\varepsilon} q_{l+1,n}$ is equivalent to the inequality $\varepsilon \geqslant \ln\left(\frac{q_{1,n}-\delta}{q_{l+1,n}}\right)$ from Theorem 4 and by Corollary 1 we get $q_{1,n} - e^{\varepsilon} q_{l+1,n} \sim \frac{p}{e}\left(1 - a_{l+1} e^{\varepsilon}\right)$.

Corollary 3. *Fix $\sigma \in \mathcal{S}_n$. Consider a metric space $(\mathcal{S}_n, d_l)$ for $l \geqslant 1$ being a constant. Let $\delta \in [0, 1]$ and $\varepsilon \geqslant 0$. If $n \geqslant 7$ and*

$$p \leqslant \begin{cases} \frac{\delta + \frac{1}{n}(e^{\varepsilon} - 1)}{r_{1,n} - e^{\varepsilon} r_{l+1,n} + \frac{1}{n}(e^{\varepsilon} - 1)} \sim \frac{e \cdot \delta}{1 - e^{\varepsilon} a_{l+1}} & \text{for} \quad \delta < r_{1,n} - e^{\varepsilon} r_{l+1,n} \\ 1 & \text{for} \quad \delta \geqslant r_{1,n} - e^{\varepsilon} r_{l+1,n}, \end{cases}$$

then the p-mix M is (ε, δ)-differentially private.

Proof. By Corollary 2 we know that for $\varepsilon \geqslant 0$ the p-mix M is (ε, δ)-differentially private if only $\delta \geqslant q_{1,n} - e^{\varepsilon} q_{l+1,n}$. By (4) we can rewrite it as

$$\delta \geqslant p \cdot r_{1,n} + \frac{1-p}{n} - e^{\varepsilon}\left(p \cdot r_{l+1,n} + \frac{1-p}{n}\right)$$

which is equivalent to

$$p \leqslant \frac{\delta + \frac{1}{n}(e^{\varepsilon} - 1)}{r_{1,n} - e^{\varepsilon} r_{l+1,n} + \frac{1}{n}(e^{\varepsilon} - 1)} \sim \frac{e \cdot \delta}{1 - e^{\varepsilon} a_{l+1}},$$

where asymptotics follows from Fact 2.

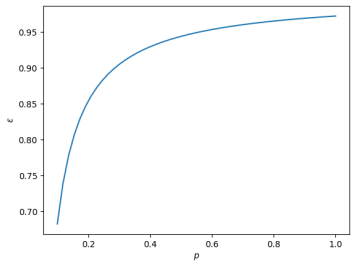

Fig. 4: Privacy parameter ε of the p-mix M for $\delta = 0.01$.

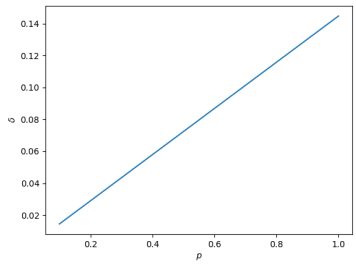

Fig. 5: Privacy parameter δ of the p-mix M for $\varepsilon = 0.5$.

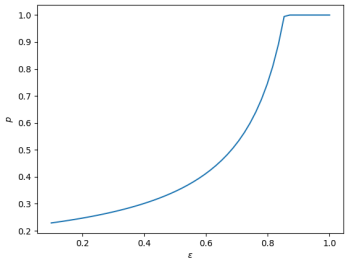

Fig. 6: Parameter p of the p-mix M for $\delta = 0.05$.

Below we present also a few explicit examples of the relations between the parameters ε, δ and p. All the examples consider asymptotic results w.r.t. metric d_1. When needed the parameters were set to $\varepsilon = 0.5$, and $\delta = 0.01$ or $\delta = 0.05$ commonly used in the analysis of differential privacy offered by a wide range of mechanisms (see e.g. [25,18]). Note however that the calibration of ε, δ parameters depends strongly on the considered scenarios. In a one-shot response mechanism such values may be adequate, however, if a given mechanism is assumed to be used many times or combined with other sources of knowledge by the adversary we often need to consider much lower values ([26]). Figure 4 shows the parameter ε as a function of p for which the p-mix M is $(\varepsilon, 0.01)$-differentially private. Note that here ε does not exceed 1 even for $p = 1$, thus when the selector plays simply the optimal algorithm. Figure 5 shows the parameter δ as a function of p for which the p-mix M is $(0.5, \delta)$-differentially private. Note that if we demand δ to be at most 0.05, then p has to be around 0.35. Figure 6 shows how small the parameter p has to be for the p-mix M to obtain $(\varepsilon, 0.05)$-differential privacy. Note that if we accept $\varepsilon \approx 0.85$ or higher, then already $p = 1$ (thus simply the optimal strategy) is sufficient.

5 Conclusions and Future Work

In this paper we have investigated the optimal stopping algorithms from the information hiding perspective. We have proposed a natural mechanism constructing a suboptimal

stopping algorithms but giving better privacy properties. We have analyzed its effectiveness and privacy, and applied it to the classical secretary problem. This work can be seen as the first step towards the differentially private stopping algorithms.

The problem we consider might, at the first glance, resemble the differentially private auction problem (see [27]). Note however, that here we do not have the scores, the preferences are given by the permutation and they are only comparable, not quantifiable, so the exponential mechanism cannot be used. Moreover, due to the online nature of the problem, we cannot have a full knowledge during the whole procedure. We also have to define a metric of similarity for inputs.

Here we have concentrated on the classical secretary problem. Nevertheless, the optimal stopping literature offers the whole variety of different models. One could, e.g., analyze the algorithm that optimizes the expected rank of the candidate (intuitively, we do not necessarily require the best candidate but at least 'good enough'). It may have better information hiding properties inherently. One could consider Gusein-Zade models with parameter k in which we win when the selected candidate is in top k. It would be challenging to work on models with different ordering, e.g., using partially ordered sets instead of linear order as a qualification ordering. However, in this case it is unclear whether a reasonable metric for the similarity of preferences could be proposed.

Our approach was to keep the final algorithm simple, so we based it on known, optimal one. Elseways, one might propose an entirely different algorithm, that is not optimal but has better information hiding properties or is more effective for given privacy parameters. Note also that essentially we have focused on hiding the information about preferences, not the participation of a specific candidate. Another venue of research could be investigating whether participation hiding is feasible in such circumstances. Even if not for all candidates, maybe it could be possible for the majority of them.

We believe that this paper opens an interesting new research area lying at the crossroads of online algorithms and differential privacy.

Acknowledgements This research was partially supported by Polish National Science Centre Grant 2018/29/B/ST6/02969.

References

1. Lindley, D.: Dynamic programming and decision theory. Appl. Stat. - J. Roy. St. C **10**(1) (1961) 39–51
2. Ferguson, T.S.: Who solved the secretary problem? Statist. Sci. **4**(3) (1989) 282–289
3. Stadje, W.: Efficient stopping of a random series of partially ordered points. Multiple Criteria Decision Making Theory and Application. Lecture Notes in Economics and Mathematical Systems **177** (1980) 430–447
4. Gnedin, A.V.: Multicriteria extensions of the best choice problem: Sequential selec- tion without linear order. Contemp. Math. **125** (1992) 153–172
5. Morayne, M.: Partial-order analogue of the secretary problem the binary tree case. Discret. Math. **184**(1-3) (1998) 165–181
6. Garrod, B., Morris, R.: The secretary problem on an unknown poset. Random Struct. Algor. **43**(4) (2013) 429–451
7. Preater, J.: The best-choice problem for partially ordered objects. Oper. Res. Lett. **25**(4) (1999) 187–190

8. Kozik, J.: Dynamic threshold strategy for universal best choice problem. Proceedings of 21st International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms (2010) 439–452
9. Freij, R., Wästlund, J.: Partially ordered secretaries. Electron. Commun. Prob. **15** (2010) 504–507
10. Babaioff, M., Immorlica, N., Kleinberg, R.: Matroids, secretary problems, and online mechanisms. In Bansal, N., Pruhs, K., Stein, C., eds.: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007, SIAM (2007) 434–443
11. Chow Y.S., Moriguti S., R.H., S.M., S.: Optimal selection based on relative rank (the "secretary problem"). Isr. J. Math. **2** (1964) 81–90
12. Kumar, R., Lattanzi, S., Vassilvitskii, S., Vattani, A.: Hiring a secretary from a poset. In Shoham, Y., Chen, Y., Roughgarden, T., eds.: Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011, ACM (2011) 39–48
13. Janson, S.: The hiring problem with rank-based strategies. Electron. J. Probab. **24** (2019) 1–35
14. Kaplan, H., Naori, D., Raz, D.: Competitive analysis with a sample and the secretary problem. In Chawla, S., ed.: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, SIAM (2020) 2082–2095
15. Correa, J., Cristi, A., Feuilloley, L., Oosterwijk, T., Tsigonias-Dimitriadis, A.: The secretary problem with independent sampling. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, SIAM (2021) 2047–2058
16. Dwork, C., McSherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: TCC. Volume 3876., Springer (2006) 265–284
17. Dwork, C.: Differential privacy. In: Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006. (2006) 1–12
18. Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science **9**(3-4) (2014) pp. 211–407
19. Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., Naor, M.: Our data, ourselves: Privacy via distributed noise generation. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer (2006) 486–503
20. Dwork, C., Lei, J.: Differential privacy and robust statistics. In: STOC. Volume 9. (2009) 371–380
21. Dwork, C., Naor, M., Pitassi, T., Rothblum, G.N.: Differential privacy under continual observation. In: Proceedings of the forty-second ACM symposium on Theory of computing, ACM (2010) 715–724
22. Narayanan, A., Shmatikov, V.: De-anonymizing social networks. In: Security and Privacy, 2009 30th IEEE Symposium on, IEEE (2009) 173–187
23. Narayanan, A., Shmatikov, V.: Myths and fallacies of personally identifiable information. Commun ACM **53**(6) (2010) 24–26
24. Rogerson, P.: Probabilities of choosing applicants of arbitrary rank in the secretary problem. J. Appl. Probab. **24**(2) (1987) 527–533
25. Zhu, T., Li, G., Zhou, W., Yu, P.S.: Differential Privacy and Applications. Volume 69 of Advances in Information Security. Springer (2017)
26. Haeberlen, A., Pierce, B.C., Narayan, A.: Differential privacy under fire. In: 20th USENIX Security Symposium, San Francisco, CA, USA, August 8-12, 2011, Proceedings, USENIX Association (2011)
27. McSherry, F., Talwar, K.: Mechanism design via differential privacy. In: 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07), IEEE (2007) 94–103

Appendix

The Appendix contains the technical results on the probabilities that the optimal stopping algorithm for the secretary problem selects the k^{th} best candidate. The proofs of Theorems 5 and 6 (however in slightly different formulations) can be found in [24].

Theorem 5. *Let M^* be the optimal stopping algorithm for the secretary problem. Fix $\sigma \in \mathcal{S}_n$. Let $k \in [n]$. The probability that M^* selects the k^{th} best candidate is given by*

$$r_{k,n} = \mathbb{P}[CM^*(\sigma,\tau) = \sigma_k] = \frac{t_n - 1}{n}\left(\sum_{i=t_n}^{n-k+1} \frac{\binom{n-k}{i-1}}{\binom{n-1}{i-1}} \frac{1}{i-1} + \frac{1}{n-1}\right),$$

where t_n is the threshold from Fact 1.

Lemma 3. *Let $r_{k,n}$ be defined as in Theorem 5. Then for $k \in [n-1]$*

$$r_{k+1,n} = r_{k,n} - \frac{t_n - 1}{n} \cdot \frac{1}{k} \cdot \frac{\binom{n-t_n+1}{k}}{\binom{n-1}{k}}.$$

Proof. Since $\frac{\binom{n-k-1}{i-2}}{\binom{n-2}{i-2}} = \frac{\binom{n-i}{k-1}}{\binom{n-2}{k-1}}$ we have

$$\begin{aligned}
r_{k,n} - r_{k+1,n} &= \frac{t_n - 1}{n} \sum_{i=t_n}^{n-k+1} \frac{\binom{n-k}{i-1} - \binom{n-k-1}{i-1}}{\binom{n-1}{i-1}} \frac{1}{i-1} \\
&= \frac{t_n - 1}{n} \sum_{i=t_n}^{n-k+1} \frac{\binom{n-k-1}{i-2}}{\binom{n-1}{i-1}} \frac{1}{i-1} = \frac{t_n - 1}{n} \sum_{i=t_n}^{n-k+1} \frac{\binom{n-k-1}{i-2}}{\binom{n-2}{i-2}} \frac{i-1}{n-1} \frac{1}{i-1} \\
&= \frac{t_n - 1}{n} \frac{1}{n-1} \sum_{i=t_n}^{n-k+1} \frac{\binom{n-i}{k-1}}{\binom{n-2}{k-1}} = \frac{t_n - 1}{n} \frac{1}{n-1} \frac{\binom{n-t_n+1}{k}}{\binom{n-2}{k-1}} = \\
&= \frac{t_n - 1}{n} \cdot \frac{1}{k} \cdot \frac{\binom{n-t_n+1}{k}}{\binom{n-1}{k}}.
\end{aligned}$$

Corollary 4. *The sequence $r_{k,n}$ is non-increasing in k.*

Proof. By Lemma 3 we get

$$r_{k,n} - r_{k+1,n} = \frac{t_n - 1}{n} \cdot \frac{1}{k} \frac{\binom{n-t_n+1}{k}}{\binom{n-1}{k}} \geqslant 0.$$

Theorem 6. *Let $k \geqslant 1$ be a constant or a function of n such that $k(n) = o(n)$. Let $r_{k,n}$ be defined as in Theorem 5. Then*

$$r_{k,n} \sim \frac{1}{e}\left(1 - \sum_{s=1}^{k-1} \frac{1}{s}\left(1 - \frac{1}{e}\right)^s\right) = \frac{1}{e}\sum_{s=k}^{\infty} \frac{1}{s}\left(1 - \frac{1}{e}\right)^s.$$

Lemma 4. *Let $r_{k,n}$ be defined as in Theorem 5. Let $k = k(n) \leqslant n$ be a function linear in n. Then*

$$r_{k,n} \sim \frac{1}{e} \cdot \frac{1}{n}.$$

Proof. Recall that

$$r_{k,n} = \frac{t_n - 1}{n} \sum_{i=t_n}^{n-k+1} \frac{\binom{n-k}{i-1}}{\binom{n-1}{i-1}} \frac{1}{i-1} + \frac{t_n - 1}{n} \frac{1}{n-1}.$$

Note that if $k > n - t_n + 1$ then

$$\sum_{i=t_n}^{n-k+1} \frac{\binom{n-k}{i-1}}{\binom{n-1}{i-1}} \frac{1}{i-1} = 0$$

and by Fact 1

$$r_{k,n} = \frac{t_n - 1}{n} \frac{1}{n-1} \sim \frac{1}{e} \cdot \frac{1}{n}.$$

Hence assume that $k \leqslant n - t_n + 1$. Since $\frac{t_n - 1}{n} \sim \frac{1}{e}$ and $\frac{t_n - 1}{n} \frac{1}{n-1} \sim \frac{1}{e} \cdot \frac{1}{n}$, we need to show that $\sum_{i=t_n}^{n-k+1} \frac{\binom{n-k}{i-1}}{\binom{n-1}{i-1}} \frac{1}{i-1}$ is asymptotically smaller than $\frac{1}{n}$. Seeing that $\frac{\binom{n-k}{i-1}}{\binom{n-1}{i-1}} = \frac{\binom{n-i}{k-1}}{\binom{n-1}{k-1}}$ and that the function $f(i) = \frac{\binom{n-i}{k-1}}{\binom{n-1}{k-1}} \frac{1}{i-1}$ is decreasing in i we have

$$\begin{aligned} \sum_{i=t_n}^{n-k+1} \frac{\binom{n-k}{i-1}}{\binom{n-1}{i-1}} \frac{1}{i-1} = \sum_{i=t_n}^{n-k+1} \frac{\binom{n-i}{k-1}}{\binom{n-1}{k-1}} \frac{1}{i-1} &\leqslant \frac{n - t_n - k + 2}{t_n - 1} \frac{\binom{n-t_n}{k-1}}{\binom{n-1}{k-1}} \\ &= \frac{n - t_n - k + 2}{t_n - 1} \frac{n}{n - t_n + 1} \frac{\binom{n-t_n+1}{k}}{\binom{n}{k}}. \end{aligned} \tag{6}$$

Since $k(n) = c \cdot n$ for some constant c we get

$$\frac{n - t_n - k + 2}{t_n - 1} \frac{n}{n - t_n + 1} \sim \frac{e(1-c) - 1}{1 - 1/e}$$

and at the same time

$$\begin{aligned} \frac{\binom{n-t_n+1}{k}}{\binom{n}{k}} = \left(1 - \frac{k}{n - t_n + 2}\right) \dots \left(1 - \frac{k}{n}\right) &\leqslant \left(1 - \frac{k}{n}\right)^{t-1} \\ &\sim (1-c)^{n/e-1} = o(1/n). \end{aligned}$$

Lemma 5. *Let $r_{k,n}$ be defined as in Theorem 5. For $k \in [n-2]$*

$$\lim_{n\to\infty} \frac{r_{k,n}}{r_{k+1,n}} \geqslant \lim_{n\to\infty} \frac{r_{k+1,n}}{r_{k+2,n}}.$$

Proof. First note that when $k = k(n)$ is a function linear in n, by Lemma 4 we obtain

$$\lim_{n\to\infty} \frac{r_{k,n}}{r_{k+1,n}} = \lim_{n\to\infty} \frac{r_{k+1,n}}{r_{k+2,n}} = 1.$$

Hence assume that k is either a constant or $k(n) = o(n)$. Let $a_k = \sum_{s=k}^{\infty} \frac{1}{s}\left(1-\frac{1}{e}\right)^s$. By Theorem 5 we can write

$$\frac{r_{k,n}}{r_{k+1,n}} \sim \frac{a_k}{a_{k+1}} \quad \text{and} \quad \frac{r_{k+1,n}}{r_{k+2,n}} \sim \frac{a_{k+1}}{a_{k+2}}.$$

Let $\beta_k = \frac{1}{k}\left(1-\frac{1}{e}\right)^k$. Thus we have to prove $\frac{a_k}{a_k-\beta_k} \geqslant \frac{a_k-\beta_k}{a_k-\beta_k-\beta_{k+1}}$ which is equivalent to $a_k \geqslant \frac{\beta_k^2}{\beta_k-\beta_{k+1}}$. Let $f(k) = \frac{\beta_k^2}{\beta_k-\beta_{k+1}}$. We need to show that

$$\sum_{s=k}^{\infty} \frac{1}{s}\left(1-\frac{1}{e}\right)^s \geqslant f(k). \tag{7}$$

One can easily verify that for all $s \geqslant 1$

$$\frac{1}{s}\left(1-\frac{1}{e}\right)^s \geqslant f(s) - f(s+1).$$

Summing both sides of the above inequality over $s \geqslant k$ we obtain

$$\sum_{s=k}^{\infty} \frac{1}{s}\left(1-\frac{1}{e}\right)^s \geqslant \sum_{s=k}^{\infty}(f(s)-f(s+1)) = f(k),$$

where the last equality follows from the fact that $f(n) \xrightarrow{n\to\infty} 0$.

Theorem 7. *Let $l \geqslant 1$ be a constant and let $k \in \{2,\ldots,n-l\}$. Let $r_{k,n}$ be defined as in Theorem 5 and let also $a_l = \sum_{s=l}^{\infty} \frac{1}{s}\left(1-\frac{1}{e}\right)^s$. Then*

$$\lim_{n\to\infty} \frac{r_{k,n}}{r_{k+l,n}} < \frac{1}{a_{l+1}}.$$

Proof. By Lemma 5 we have

$$\begin{aligned}\lim_{n\to\infty} \frac{r_{k+1,n}}{r_{k+l,n}} &= \lim_{n\to\infty} \frac{r_{k+1,n}}{r_{k+2,n}} \cdot \lim_{n\to\infty} \frac{r_{k+2,n}}{r_{k+3,n}} \cdots \lim_{n\to\infty} \frac{r_{k+l-1,n}}{r_{k+l,n}} \\ &\leqslant \lim_{n\to\infty} \frac{r_{2,n}}{r_{3,n}} \cdot \lim_{n\to\infty} \frac{r_{3,n}}{r_{4,n}} \cdots \lim_{n\to\infty} \frac{r_{l,n}}{r_{l+1,n}} = \lim_{n\to\infty} \frac{r_{2,n}}{r_{l+1,n}}.\end{aligned}$$

Additionally, for $k \geqslant 2$, by Theorem 5 and Lemma 5

$$\lim_{n\to\infty} \frac{r_{k,n}}{r_{k+1,n}} < \lim_{n\to\infty} \frac{r_{1,n}}{r_{2,n}}.$$

Together, by Theorem 5, it gives

$$\begin{aligned}\lim_{n\to\infty} \frac{r_{k,n}}{r_{k+l,n}} &= \lim_{n\to\infty} \frac{r_{k,n}}{r_{k+1,n}} \cdot \lim_{n\to\infty} \frac{r_{k+1,n}}{r_{k+l,n}} \\ &< \lim_{n\to\infty} \frac{r_{1,n}}{r_{2,n}} \lim_{n\to\infty} \frac{r_{2,n}}{r_{l+1,n}} = \lim_{n\to\infty} \frac{r_{1,n}}{r_{l+1,n}} = \frac{1}{a_{l+1}}.\end{aligned}$$