

Optimal disclosure of information to privately informed agents

Ozan Candogan* Philipp Strack†

Abstract

We study information design with multiple privately informed agents who interact in a game. Each agent's utility is linear in a real-valued state. We show that there always exists an optimal mechanism which is laminar partitional and bound its “complexity”. For each type profile, such a mechanism partitions the state space and recommends the same action profile within a partition element. Furthermore, the convex hulls of any two partition elements are such that either one contains the other or they have an empty intersection. We highlight the value of screening: the ratio of the optimal and the best payoff without screening can be equal to the number of types. Along the way, we shed light on the solutions to optimization problems over distributions subject to a mean-preserving contraction constraint and additional side constraints, which might be of independent interest.

*Booth School of Business, University of Chicago

†Department of Economics, Yale University

We want to thank Dan Adelman, Dirk Bergemann, John Birge, Alex Frankel, Nicole Immorlica, Emir Kamenica, Jacob Leshno, Brendan Lucier, Stephen Morris, Rad Niazadeh, Lars Stole, Alex Wolitzky, Kai Hao Yang, and participants of the EC conference for helpful discussions and suggestions. Ozan Candogan thanks the University of Chicago Booth School of Business for financial support. Philipp Strack gratefully acknowledges financial support by the Sloan Foundation.

1 Introduction

We study how a designer can use information about a real-valued state to influence the belief and actions of a group of agents who possess private information. For example, the agents could be competing firms who each decide on the quantity they produce. The production costs could be each firm’s private information and the state could measure the total demand for the product.

The designer can without loss restrict attention to direct recommendation mechanisms where each agent truthfully reports his type and then privately observes an action recommendation. We prove that there always exists an optimal such mechanism with a particularly simple structure: For each type profile there is a partition of the state space such that the mechanism recommends the same action profile for states that belong to the same partition element. Thus, there exists a (deterministic) function mapping the state and vector of types to action recommendations. Furthermore, the partition is laminar. This implies that the convex hulls of any two partition elements are either nested or they do not overlap. As a result of the laminar structure, an optimal partition can be completely described by the collection of (smallest) intervals containing the states that induce each action profile recommendation. This structure is valuable for tractability as it reduces the designer’s optimization problem from an uncountably infinite one to an optimization problem over the end points of the aforementioned intervals.

Finally, we provide a bound on the “depth” of optimal laminar partitions. In the single agent case the laminar partition structure has depth of at most $|\Theta| + 2$, where Θ is the set of types of the agent. That is, the interval associated with an action recommendation overlaps with at most $|\Theta| + 1$ other intervals (associated with different action recommendations). This implies that either (i) a state is perfectly revealed, or (ii) it lies in an interval in which the distribution of the posterior means admits at most a finite number mass points. In the multi-agent case a similar bound on the depth of the laminar partitions can be obtained if the number of possible actions is finite for each agent. In contrast to the single-agent case, where the bound is independent of the number of actions, this bound depends quadratically on it. This difference is driven by the fact that while in the single-agent case the action recommendation reveals the partition element in which the state lies; this is not the case when there are multiple agents.

Given that the state space is a continuum, it is not a priori clear how to obtain the optimal mechanism in a tractable way. To address this question, we focus on the finite action case.

We identify a transformation in the single agent case that leads to a finite-dimensional convex program (despite the states and the space of signals being uncountably infinite). Similarly, in the multi-agent case we derive a finite-dimensional (though not necessarily convex) program.

Furthermore, we discuss some properties of the optimal mechanism: Focusing on the single agent case, we prove that restricting attention to mechanisms that do not screen the agent (and reveal the same information to all types) can be strictly suboptimal and in general achieve only a $1/|\Theta|$ share of the optimal value for the designer. This is in contrast to Kolotilin, Mylovanov, Zapechelnyuk and Li (2017) and Guo and Shmaya (2019) who show that in the binary action single agent case there is no benefit to screening the agent.

Through an example we illustrate that unlike in classical mechanism design, “non-local” incentive compatibility constraints might bind in the optimal mechanism (even if the agent’s utility is supermodular in his actions and type). Finally, under the optimal mechanism the actions of different types need not be ordered for all states. For instance, there are states where the low and the high types take a higher action than the intermediate types.¹

As a crucial step in obtaining our results, we study optimization problems over distributions, where the objective is linear in the chosen distribution, and a distribution is feasible if it satisfies (i) a majorization constraint as well as (ii) some linear side constraints. We characterize properties of optimal solutions to such problems. In particular, we show that one can find optimal distributions that redistribute the mass in each interval where the majorization constraint does not bind, to at most $n + 2$ mass points, where n is the number of side constraints. Moreover, there exists a laminar partition of the underlying state space such that the signal based on this laminar partition “generates” the optimal distribution. Our main result is proven by decoupling the information design problem over type profiles into optimization problems under majorization and linear side constraints. Given the generality of such optimization formulations, we suspect that our results may have applications beyond the information design problem studied in the paper. We discuss some immediate applications in Section 5.

Literature Review Following the seminal work by Kamenica and Gentzkow (2011), the literature on Bayesian persuasion studies how a designer can use information to influence

¹This can be leveraged to show that “nested” information structures that are optimal in related information design settings with 2 actions are suboptimal (see e.g., Guo and Shmaya, 2019).

the action taken by an agent. This framework has proven useful to analyze a variety of economic applications, such as the design of grading systems², medical testing³, stress tests and banking regulation⁴, voter mobilization and gerrymandering⁵, as well as various applications in social networks⁶. For an excellent survey of the literature see Kamenica (2019) and Bergemann and Morris (2019).

Initial papers focused on either the case of a single agent who possesses no private information or the case where the designer uses public signals (Brocas and Carrillo, 2007; Rayo and Segal, 2010; Kamenica and Gentzkow, 2011; Gentzkow and Kamenica, 2016b). Kolotilin et al. (2017) and Guo and Shmaya (2019) extend this baseline model by considering the single-agent case where the agent possesses private information about his preferences and chooses between *two* actions. Assuming that the agent’s payoff is linear and additive in the state, Kolotilin et al. (2017) show that it is without loss to restrict attention to “public” signals, which do not screen the agent and induce the same signal realization regardless of the type of the agent. Guo and Shmaya (2019) consider a general monotone utility of the designer and the agent, but maintain the assumption of binary actions. They show that even though not every outcome that can be implemented with private signals can also be implemented with public signals, it is nevertheless true that the *designer-optimal* outcome can always be implemented with public signals. We complement this line of the literature by studying the case where the agent can potentially choose among *more than two actions* and find in contrast with the binary action case that public signals could yield a payoff that is as low as one over the number of types fraction of the optimal one.⁷

Bergemann and Morris (2013, 2016) consider information revelation to multiple agents and introduce the notion of “Bayes correlated equilibria”. Bayes correlated equilibria characterize the set of all outcomes that can be induced in a given game by revealing a private signal to each agent. Thus, Bayesian persuasion problems can be solved by maximizing over the set of Bayes correlated equilibria. While the basic concept does not allow for private information one can extend to the case with screening and private information (see

²Ostrovsky and Schwarz (2010); Boleslavsky and Cotton (2015); Onuchic and Ray (2020).

³Schweizer and Szech (2018).

⁴Inostroza and Pavan (2018); Goldstein and Leitner (2018); Orlov, Zryumov and Skrzypacz (2018).

⁵Alonso and Câmara (2016); Kolotilin and Wolitzky (2020).

⁶Candogan and Drakopoulos (2017); Candogan (2019b).

⁷Kolotilin et al. (2017) also provide an example showing that with more than 2 actions restricting attention to public signals may result in a payoff loss (see online Appendix A of their paper). We strengthen this insight and in Section 4.3 we establish that the maximal payoff loss due to focusing on public signals is one over the number of types fraction of the optimal one. Moreover, we show that this bound is tight.

Definition 2 in Bergemann and Morris, 2019). In this case, the private information is about the state, and hence an agent’s payoff depends on his private information only through the state. As the designer learns the state once it is realized, she will be better informed about the agents’ utilities than themselves. While the formulation is present in the literature, as far as we know the structural properties of the optimal mechanisms are not well understood in the multiple agent case with private information. In the present paper, we contribute to this literature in two ways: First, we allow the utility of an agent to directly depend on his private information, thereby relaxing the assumption that the designer is better informed than the agents – which might be economically restrictive in some settings. Second, we consider a continuum of states and focus on quasi-linear utilities which allows us to describe optimal mechanisms more explicitly in terms of laminar partitional signals.⁸

Without private information, the approaches in Bergemann and Morris (2016); Kolotilin (2018); Dworzak and Martini (2019), can be used to characterize the optimal information structure. These approaches lead to infinite-dimensional optimization problems even if there is a single agent with finitely many actions. When there is a single agent, an alternative approach due to Gentzkow and Kamenica (2016b) is to associate a convex function with each information structure, and cast the information design problem as an optimization problem over all convex functions that are sandwiched in between two convex functions (associated with the full disclosure and no-disclosure information structures). This also yields an infinite-dimensional optimization problem. In contrast, we provide a finite-dimensional optimization formulation that is applicable with multiple privately-informed agents and finitely many actions. This formulation is also convex when there is a single agent, thereby providing a tractable framework for obtaining optimal mechanisms.

The aforementioned “sandwiching” constraint is equivalent to a majorization constraint restricting the set of feasible posterior distributions. Arieli, Babichenko, Smorodinsky and Yamashita (2020) and Kleiner, Moldovanu and Strack (2020) characterize the extreme points of this set. As also observed in Candogan (2019a,b), this characterization implies that in the single agent case without private information one can restrict attention to signals where each state lies in an interval such that for all states in that interval at most 2 messages are sent. There are two additional critical challenges in our setting. First, unlike earlier work, one needs to deal with additional constraints that stem from the screening

⁸Quasi-linearity assumption is commonly made in the literature. See for instance Ostrovsky and Schwarz (2010); Ivanov (2015); Gentzkow and Kamenica (2016b); Kolotilin et al. (2017); Kolotilin (2018). For a more detailed discussion of this setting and its economic applications see Section 3.2 in Kamenica (2019).

problem. Second, since there are multiple agents, the information revealed to one agent can influence the actions taken by others, which intricately couples the information design problems for different agents. These challenges require a novel approach and render the information structures identified in the earlier literature suboptimal.

2 Model

We consider an information design setting in which a designer (she) tries to influence the action taken by privately informed agents (he/they), indexed by $i \in \{1, \dots, |N|\} = N$.

States and Types We call the information controlled by the designer the state $\omega \in \Omega$ and the private information of agent i his type $\theta_i \in \Theta_i$. The state ω lies in an interval $\Omega = [0, 1]$ and is distributed according to the (cumulative) distribution $F : \Omega \rightarrow [0, 1]$, with density $f \geq 0$.⁹ Each agent's type θ_i lies in a finite set Θ_i and we denote by $\phi(\theta) > 0$ the probability that the type vector equals $\theta = (\theta_1, \dots, \theta_{|N|}) \in \Theta \subseteq \prod_{i \in N} \Theta_i$. We assume that the state ω and the types θ are independently distributed, but allow for arbitrary correlation between the types of different agents.

Signals and Mechanisms A *direct mechanism* $\mu : \Theta \times \Omega \rightarrow \Delta(S)$ maps a type profile θ and a state ω to a conditional distribution $\mu^\theta(\cdot|\omega)$ over the set of signal realizations S .¹⁰ We denote by μ^θ the signal¹¹ associated with the type vector θ , i.e.,

$$\mu^\theta(\cdot|\omega) = \mathbb{P}[s \in \cdot | \omega, \theta].$$

Each signal realization $s \in S = \prod_{i \in N} S_i$ is $|N|$ -dimensional. The i -th coordinate s_i is privately observed by agent i , but we allow for the signals observed by different agents to be correlated. We restrict attention to signals for which Bayes rule is well defined,¹² and denote by $\mathbb{P}_\mu[\cdot|s] \in \Delta(\Omega)$ the posterior distribution induced over states by observing the signal realization s in the mechanism μ , and by $\mathbb{E}_\mu[\cdot|s]$ the corresponding expectation.

⁹The assumption that the state lies in $[0, 1]$ is a normalization that is without loss of generality for distributions with bounded support as we can rescale the state (without affecting the linearity of the utility function imposed subsequently). Furthermore, while it is important that F has no mass-points, all our result go through for any continuous distribution (which might not admit a density).

¹⁰Restricting attention to direct mechanisms is without loss of generality by the revelation principle.

¹¹We follow the convention of the Bayesian persuasion literature and call a Blackwell experiment a signal.

¹²Formally, this requires that $\mathbb{P}_\mu[\cdot|s]$ is a regular conditional probability.

When there are finitely many signal realizations

$$\mathbb{P}_\mu [\omega \leq x | s] = \frac{\sum_\theta \phi(\theta) \int_0^x \mu^\theta(\{s\} | \omega) dF(\omega)}{\sum_\theta \phi(\theta) \int_0^1 \mu^\theta(\{s\} | \omega) dF(\omega)}. \quad (\text{Bayes Rule})$$

The Agents' Actions and Utilities After observing his type θ_i , each agent i reports it to the mechanism. Given the reported type profile θ and the state realization, the mechanism draws a signal from the corresponding distribution, and agent i observes the i -th coordinate s_i of the signal realization. Then each agent i chooses an action a_i in a compact set A_i to maximize his expected utility

$$\max_{a_i \in A_i} \mathbb{E} [u_i(a_i, a_{-i}, \omega, \theta) | s_i, \theta_i].$$

We note that the expectation in the above expression is over the state ω , the action taken by other agents a_{-i} , and their types θ_{-i} . If we impose additional assumptions on the set of action profiles $A = \times_{i \in N} A_i$ we will explicitly mention them, and otherwise we allow it to be finite or infinite.

Recommendation Mechanisms A *direct recommendation mechanism* is a direct mechanism where the signal realization for each agent is an action recommendation, i.e., $S_i = A_i$. A direct recommendation mechanism is incentive compatible if it is optimal for each agent to report his true type θ_i and follow the action recommendation, instead of (mis)reporting his type as $\theta'_i \in \Theta_i$ and choosing an optimal action afterwards. Throughout without loss we focus on incentive compatible direct recommendation mechanisms. Formally, denoting by $\sigma_i : A_i \rightarrow A_i$ the action policy that maps an action recommendation to an action taken by agent i , the incentive compatibility requirement can be stated as follows:¹³

$$\begin{aligned} & \sum_{\theta_{-i}} \phi(\theta) \int_{\Omega} \int_{A_{-i}} u_i(a_i, a_{-i}, \omega, \theta) d\mu^\theta(a | \omega) dF(\omega) \\ & \geq \max_{\sigma_i} \sum_{\theta_{-i}} \phi(\theta) \int_{\Omega} \int_{A_{-i}} u_i(\sigma_i(a_i), a_{-i}, \omega, \theta) d\mu^{(\theta'_i, \theta_{-i})}(a | \omega) dF(\omega) \end{aligned} \quad (1)$$

for all i , $\theta_i, \theta'_i \in \Theta_i$. One challenge in this environment is that each agent can deviate by simultaneously misreporting his type and taking an action different from the one that is

¹³When $\theta_i = \theta'_i$ this constraint reduces to the *obedience* constraint, which ensures that it is optimal for agent i to follow the action recommendation.

recommended by the mechanism.

The Designer's Utility We denote by $v : A \times \Omega \times \Theta \rightarrow \mathbb{R}$ the designer's utility. For a given direct recommendation mechanism the designer's expected utility equals

$$\sum_{\theta} \phi(\theta) \int_{\Omega} \int_A v(a, \omega, \theta) d\mu^{\theta}(a|\omega) dF(\omega). \quad (2)$$

The designer's information design problem is to pick a direct recommendation mechanism that satisfies (1) to maximize (2).

To make this setting with infinitely many states tractable we further focus on preferences which are quasi-linear in the state:

Assumption 1 (Quasi-Linearity). *The agents' utilities $\{u_i\}$ and the designer's utility v are quasi-linear in the state, i.e., for $i \in N$ there exist functions $u_{i1}, u_{i2}, v_1, v_2 : A \times \Theta \rightarrow \mathbb{R}$ continuous in $a \in A$ such that*

$$\begin{aligned} u_i(a, \omega, \theta) &= u_{i1}(a, \theta)\omega + u_{i2}(a, \theta) \\ v(a, \omega, \theta) &= v_1(a, \theta)\omega + v_2(a, \theta). \end{aligned}$$

Assumption 1 is natural in many economic situations and is commonly made in the literature (c.f. Footnote 8).¹⁴ For example Kolotilin et al. (2017) assume that there is a single agent who has two actions $\{0, 1\}$, and that the agent's utility for one action is zero, and for the other action it is the sum of the type and state, which implies that $u_i(a_i, \omega, \theta) = a_i \times (\omega + \theta)$.

Remark. Our results generalize to the case where the preferences of all agents and the designer depend linearly¹⁵ on some (potentially) non-linear transformation of the state $h(\omega)$ as long as the distribution of $h(\omega)$ admits a density.¹⁶ What is crucial for our results is that the agents' belief about the state influences the preference of the designer and the agent only through the same real valued statistic.

¹⁴We also note that the continuity of the payoffs in a , and the compactness of A together with Assumption 1 ensure that the payoffs are bounded, i.e., $|u_i(a, \omega, \theta)|, |v(a, \omega, \theta)| \leq B$ for some $B < \infty$. In what follows, this mild technical condition is used to change the order of integrals that appear in the designer's and agents' problems.

¹⁵In the single agent case we could allow u, v to depend non-linearly on the agent's posterior expectation of $h(\omega)$.

¹⁶To see this note that for every function $h : \Omega \rightarrow \mathbb{R}$ we can redefine the state to be $\tilde{\omega} = h(\omega)$.

2.1 A Motivating Example

We next provide an economic example to illustrate the model. Two firms 1, 2 producing a good choose production quantities in $A_1 = A_2 = \{0, 1, 2\}$. The price of the product depends on the total production $a_1 + a_2$ by the firms, and is given by $d - (a_1 + a_2)$, where $d = (4 + 8\omega)$ is the demand for the good and $\omega \sim U([0, 1])$ is the state. The unit production cost of each firm is its private type and equals 4 or 6 with equal probability independently of each other and the state (i.e., $\Theta = \{(4, 4), (4, 6), (6, 4), (6, 6)\}$, $\phi \equiv 1/4$). The consumer surplus (CS) and total firm profits (FP) are respectively given by $CS = (a_1 + a_2)^2/2$, and

$$FP = ((4 + 8\omega) - (a_1 + a_2))(a_1 + a_2) - a_1\theta_1 - a_2\theta_2.$$

We are interested in characterizing the combinations of consumer surplus and firm profits that can be induced by a mediator that facilitates information exchange between the firms. To do so, we numerically derive the information structures maximizing different weighted combinations of CS and FP .¹⁷ The results are illustrated in Figure 1.

We highlight 6 points that are extremal in terms of achievable CS , FP , or welfare ($CS + FP$) and display the corresponding optimal information structures. As our main result establishes we can restrict attention to simple (laminar) signals where for each type profile the state space is partitioned such that in each partition element the same actions are taken by the agents. In this example this means that for each cost vector of the firms, the interval of possible demands is partitioned such that in each partition element the output vector of the firms is constant. In Figure 1 (b) and (c), we associate with each strategy profile a color, and use them to present the information structures that achieve these extremal points. Given the symmetry between firms, the strategy profiles associated with type profiles $(4, 6)$ and $(6, 4)$ are same up to a permutation of the agents' identities. To avoid redundancy we only display one of them.

A few economic observations are worth highlighting: First, when firms' costs are lower the expected production quantities are higher. While this monotonicity holds when taking the expectation over demand levels, production quantities are not monotone in the production cost for a fixed demand. For example when maximizing the CS in (6) for some of the high demand levels the total production is higher for the production costs $(6, 6)$ than it is for the production cost $(4, 6)$. Interestingly, the worst information structure in terms of welfare (4) is when no information about the cost of their competitors and the demand is re-

¹⁷See Supplementary Appendix B for details on the numerical computations for this example.

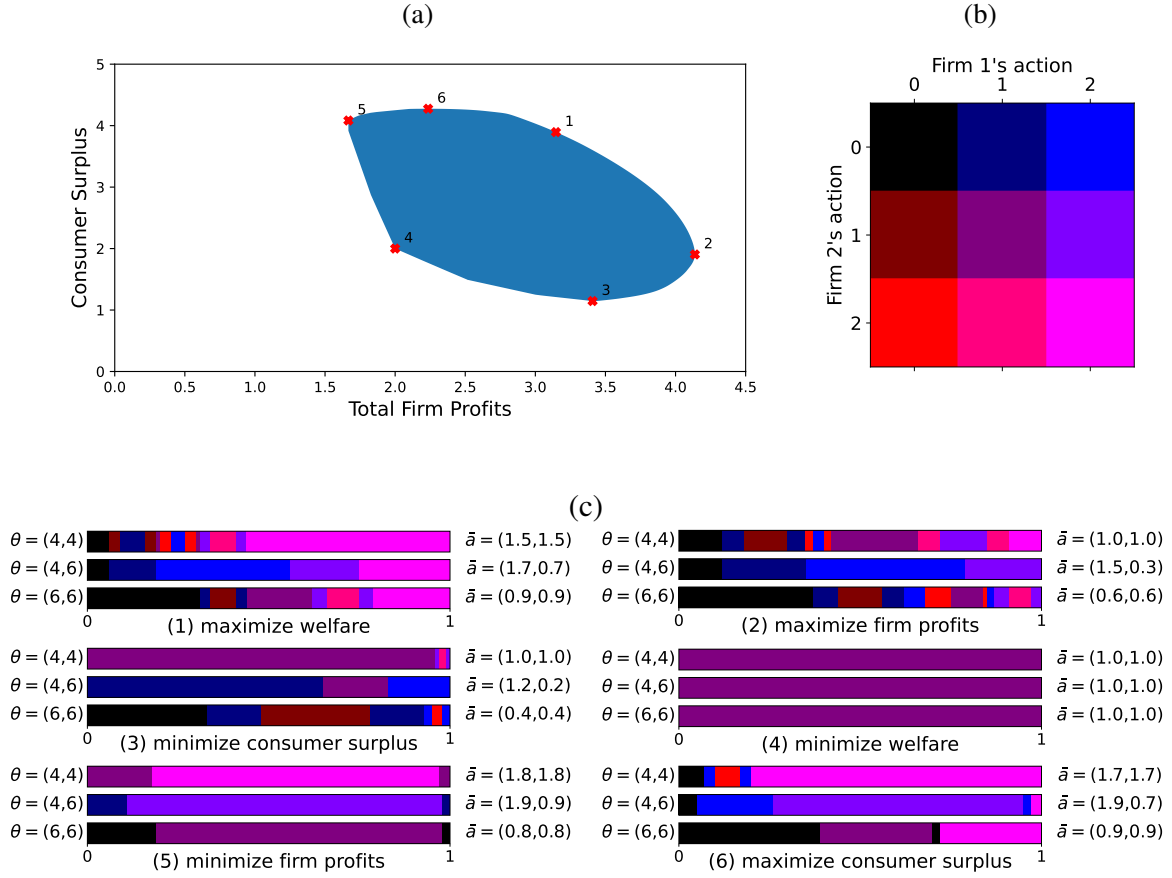


Figure 1: (a) CS and FP achievable under different information structures. We highlight 6 points on this region that achieve: (1) maximum welfare $CS + FP$, (2) maximum FP , (3) minimum CS , (4) minimum welfare $CS + FP$, (5) minimum FP , (6) maximum CS . (b) The colors assigned to different strategy profiles. (c) Optimal information structures. Here, for any type profile θ we denote by \bar{a} the vector of expected production quantities for both firms.

vealed to the firms. Conversely, maximizing welfare, FP or CS leads to nontrivial laminar partitions (1,2,6). To maximize firms' profits (2) the information structure induces more extreme asymmetric outcomes (e.g., (0, 2) or (2, 0) where one firm produces two units and the other produces zero) relative to consumer surplus maximizing information structures (where balanced outcomes such as (1, 1) become more common). This leads to a larger number of distinct signal realizations in case of profit maximization. Finally, based on the information structure, the consumer surplus and firms' profits vary significantly and there is more than a factor of two between the smallest and largest values of the aforementioned

quantities.

3 Analysis

Our analysis proceeds in several steps. First, we show that given a direct recommendation mechanism the designer can achieve the same payoff by using what we refer to as state garbling recommendation (SGR) mechanism. This reduction is consequence of our restriction to quasi-linear utilities and an auxiliary step in proving our main result. Second, we show that optimal SGR mechanisms can be characterized through problems which are decoupled across type profiles (but not across agents¹⁸). Each of the decoupled problems involves optimization over posterior mean distributions under linear side constraints. Third, we establish that solutions to such problems can always be induced by constructing a laminar partition and pooling states according to that partition. Finally, this implies our main result that there exists an optimal mechanism which for each type profile constructs a laminar partition of the state space and recommends the same action profile for states that belong to the same partition element.

3.1 State Garbling Recommendation Mechanisms

An SGR mechanism is an incentive compatible direct recommendation mechanism that for each type profile θ has the following structure:

- (i) The designer chooses an auxiliary signal ν^θ whose realization $m \in [0, 1]$ equals the induced posterior mean, i.e. $\mathbb{E}_{\nu^\theta} [\omega|m] = m$.
- (ii) For each realized posterior mean she chooses a distribution over recommended action profiles such that no action profile is recommended with positive probability at two different posterior means.

We next argue that due to our assumption of quasi-linear utilities the restriction to SGR mechanisms is without loss.¹⁹ We start with an arbitrary direct recommendation mechanism μ . Let $m_{a,\theta} = \mathbb{E}_{\mu^\theta} [\omega|a]$ denote the mean of an outside observer's posterior belief about the

¹⁸Note that there is no similar decoupling across agents, due to the strategic interactions among them.

¹⁹Note that without the restriction in (ii) the set of mechanisms described above would equal the set of direct recommendation mechanisms as the designer could always chose a fully revealing signal in (i). Due to restriction (ii), SGR mechanisms constitute a subset of the direct recommendation mechanisms.

state after observing the action profile $a \in A$ being recommended given the type profile $\theta \in \Theta$. Note that this posterior belief never becomes known to the agents as they neither observe the complete type profile nor the recommended action profile. Define $G^\theta : [0, 1] \rightarrow [0, 1]$ to be the cumulative distribution of posterior means given the type profile θ

$$G^\theta(x) = \mathbb{P}_{\mu^\theta} [m_{a,\theta} \leq x] .$$

Define $q^\theta \in \Delta(A)$ to be the distribution over action profiles conditional on type profile θ , i.e.,

$$q^\theta(B) = \int_0^1 \mu^\theta(B|\omega) dF(\omega),$$

for $B \subseteq A$. Let $q^\theta(\cdot|x) \in \Delta(A)$ be the distribution over action profiles conditional on the posterior mean $m_{a,\theta}$ associated with the action profile being equal to $x \in [0, 1]$

$$q^\theta(B|x) = \frac{\int_B \mathbf{1}_{m_{a,\theta}=x} dq^\theta(a)}{\int_A \mathbf{1}_{m_{a,\theta}=x} dq^\theta(a)} .$$

Consider the mechanism defined by the above tuple (G, q) where $\nu^\theta([0, x]) = G^\theta(x)$ and $q = (q^\theta)_\theta$ is the distribution over actions conditional on the posterior mean. In this mechanism, given the type profile θ , the designer first draws a signal realization m according to G^θ , and then recommends an action profile according to $q^\theta(\cdot|m)$. We claim that this is a valid SGR mechanism. Note that it is possibly different from the direct recommendation mechanism we started with.

In this mechanism – assuming agents follow action recommendations – the expected payoff of the designer given the type profile θ satisfies

$$\begin{aligned} \int_{\Omega} \int_A v(a, \omega, \theta) d\mu^\theta(a|\omega) dF(\omega) &= \int_A \mathbb{E}_{\mu^\theta} [v(a, \omega, \theta)|a] dq^\theta(a) \\ &= \int_A v(a, \mathbb{E}_{\mu^\theta} [\omega|a], \theta) dq^\theta(a) = \int_A v(a, m_{a,\theta}, \theta) dq^\theta(a) \\ &= \int_{\Omega} \int_A v(a, m, \theta) dq^\theta(a|m) dG^\theta(m) . \end{aligned} \quad (3)$$

The first equality leverages the boundedness of the payoffs and changes the order of integration. The second one follows from the quasi-linearity of v , the third from the definition of $m_{a,\theta}$, and the fourth from the definition of q^θ .

Using the same argument it can be readily seen that for any reported and true type pro-

files θ' , θ such that $\theta'_{-i} = \theta_{-i}$ when agents other than i follow their action recommendations we have

$$\int_{\Omega} \int_{A_{-i}} u_i(a'_i, a_{-i}, \omega, \theta) d\mu^{\theta'}(a|\omega) dF(\omega) = \int_{\Omega} \int_{A_{-i}} u_i(a'_i, a_{-i}, m, \theta) dq^{\theta'}(a|m) dG^{\theta'}(m), \quad (4)$$

where the left (right) hand side is the expected payoff of agent i from observing action recommendation a_i and taking action a'_i in the initial (new) mechanism. Since for any action recommendation the payoffs of agents coincide under the two mechanisms, it follows that the mechanism defined by the (G, q) tuple satisfies incentive compatibility, and hence is a valid SGR mechanism. Together with (3), this observation implies that the two mechanisms also yield the same payoff to the designer, and it is without loss to restrict attention to SGR mechanisms.

The expected payoff expression in the right hand side of (4) can be used to obtain a characterization of incentive compatibility of SGR mechanisms. Specifically, the SGR mechanism defined by (G, q) is incentive compatible if and only if for all $i \in N, \theta_i \in \Theta_i$

$$\begin{aligned} & \sum_{\theta_{-i} \in \Theta_{-i}} \phi(\theta) \int_{\Omega} \int_A u_i(a, m, \theta) dq^{\theta}(a|m) dG^{\theta}(m) \\ & \geq \max_{\sigma_i, \theta'_i} \sum_{\theta_{-i} \in \Theta_{-i}} \phi(\theta) \int_{\Omega} \int_A u_i(\sigma_i(a_i), a_{-i}, m, \theta) dq^{(\theta'_i, \theta_{-i})}(a|m) dG^{(\theta'_i, \theta_{-i})}(m). \end{aligned} \quad (\text{IC})$$

Feasible Posterior Mean Distributions Given that the designer's payoff and the incentive compatibility constraint can be expressed in terms of the distributions over posterior means $G = (G^{\theta})$ and distributions over action profiles $q = (q^{\theta})$ conditional on posterior means, it may be possible to reformulate the designer's problem in terms of these quantities. A natural question is thus which distributions over posterior means the designer can induce using a signal. An important notion to address this question is *mean preserving contractions* (MPC). A distribution over states $H : \Omega \rightarrow [0, 1]$ is an MPC of a distribution $\tilde{H} : \Omega \rightarrow [0, 1]$, expressed as $\tilde{H} \preceq H$, if and only if for all ω

$$\int_{\omega}^1 H(z) dz \geq \int_{\omega}^1 \tilde{H}(z) dz, \quad (\text{MPC})$$

and the inequality holds with equality for $\omega = 0$.

To see that $F \preceq G^\theta$ is necessary for G^θ to be the distribution of the posterior mean induced by some signal note that for every convex function $h : [0, 1] \rightarrow \mathbb{R}$ we have that

$$\int_0^1 h(z) dF(z) = \mathbb{E}[h(\omega)] = \mathbb{E}[\mathbb{E}_\mu[h(\omega)|s]] \geq \mathbb{E}[h(\mathbb{E}_\mu[\omega|s])] = \int_0^1 h(z) dG^\theta(z).$$

Here, the second equality is implied by the law of iterated expectations and the inequality follows from Jensen's inequality. Taking $h(z) = \max\{0, z - \omega\}$ then yields that $F \preceq G^\theta$. This condition is not only necessary, but also sufficient, see, e.g., Blackwell (1950); Blackwell and Girshick (1954); Rothschild and Stiglitz (1970) and Gentzkow and Kamenica (2016b) for an application to persuasion problems.

Lemma 1. *There exists a signal that induces the distribution G^θ over posterior means if and only if $F \preceq G^\theta$.*

This result readily implies that a vector of type profile dependent posterior mean distributions $(G^\theta)_{\theta \in \Theta}$ is feasible if and only if $F \preceq G^\theta$ for all $\theta \in \Theta$.

Optimal SGR Mechanisms Combining the characterization of incentive compatibility from (IC) and feasibility from Lemma 1 we next provide a characterization of optimal SGR mechanisms.

Proposition 1. *An SGR mechanism defined by (G, q) is incentive compatible and maximizes the designer's payoff if and only if (G, q) solve*

$$\begin{aligned} \max_{G, q} \quad & \sum_{\theta \in \Theta} \phi(\theta) \int_{\Omega} \int_A v(a, m, \theta) dq^\theta(a|m) dG^\theta(m) \\ \text{s.t.} \quad & \text{(IC)} \quad \& \quad F \preceq G^\theta \quad \forall \theta. \end{aligned} \tag{OPT}$$

One of the main challenges in this optimization problem is that even for a fixed q the incentive compatibility constraint induces a strong interdependence among the components of G , which makes it impossible to optimize over them separately. This interdependence is a natural economic feature of the multi-agent problem with private information as the designer cannot pick the action recommendation she provides to one agent and type without taking into account the fact that this might give other agents and types incentives to deviate.

3.2 Decoupling the Problem Across Type Profiles

Despite these challenges, we are able to characterize the structure of the optimal SGR mechanisms. Our approach involves decoupling the designer's problem into $|\Theta|$ sub-problems (one for each type profile θ) each involving optimization over only a single MPC constraint and linear side constraints. As the argument for doing so and the precise decomposition differ significantly in the single- and multi-agent cases we explain them separately.

3.2.1 The Single Agent Case

In this section we consider the single agent case $|N| = 1$ and thus drop the subindex indicating the agent's identity. We define $\bar{u}, \bar{v} : \Omega \times \Theta \rightarrow \mathbb{R}$ to be the agent's and designer's *indirect utility functions*, i.e. their utility at a given mean belief m if the agent takes an optimal action²⁰

$$\bar{u}(m, \theta) = \max_{a \in A} u(a, m, \theta) \quad (5)$$

$$\bar{v}(m, \theta) = \max_{a \in A(m, \theta)} v(a, m, \theta), \quad (6)$$

where $A(m, \theta) = \operatorname{argmax}_{b \in A} u(b, m, \theta)$. Since in an SGR mechanism no action is recommended at two different posterior means the agent can infer the posterior mean from the action recommendation. As any action recommendation policy q that satisfies (IC) must always recommend an action that is optimal for the agent at that posterior belief²¹, we can rewrite (IC) as

$$\int_{\Omega} \bar{u}(m, \theta) dG^{\theta}(m) \geq \max_{\theta'} \int_{\Omega} \bar{u}(m, \theta) dG^{\theta'}(m). \quad (7)$$

Let (G^*, q^*) be an optimal solution to the problem given in Proposition 1. We define the value e_{θ} type θ could achieve when deviating optimally from reporting his type truthfully

$$e_{\theta} = \max_{\theta' \neq \theta} \int_{\Omega} \bar{u}(m, \theta) dG^{*, \theta'}(m). \quad (8)$$

We also define d_{θ} to be the value the agent gets when reporting his type truthfully

$$d_{\theta} = \int_{\Omega} \bar{u}(m, \theta) dG^{*, \theta}(m). \quad (9)$$

²⁰We note that the indirect utility \bar{u} is convex in m .

²¹Formally, this means that $a \notin A(m, \theta) \Rightarrow q^{\theta}(a|m) = 0$.

We note that $e_\theta, d_{-\theta}$ do not depend on $G^{*,\theta}$. We can thus characterize $G^{*,\theta}$ by optimizing over G^θ while taking $(G^{*,\theta'})_{\theta' \neq \theta}$ as given. This leads to our next lemma.

Lemma 2. *Consider the single agent case and let e, d be the constants associated with an optimal SGR mechanism (G^*, q^*) . Then $(H^\theta, (G^{*,\theta'})_{\theta' \neq \theta}, q^*)$ is an optimal SGR mechanism if and only if for any type $\theta \in \Theta$ the distribution H^θ solves*

$$\max_{H^\theta \succeq F} \int_{\Omega} \bar{v}(s, \theta) dH^\theta(s) \quad (10)$$

$$s.t. \int_{\Omega} \bar{u}(s, \theta) dH^\theta(s) \geq e_\theta \quad (11)$$

$$\int_{\Omega} \bar{u}(s, \eta) dH^\theta(s) \leq d_\eta \quad \forall \eta \neq \theta. \quad (12)$$

In this formulation we maximize the payoff the designer receives from type θ under constraint (11). This constraint ensures that type θ does not want to deviate and report to be another type.²² Similarly, constraint (12) ensures that no other type wants to report his type as θ . We note that (11) and (12) encode the incentive constraints given in (7) in which G^θ appears.

3.2.2 The Multi-Agent Case

We next turn to the multi-agent case. Without loss we normalize here the probability of each type profile to $\phi(\theta) = 1/|\Theta|$ to make the equations easier to read.²³ The main challenge relative to the single agent case is that in an SGR mechanism the agents are in general unable to infer the posterior mean $m_{a,\theta}$ from their action recommendation. As a consequence, the action recommendations q are not determined by the (IC) constraint and we cannot omit them from the problem.

Let (G^*, q^*) be a solution to (OPT) and consider the corresponding SGR mechanism. Define $e_{i,\theta_i,\theta'_i,\sigma_i}$ to be the payoff²⁴ agent i of type θ_i gets by reporting his type as θ'_i (where

²²By considering the optimal deviation we reduced the number of incentive constraints in (8) from $(|\Theta|-1)$ to 1.

²³This is without loss of generality as given a problem instance with arbitrary $\phi(\theta)$ one can define a new utility $|\Theta| \phi(\theta) u_i(a, \omega, \theta)$ for each agent i and the designer $|\Theta| \phi(\theta) v(a, \omega, \theta)$ which entails exactly the same incentives in the original and the new problem instances. This amounts to a change of measure from $\phi(\cdot)$ to the uniform measure.

²⁴More precisely these quantities correspond to payoffs multiplied with the probability $\sum_{\eta_{-i}} \phi(\theta_i, \eta_{-i})$ of the event that agent i 's private type equals θ_i . For the subsequent discussion, this normalization does not play a role. Thus, with some abuse of terminology, we refer to such quantities as payoffs.

possibly $\theta'_i = \theta_i$) and then deviating according to action policy σ_i :

$$e_{i,\theta_i,\theta'_i,\sigma_i} = \sum_{\theta_{-i} \in \Theta_{-i}} \int_{\Omega} \int_A u_i(\sigma_i(a_i), a_{-i}, m, \theta_i, \theta_{-i}) dq^{*,(\theta'_i, \theta_{-i})}(a|m) dG^{*,(\theta'_i, \theta_{-i})}(m).$$

Letting $I(\cdot)$ denote the identity, the payoff of agent i from truthfully reporting his type and following the action recommendation equals $e_{i,\theta_i,\theta_i,I}$. Similarly, the best payoff he can achieve after misreporting his type and possibly taking an action different from the recommended one equals

$$e_{i,\theta_i} = \max_{\sigma_i, \theta'_i \neq \theta_i} e_{i,\theta_i,\theta'_i,\sigma_i}. \quad (13)$$

It will be convenient to decompose the payoff $e_{i,\theta_i,\theta'_i,\sigma_i}$ into payoff when the reported type profile equals some $\theta' \in \Theta$ and the sum of payoffs $\gamma_{i,\theta_i,\theta'_i,\sigma_i}(\theta')$ from other type profiles:²⁵

$$e_{i,\theta_i,\theta'_i,\sigma_i} = \gamma_{i,\theta_i,\theta'_i,\sigma_i}(\theta') + \int_{\Omega} \int_A u_i(\sigma_i(a_i), a_{-i}, m, \theta_i, \theta'_{-i}) dq^{*,\theta'}(a|m) dG^{*,\theta'}(m).$$

By definition for $\theta'_i \neq \theta_i$ the quantities e_{i,θ_i} and $\gamma_{i,\theta_i,\theta'_i,\sigma_i}(\theta')$ do not depend on the posterior mean distribution $G^{*,\theta}$. We next show that we can characterize $G^{*,\theta}$ by optimizing the posterior mean distribution chosen for type profile θ while taking those for other type profiles $(G^{*,\eta})_{\eta \neq \theta}$ as given.

Lemma 3. *Let e, γ be the constants associated with an optimal SGR mechanism (G^*, q^*) . Then $(H^\theta, (G^{*,\theta'})_{\theta' \neq \theta}, q^*)$ is an optimal SGR mechanism if and only if for any type $\theta \in \Theta$ the distribution H^θ solves*

$$\max_{H^\theta \succeq F} \int_{\Omega} \int_A v(a, m, \theta) dq^{*,\theta}(a|m) dH^\theta(m)$$

such that

$$\begin{aligned} & \gamma_{i,\theta_i,\theta_i,I}(\theta) + \int_{\Omega} \int_A u_i(a, m, \theta) dq^{*,\theta}(a|m) dH^\theta(m) \\ & \geq \gamma_{i,\theta_i,\theta_i,\sigma_i}(\theta) + \int_{\Omega} \int_A u_i(\sigma_i(a_i), a_{-i}, m, \theta) dq^{*,\theta}(a|m) dH^\theta(m) \quad \forall i, \sigma_i \neq I \\ & \gamma_{i,\theta_i,\theta_i,I}(\theta) + \int_{\Omega} \int_A u_i(a, m, \theta) dq^{*,\theta}(a|m) dH^\theta(m) \geq e_{i,\theta_i} \quad \forall i, \\ & \gamma_{i,\eta_i,\theta_i,\sigma_i}(\theta) + \int_{\Omega} \int_A u_i(\sigma_i(a_i), a_{-i}, m, \eta_i, \theta_{-i}) dq^{*,\theta}(a|m) dH^\theta(m) \leq e_{i,\eta_i,\eta_i,I} \quad \forall i, \sigma_i, \eta_i \neq \theta_i. \end{aligned}$$

²⁵The quantity $\gamma_{i,\theta_i,\theta'_i,\sigma_i}(\theta')$ is equivalently given by the summation defining $e_{i,\theta_i,\theta'_i,\sigma_i}$ after excluding the summand with $\theta_{-i} = \theta'_{-i}$.

This lemma can be best explained through a simple thought experiment. Suppose that we modify the initial SGR mechanism (G^*, q^*) by replacing distribution $G^{*,\theta}$ with H^θ . The left hand side of the first (and second) constraint is the payoff of agent i in the new mechanism after reporting his type truthfully and following the recommendation. The right hand side of the first constraint is the payoff achieved via truthful type report followed by a deviation to action policy σ_i . Similarly, the right hand side of the second constraint is the maximal payoff agent i can achieve in the new mechanism by misreporting his type. Note that when he misreports his type as $\theta'_i \neq \theta_i$, the signal is drawn from a distribution other than H^θ . Thus, the right hand side of the second constraint is a constant in this problem. Finally, the left hand side of the third constraint is the payoff agent i can guarantee by misreporting his type as θ_i , when his type is actually η_i . The right hand side is the payoff from truthful reporting. When H^θ satisfies these constraints, it follows that the resulting mechanism $(H^\theta, (G^{*,\theta'})_{\theta' \neq \theta}, q^*)$ still satisfies incentive compatibility, and is a valid SGR mechanism. As the objective in this optimization problem is the designer's payoff for the type profile θ , this implies the lemma.

3.3 Laminar Partitional Signals

We next describe a small class of signals, *laminar partitional signals*. We first define partitional signals:

Definition 1 (Partitional Signal). *A signal μ is partitional if for each signal realization $s \in S$ there exists a set $P_s \subseteq \Omega$ such that $\mu(\{s\}|\omega) = \mathbf{1}_{\omega \in P_s}$.*

A partitional signal partitions the state space into sets $(P_s)_s$ and reveals to the agent the set in which the state ω lies. Partitional signals are thus noiseless in the sense that the mapping from the state to the signal is deterministic. A simple example of signals which are not partitional are normal signals where the signal equals the state ω plus normal noise and thus is random conditional on the state. Denote by $\text{conv}(\cdot)$ the convex hull. The next definition further restricts the partition structure.

Definition 2 (Laminar Partitional Signal). *A partition $(P_s)_s$ is laminar if there is a partial order \triangleright on S such that $P_s = \text{conv} P_s \setminus \cup_{s' | s \triangleright s'} \text{conv} P_{s'}$ for any s . A partitional signal is laminar if its associated partition is laminar.*

This definition readily implies that $\text{conv} P_s \cap \text{conv} P_{s'} \in \{\emptyset, \text{conv} P_s, \text{conv} P_{s'}\}$ for any s, s' . The restrictions imposed by laminar partitional signals are illustrated in Figure 2.




Figure 2: The partition of the state space $\Omega = [0, 1]$ on the left is not laminar while the partition on the right is laminar as the convex hull of all pairs of sets P_1, P_2, P_3 are either nested or have an empty intersection.

Note that the elements of the laminar partition may belong to disjoint intervals (see Figure 1 for various examples of laminar partitions). We define the depth of a laminar partition as the smallest number k such that the state space can be partitioned into intervals each of which contains at most k elements of the laminar partition and every partition element is contained in one interval. Intuitively, this captures how complicated the laminar partition is: If convex hulls of partition elements are disjoint (nested) the depth is equal to 1 (the number of signal realizations $|S|$).

Definition 3 (Laminar Partitional Mechanism). *A direct recommendation mechanism is laminar partitional if it consists of laminar partitional signals, i.e., for each type profile θ there exists a laminar partition P^θ of the state space Ω such that the same action profile is recommended in each partition element.*

We next establish that there always exists a laminar partitional mechanism that is optimal. To simplify notation we denote by $P^\theta(\omega) = \{P_s^\theta : \omega \in P_s^\theta\}$ the set of states where the same signal is realized as in state ω , for a partitional signal with partition $P^\theta = (P_s^\theta)_s$.

Theorem 1. *Let $|A|$ be finite or $|N| = 1$. There exists an optimal laminar partitional mechanism. Furthermore, given the partitions partition P^θ for each θ there exists intervals $I_1^\theta, I_2^\theta, \dots$ such that*

- (i) $\omega \notin \cup_k I_k^\theta$ implies $P^\theta(\omega) = \{\omega\}$;
- (ii) $\omega \in I_k^\theta$ implies $P^\theta(\omega) \subseteq I_k^\theta$.

The proof is based on a result that characterizes the solutions to optimization problems over mean preserving contractions under linear side constraints, such as those in Section 3.2. As this result might be of independent interest we explain it in Section 3.4.

Theorem 1 simplifies the search for optimal mechanisms. First, it implies that for each type profile θ the designer needs to consider only partitional signals, which (deterministically) recommends the same action profile for all states in an element of the partition P^θ .

Theorem 1 thus implies that the designer does not need to rely on random signals whose distribution conditional on the state could be arbitrarily complex.

In general, the partitions that define a deterministic signal can be quite complex, and, for instance, each partition element can be a disjoint union of countably many subsets of states. The fact that the partition can be chosen to be laminar is thus a further important simplification. To see why, consider the case with finitely many action profiles. Since the optimal signal is partitional the signal realizations correspond to at most $|A|$ subsets of the state space. Due to the laminar structure each subset can be identified with its convex hull, which is an interval. As each interval is completely described by its endpoints it follows that each laminar partitional signal can be identified with a point in $\mathbb{R}^{2|A|}$. Thus, the problem of finding the optimal mechanism can be written as an optimization problem over $\mathbb{R}^{2|A| \times |\Theta|}$. This contrasts with the space of mechanisms, which is not finite dimensional even if one restricts to finitely many signal realizations.

The second part of the theorem implies that for each fixed type profile θ there are two types of partition elements: There are “pooling intervals” (I_k^θ) where multiple state realizations induce the same action profile recommendation. In their complement, each state is mapped to a unique action profile recommendation.²⁶ Moreover, each pooling interval can equivalently be expressed as the union of partition elements it intersects, which also constitute a laminar partition of this pooling interval. This implies that the task of constructing laminar partitions that induce optimal posterior mean distributions also decouples over pooling intervals. We next provide bounds on the depth of the laminar structure:

Proposition 2. *Consider the setting of Theorem 1.*

- (i) *If $|N| = 1$, then in each I_k^θ at most $|\Theta| + 2$ action profiles are realized.*
- (ii) *If $|A|$ is finite, then $w \in \cup_k I_k^\theta$ almost surely, and in each I_k^θ at most $\sum_{i \in N} |A_i|^2 |\Theta_i| + 2$ action profiles are realized.*

Suppose we restrict attention to SGR mechanisms (G, q) , where $q^\theta(\cdot|m)$ is degenerate and deterministically recommends an action profile for each posterior mean m and type profile θ (which is without loss for the single agent case). Then Proposition 2 follows immediately by counting the number of side constraints in Lemma 2 and 3 which by Proposition 3 given in the next section limit the depth of the laminar structure. To cover the case of nondegenerate action profile distributions, an additional compactness argument is necessary. We provide this argument in the appendix.

²⁶Signals that induce such outcomes are relevant when there is a continuum of action profiles.

Note that in Proposition 2 part (i) the number of action profiles that are realized in each interval is independent of the number of actions available to the agent, whereas this is not the case in part (ii). In fact, it is possible to construct numerical examples where the number of available actions impacts this quantity and the depth of the laminar partitional signals (see Supplementary Appendix A).

This dichotomy emerges since in the single-agent case, any action recommendation perfectly reveals to the agent the partition element containing the state and the corresponding posterior mean. As explained in Section 3.2.1 this implies that one can completely express the problem in terms of indirect utilities which makes this equivalent to a problem without action choices. In contrast, in the multi-agent case due to the uncertainty about other agents' types and action recommendations the partition element does not become common knowledge among the agents. As a consequence, it is not possible to express the designer's problem in terms of indirect utilities dropping the actions.

3.4 Maximizing under MPCs and Side Constraints

The next section derives an abstract mathematical result about optimization under MPC constraints and side constraints that implies Theorem 1 and Proposition 2. We discuss this result separately as similar mathematical problems emerge in economic applications other than Bayesian persuasion.²⁷

Consider the problem of maximizing the expectation of an arbitrary upper semicontinuous function $v : [0, 1] \rightarrow \mathbb{R}$ over all distributions G that are mean-preserving contractions of a given distribution $F : [0, 1] \rightarrow [0, 1]$ subject to $n \geq 0$ additional linear constraints

$$\begin{aligned} \max_{G \succeq F} \int_{\Omega} v(s) dG(s) \\ \text{s.t. } \int_{\Omega} u_i(s) dG(s) \geq 0 \text{ for } i \in \{1, \dots, n\}. \end{aligned} \tag{14}$$

Throughout, we assume that the functions $u_i : [0, 1] \rightarrow \mathbb{R}$ are continuous. The next result establishes conditions that need to be satisfied by any solution of problem (14). Our results extend the insights of Candogan (2019b); Arieli et al. (2020); Kleiner et al. (2020) who analyzed the problem of maximizing over mean preserving contractions *without* side-

²⁷For example Kleiner et al. (2020) discuss how optimization problems under mean preserving contraction constraints naturally arise in delegation problems. We leave the exploration of other applications of this mathematical result for future work to keep the exposition focused on the persuasion problem.

constraints. We allow for side constraints as they naturally appear as incentive constraints and in settings with multiple agents. While without side-constraints each interval is optimally contracted into a distribution with just two points in its support, we find that in general the cardinality of the support equals the number of side constraints plus two.

Proposition 3. *There exists a solution G to problem (14) and a countable collection of disjoint intervals I_1, I_2, \dots such that G equals distribution F outside the intervals, i.e.,*

$$G(x) = F(x) \text{ for } x \notin \cup_j I_j \quad (15)$$

and each interval $I_j = (a_j, b_j)$ redistributes the mass of F among at most $n + 2$ mass points $m_{1,j}, m_{2,j}, \dots, m_{n+2,j} \in I_j$

$$G(x) = G(a_j) + \sum_{r=1}^{n+2} p_{r,j} \mathbf{1}_{m_{r,j} \leq x} \quad \text{for } x \in I_j \quad (16)$$

with $\sum_{r=1}^{n+2} p_{r,j} = F(b_j) - F(a_j)$ and the same expectation $\int_{I_j} x dG(x) = \int_{I_j} x dF(x)$.

The existence of an optimal solution follows from standard arguments exploiting the compactness of the feasible set of (14). To establish the remaining claims of Proposition 3, we first fix an optimal solution, and consider an interval where the MPC constraint does not bind at this solution. As both the constraints as well as the objective function in (14) are linear functionals in the CDF we can optimize over (any subinterval of) this interval fixing the solution on the complement of this interval, to obtain another optimal solution. In this auxiliary optimization problem the MPC constraint is relaxed by a constraint fixing the conditional mean of the distribution on this interval. This problem is now a maximization problem over distributions subject to the n original constraints and an additional identical mean constraint. It was shown in Winkler (1988) that each extreme point of the set of distributions, which are subject to a given number k of linear constraints, is the sum of at most $k + 1$ mass points. For our auxiliary optimization problem, this ensures the existence of an optimal solution with $n + 2$ mass points. A challenge is to establish that the solution to the auxiliary problem is feasible and satisfies the MPC constraint. The main idea behind this step is to show that if it is not feasible, then one can construct an optimal solution where the MPC constraint binds on a larger set. However, this can never be the case if we start with an optimal solution where the set on which the MPC constraint binds is maximal (which exists by Zorn's lemma). Combining such an initial optimal solution with

the optimal solution for the auxiliary optimization problem, we obtain a new solution that satisfies the conditions of the proposition over this interval. By repeating this argument for all intervals where the MPC constraint does not bind it follows that the claim holds for the entire support.

Laminar Structure Let ω be a random variable distributed according to F . Our next result shows that each interval I_j in Proposition 3 admits a laminar partition such that when the realization of ω belongs to some I_j , revealing the partition element that contains it and simply revealing ω when it does not belong to any I_j induces a posterior mean distribution, given by G . Proposition 3 together with this result yields the optimality of partitional signals as stated in Theorem 1, as well as the depth of the corresponding laminar families presented in Proposition 2.

Proposition 4. *Consider the setting of Proposition 3 and let ω be distributed according to F . For each interval I_j there exists a laminar partition $\Pi_j = (\Pi_{r,j})_r$ such that for all $r \in \{1, \dots, n+2\}$*

$$\mathbb{P}[\omega \in \Pi_{r,j}] = p_{r,j} \quad \text{and} \quad \mathbb{E}[\omega | \omega \in \Pi_{r,j}] = m_{r,j}. \quad (17)$$

The proof of this claim relies on a partition lemma (stated in the appendix), which strengthens this result by shedding light on how the partition Π_j can be constructed. The proof of the latter lemma is inductive over the number of mass points. When G given in Proposition 3 has two mass points in I_j , the partition element that corresponds to one of these mass points is an interval and the other one is the complement of this interval relative to I_j . Moreover, it can be obtained by solving a system of equations, expressed in terms of the end points of this interval, that satisfy condition (17). As this partition is laminar this yields the result for the case where there are only 2 mass points in I_j .

When G consists of $k > 2$ mass points in I_j one can find a subinterval, such that: (i) the expected value of $\omega \sim F$ conditional on ω being inside this subinterval equals the value of the largest mass point, and (ii) the probability assigned to the interval equals the probability G assigns to the largest mass point. Conditional on ω being outside this interval, the distribution thus only admits $k - 1$ mass points and is a mean preserving contraction of the distribution F . This allows us to invoke the induction hypothesis to generate a laminar partition such that revealing in which partition element ω lies generates the desired conditional distribution of the posterior mean. Finally, as this laminar partition combined

with the subinterval associated with the largest mass point of G in I_j is again a laminar partition, we obtain the result for distributions consisting of $k > 2$ mass points.

The proof of Proposition 4 (and Lemma 11 of the Appendix) details these arguments, and also offers an algorithm for constructing a laminar partition satisfying (17). While the result is stated by focusing on the setting of Proposition 3, as can be seen from the proof, the optimality of G does not play any role. Hence, the claim continues to hold for any distribution G that satisfies only conditions (15) and (16).

4 Single Agent Case: Screening vs. No-Screening

In this section, we focus on the single agent case $|N| = 1$. Throughout we also assume that the set of actions is finite $|A| = \{1, \dots, |A|\}$ and the designer's payoff $v(a, \theta)$ depends only on the action and the agent's type. Our setting thus reduces to the problem of persuading a privately informed agent, which is of independent interest. The case of binary actions was analyzed in Kolotilin et al. (2017) and Guo and Shmaya (2019) (who analyze this problem under slightly different assumptions). We first show that in the single agent setting described above – without restricting attention to binary actions – the optimal (SGR) mechanism can be obtained by solving a *finite-dimensional convex program* (Section 4.1). Then, we exemplify the optimal mechanism and contrast it with the optimal mechanisms derived in the literature by restricting attention further to binary action settings (Sections 4.2 and 4.3).

4.1 A Convex Program for the Single-Agent Case

As a consequence of Assumption 1 there exist a partition of Ω into intervals $(B_{a,\theta})_{a \in A}$ such that action a is optimal for the agent of type θ if and only if his mean belief is in the interval $B_{a,\theta}$. By relabeling the actions for each type we can without loss assume that the intervals $B_{a,\theta} = [b_{a-1,\theta}, b_{a,\theta}]$ are ordered with respect to the actions,²⁸ and hence for all $m \in B_{a,\theta}$:

$$\bar{u}(m, \theta) = u_1(a, \theta)m + u_2(a, \theta).$$

²⁸Formally, $0 = b_{0,\theta} \leq b_{1,\theta} \leq \dots \leq b_{|A|,\theta} = 1$. If an action a is never optimal for a type θ set $b_{a-1,\theta} = b_{a,\theta} = b_{|A|,\theta} = 1$. This is without loss as no signal induces a posterior belief of 1 with strictly positive probability and the action thus plays no role in the resulting optimization problem.

Consider an SGR mechanism with posterior mean distributions (G^θ). Denote by $p_{a,\theta}$ the probability that action a is recommended to type θ and by $m_{a,\theta} \in B_{a,\theta}$ the posterior mean induced by this recommendation. The expected payoff of type θ from reporting his type as θ' equals

$$\sum_{a' \in A} p_{a',\theta'} \bar{u}(m_{a',\theta'}, \theta). \quad (18)$$

Defining

$$z_{a,\theta} = m_{a,\theta} p_{a,\theta}$$

to be the product of the posterior mean $m_{a,\theta}$ induced by the action recommendation a and the probability $p_{a,\theta}$ of that recommendation, the incentive compatibility constraint (7) for type θ can be expressed as:

$$\sum_{a \in A} u_1(a, \theta) z_{a,\theta} + u_2(a, \theta) p_{a,\theta} \geq \sum_{a' \in A} \left[\max_{a \in A} u_1(a, \theta) z_{a',\theta'} + u_2(a, \theta) p_{a',\theta'} \right] \quad \forall \theta'. \quad (19)$$

Here, the left hand side is the payoff of this type from reporting his type truthfully and subsequently following the recommendation of the mechanism, whereas the right hand side is the payoff from reporting type as θ' and taking the best possible action (possibly different than the recommendation of the mechanism) given the signal realization. Recall that the distribution G^θ is an MPC of F for all θ (Lemma 1). Our next lemma establishes that the MPC constraints also admit an equivalent restatement in terms of (p, z) .²⁹

Lemma 4. $G^\theta \succeq F$ if and only if $\sum_{a \geq \ell} z_{a,\theta} \leq \int_{1 - \sum_{a \geq \ell} p_{a,\theta}}^1 F^{-1}(x) dx$, where the inequality holds with equality for $\ell = 1$.

Our observations so far establish that the incentive compatibility and MPC constraints can both be expressed in terms of the (p, z) tuple. As a consequence of these observations we can reformulate the problem of obtaining optimal SGR mechanisms, given in Proposi-

²⁹This reformulation was first used in Candogan (2019b) and for completeness we include a proof in the appendix.

tion 1, in terms of (p, z) as follows:

$$\begin{aligned}
& \max_{\substack{p \in (\Delta^{|A|})^\Theta \\ z \in \mathbb{R}^{|A| \times |\Theta|} \\ y \in \mathbb{R}^{|A| \times |\Theta|^2}}} \sum_{\theta \in \Theta} \phi(\theta) \sum_{a \in A} p_{a,\theta} v(a, \theta) \\
& \text{s.t.} \quad \sum_{a \geq \ell} z_{a,\theta} \leq \int_{1 - \sum_{a \geq \ell} p_{a,\theta}}^1 F^{-1}(x) dx \quad \forall \theta \in \Theta, \ell > 1, \\
& \quad \sum_{a \in A} z_{a,\theta} = \int_0^1 F^{-1}(x) dx \quad \forall \theta \in \Theta, \\
& \quad u_1(a, \theta) z_{a',\theta'} + u_2(a, \theta) p_{a',\theta'} \leq y_{a',\theta,\theta'} \quad \forall \theta, \theta' \in \Theta, a, a' \in A, \\
& \quad \sum_{a' \in A} y_{a',\theta,\theta'} \leq \sum_{a \in A} (u_1(a, \theta) z_{a,\theta} + u_2(a, \theta) p_{a,\theta}) \quad \forall \theta, \theta' \in \Theta, \\
& \quad p_{a,\theta} b_{a-1,\theta} \leq z_{a,\theta} \leq p_{a,\theta} b_{a,\theta} \quad \forall \theta \in \Theta, a \in A.
\end{aligned} \tag{OPT2}$$

In this formulation, the first two constraints are the restatement of the MPC constraints (see Lemma 4). The value $y_{a',\theta,\theta'}$ corresponds to the utility the agent of type θ gets from observing the signal associated with type θ' and taking the optimal action when the recommended action is a' . It can be easily checked that $y_{a',\theta,\theta'} = \max_{a \in A} u_1(a, \theta) z_{a',\theta'} + u_2(a, \theta) p_{a',\theta'}$ at an optimal solution.³⁰ Thus, it follows that the third and fourth constraints restate the incentive compatibility constraint (19), by using $y_{a',\theta,\theta'}$ to capture the summands in the right hand side of the aforementioned constraint. Finally, the last constraint captures that the posterior mean $z_{a,\theta}/p_{a,\theta}$ must lie in $B_{a,\theta}$ for the action a to be optimal.

It is worth pointing out that (OPT2) is a finite-dimensional *convex* optimization problem. This is unlike the infinite dimensional optimization formulation of Proposition 1. (OPT2) restates the designer's problem in terms of the (p, z) tuple. Two points about this reformulation are important to highlight. First, an alternative approach would involve optimizing directly over distributions G^θ that satisfy the IC constraints (7) and have a single mass point $m_{a,\theta} \in B_{a,\theta}$ for each $a \in A$ with weight $p_{a,\theta}$. This could be formulated as a finite-dimensional problem as well (by searching over the location $m_{a,\theta}$ and weight $p_{a,\theta}$ of each mass point). However, this approach does not yield a convex optimization formulation as the set of such (p, m) tuples is *not* convex. The formulation in (OPT2) amounts to a change of variables that yields a convex program.

Second, given an optimal solution to (OPT2), the distributions $(G^\theta)_{\theta \in \Theta}$ of an optimal

³⁰This is because when $y_{a',\theta,\theta'}$ is strictly larger than the right hand side, it can be decreased to construct another feasible solution with the same objective.

SGR mechanism can be obtained straightforwardly by placing a mass point with weight $p_{a,\theta}$ at $z_{a,\theta}/p_{a,\theta}$ for each action a with $p_{a,\theta} > 0$. Moreover, as discussed in Section 3.4, an optimal mechanism that induces these distributions can be obtained by constructing a laminar partition of the state space (by following the approach in Proposition 4 and Lemma 11 of the Appendix). These observations imply our next proposition.

Proposition 5. *For every optimal solution (p, z, y) of (OPT2) the SGR mechanism which recommends the action a for type θ with probability $p_{a,\theta}$ and induces a posterior mean of $z_{a,\theta}/p_{a,\theta}$ (when $p_{a,\theta} > 0$) is an optimal mechanism. Moreover, there exists a laminar partitional mechanism implementing these distributions.*

Remark. For the multi-agent case it is possible to obtain a similar finite-dimensional optimization problem. However, in this case, there are two difficulties. First, while in the single-agent case the actions associated with different posterior mean levels are known, this is not the case for multiple agents. This issue can be circumvented by optimizing over the order of posterior mean levels associated with different action profiles. Second, unlike the formulation of this section the resulting optimization problem is non-convex. In some instances, including the one in Section 2.1, one can get around these difficulties by leveraging further structure of the problem. More generally, numerical methods for non-convex optimization can be used. See Supplementary Appendix B for details.

4.2 An Example

Section 2.1 illustrates optimal laminar partitional mechanisms in a Cournot game. We next illustrate our results through a simpler single agent example. This example generalizes the buyer-seller setting from Kolotilin et al. (2017), who assume single unit demand, to the case where the buyer can demand more than one unit and has a decreasing marginal utility in the number of units. As our example reduces to their setup for the case of a single unit, this example allows us to highlight the effects of the buyer having more than two actions.

In this example, the agent is a buyer who decides how many units of an indivisible good to purchase. He is privately informed about his type which captures his taste for the good. The designer is a seller who controls information about the quality of the good, captured by the state. We assume that prices are linear in consumption and set the price of one unit of the good to $10/3$. The utility the buyer derives from the a -th unit of the good is given by

$$(\theta + \omega) \max\{5 - a, 0\}.$$

His marginal utility of consumption decreases linearly in the number of goods, increases in the good's quality ω , and in his taste parameter θ . The quality of the good is distributed uniformly in $[0, 1]$ and the buyer's taste parameter either takes a low $\theta = 0.3$, intermediate $\theta = 0.45$, or high value $\theta = 0.6$ with equal probability. The seller commits to a laminar partitional mechanism to maximize the (expected) number of units sold. It is straightforward to see that in this problem the agent considers finitely many actions: purchasing 0, 1 and 2 units (see Supplementary Appendix C). Hence the designer's problem can be formulated and solved using the finite-dimensional convex program of Section 4.1. We solve this program, and construct the optimal laminar partitional mechanism (displayed in Figure 3).³¹




Figure 3: The optimal SGR mechanism.

In this figure, each bar represents the state space and its differently colored regions the optimal partition (for the corresponding type). For each type, the designer reveals whether the state belongs to the region(s) marked with 0, 1, 2; and the buyer finds it optimal to purchase the corresponding number of units. Under the optimal mechanism, the *expected* purchase quantity increases with the type.³² While the expected quantities are ordered, the quantities purchased by different types for a given state are *not*. For instance, for states between 0.79 and 0.83 the low and the high types purchase two units, and the medium type purchases one unit. Note that this implies that the purchase regions of buyers are not “nested” in the sense of Guo and Shmaya (2019), who establish the optimality of such a nested structure for the case of two actions $|A| = 2$. Moreover, low and medium types may end up purchasing lower quantities in some high states, than they do for lower states.

³¹In the figure, the cutoffs are reported after rounding, e.g., the cutoff for the high type is approximately at 0.06. For sake of exposition, in our discussion we stick to the rounded values.

³²This can be seen as the high type purchases two units in the states where the medium type purchases only one unit, which in turn leads to higher expected purchase. Similarly, when the low type purchases zero units, the medium type purchases zero or one units; and the size of the set of states where the medium type purchases two units is larger than that for the low type.

In fact, under the optimal mechanism, for the best and the worst states, the low type purchases zero units. Thus, in the optimal mechanism, the low and medium type of the buyer sometimes consume a *smaller* quantity of the good if it is of higher quality. This (maybe counterintuitive) feature of the optimal mechanism is a consequence of the incentive constraints: By pooling some high states with low states, one makes it less appealing for the high type to deviate and observe the signal meant for a lower type.

Remark. In case of binary actions and under some assumptions on the payoff structure,³³ Kolotilin et al. (2017) and Guo and Shmaya (2019) establish that the optimal mechanism admits a “public” implementation. For each type the corresponding laminar partitioned signal induces one action in a subinterval of the state space, and the other action in the complement of this interval. It can be shown that these intervals are nested which implies that the mechanism that reveals messages associated with different types to all agent types is still optimal. Thus, as opposed to first eliciting types and then sharing with each type the realization of the signal associated with this type, the designer can achieve the optimal outcome by sharing a signal (which encodes the information of the signals of all types) publicly with all agent types. In other words, screening is not useful. By contrast, it is straightforward to establish that the mechanism illustrated in Figure 3 does not admit a public implementation, and any public mechanism yields strictly lower payoffs to the designer. See Supplementary Appendix C for further details.

Remark. Given the mechanism of Figure 3, one can readily check which incentive compatibility constraints are binding. It turns out that both the medium and the high types are indifferent among reporting their types as low, medium, or high. Similarly, the low type is indifferent between reporting his type as low or medium, but achieves strictly lower payoff from reporting his type as high. Interestingly, these observations imply that unlike in classical mechanism design settings “non-local” incentive constraints might bind in the optimal mechanism.³⁴

³³Both papers normalize the payoff of the action 0 to zero. The assumption in Kolotilin et al. (2017) is equivalent to the assumption that for $\theta' \leq \theta$ if $\mathbb{E}[u(1, \omega, \theta')] \geq 0$ then $\mathbb{E}[u(1, \omega, \theta)] \geq 0$ under any probability measure. Guo and Shmaya (2019) establish this result under assumptions that in our setting are equivalent to $\omega \mapsto \frac{u(1, \omega, \theta)}{v(1, \omega, \theta)}$ and $\omega \mapsto \frac{u(1, \omega, \theta)}{u(1, \omega, \theta')}$ are increasing for all $\theta' \leq \theta$.

³⁴This in despite the fact that the agent’s utility is supermodular in his actions and type.

4.3 The Value of Screening and Private Signals

As discussed earlier, the optimal laminar partitional mechanism reveals different signals to different types. What if we restricted attention to public signals where all types observe the same signal? Suppose that the designer's payoff is non-negative. For any mechanism (μ^1, \dots, μ^n) where different types observe different signals, the designer can always construct a public mechanism $(\mu^\theta, \dots, \mu^\theta)$ where each type observes the signal μ^θ associated with type θ in the original mechanism. Denoting by G^θ the posterior mean distribution under μ^θ , we conclude that doing so and choosing θ optimally guarantees her at least a payoff of

$$\max_{\theta \in \Theta} \phi(\theta) \int_{\Omega} \bar{v}(s, \theta) dG^\theta(s).$$

Since the designer's payoff is nonnegative, this is at least a $1/|\Theta|$ fraction of the payoff achieved by the original mechanism:

$$\sum_{\theta \in \Theta} \phi(\theta) \int_{\Omega} \bar{v}(s, \theta) dG^\theta(s).$$

Thus, a public mechanism guarantees a $1/|\Theta|$ fraction of the payoff achieved by the optimal mechanism to the designer. We next establish that this bound is tight.

Proposition 6. *Assume that the designer's utility v is non-negative.*

- (i) *In any problem there exists a public persuasion mechanism which achieves a $1/|\Theta|$ fraction of the optimal value achievable by an optimal mechanism.*
- (ii) *In some problems no public persuasion mechanism yields more than a $1/|\Theta|$ fraction of the optimal value achievable by an optimal mechanism.*

We prove (the second part of) this proposition by explicitly constructing an example where the $1/|\Theta|$ ratio is achieved. The idea behind the example is to give all types of the agent identical preferences and chose the payoff of the designer such that she wants different types of the agent to chose different actions. In a public mechanism all agents have to choose the same action which leads to at most 1 out of $|\Theta|$ types choosing the action preferred by the designer. The example is constructed such that in a mechanism with private signals the designer can induce *all types* to chose her most preferred action. If the payoff from inducing the correct action equals 1 and the payoff from any other action to the designer equals 0, this achieves the $1/|\Theta|$ bound. The main challenge in the construction, which is handled through a careful choice of payoffs, is to ensure that all types of the agent

are indifferent between all signals to ensure that no type has incentives to misreport.

Two points are worthwhile highlighting about the example. First, it achieves the worst case $1/|\Theta|$ bound even when attention is restricted to a simple subclass of problem instances. For instance, the designer has a payoff of either 0 or 1 for different actions of the agent, and the agent has finitely many actions and type-independent utility functions. Second, by relabeling the actions one can easily modify the example such that the designer's utility is independent of the agent's type and the agent's utility depends on his type. Proposition 6 thus holds unchanged even if one restricts attention to problems where the designer's utility depends only on the agent's action, but not on his type or belief.

5 Discussion and Conclusion

Our results can be extended in various dimensions. Persuasion problems where the designer's payoff depends on the induced posterior mean, but the admissible posterior mean distributions need to satisfy additional side-constraints are naturally subsumed. Below we discuss some other economically-relevant extensions and applications of our results.

Type-Dependent Participation Constraints In our analysis we can allow each type of an agent to face a participation constraint. That is, the mechanism must provide the relevant type with at least some given expected utility. Our analysis and results carry over to this case unchanged as (IC) already encodes such an endogenous constraint capturing the value of deviating by observing the signal meant for another type. To adjust the result for this case one just needs to in addition include the value of opting out of the mechanism in the incentive constraint.

Competition among Multiple Designers Another application of our approach is to competition among multiple designers. Suppose that each designer offers a mechanism and the agents can choose to observe the signal of *one* of them.³⁵ Each designer receives a higher payoff if an agent chooses her mechanism and might have different preferences over the agents' actions. Again the designer has to ensure that the signal she provides each type of an agent with, yields a sufficiently high utility such that this type does not prefer to observe either another signal of the same designer or a signal provided by a different designer. This

³⁵Another plausible model of competition is one where the agents can observe the signals of all designers. For an analysis of this situation see Gentzkow and Kamenica (2016a).

situation corresponds to an endogenous type-dependent participation constraint which is determined in equilibrium. As our analysis works for any participation constraint it also carries over to this case.

Beyond Persuasion Problems An immediate extension is to allow the designer to influence the agents' utilities by also designing transfers. For instance, in the context of the example of Section 4.2, the seller might not only control the information she provides to the buyer, but also might charge different buyers different prices. Such settings are considered, e.g., in Wei and Green (2020); Guo, Li and Shi (2020); Yang (2020); Yamashita and Zhu (2021). As our results apply for any utility function, it is still without loss to restrict attention to laminar partitional signals. Consider the case of a single agent, who (i) has finitely many actions, and (ii) his preferences are quasi-linear in the transfers. The designer's optimal mechanism (which now determines the information structure as well as the transfers) can be formulated following an approach similar to the one in Section 4.1. Additional variables which capture transfers need to be added to the optimization formulation of that section. Due to (i) these transfers can be represented by finite-dimensional vectors, and due to (ii) the resulting problem remains convex. Thus, similar to Section 4.1 an optimal mechanism can be obtained tractably by solving a finite-dimensional convex program.

Finally, while this paper focused on persuasion problems, the mathematical result we obtain on maximization problems over mean preserving contractions under side-constraints can be applied in other economic settings which lead to similar mathematical formulations. For example as first observed in Kolotilin and Zapechelnyuk (2019) the persuasion problem is closely related to delegation problems where the agent privately observes the state and the designer commits to an action as a function of a message sent by the agent. Kleiner et al. (2020) show that this problem can be reformulated as a maximization problem under majorization constraints which is a special case of the problem we discuss in Section 3.4. Our results thus allow one to analyze delegation problems where there is a constraint on the actions taken by the designer.³⁶ For example if the agent is the manager of a subdivision of a firm and the designer is the CEO who allocates money to that subdivision depending on the

³⁶While mathematically closely related, the delegation problem is economically fundamentally different from the persuasion problem. For example the majorization constraint in the delegation problem corresponds to an incentive compatibility constraint while it corresponds to a feasibility constraint in the persuasion problem. The side constraints correspond to a feasibility constraint in the delegation problem while they correspond to an incentive compatibility constraint in the persuasion problem.

manager's report, our results allow one to analyze the case where the CEO faces a budget constraint and on average cannot allocate more than a given amount to that subdivision.

Appendix

Proof of Lemma 4. The condition $G^\theta \succeq F$ can equivalently be stated as:

$$\int_0^x (1 - G^\theta(t)) dt \geq \int_0^x (1 - F(t)) dt, \quad (20)$$

for all x , where the inequality holds with equality for $x = 1$. This inequality can be expressed in the quantile space as

$$\int_0^x (G^\theta)^{-1}(t) dt \geq \int_0^x F^{-1}(t) dt, \quad (21)$$

for all $x \in [0, 1]$, with equality at $x = 1$. Note that since G^θ is a discrete distribution, this condition holds if and only if it holds for $x = \sum_{a \leq \ell} p_{a,\theta}$ and $\ell \in A$. For such x , we have

$$\int_0^x (G^\theta)^{-1}(t) dt = \sum_{a \leq \ell} p_{a,\theta} m_{a,\theta} = \sum_{a \leq \ell} z_{a,\theta}, \quad (22)$$

and (21) becomes

$$\sum_{a \leq \ell} z_{a,\theta} \geq \int_0^{\sum_{a \leq \ell} p_{a,\theta}} F^{-1}(t) dt. \quad (23)$$

Since $\int_0^1 F^{-1}(t) dt = \int_0^1 (G^\theta)^{-1}(t) dt = \sum_{a \in A} z_{a,\theta}$, the claim follows from (23) after rearranging terms. \square

Lemma 5. *Let $|A| < \infty$ or $|N| = 1$ an optimal mechanism exists.*

Proof. We first argue that an optimal mechanism exists in the case of finitely many actions $|A| < \infty$. First, we note that set of feasible mechanisms is non-empty as the designer can always choose to reveal no information and induce a Bayes Nash equilibrium of the resulting game (which exists as there are finitely many types and actions). The action recommendations of the associated direct mechanism simply recommend to each agent the action she would take knowing only her type in a Bayes Nash equilibrium. As we have argued in Section 3.1 for every IC mechanism there exists an SGR mechanism which is

IC and achieves the same payoff for the designer. We can thus restrict attention to SGR mechanisms. As discussed in Section 3.1 these mechanisms are parametrized by $q^\theta \in \Delta(A)$ and $m_{a,\theta} \in [0, 1]$.³⁷ Thus each SGR mechanism (q, m) can be identified with a vector in $[0, 1]^{2^{|\Theta||A|}}$. Furthermore, the expected utility of the designer and an agent i can be expressed respectively as

$$\begin{aligned} & \sum_{\theta \in \Theta} \phi(\theta) \sum_{a \in A} q^\theta(a) u_i(a, m_{a,\theta}, \theta), \\ & \sum_{\theta_{-i}} \phi(\theta) \sum_{a_{-i}} q^\theta(\sigma_i(a_i), a_{-i}) u_i(\sigma_i(a_i), a_{-i}, m_{a,\theta}, \theta). \end{aligned}$$

Furthermore, the MPC constraint can be rewritten in the (q, m) parametrization as

$$\sum_{a \in A} q^\theta(a) \max\{m_{a,\theta} - r, 0\} \geq \int_r F(x) dx$$

for all $r \in [0, 1]$ and with equality at 0. As both of the objective and the constraint are continuous in (q, m) it follows that the principal maximizes a continuous function over a compact subset of $[0, 1]^{2^{|\Theta||A|}}$ and hence a maximizer exists.

We next argue existence of a maximizer for the single agent case with an arbitrary action set. As argued in Lemma 6 the set of feasible distributions G^θ is sequentially compact. As the product of finitely many sequentially compact spaces is also sequentially compact the set of vectors $(G^\theta)_{\theta \in \Theta}$ is also sequentially compact. As \bar{u}_i is continuous it follows that the IC constraint (7) is continuous in G . As $G \mapsto \sum_{\theta \in \Theta} \phi(\theta) \int_\Omega \bar{v}(s, \theta) dG^\theta(s)$ is upper semicontinuous it follows that the designer maximizes an upper hemicontinuous linear function over a compact convex set. By Bauer's maximum principle a maximizer exists. \square

Lemma 6. *Suppose $u_i : [0, 1] \rightarrow \mathbb{R}$ is a continuous function for $i \in \{1, \dots, n\}$. The set of distributions $G : [0, 1] \rightarrow [0, 1]$ that satisfy $G \succeq F$ and*

$$\int_\Omega u_i(s) dG(s) \geq 0 \text{ for } i \in \{1, \dots, n\} \quad (24)$$

is compact in the weak topology.

Proof. First, note that as u_i is continuous it is bounded on $[0, 1]$. Consider a sequence of

³⁷Note here that $q^\theta(a)$ is the probability of the action profile a given the type profile θ *not* conditioning on the state.

distributions G^k , $k \in \{1, 2, \dots\}$ that satisfy the constraints in (24). By Helly's selection theorem there exists a subsequence that converges pointwise. From now on assume that (G^k) is such a subsequence and denote by G^∞ the right-continuous representation of its point-wise limit. Thus, any sequence of random variables m^k such that $m^k \sim G^k$ converges in distribution to a random variable distributed according to G^∞ .

As (u_i) are continuous and bounded this implies that for all i we have

$$\lim_{k \rightarrow \infty} \int_{\Omega} u_i(s) dG^k(s) = \int_{\Omega} u_i(s) dG^\infty(s).$$

Furthermore, for all $x \in [0, 1]$

$$\lim_{k \rightarrow \infty} \int_x^1 G^k(s) ds = \int_x^1 G^\infty(s) ds,$$

and hence G^∞ also satisfies $G^\infty \succeq F$. We have hence established sequential compactness. As the topology of weak convergence is metrizable by the Prokhorov metric and in a metrizable space, a subset is compact if and only if it is sequentially compact the set of distributions given in the statement of the lemma is compact with respect to the weak topology. \square

Lemma 7. *Let $F, G : [0, 1] \rightarrow [0, 1]$ be CDFs and let F be continuous. Suppose that G is a mean-preserving contraction of F and for some $x \in [0, 1]$*

$$\int_x^1 F(s) ds = \int_x^1 G(s) ds.$$

Then $F(x) = G(x)$. Furthermore, G is continuous at x .

Proof. Define the function $L : [0, 1] \rightarrow \mathbb{R}$ as $L(z) = \int_z^1 F(s) - G(s) ds$. As G is a mean-preserving contraction of F we have that $L(z) \leq 0$ for all $z \in [0, 1]$. By the assumption of the lemma $L(x) = 0$. By definition L is absolutely continuous and has a weak derivative, which we denote by $L'(z) = G(z) - F(z)$. As F is continuous L' has only up-ward jumps and is right-continuous. For L to have a maximum at x we need that $\lim_{z \nearrow x} L'(z) \geq 0$ and $\lim_{z \searrow x} L'(z) \leq 0$. This implies that

$$\lim_{z \searrow x} G(z) - F(z) \leq 0 \leq \lim_{z \nearrow x} G(z) - F(z).$$

In turn, this implies that $\lim_{z \searrow x} G(z) \leq \lim_{z \nearrow x} G(z)$. As G is a CDF it is non-decreasing

and thus G is continuous at x . Consequently, L is continuously differentiable at x and as L admits a maximum at x , we have that $0 = L'(x) = G(x) - F(x)$. \square

Lemma 8. Fix an interval $[a, b] \subseteq [0, 1]$, $c \in \mathbb{R}$, upper semicontinuous $v : [0, 1] \rightarrow [0, 1]$ and continuous $\tilde{u}_1, \dots, \tilde{u}_n : [0, 1] \rightarrow \mathbb{R}$ and consider the problem

$$\max_{\tilde{G}} \int_{\Omega} v(s) d\tilde{G}(s) \quad (25)$$

$$\text{subject to } \int_{\Omega} \tilde{u}_i(s) d\tilde{G}(s) \geq 0 \text{ for } i \in \{1, \dots, n\} \quad (26)$$

$$\int_a^b G(s) ds = c \quad (27)$$

$$\int_{[a,b]} d\tilde{G}(s) = 1. \quad (28)$$

If the set of distributions that satisfy (26)-(28) is non-empty then there exists a solution to the above optimization problem that is supported on at most $n + 2$ points.

Proof. Consider the set of distributions that assign probability 1 to the set $[a, b]$. The extreme points of this set are the Dirac measures in $[a, b]$. Let \mathcal{D} be the set of distributions which satisfy (26)-(27) and are supported on $[a, b]$. By Theorem 2.1 in Winkler (1988) each extreme points of the set \mathcal{D} is the sum of at most $n + 2$ mass points as (26) and (27) specify $n + 1$ constraints. Note, that the set of the set of distributions satisfying (26)-(28) is compact. As v is upper semicontinuous the function $\tilde{G} \rightarrow \int_0^1 v(s) d\tilde{G}(s)$ is upper semi continuous and linear. Thus, by Bauer's maximum principle (see for example Result 7.69 in Aliprantis and Border 2013) there exist a maximizer at an extreme point of \mathcal{D} which establishes the result. \square

Lemma 9. Suppose that H, G are distribution that assign probability 1 to $[a, b]$. Let M be an absolutely continuous function such that $\int_x^b G(s) ds > M(x)$ for all $x \in [a, b]$, and $\int_{\hat{x}}^b H(y) dy < M(\hat{x})$ for some $\hat{x} \in [a, b]$. Then, there exists $\lambda \in (0, 1)$ such that for all $x \in [a, b]$

$$\int_x^b (1 - \lambda)G(s) + \lambda H(s) ds \geq M(x)$$

with equality for some $x \in [a, b]$.

Proof. Define

$$L_{\lambda}(x) = \int_x^b (1 - \lambda)G(y) + \lambda H(y) dy - M(x),$$

and $\phi(\lambda) = \min_{z \in [a,b]} L_\lambda(z)$. As M is continuous, by the assumptions of the lemma we have that

$$\phi(0) = \min_{x \in [a,b]} L_0(x) = \min_{x \in [a,b]} \left[\int_x^b G(s) ds - M(x) \right] > 0$$

and

$$\phi(1) = \min_{x \in [a,b]} L_1(x) = \min_{x \in [a,b]} \left[\int_x^b H(s) ds - M(x) \right] \leq \int_{\hat{x}}^b H(s) ds - M(\hat{x}) < 0.$$

Furthermore,

$$\left| \frac{\partial L_\lambda(z)}{\partial \lambda} \right| = \left| \int_x^b H(s) - G(s) ds \right| \leq b - a.$$

Hence, $\lambda \mapsto L_\lambda(z)$ is uniformly Lipschitz continuous and the envelope theorem thus implies that ϕ is Lipschitz continuous. As $\phi(0) > 0$, and $\phi(1) < 0$ there exist some $\lambda^* \in (0, 1)$ such that $\phi(\lambda^*) = 0$. This implies that for all $x \in [a, b]$

$$\int_x^b (1 - \lambda^*)G(s) + \lambda^*H(s) ds \geq M(x)$$

with equality for some $x \in [a, b]$. This completes the proof. \square

Lemma 10. *For a solution G to problem (14) denote the set of points where the mean preserving contraction constraint is binding by*

$$B_G = \left\{ z \in [0, 1] : \int_z^1 F(s) ds = \int_z^1 G(s) ds \right\}. \quad (29)$$

There exists a solution to (14) where the set B_G is maximal in set inclusion sense.

Proof. As the set of feasible distributions is compact with respect to the weak topology by Lemma 6 and the function $G \mapsto \int_0^1 v(s) dG(s)$ is upper semicontinuous in the weak topology the optimization problem (14) admits a solution. Observe that for any solution G , (29) implies that B_G is a closed set.

Let \mathcal{G}' denote the set of all solutions to (14). Denote by \mathcal{G} the subset of \mathcal{G}' such that (i) $\{B_G | G \in \mathcal{G}\} = \{B_G | G \in \mathcal{G}'\}$ and $B_{G_1} \neq B_{G_2}$ for any $G_1, G_2 \in \mathcal{G}'$ such that $G_1 \neq G_2$ (the existence of such \mathcal{G} follows from the axiom of choice). Define a partial order \succeq on \mathcal{G} : $G_1 \succeq G_2$ if $B_{G_1} \supseteq B_{G_2}$. Consider a totally ordered subset \mathcal{G}_c of \mathcal{G} .

Let $B = \cup_{G \in \mathcal{G}_c} B_G$ and denote by \bar{B} the closure of this set. Since \bar{B} is closed and bounded, it is compact. Similarly B_G is compact for each $G \in \mathcal{G}$. For $r \in \mathbb{N}_+$, consider a

solution $G_r \in \mathcal{G}_c$ such that

$$\max_{y \in \bar{B}} \min_{x \in B_{G_r}} |x - y| < 1/r,$$

where the optima are achieved due to compactness, the continuity of the argument being optimized (and the theorem of maximum). Existence of such a solution follows from the definition of B and the fact that \mathcal{G}_c is totally ordered. By Helly's selection theorem, the sequence $\{G_r\}$ has a convergent subsequence. Let G_∞ denote its limit. Since $G \mapsto \int_0^1 v(s) dG(s)$ is upper semicontinuous in the weak topology, it follows that G_∞ also solves (14). Furthermore, by construction, B_{G_∞} is dense in \bar{B} . Since B_{G_∞} is also closed, it follows that it follows that $B_{G_\infty} = \bar{B}$. This implies that $G_\infty \succeq G$ for every $G \in \mathcal{G}_c$. Zorn's lemma (see for example Section 1.12 in Aliprantis and Border 2013) implies that \mathcal{G} has a maximal element, G^* . By the definition of our partial order this implies that $B_{G^*} \supseteq B_G$ for every $G \in \mathcal{G}$ and the claim follows. \square

Proof of Proposition 3. The first part of the claim follows from Lemma 10. The lemma also implies that there exists a solution G for which the set B_G defined in (29) is maximal in set inclusion sense. Consider such a solution.

Fix a point $x \notin B_G$. We define (a, b) to be the largest interval such that the mean-preserving contraction constraint does not bind on that interval for the solution G , i.e.

$$a = \max \left\{ z \leq x : z \in B_G \right\} \quad b = \min \left\{ z \geq x : z \in B_G \right\}.$$

If G assigns probability zero to the interval $[a, b]$ there are 0 mass-points in the interval and we have thus established that there are less than $n + 2$ mass-points in that interval. Thus, assume for the rest of the proof that G assigns strictly positive mass to $[a, b]$. By Lemma 7 G assigns no mass to a or b and hence G also assigns strictly positive mass to the interior of $[a, b]$. Consider now an interval $[\hat{a}, \hat{b}] \subset (a, b)$ such that G assigns strictly positive mass to $[\hat{a}, \hat{b}]$. We define $G_{[\hat{a}, \hat{b}]} : [0, 1] \rightarrow [0, 1]$ to be the CDF of a random variable that is distributed according to G conditional on the realization being in the interval $[\hat{a}, \hat{b}]$

$$G_{[\hat{a}, \hat{b}]}(z) = \frac{G(z) - G(\hat{a}_-)}{G(\hat{b}) - G(\hat{a}_-)},$$

where $G(\hat{a}_-) = \lim_{s \nearrow \hat{a}} G(s)$. We note that $G_{[\hat{a}, \hat{b}]}$ is non-decreasing, right-continuous, and satisfies $G_{[\hat{a}, \hat{b}]}(\hat{b}) = 1$. Thus, it is a well defined CDF supported on $[\hat{a}, \hat{b}]$. As G is feasible

we get that

$$\int_{\hat{a}}^{\hat{b}} u_k(s) dG_{[\hat{a}, \hat{b}]}(s) + \frac{1}{G(\hat{b}) - G(\hat{a}_-)} \int_{\Omega \setminus [\hat{a}, \hat{b}]} u_k(s) dG(s) \geq 0 \quad \text{for } k \in \{1, \dots, n\}. \quad (30)$$

To simplify notation we define the functions $\tilde{u}_1, \dots, \tilde{u}_n$, where for all k

$$\tilde{u}_k(z) = u_k(z) + \frac{1}{G(\hat{b}) - G(\hat{a}_-)} \int_{\Omega \setminus [\hat{a}, \hat{b}]} u_k(y) dG(y). \quad (31)$$

Note that using this notation (30) can be restated as:

$$\int_{\Omega} \tilde{u}_k(s) dG_{[\hat{a}, \hat{b}]}(s) \geq 0 \quad \text{for } k \in \{1, \dots, n\}. \quad (32)$$

As G satisfies the mean-preserving contraction constraint relative to F , using the fact that $a < \hat{a}$ and $\hat{b} < b$, for $z \in [\hat{a}, \hat{b}]$ we obtain:

$$\int_z^{\hat{b}} G_{[\hat{a}, \hat{b}]}(s) ds > \frac{1}{G(\hat{b}) - G(\hat{a}_-)} \left[\int_z^1 F(s) ds - \int_{\hat{b}}^1 G(s) ds - (\hat{b} - z)G(\hat{a}_-) \right] = M(z). \quad (33)$$

Consider now the maximization problem over distributions supported on $[\hat{a}, \hat{b}]$ that satisfy the constraints derived above (after replacing the strict inequality in (33) with a weak inequality) and maximize the original objective:

$$\begin{aligned} & \max_H \int_{\Omega} v(s) dH(s) \\ & \text{subject to } \int_{\Omega} \tilde{u}_i(s) dH(s) \geq 0 \quad \text{for } i \in \{1, \dots, n\} \\ & \int_z^{\hat{b}} H(s) ds \geq M(z) \quad \text{for } z \in [\hat{a}, \hat{b}] \\ & \int_{[\hat{a}, \hat{b}]} dH(s) = 1. \end{aligned} \quad (34)$$

By (32) and (33) the conditional CDF $G_{[\hat{a}, \hat{b}]}$ is feasible in the problem above. We claim that it is also optimal. Suppose, towards a contradiction, that there exist a CDF H that is

feasible in (34) and achieves a strictly higher value than $G_{[\hat{a}, \hat{b}]}$. Consider the CDF

$$K(z) = \begin{cases} G(z) & \text{if } z \in [0, 1] \setminus [\hat{a}, \hat{b}] \\ G(\hat{a}_-) + H(z)(G(\hat{b}) - G(\hat{a}_-)) & \text{if } z \in [\hat{a}, \hat{b}], \end{cases}$$

which equals G outside the interval $[\hat{a}, \hat{b}]$ and H conditional on being in $[\hat{a}, \hat{b}]$. Using (31), the definition of $M(z)$, and the feasibility of H in (34), it can be readily verified that this CDF is feasible in the original problem (14). Moreover, it achieves a higher value than G , since H achieves strictly higher value than $G_{[\hat{a}, \hat{b}]}$ in (34). However, this leads to a contradiction to the optimality of G in (14), thereby implying that $G_{[\hat{a}, \hat{b}]}$ is optimal in (34).

Next, we establish that there cannot exist an optimal solution H to the problem (34) where for some $z \in (\hat{a}, \hat{b})$

$$\int_z^{\hat{b}} H(s) ds = M(z). \quad (35)$$

Suppose such an optimal solution exists. Then, K would be an optimal solution to the original problem satisfying $z \in B_K \supset B_G$, where B_K is defined as in (29) (after replacing G with K) and is the set of points where the mean preserving contraction constraint binds. However, this contradicts that G is a solution to the original problem that is maximal (in terms of the set where the MPC constraints bind).

We next consider a relaxed version of the optimization problem (34) where we replace the second constraint of (34) with a constraint that ensures that H has the same mean as $G_{[\hat{a}, \hat{b}]}$:

$$\begin{aligned} & \max_H \int_{\Omega} v(s) dH(s) \\ \text{subject to } & \int_{\Omega} \tilde{u}_i(s) dH(s) \geq 0 \quad \text{for } i \in \{1, \dots, n\} \\ & \int_{\hat{a}}^{\hat{b}} H(s) ds = \int_{\hat{a}}^{\hat{b}} G_{[\hat{a}, \hat{b}]}(s) ds \\ & \int_{[\hat{a}, \hat{b}]} dH(s) = 1. \end{aligned}$$

By Lemma 8 there exists a solution J to this relaxed problem that is the sum of $n + 2$ mass

points. Since $G_{[\hat{a}, \hat{b}]}$ is feasible in this problem, it readily follows that

$$\int_{\Omega} v(s) dJ(s) \geq \int_{\Omega} v(s) dG_{[\hat{a}, \hat{b}]}(s). \quad (36)$$

Suppose, towards a contradiction, that there exists $z \in [\hat{a}, \hat{b}]$ such that

$$\int_z^{\hat{b}} J(s) ds < M(z). \quad (37)$$

Then, by Lemma 9, there exists some $\lambda \in (0, 1)$ such that $(1 - \lambda)G_{[\hat{a}, \hat{b}]} + \lambda J$ satisfies

$$\int_r^{\hat{b}} (1 - \lambda)G_{[\hat{a}, \hat{b}]}(s) + \lambda J(s) ds \geq M(r), \quad (38)$$

for all $r \in [\hat{a}, \hat{b}]$, and the inequality holds with equality for some $r \in [\hat{a}, \hat{b}]$. This implies that $(1 - \lambda)G_{[\hat{a}, \hat{b}]} + \lambda J$ is feasible for the problem (34). Furthermore, by the linearity of the objective, (36), and the optimality of $G_{[\hat{a}, \hat{b}]}$ in (34), it follows that $(1 - \lambda)G_{[\hat{a}, \hat{b}]} + \lambda J$ is also optimal in (34). However, this leads to a contradiction to the fact that (34) does not admit an optimal solution where the equality in (35) holds for some $z \in [\hat{a}, \hat{b}] \subset [a, b]$.

Consequently, the inequality (37) cannot hold, and J must be feasible in problem (34). Together with (36) this implies that J is an optimal solution to (34) that assigns mass to only $n + 2$ points in the interval $[\hat{a}, \hat{b}]$. This implies that the CDF

$$\begin{cases} G(z) & \text{if } z \in [0, 1] \setminus [\hat{a}, \hat{b}] \\ G(\hat{a}_-) + J(z)(G(\hat{b}) - G(\hat{a}_-)) & \text{if } z \in [\hat{a}, \hat{b}] \end{cases} \quad (39)$$

is a solution of the original problem that assigns mass to only $n + 2$ points in the interval $[\hat{a}, \hat{b}]$. By setting $\hat{a} = a + \frac{1}{r}$ and $\hat{b} = b - \frac{1}{r}$ we can thus find a sequence of solutions (H^r) to (14) that each have at most $n + 2$ mass points in the interval $[a + \frac{1}{r}, b - \frac{1}{r}]$. As the set of feasible distributions is closed and the objective function is upper semicontinuous this sequence admits a limit point H^∞ which itself is optimal in (14). This limit distribution consists of at most $n + 2$ mass points in the interval (a, b) . Furthermore, by definition of a, b and our construction in (39) each solution H^r and hence H^∞ satisfies the MPC constraint with equality at $\{a, b\}$. Thus, Lemma 7 implies that H^∞ is continuous at these points, and $H^\infty(a) = F(a)$ and $H^\infty(b) = F(b)$.

Hence, we have established that for every solution G for which B_G is maximal, either $x \in B_G$ which by Lemma 7 implies that $G(x) = F(x)$. Or $x \notin B_G$ and then one can find a new solution \tilde{G} such that (i) \tilde{G} has at most $n + 2$ mass points in the interval (a, b) with $a = \max\{z \leq x : z \in B_G\}$ and $b = \min\{z \geq x : z \in B_G\}$, (ii) $\tilde{G}(a) = F(a)$ and $\tilde{G}(b) = F(b)$ which implies that the mass inside the interval $[a, b]$ is preserved, and (iii) \tilde{G} matches G outside (a, b) . Since every interval contains a rational number there can be at most countably many such intervals. Proceeding inductively, the claim follows. \square

To establish Proposition 4, we make use of the partition lemma, stated next:

Lemma 11 (Partition Lemma). *Suppose that distributions F, G are such that $\int_x^1 G(t)dt \geq \int_x^1 F(t)dt$ for $x \in I = [a, b]$, where the inequality holds with equality only for the end points of I . Suppose further that $G(a) = F(a)$, $G(x) = G(a) + \sum_{r=1}^K p_r \mathbf{1}_{x \leq m_r}$ for $x \in I$ where $\sum_{r=1}^K p_r = F(b) - F(a)$, (m_r) is a (weakly) increasing in r , and $m_r \in I$ for $r \in [K] \equiv \{1, \dots, K\}$.*

There exists a collection of intervals $\{J_r\}_{r \in [K]}$ such that $\{P_k\} = \{J_k \setminus \cup_{\ell \in A | \ell > k} J_\ell\}$ is a laminar partition, which satisfies:

- (a) $J_1 = I$, and if $K > 1$, then $F(\inf J_1) < F(\inf J_K) < F(\sup J_K) < F(\sup J_1)$;
- (b) $\int_{P_k} dF(x) = p_k$ for all $k \in [K]$;
- (c) $\int_{P_k} x dF(x) = p_k m_k$ for all $k \in [K]$.

Proof of Lemma 11. We prove the claim by induction on K . Note that when $K = 1$ we have $J_1 = P_1 = I$, which readily implies properties (a) and (b). In addition, the definition of p_1, m_1 implies that

$$\begin{aligned} G(b)b - G(a)a - p_1 m_1 &= G(a)(b - a) + p_1(b - m_1) = \int_a^b G(t)dt = \int_a^b F(t)dt \\ &= F(b)b - F(a)a - \int_I t dF(t) = G(b)b - G(a)a - \int_{P_1} t dF(t). \end{aligned} \tag{40}$$

Hence, property (c) also follows.

We proceed by considering two cases: $K = 2, K > 2$.

$K = 2$: Let $t_1, t_2 \in I$ be such that $F(t_1) - F(a) = F(b) - F(t_2) = p_1$. Observe that since $\int_x^1 G(t)dt \geq \int_x^1 F(t)dt$ $x \in I$ and this inequality holds with equality only at the

end points of I , we have (i) $\int_a^{t_1} F(x)dx > \int_a^{t_1} G(x)dx$ and (ii) $\int_{t_2}^b F(x)dx < \int_{t_2}^b G(x)dx$. Using the first inequality and the definition of G we obtain:

$$\begin{aligned} p_1(t_1 - m_1)^+ + G(a)(t_1 - a) &\leq \int_a^{t_1} G(x)dx < \int_a^{t_1} F(x)dx \\ &= F(t_1)t_1 - F(a)a - \int_a^{t_1} x dF(x) = (G(a) + p_1)t_1 - G(a)a - \int_a^{t_1} x dF(x). \end{aligned} \quad (41)$$

Rearranging the terms, this yields

$$p_1 m_1 \geq p_1 t_1 - p_1(t_1 - m_1)^+ > \int_a^{t_1} x dF(x). \quad (42)$$

Similarly, using (ii) and the definition of G we obtain:

$$\begin{aligned} G(b)(b - t_2) - p_1(m_1 - t_2)^+ &\geq \int_{t_2}^b G(x)dx > \int_{t_2}^b F(x)dx \\ &= F(b)b - F(t_2)t_2 - \int_{t_2}^b x dF(x) = G(b)b - (G(b) - p_1)t_2 - \int_{t_2}^b x dF(x). \end{aligned} \quad (43)$$

Rearranging the terms, this yields

$$p_1 m_1 \leq p_1 t_2 + p_1(m_1 - t_2)^+ < \int_{t_2}^b x dF(x). \quad (44)$$

Combining (42) and (44), and the fact that $F(t_1) - F(a) = F(b) - F(t_2) = p_1$ implies that there exist $\hat{t}_1, \hat{t}_2 \in \text{int}(I)$ satisfying $F(a) < F(\hat{t}_1) < F(\hat{t}_2) < F(b)$ such that $F(\hat{t}_1) - F(a) + F(b) - F(\hat{t}_2) = p_1$ and

$$\int_a^{\hat{t}_1} x dF(x) + \int_{\hat{t}_2}^b x dF(x) = p_1 m_1. \quad (45)$$

Note that

$$\begin{aligned} (b - a)G(a) + (b - m_1)p_1 + (b - m_2)p_2 &= \int_a^b G(x)dx = \int_a^b F(x)dx \\ &= bF(b) - aF(a) - \int_a^b x dF(x) = bG(b) - aG(a) - \int_a^b x dF(x). \end{aligned}$$

Since $p_1 + p_2 = G(b) - G(a)$, this in turn implies that $\int_a^b x dF(x) = p_1 m_1 + p_2 m_2$.

Combining this observation with (45), we conclude that

$$\int_{\hat{t}_1}^{\hat{t}_2} x dF(x) = p_2 m_2. \quad (46)$$

Let $J_2 = [\hat{t}_1, \hat{t}_2]$, and $J_1 = I$, and define P_1, P_2 as in the statement of the lemma. Observe that this construction immediately satisfies (a) and (b). Moreover, (c) also follows from (45) and (46). Thus, the claim holds when $K = 2$.

$K > 2$: Suppose that $K > 2$, and that the induction hypothesis holds for any $K' \leq K - 1$. Let $\hat{p}_2 = p_K$, $\hat{m}_2 = m_K$; and $\hat{p}_1 = \sum_{k \in [K-1]} p_k$, $\hat{m}_1 = \frac{1}{\hat{p}_1} \sum_{k \in [K-1]} p_k m_k$. Define a distribution \hat{G} such that $\hat{G}(x) = G(x)$ for $x \notin I$, $\hat{G}(a) = F(a)$, and $\hat{G}(x) = \hat{G}(a) + \sum_{r=1}^2 \hat{p}_r \mathbf{1}_{x \leq \hat{m}_r}$. This construction ensures that $\hat{p}_1 + \hat{p}_2 = F(b) - F(a)$ and $\hat{m}_2 > \hat{m}_1$. Moreover, \hat{G} is a mean preserving contraction of G , and hence $\int_x^1 \hat{G}(t) dt \geq \int_x^1 G(t) dt$. Since $\hat{G}(x) = G(x)$ for $x \notin I$, this in turn implies that $\int_x^1 \hat{G}(t) dt \geq \int_x^1 F(t) dt$ for $x \in I$ where the inequality holds with equality only for the end points of I . Thus, the assumptions of the lemma hold for \hat{G} and F , and using the induction hypothesis for $K' = 2$, we conclude that there exists intervals \hat{J}_1, \hat{J}_2 and sets $P_2 = \hat{J}_2$, $P_1 = \hat{J}_1 \setminus \hat{J}_2$, such that

$$(\hat{a}) \quad I = \hat{J}_1 \supset \hat{J}_2, \text{ and } F(\inf \hat{J}_1) < F(\inf \hat{J}_2) < F(\sup \hat{J}_2) < F(\sup \hat{J}_1);$$

$$(\hat{b}) \quad \int_{P_k} dF(x) = \hat{p}_k \text{ for } k \in \{1, 2\};$$

$$(\hat{c}) \quad \int_{P_k} x dF(x) = \hat{p}_k \hat{m}_k \text{ for all } k \in \{1, 2\}.$$

Note that (\hat{b}) and (\hat{c}) imply that $\hat{m}_2 \in \hat{J}_2$.

Denote by x_0, x_1 the end points of \hat{J}_2 and let $q_0 = F(x_0) > F(a)$, $q_1 = F(x_1) < F(b)$. Define a cumulative distribution function $F'(\cdot)$, such that

$$F'(x) = \begin{cases} F(x)/(1 - \hat{p}_2) & \text{for } x \leq x_0, \\ F(x_0)/(1 - \hat{p}_2) & \text{for } x_0 < x < x_1, \\ (F(x) - \hat{p}_2)/(1 - \hat{p}_2) & \text{for } x_1 \leq x. \end{cases} \quad (47)$$

Set $p'_k = p_k/(1 - \hat{p}_2)$ and $m'_k = m_k$ for $k \in [K - 1]$. Let distribution G' be such that $G'(x) = G(x)/(1 - \hat{p}_2)$ for $x \notin I$, and $G'(x) = G'(a) + \sum_{r \in [K-1]} p'_r \mathbf{1}_{x \leq m'_r}$. Observe that by construction $G'(a) = F'(a)$, $\sum_{r \in [K-1]} p'_r = F'(b) - F'(a)$, and $\{m'_r\}$ is increasing in r , where $m'_r \in I$, $m'_r \leq \hat{m}_2$ for $r \in [K - 1]$. The following lemma implies that G' and F' also satisfy the MPC constraints over I :

Lemma 12. $\int_x^1 G'(t)dt \geq \int_x^1 F'(t)dt$ for $x \in I$, where the inequality holds with equality only for the end points of I .

Proof. The definition of G' implies that it can alternatively be expressed as follows:

$$G'(x) = \begin{cases} G(x)/(1 - \hat{p}_2) & \text{for } x < \hat{m}_2, \\ (G(x) - \hat{p}_2)/(1 - \hat{p}_2) & \text{for } x \geq \hat{m}_2. \end{cases} \quad (48)$$

Since $\int_b^1 G(t)dt = \int_b^1 F(t)dt$, (47) and (48) readily imply that $\int_b^1 G'(t)dt = \int_b^1 F'(t)dt$. Similarly, using these observations and (47) we have

$$\begin{aligned} (1 - \hat{p}_2) \int_a^1 F'(t)dt &= \int_a^1 F(t)dt - \int_{x_0}^{x_1} F(t)dt + F(x_0)(x_1 - x_0) - \hat{p}_2(1 - x_1) \\ &= \int_a^1 F(t)dt - F(x_1)x_1 + F(x_0)x_0 + \hat{p}_2\hat{m}_2 + F(x_0)(x_1 - x_0) - \hat{p}_2(1 - x_1) \\ &= \int_a^1 G(t)dt - \hat{p}_2(1 - \hat{m}_2) \end{aligned} \quad (49)$$

Here, the second line rewrites $\int_{x_0}^{x_1} F(t)dt$ using integration by parts, and leverages (\hat{c}) . The third line uses the fact that $\hat{p}_2 = F(x_1) - F(x_0)$ and $\int_a^1 G(t)dt = \int_a^1 F(t)dt$. On the other hand, (48) readily implies that:

$$(1 - \hat{p}_2) \int_a^1 G'(t)dt = \int_a^1 G(t)dt - \hat{p}_2(1 - \hat{m}_2) \quad (50)$$

Together with (49), this equation implies that $\int_a^1 G'(t)dt = \int_a^1 F'(t)dt$. Thus, the inequality in the claim holds with equality for the end points of I .

Recall that $\hat{m}_2 \in \hat{I}_2$ and hence $a < x_0 \leq \hat{m}_2 = m_K \leq x_1 < b$. We complete the proof by focusing on the value x takes in the following cases: (i) $a < x \leq x_0$, (ii) $x_0 \leq x \leq \hat{m}_2$, (iii) $\hat{m}_2 \leq x \leq x_1$, (iv) $x_1 \leq x < b$.

Case (i): Using the observations $\int_x^1 G(t)dt > \int_x^1 F(t)dt$ and $\int_a^1 G(t)dt = \int_a^1 F(t)dt$ together with (47) and (48) yields

$$\int_a^x G'(t)dt = \frac{1}{1 - \hat{p}_2} \int_a^x G(t)dt < \frac{1}{1 - \hat{p}_2} \int_a^x F(t)dt = \int_a^x F'(t)dt. \quad (51)$$

Together with $\int_a^1 G'(t)dt = \int_a^1 F'(t)dt$ this implies that $\int_x^1 G'(t)dt > \int_x^1 F'(t)dt$ in case

(i).

Case (ii): Using (47) and (48) we obtain:

$$(1 - \hat{p}_2) \int_x^1 G'(t) - F'(t) dt = \int_x^1 G(t) dt - (1 - \hat{m}_2) \hat{p}_2 - \int_{x_1}^1 F(t) dt - \int_x^{x_1} F(x_0) dt + (1 - x_1) \hat{p}_2.$$

Since G is an increasing function, it can be seen that the right hand side is a concave function of x . Thus, for $x \in [x_0, \hat{m}_2]$ this expression is minimized for $x = x_0$ or $x = \hat{m}_2$. For $x = x_0$, case (i) implies that the expression is non-negative. We next argue that for $x = \hat{m}_2$ the expression remains non-negative. This in turn implies that $\int_x^1 G'(t) - F'(t) dt \geq 0$ for $x \in [x_0, \hat{m}_2]$, as claimed.

Setting $x = \hat{m}_2$, recalling that $\int_b^1 G(t) dt = \int_b^1 F(t) dt$, and observing that $G(t) = F(t)$ for $t \in [\hat{m}_2, b]$ the right hand side of the previous equation reduces to:

$$\begin{aligned} R &:= (b - \hat{m}_2)F(b) - (1 - \hat{m}_2)\hat{p}_2 - \int_{x_1}^b F(t) dt - (x_1 - \hat{m}_2)F(x_0) + (1 - x_1)\hat{p}_2 \\ &= (b - \hat{m}_2)F(b) - \int_{x_1}^b F(t) dt - (x_1 - \hat{m}_2)F(x_0) - (x_1 - \hat{m}_2)\hat{p}_2 \\ &= (b - x_1)F(b) - \int_{x_1}^b F(t) dt + (x_1 - \hat{m}_2)(F(b) - F(x_0) - \hat{p}_2). \end{aligned} \quad (52)$$

Since $F(b) \geq F(x_1) = \hat{p}_2 + F(x_0)$, we conclude:

$$R \geq (b - x_1)F(b) - \int_{x_1}^b F(t) dt \geq 0, \quad (53)$$

where the last inequality applies since F is weakly increasing. Thus, we conclude that $\int_{\hat{m}_2}^1 G'(t) - F'(t) dt \geq 0$, and the claim follows.

Case (iii): First observe that (47) and (48) imply that

$$(1 - \hat{p}_2) \int_x^1 G'(t) - F'(t) dt = \int_x^1 G(t) dt - (1 - x)\hat{p}_2 - \int_{x_1}^1 F(t) dt - \int_x^{x_1} F(x_0) dt + (1 - x_1)\hat{p}_2.$$

Similar to case (ii), the right hand side is a concave function of x . Thus, for $x \in [\hat{m}_2, x_1]$ this expression is minimized for $x = \hat{m}_2$ or $x = x_1$. When $x = \hat{m}_2$, case (ii) implies that $\int_x^1 G'(t) - F'(t) dt \geq 0$. Similarly, when $x = x_1$, case (iv) implies that $\int_x^1 G'(t) - F'(t) dt \geq 0$. Thus, it follows that $\int_x^1 G'(t) - F'(t) dt \geq 0$ for all $x \in [\hat{m}_2, x_1]$.

Case (iv): In this case, (47) and (48) readily imply that

$$(1 - \hat{p}_2) \int_x^1 G'(t) - F'(t) dt = \int_x^1 G(t) - F(t) dt > 0,$$

where the inequality follows from our assumptions on F and G . \square

Summarizing, we have established that the distribution G' and F' satisfy the conditions of the lemma. By the induction hypothesis, we have that there exists intervals $\{J'_k\}_{k \in [K-1]}$ and sets $P'_k = J'_k \setminus \cup_{\ell \in [K-1], \ell > k} J'_\ell$ for all $k \in \mathcal{A}'$ such that:

$$(a') \quad J'_1 = I, \text{ and } F(\inf J'_1) < F(\inf J'_{K-1}) < F(\sup J'_{K-1}) < F(\sup J'_1);$$

$$(b') \quad \int_{P'_k} dF'(x) = p'_k \text{ for all } k \in [K-1];$$

$$(c') \quad \int_{P'_k} x dF'(x) = p'_k m'_k \text{ for all } k \in [K-1].$$

Let $J_k = J'_k \setminus \hat{J}_2$ for $k \in [K-1]$ such that $\hat{J}_2 \not\subseteq J'_k$, and $J_k = J'_k$ for the remaining $k \in [K-1]$. Define $J_K = \hat{J}_2 = [x_0, x_1]$. For $k \in [K]$, let $P_k = J_k \setminus \cup_{\ell \in [K], \ell > k} J_\ell$. Note that the definition of the collection $\{P_k\}_{k \in [K]}$ implies that it constitutes a laminar partition of I . Observe that the construction of $\{J_k\}_{k \in [K]}$ and (\hat{a}) , (a') imply that these intervals also satisfy condition (a) of the lemma. Note that by construction we have

$$P_k \subseteq P'_k \subseteq P_k \cup J_K \quad \text{and} \quad P_k \cap J_K = \emptyset \quad \text{for } k \in [K-1]. \quad (54)$$

Since $\int_{J_K} dF'(t) = 0$ by (47) this observation implies that $\int_{P'_k} dF'(t) = \int_{P_k} dF'(t)$ for $k \in [K-1]$.

Using (47), (b') , and (54), this observation implies that

$$\int_{P_k} dF(t) = \int_{P_k} dF'(t)(1 - \hat{p}_2) = \int_{P'_k} dF'(t)(1 - \hat{p}_2) = p'_k(1 - \hat{p}_2) = p_k,$$

for $k \in [K-1]$. Similarly, by (\hat{b}) we have $\int_{P_K} dF(t) = \int_{\hat{P}_2} dF(t) = \hat{p}_2 = p_K$.

Finally, observe that by (\hat{c}) we have $\int_{P_K} t dF(t) = \int_{\hat{P}_2} t dF(t) = \hat{p}_2 \hat{m}_2 = p_K m_K$. Similarly, (47) and (54) imply that for $k \in [K-1]$, we have

$$\int_{P_k} t dF(t) = (1 - \hat{p}_2) \int_{P_k} t dF'(t) = (1 - \hat{p}_2) \int_{P'_k} t dF'(t) = (1 - \hat{p}_2) p'_k m'_k = p_k m_k.$$

These observations imply that the constructed $\{J_k\}_{k \in [K]}$ and $\{P_k\}_{k \in [K]}$ satisfy the induction hypotheses (a)–(c) for K . Thus, the claim follows by induction. \square

Proof of Proposition 4. By definition, the interval I_j in the statement of Proposition 4 satisfies the conditions of Lemma 11, (after setting $a = a_j$, $b = b_j$). The lemma defines auxiliary intervals $\{J_r\}$ and explicitly constructs a laminar partition that satisfies conditions (a)–(c). Here, conditions (b) and (c) readily imply that the constructed laminar partition satisfies the claim in Proposition 4, concluding the proof. \square

Proof of Theorem 1. The existence of an optimal mechanism follows from standard compactness arguments and is proven in Lemma 5 in the appendix. Consider now an arbitrary optimal SGR mechanism. Fix a type profile θ . By combining Lemma 2 and 3 we can replace G^θ by another solution to the respective optimization problem under MPC and linear side constraints and obtain a new SGR mechanism. By Propositions 3 and 4 there always exists a solution to this optimization problem under MPC and linear side constraints that can be implemented by a laminar partitional signal. Iterating this process over type profiles we get that there exists an optimal SGR mechanism in which each distribution G^θ can be implemented by a laminar partitional signal. Thus, we constructed an optimal laminar partitional mechanism. \square

Proof of Proposition 2. Part (i) of the proposition follows from Proposition 4 by counting the number of linear side constraints in the optimization problem stated in Lemma 2. Similarly, by counting the side constraints in the optimization problem stated in Lemma 3 it follows that for every optimal SGR mechanism (G^*, q^*) and every type profile θ there exists a laminar partitional signal that generates $G^{*,\theta}$. This proves the result if the action profile distribution $q^{*,\theta}(\cdot|m)$ is degenerate and deterministically recommends an action profile for each posterior mean m and type profile θ . We refer to SGR mechanisms associated with such action profile distributions as non-randomized SGR mechanisms (since for a given type profile and posterior mean, their recommendation is deterministic). In the multi-agent case, in contrast to the single agent case, non-randomized SGR mechanisms need not be optimal and the designer may need to use non-degenerate distributions of recommended action profiles.

To show the result when optimal SGR mechanisms require non-degenerate action profile distributions, we consider the parametrization of SGR mechanisms in terms of the mean $m_{a,\theta}$ and the *unconditional* probability of each action profile $q^\theta(a)$ (see Section 3.1). In this parametrization the set of non-randomized SGR mechanisms is dense in the space of

all SGR mechanisms (as for each randomized mechanism an arbitrarily small perturbation of all the means induces a non-randomized mechanism). The designer's payoff as well as the IC constraints are continuous in this parametrization (see the proof of Lemma 5), and hence there exists a sequence of non-randomized SGR mechanisms such that the payoff of the designer converges to the value of the optimal (randomized) SGR mechanism along the sequence. By the argument of Lemma 3.2 for each of these non-randomized SGR mechanisms there exists a laminar partitional mechanism with weakly larger payoff and partition depth of at most $\sum_{i \in N} |A_i|^2 |\Theta_i| + 2$. Since there are finitely many partial orders defining laminar partitions, there exist one that appears infinitely often along the sequence. Since for a given partial order the laminar partition is defined in terms of the end points of the convex hulls of the partition elements which belong to $[0, 1]^{|A|}$, there is a subsequence associated with this partial order that converges to a laminar partition consistent with the same partial order – which still has depth bounded by $\sum_{i \in N} |A_i|^2 |\Theta_i| + 2$. Moreover, the designer's payoff is continuous in the end points of the aforementioned intervals (since the distribution of the state is continuous). Thus, this limit point defines a new laminar partitional mechanism which achieves the optimal objective and has at most the depth stated in the claim. \square

Proof of Proposition 6. The first claim is immediate and follows as explained in the text. Here, we focus on the following example and use it to prove the second part of the claim.

Example 1. *There is a single agent, all types are equally likely, i.e., $\phi(\theta) \equiv 1/|\Theta|$ for all $\theta \in \Theta = \{1, \dots, n\}$, and the state is uniformly distributed in $[0, 1]$. For $k \in \{-2n, \dots, 2n\}$ we define intervals $B_{L,k} = [b_{L,k-1}, b_{L,k}]$, $B_{R,k} = [b_{R,k-1}, b_{R,k}]$. Here, for any integer k we let*

$$b_{L,k} = \frac{1}{4} + \frac{1}{8} \text{sgn}(k) \sqrt{\frac{|k|}{2n}} \quad b_{R,k} = \frac{3}{4} + \frac{1}{8} \text{sgn}(k) \sqrt{\frac{|k|}{2n}}.$$

All types of the agent share the same indirect utility function \bar{u} , such that $\bar{u}(m, \theta) = (m - \frac{1}{2})^2$ for all $m \in \{b_{L,k}, b_{R,k}\}$, and linearly interpolated in each $B_{L,k}$ and $B_{R,k}$ (in our construction the payoffs outside these intervals will be immaterial). The indirect utility of the designer is

$$\bar{v}(m, \theta) = \begin{cases} 1 & \text{if } m \in B_{L,2\theta} \cup B_{L,-2\theta} \cup B_{R,2n+2-2\theta} \cup B_{R,-2n-2+2\theta} \\ 0 & \text{otherwise.} \end{cases} \quad (55)$$

The agent's indirect utility functions can be generated by taking the set of actions to be $\{a_{L,k}, a_{R,k}\}$ for $k \in \{-2n - 1, \dots, +2n + 1\}$ and the utilities as a function of the action

to be

$$u(a_{.,k}, \omega, \theta) = c_{.,k-1}^2 + \frac{\omega - b_{.,k-1}}{b_{.,k} - b_{.,k-1}} (c_{.,k}^2 - c_{.,k-1}^2),$$

where $c_{.,k} = b_{.,k} - \frac{1}{2}$. Similarly, we let $v(a, \omega, \theta) = 1$ for actions $a_{L,2\theta}, a_{L,-2\theta}, a_{R,2n+2-2\theta}, a_{R,-2n-2+2\theta}$ and zero otherwise.

We begin by establishing that in the setting of Example 1 no public mechanism achieves more than $1/|\Theta|$. Note that by our construction in (55), for any posterior mean m the indirect utility of the designer equals 1 for at most a single type, i.e., $\sum_{\theta \in \Theta} \bar{v}(m, \theta) \leq 1$. As $\phi(\theta) = 1/|\Theta|$ this immediately implies that for any *type independent* distribution of the posterior mean G the designer can achieve a payoff of at most $1/|\Theta|$.

Next consider the following private mechanism: The distribution G^θ for type $\theta \in \Theta$ consists of 4 equally likely mass points at $b_{L,2\theta}, b_{L,-2\theta}, b_{R,2n+2-2\theta}, b_{R,-2n-2+2\theta}$. It is straightforward to see that the signal based on the partition $(\Pi_k)_{k=1}^4$ with $\Pi_1 = (b_{L,2\theta} - 1/8, b_{L,2\theta} + 1/8)$, $\Pi_2 = [0, 1/2] \setminus \Pi_1$, $\Pi_3 = [b_{R,2n+2-2\theta} - 1/8, b_{R,2n+2-2\theta} + 1/8]$, $\Pi_4 = (1/2, 1] \setminus \Pi_3$ induces the desired posterior mean distribution. At each of these beliefs the agent's indirect utility is given by $\bar{u}(m, \theta) = (m - \frac{1}{2})^2$. Thus, the benefit the agent of type θ' derives from observing the signal meant for type θ (relative to observing no signal) equals the variance of G^θ .

Note that the variance conditional on the posterior being less than $1/2$ equals $1/2(b_{L,2\theta} - 1/4)^2 + 1/2(b_{L,-2\theta} - 1/4)^2 = \frac{\theta}{64 \cdot n}$ and the variance conditional on the posterior being greater than $1/2$ equals $1/2(b_{R,2n+2-2\theta} - 3/4)^2 + 1/2(b_{R,-2n-2+2\theta} - 3/4)^2 = \frac{n+1-\theta}{64 \cdot n}$. By the law of the total variance the variance of G^θ thus equals $\frac{1}{2} \frac{\theta}{64 \cdot n} + \frac{1}{2} \frac{n+1-\theta}{64 \cdot n} + \frac{1}{2} \frac{1}{4} + \frac{1}{2} \frac{1}{4} = \frac{9n+1}{128 \cdot n}$. Since, this quantity does not depend on θ , we conclude that each type derives equal utility from any signal and the mechanism is incentive compatible. Each mean in the support G^θ persuades the agent to take an action that yields a payoff of 1 to the designer. Hence, this mechanism with private signals yields a payoff of 1. \square

References

- Aliprantis, Charalambos and Kim Border**, *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer-Verlag Berlin and Heidelberg GmbH & Company KG, 2013.
- Alonso, Ricardo and Odilon Câmara**, "Persuading voters," *American Economic Review*, 2016, 106 (11), 3590–3605.

- Arieli, Itai, Yakov Babichenko, Rann Smorodinsky, and Takuro Yamashita**, “Optimal Persuasion via Bi-Pooling,” *Available at SSRN*, 2020.
- Bergemann, Dirk and Stephen Morris**, “Robust predictions in games with incomplete information,” *Econometrica*, 2013, 81 (4), 1251–1308.
- **and** — , “Bayes correlated equilibrium and the comparison of information structures in games,” *Theoretical Economics*, 2016, 11 (2), 487–522.
- **and** — , “Information design: A unified perspective,” *Journal of Economic Literature*, 2019, 57 (1), 44–95.
- Blackwell, David**, “Comparison of experiments,” Technical Report, HOWARD UNIVERSITY Washington United States 1950.
- Blackwell, David A and Meyer A Girshick**, “Theory of Games and Statistical Decisions,” *John Willey and Sons, New York*, 1954.
- Boleslavsky, Raphael and Christopher Cotton**, “Grading standards and education quality,” *American Economic Journal: Microeconomics*, 2015, 7 (2), 248–79.
- Brocas, Isabelle and Juan D. Carrillo**, “Influence through ignorance,” *RAND Journal of Economics*, 2007, 38 (4), 931–947.
- Candogan, Ozan**, “Optimality of double intervals in persuasion: A convex programming framework,” *Available at SSRN 3452145*, 2019.
- , “Persuasion in Networks: Public Signals and k-Cores,” in “Proceedings of the 2019 ACM Conference on Economics and Computation” ACM 2019, pp. 133–134.
- **and Kimon Drakopoulos**, “Optimal Signaling of Content Accuracy: Engagement vs. Misinformation,” 2017.
- Dworczak, Piotr and Giorgio Martini**, “The simple economics of optimal persuasion,” *Journal of Political Economy*, 2019, 127 (5), 1993–2048.
- Gentzkow, Matthew and Emir Kamenica**, “Competition in persuasion,” *The Review of Economic Studies*, 2016, 84 (1), 300–322.
- **and** — , “A Rothschild–Stiglitz approach to Bayesian persuasion,” *American Economic Review*, 2016, 106 (5), 597–601.
- Goldstein, Itay and Yaron Leitner**, “Stress tests and information disclosure,” *Journal of Economic Theory*, 2018, 177, 34–69.
- Guo, Yingni and Eran Shmaya**, “The interval structure of optimal disclosure,” *Econometrica*, 2019, 87 (2), 653–675.
- , **Hao Li, and Xianwen Shi**, “Optimal discriminatory disclosure,” Technical Report, Working paper 2020.

- Inostroza, Nicolas and Alessandro Pavan**, “Persuasion in global games with application to stress testing,” 2018.
- Ivanov, Maxim**, “Optimal signals in Bayesian persuasion mechanisms with ranking,” *Work. Pap., McMaster Univ., Hamilton, Can*, 2015.
- Kamenica, Emir**, “Bayesian persuasion and information design,” *Annual Review of Economics*, 2019, *11*, 249–272.
- **and Matthew Gentzkow**, “Bayesian Persuasion,” *American Economic Review*, 2011, *101* (6), 2590–2615.
- Kleiner, Andreas, Benny Moldovanu, and Philipp Strack**, “Extreme Points and Majorization: Economic Applications,” *Available at SSRN*, 2020.
- Kolotilin, Anton**, “Optimal information disclosure: A linear programming approach,” *Theoretical Economics*, 2018, *13* (2), 607–635.
- **and Alexander Wolitzky**, “The Economics of Partisan Gerrymandering,” 2020.
- **and Andriy Zapechelnyuk**, “Persuasion meets delegation,” *arXiv preprint arXiv:1902.02628*, 2019.
- **, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li**, “Persuasion of a privately informed receiver,” *Econometrica*, 2017, *85* (6), 1949–1964.
- Onuchic, Paula and Debraj Ray**, “Conveying Value Via Categories,” *working paper*, 2020.
- Orlov, Dmitry, Pavel Zryumov, and Andrzej Skrzypacz**, “Design of macro-prudential stress tests,” 2018.
- Ostrovsky, Michael and Michael Schwarz**, “Information disclosure and unraveling in matching markets,” *American Economic Journal: Microeconomics*, 2010, *2* (2), 34–63.
- Rayo, Luis and Ilya Segal**, “Optimal information disclosure,” *Journal of political Economy*, 2010, *118* (5), 949–987.
- Rothschild, Michael and Joseph E Stiglitz**, “Increasing risk: I. A definition,” *Journal of Economic theory*, 1970, *2* (3), 225–243.
- Schweizer, Nikolaus and Nora Szech**, “Optimal revelation of life-changing information,” *Management Science*, 2018, *64* (11), 5250–5262.
- Wei, Dong and Brett Green**, “(Reverse) Price Discrimination with Information Design,” *Available at SSRN 3263898*, 2020.
- Winkler, Gerhard**, “Extreme points of moment sets,” *Mathematics of Operations Research*, 1988, *13* (4), 581–587.
- Yamashita, Takuro and Shuguang Zhu**, “Optimal public information disclosure by

mechanism designer,” 2021.

Yang, Kai Hao, “Selling Consumer Data for Profit: Optimal Market-Segmentation Design and its Consequences,” 2020.