

Supplementary Appendix

A Optimal Mechanisms with Large Laminar Depth

In this section we provide an example where all optimal laminar partitional mechanisms have depth exceeding $|\Theta| + 2$ (i.e., the depth in the single-agent case, see Proposition 2 (i)). There are two players $N = \{1, 2\}$ with two possible types $\Theta_1 = \Theta_2 = \{0, 1\}$ and four possible actions each $A_1 = A_2 = \{0, 1, 2, 3\}$. For convenience, we order type profiles and define a function $\delta : \Theta_1 \times \Theta_2 \rightarrow \{0, 1, 2, 3\}$ such that $\delta(0, 0) = 0$, $\delta(0, 1) = 1$, $\delta(1, 0) = 2$, $\delta(1, 1) = 3$. The players play a zero-sum game. The payoff matrix of the row player for the type profile $\theta = (\theta_1, \theta_2)$ is: $\omega(I + P_{\delta(\theta)})$, where P_k is the permutation matrix whose (ℓ_1, ℓ_2) -th entry is one if $\ell_2 - \ell_1 = k \bmod(4)$. The state ω is distributed uniformly on $[0, 1]$. The type profile distribution is such that $\phi(0, 0) = 0.1$, $\phi(0, 1) = 0.2$, $\phi(1, 0) = 0.3$, $\phi(1, 1) = 0.4$. The state and the types are distributed independently. The designer's payoff is 1 if $a_1 = a_2$ and 0 otherwise.

An optimal mechanism is given in Figure 4. As can be seen from this figure the depth of the laminar family supporting the optimal information structure is larger than $|\Theta| + 2 = 6$. We numerically verified that any other laminar partitional mechanism mechanism that is optimal also has depth greater than 6. Furthermore, when the number of actions is smaller (and the type space is the same) for any payoff structure laminar families of smaller depth suffice. Conversely, when the number of actions is larger, even with the same type space it is possible to obtain even deeper laminar families at the optimal mechanism for variants of this example.

B A Finite-Dimensional Formulation for the Multi-Agent Case

In the single-agent case, when the agent has finitely many actions Section 4.1 established that it is possible to obtain the optimal mechanism by solving a finite-dimensional convex program. This simplification was partly driven by two factors: (i) the agent can perfectly infer the posterior mean from the action recommendation (ii) the posterior mean levels that induce a given action can be characterized explicitly given the agent's payoff function. These factors allowed us to remove the recommended action from the problem and express

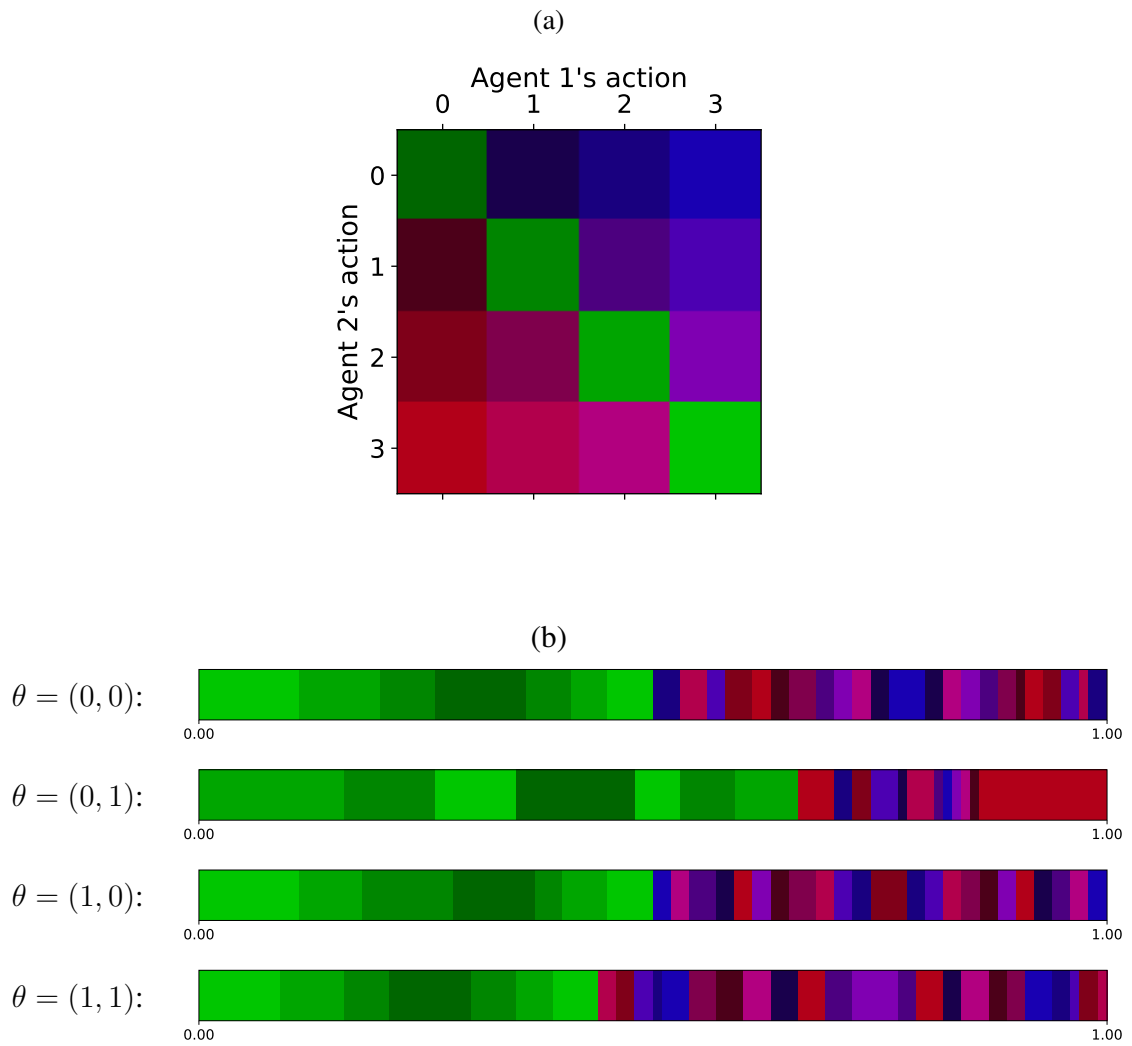


Figure 4: Optimal mechanism. (a) shows the colors associated with each strategy profile. (b) is the laminar partitional signals that constitute an optimal mechanism. Different shades of green encode the strategy profiles where the designer achieves nonzero payoff. For all type profiles, such strategy profiles are associated with smaller states. The laminar partitional signals in this example have depth 12.

it purely in terms of posterior means. As these factors are not present in the multi-agent case, it is unclear whether one can obtain the optimal mechanism through a solution of a finite-dimensional optimization problem. We next argue that indeed through the solution of finite-dimensional programs it is possible to obtain an optimal mechanism, for as long as the agents have finitely many actions.

Consider the formulation in (OPT). Note that for any given profile θ the distribution q^θ over type profiles determines the action profiles recommended at different posterior mean levels. Since no action profile is recommended with positive probability at two different posterior mean levels, it means that action profiles are ordered according to the posterior mean levels that induce them. Denote by δ^θ this order: $\delta^\theta(a) \geq \delta^\theta(a')$ if posterior mean that induces a is larger than that associated with a' when the type profile is θ .

Following an approach similar to the one in Section 4.1, we can now express the designer's problem as follows:

$$\begin{aligned}
& \max_{\{\delta^\theta\}_\theta} \max_{\substack{p \in (\Delta^{|A|})^\Theta \\ z \in \mathbb{R}_+^{|A| \times |\Theta|} \\ y_i \in \mathbb{R}^{|A_i| \times |\Theta|^2}}} \sum_{\theta \in \Theta} \phi(\theta) \sum_{a \in A} p_{a,\theta} v(a, \theta) \\
& \text{s.t.} \quad \sum_{a | \delta^\theta(a) \geq \delta^\theta(\ell)} z_{a,\theta} \leq \int_{1 - \sum_{a | \delta^\theta(a) \geq \delta^\theta(\ell)} p_{a,\theta}}^1 F^{-1}(x) dx \quad \forall \theta \in \Theta, \ell > 1, \\
& \quad \sum_{a \in A} z_{a,\theta} = \int_{\Omega} F^{-1}(x) dx \quad \forall \theta \in \Theta, \\
& \quad \sum_{\theta_{-i}} \phi(\theta) \sum_{a_{-i}} (u_{i1}(a, \theta) z_{a,\theta} + u_{i2}(a, \theta) p_{a,\theta}) \\
& \quad \geq \sum_{\theta_{-i}} \phi(\theta) \sum_{a_{-i}} (u_{i1}(a'_i, a_{-i}, \theta) z_{a,\theta} + u_{i2}(a'_i, a_{-i}, \theta) p_{a,\theta}) \quad \forall i, \theta_i, a_i, a'_i \\
& \quad y_{i,\theta_i,\theta'_i,a_i} \geq \sum_{\theta_{-i}} \phi(\theta) \sum_{a_{-i}} (u_{i1}(a'_i, a_{-i}, \theta) z_{a,\theta} + u_{i2}(a'_i, a_{-i}, \theta) p_{a,\theta}) \quad \forall i, \theta_i, \theta'_i, a_i, a'_i \\
& \quad \sum_{a_i} \sum_{\theta_{-i}} \phi(\theta) \sum_{a_{-i}} (u_{i1}(a, \theta) z_{a,\theta} + u_{i2}(a, \theta) p_{a,\theta}) \geq \sum_{a_i} y_{i,\theta_i,\theta'_i,a_i} \quad \forall i, \theta_i, \theta'_i \\
& \quad z_{a,\theta} p_{a',\theta} \geq z_{a',\theta} p_{a,\theta} \quad \forall \theta, \delta^\theta(a) \geq \delta^\theta(a') \\
& \quad z_{a,\theta} \leq p_{a,\theta} \quad \forall a, \theta.
\end{aligned}$$

In this optimization problem, $p_{a,\theta}$ denotes the probability with which strategy profile a is induced when the type profile is θ , and $m_{\theta,a} = z_{a,\theta}/p_{a,\theta}$ is the corresponding posterior mean level. Note that $\{p_{a,\theta}, z_{a,\theta}\}_a$ tuple constitutes a reparameterization of G^θ . For a given

order δ^θ on posterior mean levels, the first two constraints amount to a restatement of the MPC constraint $G^\theta \succeq F$. Note that if agents report their types truthfully and follow the action recommendations, the payoff of agent i for type profile θ and action recommendation profile a is given by $u_{i1}(a'_i, a_{-i}, \theta) z_{a, \theta'_i, \theta_{-i}} / p_{a, \theta'_i, \theta_{-i}} + u_{i2}(a'_i, a_{-i}, \theta)$. This implies that his expected payoff³⁸ is given as in the left hand side of the third constraint. Similarly, the right hand side is the payoff from taking action a'_i . Thus, the third constraint ensures that if agents report their type truthfully and agent i gets the action recommendation a_i , any deviation reduces his payoff. Suppose that agent i is of type θ_i but he misreported his type as θ'_i , and received action recommendation a_i . Assuming all agents still truthfully report their types and follow action recommendations, what is i 's payoff from taking action a'_i ? The right hand side of the fourth constraint captures this quantity. At the optimal solution, the left hand side, $y_{i, \theta_i, \theta'_i, a_i}$ equals the maximization of this quantity over a'_i , which is the best payoff i can guarantee after the type report θ'_i and action recommendation a_i . Aggregating these terms over all i yields the right hand side of the fifth constraint, which is the expected payoff of agent i from misreporting his type as θ'_i . The left hand side is the payoff from truthful reporting and following action recommendations. Thus, the fifth constraint ensures that agent i has no incentive to misreport his type. The sixth constraint can be equivalently written as $m_{\theta, a} = z_{a, \theta} / p_{a, \theta} \geq m_{\theta, a'} = z_{a', \theta} / p_{a', \theta}$. This ensures that the $\{p_{a, \theta}, z_{a, \theta}\}_a$ tuple and the associated distribution G^θ is consistent with δ^θ in terms of the ranking of the posterior means of strategy profiles. Finally, the last constraint (together with the nonnegativity of $p_{a, \theta}, z_{a, \theta}$) ensures that the posterior means are between 0 and 1.

To solve this problem, we can first fix δ^θ in the outer problem and solve the associated inner problem. Then, we can search over the orders δ^θ (of which there are finitely many) in the outer problem. There are two challenges with this approach. First, the number of orders to consider in the outer problem can be large. Second, unlike the formulation in Section 4.1, due to the sixth constraint the inner problem is not convex.

It turns out that it is possible to overcome both challenges. Let us start with the second challenge. Despite the fact that the inner problem is nonconvex, a locally optimal solution can be obtained using, e.g., gradient ascent. If at a locally optimal solution, the nonconvex constraints are not binding then, it follows that the solution is locally optimal in the problem where these constraints are relaxed. However, the latter problem is convex and local

³⁸As explained in the main text, this quantity is actually equal to the expected payoff times $\sum_{\theta_{-i}} \phi(\theta)$. With some abuse of terminology, throughout the supplementary appendix we ignore this normalization and refer to such quantities as payoffs.

optimality implies global optimality. Thus, the aforementioned solution is a globally optimal solution to the inner problem. In all our numerical experiments (including the Cournot example discussed in Section 2.1) this was the case, i.e., when we obtained a locally optimal solution using a solver, we observed that the nonconvex constraints did not bind and verified global optimality of the said solution.

The first challenge is problem specific, but the search can be drastically reduced in some cases. For instance, observe that in the Cournot example of Section 4.1, there are 9 strategy profiles, and naively there are $9!$ orders to consider. However, due to the symmetry in the problem it can be readily seen that the posterior means associated with strategy profiles (a_i, a_j) and (a_j, a_i) are identical. Furthermore, intuitively, posterior means associated with larger aggregate supply levels will be larger. That is $m_{\theta, a} > m_{\theta', a'}$ if $a_i + a_j > a'_i + a'_j$. Once this restriction imposed, together with symmetry the number of orders to consider reduces only to two (one where strategy profiles $(0, 2), (2, 0)$ are associated with higher posterior mean levels than $(1, 1)$, and one with lower). Thus, solving the inner problem for these two orders, and picking the solution that results in a higher payoff delivers the optimal mechanism. This is in fact how we obtain the optimal mechanisms in Section 2.1 (where we also numerically verify that imposing the aforementioned condition is without loss). Notably this approach allows for constructing the optimal mechanisms without discretization of the state space. Using the approach described here the optimal solution to the optimization problem in Section 2.1 is obtained in ~ 20 ms for most weighted combinations of CS and FP (using off-the-shelf interior point methods of the Knitro solver).

C Additional Details for the Example in Section 4.2

Here we revisit the example in Section 4.2. The indirect utility $\bar{u}(m, \theta)$ of the buyer in this example is given in Figure 5. When the expected quality m of the good is low, all types find it optimal to purchase zero units, yielding a payoff of zero. As the expected quality improves, the purchase quantity increases. In Figure 5, the curve for each type is piecewise linear, and the kink-points of each curve correspond to the posterior mean levels where the agent increases his purchase quantity. Since the state and hence the posterior mean belong to $[0, 1]$, the purchase quantity of each type is at most 2 units, and each curve in the figure has at most two kink-points. This is easily seen as the utility any buyer type derives from consuming the third unit of the good is bounded by $(\theta + \omega) \max\{5 - 3, 0\} \leq (0.6 + 1) \cdot 2 = 3.2$ which is less than the price of $10/3$. These observations imply that in this problem, the




Figure 5: The indirect utility of the agent

agent effectively considers finitely many actions, namely the quantities in 0, 1 and 2.

The effect of the incentive compatibility constraints on the optimal mechanism are easily seen from Figure 3. For instance, the high type's payoff from a truthful type report is strictly positive. If this were the only relevant type, the designer could choose a strictly smaller threshold than 0.06 and still ensure purchase of two units whenever state realization is above this threshold, thereby increasing the expected purchase amount of the high type. However, when the other types are also present, such a change in the signal of the high type incentivizes this type to deviate and misreport his type as low or medium. Changing the signals of the remaining types to recover incentive compatibility, reduces the payoff the designer derives from them. The mechanism in Figure 3 maximizes the designer's payoff while carefully satisfying such incentive compatibility constraints.

As discussed in Section 4.2, in the binary action case it is without loss to focus on public mechanisms (which do not elicit the agent's type). In this case, one way to obtain an optimal public mechanism is to first solve for the optimal mechanism without restriction to public ones, and then reveal to each type the signals associated with all types. By contrast, the mechanism illustrated in Figure 3 does not admit such a payoff-equivalent public implementation. For instance, under this mechanism the high type purchases two units whenever the state realization is higher than 0.06. Suppose that this type of agent had access to the signals of, for instance, the low type as well. Then, he could infer whether the state is in $[0.06, 0.16] \cup [0.94, 1]$. Conditional on the state being in this set, his expectation of the state would be approximately 0.43. This implies that the expected payoff of the high type from purchasing the second unit is $(0.43 + 0.6) \times 3 - 10/3 < 0$. Thus, for state realizations that belong to the aforementioned set, the high type finds it optimal to strictly

reduce his consumption (relative to the one in Figure 3). In other words, observing the additional signal reduces the expected purchase of the high type (and the other types). Hence, such a public implementation is strictly suboptimal. As a side note, the optimal public implementation can be obtained by replacing different types with a single “representative type” and using the framework of Section 3. More precisely, we can replace the designer’s indirect utility with $\bar{v}(m) = \sum_{\theta} \phi(\theta) \max_{a \in A(m, \theta)} v(a, m, \theta)$ and maximize $\int \bar{v}(m) dG(m)$ by choosing a distribution $G \succeq F$ (without any side constraints – since with public signals the designer does not screen the agent, and hence the IC constraints become irrelevant). We numerically conducted this exercise and also verified that restricting attention to public mechanisms yields a strictly lower expected payoff to the designer.