

Optimal Allocation with Peer Information*

Axel Niemeyer[†]

Justus Preusser[‡]

Abstract

We study allocation problems without monetary transfers where agents have correlated types, i.e., hold private information about one another. Such peer information is relevant in various settings, including science funding, allocation of targeted aid, or intra-firm allocation. Incentive compatibility requires that agents cannot improve their own allocation by misrepresenting the merits of allocating to others. We characterize optimal incentive-compatible mechanisms using techniques from the theory of perfect graphs. Optimal mechanisms improve on review panels commonly observed in practice by eliciting information directly from eligible agents and by using allocation lotteries to alleviate incentive constraints. Computational hardness results imply that exactly optimal mechanisms are impractically complex. We propose ranking-based mechanisms as a viable alternative and show that they are approximately optimal when agents are informationally small, i.e., when no single agent has information that is crucial for evaluating a large fraction of the other agents.

Keywords: Mechanism design without transfers, correlated types, peer information, extreme points, informational size, perfect graphs, peer selection

JEL codes: D82, D71

*Parts of this paper previously circulated under the title “Simple Allocation with Correlated Types.” We thank several anonymous referees, Marco Battaglini, Florian Brandl, Modibo Camara, Gregorio Curello, Francesc Dilmé, Tangren Feng, Nathan Hancart, Deniz Kattwinkel, Jan Knöpfle, Nenad Kos, Daniel Krähmer, Stephan Lauermaun, Irene Lo, Andrew Mackenzie, Benny Moldovanu, Marco Ottaviani, Alessandro Pavan, Johannes Schneider, Ludvig Sinander, Francesco Squintani, Omer Tamuz, Éva Tardos, Rakesh Vohra, Kai Hao Yang, and numerous seminar audiences for helpful comments and discussions. Niemeyer acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2126/1 – 390838866. Preusser acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B04). Preusser acknowledges financial support by the European Research Council (HEUROPE 2022 ADG, GA No. 101055295 – InfoEcoScience). Preusser is grateful for the hospitality of Yale University.

[†]Division of the Humanities and Social Sciences, California Institute of Technology, niemeyer@caltech.edu.

[‡]Department of Economics and IGIER, Bocconi University, justus.preusser@unibocconi.it.

1 Introduction

This paper studies how peer information can be used to improve efficiency in allocation problems without monetary transfers. By peer information, we mean information that agents hold about the merits of allocating to other agents. Important examples include peer review in science funding, community-based targeting of social aid, or resource allocation within organizations based on information dispersed across several divisions.¹ In these settings, monetary transfers play a limited role. For example, when allocating funds such as a research grant, the allocation itself is a monetary transfer; additional transfers would simply amount to a different allocation.

Peer information is not directly available to the social planner. Like any other piece of private information, agents must be incentivized to truthfully report what they know about their peers. For example, when asked for a peer assessment, an agent may be tempted to discredit their peers if doing so improves their own allocation. Dishonest peer assessments are documented in several empirical studies, for example, when allocating loans to entrepreneurs (Hussam et al. (2022)), identifying workers for a promotion (Huang et al. (2019)), or in lab experiments where payoffs are determined based on how subjects evaluate each other’s performance on a task (Carpenter et al. (2010) and Baliatti et al. (2016)).

In our model, a principal allocates a good among several agents. Each agent wants the good, but the principal wants to allocate the good to the agent who needs it most or who can put it to best use. The principal’s values from assigning the good to different agents are initially unobserved by all parties, but each agent has a signal—their type—that is arbitrarily correlated with the entire profile of the principal’s values. An agent’s type captures all private information that the agent has about themselves and their peers. For example, when allocating social aid in a local community, a type may summarize an agent’s noisy assessment of their own and their neighbors’ values in an underlying social network. The fact that agents’ types are correlated is the main

¹For field studies on targeting with community information, see e.g. Alatas et al. (2012), Alatas et al. (2019) or Hussam et al. (2022). The oil and gas company BP has been reported to use a peer review mechanism for capital expenditure authorizations across business units (Goold (2005)). Other examples include group members selecting a leader based on nominations (e.g. Alon et al. (2011) and Holzman and Moulin (2013)), job promotions being decided based on feedback from colleagues who are also vying for a promotion (Huang et al. (2019)), and hiring in the gig economy relying on peer assessments of other applicants (Kotturi et al. (2020)). See the survey by Olckers and Walsh (2024) for further examples.

analytical challenge in our model.

To ensure incentives for honest reports, we focus on dominant-strategy incentive-compatible (DIC) mechanisms; that is, each agent finds it optimal to report their type truthfully, regardless of what the agent believes about the other agents' types and reports. In our setting, this robust notion of incentives means that an agent's report can never increase their *own* chances of getting the good. However, an agent's report can help determine the optimal allocation among the *other* agents.

Although agents' types are plausibly correlated in many allocation problems without transfers, the existing literature has not studied optimal mechanisms for such settings and does not provide tools for finding them. On the one hand, a small but growing literature replaces transfers with non-monetary screening devices but maintains the ubiquitous assumption that types are statistically independent. This assumption is necessary for the applicability of standard tools but rules out peer information because an agent's own type is the only source of information about the value of allocating to that agent. On the other hand, the existing literature on correlated types has focused on how transfers can be used to screen agents via their beliefs about the types of others (Cr  mer and McLean (1985, 1988)); these constructions require precise knowledge of agents' beliefs and are moot in settings without transfers. In contrast, we do not exploit beliefs as a screening device and do not rely on transfers to leverage correlation. Instead, we focus on the purely informational channel of correlation—how one agent's information can help determine the optimal allocation among the other agents.²

We derive two main insights. First, we describe optimal DIC mechanisms. We find that these mechanisms are too complex to be of practical use, much in contrast to the many elegant insights of mechanism design theory in settings with independent types. This complexity suggests shifting focus to simpler, approximately optimal mechanisms. Second, we show that a simple class of DIC mechanisms is indeed approximately optimal if agents are informationally small, i.e., if no single agent has information that is crucial for evaluating a large fraction of the other agents. In the following, we explain these insights in more detail.

The principal faces a fundamental trade-off between allocation and information elicitation. By allocating to an agent at a given type profile, the principal commits to allocating to the agent at all type profiles that differ from the given profile only

²Kattwinkel et al. (2022) is the only other paper that studies correlated types in an allocation problem without transfers. We discuss the differences to this and other papers in [Section 6](#).

in the agent’s report (otherwise the agent could influence their own allocation by misreporting). Thus, the principal can either benefit from allocating to the agent at these type profiles or from using the variation in the agent’s information to decide the allocation among other agents, but not from both. Simple mechanisms that partition agents into those whose reports inform the allocation and those eligible to receive resources—like review panels commonly observed in practice—generally fail to resolve this trade-off optimally. Instead, in optimal mechanisms an agent may provide information at some type profiles and receive the good at others.

As a first step, we provide a graph-theoretic characterization of DIC mechanisms. We define an auxiliary graph—the *feasibility graph*—which encodes the trade-off between allocation and information elicitation. DIC mechanisms are represented by certain combinatorial objects, stable sets and fractional stable sets, in the feasibility graph (Lemma 1).

We use this graph-theoretic characterization to show that the set of optimal mechanisms eludes simple economic descriptions.

To make this point precise, we begin with the subset of deterministic DIC mechanisms, which are conceptually simpler and easier to implement in practice. Moreover, every deterministic DIC mechanism is an extreme point of the set of DIC mechanisms and hence an indispensable candidate for optimality. The problem of finding an optimal deterministic DIC mechanism—a stable set of maximum weight in the feasibility graph—is NP-hard (Theorem 1).³ Thus, if one could give a simple economic description of deterministic DIC mechanisms, then this description would suggest a computationally efficient procedure for finding an optimal deterministic DIC mechanism, which is impossible under the widely believed $P \neq NP$ conjecture.⁴

We then extend the analysis to stochastic mechanisms and show that the principal can benefit from randomizing the allocation. Intuitively, randomization allows the principal to resolve the trade-off between allocation and information elicitation more flexibly because an agent may win with some probability but still influence how the remaining probability is distributed among the other agents. Formally, the set of DIC mechanisms admits stochastic extreme points if and only if there are at least three

³Related hardness results are known for certain revenue-maximization problems with correlated values (Papadimitriou and Pierrakos (2011, 2015); Papadimitriou et al. (2016, 2022)).

⁴Papadimitriou et al. (2016, see their introduction) analogously refer to computational complexity to argue that there is no simple characterization of optimal ex-post IR auctions with correlated values.

agents and at least one agent has at least three possible types.⁵ Moreover, when each agent has many possible types, essentially all extreme points are stochastic: the set of deterministic DIC mechanisms, all of which are extreme, constitutes a vanishing fraction of the set of extreme points (Theorem 2). Stochastic extreme points inherit the complexity of deterministic DIC mechanisms on those parts of the type space where they allocate deterministically. On the remaining parts, the benefit of randomization can be traced to how stochastic mechanisms distribute allocation probability around certain cycles—odd holes—in the feasibility graph (Lemma 2). The proofs leverage powerful results from the theory of perfect graphs (Chvátal (1975), Chudnovsky et al. (2006)).

To address the complexity of optimal mechanisms, we identify simple mechanisms that we call *ranking-based mechanisms* and show that they are approximately optimal in natural environments. To describe these mechanisms, we define an agent’s *peer value* as the principal’s expected value of allocating to that agent based only on the reports of the other agents. A ranking-based mechanism ranks agents according to their peer values and randomly selects an agent above a given rank threshold. Before allocating the good, the mechanism checks for a conflict of interest: does the agent pass the rank threshold robustly, regardless of what the agent reports about their peers? To ensure DIC, the principal keeps the good if a conflict of interest is detected.

Ranking-based mechanisms are approximately optimal when agents are informationally small (Theorem 3). For a given rank threshold, the mechanism performs badly only at type profiles where relatively few agents are robustly above the threshold. At such a profile, the principal inefficiently withholds the good with high probability. The performance of the mechanism thus hinges on the impact that any individual agent’s report has on the rank of their own peer value. We refer to this impact as the environment’s *informational size*. For vanishing informational size, the principal inefficiently withholds the good with vanishing probability. By suitably increasing the rank threshold as informational size vanishes, ranking-based mechanisms approximate the performance of optimal DIC mechanisms.

Informational size is naturally small in many applications of interest. For example, suppose there is an underlying social network and that each agent has information only about their neighbors in this network. In the applications given earlier, this could be the researchers working on nearby subfields, the literal neighbors in a community, or the

⁵In the model, the number of possible types per agent is finite and at least two.

firm’s divisions situated at the same geographic location. In this setup, informational size is small if the network’s neighborhoods are small relative to the overall size of the network; that is, if subfields are highly specialized, one’s neighbors constitute a small fraction of the overall community, and the divisions of the firm are spread across many locations. As another example, even if each agent holds information about all other agents and can therefore influence everyone else’s peer value, informational size is still small if no agent has information that is crucial for evaluating a large fraction of the other agents. In these applications, ranking-based mechanisms thus demonstrate how to provide incentives for honest evaluations while sacrificing little economic efficiency. This insight extends to multi-unit allocation problems.

Next, we describe the model. In the first part of the analysis, we introduce the feasibility graph (Section 3) and describe optimal DIC mechanisms and discuss their complexity (Section 4). In the second part (Section 5), we turn to approximate optimality and propose ranking-based mechanisms; readers interested primarily in these mechanisms may skip the first part without loss of continuity. In Section 6, we discuss the related literature. Section 7 makes concluding remarks. All proofs are in Appendix A. The Supplementary Appendix presents ancillary results.

2 Model

Environment. A principal allocates a good among $n \geq 2$ agents.⁶ Let $[n] = \{1, \dots, n\}$ denote the set of agents. Each agent i enjoys a payoff of 1 if allocated the good and a payoff of 0 otherwise.⁷ The principal’s payoff from allocating to agent i is given by $u_i \in [-1, 1]$. The principal’s payoff from keeping the good is normalized to 0.

The payoff profile $(u_1, \dots, u_n)$ is initially unobserved by all parties. However, each agent i has a private type θ_i from a finite type space Θ_i that is informative about the payoff profile (and the types of the other agents). Each agent i has at least two possible types, $|\Theta_i| \geq 2$. The joint distribution of payoffs and types is given by a Borel probability measure μ on $\Theta \times [-1, 1]^n$, where Θ denotes the set of type profiles with typical element $\theta = (\theta_1, \dots, \theta_n)$. For our analysis, the distribution μ need not be a common prior; it may simply be interpreted as the principal’s belief about the

⁶Our main insights extend to the problem of allocating multiple units; see Section 7.

⁷Nothing in our analysis would change if an agent’s payoff from getting the good would depend on the agents’ private information, provided that this payoff never changes sign. Indeed, the good could be a bad (e.g. an undesirable task), in which case the payoff is negative.

structure of peer information.

Notation. The set of type profiles of agents other than i is denoted Θ_{-i} with typical element θ_{-i} . For a profile $\theta \in \Theta$, we write $\mu(\theta_{-i})$ and $\mu(\theta)$ to mean the marginal probabilities of θ_{-i} and θ , respectively. We also write $\mathbb{E}[u_i|\theta]$ for the conditional expectation of u_i given θ (provided $\mu(\theta) > 0$, else we arbitrarily define $\mathbb{E}[u_i|\theta] = 0$).

Mechanisms. A (*direct*) *mechanism* is a function $q: \Theta \rightarrow [0, 1]^n$ such that $\sum_{i \in [n]} q_i(\theta) \leq 1$ for all $\theta \in \Theta$. Here, q_i denotes agent i 's probability of winning the good.⁸ Whenever the sum of winning probabilities is less than 1 at some type profile, the principal keeps the good with the remaining probability.

A mechanism is *dominant-strategy incentive-compatible (DIC)* if $q_i(\theta_i, \theta_{-i}) \geq q_i(\theta'_i, \theta_{-i})$ for all $i \in [n]$, $\theta_i, \theta'_i \in \Theta_i$ and $\theta_{-i} \in \Theta_{-i}$. By swapping the roles of θ_i and θ'_i in this inequality, we see that a mechanism q is DIC if and only if each agent i 's allocation is constant in their own report; that is, $q_i(\theta_i, \theta_{-i}) = q_i(\theta'_i, \theta_{-i})$. We henceforth drop i 's type from i 's winning probability, i.e., $q_i(\theta_{-i}) = q_i(\theta)$, since we focus on DIC mechanisms throughout the paper. The set of DIC mechanisms is denoted by $\mathcal{Q}$. The Revelation Principle applies.

The principal's utility from a DIC mechanism q is denoted $U(q)$ and is given by

$$U(q) = \sum_{\theta \in \Theta} \sum_{i \in [n]} \mu(\theta) q_i(\theta_{-i}) \mathbb{E}[u_i|\theta].$$

A DIC mechanism is *optimal* if it maximizes the principal's utility across all DIC mechanisms. A mechanism q is *deterministic* if it maps to $\{0, 1\}^n$. A mechanism is *stochastic* if it is not deterministic.

Examples of peer information. Our model nests various forms of peer information.

Correlated information in mechanism design is often modeled by assuming that $u_i = \theta_i$. That is, each agent privately knows their value. Information about the other agents' values is captured entirely by correlation between the values of different agents.

In the peer selection literature (e.g. Alon et al. (2011) and Holzman and Moulin (2013)), the principal selects an agent based on nominations from their peers. We can nest the peer selection problem by assuming that each agent i 's type θ_i equals the identity j of another agent. To model a principal who aims to select an agent

⁸One can alternatively interpret the good as a divisible resource. Agent i 's share of the resource is denoted q_i , and the principal's payoff from allocating to agent i is linear in q_i .

with many nominations, we could let the payoff u_j from allocating to an agent j be given by $u_j = \sum_{i \in [n]: i \neq j} \mathbf{1}(\theta_i = j)$. Similarly, we could capture agents' submitting multiple nominations, scores, rankings, etc. See [Section 6](#) for a detailed discussion of the relation of our paper to the peer selection literature.

Finally, suppose there is an underlying social network. For each i , let $N(i)$ denote agent i 's neighbors in the network. Agent i privately observes u_i and a noisy signal $u_j + \varepsilon_{ij}$ for each of the neighbors $j \in N(i)$, where ε_{ij} is some random variable. Thus, agent i 's type θ_i is $(u_i, (u_j + \varepsilon_{ij})_{j \in N(i)})$.

Peer values. For all $i \in [n]$ and $\theta_{-i} \in \Theta_{-i}$, we define the *peer value of agent i at θ_{-i}* as $\bar{u}_i(\theta_{-i}) = \mathbb{E}[u_i | \theta_{-i}]$. The peer value captures the collective prediction of others about the value u_i of agent i . (If $\mu(\theta_{-i}) = 0$, we arbitrarily put $\bar{u}_i(\theta_{-i}) = 0$.)

The principal's utility from a DIC mechanism can be written in terms of the peer values. For a DIC mechanism q , straightforward manipulations show

$$U(q) = \sum_{\theta \in \Theta} \sum_{i \in [n]} \mu(\theta) q_i(\theta_{-i}) \bar{u}_i(\theta_{-i}). \quad (1)$$

Thus, the principal can only elicit the agents' peer information. Indeed, if agents have no peer information in the sense that for all i the peer value $\bar{u}_i(\theta_{-i})$ is constant in θ_{-i} , then (1) implies that a constant mechanism that allocates to an agent i with the highest unconditional expected value $\mathbb{E}[u_i]$ is optimal.

However, as soon as agents hold some peer information, the principal may benefit from designing non-constant DIC mechanisms. For example, a *jury mechanism* partitions agents into jurors and candidates. Jurors provide information but cannot receive the good, whereas candidates may receive the good but do not provide information. These mechanisms, which resemble review panels commonly seen in practice, ensure DIC in the most straightforward way: those who influence the allocation never receive the good. It is easy to show that jury mechanisms are optimal when there are just two agents ($n = 2$), i.e., one agent decides whether the other agent should get the good.⁹ However, our main results imply that jury mechanisms are not generally optimal.

⁹Jury mechanisms are also optimal if there are three agents ($n = 3$) and the principal must allocate the good; see the Supplementary Appendix. This extends a result of Holzman and Moulin (2013).

3 A Graph-Theoretic Characterization of DIC Mechanisms

In this section, we characterize DIC mechanisms via an auxiliary graph, which we then use to study the structure of optimal DIC mechanisms in the next section. As explained below, the graph captures a fundamental trade-off between allocation and information elicitation. Readers who are primarily interested in the proposed approximately optimal mechanisms may skip this section and the next without loss of continuity.

Recall that when allocating to an agent i at some type profile, DIC requires that the principal commits to also allocate to i at all type profiles that differ only in i 's type. Using the notation $q_i(\theta_{-i})$, where we omit i 's type from i 's allocation in a DIC mechanism q , the feasibility constraints read

$$\forall \theta \in \Theta, \quad \sum_{i=1}^n q_i(\theta_{-i}) \leq 1. \quad (\text{Feasibility})$$

The feasibility graph tracks commitments to allocate to different agents that are mutually incompatible with the feasibility constraints. For example, the principal cannot simultaneously allocate to agent i at type profile $(\theta_i, \theta_j, \theta_{-ij})$ and to another agent j at type profile $(\theta'_i, \theta'_j, \theta_{-ij})$ since the principal would have to commit to allocate to both i and j at $(\theta'_i, \theta_j, \theta_{-ij})$ to satisfy DIC. In other words, the commitments to allocate to agent i when others report θ_{-i} and to agent j when others report θ'_{-j} are incompatible if $q_i(\theta_{-i})$ and $q_j(\theta'_{-j})$ simultaneously appear in the same (Feasibility) constraint for some type profile. Define $V = \cup_{i=1}^n (\{i\} \times \Theta_{-i})$ and think of each $(i, \theta_{-i}) \in V$ as indexing the commitment to agent i at θ_{-i} .

Definition 1 (Feasibility graph). The *feasibility graph* G is the graph whose vertex set is $V = \cup_{i=1}^n (\{i\} \times \Theta_{-i})$ and such that two vertices (i, θ_{-i}) and (j, θ'_{-j}) are adjacent if and only if $i \neq j$ and there exists a type profile $\hat{\theta}$ such that $\hat{\theta}_{-i} = \theta_{-i}$ and $\hat{\theta}_{-j} = \theta'_{-j}$.

In the Supplementary Appendix, we depict the feasibility graph in an example.

The feasibility graph captures a trade-off between allocating to an agent and using the agent's information. First, suppose the principal selects a vertex (i, θ_{-i}) (i.e. commits to allocate to agent i when others report θ_{-i}). Then, the principal cannot select any vertex $(j, (\theta'_i, \theta_{-ij}))$ for any $j \neq i$ or θ'_i since these vertices are all adjacent to

(i, θ_{-i}) . Second, suppose the principal does not select (i, θ_{-i}) . Then, the principal can decide for each θ'_i which vertex $(j, (\theta'_i, \theta_{-ij}))$ to select (subject to the selected vertices' being non-adjacent to other selected vertices). In the first case, the principal benefits from allocating to i ; in the second case, from i 's information.

In graph-theoretic terms, deterministic DIC mechanisms correspond to stable (or independent) sets of the feasibility graph G , and (possibly stochastic) DIC mechanisms correspond to fractional stable sets of the feasibility graph G . A *stable set* of G is a subset of pairwise non-adjacent vertices. A stable set S of G can be mapped to a deterministic DIC mechanism q by setting $q_i(\theta_{-i}) = \mathbf{1}((i, \theta_{-i}) \in S)$ for all $(i, \theta_{-i}) \in V$. Conversely, if q is deterministic and DIC, then $\{v \in V : q(v) = 1\}$ is a stable set (where $q(v) = q_i(\theta_{-i})$ for $v = (i, \theta_{-i})$). A *clique* of G is a subset of pairwise adjacent vertices, and a clique is *maximal* if it is not contained in another clique. A *fractional stable set* of G is a function $q : V \rightarrow [0, 1]$ such that $\sum_{v \in X} q(v) \leq 1$ for all maximal cliques X in G ,¹⁰ and the *fractional stable set polytope* of G is the set of all fractional stable sets of G . Every type profile corresponds to a maximal clique since the variables appearing in one (Feasibility) constraint are all adjacent and since every two (Feasibility) constraints have at most one variable in common. We summarize:

Lemma 1. *There is a bijection between deterministic DIC mechanisms and stable sets of the feasibility graph. There is a bijection between DIC mechanisms and fractional stable sets of the feasibility graph.*

Identifying an optimal DIC mechanism is tantamount to solving a combinatorial optimization problem. The peer values and the type distribution imply weights for the vertices of the feasibility graph G ; the weight $w_i(\theta_{-i})$ on vertex (i, θ_{-i}) is $\mu(\theta_{-i})\bar{u}_i(\theta_{-i})$. The problem of finding an optimal deterministic DIC mechanism is the same as finding a stable set in G with the highest cumulative weight; this is an instance of the *maximum weight stable set (MWSS) problem*. The problem of finding an optimal (possibly stochastic) DIC mechanism is the *fractional relaxation* of this problem. Our goal is hence to understand these problems for the special case of the feasibility graph.

¹⁰The combinatorial optimization literature distinguishes between edge-constrained and clique-constrained fractional stable sets. Only the latter correspond to stochastic DIC mechanisms, and hence we focus on these. See Sections 64.5 and 64.7 in Schrijver (2003), respectively, for details.

4 Optimal DIC Mechanisms

In this section, we show that optimal DIC mechanisms have a complex structure despite resolving an intuitive economic trade-off between allocation and information elicitation.

In [Section 4.1](#), we first study deterministic mechanisms, which are conceptually simpler and easier to implement in practice. One might expect that restricting attention to deterministic mechanisms would yield a simple characterization. After all, the constraints in the principal’s problem are easy to describe. Moreover, it may be difficult to imagine why optimal mechanisms would need to depart significantly from simple schemes such as jury mechanisms. However, using computational complexity arguments, we show that there must exist many other deterministic DIC mechanisms that elude straightforward economic descriptions ([Theorem 1](#)), and these mechanisms are indispensable for optimality.

In [Section 4.2](#), we extend the analysis to stochastic mechanisms and show that the principal can benefit from randomizing the allocation. Formally, the set of DIC mechanisms generally admits stochastic extreme points, and these are prevalent among all extreme points when type spaces are rich ([Theorem 2](#)). Intuitively, randomization allows the principal to resolve the trade-off between allocation and information elicitation more flexibly: an agent may win with some probability but still influence how the remaining probability is distributed among the other agents. Importantly, we also explain why stochastic mechanisms inherit the complexity of deterministic mechanisms.

4.1 Deterministic Mechanisms

For three or more agents, the problem of finding an optimal deterministic DIC mechanism is NP-hard. This problem is defined as follows.

Definition 2 (OPTDET- n). The input consists of finite sets $\Theta_1, \dots, \Theta_n$ and weights $w_i : \Theta_{-i} \rightarrow \mathbb{Z}$ for all $i \in [n]$. The problem is to find a deterministic DIC mechanism q (for n agents with type spaces $\Theta_1, \dots, \Theta_n$) that maximizes $\sum_{i,\theta} w_i(\theta_{-i}) q_i(\theta_{-i})$ across all deterministic DIC mechanisms.¹¹

¹¹In a given environment, the weights of interest are given by $w_i(\theta_{-i}) = \mu(\theta_{-i}) \bar{u}_i(\theta_{-i})$ for all (i, θ_{-i}) . Conversely, all weight vectors can result from some environment (up to rescaling). Computation requires that weights are given as integers or rational numbers. By fixing n , we ensure that complexity

Theorem 1. *For fixed $n \geq 3$, the problem OPTDET- n is NP-hard.*

Theorem 1 implies that there is no conceptually simple description of deterministic DIC mechanisms, at least not to the extent that such a description could be leveraged to design an efficient algorithm to find an optimal mechanism (unless $P = NP$). This complexity is in contrast to many allocation problems with independent types, where optimal mechanisms often admit simple, parametric descriptions that naturally suggest efficient computational solutions.¹² Although there is no formal notion of conceptual complexity in the mechanism design literature, computational complexity is arguably the best available proxy.¹³

The proof gives further insight into the origin of this complexity. Intuitively, the allocation at one type profile constrains in a non-local, combinatorial way the allocation at other type profiles.¹⁴ Thus, it is not possible to construct an optimal mechanism based on an “index” computed from local information at each type profile, akin to Myersonian virtual values. This non-locality results from the combinatorics of the stable sets of the feasibility graph (Lemma 1). We show that the edges of any arbitrary graph can be simulated via induced paths in the feasibility graph in a way that preserves the structure and cardinality of stable sets for appropriately chosen weights. That is, we prove Theorem 1 by reducing the problem of finding a maximum cardinality stable set for an arbitrary graph (which is NP-hard) to the problem of finding a maximum weight stable set in the feasibility graph. In this

does not arise merely from increasing the number of agents.

¹²For example, consider the allocation problem without transfers in Ben-Porath et al. (2014), where the principal can verify agents’ types at a cost. Ben-Porath et al. (2014) show that optimal mechanisms are characterized by two parameters: a type cutoff and a “favored” agent. This characterization immediately yields a computationally efficient solution by exhaustive search over the two parameters. Vohra (2012) shows how a number of problems with and without transfers correspond to polymatroid optimization problems with respect to reduced-form allocations, which can be solved efficiently via a greedy algorithm. See e.g. Cai et al. (2012) and Alaei et al. (2019) for details regarding computation.

¹³Conversely, Hart and Nisan (2017, Section 1.1) argue that even computationally tractable problems may still be complex on a “conceptual” level “since, even after computing the precise solution, one may not understand its structure, what it means and represents, and how it varies with the given parameters.”

¹⁴For instance, by allocating to an agent i at a type profile $\theta = (\theta_i, \theta_j, \theta_{-ij})$, the principal can no longer allocate to any agent j at type profile $(\theta'_i, \theta'_j, \theta_{-ij})$ for all θ'_i and θ'_j . Hence, even if at θ agent i has the highest peer value $\bar{u}_i(\theta_{-i}) = \mathbb{E}[u_i|\theta_{-i}]$, we cannot conclude that the optimal mechanism allocates to i because doing so constrains the allocation that can be made at other type profiles. The allocations at these type profiles, in turn, constrain the allocations at various other type profiles, and so on. Moreover, if i has relevant peer information—i.e., if the peer values $\bar{u}_j(\theta'_i, \theta_{-ij}) = \mathbb{E}[u_j|\theta'_i, \theta_{-ij}]$ differ significantly across j and θ'_i —it is not obvious how the principal should capitalize on this information precisely because it constrains the decision at further type profiles.

way, the combinatorics of finding a maximum stable set in an arbitrary graph are replicated by the combinatorics of deciding at which type profiles which agent gets the good subject to DIC to maximize the principal’s utility. Since maximum stable sets of arbitrary graphs have no discernible structure, this proof makes precise why the economic constraints of our problem—feasibility and DIC—impose no discernible structure on optimal deterministic DIC mechanisms either, at least not without any further restrictions on the problem.

4.2 Stochastic Mechanisms

We now turn to stochastic DIC mechanisms. We focus here on stochastic extreme points of the set Q of DIC mechanisms, as these are the relevant candidates for optimal mechanisms given the principal’s linear objective.¹⁵ In the language of [Lemma 1](#), we hence study extreme points of the fractional stable set polytope of the feasibility graph G . Little is known about these extreme points for general graphs, so our results will have to leverage the special structure of the feasibility graph.¹⁶

Our main result about stochastic mechanisms is that stochastic extreme points exist outside of special cases and that they are prevalent when type spaces are rich. Let $\det Q$ denote the set of deterministic DIC mechanisms, and let $\text{ext } Q$ denote the set of extreme points of the set of DIC mechanisms Q . All deterministic DIC mechanisms are extreme points.

Theorem 2. *The set of DIC mechanisms admits stochastic extreme points if and only if $n \geq 3$ and $|\Theta_i| \geq 3$ for some agent $i \in [n]$. Moreover, for all $\varepsilon > 0$ there exists $m \in \mathbb{N}$ such that if $|\Theta_i| \geq m$ for all $i \in [n]$ and $n \geq 3$, then $|\det Q| < \varepsilon |\text{ext } Q|$.*

Next, we further characterize the structure of stochastic extreme points. These structural insights are also the basis for proving [Theorem 2](#) and for elucidating why randomization cannot simplify the description of optimal mechanisms.

We begin by observing that a stochastic DIC mechanism q can be decomposed according to the type profiles where it allocates deterministically versus stochastically.

¹⁵An extreme point of Q is a DIC mechanism that cannot be represented as a convex combination of other DIC mechanisms. The extreme points are the natural objects of interest because every (optimal) DIC mechanism can be represented as a convex combination of (optimal) extreme points. Further, in our setting every extreme point is exposed: there exists a joint distribution of types and values that renders this extreme point uniquely optimal among all DIC mechanisms.

¹⁶See Koster and Wagler ([2006](#)) and the references therein.

A *stochastic component* of q is an inclusion-wise maximal connected set of vertices v of the feasibility graph G such that $q(v) \in (0, 1)$.¹⁷ The *deterministic component* of q is the set of vertices v such that $q(v) = 1$. The support of q is thus partitioned into a deterministic component and several stochastic components.¹⁸

From this decomposition, it can be seen that allowing randomization cannot simplify the description of candidate optimal mechanisms. First, as mentioned above, every deterministic DIC mechanism is an extreme point of the set of DIC mechanisms and hence an indispensable candidate for optimality. Second, stochastic mechanisms inherit the structure and complexity of deterministic mechanisms on their deterministic components. To see this, fix any stochastic DIC mechanism with a non-empty deterministic component. Consider the subgraph of the feasibility graph induced by the deterministic component and its neighborhood. As in [Lemma 1](#), the deterministic component can be identified with a stable set of this subgraph. These stable sets obey the same combinatorics as the stable sets of the entire feasibility graph, i.e. of deterministic DIC mechanisms. For instance, take the class of DIC mechanisms where the report of a dummy agent decides whether the principal uses a stochastic extreme point or a deterministic mechanism for the remaining agents. Describing this class of n -agent extreme-point mechanisms entails describing all deterministic $(n - 1)$ -agent DIC mechanisms.

A stochastic extreme point is characterized by certain cycles—odd holes—in its stochastic components. For $k \in \mathbb{N}$, a set $\{v_1, \dots, v_k\}$ of k vertices of G is an *odd* (*k*-)*hole* if $k \geq 5$ is odd and, for all $j \in \{1, \dots, k\}$, the vertex v_j is adjacent to v_{j-1} and v_{j+1} (where $v_1 = v_{k+1}$) and to no other vertex from $\{v_1, \dots, v_k\}$. Say that a stochastic mechanism is *simple* if it takes values in $\{0, 1/2, 1\}$.

Lemma 2. *Let q be a stochastic DIC mechanism. If $q \in \text{ext } Q$, then every stochastic component of q contains an odd hole. The converse holds if q is simple.*

[Lemma 2](#) shows why stochastic mechanisms can outperform deterministic mechanisms: they can distribute allocation mass more flexibly around odd holes in the feasibility graph. For an illustration, consider [Figure 1](#), which depicts an odd hole in the feasibility graph, together with the allocation probabilities of a deterministic and a simple stochastic DIC mechanism. Recall that adjacent vertices of the feasibility graph

¹⁷Two vertices are *connected* if there is a path joining them.

¹⁸This decomposition parallels a result of Nemhauser and Trotter ([1974](#), Theorem 2.4) for edge-constrained fractional stable sets (cf. [Footnote 10](#)).




Figure 1: An odd hole in the feasibility graph. The outer labels show how a stochastic mechanism might allocate around the hole; the inner labels show how a deterministic mechanism might allocate around the hole.

are commitments to allocate to two distinct agents that are incompatible with the ([Feasibility](#)) constraints at some type profile. Thus, the deterministic DIC mechanism can select at most three out of seven vertices from the hole. However, the stochastic mechanism can select all seven vertices, each with probability $1/2$, by splitting the good between two agents at type profiles where commitments are incompatible. In particular, the stochastic mechanism distributes more total allocation mass around the hole than any deterministic DIC mechanism and, thus, cannot be written as a convex combination of deterministic DIC mechanisms. In the Supplementary Appendix, we graphically depict such an odd hole alongside a stochastic DIC mechanism.

[Lemma 2](#) relies on a result by Chvátal ([1975](#)) and the Strong Perfect Graph Theorem (Chudnovsky et al. ([2006](#))). Recall that DIC mechanisms correspond to fractional stable sets of the feasibility graph, while deterministic DIC mechanisms correspond to stable sets. Chvátal ([1975](#)) shows that, for a general graph, every fractional stable set is a convex combination of (incidence vectors of) stable sets if and only if the given graph is perfect. The Strong Perfect Graph Theorem asserts that a graph G is perfect if and only if neither G nor its complement admit an odd hole. (We define perfection and the complement of a graph in [Appendix A.1](#).) [Lemma 2](#) follows by analyzing odd cycles and their complements in the feasibility graph and leveraging this graph’s special structure.¹⁹ See [Section 6](#) for how our approach based on perfect graphs relates to other concepts from combinatorial optimization that have been used in the mechanism and market design literature.

The existence part of [Theorem 2](#) is almost immediate from [Lemma 2](#), but the

¹⁹Specifically, the feasibility graph admits no odd 5-holes; its complement admits no odd holes.

prevalence part is more involved. Using [Lemma 2](#), we build simple stochastic extreme points by inserting odd holes into stable sets of the feasibility graph (i.e., into deterministic mechanisms). At the heart of the proof is then a counting argument showing that, as type spaces become richer, the number of possible ways to insert odd holes into stable sets grows faster than the number of possible stable sets. This shows that simple stochastic extreme points asymptotically outnumber deterministic mechanisms.

5 Approximate Optimality and Ranking-based Mechanisms

Our findings in [Section 4](#) show that optimal DIC mechanisms lack closed-form descriptions and may therefore be difficult to implement and institutionalize in practice. More broadly, this lack of structure casts doubt on the ability of the optimal mechanism design framework to yield practically viable predictions in environments with peer information, even in environments where constructions à la Crémer and McLean ([1985, 1988](#)) are ruled out by the absence of transfers.

In this section, we therefore shift our focus from optimal DIC mechanisms to approximately optimal ones. As a proof of concept, we propose a simple class of mechanisms that retain key properties of optimal mechanisms: they use all agents' information, consider all agents as potential recipients of the good, and randomize the allocation to alleviate incentive constraints. We call these mechanisms *ranking-based mechanisms*. Unlike optimal mechanisms, it is easy to see when and why ranking-based mechanisms perform well. Specifically, their performance depends on the extent to which any individual agent's report influences the peer values of others. This influence is naturally small when there are many agents and, in this case, ranking-based mechanisms are approximately optimal among all DIC mechanisms.

Ranking-based mechanisms aim to allocate to agents who are ranked highly in terms of their peer value $\bar{u}_i(\theta_{-i}) = \mathbb{E}[u_i|\theta_{-i}]$. Given a type profile θ , the agent with the highest peer value receives rank 1, the second highest receives rank 2, and so on. Ties are broken lexicographically, favoring agent i over agent j if $i < j$. It is convenient to normalize all ranks to be in $[0, 1]$ by dividing by the number n of agents. Let $r_i(\theta)$ denote i 's rank at θ .²⁰

²⁰Formally, $r_i(\theta) = n^{-1}|\{j \in [n]: \bar{u}_j(\theta_{-j}) > \bar{u}_i(\theta_{-i})\}| + n^{-1}|\{j \in [n]: (\bar{u}_j(\theta_{-j}) = \bar{u}_i(\theta_{-i})) \wedge (i \geq j)\}|$.

It is not incentive compatible to naively allocate the good to an agent with a high peer value: agents can misreport their type to diminish the peer values of other agents, thereby improving their own rank. Define

$$r_i^*(\theta_{-i}) = \max_{\theta_i \in \Theta_i} r_i(\theta_i, \theta_{-i})$$

to be i 's *robust rank at θ_{-i}* . Intuitively, the robust rank reflects a hypothetical scenario where agent i appraises their competitors, thereby diminishing their own rank.

A ranking-based mechanisms can be interpreted to work in two steps. First, the principal randomly selects an agent ranked, say, in the top 10% of peer values. Second, the principal allocates to the selected agent if and only if the agent's *robust rank* is also within the top 10%. Intuitively, the principal checks for a conflict of interest: does the agent pass the bar regardless of what the agent claims about their peers?

Definition 3 (Ranking-based mechanism). Let $p \in (0, 1)$. The *ranking-based mechanism q^p with threshold p* is defined for all $i \in [n]$ and $\theta_{-i} \in \Theta_{-i}$ by:

$$q_i^p(\theta_{-i}) = \begin{cases} \frac{1}{\lfloor pn \rfloor} & \text{if } r_i^*(\theta_{-i}) \leq p \text{ and } u_i(\theta_{-i}) \geq 0;^{21} \\ 0 & \text{else.} \end{cases}$$

Ranking-based mechanisms are DIC: $q_i^p(\theta_{-i})$ does not depend on agent i 's type report. Moreover, q^p is a feasible mechanism, never allocating more than is available: the number of agents with a robust rank better than p is at most the number of agents with a rank better than p which, in turn, is at most $\lfloor pn \rfloor$ by definition of the rank.

A crucial factor for the performance of ranking-based mechanisms is the impact that any individual agent's report has on their own rank. We refer to this impact as the informational size.²² Informational size is crucial since the mechanism may inefficiently withhold the good with high probability at type profiles where many agents have a unilateral misreport that raises their rank above the rank threshold.²³

²¹ $\lfloor x \rfloor$ denotes the largest integer smaller than x .

²²Similar notions of informational size appear in earlier work on mechanism design and general equilibrium with informational asymmetries, e.g. Gul and Postlewaite (1992), McLean and Postlewaite (2002), Gerardi et al. (2009), McLean and Postlewaite (2015), and Andreyanov and Sadzik (2021). A precursor to these notions is the notion of non-exclusive information (Postlewaite and Schmeidler (1986)). In those models, the agents have private signals about a hidden state. Informational size measures the impact of individual signals on the posterior about the state conditional on the full signal profile. For our setting, it is natural to define informational size via peer values and ranks.

²³To address this inefficiency, one could attempt to modify the mechanism. Suppose the principal

Definition 4. For $\theta \in \Theta$, the *informational size at θ* is denoted $\delta(\theta)$ and is given by

$$\delta(\theta) = \max_{i \in [n], \theta'_i \in \Theta_i} |r_i(\theta_i, \theta_{-i}) - r_i(\theta'_i, \theta_{-i})|.$$

Informational size is small if no agent has a large influence on the peer values of a large fraction of other agents. Small informational size is a natural assumption in environments with many agents, as illustrated by the following two examples.

Example 1. Informational size is small if each agent only influences the peer values of few others (but possibly has a large influence on those). Suppose there is an underlying social network. For each i , let $N(i)$ denote agent i 's neighbors in the network. Agent i 's type is informative only about the values and types of their neighbors $N(i)$. If the network is sparse in the sense that the degree $|N(i)|$ of every agent i is small relative to the size n of the network, then informational size is also small since each agent's report only influences the peer values of their neighbors. (Specifically, agent i can impact their own rank by at most $|N(i)|/n$.) Besides the examples given in the introduction, a sparse network may also arise from information acquisition constraints: for example, in peer review, each researcher cannot feasibly assess all other researchers.

Example 2. Informational size may also be small if each agent can influence the peer values of many others, but influences each given peer value only to a small degree. Suppose each agent i privately observes a noisy signal $u_j + \varepsilon_{ij}$ for *every* other agent j , where ε_{ij} is some random variable. If the variables ε_{ij} are conditionally i.i.d. across i for every j , then informational size is small when there are many agents, except at a vanishing number of type profiles where peer values are highly concentrated. Alternatively, assume each agent's type specifies for every other agent a numerical rating between, say, 1 and 5. Types could be correlated across agents. If the principal's values are simply the average peer ratings, then the informational size is small whenever there are many agents, except at a vanishing number of type profiles where peer values are highly concentrated.

The next result shows that ranking-based mechanisms become approximately

randomly samples a second agent j if the first sampled agent i is above the bar but not robustly so. If j 's report affects whether i is robustly above the bar (but not whether i is above the bar), then j would have a profitable misreport as they get a second chance of being selected. Thus, naive resampling need not be incentive-compatible and the principal may have to withhold the good from the resampled agent j to restore DIC. Thus, the efficiency gain from resampling is small in most cases.

optimal as informational size vanishes probabilistically. The fact that informational size is only required to vanish probabilistically allows for a small number of type profiles where informational size is large, such as type profiles where peer values are highly concentrated. The result uses a mild technical condition.

Definition 5 (Regularity). A sequence of environments $(n, \Theta^n, \mu^n)_{n \in \mathbb{N}}$ is *regular* if for all $\varepsilon > 0$ there exists $p \in (0, 1)$ such that

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{\theta \sim \mu^n} [\bar{u}_{[1]}^n(\theta) - \bar{u}_{[pn]}^n(\theta)] < \varepsilon,$$

where $\bar{u}_{[\ell]}^n(\theta)$ is the ℓ 'th highest peer value at profile θ .

Regularity is a continuity requirement on the distribution of peer values in the many-agent limit, requiring that the expectation of order statistics is continuous at the top. That is, the top peer-ranked agent is comparable to the runner-ups. Regularity is consistent with the narrative that it is difficult to distinguish the very best agents by peer evaluations alone (see e.g. the discussion of Fang and Casadevall (2016) in the context of science funding). Nevertheless, even if regularity is violated but informational size is small, ranking-based mechanism with a threshold p close to 0 still allocate to the top agents with overwhelming probability—they may just fail to pick up on the extraordinary value of the very best agent.

Theorem 3. Let $(n, \Theta^n, \mu^n)_{n \in \mathbb{N}}$ be a regular sequence of environments. Let the associated informational size $(\delta^n)_{n \in \mathbb{N}}$ converge to 0 in probability; that is, for all $d > 0$,

$$\lim_{n \rightarrow \infty} \mu^n \{\theta \in \Theta^n : \delta^n(\theta) > d\} = 0.$$

Then, the difference between the principal's expected utility from an optimal ranking-based mechanism and an optimal DIC mechanism vanishes as $n \rightarrow \infty$.

The proof goes as follows. Fixing a threshold $p \in (0, 1)$ and letting the informational size vanish in probability, we show that the ex-ante probability that the ranking-based mechanism q^p withholds the good vanishes. Conditional on allocating, q^p allocates to an agent with rank better than p . Under regularity, if p is close to 0, the peer value of such an agent is close to the highest peer value (in expectation). Recall from eq. (1) that the principal's expected utility in a DIC mechanism is determined by the peer values. In particular, no DIC mechanism does better than always allocating to

the highest peer value. This upper bound is approximately attained by any optimal ranking-based mechanism, which completes the proof.²⁴ For every fixed number n of agents, the proof also implies lower bounds on the performance of ranking-based mechanisms in terms of p and informational size.

A broader insight is that incentives matter little for eliciting peer information when agents are informationally small. For fixed informational size, it is typically impossible to identify and allocate to the agent with the highest peer value. The reason is that, by DIC, the principal cannot elicit an agent’s information about how the agent is ranked relative to others and then use this information to decide whether to allocate to the agent. Optimally resolving this tension requires complex DIC mechanisms (Section 4). But, if informational size is small, then an individual agent’s information is less important for how the agent is ranked relative to others, and hence there is little loss from ignoring this information when deciding the agent’s allocation. Thus, simple mechanisms such as ranking-based mechanisms can approximately attain the upper bound of always allocating to the agent with the highest peer value.²⁵

Ranking-based mechanisms are also approximately optimal for allocating *multiple* units of a good to agents with unit demand. Theorem 3 generalizes not only to sequences of environments along which the ratio of the number of units to the number of agents vanishes, but also to sequences along which this ratio is bounded away from zero (see the working paper version, Niemeyer and Preusser (2025)).

6 Related Literature

This paper contributes to the literature on correlation in mechanism design, on allocation without transfers, on peer selection, and on randomization in mechanism and market design.

In settings with transfers, Crémer and McLean (1985, 1988) and McAfee and Reny (1992) show how the principal can leverage belief heterogeneity across types to extract all private information without leaving information rents to the agents.²⁶

²⁴Note that this result does not imply that the principal approaches the first best, i.e., the utility from always allocating to an agent with the highest value given the *entire* type profile, $\max_{i \in [n]} \mathbb{E}[u_i | \theta]$.

²⁵In the Supplementary Appendix, we show that the even simpler jury mechanisms described in Section 2 need not be approximately optimal with many agent if informational size vanishes. However, we provide strong symmetry conditions under which jury mechanisms are approximately optimal.

²⁶See also Lopomo et al. (2022) for a recent treatment.

Their constructions are infeasible in our setting since there are no monetary transfers and since we demand robustness to beliefs.²⁷

Without transfers, Kattwinkel et al. (2022) characterize Bayesian IC (BIC) mechanisms with two agents and correlated types when the principal must allocate the good. By contrast, we focus on dominant-strategy IC (DIC) mechanisms with many agents. Since DIC implies BIC, our complexity results suggest that a simple characterization of BIC mechanisms may be difficult to find for a general number of agents. We are unaware of other papers on allocation problems without transfers where information is correlated across agents. Other papers consider correlation in the form of an external signal that the principal can use to cross-check an agent’s type report (Kattwinkel (2020); Bloch et al. (2023); Kattwinkel and Knoepfle (2023); Pereyra and Silva (2023)).

A recent literature on allocation problems without transfers studies non-monetary screening devices in settings with independent types (i.e. there is no peer information). These devices include verification,²⁸ ex-post punishments,²⁹ promises of future allocations,³⁰ evidence,³¹ costly signaling,³² heterogeneous risk attitudes,³³ investment and falsification,³⁴ and allocative externalities.³⁵ An important direction for future research is to combine peer information (i.e. correlated types) with such non-monetary devices. Towards this goal, we contribute by analyzing the natural theoretical benchmark with correlation and without other devices and by developing novel tools for this analysis.

Moreover, we contribute to the literature on peer selection at the intersection of economics and computer science (for a survey, see Olckers and Walsh (2024)). In this literature, a principal aims to select an agent based on nominations from their peers (or grades, rankings, etc.). Our model differs from this framework in that nominations are not taken as primitives. Instead, the principal’s payoff from an allocation depends on arbitrary private information that agents hold about one another; nominations are

²⁷Cr mer and McLean (1985, 1988) also consider full surplus extraction in *dominant-strategy* IC mechanisms, but there they still allow agents’ participation decisions to depend on their beliefs.

²⁸See Ben-Porath et al. (2014), Epitropou and Vohra (2019), Erlanson and Kleiner (2019), Halac and Yared (2020), Erlanson and Kleiner (2024), and Patel and Urgan (2025).

²⁹See Mylovanov and Zapechelnyuk (2017) and Li (2020).

³⁰See Kov c et al. (2013), Lipnowski and Ramos (2020), Guo and H rner (2021), and Li and Libgober (2024).

³¹See Ben-Porath et al. (2019, 2023).

³²See Condorelli (2012), Chakravarty and Kaplan (2013), and Akbarpour et al. (2023).

³³See Ortoleva et al. (2021).

³⁴See Perez-Richet and Skreta (2022), Augias and Perez-Richet (2023), Perez-Richet and Skreta (2023), and Li and Qiu (2024).

³⁵See Bhaskar and Sadler (2019) and Goldl cke and Tr ger (2025).

a special case in which an agent’s nomination equals their type.

We also depart from this literature by taking an optimal mechanism design approach. One strand of the literature, following de Clippel et al. (2008) and Holzman and Moulin (2013), takes an axiomatic approach. The central axiom—*impartiality*—requires that an agent’s submitted nomination does not affect their chances of being selected. Impartiality is equivalent to DIC in our setting. Our results therefore provide new insights about the set of impartial nomination rules, including a graph-theoretic characterization of the set and its extreme points, and a complexity-theoretic result about deterministic rules. Another strand, following Alon et al. (2011), evaluates mechanisms on their approximation guarantees relative to various benchmarks. We focus on expected values, leading to different results and calling for different techniques.³⁶

Kurokawa et al. (2015) propose the *credible subset mechanism*, which shares similarities with our ranking-based mechanisms. In their model, motivated by conference peer review, n agents apply to k slots and each agent provides a numerical grade for m other agents. The credible subset mechanism selects k agents from those that would make it into the top k based on average peer grades if agents gave the worst possible grades for others. Similar to ranking-based mechanisms, incentive compatibility requires that sometimes no agent be selected. The credible subset mechanism provides a good approximation guarantee relative to naively selecting the k best agents if the number of slots k is large relative to the number of grades m per agent. From the perspective of our paper, this criterion for the approximate optimality of the credible subset mechanism is a specific instance of informational smallness. Our ranking-based mechanisms, in conjunction with our notion of informational size, can be regarded as generalizing this logic to allocation problems with general forms of peer information and without restrictions on the number of slots or ratings per agent.

Bloch and Olckers (2021, 2022) and Baumann (2026) study models in which agents hold information about their neighbors’ allocation values in a social network. Bloch and Olckers (2021, 2022) assume that agents observe ordinal comparisons between their neighbors’ values. They ask for which networks the principal can reconstruct the ordinal ranking of agents from their reports if each agent desires a high rank. Baumann (2026) assumes that each agent perfectly observes their neighbors’ values

³⁶E.g., jury mechanisms can be optimal in our setup. The *2-partition mechanism* of Alon et al. (2011), an analogue of jury mechanisms, is not optimal by their criterion (Fischer and Klimm (2015)). Our focus on expected values calls for an investigation of extreme points, which have gone unstudied.

and that information is partially verifiable. Baumann asks for which networks the principal can fully implement the first-best of selecting the highest value agent (under various equilibrium notions).

Finally, we contribute to the literature on randomization in mechanism and market design. Several papers provide conditions under which it is without loss to focus on deterministic mechanisms (e.g. Kováč and Mylovanov (2009), Jarman and Meisner (2017), Chen et al. (2019), and Rivera Mora (2022)). Our setting with peer information is different as deterministic mechanisms typically do not suffice for implementation and optimality. The market design literature has also used techniques from combinatorial optimization to study whether random allocations can be realized as mixtures of deterministic ones. For comparison with integrality criteria used in that literature, note that the set of DIC mechanisms can be written as the set of solutions to $Aq \leq 1$ and $q \geq 0$, where A is a $\{0, 1\}$ -valued matrix encoding (Feasibility). The extreme points of this set are all integral (i.e. represent deterministic mechanisms) if and only if A is a perfect matrix. Perfection of A is equivalent to perfection of the feasibility graph G . Perfection of A is weaker than a condition called balancedness, which in turn is weaker than total unimodularity. For these definitions and implications, see e.g. Conforti et al. (2001). The literature has applied total unimodularity and balancedness to problems in which the constraints are of the form $c \leq Aq \leq b$ for varying vectors c and b (Budish et al. (2013); Pycia and Ünver (2015)). For us, the exact condition characterizing integrality is perfection since the right-hand side of our program is the all-ones vector. Moreover, while the market design literature seeks to find large domains under which all extreme points are integral, our main insight is that randomization is the norm in our setting.

7 Concluding Remarks

We conclude by discussing extensions and shortcomings of our analysis, including mandatory allocation, the allocation of multiple goods, restricted environments, costly information acquisition, nepotism, and coalition-proofness.

Mandatory allocation. We have considered mechanisms where the principal may withhold the good from the agents. The principal may wish to do so when the (opportunity) cost of providing the good exceeds the benefit from allocating or because

withholding the good helps address incentive constraints. However, the principal may sometimes be constrained to allocate the good: it may be untenable to withhold social aid or the assignment of an important task, or resources are fully allocated due to ratchet effects in a broader organizational context (e.g. Liebman and Mahoney (2017)).

The working paper (Niemeyer and Preusser (2025)) establishes analogous results for mandatory allocation. For the characterization and complexity results (Section 4), the only change is to raise the threshold for the number of agents by one. For example, if there are at least *four* agents and each agent has many possible types, then stochastic extreme-point mechanisms that always allocate are prevalent among the extreme points of the set DIC mechanisms that always allocate. For an intuition, think of the principal keeping the object as the principal allocating to a default agent $n + 1$ whose report never influences the allocation. (This intuition is incomplete since not all DIC mechanisms that always allocate admit a default agent.) Similarly, ranking-based mechanisms can be modified by designating a default agent.

Multiple Goods. Our qualitative insights extend to a setting in which the principal has k units available (e.g. several research grants) and allocates to each agent at most one unit (Niemeyer and Preusser (2025)). Ranking-based mechanisms remain approximately optimal. In the graph-theoretic characterization of DIC mechanisms, stable sets are replaced with subgraphs that contain no clique of size greater than k . Finding an optimal deterministic mechanism remains NP-hard if $k \leq n - 2$. This result implies the existence of stochastic extreme points.

Restricted environments. Without restrictions on the environment—the joint distribution of the agents’ types and the principal’s payoffs—all extreme points of the set of DIC mechanisms are candidates for optimality. Since these mechanisms admit no simple description, it is worth asking if there are meaningful restricted environments with a more tractable subset of candidate mechanisms. We mention two directions that do not seem promising.

First, peer information is uninteresting in the following standard symmetric environment: each agent knows their own allocation value and these values are correlated, exchangeable random variables. If the good must be allocated, it is optimal to ignore all reports and pick a winner at random (Niemeyer and Preusser (2025)).

Second, one could restrict type spaces to be binary. This is a special case where deterministic DIC mechanisms are optimal among all DIC mechanisms. For binary

types, finding an optimal deterministic DIC mechanism corresponds to finding an optimal weighted hypercube matching. Hypercube matchings are investigated in the mathematical literature; simple descriptions are not to be expected. For example, the number of these matchings is unknown except for small n (Östergård and Pettersson (2013)). There is a description of all (perfect) matchings for $n \leq 4$ in Fink (2009).

Costly information acquisition. Agents may incur significant costs when acquiring information to evaluate their peers (e.g., when evaluating a grant proposal). These costs may strengthen the case for ranking-based mechanisms: if each agent can feasibly acquire information only about a few other agents, then each agent is informationally small if there are many agents.

Recent discussions in science funding suggest to randomize the allocation among agents receiving the best peer reviews since it can be too costly for reviewers to further distinguish among the best proposals (Fang and Casadevall (2016)). Although our model does not feature such costs, some mechanisms, such as ranking-based ones, may still randomize among agents to ensure DIC. Thus, providing incentives for honest evaluations can be seen as a supporting argument for randomization.

Nepotism. We have assumed that each agent only cares about their own allocation. In practice, however, agents may also care about the allocations of others, such as their friends and family when it comes to the allocation of targeted aid (for empirical evidence, see e.g. the survey by Olckers and Walsh (2024)). In this case, an analogous notion of DIC might require that each agent must influence neither their own allocation nor the allocation of their friends. One could then analyze optimal mechanisms based on a modified feasibility graph.³⁷ Ranking-based mechanisms would continue to work well if each agent has few friends relative to the total number of agents and is also informationally small for the ranks of their friends.

Coalition-proofness. In our model, only constant mechanisms are coalition-proof, where coalition-proofness means that no subset of agents can improve its total allocation probability via a joint misreport.³⁸ This result is not surprising as the principal has

³⁷In this graph, each vertex specifies an agent i and types of all agents except i and i 's friends. As in the feasibility graph, such a vertex should be viewed as a commitment to allocate to agent i . Vertices are adjacent if the commitments potentially violate feasibility.

³⁸Consider a coalition of two agents. Coalition-proofness implies DIC, meaning no agent can change their own allocation. The total allocation probability in a two-person coalition must be fixed, and therefore no agent can change the allocation of their coalition partner either.

no screening devices beyond the allocation itself.

A Proofs

A.1 Preliminaries

We recall some definitions and present some basic facts about the feasibility graph.

Definitions. Here, let G be a simple undirected graph with vertices V and edges E .

Given a subset $V' \subset V$, the *subgraph of G induced by V'* , denoted $G[V']$, is the graph on vertices V' where two vertices are adjacent if and only if they are adjacent in G . A graph is an *induced subgraph of G* if it is induced by a subset of vertices.

We next define perfection. The *clique number* of G is the maximum cardinality $|C|$ across cliques C of G . For an integer k , a partition of V into k non-empty stable sets is a *k -node-colouring*. The *chromatic number* of G is the smallest integer k such that G admits a k -node-colouring. The graph G is *perfect* if for all induced subgraphs (V', E') of G the clique number of (V', E') equals the chromatic number of (V', E') .

The complement $\bar{G}$ of G is the graph on the same set of vertices as G but where a pair of vertices is adjacent in $\bar{G}$ if and only if the pair is non-adjacent in G .

The *fractional stable set polytope* of G , denoted $QSTAB(G)$, is the set of functions $q: V \rightarrow [0, 1]$ such that all maximal cliques X satisfy $\sum_{v \in X} q(v) \leq 1$. The *stable set polytope* of G , denoted $STAB(G)$, is the convex hull of the set of vectors that are incidence vectors of stable sets of G .

Notation and basic facts. In what follows, G refers to the feasibility graph.

For some proofs with three or more agents it will be convenient to denote vertices as follows. Given a type profile θ , we write θ_{-123} to mean the agents other than 1 to 3 (if such other agents exist). We also write $(\cdot, \theta_2, \theta_3, \theta_{-123})$ to mean the vertex $(1, \theta_{-1})$, we write $(\theta_1, \cdot, \theta_3, \theta_{-123})$ to mean $(2, \theta_{-2})$, and so on.

We say a type profile θ *contains* the vertices $\{(i, \theta_{-i}): i \in [n]\}$. If (i, θ_{-i}) and (j, θ'_{-j}) are adjacent vertices, then $\hat{\theta} = (\theta'_i, \theta_j, \theta_{-ij})$ is the unique type profile containing (i, θ_{-i}) and (j, θ'_{-j}) . The map $\theta \mapsto \{(i, \theta_{-i}): i \in [n]\}$ is a bijection from types profiles to maximal cliques. It follows that for all two adjacent vertices v and v' , there is a unique maximal clique containing both v and v' .

A.2 Proof of Theorem 1

To show that $\text{OPTDET-}n$ is NP-hard if $n \geq 3$, it suffices to show that the following decision problem, DET-3 , is NP-complete.

Definition 6 ($\text{DET-}n$). The input consists of an integer k , finite sets $\Theta_1, \dots, \Theta_n$ and weights $w_i : \Theta_{-i} \rightarrow \{0, 1\}$ for all $i \in [n]$. The problem is to decide whether there is a deterministic DIC mechanism q (for n agents with respective type spaces $\Theta_1, \dots, \Theta_n$) such that $\sum_{i, \theta} w_i(\theta_{-i}) q_i(\theta_{-i}) \geq k$.

Recall that STABLESET refers to the following decision problem. The input is an integer $\hat{k}$ and a (simple undirected) graph $\hat{G}$. The problem is to determine whether $\hat{G}$ admits a stable set of cardinality $\hat{k}$ or greater. STABLESET is NP-complete (Korte and Vygen (2018, Theorem 15.23)). To show that DET-3 is NP-complete, we show that STABLESET polynomially reduces to DET-3 .

Let $(\hat{G}, \hat{k})$ be an instance of STABLESET . We next construct our candidate instance of DET-3 . Let $\hat{V}$ and $\hat{E}$ denote the vertex and edge sets of $\hat{G}$. Fix an arbitrary linear order $<$ on $\hat{V}$ (say, by enumeration). Define $\Theta_1 = \hat{E}$ and $\Theta_2 = \Theta_3 = \hat{V}$. The resulting feasibility graph is denoted G , with vertices V and edges E . Next, for all edges e of $\hat{G}$, define a set $P(e)$ of vertices in G as follows: denote the edge e as $e = vv'$ such that $v < v'$, and define

$$P(e) = ((\cdot, v, v), (e, v, \cdot), (e, \cdot, v'), (\cdot, v', v')).$$

Notice that $P(e)$ is an induced path in G from $(\cdot, v, v)$ to $(\cdot, v', v')$. We refer to $(\cdot, v, v)$ and $(\cdot, v', v')$ as the *endpoints* of $P(e)$; the vertices $(e, v, \cdot)$ and $(e, \cdot, v')$ are the *interior* vertices of $P(e)$. (Likewise, a vertex is an *endpoint* if for some e it is an endpoint of $P(e)$; a vertex is *interior* if for some e it is an interior vertex of $P(e)$.) Next, define $w : V \rightarrow \{0, 1\}$ as follows. For all $v \in V$, let $w(v) = 1$ if there is an edge e in $\hat{E}$ such that $v \in P(e)$; else, let $w(v) = 0$. Finally, let $k = \hat{k} + |\hat{E}|$.

To complete the proof, we argue that $\hat{G}$ admits a stable set of cardinality $\hat{k}$ or greater if and only if G admits a weighted stable set of weight k or greater.

Let $\hat{S}$ be a stable set in $\hat{G}$ of cardinality $\hat{k}$ or greater. Construct a stable set S in G as follows. Denote $T = \{(\cdot, v, v) : v \in \hat{S}\}$, and include T in S . Next, for all edges $e \in \hat{E}$, find an interior vertex ω of $P(e)$ that is non-adjacent to all vertices in T , and include ω in $\hat{S}$; such a vertex ω exists since $\hat{S}$ is stable, meaning that the two endpoints of $P(e)$ are not both included in T . Clearly, S has weight of at least $k = \hat{k} + |\hat{E}|$.

Moreover, S is stable since T is stable, since the interior vertices are chosen to be non-adjacent from T , and since interior vertices of distinct paths are non-adjacent.

Now suppose G admits a stable set S of weight $\hat{k} + |\hat{E}|$ or greater.

In an auxiliary step, we argue that G admits a stable set having weight $\hat{k} + |\hat{E}|$ or greater and such that for all $e \in \hat{E}$ the stable set contains at most one endpoint of $P(e)$. Indeed, suppose for some $e \in \hat{E}$ both endpoints of $P(e)$ are in a stable set of G . Note that none of the interior vertices of $P(e)$ are in the stable set, by stability. Now consider the following adjustment to the stable set: choose one of the endpoints of $P(e)$, find the adjacent interior vertex in $P(e)$, and in the stable set replace the chosen endpoint by the interior vertex. The resulting set is stable since in G the chosen interior vertex is only adjacent to the other interior vertex of $P(e)$ (which we already noted is not in the stable set) and to the removed endpoint of $P(e)$. Further, the adjustment leaves the total weight unchanged. By repeating this adjustment a finite number of times, we obtain a stable set with the claimed properties.

Thus suppose G admits a stable set with weight $\hat{k} + |\hat{E}|$ or greater and such that for all $e \in \hat{E}$ the stable set S contains at most one endpoint of $P(e)$. Note that S contains at most $|\hat{E}|$ interior vertices. Thus S contains at least $\hat{k}$ endpoint vertices. Now define $\hat{S} = \{v \in \hat{V} : (\cdot, v, v) \in S\}$. The set $\hat{G}$ is stable in $\hat{G}$ since, for all $e \in E$, at most one endpoint of $P(e)$ is in S . Moreover, since S contains at least $\hat{k}$ endpoint vertices, we conclude $\hat{S}$ has cardinality of at least $\hat{k}$. $\square$

A.3 Proofs for Section 4.2

In this part of the appendix, we prove [Theorem 2](#) and [Lemma 2](#).

A.3.1 Auxiliary results

As explained in [Section 4.2](#), the existence of stochastic extreme DIC mechanisms is connected to odd holes in the feasibility graph. The goal of the next auxiliary results is to establish that all extreme DIC mechanisms are deterministic if and only if the feasibility graph does not have an odd hole of length 7 or greater ([Lemma 5](#)).

For $i \in [n]$, define $V_i = \{i\} \times \Theta_{-i}$ as the set of i -vertices. Given a vertex $v = (i, \theta_{-i})$ and j distinct from i , we say θ_j is the *type of j at v* .

Lemma 3. *Let v, v', v'' be distinct vertices, where v' is an i -vertex. Let v' be adjacent to v and v'' . If v and v'' are non-adjacent, then i 's type at v differs from i 's type at v'' .*

Proof of Lemma 3. The type profile containing v and v' coincides with the type profile containing v' and v'' in the types of all agents other than i . If these two profiles were to also agree in i 's type, then v and v'' would either be adjacent or coincide. $\square$

Lemma 4. *The feasibility graph does not admit an odd 5-hole. The complement of the feasibility graph does not admit an odd hole.*

Proof of Lemma 4. Towards a contradiction, let G admit an odd 5-hole $(v_1, \dots, v_5)$. Recall that two vertices are adjacent only if they belong to different agents. Since the hole contains five vertices, there is an agent i such that the hole contains exactly one vertex of this agent. Without loss, let the vertex belonging to i be the vertex v_2 . Since (v_3, v_4, v_5, v_1) is a path that contains no i -vertex, the type of agent i is constant across (v_3, v_4, v_5, v_1) . However, since H is a hole, vertex v_2 is adjacent to v_1 and v_3 while v_1 and v_3 are non-adjacent. Thus Lemma 3 implies that i 's type at v_1 differs from i 's type at v_3 . Contradiction. Thus G does not admit an odd 5-hole.

Towards a contradiction, let the complement of G admit an odd hole $(v_1, \dots, v_k)$ for some $k \geq 5$. If $k = 5$, then $(v_1, v_3, v_5, v_2, v_4)$ is an odd 5-hole in G ; contradiction. Thus let $k \geq 7$. By definition of the complement, $\{v_1, v_4, v_6\}$ and $\{v_2, v_4, v_6\}$ are cliques in G . We know that there is a unique maximal clique of G containing v_4 and v_6 . Thus this clique also contains both v_1 and v_2 . In particular, either $v_1 = v_2$, or v_1 and v_2 are adjacent in G ; in either case, we have a contradiction to the assumption that $(v_1, \dots, v_k)$ is an odd hole in the complement of G . Thus the complement of G does not admit an odd hole. $\square$

Lemma 5. *All extreme DIC mechanisms are deterministic if and only if the feasibility graph does not admit an odd hole of length seven or more.*

Proof of Lemma 5. Recall that the set of DIC mechanisms equals $QSTAB(G)$, while the convex hull of the set of deterministic DIC mechanisms equals $STAB(G)$. Theorem 3.1 of Chvátal (1975) implies that $STAB(G) = QSTAB(G)$ holds if and only if G is perfect. According to the Strong Perfect Graph Theorem (Chudnovsky et al. (2006)), the feasibility graph is perfect if and only if neither G nor its complement admit an odd hole. The claim follows from Lemma 4. $\square$

A.3.2 Proof of Lemma 2

First, let q be extreme and let K be a non-empty stochastic component of q .

Note, if $v \in K$, $v' \notin K$ and $q(v') > 0$, then v and v' are non-adjacent. Indeed, if v and v' are adjacent and $q(v') \in (0, 1)$, then $v' \notin K$ contradicts the fact that K is a component; if v and v' are adjacent and $q(v') = 1$, then feasibility is contradicted.

It follows that the restriction of q to K is an extreme point of the fractional stable set polytope of $G[K]$, where $G[K]$ is the subgraph induced by K . By the Strong Perfect Graph Theorem and Theorem 3.1 of Chvátal (1975), the graph $G[K]$ or the complement of $G[K]$ admit an odd hole. Lemma 4 implies the complement of $G[K]$ does not admit an odd hole. Thus $G[K]$ admits an odd hole.

It remains to show that q is extreme if q only takes values in $\{0, 1/2, 1\}$ and every stochastic component of q contains an odd hole. Let q' be a DIC mechanism receiving non-zero weight in a convex combination that equals q . Clearly, q and q' agree at the vertices to which q assigns 0 or 1. Now consider a stochastic component K of q . Thus $q(v) = 1/2$ for all $v \in K$. We show $q'(v) = 1/2$ for all $v \in K$. We proceed in two steps.

First, by assumption, the component K contains an odd hole $(v_1, \dots, v_k)$ for some integer k . By construction, we have $q(v_\ell) + q(v_{\ell+1}) = 1$ for all $\ell \in \{1, \dots, k\}$ (where $v_{k+1} = v_1$). Since v_ℓ and $v_{\ell+1}$ are adjacent, all DIC mechanisms q'' from the convex combination satisfy $q''(v_\ell) + q''(v_{\ell+1}) \leq 1$. Thus $q'(v_\ell) + q'(v_{\ell+1}) = 1$ for all $\ell \in \{1, \dots, k\}$. Using that k is odd, we infer $q'(v_1) = 1 - q'(v_2) = \dots = q'(v_k) = 1 - q'(v_1)$. Thus q' equals $1/2$ on the hole $(v_1, \dots, v_k)$.

Second, since K is a component and since there is an odd hole in K where q and q' agree, it suffices to show the following: if v and v' in K are adjacent and $q'(v) = 1/2$, then $q'(v') = 1/2$. To that end, note $q(v) = q(v') = 1/2$. Clearly, all DIC mechanisms q'' from the convex combination satisfy $q''(v) + q''(v') \leq 1$. Since $q(v) + q(v') = 1$, we infer $q'(v) + q'(v') = 1$. Since $q'(v) = 1/2$ by assumption, we conclude $q'(v') = 1/2$. $\square$

A.3.3 Proof of Theorem 2

Steps 1 and 2 prove the necessary and sufficient conditions for a stochastic extreme point to exist. Step 3 shows that most extreme points are stochastic if type spaces are rich and there are at least three agents.

Step 1. *All extreme points are deterministic if $n = 2$ or if all type spaces are binary.*

Let $n = 2$. Recall that two vertices are adjacent only if they belong to distinct agents. It follows that G is bipartite, and hence does not admit an odd hole. Thus Lemma 5 implies that all extreme DIC mechanisms are deterministic.

Let all type spaces be binary.³⁹ We show that G does not admit an odd hole, so that Lemma 5 implies that all extreme DIC mechanisms are deterministic. Let $(v_1, \dots, v_k)$ be an induced cycle; that is, for all $\ell \in \{1, \dots, k\}$, vertex v_ℓ is adjacent to $v_{\ell-1}$ and $v_{\ell+1}$, and non-adjacent to all other vertices of the cycle (where $v_{k+1} = v_1$ is understood). We show k is even. It suffices to show that for all agents i the cycle contains an even number of i -vertices. We use two observations, valid for all $\ell \in \{1, \dots, k\}$. First, if none of $v_{\ell-1}$, v_ℓ and $v_{\ell+1}$ are i -vertices, then i 's type is constant across these vertices. Second, if v_ℓ is an i -vertex, then Lemma 3 implies that the type of i at $v_{\ell-1}$ differs from the type of i at $v_{\ell+1}$. Since agent i has two possible types, the two observations together imply that the cycle contains an even number of i -vertices.

Step 2. *A stochastic extreme point exists if $n \geq 3$ and at least one type space is non-binary.*

We show there exists an odd hole, so that the claim follows from Lemma 5. By possibly relabeling, let Θ_1 and Θ_2 contain $\{0, 1\}$, and let Θ_3 contain $\{0, 1, 2\}$; such agents and type spaces exist by assumption. Fix an arbitrary profile θ_{-123} of agents other than 1, 2 and 3 (if such agents exists). The following seven vertices form an odd hole:

$$\begin{aligned} v_1 &= (\cdot, 0, 0, \theta_{-123}), & v_5 &= (0, 1, \cdot, \theta_{-123}), \\ v_2 &= (1, \cdot, 0, \theta_{-123}), & v_6 &= (0, \cdot, 2, \theta_{-123}), \\ v_3 &= (1, 1, \cdot, \theta_{-123}), & v_7 &= (0, 0, \cdot, \theta_{-123}). \\ v_4 &= (\cdot, 1, 1, \theta_{-123}), \end{aligned}$$

Step 3. *For all $\varepsilon > 0$ there exists $m \in \mathbb{N}$ such that if $|\Theta_i| \geq m$ for all $i \in [n]$ and $n \geq 3$, then $|\det Q| < \varepsilon |\text{ext } Q|$.*

Let $\mathcal{S}$ denote the set of stable sets of the feasibility graph. Recall that there is a bijection between stable sets and deterministic DIC mechanisms. Thus $|\mathcal{S}| = |\det Q|$.

Fix $n \geq 3$. Let $m = \min_{i \in [n]} |\Theta_i|$. By possibly relabeling the type spaces, let each type space contain $\{1, \dots, m\}$. To prove the claim, we prove that if $m \geq 6$, then

$$|\text{ext } Q| \geq |\mathcal{S}| \left(1 + \frac{m-2}{3(n+1)^9} \right).$$

³⁹An alternative proof is to note that DIC mechanisms are isomorphic to fractional matchings in a hypercube graph. Hypercubes are bipartite, and the matching polytope for bipartite graphs is integral.

Let m' denote the largest integer multiple of 3 that is less than m . Note $m' \geq m - 2$.

We begin by constructing a family of odd holes. Fix an arbitrary type profile $\theta_{-123} = (\theta_4, \dots, \theta_n)$ of agents other than 1, 2 and 3 (if such agents exist). For $k \in \{3, 6, \dots, m'\}$, let $H(k)$ denote the set consisting of the following nine vertices:

$$\begin{aligned} v_1 &= (\cdot, k-2, k-2, \theta_{-123}), & v_6 &= (k, k, \cdot, \theta_{-123}), \\ v_2 &= (k-1, \cdot, k-2, \theta_{-123}), & v_7 &= (\cdot, k, k, \theta_{-123}), \\ v_3 &= (k-1, k-1, \cdot, \theta_{-123}), & v_8 &= (k-2, \cdot, k, \theta_{-123}), \\ v_4 &= (\cdot, k-1, k-1, \theta_{-123}), & v_9 &= (k-2, k-2, \cdot, \theta_{-123}), \\ v_5 &= (k, \cdot, k-1, \theta_{-123}), \end{aligned}$$

Note, $H(k)$ is an odd hole. Let $\mathcal{H} = \{H(k) : k \in \{3, 6, \dots, m'\}\}$. Note $|\mathcal{H}| = m'/3$.

We use further auxiliary definitions. Recall that, for the feasibility graph, for each pair of adjacent vertices there is a unique maximal clique that contains both of them. In particular, for all $H \in \mathcal{H}$, for each of the nine pairs of adjacent vertices in H , there is a unique maximal clique containing the pair. Let $\mathcal{X}_H$ denote the nine maximal cliques obtained in this way, and let $V_H = \cup_{X \in \mathcal{X}_H} X$ denote the vertices contained in these cliques (including vertices that are not themselves in H). Let $N(H)$ denote the neighborhood of H , i.e., the vertices adjacent to at least one vertex in H . Note $V_H \subseteq N(H)$. Let L_H denote the number of stable sets in the subgraph induced by V_H .

We next construct a family of stochastic extreme points. Given $H \in \mathcal{H}$ and $S \in \mathcal{S}$, define $q_{H,S} : V \rightarrow [0, 1]$ as follows: for all $v \in H \cup (N(H) \cap (S \setminus V_H))$, let $q_{H,S}(v) = 1/2$; for all $v \in S \setminus N(H)$, let $q_{H,S}(v) = 1$; else, let $q_{H,S}(v) = 0$. We later verify that $q_{H,S}$ is a stochastic extreme DIC mechanism for all $H \in \mathcal{H}$ and $S \in \mathcal{S}$.

We next provide a number of estimates for the number $|\{q_{H,S} : H \in \mathcal{H}, S \in \mathcal{S}\}|$ of stochastic extreme points just constructed.

Fixing an arbitrary $H \in \mathcal{H}$, let $Q_H = \{q_{H,S} : S \in \mathcal{S}\}$. Given $q \in Q_H$, let $q_H^{-1}(q) = \{S \in \mathcal{S} : q_{H,S} = q\}$. We claim that $|\mathcal{S}| \leq |Q_H| L_H$ holds, where we recall L_H that denotes the number of stable sets in the subgraph induced by V_H . To that end, we make two observations. First, inspecting the definition of $q_{H,S}$, if S and S' are two stable sets such that $S \setminus V_H = S' \setminus V_H$, then $q_{H,S} = q_{H,S'}$. Second, since S is stable, the set $S \cap V_H$ is itself a stable set on the subgraph induced by V_H . The two observations imply that $|q_H^{-1}(q)| \leq L_H$ holds for all $q \in Q_H$. Finally, we note that $\{q_H^{-1}(q) : q \in Q_H\}$ partitions $\mathcal{S}$. Hence $|\mathcal{S}| \leq |Q_H| L_H$.

We next claim that all $H \in \mathcal{H}$ satisfy $L_H \leq (n+1)^9$. Recall the definition $V_H = \cup_{X \in \mathcal{X}_H} X$, where $\mathcal{X}_H$ is a set of nine maximal cliques. Each maximal clique of the feasibility graph contains n vertices, and a stable set selects from each maximal clique at most one vertex (or selects nothing). Hence $L_H \leq (n+1)^9$.

Lastly, we claim that if H and H' in $\mathcal{H}$ are distinct, then $Q_H \cap Q_{H'} = \emptyset$. Equivalently, for arbitrary $S, S' \in \mathcal{S}$ and $k, k' \in \{3, \dots, m'\}$, we have $q_{H(k),S} = q_{H(k'),S'}$ only if $k = k'$. To that end, let $S, S' \in \mathcal{S}$ and $k, k' \in \{3, \dots, m'\}$ and $q_{H(k),S} = q_{H(k'),S'}$. Fix two adjacent vertices v and $\hat{v}$ in H . Thus $q_{H(k),S}(v) = q_{H(k),S}(\hat{v}) = q_{H(k'),S'}(v) = q_{H(k'),S'}(\hat{v}) = 1/2$. Note at least one of v and $\hat{v}$ must be in $H(k')$; indeed, else $q_{H(k'),S'}(v) = q_{H(k'),S'}(\hat{v}) = 1/2$ requires that both be in S' , contradicting stability of S' . Without loss, let $v \in H(k')$. In particular, $v \in H(k) \cap H(k')$. Suppose v is an i -vertex, take $j \in \{1, 2, 3\} \setminus \{i\}$ and consider j 's type at v .⁴⁰ Since $v \in H(k) \cap H(k')$, agent j 's type is in both $\{k, k+1, k+2\}$ and $\{k', k'+1, k'+2\}$. Since k and k' are integer multiples of 3, we conclude $k = k'$.

Collecting our work, we get

$$\begin{aligned} |\{q_{H,S} : H \in \mathcal{H}, S \in \mathcal{S}\}| &= \sum_{H \in \mathcal{H}} |Q_H| \geq \sum_{H \in \mathcal{H}} \frac{|\mathcal{S}|}{L_H} \geq \sum_{H \in \mathcal{H}} \frac{|\mathcal{S}|}{(n+1)^9} = \frac{m'}{3} \frac{|\mathcal{S}|}{(n+1)^9} \\ &\geq \frac{m-2}{3} \frac{|\mathcal{S}|}{(n+1)^9}. \end{aligned}$$

Finally, using that for all $H \in \mathcal{H}$ and $S \in \mathcal{S}$ the DIC mechanism $q_{H,S}$ is stochastic and extreme, we infer the promised inequality

$$|\text{ext } Q| \geq |\mathcal{S}| + |\{q_{H,S} : H \in \mathcal{H}, S \in \mathcal{S}\}| \geq |\mathcal{S}| \left(1 + \frac{m-2}{3(n+1)^9}\right).$$

Finally, we show that $q_{H,S}$ is an extreme DIC mechanism for all $H \in \mathcal{H}$ and $S \in \mathcal{S}$.

We first show $q_{H,S}$ is a well-defined DIC mechanism. To show $q_{H,S}$ is well-defined, we have to argue that all maximal cliques X satisfy $q(X) \leq 1$. First, suppose X contains a vertex v in $S \setminus N(H)$. It follows that X contains no vertex in H (else v would be in $N(H)$) nor a vertex in $(N(H) \cap (S \setminus V_H))$ (by stability of S). Hence, in this case, $q_{H,S}$ assigns 1 to v and 0 to all other vertices in X . Next, suppose X only contains vertices in $H \cup (N(H) \cap (S \setminus V_H))$. Now X can contain at most two vertices from H (since H is an odd hole) and at most one vertex in $N(H) \cap (S \setminus V_H)$ (by stability of S). Further, if

⁴⁰Recall, for $i \in [n]$ we call $V_i = \{i\} \times \Theta_{-i}$ the set of i -vertices, and j 's type at (i, θ_{-i}) means θ_j . E.g., if $v = (\cdot, k-2, k-2, \theta_{-123})$, then v is a 1-vertex and the types of agents 2 and 3 are $k-2$ at v .

X contains a vertex in $N(H) \cap (S \setminus V_H)$, then it contains at most one vertex in H (by definition of V_H). Since all vertices in $H \cup (N(H) \cap (S \setminus V_H))$ are assigned $1/2$, we infer $q(X) \leq 1$. Lastly, if X contains no vertices in $(S \setminus N(H)) \cup (H \cup (N(H) \cap (S \setminus V_H)))$, then clearly $q(X) = 0$ holds.

[Lemma 2](#) implies $q_{H,S}$ is extreme since $q_{H,S}$ maps to $\{0, 1/2, 1\}$, and since $H \cup (N(H) \cap (S \setminus V_H))$ is the unique stochastic component of $q_{H,S}$ and contains the odd hole H . $\square$

A.4 Proof of [Theorem 3](#)

Fix $n \in \mathbb{N}$ and consider the environment (n, Θ^n, μ^n) . For all $i \in [n]$, let $\bar{u}_i^n$ denote i 's peer value in this environment. Let r_i^n and $r_i^{n,*}$, respectively, denote i 's rank and robust rank, respectively. For $p \in (0, 1)$, let $q^{n,p}$ denote the ranking-based mechanism with threshold p . Let U^n denote the principal's expected utility (as a function on the set of DIC mechanisms). For all $\ell \in [n]$, let $\bar{u}_{[\ell]}^n(\theta)$ be the ℓ 'th highest peer value.

We show that for all $\varepsilon > 0$ there exists $p \in (0, 1)$ such that

$$\left| U^n(q^{n,p}) - \sum_{\theta \in \Theta^n} \mu^n(\theta) \max\{0, \bar{u}_{[1]}^n(\theta)\} \right| < \varepsilon$$

for all but finitely many n . This claim proves the theorem since, for every environment, no DIC mechanism yields a higher expected utility than the utility from always allocating to an agent with the highest positive peer value; indeed, recall from [\(1\)](#) that every n and every DIC mechanism q in the n 'th environment, we have

$$U^n(q) = \sum_{\theta \in \Theta^n} \sum_{i \in [n]} \mu^n(\theta) q_i(\theta_{-i}) \bar{u}_i^n(\theta_{-i}) \leq \sum_{\theta \in \Theta^n} \sum_{i \in [n]} \mu^n(\theta) \max\{0, \bar{u}_{[1]}^n(\theta)\}.$$

In the n 'th environment, given a type profile $\theta \in \Theta^n$ and $p \in (0, 1)$, recall that $\bar{u}_{[pn]}^n(\theta)$ is the $[pn]$ 'th highest peer value. At every type profile θ , every agent who is allocated the good by $q^{n,p}$ has a peer value of at least $\bar{u}_{[pn]}^n(\theta)$.

For all $n \in \mathbb{N}$ and $p \in (0, 1)$ and $\theta \in \Theta^n$, we have the following lower bound on the

principal's payoff at θ from $q^{p,n}$:

$$\begin{aligned}
U^n(q^{p,n}) &= \sum_{\theta \in \Theta^n} \sum_{i \in [n]} \mu^n(\theta) \max\{0, \bar{u}_i^n(\theta_{-i})\} \frac{\mathbf{1}_{(r_i^{n,*}(\theta_{-i}) \leq p)}}{np} \\
&\geq \sum_{\theta \in \Theta^n} \mu^n(\theta) \max\{0, \bar{u}_{\lfloor pn \rfloor}^n(\theta)\} \sum_{i \in [n]} \frac{1}{np} \mathbf{1}_{(r_i^{n,*}(\theta_{-i}) \leq p)} \\
&\geq \sum_{\theta \in \Theta^n} \mu^n(\theta) \max\{0, \bar{u}_{\lfloor pn \rfloor}^n(\theta)\} \sum_{i \in [n]} \frac{1}{np} \mathbf{1}_{(r_i^n(\theta) \leq p - \delta^n(\theta))} \\
&\geq \sum_{\theta \in \Theta^n} \mu^n(\theta) \max\{0, \bar{u}_{\lfloor pn \rfloor}^n(\theta)\} - 1 + \sum_{\theta \in \Theta^n} \mu^n(\theta) \sum_{i \in [n]} \frac{1}{np} \mathbf{1}_{(r_i^n(\theta) \leq p - \delta^n(\theta))},
\end{aligned}$$

where the last line uses that all peer values are in $[-1, 1]$.

Since no two agents have the same rank (the rank entails a tie-breaking rule),

$$\sum_{\theta \in \Theta^n} \mu^n(\theta) \sum_{i \in [n]} \frac{1}{np} \mathbf{1}_{(r_i^n(\theta) \leq p - \delta^n(\theta))} = \sum_{\theta \in \Theta^n} \mu^n(\theta) \frac{\lfloor n(p - \delta^n(\theta)) \rfloor}{np}.$$

Since $(\delta^n)_{n \in \mathbb{N}}$ converges to 0 in probability, for fixed $p \in (0, 1)$ and as $n \rightarrow \infty$ we have

$$\sum_{\theta \in \Theta^n} \mu^n(\theta) \frac{\lfloor n(p - \delta^n(\theta)) \rfloor}{np} \rightarrow 1.$$

Returning to the lower bound for $U^n(q^{p,n})$, we may complete the proof by showing that for all $\varepsilon > 0$ there exists $p \in (0, 1)$ such that

$$\limsup_{n \rightarrow \infty} \sum_{\theta \in \Theta^n} \mu^n(\theta) (\max\{0, \bar{u}_{[1]}^n(\theta)\} - \max\{0, \bar{u}_{\lfloor pn \rfloor}^n(\theta)\}) < \varepsilon. \quad (2)$$

Recall that $\bar{u}_{[1]}^n(\theta)$ is the maximal peer value. In particular, $\bar{u}_{[1]}^n(\theta) \geq \bar{u}_{\lfloor pn \rfloor}^n(\theta)$. Hence, an upper bound on the sum in (2) equals

$$\sum_{\theta \in \Theta^n} \mu^n(\theta) (\bar{u}_{[1]}^n(\theta) - \bar{u}_{\lfloor pn \rfloor}^n(\theta)).$$

By regularity, there exists p such that this upper bound is at most ε for all sufficiently large n . Hence, (2) follows. $\square$

References

- Akbarpour, Mohammad, Piotr Dworczak, and Frank Yang (2023). “Comparison of Screening Devices”. In: *Proceedings of the 24th ACM Conference on Economics and Computation*, pp. 60–60.
- Alaei, Saeed, Hu Fu, Nima Haghpahan, Jason Hartline, and Azarakhsh Malekian (2019). “Efficient computation of optimal auctions via reduced forms”. In: *Mathematics of Operations Research* 44.3, pp. 1058–1086.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi (2019). “Does elite capture matter? Local elites and targeted welfare programs in Indonesia”. In: *AEA Papers and Proceedings*. Vol. 109, pp. 334–39.
- Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias (2012). “Targeting the poor: evidence from a field experiment in Indonesia”. In: *American Economic Review* 102.4, pp. 1206–40.
- Alon, Noga, Felix Fischer, Ariel Procaccia, and Moshe Tennenholtz (2011). “Sum of us: Strategyproof selection from the selectors”. In: *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 101–110.
- Andreyanov, Pasha and Tomasz Szadzik (2021). “Robust mechanism design of exchange”. In: *The Review of Economic Studies* 88.2, pp. 521–573.
- Augias, Victor and Eduardo Perez-Richet (2023). “Non-Market Allocation Mechanisms: Optimal Design and Investment Incentives”. Working paper.
- Balietti, Stefano, Robert L. Goldstone, and Dirk Helbing (2016). “Peer review and competition in the Art Exhibition Game”. In: *Proceedings of the National Academy of Sciences* 113.30, pp. 8414–8419.
- Baumann, Leonie (2026). “Robust implementation with peer mechanisms and evidence”. In: *Games and Economic Behavior* 155, pp. 192–211.
- Ben-Porath, Elchanan, Eddie Dekel, and Barton L. Lipman (2014). “Optimal allocation with costly verification”. In: *American Economic Review* 104.12, pp. 3779–3813.
- (2019). “Mechanisms with evidence: Commitment and robustness”. In: *Econometrica* 87.2, pp. 529–566.
- (2023). “Sequential mechanisms for evidence acquisition”. Working paper.
- Bhaskar, Dhruva and Evan Sadler (2019). “Resource Allocation with Positive Externalities”. Working paper.

- Bloch, Francis, Bhaskar Dutta, and Marcin Dziubiński (2023). “Selecting a winner with external referees”. In: *Journal of Economic Theory* 211, p. 105687.
- Bloch, Francis and Matthew Olckers (2021). “Friend-based ranking in practice”. In: *AEA Papers and Proceedings*. Vol. 111, pp. 567–71.
- (2022). “Friend-based ranking”. In: *American Economic Journal: Microeconomics* 14.2, pp. 176–214.
- Budish, Eric, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom (2013). “Designing random allocation mechanisms: Theory and applications”. In: *American economic review* 103.2, pp. 585–623.
- Cai, Yang, Constantinos Daskalakis, and S. Matthew Weinberg (2012). “An algorithmic characterization of multi-dimensional mechanisms”. In: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 459–478.
- Carpenter, Jeffrey, Peter Hans Matthews, and John Schirm (2010). “Tournaments and office politics: Evidence from a real effort experiment”. In: *American Economic Review* 100.1, pp. 504–517.
- Chakravarty, Surajeet and Todd R Kaplan (2013). “Optimal allocation without transfer payments”. In: *Games and Economic Behavior* 77.1, pp. 1–20.
- Chen, Yi-Chun, Wei He, Jiangtao Li, and Yeneng Sun (2019). “Equivalence of Stochastic and Deterministic Mechanisms”. In: *Econometrica* 87.4, pp. 1367–1390.
- Chudnovsky, Maria, Neil Robertson, Paul Seymour, and Robin Thomas (2006). “The strong perfect graph theorem”. In: *Annals of Mathematics*, pp. 51–229.
- Chvátal, Vášek (1975). “On certain polytopes associated with graphs”. In: *Journal of Combinatorial Theory, Series B* 18.2, pp. 138–154.
- Condorelli, Daniele (2012). “What money can’t buy: Efficient mechanism design with costly signals”. In: *Games and Economic Behavior* 75.2, pp. 613–624.
- Conforti, Michele, Gérard Cornuéjols, Ajai Kapoor, and Kristina Vučković (2001). “Perfect, ideal and balanced matrices”. In: *European Journal of Operational Research* 133.3, pp. 455–461.
- Crémer, Jacques and Richard P. McLean (1985). “Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent”. In: *Econometrica: Journal of the Econometric Society*, pp. 345–361.
- (1988). “Full extraction of the surplus in Bayesian and dominant strategy auctions”. In: *Econometrica: Journal of the Econometric Society*, pp. 1247–1257.

- de Clippel, Geoffroy, Hervé Moulin, and Nicolaus Tideman (2008). “Impartial division of a dollar”. In: *Journal of Economic Theory* 139.1, pp. 176–191.
- Epitropou, Markos and Rakesh Vohra (2019). “Optimal On-Line Allocation Rules with Verification”. In: *International Symposium on Algorithmic Game Theory*. Springer, pp. 3–17.
- Erlanson, Albin and Andreas Kleiner (2019). “A note on optimal allocation with costly verification”. In: *Journal of Mathematical Economics* 84, pp. 56–62.
- (2024). “Optimal Allocations with Capacity Constrained Verification”. Working paper.
- Fang, Ferric C. and Arturo Casadevall (2016). “Research Funding: the Case for a Modified Lottery”. In: *mBio* 7.2, 10.1128/mbio.00422–16.
- Fink, Jiří (2009). “Connectivity of matching graph of hypercube”. In: *SIAM Journal on Discrete Mathematics* 23.2, pp. 1100–1109.
- Fischer, Felix and Max Klimm (2015). “Optimal impartial selection”. In: *SIAM Journal on Computing* 44.5, pp. 1263–1285.
- Gerardi, Dino, Richard McLean, and Andrew Postlewaite (2009). “Aggregation of expert opinions”. In: *Games and Economic Behavior* 65.2, pp. 339–371.
- Goldlücke, Susanne and Thomas Tröger (2025). “The multiple-volunteers principle”. In: *Journal of Economic Theory*, p. 106114.
- Goold, Michael (2005). “Making peer groups effective: lessons from BP’s experiences”. In: *Long Range Planning* 38.5, pp. 429–443.
- Gul, Faruk and Andrew Postlewaite (1992). “Asymptotic efficiency in large exchange economies with asymmetric information”. In: *Econometrica: Journal of the Econometric Society*, pp. 1273–1292.
- Guo, Yingni and Johannes Hörner (2021). “Dynamic allocation without money”. Working paper.
- Halac, Marina and Pierre Yared (2020). “Commitment versus flexibility with costly verification”. In: *Journal of Political Economy* 128.12, pp. 4523–4573.
- Hart, Sergiu and Noam Nisan (2017). “Approximate revenue maximization with multiple items”. In: *Journal of Economic Theory* 172, pp. 313–347.
- Holzman, Ron and Hervé Moulin (2013). “Impartial nominations for a prize”. In: *Econometrica* 81.1, pp. 173–196.
- Huang, Yifei, Matt Shum, Xi Wu, and Jason Zezhong Xiao (2019). “Discovery of bias and strategic behavior in crowdsourced performance assessment”. Working paper.

- Hussam, Reshmaan, Natalia Rigol, and Benjamin N Roth (2022). “Targeting high ability entrepreneurs using community information: Mechanism design in the field”. In: *American Economic Review* 112.3, pp. 861–98.
- Jarman, Felix and Vincent Meisner (2017). “Deterministic mechanisms, the revelation principle, and ex-post constraints”. In: *Economics Letters* 161, pp. 96–98.
- Kattwinkel, Deniz (2020). “Allocation with Correlated Information: Too good to be true”. In: *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 109–110.
- Kattwinkel, Deniz and Jan Knoepfle (2023). “Costless information and costly verification: A case for transparency”. In: *Journal of Political Economy* 131.2, pp. 504–548.
- Kattwinkel, Deniz, Axel Niemeyer, Justus Preusser, and Alexander Winter (2022). “Mechanisms without transfers for fully biased agents”. Working paper.
- Korte, Bernhard and Jens Vygen (2018). *Combinatorial Optimization: Theory and Algorithms*. 6th ed. 2018. Algorithms and Combinatorics; 21. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Koster, Arie M.C.A. and Annegret K. Wagler (2006). “The extreme points of QSTAB (G) and its implications”. Working paper.
- Kotturi, Yasmine, Anson Kahng, Ariel Procaccia, and Chinmay Kulkarni (2020). “Hirepeer: Impartial peer-assessed hiring at scale in expert crowdsourcing markets”. In: *Proceedings of the AAAI conference on artificial intelligence*. Vol. 34. 03, pp. 2577–2584.
- Kováč, Eugen, Daniel Krähmer, and Tymon Tatur (2013). “Optimal stopping in a principal-agent model with hidden information and no monetary transfers”. Working paper.
- Kováč, Eugen and Tymofiy Mylovanov (2009). “Stochastic mechanisms in settings without monetary transfers: The regular case”. In: *Journal of Economic Theory* 144.4, pp. 1373–1395.
- Kurokawa, David, Omer Lev, Jamie Morgenstern, and Ariel D Procaccia (2015). “Impartial Peer Review.” In: *IJCAI*. Vol. 15. Citeseer, pp. 582–588.
- Li, Yingkai and Xiaoyun Qiu (2024). “Screening Signal-Manipulating Agents via Contests”. Working paper.
- Li, Yunan (2020). “Mechanism design with costly verification and limited punishments”. In: *Journal of Economic Theory* 186, p. 105000.

- Li, Zihao and Jonathan Libgober (2024). “The Dynamics of Verification when Searching for Quality”. Working paper.
- Liebman, Jeffrey B. and Neale Mahoney (2017). “Do expiring budgets lead to wasteful year-end spending? Evidence from federal procurement”. In: *American Economic Review* 107.11, pp. 3510–3549.
- Lipnowski, Elliot and Joao Ramos (2020). “Repeated delegation”. In: *Journal of Economic Theory* 188, p. 105040.
- Lopomo, Giuseppe, Luca Rigotti, and Chris Shannon (2022). “Detectability, duality, and surplus extraction”. In: *Journal of Economic Theory* 204, p. 105465.
- McAfee, R. Preston and Philip J. Reny (1992). “Correlated information and mechanism design”. In: *Econometrica: Journal of the Econometric Society*, pp. 395–421.
- McLean, Richard and Andrew Postlewaite (2002). “Informational size and incentive compatibility”. In: *Econometrica* 70.6, pp. 2421–2453.
- (2015). “Implementation with interdependent valuations”. In: *Theoretical Economics* 10.3, pp. 923–952.
- Mylovanov, Tymofiy and Andriy Zapechelnjuk (2017). “Optimal allocation with ex post verification and limited penalties”. In: *American Economic Review* 107.9, pp. 2666–94.
- Nemhauser, George L. and Leslie E. Trotter (1974). “Properties of vertex packing and independence system polyhedra”. In: *Mathematical programming* 6.1, pp. 48–61.
- Niemeyer, Axel and Justus Preusser (2025). *Optimal Allocation with Peer Information*.
- Olckers, Matthew and Toby Walsh (2024). “Manipulation and Peer Mechanisms: A Survey”. In: *Artificial Intelligence*, p. 104196.
- Ortoleva, Pietro, Evgenii Safonov, and Leeat Yariv (2021). “Who Cares More? Allocation with Diverse Preference Intensities”. Working paper.
- Östergård, Patric R.J. and Ville H. Pettersson (2013). “Enumerating perfect matchings in n-cubes”. In: *Order* 30.3, pp. 821–835.
- Papadimitriou, Christos and George Pierrakos (2011). “On optimal single-item auctions”. In: *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 119–128.
- (2015). “Optimal deterministic auctions with correlated priors”. In: *Games and Economic Behavior* 92, pp. 430–454.

- Papadimitriou, Christos, George Pierrakos, Alexandros Psomas, and Aviad Rubinstein (2022). “On the complexity of dynamic mechanism design”. In: *Games and Economic Behavior* 134, pp. 399–427.
- Papadimitriou, Christos, George Pierrakos, Christos-Alexandros Psomas, and Aviad Rubinstein (2016). “On the complexity of dynamic mechanism design”. In: *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM, pp. 1458–1475.
- Patel, Rohit and Can Ugun (2025). “Costly Verification and Money Burning”. Working paper.
- Pereyra, Juan S. and Francisco Silva (2023). “Optimal assignment mechanisms with imperfect verification”. In: *Theoretical Economics* 18.2, pp. 793–836.
- Perez-Richet, Eduardo and Vasiliki Skreta (2022). “Test design under falsification”. In: *Econometrica* 90.3, pp. 1109–1142.
- (2023). “Fraud-proof non-market allocation mechanisms”. Working paper.
- Postlewaite, Andrew and David Schmeidler (1986). “Implementation in differential information economies”. In: *Journal of Economic Theory* 39.1, pp. 14–33.
- Pycia, Marek and M. Utku Ünver (2015). “Decomposing random mechanisms”. In: *Journal of Mathematical Economics* 61, pp. 21–33.
- Rivera Mora, Ernesto (2022). “Deterministic Mechanism Design”. Working paper.
- Schrijver, Alexander (2003). *Combinatorial optimization: polyhedra and efficiency*. Vol. 24. 2. Springer.
- Vohra, Rakesh (2012). “Optimization and mechanism design”. In: *Mathematical programming* 134, pp. 283–303.