Optimal Allocation with Peer Information*

Axel Niemeyer† Justus Preusser‡

This version: September 10, 2024

Abstract

We study an allocation problem without monetary transfers where agents hold private information about one another, modeled as a general form of correlated information. Such peer information is relevant in a number of settings, including science funding, allocation of targeted aid, or intra-firm allocation. We study optimal dominant-strategy incentive-compatible (DIC) mechanisms and provide characterizations using techniques from the theory of perfect graphs. Optimal DIC mechanisms tend to be complex and involve allocation lotteries that cannot be purified without upsetting incentives. In rich type spaces, nearly all extreme points of the set of DIC mechanisms are stochastic. Finding an optimal deterministic DIC mechanism is NP-hard. We propose the simple class of ranking-based mechanisms and show that they are approximately optimal when agents are informationally small. These mechanisms allocate to agents ranked highly by their peers but strategically deny the allocation to agents suspected of having evaluated their peers dishonestly.

Keywords: mechanism design without transfers, correlated types, peer information, extreme points, informational size, ranking-based mechanisms, perfect graphs, peer selection

JEL codes: D82, D71, C65

*Parts of this paper were previously circulated under the title "Simple Allocation with Correlated Types." We thank Marco Battaglini, Florian Brandl, Modibo Camara, Gregorio Curello, Francesc Dilmé, Tangren Feng, Nathan Hancart, Deniz Kattwinkel, Jan Knöpfle, Nenad Kos, Daniel Krähmer, Stephan Lauermann, Irene Lo, Andrew Mackenzie, Benny Moldovanu, Marco Ottaviani, 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 gratefully acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B04). Preusser gratefully acknowledges financial support from the European Research Council (HEUROPE 2022 ADG, GA No. 101055295 – InfoEcoScience). Preusser is grateful for the hospitality of Yale University, where parts of this paper were written.

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

‡Department of Economics, Bocconi University, *justus.preusser@unibocconi.it*.

1 Introduction

Consider the problem of a principal allocating a good among several agents without using monetary transfers. The principal's payoffs from the allocation depend on the agents' private information, but each agent wants the good for themselves. Taking a mechanism design approach, we study how the principal can use *peer information*—information that agents hold about one another—to implement desirable allocations. Important examples of the use of peer information in practice include peer review in science funding, community-based targeting of financial aid or credit, or the allocation of resources within organizations based on information dispersed among several divisions.[1]

We model peer information as a form of correlated information. 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. Agents may or may not observe their own values, and values may or may not coincide with their willingness to pay. By studying peer information, we depart from much of the mechanism design literature, which typically assumes that each agent knows only their own value and that these values are statistically independent. The absence of transfers, which play a limited role in the examples above, is another key departure from the literature.[2]

Agents may be tempted to misreport the information that they hold about themselves or others if doing so increases their own chances of receiving the good. These incentive problems arise even when an agent is asked to only provide information about the other agents, as they might discredit their peers to improve their own standing. 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; Balietti et al., 2016).

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, re-

[1]For 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 used a peer review mechanism for capital expenditure authorizations across business units (Goold, 2005). Other examples of peer information include group members selecting a leader based on nominations (e.g. Alon et al., 2011; 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.

[2]To understand peer information per se, we also abstract away from non-monetary screening devices that may be relevant in practice, such as costly verification (e.g. Ben-Porath et al., 2014), ex-post punishments (e.g. Mylovanov and Zapechelnyuk, 2017), or promises of future allocations (e.g. Guo and Hörner, 2021).

gardless of what the agent believes about the other agents' types or reports. In our setting, DIC precisely means that an agent's report can never increase their own chances of getting the good. However, an agent's report can still be used to determine the allocation among the other agents.

Since DIC mechanisms provide incentives that are robust to beliefs, peer information cannot be used to screen an agent via their beliefs about the types of their peers. The belief-screening channel is at the heart of the classical constructions of Crémer and McLean (1985, 1988) and McAfee and Reny (1992) in settings with monetary transfers. The channel is also crucial in a few recent papers without transfers that show how to cross-check an agent's report with external signals that are correlated with the agent's type (Kattwinkel, 2019; Bloch et al., 2023; Kattwinkel and Knoepfle, 2023; Pereyra and Silva, 2023).[3] In contrast, our focus is entirely on how peer information can be used to guide the allocation; we do not exploit correlation as a screening device.

The following class of *jury mechanisms* is an instructive starting point. Before consulting anyone, the principal assigns some agents the role of "jurors," the others the role of "candidates." The jurors' reports determine which candidate gets the good. Jurors never win, and candidates are never consulted. Jury mechanisms are DIC and resemble mechanisms used in practice. For example, scientific peer review is founded on the idea that reviewers should have no stake in the decision. But are jury mechanisms optimal? For example, is it possible to improve the mechanism by also eliciting information from the candidates? We show that jury mechanisms are optimal in two special cases: if there are two agents and the principal can keep the good, or if there are up to three agents and the principal must allocate the good. In general, however, other DIC mechanisms can outperform jury mechanisms.

We contribute two sets of results. First, we characterize DIC mechanisms, showing that optimal DIC mechanisms tend to be complex and involve randomization. Second, we present a simple, interpretable class of DIC mechanisms that is approximately optimal when agents are informationally small. Informational smallness roughly means that no agent alone has information that is crucial for evaluating a large fraction of the other agents. In the main text, we focus on the problem where the principal can keep the good, but we provide analogues for all main results if the good must be allocated.

For the first set of results, we characterize DIC mechanisms using an auxiliary undirected graph. This graph encodes the main trade-off in our problem: by allocating to an agent at a given type profile, the principal commits to ignoring the agent's information and must

[3] Kattwinkel et al. (2022) is the only other paper that studies correlated information *across* agents, i.e., correlated types, in an allocation problem without transfers. The key difference is that Kattwinkel et al. (2022) consider Bayesian IC mechanisms. We discuss this paper in more detail in Section 8.

therefore allocate to the agent at all type profiles that differ from the given profile only in the agent's report (otherwise the agent could influence their allocation by misreporting). Vertices of the graph represent commitments to allocate to an agent across type profiles, and edges represent commitments that are incompatible with feasibility (there is only one good to allocate). A mechanism must make non-conflicting commitments, i.e. select non-adjacent vertices. Thus, in graph-theoretic terms, deterministic DIC mechanisms are stable sets of the graph, and stochastic DIC mechanisms are fractional stable sets.

We use this graph-theoretic perspective to investigate the extreme points of the set of DIC mechanisms. Extreme points are crucial for two reasons. First, every DIC mechanism can be represented as a randomization over extreme points. Second, for every extreme point, there is an environment that renders this extreme point uniquely optimal. Thus no extreme point can be disregarded a priori as a potential candidate for optimality.[4]

Using results from the theory of perfect graphs (Chvátal, 1975; Chudnovsky et al., 2006), we find that the set of DIC mechanisms typically admits stochastic extreme points, i.e., mechanisms that for some reports allocate randomly and where this randomization cannot be purified without upsetting incentives. More precisely, stochastic extreme points exist if and only if there are at least three 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, constitute a vanishing fraction of the set of extreme points. We interpret the prevalence of stochastic extreme points to mean that randomization permits a more flexible aggregation of peer information: stochastic DIC mechanisms can better resolve the trade-off between allocating to an agent and using the agent's information since the agent may win with some probability but still influence how the remaining probability is distributed among the other agents. From a technical perspective, we trace this effect to how stochastic mechanisms distribute allocation probability around certain cycles—odd holes—in the auxiliary graph.

Beyond the prevalence of stochastic extreme points, it is difficult to obtain qualitative economic insights about optimal DIC mechanisms. To make this point, we leverage the representation of deterministic DIC mechanisms as stable sets of our auxiliary graph to show that the problem of finding an optimal deterministic DIC mechanism is NP-hard.[6] While this

[4]The extreme points approach has gained traction in the mechanism design literature in recent years (e.g. Manelli and Vincent, 2007; Manelli and Vincent, 2010; Kleiner et al., 2021). However, none of the existing results can be applied to our specific setting without transfers and with correlation.

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

[6]Related hardness results are known for canonical mechanism design problems with transfers and correlation. Papadimitriou and Pierrakos (2011, 2015) show hardness of finding an optimal deterministic ex-post IC and IR mechanism for auctioning a single good to bidders with correlated private values. Papadimitriou et al. (2016, 2022) show hardness of finding an optimal deterministic IC mechanism for selling to a single

result has straightforward implications for computational tractability, our primary interest lies in the economic implications. Specifically, if one could give a simple economic description of deterministic DIC mechanisms, then this description should suggest a computationally efficient procedure for finding an optimal deterministic DIC mechanism.[7] According to our result, however, such an efficient procedure does not exist under the widely believed $P \neq NP$ conjecture. Although there is no formal notion of conceptual complexity in the mechanism design literature, establishing computational complexity is arguably the best available proxy.[8] Note that the lack of a simple economic description extends to all (possibly stochastic) optimal DIC mechanisms since every deterministic DIC mechanism is uniquely optimal in some environment, meaning any description of optimal DIC mechanisms a fortiori requires a description of deterministic DIC mechanisms.

For our second set of results and motivated by the conceptual and computational complexity of optimal DIC mechanisms, we turn away from exact optimality. Instead, we identify interpretable DIC mechanisms that are approximately optimal among all DIC mechanisms. To that end, we define an agent's *peer value* as the principal's expected value of allocating to the agent based only on the reports of the other agents. No DIC mechanism does better than always allocating to an agent with the highest peer value. We therefore investigate mechanisms that allocate to agents who are ranked highly according to their peer values.

We propose *ranking-based mechanisms*, which work as follows. Since agents can diminish the peer values of the other agents, it is not incentive compatible to naively rank agents according to their peer values and randomly select a winner above a given rank threshold. In a ranking-based mechanisms, instead, if an agent above the threshold has been selected to receive the good, the principal first checks for a conflict of interest: does the agent pass the rank threshold robustly, regardless of the agent's report about their peers? Intuitively, the agent must pass the bar in a hypothetical scenario where the agent appraises their competitors (thereby diminishing their own rank). If the agent passes the check, they receive the good. Otherwise, if a conflict of interest is detected, the principal keeps the good.

Ranking-based mechanisms are approximately optimal when agents are informationally small. Indeed, 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

buyer in a repeated interaction over two periods when valuations are correlated across periods.

[7]Papadimitriou et al. (2016, see their introduction) use computational complexity to draw an analogous conclusion regarding the non-existence of a simple characterization of optimal ex-post IR auctions with correlated values.

[8]Conversely, in the context of the multi-good monopoly problem, 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."

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.[9] The performance of a ranking-based mechanism admits a lower bound that depends on informational size and that holds type profile-by-type profile. For vanishing informational size, the principal inefficiently withholds the good only 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 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 narrow, 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. Relative to jury mechanisms, ranking-based mechanisms also have the appeal of not outright excluding any agent: everyone's information is collected, and everyone is a potential recipient of the good.

Outline. In the next section, we describe the model and make some preliminary observations. In Section 3, we show that jury mechanisms are optimal if there are few agents. In Section 4, we introduce the auxiliary graph. In Section 5, we characterize extreme points of the set of DIC mechanisms, emphasizing the prevalence of stochastic mechanisms. In Section 6, we discuss the structure of deterministic DIC mechanisms. In Section 7, we turn to approximate optimality and ranking-based mechanisms. In Section 8, we discuss the related literature in greater detail, and Section 9 collects some concluding remarks. All proofs are in Appendix A. Appendix B collects supplementary material.

[9] Similar notions of informational size appear in earlier work on mechanism design and general equilibrium with information asymmetries (e.g., Gul and Postlewaite, 1992; McLean and Postlewaite, 2002, 2015; Gerardi et al., 2009; Andreyanov and Sadzik, 2021).

2 Model and Preliminaries

2.1 Model

Environment. A principal allocates a good among a number n of agents, where $n \geq 2$. 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.[10] 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). 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 $\theta = (\theta_1, \dots, \theta_n)$. Each agent i has at least two possible types, $|\Theta_i| \geq 2$.

Notation. As usual, 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 $q \colon \Theta \to [0, 1]^n$ such that $\sum_{i=1}^{n} q_i(\theta) \leq 1$ for all $\theta \in \Theta$. Here, q_i denotes agent i's winning probability.[11] Since there is only one good to allocate, the sum of winning probabilities is less than 1. The principal keeps the good whenever the sum is strictly less than 1 at some type profile.

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

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

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

[10]Nothing in our analysis would change if the agent's payoff from being allocated 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.

[11]One 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.

A DIC mechanism is *optimal* if it maximizes the principal's utility across all DIC mechanisms. A mechanism q is *deterministic* if $q_i(\theta) \in \{0, 1\}$ for all $i \in [n]$ and $\theta \in \Theta$. A mechanism is *stochastic* if it is not deterministic.

2.2 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, and 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; 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 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 8 for a detailed discussion of the relation of our paper to the peer selection literature.
- 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 $(u_i, (u_j + \varepsilon_{ij})_{j \in N(i)})$.

2.3 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

$$U(q) = \sum_{\theta \in \Theta} \sum_{i \in [n]} \mu(\theta) q_i(\theta_{-i}) \bar{u}_i(\theta_{-i}). \tag{2.1}$$

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 (2.1) implies that a constant mechanism that allocates to an agent i with the highest unconditional expected value $\mathbb{E}[u_i]$ is optimal.

3 Jury mechanisms

As a benchmark, we briefly discuss jury mechanisms, the simplest class of DIC mechanisms that meaningfully elicit information from the agents.

Definition 1 (Jury mechanisms)**.** A mechanism q is a *jury mechanism* if there is a partition of the set of agents into *jurors* J and *candidates* C such that:

(1) jurors are never allocated the good ($q_i = 0$ for all $i \in J$);

(2) candidates' reports never influence the allocation (q is constant in θ_i for all $i \in C$).

Jury mechanisms ensure DIC in the most straightforward manner: the agents that can influence the allocation can never receive the good.

Theorem 3.1. *If $n = 2$, then every DIC mechanism is a randomization over jury mechanisms; in particular, a jury mechanism is optimal.*

If $n \leq 3$, then every DIC mechanism q that always allocates the good (i.e., $\sum_{i \in [n]} q_i = 1$) is a randomization over jury mechanisms; in particular, a jury mechanism is optimal among mechanisms that always allocate.

With two agents, a jury mechanism involves one juror who recommends whether the other agent (the candidate) should be allocated the good or whether the principal should keep the good. With three agents in the must-allocate case, there is one juror who decides between the two remaining agents (the candidates). Beyond these special cases, we shall see that jury mechanisms are *not* generally optimal: the principal may do better by considering all agents as sources of information and as potential recipients of the good.

Henceforth, we focus on mechanisms in which the principal may keep the good. The one-agent difference in Theorem 3.1 holds more generally since, intuitively, mechanisms with n agents can be thought of as mechanisms with $n + 1$ agents that always allocate and where one agent has no private information. We discuss the mandatory allocation case in more detail in Section 9.

4 A Graph-Theoretic Characterization of DIC Mechanisms

In this section, we characterize DIC mechanisms in terms of an auxiliary graph—the feasibility graph—which we will use to gain insights into the structure of optimal DIC mechanisms.

Recall that by allocating to an agent, in order to satisfy DIC, the principal commits to also allocate to the agent at all type profiles that only differ in the agent'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

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

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 $(\theta'_i, \theta_j, \theta_{-ij})$ in order 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 2 (Feasibility graph)**.** The *feasibility graph* G is the undirected 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 $\hat{\theta}_{-j} = \theta'_{-j}$.

The feasibility graph can be interpreted in terms of the principal's trade-off between allocating to an agent and using the agent's information. Suppose the principal selects vertex (i, θ_{-i}) (i.e. commits to give the good to agent i when others report θ_{-i}). Thus, 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}). Conversely, if 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 vertices that might be selected). In the first case, the principal's benefits from allocating to i; in the second case, from i's information. In Appendix B.1, we depict the feasibility graph in an example.

In graph-theoretic language, 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 *fractional stable set* of G is a function $q : V \to [0, 1]$ such that $\sum_{v \in C} q(v) \leq 1$ for all maximal cliques C in G. Recall that a *clique* of G is a subset of pairwise adjacent vertices, and a clique is *maximal* if it is not contained in another clique. 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 4.1. *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. Specifically, the peer values and the distribution of types imply weights for the vertices of G; the weight 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 the feasibility graph G with the highest cumulative weight; this is an instance of the *maximum weight stable set problem (MWSS)*. The problem of finding an optimal (possibly stochastic) DIC mechanism is the *fractional relaxation* of this problem.

Next, we derive properties of optimal DIC mechanisms as consequences of this graph-theoretic characterization.

5 Extreme Points and Stochastic Mechanisms

The main insight of this section is that the extreme points of the set of DIC mechanisms are typically stochastic.[12] That is, at some type profiles the object is allocated randomly, and this randomization cannot be purified without violating DIC.

Extreme points are crucial for two reasons. First, the set Q of DIC mechanisms is given by the convex hull of its extreme points since it is a polytope in Euclidean space (here, we use the finiteness of the type spaces). Second, the maximum of the principal's utility U over Q is attained at an extreme point of Q since U is a linear functional. In fact, for each extreme point, there exists a distribution of types and values for which this extreme point is uniquely optimal.[13] Therefore, extreme points characterize the set of DIC mechanisms, and no extreme point can a priori be dismissed as irrelevant for optimization.

For conciseness, we say that a DIC mechanism is *extreme* if it is an extreme point of Q. All deterministic DIC mechanisms are extreme. The next theorem addresses the converse: are all extreme DIC mechanisms deterministic?

[12] An *extreme point* of Q is a DIC mechanism that cannot be expressed as a convex combination of two other DIC mechanisms.

[13] The argument is as follows. Let q be an extreme DIC mechanism. Since the set of DIC mechanisms is a polytope in Euclidean space, the mechanism q is exposed; that is, there exists $p\colon [n] \times \Theta \to \mathbb{R}$ such that $p \cdot q > p \cdot q'$ for all DIC mechanisms q' distinct from q. By possibly scaling p, assume p maps to $[-1, 1]$. Let the distribution of type profiles be uniform over all type profiles. For each agent i and profile θ, conditional on θ, let u_i be the degenerate random variable equal to $p_i(\theta)$. In the resulting environment, q is uniquely optimal.

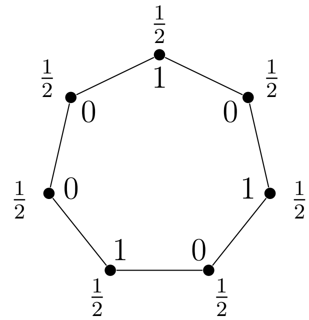

Figure 1: An odd hole in the feasibility graph. The inner vertex labels show how a deterministic mechanism might allocate around the hole; the outer vertex labels show how a stochastic mechanism might allocate around the hole.

Theorem 5.1. *All extreme DIC mechanisms are deterministic if and only if at least one of the following is true:*

(1) there are at most two agents ($n = 2$);

(2) all types spaces are binary ($|\Theta_i| = 2$ for all $i \in [n]$).

(Recall that the model assumes $n \geq 2$ and $|\Theta_i| \geq 2$ for all $i \in [n]$.) Thus, except in two special cases, deterministic mechanisms do not suffice for implementation and optimality. But just how prevalent are stochastic mechanisms among the extreme points?

When type spaces are large, then all but a vanishing fraction of extreme DIC mechanisms are stochastic. To state the result, let $\det Q$ denote the set of deterministic DIC mechanisms, and let $\operatorname{ext} Q$ denote the set of extreme DIC mechanisms.

Theorem 5.2. *Fix $n \geq 3$. For all $\varepsilon > 0$ there exists $m \in \mathbb{N}$ such that if $|\Theta_i| \geq m$ for all $i \in [n]$, then $|\det Q| < \varepsilon\, |\operatorname{ext} Q|$.*

Stochastic extreme DIC mechanisms are characterized by certain cycles—odd holes—in the feasibility graph G. For $k \in \mathbb{N}$, a set $\{v_1, \ldots, v_k\}$ of k vertices of G is an *odd (k-)hole* if $k \geq 5$ is odd and, for all $j \in \{1, \ldots, k\}$, in G 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 $\{v_1, \ldots, v_k\}$. To state the next result, call a *stochastic component of q* an inclusion-wise maximal connected set of vertices v of G such that $q(v) \in (0, 1)$.[14]

Theorem 5.3. *Let q be a stochastic DIC mechanism. If q is extreme, then every stochastic component of q contains an odd hole. The converse holds if $q(v) \in \{0, \frac{1}{2}, 1\}$ for all $v \in V$.*

Theorem 5.3 shows why stochastic mechanisms can outperform deterministic mechanisms: they can distribute allocation mass more flexibly around odd holes in the feasibility

[14] Two vertices are *connected* if there is a path joining them.

graph. For an illustration, consider Figure 1, which depicts an odd hole in the feasibility graph. Each vertex of the hole is a commitment to allocate to a certain agent given a combination of type reports of the other agents. Adjacent vertices are commitments that are incompatible with the (Feasibility) constraints (and therefore cannot be selected simultaneously by a deterministic mechanism). A deterministic DIC mechanisms can select at most three out of seven vertices from the hole. However, a stochastic DIC mechanism may select all seven vertices, each with probability $1/2$, and therefore distribute more total allocation mass around the hole than any deterministic DIC mechanism. In particular, such a stochastic DIC mechanism cannot be written as a convex combination of deterministic DIC mechanisms. In Appendix B.1, we further elaborate on this example, showing in particular how odd holes are embedded in the feasibility graph.

The proofs in this section are based on a theorem 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 the complement of G admit an odd hole. (We define perfection and the complement of a graph in Appendix A.1.1.)

We apply these technical results as follows, leveraging the special structure of the feasibility graph. We first show that the feasibility graph admits no odd 5-holes and that its complement does not admit odd holes at all. We obtain Theorem 5.1 (existence) by checking the feasibility graph for odd holes of length 7 or more. We obtain Theorem 5.3 (characterization) by applying the technical results to each stochastic component separately. The proof of Theorem 5.2 (prevalence) is more difficult since the number of stable sets diverges as the number of possible types increases. However, we can show that the number of certain stable sets in the feasibility graph increases more slowly than the number of odd holes, which is sufficient to establish the result.

Further understanding extreme points of the set of DIC mechanisms requires understanding deterministic DIC mechanisms, i.e., stable sets of the feasibility graph. One reason is that every deterministic DIC mechanism is extreme. Another reason is that stochastic extreme DIC mechanisms may also allocate deterministically on parts of the type space. For example, consider a mechanism with a dummy agent whose report decides whether the principal uses a stochastic extreme DIC mechanism or a deterministic DIC mechanism for the remaining agents. A description of this class of n-agent mechanisms, in particular, requires a description of all $(n-1)$-agent deterministic DIC mechanisms. More generally, the deterministic part of a stochastic extreme DIC mechanism q can be an arbitrary stable set of vertices that are all non-adjacent to the stochastic components of q. These stable sets are themselves stable sets of the entire feasibility graph.

6 The Structure of Deterministic DIC Mechanisms

In this section, we study deterministic allocations. With two agents, deterministic DIC mechanisms are jury mechanisms (which are also the extreme points of the set of DIC mechanisms by Theorem 3.1). Is there an equally straightforward characterization of deterministic DIC mechanisms when there are three or more agents?

If $n \geq 3$, then the problem of finding an optimal deterministic DIC mechanism is NP-hard, suggesting significant hurdles to any simple description of deterministic DIC mechanisms. Before discussing the implications in more detail, let us state the formal result.

Definition 3 (OptDet-n)**.** For $n \in \mathbb{N}$, let OptDet-n be the following optimization problem. The input consists of finite sets $\Theta_1, \dots, \Theta_n$ and weights $w_i : \Theta_{-i} \to \mathbb{Z}$ for all $i \in [n]$. The problem is to find a deterministic DIC mechanism q (for n agents with respective 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.[15]

Theorem 6.1. *If $n \geq 3$, then* OptDet-n *is NP-hard.*

This complexity is in contrast to many allocation problems with independent types where optimal mechanisms can be succinctly described and interpreted in economic terms, often immediately suggesting an efficient computational solution.[16]

For the proof, recall that deterministic DIC mechanisms correspond to stable sets of the feasibility graph; the problem of finding an optimal deterministic DIC mechanisms is the maximum-weight stable set problem (MWSS) on the feasibility graph. For general graphs, it is well-known that MWSS is NP-hard. The content of Theorem 6.1 is that the same conclusion holds even for feasibility graphs. Roughly speaking, we can simulate adjacencies in any graph via induced paths in a feasibility graph for appropriately chosen type spaces and weights.

[15]In a given environment, the weights of interest are given by $w_i(\theta_{-i}) = \mu(\theta_{-i}) u_i(\theta_{-i})$ 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.

[16]For example, consider the allocation problem without transfers in Ben-Porath et al. (2014), where the principal can verify agents' types at a cost. They characterize optimal mechanisms in terms of two parameters: a type cutoff and a "favored" agent. This characterization immediately yields a computationally efficient solution by exhaustive search over the two parameters. For efficient computational solutions in settings with transfers, see e.g. Cai et al. (2012) and Alaei et al. (2019).

Remark 6.2. The conclusion of Theorem 6.1 obtains if one restricts attention to the relatively simple instances of OptDet-n where all weights are in $\{0, 1\}$.

Beyond the straightforward implications regarding computation, Theorem 6.1 suggests that there is no conceptually simple description of optimal deterministic DIC mechanisms or of how they depend on the underlying distribution of types and values. This interpretation rests on the premise that a conceptually simple description would imply an efficient algorithm for finding an optimal deterministic DIC mechanism. Unless $P = NP$, however, such an efficient algorithm does not exist. Although there is no formal notion of conceptual complexity in the mechanism design literature, establishing computational complexity is arguably the best available proxy.

The lack of a conceptually simple description of optimal deterministic DIC mechanisms extends to optimal (possibly stochastic) DIC mechanisms for two reasons.[17] First, every deterministic DIC mechanisms is extreme, and hence an indispensable candidate for optimality; thus, a description of optimal DIC mechanisms entails a description of deterministic ones. Second, as discussed in Section 5, stochastic optimal DIC mechanisms may allocate deterministically on parts of the type space and describing these deterministic parts again entails a description of deterministic DIC mechanisms.

Remark 6.3. One may wonder if more can be said about deterministic mechanisms if all agents' type spaces are binary. This is the special case identified in Theorem 5.1 where deterministic DIC mechanisms are optimal among all DIC mechanisms. With n agents and binary types, the feasibility graph is the line graph of the n-dimensional hypercube. In particular, the problem of determining an optimal deterministic DIC mechanism corresponds to finding an optimal weighted hypercube matching. Hypercube matchings have been 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). One can find a description of all (perfect) matchings for $n \leq 4$ in Fink (2009).

7 Approximate Optimality and Ranking-based Mechanisms

In this section, we shift our focus from optimal DIC mechanisms to approximately optimal ones. Optimal DIC mechanisms lack closed-form descriptions, and it is unclear how these

[17] This claim is about conceptual complexity, not computational complexity (see Footnote 8). The problem of determining an optimal (possibly stochastic) DIC mechanism is solvable in polynomial time by linear programming, although the number of variables and constraints is exponential in the number of agents.

mechanisms depend on the underlying distribution of types and values, making it difficult to institutionalize them in practice. In contrast, the *ranking-based mechanisms* that we propose here are simple, and it is easy to see when and why they perform well. Specifically, we link the performance of ranking-based mechanisms to the impact that any individual agent's report has on the rank of their own peer value. As we will argue, this impact is likely to be small if there many agents.

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*.[18]

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 is obtained in a hypothetical scenario where i appraises their competitors, thereby diminishing their own rank.

A ranking-based mechanism randomly selects an agent ranked, say, in the top 10% of peer values. The principal then allocates to the selected agent if and only if the agent's robust rank is still 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 4 (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:

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

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 pn 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

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

size. 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.

Definition 5. For $\theta \in \Theta$, the *informational size at* θ is denoted $\delta(\theta)$ and given by

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

The 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. We illustrate informally with 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 privately observes u_i and a noisy signal $u_j + \varepsilon_{ij}$ for each of their neighbors $j \in N(i)$, where ε_{ij} is some random variable. If the network is sparse in the sense that the degree $|N(i)|$ of every agent is small relative to the size n of the network, then the informational size is also small since each agent's report only influences the peer values of their neighbors. 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 u_i and 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 the informational size of each agent i is typically small when there are many agents, except at type profiles where peer values are highly concentrated.[19] 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 typically small whenever there are many agents, except at type profiles where peer values are highly concentrated. ▲

Remark 7.1. Other 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, 2015; Gerardi et al., 2009; Andreyanov and Sadzik, 2021). In those models, the agents have private signals about an underlying unobserved state

[19] With unboundedly informative signals (e.g. normally distributed noise), informational size may also be large at type profiles where extreme signal realizations significantly move the principal's posterior.

that is payoff-relevant for all agents; informational size measures the impact of individual signals on the posterior about the state conditional on the full profile of private signals. For our specific allocation problem, it is instead natural to define informational size in terms of peer values and ranks.

The next result shows that ranking-based mechanisms become approximately 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 the informational size is large, such as type profiles where peer values are highly concentrated.

The result uses a minor technical assumption.

Definition 6 (Regularity)**.** A sequence $(n, \Theta^n, \mu^n)_{n\in\mathbb{N}}$ of environments with associated peer values $(\bar{u}^k)_{k\in\mathbb{N}}$ is *regular* if for all $\varepsilon > 0$ there exists $\eta > 0$ such that

$$\lim_{n\to\infty} \mu^n \left\{\theta \in \Theta^n \colon \frac{1}{n} \left| \left\{ j \in [n] \colon \bar{u}^n_j(\theta_{-j}) + \varepsilon \geq \max_{i\in[n]} \bar{u}^n_i(\theta_{-i}) \right\} \right| \geq \eta \right\} = 1. \tag{7.1}$$

Regularity says that the number of agents who are competitive with the top-ranked agent is proportional to the total number of agents. 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 7.2. *Let $(n, \Theta^n, \mu^n)_{n\in\mathbb{N}}$ be a regular sequence of environments. Suppose the associated informational size $(\delta^n)_{n\in\mathbb{N}}$ converges to 0 in probability; that is, for all $d > 0$,*

$$\lim_{n\to\infty} \mu^n \{\theta \in \Theta^n \colon \delta^n(\theta) > d\} = 0. \tag{7.2}$$

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

The proof goes as follows. Fixing $p \in (0, 1)$ and letting the informational size vanish in probability, we show that the ex-ante probability that q^p withholds the good vanishes. Conditional on being allocated, the good is allocated to an agent with rank below p. Under regularity, if p is close to 0, the peer value of such an agent is close to the highest peer value (with high probability). But no DIC mechanism can do better than always allocating to the highest peer value, and this upper bound is approximately attained by the optimal ranking-based mechanism. For every fixed environment, 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 a 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 (Sections 4 to 6). However, if informational size is small, then an individual agent's information is less important for determining their value relative to others, and hence there is little loss from ignoring this information when deciding the agent's allocation. Therefore, simple ranking-based mechanisms can approximately obtain the upper bound of always allocating to the agent with the highest peer value.

Remark 7.3. The performance of jury mechanisms (Section 3) in allocation problems with many agents is less clear: vanishing informational size alone does not imply that jury mechanisms are approximately optimal (Appendix B.2.1). Jury mechanisms may underperform since the roles of jurors and candidates are assigned before any information is elicited, but whether agents fit their assigned roles may vary across type profiles. Nevertheless, when there are many agents, there is intuitively little loss from ignoring the candidates' reports if, from an ex-ante perspective, agents are exchangeable as suppliers of information about others. Likewise, there is little loss from excluding the jurors as potential recipients of the good if, from an ex-ante perspective, agents are exchangeable as recipients of the good. In Appendix B.2.2, we formalize two notions of exchangeability under which jury mechanisms are indeed approximately optimal with many agents.

8 Related Literature

This paper relates to the literature on correlation in mechanism design, on allocation without transfers, on peer selection, and on the role of randomization in mechanism design.

In settings with transfers, Crémer and McLean (1985, 1988) and McAfee and Reny (1992) show how the principal can exploit correlation to extract all private information without leaving information rents to the agents.[20] The principal offers each agent a menu of side-bets regarding the types of the other agents. Different types of a given agent hold different beliefs about the types of others, leading different types to self-select into different lotteries (under an assumption on the correlation structure). These constructions are infeasible in our

[20] See also Lopomo et al. (2022) for a recent treatment.

problem since there are no monetary transfers and since we demand robustness to beliefs.[21]

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 mechanisms with many agents. Since all DIC mechanisms are 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 not aware of other papers on allocation problems without transfers where information is correlated across agents. Other papers instead consider correlation in the form of an external signal that the principal can use to cross-check an agent's type report (Kattwinkel, 2019; Bloch et al., 2023; Kattwinkel and Knoepfle, 2023; Pereyra and Silva, 2023).

The literature on allocation problems without transfers typically maintains the assumption of independent types and studies non-monetary screening devices, including verification (Ben-Porath et al., 2014; Epitropou and Vohra, 2019; Erlanson and Kleiner, 2019, 2024; Patel and Urgun, 2022), ex-post punishments (Mylovanov and Zapechelnyuk, 2017; Li, 2020), promises of future allocations (Kovác et al., 2013; Lipnowski and Ramos, 2020; Guo and Hörner, 2021; Li and Libgober, 2023), evidence (Ben-Porath et al., 2019, 2023), costly signalling (Condorelli, 2012; Chakravarty and Kaplan, 2013; Akbarpour et al., 2023), heterogeneous risk attitudes (Ortoleva et al., 2021), investment and falsification (Perez-Richet and Skreta, 2022; Augias and Perez-Richet, 2023; Perez-Richet and Skreta, 2023; Li and Qiu, 2024), or allocative externalities (Bhaskar and Sadler, 2019; Goldlücke and Tröger, 2023).

We also contribute to the literature on peer selection at the intersection of economics and computer science (for a survey with detailed references, see Olckers and Walsh, 2024). In this literature, the principal aims to select an agent based on nominations (or grades, rankings, etc.) from their peers. One difference to this literature is that we do not start with nominations as primitives. Instead, the principal's payoffs from the allocation are based on the agents' private information. We allow arbitrary peer information, and we can nest nominations as a special case by modeling an agent's nomination as their type. The other key difference is that we take an optimal mechanism design approach.

- One substrand, following de Clippel et al. (2008) and Holzman and Moulin (2013), takes an axiomatic approach. The central axiom—*impartiality*—is equivalent to DIC: one's own submitted nomination should have no impact on one's own chances of being selected. We contribute to this literature via our characterization of DIC mechanisms, showing how the set of impartial nomination rules can be understood in terms of stable sets of the feasibility graph. Moreover, our Theorem 3.1 is related to results by Holzman

[21] Crémer 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.

and Moulin (2013, Proposition 2.i) and Kato and Ohseto (2004, Theorem 5) (obtained in the context of pure exchange economies). Their results can be interpreted in our context as a characterization of three-agent *deterministic* DIC mechanisms that always allocate the good as jury mechanisms. We extend this characterization to stochastic mechanisms. As we highlight, there is often a considerable gap between deterministic and stochastic DIC mechanisms, with the special cases in which jury mechanisms are optimal being among the few exceptions.

- Another substrand, following Alon et al. (2011), provides approximation guarantees relative to various benchmarks in various peer selection contexts. We instead focus on expected values, leading to different results and calling for different techniques.[22] The paper by Kurokawa et al. (2015) deserves special mention; their *credible subset mechanism* shares similarities with our ranking-based mechanisms. In their model, motivated by conference peer review, n agents apply to k slots and each agents 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. In contrast, our ranking-based mechanisms works well for general models of peer information and even when there is only one slot to allocate, as long as agents are informationally small.

Bloch and Olckers (2021, 2022) and Baumann (2023) consider settings where agents hold information about their neighbors' allocation values in an underlying social network. Bloch and Olckers (2021, 2022) assume that agents observe ordinal comparisons between their neighbors' values and ask for which networks the principal can reconstruct the ordinal ranking of agents from their reports when each agent wants to be ranked highly. Baumann (2023) assumes that each agent perfectly observes their neighbors' values 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 the gap between stochastic and deterministic mechanisms (e.g., Kováč and Mylovanov, 2009; Budish et al., 2013; Pycia and Ünver, 2015; Jarman and Meisner, 2017; Chen et al., 2019; Rivera Mora, 2022) by showing that, in

[22]For example, while jury mechanisms can be optimal in our setup, the 2-*partition mechanism* of Alon et al. (2011), which is a natural analogue of jury mechanisms, is not optimal according to their criterion (Fischer and Klimm, 2015). Moreover, our focus on expected values justifies a focus on the extreme points of the set of DIC mechanisms; these have so far gone unstudied.

our setting with correlation and no transfers, deterministic mechanisms typically do not suffice for implementation and optimality. Our tools based on the theory of perfect graphs may be useful for understanding stochastic mechanisms in other settings where the relevant constraints have a combinatorial flavor.

9 Concluding Remarks

We have studied an allocation problem without monetary transfers where agents have private information about their peers, modeled via correlated types, and must be incentivized to report this information truthfully. The broader takeaway is that there is no simple, one-size-fits-all solution for such problems if one insists on exactly optimal mechanisms. Achieving optimality may require the use of allocation lotteries that cannot be purified without upsetting the agents' incentive constraints. At the same time, the combinatorial nature of the incentive constraints implies that finding an optimal deterministic mechanism is NP-hard. This complexity result suggests significant hurdles to any simple economic description of optimal mechanisms. However, incentive constraints matter less for eliciting peer information when agents are informationally small, and simple mechanisms such as ranking-based mechanisms become approximately optimal. We conclude the paper by discussing extensions and shortcomings of our analysis.

Mandatory allocation. We have considered mechanisms where the principal may withhold the good from the agents. The principal may wish to do so for two reasons. First, the (opportunity) cost of providing the good may exceed the benefit from allocating. Second, allocating may be efficient but withholding the good can help address incentive constraints. For example, when there are three agents, jury mechanisms are optimal among DIC mechanisms that always allocate (Theorem 3.1), while optimal DIC mechanisms that withhold the good can be markedly more complex and involve randomization (Theorems 5.1 and 6.1).

Nevertheless, in some applications, the principal may be constrained to always allocate the good to the agents. For example, it may untenable to withhold social aid or the assignment of an important task. Another example is a ratchet effect where the principal allocates all available resources since leftover resources may be interpreted as a signal of reduced needs by the principal's superiors (see e.g. Liebman and Mahoney, 2017).

We provide analogues for our results when the principal must allocate; see Appendix B.4. For the analogues of the characterization and complexity results (Sections 5 and 6), the only change is to increase the threshold values 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 DIC 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's keeping the object as the principal allocating to a default agent $n + 1$ whose report never influences the allocation. (This intuition is not exact since the principal need not designate such a default agent.) In the same vein, ranking-based mechanisms can be modified by designating a default agent.

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. We have argued that these mechanisms admit no simple closed-form description. Hence it may be worth investigating whether there are meaningful restricted environments with a more tractable subset of candidate mechanisms (i.e. a restricted set of objective functionals that only exposes a subset of the extreme points). We note that peer information can be uninteresting if the environment is "too simple." For example, consider a standard symmetric environment in which each agent knows their own allocation value, and these values are exchangeable random variables: if the good must allocated, then it is optimal to ignore all reports and pick a winner at random (Appendix B.3).

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. Interestingly, 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 information acquisition costs, ranking-based mechanisms still randomize among the top agents to ensure DIC. Thus, alleviating incentives for dishonest evaluations can be seen as a supporting argument for randomization in science funding. It would be interesting to include information acquisition costs, perhaps even participation costs (e.g., the costs of preparing an application), more explicitly in the analysis of peer mechanisms.

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, say their friends when it comes to the allocation of targeted aid or their co-authors when it comes to the allocation of science funding (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.[23] Ranking-based mechanisms would continue to work well if each agent has relatively few friends and is also informationally small for the ranks of their friends.

Coalition-proofness. In our model, it is easy to show that only constant mechanisms are coalition-proof in the sense that no coalition of agents can improve its total allocation probability via a joint misreport.[24] This is to be expected since we have abstracted from all screening devices that may be available to the principal (see the discussion in Section 8 for examples of such devices). To understand collusion more generally, it would be interesting to consider a dynamic allocation problem where collusion can be sustained via repeated interactions (see Shah (2022) for a discusion of collusion rings in peer review).

A Proofs

A.1 Preliminaries

In this section, we recall some definitions from graph theory and present some basic facts about the feasibility graph (Section 4).

A.1.1 Definitions

For a moment, let G with vertices V and edges E denote an arbitrary simple undirected graph.

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

We next define perfection. The *clique number* of G equals $\max\{|C|: C \text{ is a clique of } G\}$. For an integer k, a partition of V into k non-empty stable sets is a *k-node-colouring*. The

[23] In 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.

[24] Consider a coalition of just two agents. Coalition-proofness, in particular, implies DIC, i.e., no agent can change their own allocation probability by misreporting their type. However, the total allocation probability in a two-person coalition must be fixed, and therefore no agent can change the allocation probability of their coalition partner either. This notion of coalition-proofness posits that agents run a lottery among themselves to reallocate any gained allocation probability. The impossibility obtains analogously for coalition-proofness based on Pareto improvements.

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 every 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 $q\colon V \to [0,1]$ 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.

A.1.2 Notation and basic facts

In what follows, G refers to the feasibility graph, a vertex means a vertex of the feasibility graph, etc.

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 $\{(i, \theta_{-i})\colon i \in [n]\}$ 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 $\theta \mapsto \{(i, \theta_{-i})\colon i \in [n]\}$ 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 3.1

If $n = 2$ and q is a DIC mechanism that always allocates, then feasibility requires $q_1(\theta_2) + q_1(\theta_1) = 1$. In particular, q is constant. A constant mechanism is necessarily a convex combination of constant deterministic mechanisms. All constant mechanisms are jury mechanisms.

For $n = 2$, a DIC mechanism (that need not always allocate) is a special DIC mechanism with three agents that always allocates and where one agent is a default agent whose report does not influence the allocation but absorbs all residual allocation probability.

Thus, it suffices to prove the claim for three agents and DIC mechanisms that always allocate. In what follows, let $n = 3$, and let $\bar{Q}$ denote set of DIC mechanisms that always allocate.

In our terminology, Holzman and Moulin (2013, Proposition 2.ii) show that all deterministic mechanisms in $\bar{Q}$ are jury mechanisms. Hence, to prove that all mechanisms in $\bar{Q}$ are convex combinations of jury mechanisms, it suffices to show that all extreme points of $\bar{Q}$ are deterministic. Thus, let $q \in \bar{Q}$ be stochastic. We show q is not an extreme point of $\bar{Q}$.

Let $\tilde{V}$ denote the subset of vertices v such that $q(v) \in (0,1)$. We construct non-empty disjoint subsets R (red) and B (blue) of $\tilde{V}$ such that, for all maximal cliques X, either $X \cap (R \cup B) = \emptyset$ or $|X \cap R| = |X \cap B| = 1$. Recall that we may identify maximal cliques with type profiles. If $X \cap (R \cup B) = \emptyset$, we say that the type profile (associated with X) is *uncolored*; if $|X \cap R| = |X \cap B| = 1$, we say that the type profile is *two-colored*.

The existence of such sets R and B imply that q is not an extreme point. Indeed, for a sufficiently small number $\varepsilon > 0$, both $q + \varepsilon(\mathbf{1}_R - \mathbf{1}_B)$ and $q - \varepsilon(\mathbf{1}_R - \mathbf{1}_B)$ are in $\bar{Q}$, and q is a convex combination of $q + \varepsilon(\mathbf{1}_R - \mathbf{1}_B)$ and $q - \varepsilon(\mathbf{1}_R - \mathbf{1}_B)$. (Here, $\mathbf{1}_R$ and $\mathbf{1}_B$, respectively, denote the indicator function for R and B, respectively.)

We assume without loss there exists $(\theta_1^0, \theta_2^0, \theta_3^0)$ such that $(\cdot, \theta_2^0, \theta_3^0)$ and $(\theta_1^0, \cdot, \theta_3^0)$ are in $\tilde{V}$ (since q always allocates). Define $T_1 = \{\theta_1 \in \Theta_1 \colon (\theta_1, \cdot, \theta_3^0) \notin \tilde{V}\}$ and $T_2 = \{\theta_2 \in \Theta_2 \colon (\cdot, \theta_2, \theta_3^0) \notin \tilde{V}\}$. Let $T_1^c = \Theta_1 \setminus T_1$ and $T_2^c = \Theta_2 \setminus T_2^c$. Both T_1^c and T_2^c are non-empty since $(\cdot, \theta_2^0, \theta_3^0)$ and $(\theta_1^0, \cdot, \theta_3^0)$ are in $\tilde{V}$.

First, if T_1 and T_2 are both empty, then the sets $R = \{(\theta_1, \cdot, \theta_3^0) \colon \theta_1 \in \Theta_1\}$ and $B = \{(\cdot, \theta_2, \theta_3^0) \colon \theta_2 \in \Theta_2\}$ have the desired properties.

Henceforth, let $T_1 \neq \emptyset$ (the case $T_2 \neq \emptyset$ being analogous). We distinguish two cases.

Case 1. *Let $T_2 \neq \emptyset$.* Define $\tilde{\Theta}_3$ as the set of $\theta_3 \in \Theta_3$ such that for all $\theta_1 \in T_1^c$ and $\theta_2 \in T_2^c$ we have $(\theta_1, \cdot, \theta_3) \in \tilde{V}$ and $(\cdot, \theta_2, \theta_3) \in \tilde{V}$. Note $\tilde{\Theta}_3$ is non-empty since $\theta_3^0 \in \tilde{\Theta}_3$. We make the following observations.

- **Observation 1.1.** *If $(\theta_1, \theta_2) \in T_1 \times T_2$, then $(\theta_1, \theta_2, \cdot) \notin \tilde{V}$.*
  Indeed, if $(\theta_1, \theta_2, \cdot) \in \tilde{V}$, then agent 3 is the only agent enjoying an interior winning probability at the profile $(\theta_1, \theta_2, \theta_3^0)$, which is impossible since the object is always allocated.
- **Observation 1.2.** *If $(\theta_1, \theta_2) \in (T_1 \times T_2^c) \cup (T_1^c \times T_2)$, then $(\theta_1, \theta_2, \cdot) \in \tilde{V}$.*
  Indeed, the assumption $(\theta_1, \theta_2) \in (T_1 \times T_2^c) \cup (T_1^c \times T_2)$ implies that either agent 1 or agent 2 enjoys an interior winning probability at the profile $(\theta_1, \theta_2, \theta_3^0)$; hence agent 3 must also enjoy an interior winning probability at this profile.
- **Observation 1.3.** *If $\theta_3 \notin \tilde{\Theta}_3$, then all $\theta_1 \in T_1$ and $\theta_2 \in T_2$ satisfy $(\theta_1, \cdot, \theta_3) \in \tilde{V}$ and $(\cdot, \theta_2, \theta_3) \in \tilde{V}$.*
  Indeed, let $\theta_3 \notin \tilde{\Theta}_3$. By definition of $\tilde{\Theta}_3$, there exists $\theta_1' \in T_1^c$ or $\theta_2' \in T_2^c$ such that $(\theta_1', \cdot, \theta_3) \notin \tilde{V}$ or $(\cdot, \theta_2', \theta_3) \notin \tilde{V}$. We first prove the claim in the case in which there is $\theta_2' \in T_2^c$ such that $(\cdot, \theta_2', \theta_3) \notin \tilde{V}$.
  Let $\theta_1 \in T_1$. We show $(\theta_1, \cdot, \theta_3) \in \tilde{V}$. Since $(\theta_1, \theta_2') \in T_1 \times T_2^c$, we have $(\theta_1, \theta_2', \cdot) \in \tilde{V}$ (Observation 1.2). Since also $(\cdot, \theta_2', \theta_3) \notin \tilde{V}$ (by assumption), feasibility at $(\theta_1, \theta_2', \theta_3)$

requires $(\theta_1, \cdot, \theta_3) \in \tilde{V}$, as desired.

Now let $\theta_2 \in T_2$. We show $(\cdot, \theta_2, \theta_3) \in \tilde{V}$. By assumption, T_1 is non-empty. Find $\theta_1 \in T_1$. By the previous paragraph, $(\theta_1, \cdot, \theta_3) \in \tilde{V}$. Since $(\theta_1, \theta_2) \in T_1 \times T_2$, we know $(\theta_1, \theta_2, \cdot) \notin \tilde{V}$ (Observation 1.1). Thus, since $(\theta_1, \cdot, \theta_3) \in \tilde{V}$, feasibility at $(\theta_1, \theta_2, \theta_3)$ requires $(\cdot, \theta_2, \theta_3) \in \tilde{V}$, as desired.

We have proven the claim in the case in which there is $\theta_2' \in T_2^c$ such that $(\cdot, \theta_2', \theta_3) \notin \tilde{V}$. If there exists $\theta_1' \in T_1^c$ such that $(\theta_1', \cdot, \theta_3) \notin \tilde{V}$, the argument is analogous with the roles of agents 1 and 2 switched. Switching the roles is valid since all auxiliary claims proven up to this point hold symmetrically for both agents 1 and 2 and since both T_1 and T_2 are non-empty.

We next define our candidates for R and B. We first define the following:

$$\forall \theta_3 \in \tilde{\Theta}_3, \quad R'(\theta_3) = \{(\cdot, \theta_2, \theta_3) \colon \theta_2 \in T_2^c\} \quad \text{and} \quad B'(\theta_3) = \{(\theta_1, \cdot, \theta_3) \colon \theta_1 \in T_1^c\};$$
$$\forall \theta_3 \notin \tilde{\Theta}_3, \quad R'(\theta_3) = \{(\theta_1, \cdot, \theta_3) \colon \theta_1 \in T_1\} \quad \text{and} \quad B'(\theta_3) = \{(\cdot, \theta_2, \theta_3) \colon \theta_2 \in T_2\}.$$

Let $R_0 = \{(\theta_1, \theta_2, \cdot) \colon \theta_1 \in T_1^c, \theta_2 \in T_2\}$ and $B_0 = \{(\theta_1, \theta_2, \cdot) \colon \theta_1 \in T_1, \theta_2 \in T_2^c\}$. Finally, let $R = R_0 \cup (\cup_{\theta_3 \in \Theta_3} R'(\theta_3))$ and $B = B_0 \cup (\cup_{\theta_3 \in \Theta_3} B'(\theta_3))$.

We next verify $R \cup B \subseteq \tilde{V}$. Observation 1.2 asserts $R_0 \cup B_0 \subseteq \tilde{V}$. Next, for all $\theta_3 \in \tilde{\Theta}_3$, the definition of $\tilde{\Theta}_3$ immediately implies $R'(\theta_3) \cup B'(\theta_3) \subseteq \tilde{V}$, and for all $\theta_3 \notin \tilde{\Theta}_3$ Observation 1.3 implies $R'(\theta_3) \cup B'(\theta_3) \subseteq \tilde{V}$. The sets R and B are disjoint and non-empty (since T_1^c, T_2^c and $\tilde{\Theta}_3$ are all non-empty).

To complete the proof in Case 1, it remains to verify all type profiles are two-colored or uncolored. By inspection, one may verify the following: If $\theta_3 \in \tilde{\Theta}_3$, then all type profiles in $T_1 \times T_2 \times \{\theta_3\}$ are uncolored, whereas all other type profiles in $\Theta_1 \times \Theta_2 \times \{\theta_3\}$ are two-colored. Conversely, if $\theta_3 \notin \tilde{\Theta}_3$, then all type profiles in $T_1^c \times T_2^c \times \{\theta_3\}$ are uncolored, whereas all other type profiles in $\Theta_1 \times \Theta_2 \times \{\theta_3\}$ are two-colored.

Case 2. *Let $T_2 = \emptyset$.* Thus, all $\theta_2 \in \Theta_2$ satisfy $(\cdot, \theta_2, \theta_3^0) \in \tilde{V}$. Define $\hat{\Theta}_3$ as the set of $\theta_3 \in \Theta_3$ such that all $\theta_2 \in \Theta_2$ satisfy $(\cdot, \theta_2, \theta_3) \in \tilde{V}$. Define $\hat{T}_1$ as the set of $\theta_1 \in \Theta_1$ for which there exists $\theta_3 \in \hat{\Theta}_3$ such that $(\theta_1, \cdot, \theta_3) \notin \tilde{V}$. Notice that $\hat{\Theta}_3$ is non-empty since $\theta_3^0 \in \hat{\Theta}_3$; the set $\hat{T}_1$ may or may not be empty.

We observe the following:

- **Observation 2.1.** *If $(\theta_1, \theta_2) \in \hat{T}_1 \times \Theta_2$, then $(\theta_1, \theta_2, \cdot) \in \tilde{V}$.*
  Indeed, by the definition of $\hat{T}_1$, there is $\theta_3 \in \hat{\Theta}_3$ such that $(\theta_1, \cdot, \theta_3) \notin \tilde{V}$. By definition of $\hat{\Theta}_3$, we have $(\cdot, \theta_2, \theta_3) \in \tilde{V}$. Hence feasibility at $(\theta_1, \theta_2, \theta_3)$ requires $(\theta_1, \theta_2, \cdot) \in \tilde{V}$.
- **Observation 2.2.** *If $\theta_3 \notin \hat{\Theta}_3$, then all $\theta_1 \in \hat{T}_1$ satisfy $(\theta_1, \cdot, \theta_3) \in \tilde{V}$.*
  Indeed, by definition of $\hat{\Theta}_3$, there exists $\theta_2 \in \Theta_2$ such that $(\cdot, \theta_2, \theta_3) \notin \tilde{V}$. By the previ-

ous claim, we have $(\theta_1, \theta_2, \cdot) \in \tilde{V}$. Hence feasibility at $(\theta_1, \theta_2, \theta_3)$ requires $(\theta_1, \cdot, \theta_3) \in \tilde{V}$. We next define our candidates for R and B. We first define the following:

$$\begin{aligned}
\forall \theta_3 \in \hat{\Theta}_3, \quad & R'(\theta_3) = \{(\cdot, \theta_2, \theta_3) \colon \theta_2 \in \Theta_2\}; \\
\forall \theta_3 \notin \hat{\Theta}_3, \quad & R'(\theta_3) = \{(\theta_1, \cdot, \theta_3) \colon \theta_1 \in \hat{T}_1\}; \\
& R = \cup_{\theta_3 \in \Theta_3} R'(\theta_3); \\
& B = \{(\theta_1, \cdot, \theta_3) \colon \theta_1 \notin \hat{T}_1, \theta_3 \in \hat{\Theta}_3\} \cup \{(\theta_1, \theta_2, \cdot) \colon \theta_1 \in \hat{T}_1, \theta_2 \in \Theta_2\}.
\end{aligned}$$

Observations 2.1 and 2.2 and the definition of $\hat{\Theta}_3$ imply that $\cup_{\theta_3} R'(\theta_3)$ and $\{(\theta_1, \theta_2, \cdot) \colon \theta_1 \in \hat{T}_1, \theta_2 \in \Theta_2\}$ are contained in $\tilde{V}$. We infer from the definition of $\hat{T}_1$ that $\{(\theta_1, \cdot, \theta_3) \colon \theta_1 \notin \hat{T}_1, \theta_3 \in \hat{\Theta}_3\}$ is also in $\tilde{V}$. Lastly, it is clear that R and B are disjoint, and they are non-empty since $\hat{\Theta}_3$ is non-empty.

To complete the proof in Case 2, it remains to show all type profiles are two-colored or uncolored. By inspection of R and B, one may verify the following: If $\theta_3 \in \hat{\Theta}_3$, then all type profiles in $\Theta_1 \times \Theta_2 \times \{\theta_3\}$ are two-colored; if $\theta_3 \notin \hat{\Theta}_3$, then all type profiles in $\hat{T}_1 \times \Theta_2 \times \{\theta_3\}$ are two-colored whereas all type profiles in $\hat{T}_1^c \times \Theta_2 \times \{\theta_3\}$ are uncolored. □

A.3 Proofs for Section 5

In this part of the appendix, we prove Theorems 5.1 to 5.3. We shall use Theorem 5.3 in the proof of Theorem 5.2, and hence we present the proof of Theorem 5.2 last.

A.3.1 Auxiliary results

As indicated in the main text, the existence of stochastic extreme DIC mechanisms is connected to odd holes in the feasibility graph. The goal of the next few 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 A.3).

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 A.1. *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 A.1. 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. □

Lemma A.2. *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 A.2. Towards a contradiction, suppose G admits an odd 5-hole $(v_1, \ldots, v_5)$. Recall that two vertices are adjacent only if they belong to different agents. Since the hole contains five vertices, there is an agent such that the hole contains exactly one vertex of this agent. Without loss, let the vertex belonging to this agent 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 A.1 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 $(v_1, \ldots, v_k)$ 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 $(v_1, \ldots, v_k)$ is an odd hole in the complement of G. Thus the complement of G does not admit an odd hole. □

Lemma A.3. *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 A.3. 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) (see, e.g., Korte and Vygen, 2018, Theorem 16.21) 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 A.2. □

A.3.2 Proof of Theorem 5.1

Let $n = 2$. Recall that two vertices are adjacent only if they belong to distinct agents. It follows that G is bi-partite, and hence does not admit an odd hole. Thus Lemma A.3 implies that all extreme DIC mechanisms are deterministic.

Let all type spaces be binary. We show that G does not admit an odd hole, so that Lemma A.3 implies that all extreme DIC mechanisms are deterministic.[25] Without loss of generality, relabel the type spaces such that $\Theta_i = \{0, 1\}$ for all $i \in N$. Let $(v_1, \ldots, v_k)$ be

[25] An alternative argument is as follows. If all type spaces are binary, then the feasibility graph corresponds to the line graph of the n-dimensional hypercube. The hypercube is bi-partite, and it is a fact that any line

an induced cycle; that is, for all $\ell \in \{1, \ldots, k\}$, 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 $\ell \in \{1, \ldots, k\}$. 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 implies Lemma A.1 that the type of i at $v_{\ell-1}$ differs from 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.

Now let $n \geq 3$ and suppose at least one type space is non-binary. We show there exists an odd hole. By possibly relabeling the agents and the type spaces, suppose Θ_1 and Θ_2 contains $\{0, 1\}$, and that Θ_3 contains $\{0, 1, 2\}$. 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:

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

□

A.3.3 Proof of Theorem 5.3

First, let q be extreme and let K be a non-empty stochastic component of q. Let $G[K]$ denote the subgraph induced by K.

Note that, 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 we have a contradiction to feasibility.

It follows that the restriction of q to K is an extreme point of the fractional stable set polytope of $G[K]$. According to 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 A.2 implies the complement of $G[K]$ does not admit an odd hole. Thus $G[K]$ admits an odd hole.

graph of a bi-partite graph is perfect. That is, the feasibility graph is perfect, and so the claim follows from the theorem of Chvátal (1975). Another alternative argument proceeds by observing that, when all type spaces are binary, the fractional stable set polytope coincides with the fractional matching polytope of the n-hypercube; it is well-known that this polytope has integer-valued extreme points.

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

(1) By assumption, the component K contains an odd hole $(v_1, \ldots, v_k)$ for some integer k. By construction, we have $q(v_\ell) + q(v_{\ell+1}) = 1$ for all $\ell \in \{1, \ldots, k\}$ (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 $\ell \in \{1, \ldots, k\}$. Using that k is odd, we infer $q'(v_1) = 1 - q'(v_2) = \ldots = q'(v_k) = 1 - q'(v_1)$. Thus q' equals $\frac{1}{2}$ on the hole $(v_1, \ldots, v_k)$.

(2) 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 $q'(v) = \frac{1}{2}$, then $q'(v') = \frac{1}{2}$. To that end, note $q(v) = q(v') = \frac{1}{2}$. 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 $q'(v) = \frac{1}{2}$ by assumption, we conclude $q'(v') = \frac{1}{2}$. □

A.3.4 Proof of Theorem 5.2

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 $n \geq 3$. Denote $m = \min_{i \in [n]} |\Theta_i|$. By possibly relabelling the agents' type spaces, assume that each type space contains $\{1, \ldots, m\}$. We prove that if $m \geq 6$, then

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

which implies Theorem 5.2. 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 $\theta_{-123} = (\theta_4, \ldots, \theta_n)$ of agents other than 1, 2 and 3 (if such agents exist). For $k \in \{3, 6, \ldots, m'\}$, let

$H(k)$ denote the set consisting of the following nine vertices:

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

One may verify that $H(k)$ is an odd hole. Let $\mathcal{H} = \{H(k)\colon k \in \{3, 6, \dots, m'\}\}$. Note $|\mathcal{H}| = \frac{m'}{3}$.

We use some 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; that is, the vertices adjacent to at least one vertex in H. Note $V_H \subseteq N(H)$. Further, 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 $q_{H,S}\colon V \to [0,1]$ as follows:

$$
\forall v \in V, \quad q_{H,S}(v) = \begin{cases} \frac{1}{2}, & \text{if } v \in H \cup (N(H) \cap (S \setminus V_H)); \\ 1, & \text{if } v \in S \setminus N(H); \\ 0, & \text{else.} \end{cases}
$$

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 $|\{q_{H,S}\colon H \in \mathcal{H}, S \in \mathcal{S}\}|$ of stochastic extreme points just constructed.

Fixing an arbitrary $H \in \mathcal{H}$, let $Q_H = \{q_{H,S}\colon S \in \mathcal{S}\}$. Given $q \in Q_H$, let $q_H^{-1}(q) = \{S \in \mathcal{S}\colon q_{H,S} = q\}$. We claim that $|\mathcal{S}| \le |Q_H| L_H$ 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 $\{q_H^{-1}(q) \colon q \in Q_H\}$ 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 $k, k' \in \{3, \ldots, m'\}$, we have $q_{H(k),S} = q_{H(k'),S'}$ only if $k = k'$. To that end, let $S, S' \in \mathcal{S}$ and $k, k' \in \{3, \ldots, m'\}$ and $q_{H(k),S} = q_{H(k'),S'}$. Fix two adjacent vertices v and $\hat{v}$ in H. Thus $q_{H(k),S}(v) = q_{H(k),S}(\hat{v}) = q_{H(k'),S'}(v) = q_{H(k'),S'}(\hat{v}) = \frac{1}{2}$. Note at least one of v and $\hat{v}$ must be in $H(k')$; indeed, else $q_{H(k'),S'}(v) = q_{H(k'),S'}(\hat{v}) = \frac{1}{2}$ 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.[26] 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

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

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

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

as promised.

It remains to prove that for all $H \in \mathcal{H}$ and $S \in \mathcal{S}$ the function $q_{H,S}$ is a stochastic extreme DIC mechanism.

We first show $q_{H,S}$ is indeed 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

[26] Recall that 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. Thus, for example, if $v = (\cdot, k-2, k-2, \theta_{-123})$, then v is an i-vertex and the types of agents 2 and 3 are both $k-2$ at v.

$N(H)$) 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 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 $\frac{1}{2}$, 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.

Finally, Theorem 5.3 implies that $q_{H,S}$ is extreme since $q_{H,S}$ takes values in $\{0, \frac{1}{2}, 1\}$, 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. □

A.4 Proof of Theorem 6.1

Consider the following decision problem:

Definition 7 (Det-n)**.** The input consists of an integer k, finite sets $\Theta_1, \ldots, \Theta_n$ and weights $w_i : \Theta_{-i} \to \{0, 1\}$ 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 $\Theta_1, \ldots, \Theta_n$) such that $\sum_{i,\theta} w_i(\theta_{-i}) q_i(\theta_{-i}) \geq k$.

To show that OptDet-n is NP-hard if $n \geq 3$, it suffices to show that Det-3 is NP-complete.

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 Det-3 polynomially reduces to StableSet.

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

$$P(e) = ((\cdot, v, v), (e, v, \cdot), (e, \cdot, v'), (\cdot, v', v')) . \tag{A.2}$$

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 $w\colon V \to \{0,1\}$ 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 $T = \{(\cdot, v, v) \colon v \in \hat{S}\}$, 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 $P(e)$. Moreover, the adjustment clearly 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 $\hat{S} = \{v \in \hat{V} \colon (\cdot, v, v) \in S\}$. 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 $\hat{k}$. □

A.5 Proof of Theorem 7.2

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, and let U^n denote the principals' expected utility (as a function on the set of DIC mechanisms).

Under the hypothesis that the sequence of environments is regular and informational size vanishes in probability, we show that for all $\varepsilon > 0$ there exists $p \in (0,1)$ such that

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

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 one from always allocating to an agent with the highest positive peer value. Indeed, recall from (2.1) that every n and every DIC mechanism q in the n'th environment, we have

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

In the n'th environment, given a type profile $\theta \in \Theta^n$ and $p \in (0,1)$, let $\bar{u}^n(p,\theta)$ denote the smallest peer value among the agents i whose rank $r_i(\theta)$ is below p; that is, $\bar{u}^n(p,\theta) = \min\{\bar{u}_i^n(\theta_{-i})\colon i \in [n] \wedge r_i(\theta) \leq p\}$. At every type profile θ, every agent who is allocated the good by $q^{n,p}$ has a peer value of at least $\bar{u}^n(p,\theta)$.

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 $q^{p,n}$:

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

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

Since no two agents have the same rank,[27]

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

[27] For a real number x, we write $\lfloor x \rfloor$ for the largest integer weakly smaller than x.

Since $(\delta^n)_{n\in\mathbb{N}}$ converges to 0 in probability, we have

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

as $n \to \infty$ for all fixed $p \in (0,1)$.

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

$$\lim_{n\to\infty} \sum_{\theta\in\Theta^n} \mu^n(\theta) \left(\max\left(0, \max_{i\in[n]} \bar{u}_i^n(\theta_{-i})\right) - \max\left(0, \bar{u}^n(p,\theta)\right)\right) < \varepsilon. \tag{A.3}$$

Let $\varepsilon > 0$. Regularity implies that there exists $p \in (0,1)$ such that

$$\lim_{k\to\infty} \mu^n \left\{\theta \in \Theta^n \colon \frac{1}{n}\left|\left\{j \in [n] \colon u_j^k(\theta_{-j}) + \varepsilon \geq \max_{i\in[n]} u_i^n(\theta_{-i})\right\}\right| \geq p\right\} = 1.$$

Given a profile θ, recall that $\bar{u}^n(p,\theta)$ denotes the smallest peer value among the agents i whose rank is below p. If θ is a type profile such that at least np agents have a peer value within ε of the highest peer value $\max_{i\in[n]} u_i^n(\theta_{-i})$, then $\bar{u}^n(p,\theta) \geq \max_{i\in[n]} u_i^n(\theta_{-i}) - \varepsilon$. Hence

$$\lim_{n\to\infty} \mu^n \left\{\theta \in \Theta^n \colon \bar{u}^n(p,\theta) \geq \max_{i\in[n]} u_i^n(\theta_{-i}) - \varepsilon\right\} = 1.$$

Since all peer values are in $[-1,1]$, for all $\varepsilon > 0$ there exists $p < 1$ such that (A.3) holds. □

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 Haghpanah, 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 Sadzik (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".

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 (2023). "Robust implementation with peer mechanisms and evidence". In: *Available at SSRN 2873888*.

Ben-Porath, Elchanan, Eddie Dekel, and B Lipman (2023). "Sequential mechanisms for evidence acquisition".

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.

Bhaskar, Dhruva and Evan Sadler (2019). "Resource Allocation with Positive Externalities". Available at SSRN 2853085.

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, Vašek (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.

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".

Fang, Ferric C. and Arturo Casadevall (2016). "Research Funding: the Case for a Modified Lottery". In: *mBio* 7.2, 10.1128/mbio.00422–16. DOI: 10.1128/mbio.00422-16. eprint: https://journals.asm.org/doi/pdf/10.1128/mbio.00422-16. URL: https://journals.asm.org/doi/abs/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 (2023). "The Multiple-Volunteers Principle for Assigning Unpleasant and Pleasant Tasks". URL: https://EconPapers.repec.org/RePEc:bon:boncrc:crctr224_2023_464.

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". URL: https://hal.archives-ouvertes.fr/hal-03187506.

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. Publisher: Wiley Online Library, pp. 173–196.

Huang, Yifei, Matt Shum, Xi Wu, and Jason Zezhong Xiao (2019). "Discovery of bias and strategic behavior in crowdsourced performance assessment". In: *arXiv preprint arXiv:1908.01718*.

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.

Kato, Miki and Shinji Ohseto (2004). "Non-Dummy Agents in Pure Exchange Economies". In: *The Japanese Economic Review* 55.2, pp. 212–220.

Kattwinkel, Deniz (2019). "Allocation with Correlated Information: Too good to be true". URL: https://sites.google.com/view/kattwinkel.

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". URL: https://arxiv.org/abs/2205.10910.

Kleiner, Andreas, Benny Moldovanu, and Philipp Strack (2021). "Extreme points and majorization: Economic applications". In: *Econometrica* 89.4, pp. 1557–1593.

Korte, Bernhard and Jens Vygen (2018). *Combinatorial Optimization: Theory and Algorithms*. 6th ed. 2018. Algorithms and Combinatorics; 21. Berlin, Heidelberg: Springer Berlin Heidelberg.

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.

Kovác, Eugen, Daniel Krähmer, and Tymon Tatur (2013). “Optimal stopping in a principal-agent model with hidden information and no monetary transfers”. In: *Unpublished manuscript*.

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*. arXiv: 2302.09168 [econ.TH]. URL: https://arxiv.org/abs/2302.09168.

Li, Yunan (2020). “Mechanism design with costly verification and limited punishments”. In: *Journal of Economic Theory* 186, p. 105000.

Li, Zihao and Jonathan Libgober (2023). “The Dynamics of Verification when Searching for Quality”. In: *Available at SSRN 4430184*.

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.

Manelli, Alejandro M and Daniel R Vincent (2007). “Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly”. In: *Journal of Economic theory* 137.1, pp. 153–185.

— (2010). “Bayesian and dominant-strategy implementation in the independent private-values model”. In: *Econometrica* 78.6, pp. 1905–1938.

McAfee, R. Preston and Philip J. Reny (1992). “Correlated information and mecanism 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 Zapechelnyuk (2017). “Optimal allocation with ex post verification and limited penalties”. In: *American Economic Review* 107.9, pp. 2666–94.

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. DOI: 10.3386/w29208. URL: http://www.nber.org/papers/w29208.

Östergård, Patric RJ 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 Urgun (2022). “Costly Verification and Money Burning”. In.

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”.

Prömel, Hans Jürgen (2013). *Ramsey theory for discrete structures*. Vol. 3. Springer.

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”.

Shah, Nihar B (2022). “Challenges, experiments, and computational solutions in peer review”. In: *Communications of the ACM* 65.6, pp. 76–87.

B Supplementary Material (Online Appendix)

B.1 Examples

Example showing the feasibility graph. Figure 2 shows the feasibility graph in an example with three agents, where agents 1 and 2 each have two possible types, and agent 3 has three possible types. The left panel shows the set of type profiles, here depicted as a $2 \times 2 \times 3$ grid. Agent 1's type varies horizontally, agent 2's type vertically, and agent 3's type diagonally. The right panel shows the feasibility graph. To understand the connection between the panels, think of a vertex (i, θ_{-i}) of the feasibility graph as a set of type profiles along which the types of agents other than i are fixed at θ_{-i} while i's type varies across all of Θ_i. In the left panel, for example, the vertex $(1, \theta_{-1})$ is thus depicted as the line connecting the profiles $\theta = (\theta_1, \theta_2, \theta_3)$ and $(\theta_1', \theta_2, \theta_3)$. Two vertices are adjacent if and only if the corresponding sets of type profiles intersect; for example, the vertices $(1, \theta_{-1})$, $(2, \theta_{-2})$ and $(3, \theta_{-3})$ are all adjacent as the corresponding sets of type profiles intersect at θ. In the right panel, vertices are depicted as circles (instead of lines) and lines indicate adjacencies.

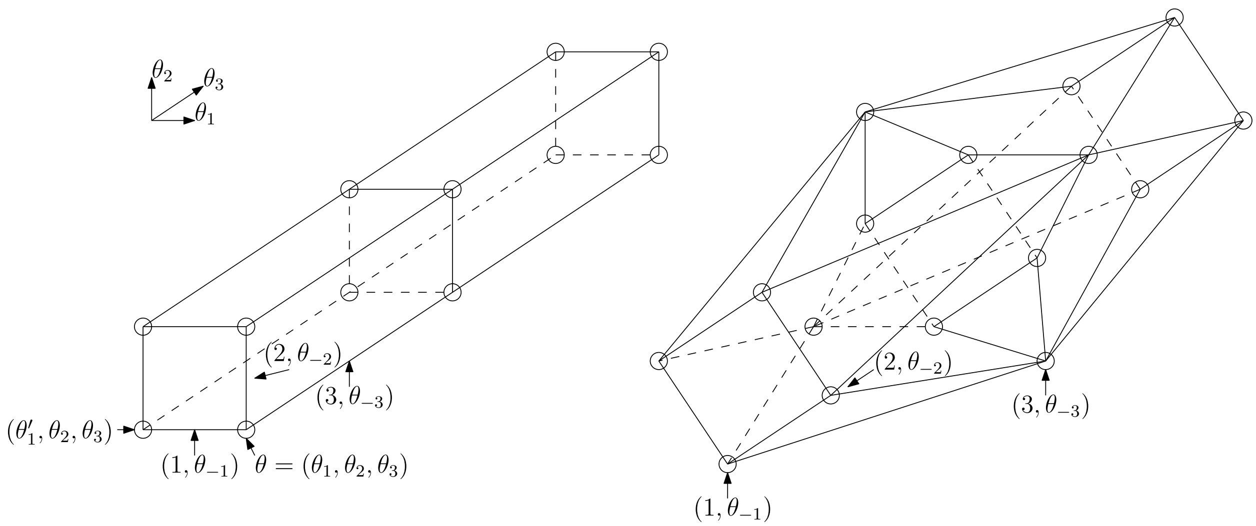

Figure 2: The set of type profiles (left) and the feasibility graph (right) in an example with $n = 3$, $|\Theta_1| = |\Theta_2| = 2$ and $|\Theta_3| = 3$.

Example showing a stochastic extreme DIC mechanism. Continuing with the example from Appendix B.1, suppose there are three agents, where agents 1 and 2 each have two possible types, and agent 3 has three possible types. The left panel of Figure 3 shows the space of type profiles. Agent 1's type varies horizontally, agent 2's type vertically, and agent 3's type diagonally.

The right panel of Figure 3 shows an odd hole $v_1, \dots, v_7$ in the feasibility graph. Let us say that a profile θ *contains* the vertices $\{(i, \theta_{-i}) \colon i \in [n]\}$. For each $\ell \in \{1, \dots, 7\}$, the type profile θ^ℓ indicated in the left panel of the figure contains the vertices $v_{\ell-1}$ and v_ℓ (where $v_0 = v_7$ is understood). For example, type profile θ^1 contains v_7 and v_1, and θ^2 contains v_1 and v_2, and so on.

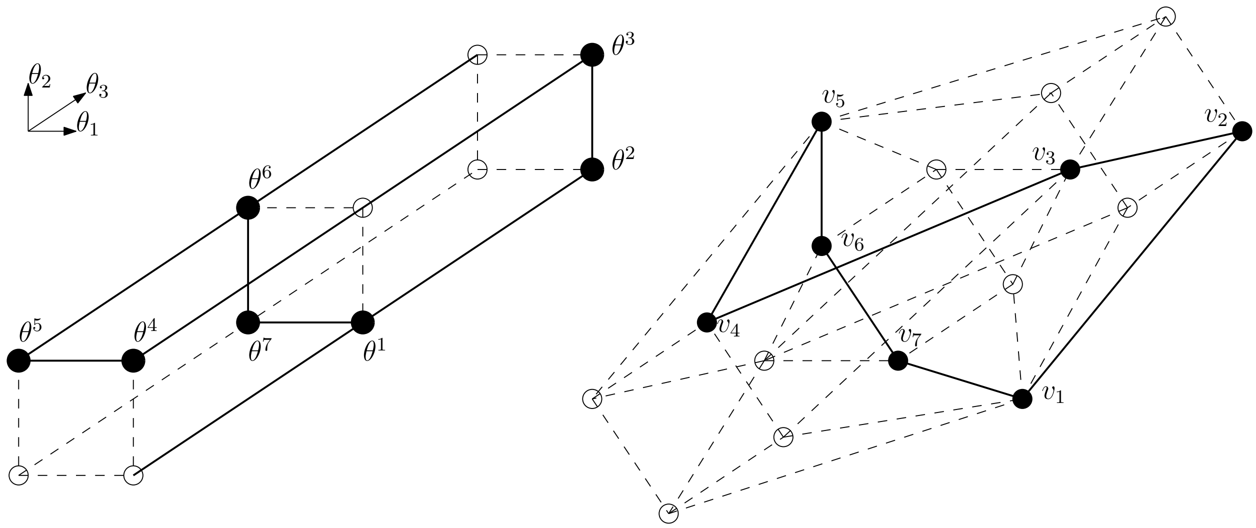

Figure 3: The right panel shows an odd hole $v_1, \dots, v_7$ in the feasibility graph. The left panel shows the seven type profiles $\theta^1, \dots, \theta^7$ that each contain two vertices of the hole.

Our candidate stochastic extreme DIC mechanism q assigns $\frac{1}{2}$ to all vertices $v_1, \dots, v_7$ of the odd hole, and 0 to all other vertices. Translated to the left panel, this means that, for each $\ell \in \{1, \dots, 7\}$, at θ^ℓ the mechanism flips a coin between the two agents who have vertices from the hole which are contained in θ^ℓ. For example, at θ^1 the mechanism flips a coin between agents 1 and 3, at θ^2 between agents 3 and 2, and so on.

In the main text, we intuited that randomization helps resolve the trade-off between allocating to an agent and using the agent's information since the agent can simultaneously win with some probability but nevertheless influence how the remaining probability is distributed among the other agents. We see this in the stochastic mechanism q. For example, agent 3 wins with probability 1/2 at both θ^1 and θ^2. These two type profiles differ only in agent 3's type. Depending on agent 3's report, the remaining probability 1/2 is assigned to agent 1 (at θ^1) or to agent 2 (at θ^2). In this sense, the principal enjoys both the benefit from allocating to agent 3 and agent 3's information.

To sharpen the intuition, suppose the weights $w_i(\theta_{-i}) = \mu(\theta_{-i}) u_i(\theta_{-i})$ are such that $w_i(\theta_{-i})$ equals 1 if the vertex (i, θ_{-i}) is in the odd hole $\{v_1, \dots, v_7\}$, and equals -1 otherwise. In particular, at θ^1 the principal wishes to select agent 1 or 3, while at θ^2 the principal wishes to select agent 2 or 3. Thus, here it is valuable to allocate to agent 3, and agent 3's type also contains valuable information.

Let us show that for these weights the stochastic DIC mechanism is in fact uniquely optimal (and therefore extreme). The principal's expected utility from q equals 7/2. Now let q' be an optimal DIC mechanism. Thus $\sum_v w(v)q'(v) \geq 7/2$. We show $q = q'$. For vertices outside $\{v_1, \ldots, v_7\}$, the weight w is strictly negative. Since decreasing the allocation probability at a vertex only slackens feasibility, we must have $q'(v) = 0$ for all v outside $\{v_1, \ldots, v_7\}$. Thus $\sum_v w(v)q'(v) = \sum_{\ell=1}^{7} q'(v_\ell)$ and it suffices to show $q'(v_\ell) = 1/2$ for all $\ell \in [7]$. For all $\ell \in [7]$, feasibility requires $q'(v_{\ell-1}) + q'(v_\ell) \leq 1$ since v_ℓ and $v_{\ell-1}$ are adjacent (where again $v_0 = v_7$). Thus $7/2 \leq q'(v_1) + \ldots + q'(v_7)$ requires $7/2 \leq q'(v_\ell) + 3$, and thus $1/2 \leq q'(v_\ell)$ for all $\ell \in [7]$. Using again $q'(v_{\ell-1}) + q'(v_\ell) \leq 1$, we conclude $q(v_\ell) = 1/2$ for all $\ell \in [7]$, as desired.

B.2 Approximate Optimality of Jury Mechanisms

Here, we discuss when jury mechanisms (Section 3) are approximately optimal with many agents. Recall that a jury mechanism splits the agents into jurors and candidates, using only the reports of the jurors to select one of the candidates. These mechanisms resemble peer review procedures used in practice. Therefore, it is interesting to understand when jury mechanisms perform well and how their performance compares to ranking-based mechanisms.

In Appendix B.2.1, we show that vanishing informational size alone does not imply that jury mechanisms are approximately optimal with many agents. Then, in Appendix B.2.2, we formalize two notions of exchangeability under which jury mechanisms are approximately optimal with many agents.

B.2.1 Vanishing Informational Size is Insufficient for Approximate Optimality of Jury Mechanisms

In this section, we show that jury mechanisms need not be approximately optimal for vanishing informational size. The building block is an environment with three agents that we later expand to have many agents. The key properties of the environment are as follows:

(1) at each type profile in the support of the type distribution, exactly two agents have a peer value of 1, while the other agent has a peer value of 0. Further, an agent's peer value is 0 if the reports of the others is outside the support.

(2) The information of an individual agent correctly identifies only with probability 2/3 which of the other two agents has a peer value of 1 (possibly both).

(3) Each individual agent has a peer value of 1 at only 2/3 of all type profiles in the support of the type distribution.

In such an environment with three agents, the argument for why ranking-based mechanisms outperform jury mechanisms is as follows. Informational size is at most 1/3 at every profile since agents influence the ranks only through the tie-breaking rule. Property 1 implies that the ranking-based mechanism with threshold $p = 2/3$ always selects an agent with peer value 1; indeed, an agent with peer value 1 has rank 2/3 or less regardless of what this agent reports, while an agent with peer value 0 has rank 1; thus at each profile the two agents with peer value 1 both enjoy a winning probability $1/(np) = 1/2$. However, properties 2 and 3 imply that each jury mechanism obtains at most 2/3. Indeed, if there is a single juror, then property (2) implies the juror finds the best candidate only with probability 2/3. For fewer or more jurors, the jury mechanism is constant, and hence property (3) implies that the constantly chosen candidate is the best only with probability 2/3.

An explicit environment with the above three properties is as follows. For each of the three agents, the type space is $\{1, 2, 3\}$. Consider the following odd 9-hole H in the feasibility graph:

$$(\cdot, 1, 1), (2, \cdot, 1), (2, 2, \cdot), (\cdot, 2, 2), (3, \cdot, 2), (3, 3, \cdot), (\cdot, 3, 3), (1, \cdot, 3), (1, 1, \cdot).$$

Let T denote the set of type profiles traversed by this odd hole; that is,

$$T = \{(1, 1, 1), (2, 1, 1), (2, 2, 1), (2, 2, 2), (3, 2, 2), (3, 3, 2), (3, 3, 3), (1, 3, 3), (1, 1, 3)\}.$$

The distribution of type profiles is uniform over T (i.e., $\mu(\theta) = \frac{1}{9}\mathbf{1}(\theta \in T)$ for all $\theta \in \Theta$) and peer values are given by $\bar{u}_i(\theta_{-i}) = \mathbf{1}((i, \theta_{-i}) \in H)$. Properties (1) and (3) are immediate from the construction. For property (2), it suffices to note that for every distinct agents i and j and type θ_i, agent j has a peer value of 1 at no more than two of the three type profiles where i's type is θ_i; for example, for $i = 3$ and $\theta_3 = 1$, agent 1 has a peer value of 1 only at $(1, 1, 1)$ and $(2, 1, 1)$ (but not at $(2, 2, 1)$, while agent 2 has a peer value of 1 only at $(2, 1, 1)$ and $(2, 2, 1)$ (but not at $(1, 1, 1)$.

We now expand the above example to have many agents and so that informational size vanishes as the number of agents diverges and regularity holds. Let the number of agents n be an integer multiple ℓ of 3. Partition the agents into ℓ groups of 3 so that within each group the agents are labeled with three consecutive integers (the labeling only matters for the tie-breaking rule). For each group, fix an auxiliary environment for three agents as above. Define the environment (for n agents) by letting all variables from the auxiliary environments be independent across groups. (Thus, for example, Nature takes ℓ independent uniform draws from the set T defined above.) Informational size is at most $2/n$ at every type profile in the support since each agent cannot influence the peer values of the agents belonging to other groups, and since within each group each agent influences the rank only through the tie-breaking rule.[28] Further, as $n \to \infty$ the sequence of environments is regular since two thirds of the agents all have the highest peer value of 1. As before, a ranking-based mechanism (with, say, threshold 2/3) always selects an agent with the highest peer value. Now consider a jury mechanism. Fix a realization of the types of the jurors. For each group ℓ, the jurors find a candidate from group ℓ with peer value 1 only with probability 2/3; this is because the types of jurors outside of group ℓ are uninformative about candidate in group ℓ, and because the jurors from group ℓ only find such a candidate with probability 2/3. There is only one good to allocate, and thus regardless from which group the principal selects a candidate using the types of the jurors, the principal finds a candidate with peer value 1 only with probability 2/3. Thus, the expected utility of each jury mechanism is at most 2/3.

B.2.2 Sufficient Conditions for Approximate Optimality of Jury Mechanisms

In this part of the appendix, we formalize two notions of exchangeability under which jury mechanisms are approximately optimal with many agents. We focus on environments in which agents receive conditionally independent signals about the values of each other.

An *information structure with conditionally independent signals* specifies the following for every $i \in \mathbb{N}$:

(1) a distribution f_i whose support $\operatorname{supp} f_i$ is finite and contained in $[-1, 1]$;

(2) for every $j \in \mathbb{N}$ distinct from i a finite set S_i^j and a function g_i^j that assigns to each $u_i \in \operatorname{supp} f_i$ a distribution over S_i^j.

Given an information structure with conditionally independent signals and $n \in \mathbb{N}$, the environment with n agents works as follows. The values $u_1, \dots, u_n$ are independent across agents 1 to n with respective marginals $f_1, \dots, f_n$. Conditional on u_i, each agent j other than i observes a signal σ_i^j about u_i drawn from $g_i^j(\cdot|u_i)$. Thus, each agent j's type is $(\sigma_i^j)_{i \in [n]\colon i \neq j}$. Conditional on $(u_1, \dots, u_n)$, the signals $(\sigma_i^j)_{i \neq j}$ are independent. Thus, the probability of a profile $((u_i)_{i \in [n]}, (\sigma_i^j)_{i,j \in [n]\colon i \neq j})$ is $\prod_{i,j \in [n]\colon i \neq j} f_i(u_i) g_i^j(\sigma_i^j|u_i)$.

Nothing in the analysis would change if an agent i also observed their own value since this information cannot be used to determine agent i's allocation (by DIC) and since agent i's value is uninformative about the other agents' values (by independence of the values).

Fix an information structure with conditionally independent signals.

- *Agents are exchangeable as suppliers of information* if for all $i \in \mathbb{N}$ the set of possible signals S_i^j and the conditional distributions $(g_i^j(\cdot|u_i))_{u_i \in \operatorname{supp} f_i}$ do not depend on j.

[28] For this argument, the partition of the agents into groups should be chosen so that the agents within a group are labeled with three consecutive integers. Since ties are broken in favor of agent i versus agent j if and only if $i < j$, it follows that each agent can influence their own rank by at most $2/n$.

Intuitively, for all i and j, the information that j can potentially provide about i does not depend systematically on j's identity.

- *Agents are exchangeable as recipients of the good* if the marginal distributions f_i do not depend on i and for all $j \in \mathbb{N}$ the set of possible signals S_i^j and all conditional distributions $(g_i^j(\cdot|u))_{u \in \operatorname{supp} f_i}$ do not depend on i.
  Intuitively, for all i and j, neither the value of allocating to i nor the information that j can potentially provide about i depend systematically on i's identity.

Theorem B.1. *Consider an information structure with conditionally independent signals in which agents are exchangeable suppliers of information or in which agents are exchangeable as recipients of the good. As the number of agent diverges, the difference between the principal's expected utility from an optimal jury mechanism and an optimal DIC mechanism vanishes.*

We provide a proof sketch that is straightforward but notationally tedious to verify. For this proof, it is crucial that the information structure is fixed. (By contrast, our result for ranking-based mechanisms, Theorem 7.2, does not require a fixed information structure.)

Suppose agents are exchangeable as suppliers of information. Let $n \in \mathbb{N}$. In the environment with n agents, let $\bar{U}^n$ denote the expected utility from always allocating to an agent with the highest peer value. Each DIC mechanism in the environment with n agents obtains at most $\bar{U}^n$. We show that $\bar{U}^n$ is attainable in the environment with $2n$ agents using a jury mechanism. The sequence $(\bar{U}^n)_{n\in\mathbb{N}}$ converges,[29] and hence the theorem follows. For all n, the value $\bar{U}^n$ is generated by the following mechanism q^n (which need not be DIC). The principal collects the entire profile of signals $(\sigma_i^j)_{i,j\in[n]\colon i\neq j}$, and then allocates to an agent i for which the expected value of u_i conditional on $(\sigma_i^j)_{j\in[n]\colon i\neq j}$ is maximal across $i \in [n]$. Now consider the following jury mechanism with $2n$ agents: agents 1 to n are the candidates, agents $n+1$ to $2n$ are the jurors. For each $j \in [n]$, the principal collects from juror $n+j$ only the signals $(\sigma_i^{n+j})_{i\in[n]\colon i\neq j}$. The principal treats the signals thus obtained in the same way that q^n would treat the signals of agent j. That is, the principal selects a candidate i with the highest expected value u_i conditional on $(\sigma_i^{n+j})_{i\in[n]\colon i\neq j}$. The only difference between this jury mechanism and q^n is that for each $i \in [n]$ the conditional expected value of agent i is computed using $n-1$ signals originating from agents $n+1$ to $2n$ instead of $n-1$ signals originating from $[n]\setminus\{i\}$. But agents are exchangeable as suppliers of information, meaning that the signals about u_i are conditionally iid. across *all* agents other than i. Thus, the origin of the signals is irrelevant for determining the conditional expected value of u_i. Thus, the described jury mechanism achieves the same expected utility as q^n.

[29]The sequence converges since it is bounded and increasing. Boundedness is immediate since values are in $[-1,1]$. Monotonicity follows since the information structure is held fixed as the number of agents increases. Thus, the principal has the option of ignoring agents who are added to the environment.

If agents are exchangeable as recipients of the good, a similar argument applies. We now construct a jury mechanism in which agents $n+1$ to $2n$ substitute for agents 1 to n whenever one of the latter is selected to receive the good by q^n. That is, agents 1 to n are the jurors, agents $n+1$ to $2n$ are the candidates. For all $j \in [n]$, the principal collects from juror j only the signals $(\sigma_i^j)_{i \in \{n+1,\dots,2n\}:\, i \neq n+j}$.[30] The principal treats the signal of $j \in \{1, \dots, n\}$ about $i \in \{n+1, \dots, n\}$ in the same way that q^n would treat the signal of j about $n-i$. Since agents are exchangeable as recipients, the conditional value of allocating to agent $i \in \{n+1, \dots, n\}$ given the elicited signal profile equals the conditional value of allocating to agent $n-i$ if the same profile of signals had been reported about about $n-i$. Thus, the described jury mechanism achieves the same expected utility as q^n.

B.3 A Strong Symmetry Condition Makes Peer Information Uninteresting

In this part of the appendix, we show that no DIC mechanisms that always allocates the good can meaningfully elicit information from the agents if the environment meets a strong symmetry condition.

In this part of the appendix, for every agent i, the value u_i has only finitely many possible realizations. We write $\mu(u_1, \dots, u_n, \theta_1, \dots, \theta_n)$ for the joint probability of a profile $(u_1, \dots, u_n)$ of values and a profile $(\theta_1, \dots, \theta_n)$ of types.

Let Ξ denote the set of permutations of $[n]$. The environment is *symmetric* if all type spaces agree ($\Theta_1 = \ldots = \Theta_n$) and all permutations $\xi \in \Xi$ satisfy

$$\mu(u_1, \dots, u_n, \theta_1, \dots, \theta_n) = \mu(u_{\xi(1)}, \dots, u_{\xi(n)}, \theta_{\xi(1)}, \dots, \theta_{\xi(n)}).$$

In the main text (Section 9), we alluded to the following special symmetric environment: each agent's type equals their value ($\theta_i = u_i$ with probability 1) and the distribution of values is invariant to permutations.

Theorem B.2. *In a symmetric environment, the principal's expected utility equals* $\mathbb{E}[u_1]$ *in every DIC mechanism that always allocates the good.*

Intuitively, symmetry is too strong for peer selection since it implies that even the information of $n-2$ agents is insufficient for distinguishing between the remaining two agents, i.e. $\mathbb{E}[u_i|\theta_{-ij}] = \mathbb{E}[u_j|\theta_{-ij}]$ for all distinct i and j and types θ_{-ij} of the others. (This intuition

[30] We exclude the signal if $i = n+j$ since q^n computes each agent's conditional expected value using $n-1$ signals. Thus, to replicate q^n, we delibaretely choose the jury mechanism to also use only $n-1$ signals for each conditional expected value.

is incomplete, though, since j's information could also be used to determine whether i wins but not whether j wins.)

The proof is based on the following stronger claim: if q is DIC, always allocates, and is such that, for all i, the allocation $q_i(\theta_{-i})$ is invariant with respect to permutations of θ_{-i}, then q is constant. This claim implies Theorem B.2 since, as we show, it is without loss to focus on such permutation-invariant mechanisms if the environment is symmetric.

Proof of Theorem B.2. For a moment, let us denote each agent's allocation as a function of the entire type profile. Say a DIC mechanism q is *symmetric* if all permutation $\xi \in \Xi$ satisfy

$$q_i(\theta_1, \ldots, \theta_n) = q_{\xi^{-1}(i)}(\theta_{\xi(1)}, \ldots, \theta_{\xi(n)}). \tag{B.1}$$

Here, $(\theta_{\xi(1)}, \ldots, \theta_{\xi(n)})$ is the profile where the type of an agent i is $\theta_{\xi(i)}$, which in turn is the type of agent $\xi(i)$ in the original profile $(\theta_1, \ldots, \theta_n)$. Symmetry says that agent i's winning probability at the permuted profile equals $\xi(i)$'s winning probability in the original profile.

We first verify the expected result that in a symmetric environment, given an arbitrary DIC mechanism q that always allocates, there is a symmetric DIC mechanisms that always allocates and that generates the same expected utility as q. Indeed, consider the mechanism q' define for all $i \in [n]$ and $(\theta_1, \ldots, \theta_n) \in \Theta$ by

$$q'_i(\theta_1, \ldots, \theta_n) = \frac{1}{n!} \sum_{\xi \in \Xi} q_{\xi^{-1}(i)}(\theta_{\xi(1)}, \ldots, \theta_{\xi(n)}).$$

One may verify that q' is DIC and always allocates. By construction, q' is symmetric. We show that q and q' generate the same expected utility for the principal. It holds:

$$\begin{aligned} U(q) &= \sum_{u,\theta,i} \mu(u_1, \ldots, u_n, \theta_1, \ldots, \theta_n) q_i(\theta_1, \ldots, \theta_n) u_i \\ &= \sum_{\xi \in \Xi} \frac{1}{n!} \sum_{u,\theta,i} \mu(u_1, \ldots, u_n, \theta_1, \ldots, \theta_n) q_i(\theta_1, \ldots, \theta_n) u_i \\ &= \sum_{\xi \in \Xi} \frac{1}{n!} \sum_{u,\theta,i} \mu(u_{\xi(1)}, \ldots, u_{\xi(n)}, \theta_{\xi(1)}, \ldots, \theta_{\xi(n)}) q_i(\theta_{\xi(1)}, \ldots, \theta_{\xi(n)}) u_{\xi(i)}; \end{aligned}$$

here the first equality is by definition, the second is clear, and the third is a change of variables in the summation (since for each fixed permutation ξ the map $(u_1, \ldots, u_n, \theta_1, \ldots, \theta_n) \mapsto (u_{\xi(1)}, \ldots, u_{\xi(n)}, \theta_{\xi(1)}, \ldots, \theta_{\xi(n)})$ ranges over all profiles of values and types). Since the environment is symmetric, we get

$$U(q) = \sum_{\xi \in \Xi} \frac{1}{n!} \sum_{u,\theta,i} \mu(u_1, \ldots, u_n, \theta_1, \ldots, \theta_n) q_i(\theta_{\xi(1)}, \ldots, \theta_{\xi(n)}) u_{\xi(i)}.$$

By another change of variables, for each fixed ξ, u, θ, we have

$$\sum_i q_i(\theta_{\xi(1)}, \ldots, \theta_{\xi(n)}) u_{\xi(i)} = \sum_i q_{\xi^{-1}(i)}(\theta_{\xi(1)}, \ldots, \theta_{\xi(n)}) u_i.$$

Thus

$$U(q) = \sum_{\xi \in \Xi} \frac{1}{n!} \sum_{u,\theta,i} \mu(u_1, \ldots, u_n, \theta_1, \ldots, \theta_n) q_{\xi^{-1}(i)}(\theta_{\xi(1)}, \ldots, \theta_{\xi(n)}) u_i.$$

The right side is simply the expected utility from q'.

To complete the proof, it suffices to show that every symmetric DIC mechansims that always allocates yields an expected utility of $\mathbb{E}[u_1]$. We show that every symmetric DIC mechanism q that always allocates is actually constant. The claim then follows since, by symmetry of the environment, every constant mechanism yields $\mathbb{E}[u_1]$.

For each agent i, we again drop i's report from i's allocation. Symmetry of q implies that $q_i(\theta_{-i})$ is invariant to permutations of θ_{-i} (consider (B.1) for permutations ξ such $\xi(i) = i$).

For this proof, we introduce some special notation. Recall that the agents share the common type space Θ_1. A generic type in Θ_1 is denoted t. A generic type profile of $n-1$ types is denoted $\boldsymbol{t}$. The space T of such $\boldsymbol{t}$ equals the $(n-1)$-fold product of Θ_1. Thus, agent i's allocation when the others report $\boldsymbol{t}$ equals $q_i(\boldsymbol{t})$, and this allocation is invariant to permutations of t.

Given $\boldsymbol{t} = (t_1, \ldots, t_{n-1}) \in T$ and $j \in [n-1]$, we denote by $\boldsymbol{t}_{-j}$ the profile of types in $\boldsymbol{t}$ other than t_j; that is, $\boldsymbol{t}_{-j} = (t_1, \ldots, t_{j-1}, t_{j+1}, \ldots, t_{n-1})$. For a type $t' \in \Theta_1$, we then denote by $(t', \boldsymbol{t}_{-j})$ the profile of $n-1$ types obtained from $\boldsymbol{t}$ by replacing t_j by t'; that is, $(t', \boldsymbol{t}_{-j}) = (t_1, \ldots, t_{j-1}, t', t_{j+1}, \ldots, t_{n-1})$.

For the proof, it shall be useful to consider the behavior of an agent's winning probability at profiles of the form $(t', \boldsymbol{t}_{-j})$ as we vary j. We establish the following auxiliary lemma.

Lemma B.3. *Let $i \in [n]$. All $t', t'' \in \Theta_1$ and $\boldsymbol{t} \in T$ satisfy*

$$\sum_{j=1}^{n-1} (q_i(t', \boldsymbol{t}_{-j}) - q_i(t'', \boldsymbol{t}_{-j})) = 0.$$

Proof of Lemma B.3. For notational simplicity, we only prove the claim for $i = n$ (later commenting on how to adapt the notation for the general case). Denote $\boldsymbol{t} = (t_1, \dots, t_{n-1})$.

In an intermediate step, let j be an agent distinct from agent n. We show

$$q_n(t'', \boldsymbol{t}_{-j}) - q_n(t', \boldsymbol{t}_{-j}) = q_j(t'', \boldsymbol{t}_{-j}) - q_j(t', \boldsymbol{t}_{-j}). \tag{B.2}$$

Note that $\boldsymbol{t}_{-j}$ is a profile of $n-2$ types. Suppose the $n-2$ agents other than i and j report $\boldsymbol{t}_{-j} = (t_1, \dots, t_{j-1}, t_{j+1}, \dots, t_{n-1})$. Let θ' be the profile where agent n reports t', agent j reports t'', and agents other than n and j report $\boldsymbol{t}_{-j}$. Let θ'' be the profile obtained from θ' by permuting n's and j's reports. Since θ' and θ'' differ only in a permutation of n's and j's reports, the allocation of agents other than n and j is the same across θ' and θ''. Since the object is always allocated, the sum of n's and j's allocation is also the same across θ' and θ''. Thus:

$$q_n(t'', \boldsymbol{t}_{-j}) + q_j(t', \boldsymbol{t}_{-j}) = q_n(t', \boldsymbol{t}_{-j}) + q_j(t'', \boldsymbol{t}_{-j});$$

the left side is the sum of n's and j's allocation at θ', the right side the sum at θ''. Rearranging yields (B.2).

Next, we sum (B.2) across all $j \in [n]$ distinct from i. Recall again that $\boldsymbol{t} = (t_1, \dots, t_{n-1})$ is a profile of $n-1$ types. Consider the profile where agents 1 to $n-1$ report $\boldsymbol{t}$, and agent n reports t'. At this profile, for all j distinct from, agent j's allocation equals $q_j(t', \boldsymbol{t}_{-j})$, while n's allocation is $q_n(\boldsymbol{t})$. Since the object is allocated, we have $\sum_{j \neq i} q_j(t', \boldsymbol{t}_{-j}) = 1 - q_n(\boldsymbol{t})$. By a similar argument, we have $\sum_{j \neq i} q_j(t'', \boldsymbol{t}_{-j}) = 1 - q_n(\boldsymbol{t})$. Thus summing (B.2) across $j \in [n] \setminus \{i\}$ yields $\sum_{j=1}^{n-1} (q_n(t', \boldsymbol{t}_{-j}) - q_n(t'', \boldsymbol{t}_{-j})) = 0$, as desired.

For the general case (where i not need equal n), the argument is analogous. One should now think of $\boldsymbol{t} = (t_1, \dots, t_{n-1})$ as the reports of all agents other than i (for some fixed assignment of these agents to $t_1, \dots, t_{n-1}$). The profile $\boldsymbol{t}_{-j}$ should be thought of as the profile $\boldsymbol{t}$ except j's type. □

We next use Lemma B.3 to complete the proof. Let $i \in [n]$ be arbitary. We show i's winning probability is constant in the reports of others. To that end, let us fix an arbitrary type $t^* \in T$. For all $k \in \{0, \dots, n-1\}$, let T_k denote the subset of profiles in T where exactly k-many entries are distinct from t^*. Let p_i denote i's winning probability when all other agents report t^*. We will show via induction over k that i's winning probability is equal to p_i whenever the others report a profile in T_k. This completes the proof since $T = \cup_{k=0}^{n-1} T_k$.

Base case $k = 0$. Immediate from the definitions of p_i and T_0.

Induction step. Let $k \geq 1$. Suppose all $\hat{\boldsymbol{t}} \in \cup_{\ell=0}^{k-1} T_\ell$ satisfy $q_i(\hat{\boldsymbol{t}}) = p_i$. Letting $\boldsymbol{t} \in T_k$ be arbitrary, we show $q_i(\boldsymbol{t}) = p_i$.

Denote $\boldsymbol{t} = (t_1, \ldots, t_n)$. Since $q_i(\boldsymbol{t})$ is invariant to permutations of $\boldsymbol{t}$, we may assume that exactly the first k entries of $\boldsymbol{t}$ are distinct from t^*, while all other entries equal t^*. That is, $\boldsymbol{t} = (t_1, \ldots, t_k, t^*, \ldots, t^*)$.

Let $\tilde{\boldsymbol{t}} = (t_1, \ldots, t_{k-1}, t^*, \ldots, t^*)$ be the profile obtained from $\boldsymbol{t}$ by replacing t_k by t^*. We now invoke Lemma B.3 to infer

$$\sum_{j=1}^{n-1} q_i(t_k, \tilde{\boldsymbol{t}}_{-j}) = \sum_{j=1}^{n-1} q_i(t^*, \tilde{\boldsymbol{t}}_{-j}). \tag{B.3}$$

Consider the profiles appearing in the sum on the left of (B.3) as j varies from 1 to $n-1$.

(1) Let $j \le k-1$. Since exactly the first $k-1$ entries of $\tilde{\boldsymbol{t}}$ are distinct from t^*, it follows that $(t_k, \tilde{\boldsymbol{t}}_{-j})$ is another profile where exactly $k-1$ entries differ from t^*. Hence the induction hypothesis implies $q_i(t_k, \tilde{\boldsymbol{t}}_{-j}) = p_i$.

(2) Let $j > k-1$. In the profile $(t_k, \tilde{\boldsymbol{t}}_{-j})$, the first $k-1$ entries are $t_1, \ldots, t_{k-1}$, the j'th entry is t_k, and all remaining entries are t^*. Hence $(t_k, \tilde{\boldsymbol{t}}_{-j})$ is a permutation of $\boldsymbol{t}$. Thus $q_i(t_k, \tilde{\boldsymbol{t}}_{-j}) = q_i(\boldsymbol{t})$.

Hence the sum on the left-hand side of (B.3) is given by $\sum_{j=1}^{n-1} q_i(t, \tilde{\boldsymbol{t}}_{-j}) = (k-1)p_i + (n-k)q_i(\boldsymbol{t})$

Now consider the sum on the right-hand side of (B.3). For all j, the profile $(t^*, \tilde{\boldsymbol{t}}_{-j})$ contains at most $(k-1)$-many entries different from t^* since $\boldsymbol{t}$ contains k entries different from t^*. By the induction hypothesis, therefore, the sum on the right-hand side of (B.3) equals $(n-1)p_i$.

Equation (B.3) thus simplifies to $(k-1)p_i + (n-k)q_i(\boldsymbol{t}) = (n-1)p_i$. Equivalently, $(n-k)(q_i(\boldsymbol{t}) - p_i) = 0$. Since $k \le n-1$, we conclude $q_i(\boldsymbol{t}) = p_i$, as promised. □

B.4 Mandatory Allocation

B.4.1 Preliminaries

Let $\bar{Q}$ denote the set of DIC mechanisms q that always allocate, i.e. $\sum_{i \in [n]} q_i(\theta_{-i}) = 1$ holds for all $\theta \in \Theta$. Let $\operatorname{ext} \bar{Q}$ denote the set of extreme points of $\bar{Q}$.

It holds $\operatorname{ext} \bar{Q} \subseteq \operatorname{ext} Q$ since a DIC mechanism that always allocates can only be represented as a convex combination of other DIC mechanisms that always allocate.

As before, let G denote the feasibility graph, and let $V = \cup_{i \in [n]} (\{i\}) \times \Theta_{-i}$ denote its vertices. Let $\bar{\mathcal{S}}$ denote the set of stable sets S of G such that $|S \cap X| = 1$ holds for all maximal cliques X. The map $q \mapsto \{v \in V : q(v) = 1\}$ is a bijection from $\bar{Q}$ to $\bar{\mathcal{S}}$.

B.4.2 Results

The next theorem characterizes when all extreme points of $\bar{Q}$ are deterministic. The only difference to the characterization for Q (Theorem 5.1) is in the threshold value for the number of agents.

Theorem B.4. *All extreme points of $\bar{Q}$ are deterministic if and only if at least one of the following is true:*

(1) there are at most three agents ($n \leq 3$);

(2) all type spaces are binary ($|\Theta_i| \leq 2$ for all $i \in [n]$).

(Recall that the model assumes $n \geq 2$ and $|\Theta_i| \geq 2$ for all $i \in [n]$.

Proof of Theorem B.4. The proof of Theorem 3.1 shows that all extreme points are deterministic if $n \leq 3$. If all type spaces are binary, the claim follows from Theorem 5.1 and the inclusion $\operatorname{ext}\bar{Q} \subseteq \operatorname{ext} Q$. Finally, suppose there are at least four agents and at least one type space is non-binary. Without loss, let agent 1 have a non-binary type space. Theorem 5.1 implies there exists a stochastic extreme point of the set of DIC mechanisms (that need not always allocate) for agents 1, 2, and 3. View this mechanism as a DIC mechanism that always allocates for agents 1 to n where all reports of agents 4 to n are ignored. This mechanisms is an extreme point of $\bar{Q}$. □

Remark B.5. In Lemma A.3, we showed that Q admits a stochastic extreme point if and only if G admits an odd hole of length 7 or greater. Only one direction of this equivalence carries over to $\bar{Q}$. Namely, if $\bar{Q}$ admits a stochastic extreme point, then G admits an odd hole of length 7 or greater. However, if $n = 3$, then all extreme points of $\bar{Q}$ are deterministic even though G admits an odd hole of length 7 or greater if at least one type space is non-binary.

The next theorem provides an of analogue Theorem 5.2: essentially all extreme points of $\bar{Q}$ are stochastic if $n \geq 4$ and type spaces are large. For the sake of simplicity, we focus on the case where the agents' type spaces all have the same cardinality.

Theorem B.6. *Fix $n \geq 4$. Suppose the agents have a common type space ($\Theta_1 = \ldots = \Theta_n$) which has cardinality $m \in \mathbb{N}$. For all $\varepsilon > 0$ there exists $m_\varepsilon \in \mathbb{N}$ such that if $m \geq m_\varepsilon$, then* $\left|\det \bar{Q}\right| < \varepsilon \left|\operatorname{ext} \bar{Q}\right|$.

The proof, presented further below, is more challenging than the one for Theorem 5.2.

Next, we provide a characterization of stochastic extreme points via odd holes, analogously to Theorem 5.3. As in the main text, for a mechanism $q \in \bar{Q}$ a *stochastic component of q* is an inclusion-wise maximal connected set of vertices v of G such that $q(v) \in (0, 1)$.

Theorem B.7. *Let q be a stochastic DIC mechanism that always allocates. If q is an extreme point of $\bar{Q}$, then every stochastic component of q contains an odd hole. The converse holds if $q(v) \in \{0, \frac{1}{2}, 1\}$ for all $v \in V$.*

Proof of Theorem B.7. First, let $q \in \operatorname{ext} \bar{Q}$. Since $\operatorname{ext} \bar{Q} \subseteq \operatorname{ext} Q$, we can invoke Theorem 5.3 to infer that each stochastic component of q admits an odd hole.

Second, suppose q only takes values in $\{0, \frac{1}{2}, 1\}$ and that each of its components admits an odd hole. Since $q \in \bar{Q}$ and $\bar{Q} \subseteq Q$, Theorem 5.3 implies q is an extreme point of Q. Using again $q \in \bar{Q}$ and $\bar{Q} \subseteq Q$, it follows that q is also an extreme point of $\bar{Q}$. □

Next, we show that the problem of determining an optimal deterministic DIC mechanisms that always allocates is NP-hard if $n \geq 4$. This result follows immediately from Theorem 6.1, but we include the definitions for the sake of completeness.

Definition 8 (OptDetMA-n)**.** For $n \in \mathbb{N}$, let OptDetMA-n be the following optimization problem ("MA" stands for mandatory allocation). The input consists of finite sets $\Theta_1, \ldots, \Theta_n$ and weights $w_i : \Theta_{-i} \to \mathbb{Z}$ for all $i \in [n]$. The problem is to find a deterministic DIC mechanism q (for n agents with respective type spaces $\Theta_1, \ldots, \Theta_n$) that always allocates and that maximizes $\sum_{i,\theta} w_i(\theta_{-i}) q_i(\theta_{-i})$ across all deterministic DIC mechanisms that always allocate.

Theorem B.8. *If $n \geq 4$, then* OptDetMA-n *is NP-hard.*

Proof of Theorem B.8. Consider the instances of OptDetMA-n where at least one agent has a singleton type space and a weight constantly equal to 0. Each such instance corresponds to an instance of OptDet-$(n-1)$. Since $n \geq 4$, Theorem 6.1 implies that OptDet-$(n-1)$ is NP-hard. Thus OptDetMA-n is also NP-hard. □

Remark B.9. If $n \leq 3$, then optimal deterministic must-allocate DIC mechanisms are computable in polynomial time. Indeed, according to Theorem 3.1 it suffices to restrict attention to jury mechanisms that always allocate. It is easy to see that the optimal juror can be found in polynomial time.

B.4.3 Proof of Theorem B.6

We introduce some useful terminology. For all $i \in [n]$, the set $V_i = \{i\} \times \Theta_{-i}$ is 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*. Given distinct i and j, two i-vertices and are *j-translates* if there is a j-vertex that is adjacent to both of them; equivalently, the i-vertices coincide exactly except for possibly in the type of agent j. To be sure, an i-vertex is a j-translate of itself. A j-vertex is adjacent to an i-vertex v only if it is adjacent to all j-translates of v.

Recall $|\bar{\mathcal{S}}| = |\det \bar{Q}|$. Recall also that the agents have a common type space with cardinality m. By possibly relabelling types, thus $\Theta_i = \{1, \ldots, m\}$ for all $i \in [n]$.

We show that for all integers $L \geq 4$ there exists $m_L \in \mathbb{N}$ such that if $m \geq m_L$, then $|\bar{\mathcal{S}}| \left(1 + \frac{L}{n}\right) \leq |\operatorname{ext} \bar{Q}|$, which proves Theorem B.6. In what follows, fix an integer $L \geq 4$.

We now define the central notion of the proof—*regularity*—which identifies odd holes and stable sets that intersect in a particular way. (This regularity notion has no connection to the regularity notion which we discussed in the context of ranking-based mechanisms.)

Definition 9 (Regularity)**.** Let $S \in \bar{\mathcal{S}}$. Let H be a set of vertices. Let $(v_1, \ldots, v_9)$ be an odd hole. Let $i, j, k, \ell \in [n]$ be distinct. The tuple $(S, H, v_1, \ldots, v_9, i, j, k, \ell)$ is *regular* if

(1) H is the set of ℓ-translates of $\{v_1, \ldots, v_9\}$, and $H \subseteq V_i \cup V_j \cup V_k$ holds.

(2) For all vertices ω, if ω is in S and adjacent to two distinct vertices in H, then ω is an ℓ-vertex.

(3) For all vertices ω, if ω is an ℓ-vertex that is contained in a maximal clique that also contains two vertices in H, then $\omega \in S$.

The triple $\boldsymbol{r} = (S, H, \ell)$ is *regular* if there exist $\{v_1, \ldots, v_9\}$ and $i, j, k \in [n]$ such that $(S, H, v_1, \ldots, v_9, i, \ldots, \ell)$ is regular.

If (S, H, ℓ) is regular, the choice $(v_1, \ldots, v_9, i, j, k)$ is unique up to the type of agent ℓ (which is held constant across $v_1, \ldots, v_9$), and the labelling of the vertices $v_1, \ldots, v_9$.

We next address the existence of regular tuples.

Lemma B.10. *There exists $m_L \in \mathbb{N}$ such that if $m \geq m_L$, then for all $S \in \bar{\mathcal{S}}$ there exists $\ell \in [n]$ and at least L distinct sets H such that (S, H, ℓ) is regular.*

Proof of Lemma B.10. We use the following special case of Theorem 9.2 of Prömel (2013).

Lemma B.11. *There exists $m_L \in \mathbb{N}$ such that, if $m \geq m_L$, then for all $q\colon \Theta \to [n]$ there exist $\Theta_1^* \subseteq \Theta_1, \ldots \Theta_n^* \subseteq \Theta_n$ such that q is constant on $\times_{i=1}^n \Theta_i^*$, and such that $\min\{|\Theta_1^*|, \ldots, |\Theta_n^*|\} \geq L$ holds.*

Recall that there is bijection between $\bar{Q}$ and $\bar{\mathcal{S}}$. Hence:

Corollary B.12. *There exists $m_L \in \mathbb{N}$ such that, if $m \geq m_L$, then for all $S \in \bar{\mathcal{S}}$ there exists $\ell \in [n]$ and $\Theta_1^* \subseteq \Theta_1, \ldots \Theta_n^* \subseteq \Theta_n$ such that $\min\{|\Theta_1^*|, \ldots, |\Theta_n^*|\} \geq L$ and such that all $\theta_{-\ell} \in \times_{i \neq \ell} \Theta_i^*$ satisfy $(\ell, \theta_{-\ell}) \in S$.*

Let m_L meet the conclusion of Corollary B.12, and let $m \geq m_L$. Let $S \in \bar{\mathcal{S}}$, and let $\Theta_1^* \times \ldots \times \Theta_n^*$ be as in the conclusion of Corollary B.12. Now let i, j, and k be three agents distinct from ℓ (recall $n \geq 4$). Let $\theta_- \in \times_{\iota \notin \{i,j,k,\ell\}} \Theta^*_\iota$ be an arbitrary profile of types of agents other than i, j, k, and ℓ (assuming such agents exist). Let $\theta_\ell \in \Theta_\ell$ be arbitrary.

We next describe the construction of a family of regular tuples. Since $L \geq 4$, for all $\iota \in \{i, j, k\}$, we may select three distinct types $\theta^1_\iota, \theta^2_\iota, \theta^3_\iota$ from Θ^*_ι. Now consider the following set of vertices (here, we denote, say, vertices of agent i by $(\cdot, \theta^1_j, \theta^1_k, \theta_\ell, \theta_-)$, etc.):

$$\begin{aligned}
v_1 &= (\cdot, \theta^1_j, \theta^1_k, \theta_\ell, \theta_-) \\
v_2 &= (\theta^2_i, \cdot, \theta^1_k, \theta_\ell, \theta_-) \\
v_3 &= (\theta^2_i, \theta^2_j, \cdot, \theta_\ell, \theta_-) \\
v_4 &= (\cdot, \theta^2_j, \theta^2_k, \theta_\ell, \theta_-) \\
v_5 &= (\theta^3_i, \cdot, \theta^2_k, \theta_\ell, \theta_-) \\
v_6 &= (\theta^3_i, \theta^3_j, \cdot, \theta_\ell, \theta_-) \\
v_7 &= (\cdot, \theta^3_j, \theta^3_k, \theta_\ell, \theta_-) \\
v_8 &= (\theta^1_i, \cdot, \theta^3_k, \theta_\ell, \theta_-) \\
v_9 &= (\theta^1_i, \theta^1_j, \cdot, \theta_\ell, \theta_-).
\end{aligned} \tag{B.4}$$

By inspection, $(v_1, \ldots, v_9)$ is an odd hole. Let H denote this odd hole and all its ℓ-translates. We show that the tuple $(S, H, v_1, \ldots, v_9, i, j, k, \ell)$ is regular for each selection of $\theta^1_i, \theta^2_i, \theta^3_i, \theta^1_j, \theta^2_j, \theta^3_j, \theta^1_k, \theta^2_k, \theta^3_k$. That is, we show

(1) H is the set of ℓ-translates of $\{v_1, \ldots, v_9\}$, and $H \subseteq V_i \cup V_j \cup V_k$ holds.

(2) For all vertices ω, if ω is in S and adjacent to two distinct vertices in H, then ω is an ℓ-vertex.

(3) For all vertices ω, if ω is an ℓ-vertex that is contained in a maximal clique that also contains two vertices in H, then $u \in S$.

Property (1) is immediate from the construction of H.

Consider property (2). Consider a vertex ω that is not an ℓ-vertex but that is adjacent to two distinct vertices v and v' in H. We show $u \notin S$. Assume v is an i-vertex (the other cases being similar). We first claim v' is not an i-vertex. Towards a contradiction, suppose v' is an i-vertex. Since v and v' are ℓ-translates of vertices in $v_1, \ldots, v_9$, there are two distinct i-vertices $\tilde{v}$ and $\tilde{v}'$ in $v_1, \ldots, v_9$ that are both adjacent to some ℓ-translate of ω. However, since $\tilde{v}$ and $\tilde{v}'$ are both i-vertices, inspection of (B.4) shows that $\tilde{v}$ and $\tilde{v}'$ must differ in at least two types (of agents other than i), implying that there is no vertex adjacent to both of them; contradiction. Thus v' is not an i-vertex. Assume v' is a j-vertex (the other case where v' is a k-vertex being similar). Now let ω be a $\tilde{k}$-vertex. We know $\tilde{k} \neq i$ (since ω would fail to be adjacent to v) and $\tilde{k} \neq j$ (since ω would fail to be adjacent to v'). Thus the types of all agents other than i and $\tilde{k}$ agree at ω and v, and the types of all other agents other than j and $\tilde{k}$ agree at ω and v'. In particular, the types of agents other than $\tilde{k}$ are in the respective spaces $\Theta_1^*, \ldots, \Theta_n^*$. We also have $\tilde{k} \neq \ell$ by assumption. It follows that ω is adjacent to an ℓ-vertex in which the agents other than ℓ all have types in the respective spaces $\Theta_1^*, \ldots, \Theta_n^*$; in particular, this ℓ-vertex is in S. Since S is stable, we conclude that ω is not in S.

Turning to property (3), consider an ℓ-vertex ω adjacent to two adjacent vertices in H. Since H includes all its ℓ-translates, ω is adjacent to two vertices in $\{v_1, \ldots, v_9\}$. Suppose these are the vertices v_1 and v_2; the arguments for the other cases are similar. The unique ℓ-vertex adjacent to v_1 and v_2 is the vertex$(\theta_i^2, \theta_j^1, \theta_k^1, \cdot, \theta_{-i})$. Given the choice of the sets $\Theta_1^*, \ldots, \Theta_n^*$, this ℓ-vertex is in S, as promised.

Lastly, by varying the initial choices of $(\theta_i^1, \theta_i^2, \theta_i^3)$, $(\theta_j^1, \theta_j^2, \theta_j^3)$, and $(\theta_k^1, \theta_k^2, \theta_k^3)$, respectively, from Θ_i^*, $\Theta_{j,m}^*$, and $\Theta_{k,m}^*$, respectively, and using that each of Θ_i^*, $\Theta_{j,m}^*$, and $\Theta_{k,m}^*$ contains at least L elements, we obtain L distinct sets H (in fact, we obtain more than L, but L is all we need for the argument). □

Let m_L meet the conclusion of Lemma B.10. As described earlier, we complete the proof by showing that if $m \geq m_L$, then

$$|\bar{\mathcal{S}}|\left(1 + \frac{L}{n}\right) \leq |\operatorname{ext} \bar{Q}|.$$

In what follows, let $m \geq m_L$.

For a later part of the argument, it shall be more convenient to work with regular triples (S, H, ℓ) where ℓ is fixed to be agent n. To that end, define $\bar{\mathcal{S}}^*$ as the set of $S \in \bar{\mathcal{S}}$ for which there exist at least L distinct sets H such that (S, H, n) is regular. If (S, H, ℓ) is some regular triple, then permuting the roles of agents ℓ and n gives another regular triple; this step uses that all agents have the same type space. Thus:

$$|\bar{\mathcal{S}}| \leq n|\bar{\mathcal{S}}^*|. \tag{B.5}$$

We make some more auxiliary definitions. Let $\boldsymbol{r} = (S, H, \ell)$ be regular.

(1) Let $N(H)$ denote the set of vertices that are adjacent to at least one vertex in H.
(2) Let $\Omega_{\boldsymbol{r}}$ denote the ℓ-vertices that are contained in a maximal clique that also contains two vertices in H; that is, the ℓ-vertices considered in point (3) of the definition of regularity. Regularity requires $\Omega_{\boldsymbol{r}} \subseteq S$.
(3) Let $V_\ell \setminus N(H)$ be the set of ℓ-vertices that are non-adjacent to all vertices in H.
(4) Let $V_{\boldsymbol{r}}^* = (S \setminus \Omega_{\boldsymbol{r}}) \cup H \cup (V_\ell \setminus N(H))$, and let $G[V_{\boldsymbol{r}}^*]$ be the subgraph induced by $V_{\boldsymbol{r}}^*$.

(5) Let $\mathcal{K}_{\boldsymbol{r}}$ denote the set of connected components of $G[V_{\boldsymbol{r}}^*]$. Define $K_{\boldsymbol{r}} = \cup\{K' \in \mathcal{K}_{\boldsymbol{r}} \colon K' \cap H \neq \emptyset\}$ as the set of vertices in those connected components that have a non-empty intersection with H.

We next describe a family of candidate stochastic extreme points.

Definition 10. For all regular $\boldsymbol{r} = (S, H, \ell)$, let $q_{\boldsymbol{r}} \colon V \to [0,1]$ be the function defined by

$$\forall_{v \in V}, \quad q_{\boldsymbol{r}} = \begin{cases} \frac{1}{2}, & \text{if } v \in K_{\boldsymbol{r}} \\ 1, & \text{if } v \in S \setminus (\Omega_{\boldsymbol{r}} \cup K_{\boldsymbol{r}}) \\ 0, & \text{else.} \end{cases}$$

Lemma B.13. *If $\boldsymbol{r} = (S, H, \ell)$ is regular, then $q_{\boldsymbol{r}}$ is feasible and a stochastic extreme point of $\bar{Q}$.*

Proof of Lemma B.13. To begin with, provided we can show that $q_{\boldsymbol{r}}$ is in $\bar{Q}$, it follows immediately from Theorem B.7 that $q_{\boldsymbol{r}}$ is a stochastic extreme point of $\bar{Q}$. Thus we show $q_{\boldsymbol{r}} \in \bar{Q}$.

Let X be a maximal clique. We have to show $q_{\boldsymbol{r}}(X) = 1$, where we denote $q_{\boldsymbol{r}}(X) = \sum_{v \in X} q(v)$.

We first observe:

(1) it holds $|X \cap H| \leq 2$. Indeed, suppose towards a contradiction $|X \cap H| \geq 3$. By definition of H, there is an odd hole $\tilde{H}$ such that H is exactly the set of ℓ-translates of $\tilde{H}$. Hence we infer from $|X \cap H| \geq 3$ that $\tilde{H}$ contains a clique of three vertices, contradicting that $\tilde{H}$ is a hole.

(2) $H \cap S = \emptyset$. Indeed, each vertex in H is adjacent to another vertex in H. Hence $H \cap S = \emptyset$ follows since each maximal clique contains an ℓ-vertex, the definition of $\Omega_{\boldsymbol{r}}$, and the inclusion $\Omega_{\boldsymbol{r}} \subseteq S$.

Since $S \in \bar{\mathcal{S}}$, there exists $v \in X \cap S$. To show $q_{\boldsymbol{r}}(X) = 1$, we proceeds in several steps.

First, suppose $v \in \Omega_{\boldsymbol{r}}$. By definition of $\Omega_{\boldsymbol{r}}$, there are distinct vertices $u, u' \in H$ adjacent to v. Since H contains all its ℓ-translates, we may assume ω and ω' are both in X. Further, we have $X \cap (S \setminus \{v\}) = \emptyset$ (since $v \in S$ and S is stable) and $X \cap (H \setminus \{u, u'\}) = \emptyset$ (since $|X \cap H| \leq 2$) and $X \cap (V_\ell \setminus N(H)) = \emptyset$ (since $X \cap H \neq \emptyset$). Hence $q_{\boldsymbol{r}}$ assigns 0 to all vertices in $X \setminus \{u, u'\}$. By definition of $K_{\boldsymbol{r}}$, we have $u, u' \in K_{\boldsymbol{r}}$, and hence $q_{\boldsymbol{r}}(X) = 1$.

In what follows, suppose $v \notin \Omega_{\boldsymbol{r}}$.

As an intermediate claim, we show that $X \cap (V_\ell \setminus N(H)) \neq \emptyset$ holds if and only if $X \cap H = \emptyset$ holds. Indeed, if $X \cap (V_\ell \setminus N(H)) \neq \emptyset$, then we definitionally have $X \cap H = \emptyset$. Conversely, if $X \cap H = \emptyset$, then the ℓ-vertex in X (each maximal clique contains a unique ℓ-vertex) is in $V_\ell \setminus N(H)$ since H is closed with respect to ℓ-translations.

The intermediate claim implies there is a vertex $\omega \in X \cap (H \cup (V_\ell \setminus N(H)))$.

As a second intermediate claim, we show

$$X \cap ((S \setminus \Omega_{\boldsymbol{r}}) \cup H \cup (V_\ell \setminus N(H))) = \{\omega, v\}.$$

Let ω' in X be a vertex distinct from ω and v. We know $\omega' \notin S$ (since $v \in S$ and S is stable). We know $\omega' \notin H$; for if $\omega' \in H$, then ω would also be in H (by the intermediate claim), implying that X contains two vertices in H; this contradicts $v \notin \Omega_{\boldsymbol{r}}$. Finally, $\omega' \notin V_\ell \setminus N(H)$; for if $\omega' \in V_\ell \setminus N(H)$, then ω would also be in $V_\ell \setminus N(H)$ (by the first intermediate claim), implying that X contains two distinct ℓ-vertices, which is impossible. Thus $X \cap ((S \setminus \Omega_{\boldsymbol{r}}) \cup H \cup (V_\ell \setminus N(H))) = \{\omega, v\}$.

We know show $q_{\boldsymbol{r}}(\omega) + q_{\boldsymbol{r}}(v) = 1$. This establishes $q_{\boldsymbol{r}}(X) = 1$ since

$$X \cap ((S \setminus \Omega_{\boldsymbol{r}}) \cup H \cup (V_\ell \setminus N(H))) = \{\omega, v\}.$$

If $\omega \neq v$, then since ω and v are adjacent and in $G[V^*_{\boldsymbol{r}}]$ it must be that ω and v are in the same connected component of $G[V^*_{\boldsymbol{r}}]$. (If $\omega = v$, then of course ω and v are also in the same connected component.) In particular, we have $\omega \in K_{\boldsymbol{r}}$ if and only if $v \in K_{\boldsymbol{r}}$. Therefore, if $\omega \neq v$ and $v \in K_{\boldsymbol{r}}$, then $q_{\boldsymbol{r}}(X) = q_{\boldsymbol{r}}(\omega) + q_{\boldsymbol{r}}(v) = \frac{1}{2}$; if $\omega \neq v$ and $v \notin K$, then $q_{\boldsymbol{r}}(X) = q_{\boldsymbol{r}}(v) = 1$. Lastly, consider the case $\omega = v$. From the choice of ω we get $v \in S \cap (H \cup (V_\ell \setminus N(H))$. We know $S \cap H = \emptyset$ holds, and hence $v \in V_\ell \setminus N(H)$. It follows that v is non-adjacent to all other vertices in S (by stability), non-adjacent to all other vertices in $(V_\ell \setminus N(H))$ (since all ℓ-vertices are non-adjacent), and non-adjacent to all other vertices in H (since $v \in V_\ell \setminus N(H)$). Hence $\{v\}$ is a connected component of $G[V^*_{\boldsymbol{r}}]$. In particular, v cannot be contained in a component that intersects H since such a component must contain more than one vertex (e.g., it contains an ℓ-translate of $v_1, \dots, v_9$). Hence $q_{\boldsymbol{r}}(X) = q_{\boldsymbol{r}}(v) = 1$. □

We next establish that the mapping $\boldsymbol{r} \mapsto q_{\boldsymbol{r}}$ is injective for regular $\boldsymbol{r}$ that use the same agent ℓ.

Lemma B.14. *Let $\boldsymbol{r} = (S, H, \ell)$ and $\boldsymbol{r}' = (S', H', \ell')$ be regular. If $q_{\boldsymbol{r}} = q_{\boldsymbol{r}'}$ and $\ell = \ell'$, then $\boldsymbol{r} = \boldsymbol{r}'$.*

Proof of Lemma B.14. We first prove $H = H'$. We prove $H \subseteq H'$, the other inclusion being analogous. Towards a contradiction, let $v \in H \setminus H'$. Recall that H consists exactly of all ℓ-translates of some odd hole, and that H contains no ℓ-vertices. In particular, there are vertices $\omega_1, \dots, \omega_9$ in H that form an odd hole and such that $v = \omega_1$ holds. Note $q_{\boldsymbol{r}} = q_{\boldsymbol{r}'}$

assigns $\frac{1}{2}$ to all vertices $\omega_1, \dots, \omega_9$. Hence $\omega_1, \dots, \omega_9$ are all in $S' \cup H' \cup (V'_\ell \setminus N(H'))$. Since $\omega_1, \dots, \omega_9$ are also all in H, and since $\ell = \ell'$ and $H \cap V_\ell = \emptyset$ hold, we infer that $\omega_1, \dots, \omega_9$ are all in $S' \cup H'$. Now $\omega_1 = v \notin H'$ requires $\omega_1 \in S'$, and hence $\omega_2, \omega_9 \in H'$. By regularity of $\boldsymbol{r}'$, we conclude that ω_1 is an ℓ'-vertex. In view of $\omega_1 = v \in H$ and $\ell = \ell'$, this gives a contradiction since H contains no ℓ-vertices.

Lastly, we prove $S = S'$. The assumption $q_{\boldsymbol{r}} = q_{\boldsymbol{r}'}$ implies $S \setminus (\Omega_{\boldsymbol{r}} \cup K_{\boldsymbol{r}}) = S' \setminus (\Omega_{\boldsymbol{r}'} \cup K_{\boldsymbol{r}'})$ and $K_{\boldsymbol{r}} = K_{\boldsymbol{r}'}$. Using $H = H'$, a moment's thought reveals $\Omega_{\boldsymbol{r}} = \Omega_{\boldsymbol{r}'}$. Recalling also the definitions of $K_{\boldsymbol{r}}$ and $K_{\boldsymbol{r}'}$, the equalities $H = H'$ and $K_{\boldsymbol{r}} = K_{\boldsymbol{r}'}$ and $\Omega_{\boldsymbol{r}} = \Omega_{\boldsymbol{r}'}$ also imply $(S \cap K_{\boldsymbol{r}}) \setminus \Omega_{\boldsymbol{r}} = (S' \cap K_{\boldsymbol{r}}) \setminus \Omega_{\boldsymbol{r}}$. Since $K_{\boldsymbol{r}}$ contains no vertex in $\Omega_{\boldsymbol{r}}$, we also have $(S \cap \Omega_{\boldsymbol{r}}) \setminus K_{\boldsymbol{r}} = S \cap \Omega_{\boldsymbol{r}}$ and $(S' \cap \Omega_{\boldsymbol{r}}) \setminus K_{\boldsymbol{r}} = S' \cap \Omega_{\boldsymbol{r}}$ and $S \cap \Omega_{\boldsymbol{r}} \cap K_{\boldsymbol{r}} = S' \cap \Omega_{\boldsymbol{r}} \cap K_{\boldsymbol{r}} = \emptyset$. Since also $S \setminus (\Omega_{\boldsymbol{r}} \cup K_{\boldsymbol{r}}) = S' \setminus (\Omega_{\boldsymbol{r}'} \cup K_{\boldsymbol{r}'})$ and $(S \cap K_{\boldsymbol{r}}) \setminus \Omega_{\boldsymbol{r}} = (S' \cap K_{\boldsymbol{r}}) \setminus \Omega_{\boldsymbol{r}}$ (as argued earlier), we conclude $S = S'$. □

We are now ready to prove

$$|\bar{\mathcal{S}}| \left(1 + \frac{L}{n}\right) \leq |\operatorname{ext} \bar{Q}|.$$

Recall that we defined $\bar{\mathcal{S}}^*$ as the set of $S \in \bar{\mathcal{S}}$ for which there exist at least L distinct sets H such that (S, H, n) is regular. Consider the correspondence R that assigns to each $S \in \bar{\mathcal{S}}^*$ such a set of L distinct tuples H. Lemmata B.13 and B.14 imply that the mapping $\boldsymbol{r} \mapsto q_{\boldsymbol{r}}$ is an injection from the graph of R to the set of stochastic extreme points. We observed in (B.5) that $|\bar{\mathcal{S}}| \leq n|\bar{\mathcal{S}}^*|$ holds. Hence

$$\begin{aligned} \left|\operatorname{ext} \bar{Q}\right| \geq |\bar{\mathcal{S}}| + |\{\boldsymbol{r} \colon (S, \boldsymbol{r}) \in \operatorname{graph} R\}| &= |\bar{\mathcal{S}}| + \sum_{S \in \bar{\mathcal{S}}^*} |R(S)| \\ &\geq |\bar{\mathcal{S}}| + \sum_{S \in \bar{\mathcal{S}}^*} L \\ &\geq |\bar{\mathcal{S}}| \left(1 + \frac{L}{n}\right), \end{aligned}$$

as desired.