

Robust Market Design with Opaque Announcements

Aram Grigoryan* Markus Möller†

August 2024

Abstract

We introduce a framework where the announcements of a clearinghouse about the allocation process are *opaque* in the sense that there can be more than one outcome compatible with a realization of type reports. We ask whether desirable properties can be ensured under opacity in a robust sense. A property can be guaranteed under an opaque announcement if every mechanism compatible with it satisfies the property. We find an impossibility result: strategy-proofness cannot be guaranteed under *any* level of opacity. In contrast, in some environments, weak Maskin monotonicity and non-bossiness can be guaranteed under opacity.

1 Introduction

Public announcements of clearinghouses often comprise a description of the allocation process (e.g., a pseudo-algorithm, a list of rules, or illustrative examples). Typically, such announcements aim to outline the outcomes one may encounter in different scenarios. Yet some details of the process often remain unspecified. For instance, Google

*Department of Economics, University of California - San Diego, a2grigoryan@ucsd.edu

†Department of Economics, University of Bonn, mmoelle2@uni-bonn.de

did not entirely explain the determination of price floors in its online ads auctions and then deployed hidden programs that adjusted these floors in their favor.^{1,2}

While opacity in announcements may be intentional or self-serving for the clearing-house, another reason might just be that the process is complicated, poorly understood, or certain variables are unclear. Such scenarios may, for instance, occur in the context of reserves-based policies, which usually do not communicate the seat-processing order of reserves. In fact, Boston Public Schools faced controversies when their processing orders of reserved- and non-reserved slots were not specified. At the same time, the school district officials may not have been aware that this processing order affected the outcomes (Dur, Kominers, Pathak, & Sönmez, 2018).³ Another common source of opacity is unspecified capacity in the context of matching applications such as college admission, school choice, and course allocation. Here, the exact numbers of school-, college- or course seats are usually not communicated to applicants. Even the school districts or the colleges frequently do not know their final capacities *a priori*. Interestingly, the capacities may also be adjusted within the admissions cycles depending on applicants' reports (Bobbio, Carvalho, Lodi, Rios, & Torrico, 2023; Dur & Van der Linden, 2021).⁴

In light of these observations, it seems interesting to ask whether specific desirable properties can still be assured under opacity. Indeed, some holistic statements about an institution (or process) may no longer be reliable in the presence of opacity. For instance, as recently put by a Google employee about *Dynamic Revenue Share (DRS)* and *Reserve Price Optimization (RPO)* programs.⁵

¹Allegedly, Google used advertisers' current or historic bid observations without the publishers' or advertisers' knowledge to adapt the price floors and effective bids.

²<https://www.wired.com/story/google-antitrust-ad-market-lawsuit/>. Last accessed, July 2024.

³The lack of transparency and misconceptions in the reserves-processing order has been highlighted in other applications including US H-1B visa allocation (Pathak, Rees-Jones, & Sönmez, 2022) and admission to technical universities in India (Aygün & Turhan, 2022; Sönmez & Yenmez, 2019).

⁴For instance, the *Course Match* program that assigns MBA students to courses (Budish, Cachon, Kessler, & Othman, 2017) has used students' preference reports to cancel entire courses or increase course capacities *on-line*, and re-run its algorithm to reoptimize the allocation for all students (Dur & Van der Linden, 2021).

⁵<https://www.wired.com/story/google-antitrust-ad-market-lawsuit/>. Last accessed, July

“One known issue with the current DRS is that it makes the auction untruthful. ... Doesn’t that [RPO] undermine the whole idea of second-price auctions? It’ll transform the system into a 1st price auction where the bidder has a strong incentive to bid LESS than he’s willing to pay.”

A central desirable property in various allocation problems is strategy-proofness. Several school districts using the *Student-Proposing Deferred Acceptance* (Gale & Shapley, 1962) essentially mention on their websites that their admission system is ‘strategy-proof’. The *New York City Public School 2024 Admission Guides* (NYC-Public-Schools, 2024) writes:

“... Place the programs on your application in your true order of preference, with your favorite program as #1. There is no better strategy!”

A similar practice is common in other matching applications such as the Israel “*Mechinot*” match (Gonczarowski, Kovalio, Nisan, & Romm, 2019), and *National Resident Matching Program (NRMP)* (NRMP, 2023). Motivated by these examples, we take a robust approach to examine whether certain desirable properties, such as strategy-proofness, can be guaranteed under opaque announcements.

We employ a general social choice problem where a central planner has to select an outcome for a set of individuals for each realization of reports (which we call a problem). The central planner makes an *announcement*, which is a description of how problems translate into outcomes. The announcement may be *opaque* in that the description may specify more than one possible outcome for some problems. We entertain the possibility that the central planner may induce any direct mechanism compatible with the description. Concretely, a mechanism is *possible* if, for every problem, the mechanism’s chosen outcome is in the set of possible outcomes specified by the announcement for that problem. Under this robust approach, an opaque announcement *guarantees* a property if every possible mechanism satisfies the property.⁶

2024.

⁶We treat holistic statements about desirable properties of the central authority as pure cheap talk. Hence, in our framework, it is unnecessary to encompass holistic statements about desirable properties themselves. In fact, a clear disentanglement from the announcement is what allows us to better isolate the effect of opacity on the reliability of such holistic statements.

To illustrate our opacity framework, consider a simple problem of allocating two objects, 1 and 2, to two individuals, Ann (A) and Bob (B). There are only two possible outcomes: let x denote the outcome where A receives object 1 and B receives object 2, and let y denote the outcome where A receives object 2 and B receives object 1. There are four potential problems, (1) both A and B prefer x to y , (2) A prefers x to y , and B prefers y to x , (3) A prefers y to x , and B prefers x to y , (4) both A and B prefer y to x .

In a typical market design framework, participants would know the exact mechanism and how it translates each of the four potential problems to a unique outcome. By contrast, in our framework, announcements can be opaque, meaning that for some problems, the description allows a non-singleton set of possible outcomes. Concretely, suppose that in the simple problem described above, the announcement is: “Our algorithm maximizes the number of individuals who obtain their highest-ranked object.” Then, the set of possible outcomes are the following:

1. if both A and B prefer x to y , then the unique possible outcome is x ,
2. if A prefers x to y , and B prefers y to x , then the set of possible outcomes is $\{x, y\}$,
3. if A prefers y to x , and B prefers x to y , then the set of possible outcomes is $\{x, y\}$,
4. if both A and B prefer y to x , then the unique possible outcome is y .

We now entertain the possibility that any mechanism that is consistent with the announcement may be induced and then ask whether strategy-proofness can be guaranteed for each of them.

Specifically, there are four possible mechanisms corresponding to the different selections of possible outcomes at problems (2) and (3). Let us call the four mechanisms $(x - x)$, $(x - y)$, $(y - x)$, and $(y - y)$. Here, $(x - x)$ denotes the mechanism that chooses x at both problems (2) and (3), $(x - y)$ denotes the mechanism that chooses x at problem (2) and chooses y at problem (3), and so on. One can now verify that

all four possible mechanisms are strategy-proof. In fact, a closer investigation reveals that $(x - x)$ is a *Top Trading Cycles (TTC)* mechanism (Shapley & Scarf, 1974) where A owns object 1 and B owns object 2, $(x - y)$ is a *Serial Dictatorship (SD)* (Satterthwaite & Sonnenschein, 1981; Svensson, 1999) where A is the first dictator, $(y - x)$ is a *SD* where B is the first dictator, and $(y - y)$ is a *TTC* where A owns 2 and B owns 1. Hence, in this example, we have a robust guarantee of strategy-proofness. Yet, the fact that we have only two outcomes plays a crucial role in this possibility.

This is shown in our main impossibility result: In a setting with at least three outcomes, strategy-proofness cannot be guaranteed for *any* level of opacity (Theorem 1). The result highlights the importance of full transparency for robust guarantees of strategy-proofness in real-life social choice problems. We also show that *weak Maskin monotonicity* (Kojima & Manea, 2010) can be guaranteed under opaque announcements (Theorems 2). The same holds for the well-known invariance property *non-bossiness* (Satterthwaite & Sonnenschein, 1981), as established in Theorem 3.

The paper is organized as follows: In section 2, we introduce the basic opacity framework. In section 3, we present our main result. Section 4 provides the possibility results for other properties. Section 5 discusses the relation to the literature. Section 6 concludes.

2 Framework

There is a finite set of individuals I and a finite set of outcomes X where $|X| = N \geq 3$. Each individual $i \in I$ has a weak preference ranking R_i , a complete and transitive binary relation over X . For two outcomes $x, y \in X$ and an individual $i \in I$, we use xP_iy to denote xR_iy and $\neg yR_ix$.

Let \mathcal{R}_X be the space of all possible preference profiles $R = (R_i)_{i \in I}$ over X that satisfy the following condition: for any $R \in \mathcal{R}_X$ and any $x, y \in X, x \neq y$, there is an $i \in I$ such that xP_iy or yP_ix . In other words, we assume that there are no two outcomes x and y such that every individual is indifferent between x and y .⁷

⁷This is without loss of generality since, in the case of a complete indifference between x and y ,

We say that a subset $\mathcal{R} \subseteq \mathcal{R}_X$ is *rich*, if whenever there is a preference profile $R \in \mathcal{R}$, an individual $i \in I$, and two outcomes $x, y \in X$, such that $x P_i y$, then there is another preference ranking $R'_i \neq R_i$ of i , such that $(R'_i, R_{-i}) \in \mathcal{R}$, and $x P_i z \implies x P'_i z, \forall z \in X$.

An *environment* is a triplet (I, X, \mathcal{R}) , where $\mathcal{R} \subseteq \mathcal{R}_X$ is an arbitrary rich subset of preference profiles.⁸ We refer to an element $R \in \mathcal{R}$ as a problem. Note that this model covers a wide range of applications, including the classic social choice setup (Arrow, 1951), matching with contracts (Hatfield & Milgrom, 2005), and constrained object allocation (Ahn & Root, 2023), among many others.⁹

The following definitions will be the building blocks of our framework. Fix an arbitrary environment (I, X, \mathcal{R}) . An **announcement** $\Pi : \mathcal{R} \mapsto 2^X \setminus \{\emptyset\}$ is a correspondence that describes the set of possible outcomes that can be induced. A (direct) mechanism is a mapping $\varphi : \mathcal{R} \rightarrow X$ that specifies an outcome for each problem. Given an announcement Π , a **possible mechanism** is a mechanism φ such that $\varphi(R) \in \Pi(R)$ for all $R \in \mathcal{R}$. Let $\Phi(\Pi)$ denote the set of all possible mechanisms given Π . We are ready to define what we mean by opacity in this setup.

We say that an announcement $\Pi : \mathcal{R} \mapsto 2^X \setminus \{\emptyset\}$ is **fully transparent** if $\Pi(R)$ is a singleton for every $R \in \mathcal{R}$. If Π is not fully transparent, then we say that Π is **opaque**. If Π is fully transparent, then $\Phi(\Pi)$ contains a single mechanism. Conversely, if an announcement Π is opaque, then $|\Phi(\Pi)| > 1$. We say that an announcement Π **guarantees** a property if every $\varphi \in \Phi(\Pi)$ satisfies the property. The following property will be central to our analysis. A mechanism $\varphi : \mathcal{R} \rightarrow X$ is **strategy-proof** if for every $i \in I$, and every $R, R' \in \mathcal{R}$, where R' differs from R only with respect to the preference ranking of i , we have that $\varphi(R) R_i \varphi(R')$.

we can relabel x and y , and treat them as a single outcome. In that case, opacity would mean that there should be at least two *distinct* possible outcomes at some problem. In this scenario our results would carry through.

⁸We require a richness assumption to exclude the possibility of trivial settings. For example, when \mathcal{R} is a singleton, strategy-proofness and all other studied properties hold trivially. Also, note that all studied social choice environments with at least three alternatives satisfy our richness condition.

⁹For example, in the matching with contracts environment (Hatfield & Milgrom, 2005), \mathcal{R} corresponds to all preferences profiles where every agent is indifferent between two outcomes that give them the same match and the same contract. Matching with contracts covers various applications, such as object allocation without money or auctions with discrete-valued bids.

In the rest of this work, we study whether some desirable properties of mechanisms can be guaranteed under opaque announcements.

3 Main Result

This section presents our main result, which shows that strategy-proofness cannot be guaranteed under any opaque announcement.

Theorem 1. *Let (I, X, \mathcal{R}) be an arbitrary environment. An announcement Π guarantees strategy-proofness if and only if Π is fully transparent, and the unique possible mechanism given Π is strategy-proof.*

Proof. The ‘if’ direction is trivial. In what follows, we prove the ‘only if’ direction.

Suppose that an announcement Π guarantees strategy-proofness. If Π is fully transparent, and the unique possible mechanism given Π is not strategy-proof, then Π does not guarantee strategy-proofness. Therefore, if Π is fully transparent, the unique possible mechanism should be strategy-proof. Hence, to prove the ‘only if’ direction, it is left to show that when Π guarantees strategy-proofness, it is fully transparent.

The proof is by contraposition. We will assume that Π is opaque and show that Π does not guarantee strategy-proofness.

Since Π is opaque, there is a problem R , such that $\Pi(R)$ is non-singleton. Let x and y be two different outcomes in $\Pi(R)$. Let i be an individual who is not indifferent between x and y . Such an individual exists because we assumed that there are no two outcomes towards which all individuals are indifferent. Without loss of generality, suppose that xP_iy . Consider two possible mechanisms $\varphi, \varphi' \in \Phi(\Pi)$ that only differ from each other by that $\varphi(R) = x$ and $\varphi'(R) = y$, and they agree for all other problems, i.e., for all $R' \neq R$, $\varphi(R') = \varphi'(R')$. To prove that Π does not guarantee strategy-proofness, it is sufficient to show that at least one of the mechanisms in $\{\varphi, \varphi'\}$ is not strategy-proof. Without loss of generality, we will assume that φ' is strategy-proof, and we will show that φ cannot be strategy-proof.

Consider an arbitrary preference ranking $R' \neq R$ that differs from R only with respect to the ranking of individual i , and it does not rank any new objects weakly above x . That is, for all $z \in X$, $xP_i z \implies xP'_i z$. By our richness assumption, such an R' exists.

Since φ' is strategy-proof, and $xP_i y = \varphi'(R)$, it should be that $xP_i \varphi'(R')$.

Because $xP_i \varphi'(R')$, and by our choice of R' , we have that $xP'_i \varphi'(R')$. Moreover, since φ and φ' agree on all problems other than R , we have that $\varphi'(R') = \varphi(R')$. Thus,

$$\varphi(R) = xP'_i \varphi'(R') = \varphi(R'),$$

which implies that φ is not strategy-proof. This completes the proof of Theorem 1. \square

4 Other Properties

This section shows that some monotonicity and invariance properties can be guaranteed under opaque announcements.

We start with **weak Maskin monotonicity** (Kojima & Manea, 2010). We say a problem $R' \in \mathcal{R}$ is a *monotonic transformation* of problem $R \in \mathcal{R}$ at outcome $x \in X$, if for all $i \in I$ and all $y \in X$, we have that $xP_i y \implies xP'_i y$. We say a mechanism $\varphi : \mathcal{R} \rightarrow X$ is **weakly Maskin monotonic**, if for any two problems $R, R' \in \mathcal{R}$, if R' is a monotonic transformation of R at outcome $\varphi(R)$, then $\varphi(R')R'_i \varphi(R)$ for all $i \in I$.

Theorem 2. *For some environments, there are opaque announcements that guarantee weak Maskin monotonicity.*

Proof. Consider an environment with a single individual and strict preferences. Since the preferences are strict for this individual, we will refer to a generic problem with P rather than R .

Suppose the outcomes are indexed, i.e., $X = \{x_1, x_2, \dots, x_N\}$. Throughout the proof, let \bar{P} denote the preference ranking with $x_1 \bar{P} x_2 \bar{P} \dots \bar{P} x_N$.

Take any P , such that $P \neq \bar{P}$. Then, at least one object changes its rank under P relative to \bar{P} . We introduce the following notation. Let $X(P)$ be the set of outcomes

that strictly improve their rank at P compared to \bar{P} , that is,

$$X(P) = \left\{ x \in X : |\{y \in X : yPx\}| < |\{y \in X : y\bar{P}x\}| \right\}.$$

Let $x(P)$ be the highest ranked outcome according to P at $X(P)$, that is,

$$x(P) = \arg \min_{x \in X(P)} |\{y \in X : xPy\}|.^{10}$$

Our goal is to construct an opaque announcement that guarantees weak Maskin monotonicity.

Consider the opaque announcement $\Pi : \mathcal{R} \mapsto 2^X \setminus \emptyset$, where

- for $P = \bar{P}$, we set $\Pi(P) = \{x_{N-1}, x_N\}$ and,
- for any $P \neq \bar{P}$, we set $\Pi(P) = \{x(P)\}$.

This implies there are only two possible mechanisms given Π , i.e., $\Phi(\Pi) = \{\varphi, \psi\}$. Moreover, these two mechanisms differ only with respect to the outcome they choose at \bar{P} . Specifically, let φ denote the mechanism with $\varphi(\bar{P}) = x_{N-1}$ and let ψ denote the mechanism with $\psi(\bar{P}) = x_N$. To finish the prove, we must establish that both φ and ψ are weakly Maskin monotonic.

Here, we only prove the weak Maskin monotonicity of φ . The proof for weak Maskin monotonicity of ψ is almost identical and therefore omitted.

Take any pair $P, P' \in \mathcal{R}$. To show that φ is weakly Maskin monotonic, we consider two cases.

Case 1. $P \neq \bar{P}$ and $P' \neq \bar{P}$.

Without loss of generality, suppose that P' is a monotonic of transformation of P at outcome $\varphi(P)$.

Let $x_n = \varphi(P) = x(P)$. By definition of $x(P)$, it should be that either (i) $x_n \neq x_1$ and x_n is the highest ranked outcome according to P , or (ii) there is an index $m < n - 1$,

¹⁰For example, if $x_1Px_5Px_4Px_2Px_3$, then $X(P) = \{x_5, x_4\}$ and $x(P) = x_5$.

such that for all $m' < m'' \leq m$, we have that $x_{m'}Px_{m''}Px_n$, and for all $m''' > m$, we have that $x_nPx_{m'''}$. To put it simply, in the latter case, the preference ranking P is of the form $x_1Px_2 \dots Px_{m-1}Px_mPx_n \dots$ for some $m < n - 1$.

We consider cases (i) and (ii) separately.

(i) Suppose $x_n \neq x_1$ and x_n is the highest ranked outcome according to P . Since P' is a monotonic transformation of P at outcome x_n , it should be that x_n is also the highest ranked outcome according to P' . Hence, by definition, $x(P') = x_n$, which is trivially weakly more preferred than x_n according to P' .

(ii) Suppose there is an index $m < n - 1$, such that for all $m' < m'' \leq m$, we have that $x_{m'}Px_{m''}Px_n$, and for all $m''' > m$, we have that $x_nPx_{m'''}$.

Since P' is a monotonic transformation of P , it should be that no outcome with index strictly larger than m is ranked higher than x_n under P' . Consider the set $X(P')$. By the previous observation, x_n can only improve its rank when we go from P to P' , and therefore $x_n \in X(P)$ implies that $x_n \in X(P')$. By definition, $\varphi(P') = x(P')$ is the highest ranked outcome at $X(P')$ according to P' , and therefore, it is weakly more preferred than x_n according to P' .

Case 2. Consider two problems P and P' , such that $P = \bar{P}$ and $P' \neq \bar{P}$.

We first show that $P = \bar{P}$ cannot be a monotonic transformation of P' at $\varphi(P')$.

Let $x_n = \varphi(P') = x(P')$. By definition of $x(P')$, we have that x_n is ranked at a strictly higher position at P' than under \bar{P} . Therefore, it is immediate that \bar{P} is not a monotonic transformation of P' at outcome x_n .

Finally, suppose that P' is a monotonic transformation of $P = \bar{P}$ at $\varphi(\bar{P})$. By definition of φ , we have that $\varphi(\bar{P}) = x_{N-1}$.

Let $x_n = \varphi(P') = x(P')$. By definition of $x(P')$, it should be that either (i) $x_n \neq x_1$ and x_n is the highest ranked outcome according to P' , or (ii) there is an index $m < n - 1$, such that for all $m' < m'' \leq m$, we have that $x_{m'}P'x_{m''}P'x_n$, and for all $m''' > m$, we have that $x_nP'x_{m'''}$, that is, the preference ranking P' is of the form

$x_1 P' x_2 \dots P' x_{m-1} P' x_m P' x_n \dots$ for some $m < n - 1$.

In either case, it is immediate that x_n is weakly more preferred than x_{N-1} according to P' . (Note that it can also be that $x_n = x_{N-1}$). This completes the proof of Theorem 2. \square

Finally, we adapt non-bossiness as defined by [Satterthwaite and Sonnenschein \(1981\)](#) to our general setup. A mechanism φ is **non-bossy**, if for any two problems $R, R' \in \mathcal{R}$, that only differ with respect to the preferences of an individual $i \in I$, it should be that whenever $\varphi(R) P_j \varphi(R')$ or $\varphi(R') P_j \varphi(R)$ for some $j \neq i$, then $\varphi(R) P_i \varphi(R')$ or $\varphi(R') P_i \varphi(R)$. In other words, non-bossiness ensures that by changing her preference report, an individual cannot harm or benefit others without harming or benefiting herself.¹¹

Theorem 3. *For some environments, there are opaque announcements that guarantee non-bossiness.*

Proof. First, note that non-bossiness holds for any environment with a single individual. In what follows, we prove a stronger statement. Namely, we show that opaque announcements can guarantee non-bossiness in environments with any number of individuals. Consider an environment (I, X, \mathcal{R}) that satisfies the following condition: there are some two outcomes $x, y \in X$, such that for any $R \in \mathcal{R}$ and for any $i \in I$, we have that $x P_i y$ or $y P_i x$.¹²

Consider an arbitrary opaque announcement $\Pi : \mathcal{R} \mapsto 2^X \setminus \emptyset$, such that $\Pi(P) \subseteq \{x, y\}$ for all $P \in \mathcal{R}$. Then, by construction, Π is non-bossy. \square

¹¹In the object allocation setup with strict preferences, our definition boils down to the non-bossiness condition by [Satterthwaite and Sonnenschein \(1981\)](#) and [Svensson \(1999\)](#).

¹²For example, the object allocation setup with strict preferences satisfies this condition. There, x and y need to be two allocations that assign different objects to all individuals.

5 Related Literature

To the best of our knowledge, ours is the first paper to formalize opacity and its conflict with guaranteeing other normative properties in social choice and market design. In that regard, our research question and results have no close counterpart in the existing economics literature.

Our work is related to the mechanism design literature on partial or limited commitment. Given a partial description of a mechanism, the designer or central planner cannot commit to a particular selection rule and can vary the selection of the mechanism based on agents' type reports. Unlike the mechanism design literature on limited commitment (Baliga, Corchon, & Sjöström, 1997; Bester & Strausz, 2000, 2001; Doval & Skreta, 2022), we do not study central planner's preferences and credible selections. Instead, we take a robust approach where the agents entertain the possibility of every mechanism selection by the planner. Our main question is different from those in all these other works; we are interested in guaranteeing a property of a mechanism in a robust sense, namely, that every possible selection from a set of mechanisms satisfies the property.

Other notions of opacity have been studied in some economic environments. Haupt and Hitzig (2023) consider opacity of contractual terms in a moral hazard model. Stahl and Strausz (2017) examine some form of opacity in the context of product quality with and without third-party certification. In a recent empirical paper, Kapor (2024) investigates the role of opacity in college admission criteria for different applicant types.

More broadly, our paper contributes to the literature on transparency in allocation and social choice problems (Akbarpour & Li, 2020; Grigoryan & Möller, 2023, 2024; Hakimov & Raghavan, 2023; Li, 2017; Möller, 2022). There is a publicly announced and fully known mechanism in these papers, and there are concerns that the designer or central planner may *deviate* from the mechanism subject to not being detected by the participants. In our setup, the announcement is not a mechanism; hence, our paper captures a different dimension of non-transparency in these problems.

Finally, our work is also related to the literature on simplicity in market design and

matching (Li, 2017; Pycia & Troyan, 2023). These papers analyze incentives when agents have a limited understanding of the mechanism’s working. Interestingly, the variants of strategy-proofness considered in these papers are compatible with these frameworks’ respective forms of uncertainty and partial commitment. By contrast, our robust notion of strategy-proofness fails for any opacity level.

6 Conclusion

We introduce a framework of opacity in social choice and market design and show that strategy-proofness cannot be robustly guaranteed whenever there is any level of opacity. Hence, in our framework, full transparency is necessary to ensure the optimality of participants’ truth-telling. We also show that some well-known monotonicity and invariance properties are compatible with opacity. We have established this compatibility in specific environments but not generally. As a potential future research question, it will be interesting to understand the strongest incentive-, monotonicity-, or invariance properties that can be guaranteed under opacity in general environments. For now, we leave this question for future research.

References

- Ahn, D. S., & Root, J. (2023). Incentives and efficiency in constrained allocation mechanisms. *arXiv preprint arXiv:2006.06776*.
- Akbarpour, M., & Li, S. (2020). Credible auctions: A trilemma. *Econometrica*, 88(2), 425–467.
- Arrow, K. (1951). Social choice with individual values. *Monograph*.
- Aygün, O., & Turhan, B. (2022). How to de-reserve reserves: Admissions to technical colleges in India. *Management Science*.
- Baliga, S., Corchon, L. C., & Sjöström, T. (1997). The theory of implementation when the planner is a player. *Journal of Economic Theory*, 77(1), 15–33.
- Bester, H., & Strausz, R. (2000). Imperfect commitment and the revelation principle: the multi-agent case. *Economics Letters*, 69(2), 165–171.

- Bester, H., & Strausz, R. (2001). Contracting with imperfect commitment and the revelation principle: the single agent case. *Econometrica*, 69(4), 1077–1098.
- Bobbio, F., Carvalho, M., Lodi, A., Rios, I., & Torrico, A. (2023). *Capacity planning in stable matching*. (EC’23: Proceedings of the 24th ACM Conference on Economics and Computation)
- Budish, E., Cachon, G. P., Kessler, J. B., & Othman, A. (2017). Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Operations Research*, 65(2), 314–336.
- Doval, L., & Skreta, V. (2022). Mechanism design with limited commitment. *Econometrica*, 90(4), 1463–1500.
- Dur, U., Kominers, S. D., Pathak, P. A., & Sönmez, T. (2018). Reserve Design: Unintended Consequences and the Demise of Boston’s Walk Zones. *Journal of Political Economy*, 6(6), 2457–2479.
- Dur, U., & Van der Linden, M. (2021). Capacity design in school choice. *Available at SSRN 3898719*.
- Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1), 9–15.
- Gonczarowski, Y. A., Kovalio, L., Nisan, N., & Romm, A. (2019). Matching for the israeli “mechinot” gap-year programs: Handling rich diversity requirements. *arXiv preprint arXiv:1905.00364*.
- Grigoryan, A., & Möller, M. (2023). *A theory of auditability for allocation and social choice mechanisms*. (arXiv preprint arXiv:2305.09314)
- Grigoryan, A., & Möller, M. (2024). Auditability in school choice. *AEA Papers and Proceedings*, 114, 492–496.
- Hakimov, R., & Raghavan, M. (2023). *Improving transparency and verifiability in school admissions: Theory and experiment* (Tech. Rep.). Discussion Paper.
- Hatfield, J. W., & Milgrom, P. (2005). Matching with Contracts. *American Economic Review*, 95, 913–935.
- Haupt, A., & Hitzig, Z. (2023). Opaque contracts. *arXiv preprint arXiv:2301.13404*.
- Kapor, A. (2024). Transparency and percent plans. *NBER Working Paper*, 32372. doi: 10.3386/w32372
- Kojima, F., & Manea, M. (2010). Axioms for deferred acceptance. *Econometrica*, 78(2), 633–653.

- Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11), 3257–87.
- Möller, M. (2022). Transparent matching mechanisms. *Available at SSRN 4073464*.
- NRMP. (2023). *The Match: Annual Report 2022-2023*. (<https://annualreport.nrmp.org/>, last accessed in July 2024)
- NYC-Public-Schools. (2024). *2024 NYC Public Schools, Admissions Guide*. (<https://www.schools.nyc.gov/docs/default-source/default-document-library/enrollment/2024-nyc-public-schools-admissions-guide>, last accessed in July 2024)
- Pathak, P. A., Rees-Jones, A., & Sönmez, T. (2022). Immigration lottery design: Engineered and coincidental consequences of h-1b reforms. *Review of Economics and Statistics*, 1–43.
- Pycia, M., & Troyan, P. (2023). A theory of simplicity in games and mechanism design. *Econometrica*, 91(4), 1495–1526.
- Satterthwaite, M. A., & Sonnenschein, H. (1981). Strategy-proof allocation mechanisms at differentiable points. *The Review of Economic Studies*, 48(4), 587–597.
- Shapley, L., & Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, 1(1), 23–37.
- Stahl, K., & Strausz, R. (2017). Certification and market transparency. *The Review of Economic Studies*, 84(4), 1842–1868.
- Svensson, L.-G. (1999). Strategy-Proof Allocation of Indivisible Goods. *Social Choice and Welfare*, 16.
- Sönmez, T., & Yenmez, M. B. (2019). *Affirmative action in India via vertical and horizontal reservations*. (No. 977. Boston College Department of Economics)