



WORKING PAPERS

N° 1718

February 2026

“Calibrated Mechanism Design”

Laura Doval and Alex Smolin

Calibrated Mechanism Design*

Laura Doval[†] Alex Smolin[‡]

February 18, 2026

Abstract

We study mechanism design when a designer repeatedly uses a fixed mechanism to interact with strategic agents who learn from observing their allocations. We introduce a static framework, *calibrated mechanism design*, requiring mechanisms to remain incentive compatible given the information they reveal about an underlying state through repeated use. In single-agent settings, we prove implementable outcomes correspond to two-stage mechanisms: the designer discloses information about the state, then commits to a state-independent allocation rule. This yields a tractable procedure to characterize calibrated mechanisms, combining information design and mechanism design. In private values environments, full transparency is optimal and correlation-based surplus extraction fails. We provide a microfoundation by showing calibrated mechanisms characterize exactly what is implementable when an infinitely patient agent repeatedly interacts with the same mechanism. Dynamic mechanisms that condition on histories expand implementable outcomes only by weakening incentive constraints, but not by enriching the designer's ability to obfuscate learning.

*The latest version of the paper can be found [here](#). We thank Dirk Bergemann, James Best, Elliot Lipnowski, Emir Kamenica, Navin Kartik, Stephen Morris, Phil Reny, Marzena Rostek, and [Refine.ink](#) for valuable comments and suggestions, as well as seminar audiences at EARIE 2024, WUSTL Economic Theory Conference 2024, MIT IDSS Distinguished Speaker Series, the Workshop in Market Design, Bonn, Berlin HU, CERGE-EI, Gerzensee, UCL, Naples, Stanford University, MIT Sloan, and the CEPR Workshop on Contracts, Incentives and Information at Collegio Carlo Alberto. Laura Doval gratefully acknowledges financial support from the Sloan Foundation. Alex Smolin gratefully acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future program (grant ANR-17-EURE-0010) and the AI Interdisciplinary Institute ANITI (grant ANR-23-IACL-0002). This paper was partly written while the first author was visiting Stanford University and the second author was visiting Northwestern University and Columbia Business School; we are grateful for their hospitality.

[†]Columbia Business School and CEPR. E-mail: laura.doval@columbia.edu.

[‡]Toulouse School of Economics and CEPR. E-mail: alexey.v.smolin@gmail.com.

1 Introduction

Many economic institutions rely on mechanisms that remain fixed while agents interact with them repeatedly. Online platforms commit to stable auction formats for advertising slots, lenders use persistent scoring algorithms for loan decisions, and regulators establish durable rules for market participants. When the mechanism's operation depends on information known only to the designer—such as the platform's data about match values, the lender's assessment of credit market conditions, or the regulator's understanding of market fundamentals—participants may infer this information by observing their outcomes across repeated interactions. This learning creates a fundamental constraint: the information a mechanism reveals through repeated use limits what outcomes it can implement in the long run. Participants can use the information gleaned from past interactions when deciding whether and how to participate, tightening the designer's incentive constraints. A lender whose approval decisions depend on unobserved credit market conditions will gradually reveal these conditions to borrowers through his lending decisions, constraining the lender's ability to provide credit efficiently. We study how this endogenous information leakage shapes the set of implementable outcomes in mechanism design.

A simple example illustrates how learning prevents the designer from exploiting his information. Consider a seller who repeatedly offers a good whose demand depends on an unobserved state, which can be either low (L) or high (H). Each state is equally likely. The seller faces a buyer whose value for the good can take one of two values, $1/2$ or 1 . The probability that the buyer's value is 1 is higher when the demand state is high. [Table 1](#) summarizes the value distribution conditional on the demand state:

	$v = 1/2$	$v = 1$
L	$2/3$	$1/3$
H	$1/3$	$2/3$

Table 1: Value distribution conditional on demand state.

Suppose the seller can design the terms of trade, that is, the probability with which he allocates the good to the buyer ($q \in [0, 1]$) and the payment the buyer makes to the seller ($t \in \mathbb{R}$). The buyer's payoff is $vq - t$, and the seller's is t . The buyer can always choose to not trade with the seller and ensure a payoff of 0.

Suppose first the buyer and the seller interact only once. [Table 2](#) depicts an optimal mechanism for the seller in this case:

	$v = 1/2$	$v = 1$
L	$(1, 0)$	$(1, 0)$
H	$(1, 3/2)$	$(1, 3/2)$

Table 2: Trade probabilities and payments as a function of buyer's value and demand state.

In this mechanism, the buyer gets the good for free when the demand state is L and pays a price of $3/2$ when it is H . If this mechanism were offered once without the buyer observing the demand state, the buyer obtains a payoff of 0 from participating and truthfully reporting her type. Unsurprisingly,

the seller extracts the buyer's surplus: the seller knows the demand state, which is correlated with the buyer's type, and exploits this information in the design of his mechanism (cf. Crémer and McLean, 1988).

Suppose now the buyer interacts repeatedly with the mechanism, but the state remains fixed. If the buyer observes nothing from her interaction with the mechanism, the buyer is willing to participate and truthfully report her value into the mechanism, no matter how many times it is offered: In each period, she anticipates getting a (continuation) payoff of 0 from engaging with the mechanism. Suppose, instead, the buyer observes her allocation in the mechanism. If the demand state is L , the buyer gets the good for free at the end of the first period, and from now on knows this is what she will get in the mechanism. If the demand state is H , the buyer gets the good and pays a price of $3/2$ as she agreed to when she decided to participate in period 1, but anticipating a price of $3/2$ from then onwards, never again participates in the mechanism. Thus, whereas the seller can implement the outcomes in [Table 2](#) when the buyer does not observe her allocations, this is no longer the case when she can.

This paper develops a framework for mechanism design in which agents' ability to learn about the designer's information from repeatedly playing a mechanism constrains implementable outcomes. In our framework, allocations depend on agents' reports and on a state known only to the designer. Through repeated participation, agents observe their allocations and gradually learn about this state. A mechanism therefore serves a dual role: it determines allocations based on reports, and it acts as an information structure that reveals the underlying state. The more the mechanism conditions on the state, the more information it leaks, and the tighter the constraints on implementable outcomes.

We approach our analysis in two steps. First, we introduce a static solution concept for mechanism design that directly models the feedback between the mechanism, the information it reveals, and participants' behavior. This solution concept allows us to tractably capture the limits on the set of implementable outcomes implied by agents' learning, while abstracting from the dynamics of experimentation. Second, we provide a dynamic microfoundation showing this static solution concept precisely captures the implementable outcomes when an infinitely patient agent repeatedly interacts with the same mechanism.

In [Section 2](#), we introduce a static solution concept—calibrated mechanism design—requiring that mechanisms remain incentive compatible and individually rational given the information they reveal about the state through their allocations. We formalize this requirement through the notion of a *calibrated mechanism*. We couple each mechanism with an information structure that describes what participants learn about the state from the mechanism. The information structure reveals to each agent an interim allocation rule—the mapping from her type reports to lotteries over her allocations—capturing what she would learn from repeatedly observing her outcomes in the mechanism. We require the information structure to be calibrated in the sense of [Foster and Vohra \(1997\)](#): the interim allocation rule each agent observes must accurately describe the allocation probabilities she faces. Throughout the paper, we study calibrated mechanism design: the designer chooses a mechanism that remains incentive compatible and individually rational when participants have access to the mechanism's calibrated information structure before playing. Calibration imposes a constraint

on the designer relative to standard mechanism design: the more the allocation rule depends on the state, the more informative the calibrated information structure becomes, and hence the more incentive and participation constraints the designer must satisfy.

In private values environments, the constraint that the mechanism must remain incentive compatible and individually rational given the information it reveals about the state pushes the designer to full transparency. We show in [Theorem 1](#) that, under the calibration constraint, the designer can do no better than inducing in each state the optimal direct mechanism when there is common knowledge of that state. In particular, in settings with transferable utility in which the designer has statistical information about the agents' types, [Theorem 1](#) implies the designer cannot extract full surplus.

In [Section 3](#), we characterize optimal calibrated mechanisms through a tractable class we dub *two-stage* mechanisms. In a two-stage mechanism, the designer first discloses information about the state to the agent—inducing a belief about the state—then commits to an allocation rule that depends only on the agent's report, not the state itself. [Theorem 2](#) shows that in single-agent settings, calibrated mechanisms and two-stage mechanisms implement exactly the same outcome distributions. This equivalence yields a practical algorithm for finding optimal calibrated mechanisms, combining tools from information design and mechanism design: for each possible belief the designer might induce, solve a standard mechanism design problem given that belief; then choose the optimal information disclosure by concavifying the resulting value function.

In the case of multiple agents, [Proposition 1](#) shows calibrated mechanisms admit a similar representation via *generalized* two-stage mechanisms: Like two-stage mechanisms, the designer individually discloses to each agent a belief about the state and offers an incentive compatible and individually rational interim allocation rule that no longer conditions on the state. Whereas the designer observes the disclosed belief profile, each agent only observes the belief disclosed to her.¹ Moreover, each agent learns only her own interim allocation rule—how her reports map to her allocations—rather than the complete mapping from type profiles to allocations. This partial observability requires additional consistency conditions to ensure agents' interim allocation rules are mutually compatible. In contrast to the single-agent case, not every generalized two-stage mechanism induces a calibrated mechanism, as generalized two-stage mechanisms may reveal strictly less information than calibrated mechanisms.

In [Section 4](#), we study optimal calibrated mechanism design in the canonical setting of quasilinear utilities, single-dimensional types and allocations. In [Section 4.1](#), we study the single-agent case. We show that if the order of types is state independent, then optimal two-stage mechanisms fully reveal the state, whereas this conclusion can be reversed when the order of types is state-dependent. In [Section 4.2](#), we compare optimal calibrated mechanism design against the Myersonian benchmark. We provide sufficient conditions under which the designer realizes the payoff of the Myersonian benchmark under the calibration constraint; under these conditions, the optimal Myersonian mechanism satisfies the agent's incentive constraints state-by-state. Building on that result, we analyze multi-agent applications in [Section 4.3](#).

¹In (generalized) two-stage mechanisms, the designer communicates with the agents *before* the agents communicate with the mechanism. Whereas this communication is a restriction on the set of implementable outcomes relative to the single-designer Myersonian benchmark, [Attar et al. \(2025\)](#) show that allowing competing principals to first communicate with agents expands the set of implementable outcomes.

[Section 5](#) provides a microfoundation for calibrated mechanism design. We analyze an infinite-horizon game where an infinitely patient agent repeatedly plays the same mechanism.² The state remains fixed, but the agent's type is redrawn each period independently of the state.³ Our notion of implementation is based on the long-run expected frequency of allocation-type-state tuples when the agent best responds to the mechanism. [Theorem 3](#) shows that the implementable outcome distributions are precisely those induced by incentive compatible two-stage mechanisms. This result validates our static framework: calibrated mechanism design captures exactly what is implementable through repeated play.

We then ask whether giving the designer additional flexibility helps. In a dynamic mechanism, the designer can condition each period's allocation on the complete history of past reports and allocations, rather than using the same mechanism repeatedly. [Theorem 4](#) shows that dynamic mechanisms expand implementable outcomes in a specific way: they correspond to two-stage mechanisms with weaker incentive compatibility and individual rationality conditions. The designer can now exploit the ability to monitor the frequency of type reports over time, which allows him to punish detectable deviations—reporting strategies whose frequency distribution differs from the true type distribution. Instead, the mechanism must be robust to *undetectable* ones. Importantly, in environments with transferable utility, this distinction vanishes: As shown in [Rahman \(2024\)](#), eliminating profitable undetectable deviations is equivalent to incentive compatibility, so dynamic mechanisms implement exactly the same distributions over physical allocations, types, and states as our static calibrated mechanisms.

Related Literature The paper lies at the intersection of four literatures: rational expectations equilibria, (public) information disclosure in mechanism design, the computer science literature on learning in repeated auctions, and dynamic implementation.

The definition of a calibrated mechanism is in the spirit of rational expectations equilibria ([Radner, 1979](#); [Green, 1977](#); [Kreps, 1977](#)). Indeed, requiring a mechanism to remain incentive compatible given the information it reveals about the state mirrors the rational-expectations requirement that prices clear markets given the information they convey. Unlike rational expectations equilibrium, where the only role of prices is to clear the market, calibrated mechanisms are chosen by a designer who understands the incentive implications of the mechanism's information leakage and trades this off against the value of conditioning the mechanism on the state. Similar to our analysis in [Section 5](#), some papers in the literature have studied the question of whether rational expectations equilibria emerge from learning dynamics (see, for instance, [Milgrom, 1981](#); [Blume et al., 1982](#)).

Following [Milgrom and Weber \(1982\)](#), a literature has studied whether a designer should publicly disclose information he knows before a mechanism is played. [Ottaviani and Prat \(2001\)](#) show revealing a signal affiliated with the buyer's value is optimal in a single-agent screening problem. When considering the case of an informed principal, they consider what we call two-stage mechanisms to bound the monopolist's profits. [Szabadi \(2018\)](#) and [Yamashita \(2018\)](#) study the optimal release of

²We assume the agent has limit-of-means preferences, so we can pass to the $\delta \rightarrow 1$ limit without approximation.

³[Theorem 3](#) holds when the agent's type is drawn once at the beginning without further assumptions on the distribution. As we explain in [Section 5](#), we choose the i.i.d. specification for the evolution of the agent's private information to put repeated and dynamic mechanisms on a more equal footing.

public information followed by an optimal mechanism conditional on that disclosure, while Fu et al. (2012) study this question in the context of a second price auction. In those papers, the restriction to public disclosure and the independence of the mechanism on information other than the disclosed one is a constraint on the class of mechanisms the designer can use. Instead, we show this class of mechanisms is without loss when the designer faces our calibration constraint in the single-agent case, but it may not be in the multi-agent case. Note, however, that when full or no disclosure are optimal in the Myersonian benchmark the distinction between private and public disclosure is immaterial. For that reason, the results on the achievability of the Myersonian benchmark are similar across their and our work. Daskalakis et al. (2016) lift the restriction to public disclosure and study the Myersonian benchmark in an auction setting, showing that the complexity of that problem is the same as that of a multi-product monopolist (cf. Guesnerie and Laffont, 1984).⁴

Motivated by the prevalence of fixed auction formats with which bidders interact repeatedly, a literature in computer science studies the properties of bidder learning algorithms and the implications for the auctioneer (see, for instance, Golrezaei et al., 2019; Nedelec et al., 2019; Kanoria and Nazerzadeh, 2020, and Nedelec et al., 2022 for a survey treatment). A common finding is that learning bidders can take advantage of “naive” auction formats which are no longer incentive compatible when bidders learn. Inspired by this literature, we develop a framework which allows us to systematically study the question of optimal mechanism design in the presence of learning agents.

Our dynamic implementation results relate to the literature that studies whether a mechanism can be implemented either by linking decisions (Jackson and Sonnenschein, 2007; Ball and Kattwinkel, 2023) or in the patient limit of a repeated interaction (Renou and Tomala, 2015; Margaria and Smolin, 2018; Meng, 2021). Both strands identify cyclical monotonicity as the condition for implementation (cf. Rochet, 1987). Rahman (2024) shows that cyclical monotonicity is equivalent to the absence of profitable undetectable deviations.

By focusing on what agents learn from the designer’s information, our paper is distinct from the literature on mechanism design with interdependent payoffs which focuses on agents’ learning about others’ types through their actions in the mechanism (Green and Laffont, 1987; Niemeyer, 2022; Häfner et al., 2025). Moreover, by focusing in the case of a designer with commitment, we are distinct from the literature on the informed principal (Myerson, 1983; Maskin and Tirole, 1990).

Lastly, our paper contributes to two literatures. First, by studying the informational role of the mechanism, we contribute to the literature on feedback in auctions, which analyzes how different feedback rules affect bidders’ information about other agents, and ultimately behavior in first price auctions (see, for instance, Esponda, 2008; Bergemann and Hörner, 2018; Cesa-Bianchi et al., 2024). Second, by showing the designer’s problem involves solving information and mechanism design problems, our paper joins a recent literature that highlights the dual role of the mechanism as an information structure and an allocation rule (Calzolari and Pavan, 2006; Dworczak, 2020; Doval and Skreta, 2022).

⁴There is also a literature that studies a designer’s disclosure of information that must be elicited from the agents (Eső and Szentes, 2007; Bergemann and Pesendorfer, 2007; Li and Shi, 2017; Krähmer, 2020; Bergemann et al., 2022a,b; Smolin, 2023). By contrast, the designer knows the realization of the state and also what information the two-stage mechanism discloses to the agents, so he need not elicit this information.

2 Calibrated Mechanism Design

In this section, we introduce the static setting and solution concept that captures the impact of agents' learning from the mechanism on the set of implementable outcomes. We defer to Section 5 the analysis of the dynamic game whose outcomes our static solution concept captures.

Primitives A designer (he) interacts with N privately informed agents (she) to determine an allocation. Let Θ_i denote the set of types of agent i , and $\Theta \equiv \times_{i=1}^N \Theta_i$. Each agent knows her type, but not those of other agents. The allocation space is given by $A \equiv \times_{i=1}^N A_i$.⁵ Finally, let Ω denote a set of states, which are known to the designer, but not to the agents. The sets Θ_i , A_i , and Ω are assumed to be finite throughout.⁶ Agent i 's payoffs are given by $u_i : A_i \times \Theta_i \times \Omega \rightarrow \mathbb{R}$. That is, agent i cares about her dimension of the allocation, her type, and the state, and not about other agents' allocations or types.

Denote by μ_0 the distribution over Ω . For each $\omega \in \Omega$, let $f(\cdot|\omega) \in \Delta(\Theta)$ denote the type distribution. We assume throughout the types are independently distributed conditional on the state, that is,

$$f(\theta|\omega) = \prod_{i=1}^N f_i(\theta_i|\omega), \quad (1)$$

for all $\theta \in \Theta$ and $\omega \in \Omega$. Together with the assumption on agents' payoffs, the assumption on $f(\cdot|\omega)$ allows us to isolate the effect of learning about the state from that of learning about others' types (perhaps because others' types provide additional information about the state).

Mechanisms We model mechanisms as mappings

$$\phi : \Theta \times \Omega \times [0, 1] \rightarrow \Delta(A), \quad (2)$$

where $\varepsilon \in [0, 1]$ is a uniformly distributed random variable, which we refer to as the randomization device.

Several comments are in order. First, to understand how a mechanism works, the timing of when the different random variables is drawn is important. In particular, we assume that both the state ω and the realization of the randomization device ε are independently drawn at the beginning, but not observed by the agents. This determines the direct mechanism $\phi(\cdot, \omega, \varepsilon) : \Theta \rightarrow \Delta(A)$ to which the agents send type reports, which in turn determines the lottery from which the allocation is drawn. Thus, the allocation is random in our setting for two reasons: on the one hand, the agents do not know the realization of (ω, ε) , and hence the direct mechanism $\phi(\cdot, \omega, \varepsilon)$ they face. Second, conditional on (ω, ε) , the allocation may be drawn at random. Mathematically, we could have subsumed *all* sources of randomness in the allocation into the randomization device. However, as we explain next, the definition in Equation 2 allows us to distinguish the source of randomness in the allocation that is informative about the state from that which is not.

⁵Assuming the allocation space is a product space is without loss of generality. Any restriction on the allocations, such as all agents must receive the same allocation, can be incorporated as restrictions on the support of the mechanism.

⁶Because we allow for lotteries over allocations, that the allocation space is finite does not preclude the case of transferable utility. Indeed, we could let each $A_i = \tilde{A}_i \times \{-K, K\}$ for some large enough $K > 0$.

Second, it is useful to consider the reason for the randomization device in the definition of a mechanism. For simplicity, consider the case of the designer facing a single agent. If the agent had repeated access to the mechanism, the agent would stand to learn the mapping $\phi(\cdot, \omega, \varepsilon) : \Theta \rightarrow \Delta(A)$ by experimenting with different reports into the mechanism and observing the resulting allocations.⁷ Without the randomization device, the agent would stand to learn a partition of the set of states, where states in the same cell of the partition induce the same direct mechanism $\phi(\cdot, \omega, \varepsilon)$. By allowing the designer to rely on the randomization device, we allow him to *obfuscate* the agent's learning beyond a simple partitional structure. Contrast this with the Myersonian benchmark in which without loss of generality the designer would offer mechanisms that do not rely on such devices, that is, $\phi_{\text{My}} : \Theta \times \Omega \rightarrow \Delta(A)$. Indeed, the Myersonian designer is not concerned with the agents' learning: without loss of generality, he does not disclose anything about the state to the agents, so that the question of how to optimally release information about the state is moot.

Lastly, note that we assume the mechanism asks the agents for type reports. In [Appendix D](#), we show that the revelation principle holds in the setting of this section: it is without loss of generality to focus on direct and incentive compatible mechanisms that induce full participation.

Calibrated information structures We now describe how a mechanism induces an information structure, which we define using the language in [Green and Stokey \(2022\)](#) and [Gentzkow and Kamenica \(2017\)](#). An information structure is a mapping⁸

$$\pi : \Omega \times [0, 1] \rightarrow S_1^* \times \cdots \times S_N^*,$$

where $\varepsilon \in [0, 1]$ is a uniformly distributed random variable—in fact, it is the same as in the definition of a mechanism—and

$$S_i^* = \Delta(A_i)^{\Theta_i},$$

is the set of agent i 's interim allocation rules.⁹ We choose this language for the information structure to capture the idea that if agent i plays the mechanism repeatedly, she stands to learn how her reports influence her allocation probabilities, i.e., her interim allocation rule. The interim allocation rule, in turn, depends on the mechanism and the strategies of others. Below, we require the interim allocation rule is well-calibrated with the mechanism and others' strategies:

Definition 1 (Calibrated information structures). *We say that the information structure is calibrated to mechanism ϕ if for all $(\omega, \varepsilon) \in \Omega \times [0, 1]$ such that $\pi(\omega, \varepsilon) = (s_1^*, \dots, s_N^*)$ we have that for all $i \in \{1, \dots, N\}$,*

⁷When the agent is infinitely patient as in [Section 5](#), we can exhibit a sequence of strategies under which the agent (approximately) learns this mapping. See the proof of [Lemma C.4](#) in [Appendix D](#).

⁸[Gentzkow and Kamenica \(2017\)](#) highlight that the language in [Green and Stokey \(2022\)](#) allows one to describe the correlation across signal structures. This is exactly what we need to allow the designer to obfuscate the agents' ability to learn. It is again instructive to consider the single-agent case. For each type report $\theta \in \Theta$, the mechanism can be seen as an information structure $\phi(\theta, \cdot) : \Omega \times [0, 1] \rightarrow \Delta(A)$. Thus, the randomization device allows the designer to control the correlation across these different signal structures, which in turn disciplines what the agent stands to learn when experimenting with different reports.

⁹The terminology is by analogy to reduced form auctions where the map from own types to own probabilities of being allocated the good are referred to as the interim allocation.

all $\theta_i \in \Theta_i$, and all $a_i \in A_i$

$$s_i^*(a_i|\theta_i) = \mathbb{E}_{\tilde{\theta}_{-i} \sim f_{-i}(\cdot|\omega)} \left[\sum_{a_{-i} \in A_{-i}} \phi(\theta_i, \tilde{\theta}_{-i}, \omega, \varepsilon)(a_i, a_{-i}) \right]. \quad (3)$$

We denote by π_ϕ the information structure calibrated to mechanism ϕ .

In words, the information structure is calibrated if whenever agent i observes that her interim allocation rule in the mechanism is s_i^* , then s_i^* describes the true probabilities with which agent i gets different allocations a_i as a function of her different type reports θ'_i in the mechanism. As the right hand side of [Equation 3](#) shows, these probabilities depend on: (i) the mechanism $\phi(\cdot, \omega, \varepsilon)$, and (ii) others' type reports. Implicit in the definition is that other agents are submitting their reports truthfully. While this is a simplification,¹⁰ it turns out to not be an issue because we study incentive compatible and individually rational mechanisms in the sense we define next.

Information leakage from a mechanism To close our model, we consider how the mechanism and its induced information structure affect agents' incentives. The mechanism ϕ and the calibrated information structure π_ϕ induce the following game of incomplete information among the agents, where we use Bayes Nash equilibrium as the solution concept. In this game, nature draws (i) the state ω from distribution μ_0 , (ii) $\varepsilon \in [0, 1]$ according to the uniform distribution, and (iii) the type profile θ from $f(\cdot|\omega)$. Then, each agent i observes her type θ_i and her signal $s_i^* = \pi_{\phi,i}(\omega, \varepsilon)$. Finally, agents simultaneously decide whether to participate in the mechanism, and conditional on participating what type report to send. Conditional on an agent choosing not to participate, each agent i gets outside option $a_{i\emptyset}$.¹¹

Formally, given the mechanism ϕ and its calibrated information structure π_ϕ , we say that the mechanism is incentive compatible if for all agents i , types $\theta_i \in \Theta_i$, signals $s_i^* \in S_i^*$ on the support of $\pi_{\phi,i}$, the following holds:

$$\theta_i \in \arg \max_{\theta'_i \in \Theta_i} \mathbb{E}_{(\omega, \varepsilon, \theta_{-i})} [u_i(\phi(\theta'_i, \theta_{-i}, \omega, \varepsilon), \theta_i, \omega) | (\theta_i, s_i^*)], \quad (\text{IC}(\theta_i, s_i^*))$$

where we abuse notation and implicitly (linearly) extend the agent's payoff function to account for lotteries over allocations (conditional on $(\theta_{-i}, \omega, \varepsilon)$). Furthermore, we say that the mechanism is individually rational if for all agents i , types $\theta_i \in \Theta_i$, and signals $s_i^* \in S_i^*$ on the support of $\pi_{\phi,i}$, the following holds:

$$\mathbb{E}_{(\omega, \varepsilon, \theta_{-i})} [u_i(\phi(\theta_i, \theta_{-i}, \omega, \varepsilon), \theta_i, \omega) - u_i(a_{i\emptyset}, \theta_i, \omega) | (\theta_i, s_i^*)] \geq 0. \quad (\text{IR}(\theta_i, s_i^*))$$

Importantly, the agents' incentive and participation constraints must hold for each of their types and each of their private signals, reflecting the agents have access to the information leaked by the mechanism before they play in it. Note, however, the mechanism need not elicit the agents' observed

¹⁰When we consider mechanisms with arbitrary message spaces in [Appendix D.1](#), the requirement of calibration is relative to both the mechanism and agents' equilibrium participation and reporting strategies.

¹¹Thus, we are assuming that $a_\emptyset = (a_{i\emptyset})_{i \in N}$ is an element of A . In [Appendix D.1](#), we consider more general participation decisions, allowing the mechanism to condition on the set of participating agents, but even with this extra generality, it is still without loss to restrict attention to mechanisms that induce full participation.

signals, as the mechanism “knows” each agent’s signal realization.

Calibrated Mechanism Design In the rest of the paper, we study the problem of *calibrated mechanism design* in which the designer selects a mechanism ϕ that satisfies Equations $\text{IC}(\theta_i, s_i^*)$ and $\text{IR}(\theta_i, s_i^*)$ for all (i, θ_i, s_i^*) , when the signals are drawn according to the calibrated information structure π_ϕ .

Definition 2 (Calibrated Mechanism Design). *Let $w : A \times \Theta \times \Omega \rightarrow \mathbb{R}$ denote the designer’s payoff and let \mathcal{M}_{cal} denote the set of mechanisms that are incentive compatible and individually rational when agents have access to the calibrated information structure. The calibrated mechanism design problem is as follows:*

$$\max_{\phi \in \mathcal{M}_{\text{cal}}} \mathbb{E}_{(\omega, \varepsilon, \theta)} [w(\phi(\theta, \omega, \varepsilon), \theta, \omega)]. \quad (\text{OPT}_{\text{cal}})$$

We refer to elements of \mathcal{M}_{cal} as calibrated mechanisms and the solution to OPT_{cal} as the optimal calibrated mechanism.

Three comments are in order:

First, calibration imposes a constraint on the designer vis-à-vis standard mechanism design. After all, the incentive and participation constraints faced by the designer are *endogenous* to the mechanism. The more the designer’s mechanism depends on the state, the more informative the calibrated information structure is, and the more incentive constraints the designer faces. Only when each agent’s interim allocation rule is constant in ω does the mechanism not leak information and the incentive and participation constraints reduce to the standard ones.

Second, in the single-agent setting, the calibration constraint admits two complementary interpretations. Throughout the paper, we emphasize the learning-by-experimentation interpretation: the calibrated information structure represents what the agent can ultimately infer by repeatedly interacting with the mechanism. Accordingly, the designer should ensure incentive compatibility with respect to the full information the agent eventually obtains. At the same time, calibration can also be interpreted as a *transparency* requirement. Indeed, upon observing signal $s^* : \Theta \rightarrow \Delta(A)$, the agent knows the consequences of her choices in the mechanism, even if she does not know the state.¹²

With multiple agents, these interpretations differ. The natural extension of the transparency requirement is that agents learn the mapping from profiles of type reports to lotteries over profiles of allocations before playing the mechanism. By contrast, the calibrated information structure reveals to each agent her interim allocation rule, that is, the mappings from her own reports to lotteries over her own allocations. As we discuss in the next section, the gap between these two interpretations is the gap between the designer publicly or privately disclosing information about the state to the agents.

Lastly, the definition of calibration assumes agents only learn about the state through their allocations

¹²It is common for online platforms to inform agents of the consequences of their choices: marketplaces inform sellers of their probability of sale at different posted prices, and transportation providers inform riders of their probability of receiving a seat upgrade at different bid levels. Even insurance companies provide consumers with projected expenditures under different plan choices.

in the mechanism, and not their payoffs.^{13,14} This assumption allows us to focus on the information that the mechanism leaks *regardless of payoff assumptions*. This allows us to avoid situations in which the mechanism does not condition the allocation on the state, but the agents learn because they have different payoffs from the same allocation in different states; or the mechanism conditions on the state, but this information is not payoff relevant to (some types of) the agent. Our microfoundation in Section 5.1 in fact deals with this last wrinkle: We show that even if the agent extracts less information than that in the calibrated information structure, she learns enough that her payoff is *as if* she had access to the calibrated information structure.

We conclude this section by illustrating how our static solution concept captures the dynamics we alluded to in the introductory example:

Example 1 (Selling a good under demand uncertainty). *Consider again the example in the introduction, in which a buyer with binary values $v \in \{1/2, 1\}$ faces a seller who knows whether demand is high ($\omega = H$) or low ($\omega = L$). The left panel of Table 3 describes the probabilities of trade and payments of the optimal (Myersonian) mechanism. In the introduction, we discussed this mechanism fails to extract full surplus in the long run as the buyer would quit the mechanism after seeing her allocation is $(1, 3/2)$. We now describe this in the language of calibration.*

The right panel of Table 3 describes the information structure induced by the surplus extraction mechanism. Because in this mechanism the buyer's allocation does not depend on her values, we describe signals as allocations. The calibrated information structure is fully informative: when the state is L , the buyer sees signal $(1, 0)$ with probability 1, and when the state is H , she sees signal $(1, 3/2)$ with probability 1.

		$v = 1/2$	$v = 1$		
		$\omega = L$	$\omega = H$	$\omega = L$	$\omega = H$
$\omega = L$	$v = 1/2$	(1, 0)	(1, 0)	1	0
	$v = 1$	(1, 3/2)	(1, 3/2)	0	1

Table 3: Trade probabilities and payments in optimal Myersonian mechanism (left); calibrated information structure (right). We describe signals as allocations, because the mechanism does not screen the buyer's values.

When the buyer has access to the calibrated information structure before playing the mechanism, the surplus extraction mechanism does not satisfy the buyer's participation constraints, which must hold for each buyer value and each signal she observes. In particular, when the buyer sees signal $(1, 3/2)$, she knows her payoff in the mechanism is negative and quits. Thus, the calibration constraint prevents the seller from extracting the buyer's surplus. In this case, the restriction induced by calibration endogenously provides the buyer with withdrawal rights, which, as Haberman and Jagadeesan (2025) show, prevent sellers from employing Crémer-McLean-style schemes.¹⁵

¹³Indeed, in the analysis of the dynamic interaction in Section 5, we assume the agent only observes her type and her allocation, but not her payoffs.

¹⁴This assumption is routinely made in dynamic settings. See Pavan et al. (2014) and Cesa-Bianchi et al. (2024) for two examples in the context of agents' behavior within mechanisms.

¹⁵In Section 3, we provide an example in which when agents have access to the calibrated information structure the optimal Myersonian mechanism fails to be incentive compatible.

Consider now the mechanism in the left panel of Table 4, which corresponds to posting a price of 1/2 when the state is L and a price of 1 when the state is H . The right panel of Table 4 depicts the calibrated information structure. Note that when the state is H , the information structure sends with probability 1 the interim allocation rule $\{(1/2, (0, 0)), (1, (1, 1))\}$, representing that if the buyer reports her value is 1/2 she gets nothing and pays nothing, whereas if her report is 1, she obtains the good at a price of 1.

	$v = 1/2$	$v = 1$		$\{(1, 1/2)\}$	$\{(1/2, (0, 0)), (1, (1, 1))\}$
$\omega = L$	(1, 1/2)	(1, 1/2)	$\omega = L$	1	0
$\omega = H$	(0, 0)	(1, 1)	$\omega = H$	0	1

Table 4: Trade probabilities and payments in optimal calibrated mechanism (left); calibrated information structure (right)

Note that the mechanism is incentive compatible and individually rational when the buyer has access to the calibrated information structure. As the results that follow allow us to establish, this is indeed the optimal calibrated mechanism.

Private value environments A natural case to consider is that when agents' payoffs are state independent, that is, for each agent i , the agent's utility function can be written as $u_i(a_i, \theta_i)$. Under private values, the state describes either statistical information about the agents' types as in Example 1, or a payoff-relevant variable for the designer.

Theorem 1 collects our main characterization result for this case. To state it, let ϕ_{full} denote the following mechanism: For each $(\omega, \varepsilon) \in \Omega \times [0, 1]$, $\phi_{full}(\cdot, \omega, \varepsilon) : \Theta \rightarrow \Delta(A)$ is the designer optimal incentive compatible and individually rational direct mechanism when it is common knowledge that the state is ω .

Theorem 1 (Private values). *Under private values, the designer's payoff under the optimal calibrated mechanism is the same payoff he would obtain by choosing ϕ_{full} .*

That is, in private values environments, the calibration constraint pushes the designer toward full transparency. In particular, in settings with transferable utility in which the designer possesses statistical information about the agents' types, Theorem 1 implies the designer cannot engage in Crémer-McLean style schemes under calibration, and hence extract full surplus. Whereas the implication of calibrated mechanism design in private values environments is powerful, the result is fairly intuitive: The designer benefits from making the mechanism opaque by pooling states inasmuch as it weakens the incentive or participation constraints of the agents. Under private values, however, agents' incentive constraints depend on the state only through the mechanism, and calibration imposes constraints on the mechanism state-by-state.¹⁶

¹⁶A tempting comparison is Maskin and Tirole (1990, Prop. 11): with private values and quasilinear utilities, the informed principal's unique equilibrium payoff coincides with the state-by-state optimum. The authors show this conclusion depends on quasilinearity: absent this assumption, an informed principal can benefit from concealing his information in the case of private values. Instead, Theorem 1 relies neither on quasilinearity nor on the designer's lack of commitment.

3 Two-stage mechanisms

In this section, we introduce an alternative representation of calibrated mechanisms that we use throughout our illustrations. We introduce it first for the case of a single agent and then for multiple agents.

Single-agent case and two-stage mechanisms We find it instructive to first consider the case $N = 1$, and for simplicity drop the subscripts 1 from the notation. Consider a mechanism ϕ and its calibrated information structure π_ϕ . When the agent of type θ observes signal s^* , two things happen: On the one hand, the agent updates her prior, $\mu_0(\omega|\theta)$,¹⁷ to some belief $\mu(\theta, s^*) \in \Delta(\Omega)$. On the other hand, the agent learns that she faces allocation rule s^* in the mechanism. Thus, her payoff in the mechanism when her type is θ , observes signal s^* , and reports θ' can be written as follows:

$$\mathbb{E}_{(\omega, \varepsilon)} [u(\phi(\theta', \omega, \varepsilon), \theta, \omega) | (\theta, s^*)] = \sum_{a \in A} s^*(a|\theta') \left(\sum_{\omega \in \Omega} \mu(\omega|\theta, s^*) u(a, \theta, \omega) \right). \quad (4)$$

In other words, the information structure π_ϕ provides the agent with all the necessary information to evaluate her payoffs in the mechanism: her belief about the state *and* her allocation rule. This allocation rule $s^* : \Theta \rightarrow \Delta(A)$ satisfies two properties. First, because under calibration s^* is the true interim allocation rule faced by the agent, she learns no further information about the state beyond that contained in $\mu(\theta, s^*)$. Second, Equations $\text{IC}(\theta_i, s_i^*)$ and $\text{IR}(\theta_i, s_i^*)$ imply the allocation rule is incentive compatible and individually rational when the agent holds belief $\mu(\theta, s^*)$.

The above discussion suggests an alternative representation of a calibrated mechanism, which we dub a two-stage mechanism and define as follows:

Definition 3 (Two-stage mechanisms). *A two-stage mechanism is a mapping $\psi : \Theta \times \Omega \rightarrow \Delta(A \times \Delta(\Omega))$ such that a Bayes plausible Blackwell experiment $\beta : \Omega \rightarrow \Delta(\Delta(\Omega))$ and an allocation rule $\alpha : \Theta \times \Delta(\Omega) \rightarrow \Delta(A)$ exist such that for all $(\theta, \omega) \in \Theta \times \Omega$ and all measurable subsets $\tilde{\Delta} \subset \Delta(\Omega)$,*¹⁸

$$\psi(\{a\} \times \tilde{\Delta} | \theta, \omega) = \int_{\tilde{\Delta}} \alpha(a|\theta, \mu) \beta(d\mu|\omega).$$

We say the two-stage mechanism is incentive compatible and individually rational if on the support of $\mu_0 \otimes \beta$, the allocation rule $\alpha(\cdot|\cdot, \mu) : \Theta \rightarrow \Delta(A)$ is incentive compatible and individually rational conditional on the agent observing μ .

In a two-stage mechanism, the designer first discloses information about ω in the form of a belief μ about Ω , and conditional on that belief—but not the state—offers a direct mechanism $\alpha(\cdot|\cdot, \mu) : \Theta \rightarrow \Delta(A)$. Two aspects of two-stage mechanisms are worth highlighting: First, the disclosure is type-independent. The designer discloses information to the agent without first communicating

¹⁷Formally,

$$\mu_0(\omega|\theta) = \frac{\mu_0(\omega) f(\theta|\omega)}{\sum_{\omega' \in \Omega} \mu_0(\omega') f(\theta|\omega')}.$$

¹⁸We refer the reader to the appendix for our mathematical conventions, in particular, the definition of the corresponding σ -algebras.

with the agent. Second, because the direct mechanism $\alpha(\cdot|\cdot, \mu)$ does not depend on ω , observing the allocation reveals no further information about the state.

Lastly, when we say the experiment β is Bayes plausible, we mean that the distribution of posteriors induced by β has mean μ_0 , and hence we can interpret μ as the designer's belief about the state conditional on observing μ .¹⁹ Whereas the designer and the agent do not necessarily have the same beliefs about the state, the agent's beliefs about the state conditional on observing μ obtain from a known transformation from those of the designer (Alonso and Câmara, 2016; Laclau and Renou, 2017). Thus, ensuring Bayes plausibility with respect to μ_0 suffices.

Theorem 2 shows that (incentive compatible and individually rational) calibrated mechanisms and two-stage mechanisms implement the same distributions over outcomes $\vartheta \in \Delta(A \times \Theta \times \Omega)$:

Theorem 2 (Two-stage and calibrated mechanisms). *Suppose $N = 1$. An outcome distribution $\vartheta \in \Delta(A \times \Theta \times \Omega)$ is implementable by an incentive compatible and individually rational calibrated mechanism if and only if it is implementable by an incentive compatible and individually rational two-stage mechanism. That is, if and only if*

$$\vartheta(a, \theta, \omega) = \mu_0(\omega) f(\theta|\omega) \int_{\Delta(\Omega)} \alpha(a|\theta, \mu) \beta(d\mu|\omega), \quad (5)$$

for some Bayes plausible $\beta : \Omega \rightarrow \Delta(\Delta(\Omega))$ and incentive compatible and individually rational $\alpha : \Theta \times \Delta(\Omega) \rightarrow \Delta(A)$.

The proof of this and all results in this section can be found in [Appendix B](#).

In the single-agent case, [Theorem 2](#) shows that the calibrated mechanism design problem is equivalent to a standard mechanism design problem in which we restrict the designer to using a specific class of mechanisms; namely, incentive compatible and individually rational two-stage mechanisms. As we explained above, a mechanism ϕ and its calibrated information structure π_ϕ can be seen as actually inducing a joint distribution over $A \times \Theta \times \Omega \times \Delta(\Omega)$. [Theorem 2](#) implies this joint distribution admits two conditional independence properties. First, the allocation is conditionally independent of the state, conditional on the agent's type and the induced belief.²⁰ This follows from the signals s^* carrying no further information about the state than that what is contained in the agent's belief. Second, the designer disclosed belief is conditionally independent of the agent's type conditional on the state. In the static setting of [Section 2](#), this is because the calibrated information structure discloses information to the agent uniformly across her types. In the dynamic setting of [Section 5.1](#), this type-independent disclosure arises endogenously because the agent's experimentation opportunities are independent of her type.

Two-stage mechanisms solve calibrated mechanism design [Theorem 2](#) is of practical import as it provides a recipe of sorts for characterizing the designer's optimal calibrated mechanism (see the applications in [Section 4](#)). For each $\mu \in \Delta(\Omega)$, the designer chooses a mechanism $\alpha(\cdot|\cdot, \mu) : \Theta \rightarrow \Delta(A)$ that maximizes his expected payoff when the designer believes μ is the distribution of states, and

¹⁹Formally, define the belief distribution induced by β , $\tau_\beta = \mu_0 \otimes \beta$. The claim is that $\mathbb{E}_{\tau_\beta} [\mu] = \mu_0$.

²⁰This is a consequence of Bayes rule: beliefs are a sufficient statistic for ω . Hence, conditional on (θ, μ) , the allocation rule carries no more information about the state.

subject to the agent's incentive compatibility and individually rational constraints conditional on the designer's belief being μ . Proceeding in this way, we obtain the designer's value function $W : \Delta(\Omega) \rightarrow \mathbb{R}$. The optimal Blackwell experiment obtains from the concavification of W . We illustrate this procedure with two examples:

Example 1 (continued). *Consider again the seller-buyer example, in which the buyer is privately informed about her value for the good and the seller knows the demand state. By Theorem 2, we can find the seller's optimal calibrated mechanism as follows. First, equate μ with the probability that the state is H . For each $\mu \in [0, 1]$, consider the following problem:*

$$W(\mu) \equiv \max_{(q, t) : V \rightarrow [0, 1] \times \mathbb{R}} \mu \left(\frac{2}{3} t(1) + \frac{1}{3} t(1/2) \right) + (1 - \mu) \left(\frac{1}{3} t(1) + \frac{2}{3} t(1/2) \right) \quad (6)$$

$$\text{s.t. } \begin{cases} (\forall v \in \{1/2, 1\}) & vq(v) - t(v) \geq 0 \\ (\forall v, v' \in \{1/2, 1\}, v \neq v') & vq(v) - t(v) \geq vq(v') - t(v') \end{cases}.$$

That is, the seller chooses an incentive compatible and individually rational selling mechanism that maximizes his expected revenue when his belief is μ . Because ω is not payoff relevant to the buyer—it is just statistical information about the buyer's valuation—and the mechanism does not depend on state, the buyer's belief about ω does not enter her incentive constraints.

The solution to the seller's problem in Equation 6 is simple: the seller posts a price of $1/2$ when $\mu \leq 1/2$ and a price of 1 when $\mu > 1/2$. Hence, the seller's value function is given by

$$W(\mu) = \max \left\{ \frac{1}{2}, \mu \frac{2}{3} + (1 - \mu) \frac{1}{3} \right\},$$

and is illustrated by the solid line in blue on Figure 1. In words, the seller either sells the good at a price of $1/2$ and the buyer buys with probability 1, or he sells the good at a price of 1 and the buyer buys whenever her value is 1, which happens with the probability in the second argument of the max.

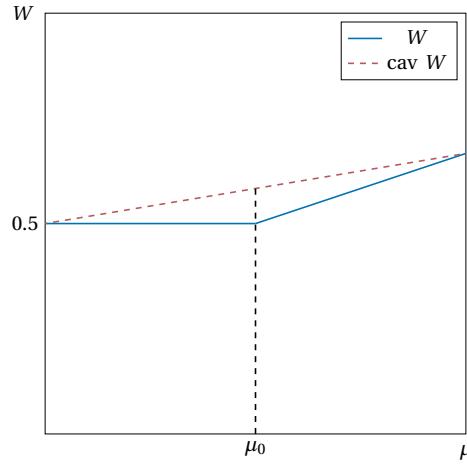


Figure 1: Seller's payoff in Example 1.

The optimal calibrated mechanism can be read from the concavification of W , which is the dashed, red line in Figure 1: The seller first reveals the state to the agent, and offers a price of $1/2$ when $\omega = L$ and a price of 1 when $\omega = H$.

Example 1 illustrates a more general principle that provides additional intuition for **Theorem 1**. In the private values case and when $N = 1$, the designer's value function $W : \Delta(\Omega) \rightarrow \mathbb{R}$ is convex. As Equation 6 illustrates, the designer maximizes a linear function in beliefs subject to constraints that do not depend on the induced belief. Convexity of W implies full disclosure is (weakly) optimal, and **Theorem 1** follows.

Example 2 (Horizontal differentiation). Consider a seller who owns a good of unknown type, $\omega \in \{L, R\}$, and a buyer whose private information is indexed by $\Theta = \{\theta_1, \theta_2, \theta_3\}$. Assume the good's type (the state) and the buyer's types are independent, and equally likely. Table 5 describes the buyer's value for the seller's good as a function of hers and the good's type, $v(\theta, \omega)$. When the good is $\omega = L$, the buyer of type θ_3 has the highest value for the good, whereas when the good is $\omega = R$, the buyer of type θ_3 has the lowest value for the good.

	θ_1	θ_2	θ_3
$\omega = L$	1	2	3
$\omega = R$	2	2	1

Table 5: Buyer's values.

Suppose the buyer's utility is quasilinear, that is, $u(q, t, \theta, \omega) = qv(\theta, \omega) - t$, and the seller wishes to maximize his revenue. Furthermore, assume the buyer's outside option is no trade.

Consider first the optimal mechanism the designer would offer absent the calibration constraint, depicted in the top panel of Table 6. This mechanism asks types θ_2 and θ_3 for a payment of 2 and allocates the good with probability 1, regardless of its kind. Instead, it asks the buyer of θ_1 to pay 1 in exchange for getting the good only when it is of her favorite kind ($\omega = R$).

	θ_1	θ_2	θ_3
$\omega = L$	(0, 1)	(1, 2)	(1, 2)
$\omega = R$	(1, 1)	(1, 2)	(1, 2)

	$\{(\theta_1, (0, 1)), (\theta_2, (1, 2)), (\theta_3, (1, 2))\}$	$\{(\theta_1, (1, 1)), (\theta_2, (1, 2)), (\theta_3, (1, 2))\}$
$\omega = L$	1	0
$\omega = R$	0	1

Table 6: Trade probabilities and transfers in the optimal mechanism (top); calibrated information structure (bottom).

The bottom panel of Table 6 depicts the information structure calibrated to the optimal mechanism. It sends two signals: when the good is L , the buyer can choose to either not get the good and pay 1, or get the good and pay 2. Instead, when the good is R , the buyer is choosing between paying 1 or 2 to obtain the good with probability 1.

Under the calibrated information structure, the optimal mechanism is neither incentive compatible nor individually rational. When the good is L , the buyer would prefer to choose (1, 1) regardless of her type. Instead, when the good is R , the buyer of θ_1 would quit the mechanism instead of paying 1 and getting nothing.

To characterize the optimal calibrated mechanism, we rely again on two-stage mechanisms. Equate μ with the probability that the good is R . Note that because states and types are independent, if the seller assigns probability μ to the state being R , so does the buyer (and vice versa). For each $\mu \in [0, 1]$, the seller solves the following problem

$$W(\mu) \equiv \max_{(q, t): \Theta \rightarrow [0, 1] \times \mathbb{R}} \sum_{\theta \in \Theta} \frac{1}{3} t(\theta) \quad (7)$$

$$\text{s.t. } \begin{cases} (\forall \theta \in \{\theta_1, \theta_2, \theta_3\}) & q(\theta) \mathbb{E}_\mu v(\theta, \cdot) - t(\theta) \geq 0 \\ (\forall \theta, \theta' \in \{\theta_1, \theta_2, \theta_3\}, \theta' \neq \theta) & q(\theta) \mathbb{E}_\mu v(\theta, \cdot) - t(\theta) \geq q(\theta') \mathbb{E}_\mu v(\theta', \cdot) - t(\theta') \end{cases}.$$

In this case, the seller's objective function does not depend on the induced belief μ as types and states are independent. Instead, the buyer's incentive and individual rationality constraints do depend on μ as the state is payoff relevant. The solution to the problem in Equation 7 is a posted price, whose value depends on μ . For instance, when $\mu \in \{0, 1\}$, the optimal price is 2 and the seller's revenue is $4/3$. Instead, when $\mu = 2/3$, the optimal price is $5/3$ and profits are maximal and equal to $5/3$. Indeed, when $\mu = 2/3$, the heterogeneity across buyer types is minimized (and hence, their rents), and by setting $p = 5/3$ all buyer types buy. The blue line in Figure 2 depicts the seller's expected profit as a function of his belief μ .

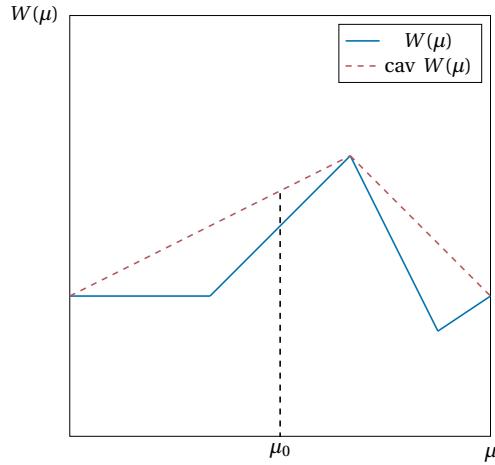


Figure 2: Seller's profit in the two-stage mechanism

The optimal calibrated mechanism can be read from the concavification of W at $\mu_0 = 1/2$, depicted by the dashed red line in Figure 2. The seller provides the buyer with partial information about the good: He either reveals the good is L and sells the good at a price of 2, or he obfuscates the good—inducing a belief of $2/3$ —and sets a price of $5/3$.

Another consequence of Theorem 2 is that without loss of generality, we can focus on calibrated mechanisms with finite calibrated information structures:

Corollary 1 (Support of calibrated information structures). *It is without loss of generality to restrict attention to two-stage mechanisms that induce at most $|\Omega|$ beliefs.*

In other words, it is without loss of generality to focus on calibrated mechanisms that induce at most $|\Omega|$ allocation rules.

Multiple agents and generalized two-stage mechanisms In the case of multiple agents, we can also interpret a calibrated mechanism as conveying to each agent i both the information she should have about the state upon seeing signal s_i^* , $\mu_i(\theta_i, s_i^*)$, and her interim allocation rule, $s_i^* : \Theta_i \rightarrow \Delta(A_i)$. However, two differences arise relative to the single-agent case: First, each agent i receives her information privately from that of other agents. Second, even if the agents put together the information they receive, this is not enough to learn the ex-post allocation rule, that is, the map from type profiles to allocations. After all, each agent i observes her interim allocation rule alone. These differences are natural when we think of calibrated mechanisms as capturing the information agents stand to learn from experimenting with the mechanism: There is no reason all agents will learn the same information, and from observing her own allocations, and not those of others, an agent can only learn about her interim allocation rule, not the ex-post one.

These observations together imply that to describe the analogue of a two-stage mechanism in multi-agent settings we need to (i) allow for agent-by-agent information disclosure, and (ii) keep track that the interim allocation rules are consistent with the same ex-post allocation rule. These considerations motivate the following generalization of a two-stage mechanism:

Definition 4 (Generalized two-stage mechanism). *A generalized two-stage mechanism is a mapping $\psi : \Theta \times \Omega \rightarrow \Delta(\Delta(\Omega)^N \times A)$ for which a tuple of mappings*

$$\beta : \Omega \rightarrow \Delta(\Delta(\Omega)^N), \quad \alpha_i : \Theta_i \times \Delta(\Omega) \rightarrow \Delta(A_i), \quad \alpha : \Theta \times \Omega \times \Delta(\Omega)^N \rightarrow \Delta(A),$$

exist such that:

1. *For all $(\theta, \omega) \in \Theta \times \Omega$, and all measurable subsets $(\tilde{\Delta}_i)_{i=1}^N \subset \Delta(\Omega)^N$, we have*

$$\psi(\times_{i=1}^N \tilde{\Delta}_i \times \{a\} | \theta, \omega) = \int_{\times_{i=1}^N \tilde{\Delta}_i} \alpha(a | \theta, \omega, \mu_1, \dots, \mu_N) \beta(d(\mu_1, \dots, \mu_N) | \omega)$$

2. *The Blackwell experiment β is Bayes plausible,*

3. *For all $i \in \{1, \dots, N\}$, the interim allocation rule α_i satisfies that for all measurable subsets $\tilde{\Delta}$ of $\Delta(\Omega)$ and all $(a_i, \theta_i, \omega) \in A_i \times \Theta_i \times \Omega$*

$$\int_{\tilde{\Delta} \times \Delta(\Omega)^{N-1}} \left\{ \alpha_i(a_i | \theta_i, \mu_i) - \mathbb{E}_{f_{-i}(\cdot | \omega)} \left[\sum_{a_{-i} \in A_{-i}} \alpha(a_i, a_{-i} | \theta_i, \theta_{-i}, \omega, \mu_i, \mu_{-i}) \right] \right\} \beta(d(\mu_i, \mu_{-i}) | \omega) = 0.$$

We say the generalized two-stage mechanism is incentive compatible and individually rational if for all $i \in \{1, \dots, N\}$, on the support of $\mu_0 \otimes \beta$, $\alpha_i(\cdot | \cdot, \mu_i)$ is incentive compatible and individually rational for agent i when she learns μ_i .

As anticipated, generalized two-stage mechanisms differ from two-stage mechanisms in three ways when $N > 1$. First, because disclosures are private, the experiment β now outputs a profile of beliefs, one for each agent. As shown in Arieli et al. (2024), β is Bayes plausible if and only if for each agent i , the marginal Blackwell experiment β_i is Bayes plausible. Second, while the individual interim allocation rule α_i only depends on the disclosed belief to agent i , μ_i , and not the state, the ex-post

allocation rule α may depend on the state, even conditional on the belief profile (μ_1, \dots, μ_N) . The reason is that this belief profile is no longer a sufficient statistic for the ex-post allocation rule as each agent i only observes their interim allocation. Third and relatedly, we need to keep track of both the interim allocation rules $(\alpha_i)_{i=1}^N$ and the ex-post allocation rule α to check that the interim allocation rules are consistent with the same mechanism. An interesting question for future work would be to characterize which interim allocation rules $(\alpha_i)_{i=1}^N$ are consistent with some ex-post allocation rule α , so that one could focus on the interim allocation rules alone.

As we show in [Proposition 1](#), a calibrated mechanism induces a generalized two-stage mechanism:

Proposition 1. *If outcome distribution $\vartheta \in \Delta(A \times \Theta \times \Omega)$ is implementable by an incentive compatible and individually rational calibrated mechanism, then it is implementable by an incentive compatible and individually rational generalized two-stage mechanism.*

In contrast to the single-agent case, not every outcome distribution implemented by a generalized two-stage mechanism can be implemented by a calibrated mechanism. On the one hand, no agent's beliefs are a sufficient statistic for the information the mechanism leaks about the state, so that the allocation rule $\bar{\alpha}$ may still leak information about the state or others' beliefs, which in turn leak information about the state. On the other hand, because in a calibrated mechanism each agent learns her interim allocation rule conditional on (ω, ε) , the incentive and participation constraints associated to a generalized two-stage mechanism are weaker than those implied by a calibrated mechanism whenever multiple interim allocation rules underlie the same belief: Even if the average interim allocation rule α_i is incentive compatible and individually rational, each of the interim allocation rules underlying that average need not be.

4 Applications

In this section, we study optimal calibrated mechanism design in canonical mechanism design settings with quasilinear utilities. We first consider the case of a single agent, with single-dimensional types and allocations, and supermodular payoffs. In [Section 4.1](#), we show that if the order of types is state independent, then optimal two-stage mechanisms fully reveal the state, whereas this conclusion can be reversed when the order of types is state-dependent. In [Section 4.2](#), we compare optimal calibrated mechanism design against the Myersonian benchmark. Lastly, we analyze a multi-agent application in [Section 4.3](#).

4.1 Calibrated Screening

We consider the following version of the model in [Section 2](#). Suppose $N = 1$ and let $\Theta = [\underline{\theta}, \bar{\theta}]$ denote the set of types. Assume θ is distributed according to a full support distribution F with density f . Hence, throughout, we consider the case in which the agent's type is independent of ω . Denote the set of allocations by $A = [0, \bar{q}] \times \mathbb{R}$, where $q \in [0, \bar{q}]$ is the (physical) allocation and $t \in \mathbb{R}$ is a payment from the agent to the designer.²¹

²¹In contrast to the model of [Section 2](#), we are assuming the set of types and allocations to be intervals in the real line. The results in the previous sections go through with richer type and allocation spaces, at the cost of more notation.

The agent's and the designer's payoffs are given by $u(q, \theta, \omega) - t$ and $w(q, \theta, \omega) + t$, respectively. Assume that if the agent does not participate, then the outside option is $a_\emptyset = (0, 0)$, and that this yields a payoff of 0 to both the designer and the agent. Throughout, we assume that for each $\omega \in \Omega$, the family of functions $\{\theta \mapsto u(q, \theta, \omega) : q \in [0, \bar{q}]\}$ is equi-Lipschitz on Θ : a positive constant L_ω exists such that for all $\theta, \theta' \in \Theta$ and $q \in [0, \bar{q}]$, $|u(q, \theta, \omega) - u(q, \theta', \omega)| \leq L_\omega |\theta - \theta'|$.²² Furthermore, the analysis that follows restricts attention to mechanisms that do not randomize on the allocation (beyond the inherent randomness of $\Omega \times [0, 1]$). [Remark 1](#) at the end of this section discusses settings in which this is not a restriction and how to generalize the observations herein when random allocations are allowed.

Our goal is to characterize the designer optimal calibrated mechanism and how its properties depend on how the state affects the order of types.

State-independent type ranking We consider first the case in which the order of types is independent of the state. Formally, assume that for all $\omega \in \Omega$, the function $u(\cdot, \omega)$ is supermodular in (q, θ) . That is, in all states, the agent with higher value of θ values q more. These assumptions are satisfied, for instance, for $u(q, \theta, \omega) = \theta \omega q$ or $u(q, \theta, \omega) = (\theta + \omega)q$.

By [Theorem 2](#), we can characterize the optimal calibrated mechanism via two-stage mechanisms. To do so, we solve the problem “backward”: For each $\mu \in \Delta(\Omega)$ the designer may induce about the state, the designer chooses an optimal direct mechanism $(q_\mu, t_\mu) : \Theta \rightarrow A$. This determines the designer's value function $W : \Delta(\Omega) \rightarrow \mathbb{R}$. We obtain the designer's optimal Blackwell experiment by studying the properties of W .

Given belief μ , define the agent's and the designer's (expected) payoff at (q, t, θ) as follows:

$$u(q, \theta | \mu) \equiv \sum_{\omega \in \Omega} \mu(\omega) u(q, \theta, \omega), \quad w(q, \theta | \mu) \equiv \sum_{\omega \in \Omega} \mu(\omega) w(q, \theta, \omega).$$

Thus, conditional on inducing belief μ , the designer's problem can be written as follows:

$$W(\mu) \equiv \max_{(q, t) : \Theta \rightarrow A} \int_{\Theta} [w(q(\theta), \theta, \mu) + t(\theta)] F(d\theta) \quad (8)$$

$$\text{s.t. } \begin{cases} (\forall \theta \in \Theta) & u(q(\theta), \theta | \mu) - t(\theta) \geq 0 \\ (\forall \theta, \theta' \in \Theta) & u(q(\theta), \theta | \mu) - t(\theta) \geq u(q(\theta'), \theta | \mu) - t(\theta') \end{cases}.$$

Our assumptions imply that $u(\cdot | \mu)$ is supermodular in (q, θ) . It follows that the designer can only choose among those $q : \Theta \rightarrow [0, \bar{q}]$ that are (weakly) increasing in θ . Let Q_1 denote the set of all such $q(\cdot)$. Furthermore, at the optimum, the participation constraint of $\theta = \underline{\theta}$ binds.

Define the virtual surplus at (q, θ, ω) as follows:

$$J((q, \theta, \omega); F) = w(q, \theta, \omega) + u(q, \theta, \omega) - u_2(q, \theta, \omega) \frac{1 - F(\theta)}{f(\theta)},$$

where u_2 is the derivative of u against its second coordinate; the equi-Lipschitz assumption implies it exists almost everywhere. Then, conditional on inducing belief μ , the designer's payoff can be written

²²This assumption ensures the Lipschitz continuity of the agent's indirect utility function when the allocation space is infinite. Because Ω is finite, requiring the condition to hold state-by-state suffices.

as follows:

$$W(\mu) = \max_{q \in Q \uparrow} \int_{\Theta} \mathbb{E}_{\mu} [J((q(\theta), \theta, \omega); F)] F(d\theta). \quad (9)$$

Note the objective is linear in μ and the constraint set is independent of μ . We conclude that W is convex, as it is the maximum of linear functionals in μ . It follows that full disclosure is an optimal experiment for the designer. Equivalently, an optimal calibrated mechanism exists in which the designer chooses the mechanism ϕ_{full}^D , where for all $\omega \in \Omega$, $\phi_{full}^D(\cdot, \omega, \cdot) : \Theta \times [0, 1] \rightarrow A$ is the optimal deterministic mechanism when it is common knowledge that the state is ω .

Proposition 2 summarizes the above discussion:

Proposition 2 (State-Independent Type Ranking). *In a single-dimensional screening problem with state-independent type ranking, the designer can do no better than choosing ϕ_{full}^D among deterministic mechanisms.*

By **Proposition 2**, in screening problems with state-independent ranking of types across states, the calibration constraint makes any pooling of mechanisms across states unprofitable.²³ Remarkably, this result holds for any designer objective, such as profit, revenue, or efficiency. It also requires no regularity assumptions on the type distribution, as we do not obtain the result by looking at the relaxed problem. Instead, our argument relies on the restriction to deterministic mechanisms (conditional on the induced belief), which in turn delivers that the set of implementable allocations does not depend on the induced belief. **Remark 1** discusses conditions under which (i) the restriction to deterministic mechanisms is without loss of optimality, and (ii) the set of implementable allocations does not depend on the induced belief, even when randomized mechanisms are allowed. Readers interested in the case of state-dependent ranking can skip this remark with little loss of continuity.

Remark 1 (**Proposition 2** without deterministic mechanisms). *Under our assumptions, deterministic mechanisms are without loss of optimality if the agent's payoff is linear in q and the designer's payoff is concave in q . (See [Section 4.2](#) for yet another condition.) However, the driving force behind **Proposition 2** is that the designer's constraint set does not depend on the induced belief. The state-by-state supermodularity assumption and the restriction to deterministic mechanisms is one way to ensure this is the case. We now discuss two other cases in which the designer's constraint set does not depend on the induced belief and thus **Proposition 2** holds for the optimal (not necessarily deterministic) calibrated mechanism.*

First, suppose the agent's payoff is linear in q , so that $u(q, \theta, \omega) = q v(\theta, \omega)$, where $v(\cdot, \omega)$ is increasing for all ω . Then, the set of implementable lotteries over q when the belief is μ is given by:

$$Q \uparrow, \text{random} = \{\xi : \Theta \rightarrow \Delta([0, \bar{q}]) : \mathbb{E}_{\xi(\theta)} [q] \text{ is increasing in } \theta\}.$$

In this case, we obtain that the designer cannot do any better than choosing ϕ_{full} , which is the

²³The calibrated mechanism which state-by-state implements the optimal direct mechanism under common knowledge of the state may reveal less than full information about the state, e.g., because at ω and ω' the same direct mechanism is optimal. The point is that pooling those states does not weaken the incentive constraints of the agent, so it is as if the designer were forced to reveal the state.

mechanism that implements in each state ω the optimal mechanism under common knowledge that the state is ω . This relates to the results in Szabadi (2018) and Yamashita (2018), who study the optimal mechanism design preceded by public information disclosure. Both papers consider settings in which the agent's payoff is linear in q and obtain that full disclosure is optimal when the ranking of types is independent of the state.²⁴

Second, suppose the agent's payoff has the form

$$u(q, \theta, \omega) = b(\theta)c(\omega)v(q) + k_1(q, \omega) + k_2(\theta, \omega),$$

where b is increasing in θ and $c(\cdot)$ does not change sign on Ω . Under this assumption, $u(q, \theta|\mu)$ satisfies monotonic expectational differences for all $\mu \in \Delta(\Omega)$ (see, e.g., Kartik et al., 2024). Consequently, one can define a linear order \succeq over $\Delta([0, \bar{q}])$ as follows: $\xi \succeq \xi'$ if $u(\xi, \theta|\mu) - u(\xi', \theta|\mu)$ is increasing in θ , where $u(\xi, \theta|\mu)$ is the linear extension of $u(\cdot, \theta|\mu)$ to $\Delta([0, \bar{q}])$. This linear order implies the ranking of types is state independent. Indeed, the analog of $Q_{\uparrow, \text{random}}$ is the set of all $\xi : \Theta \rightarrow \Delta([0, \bar{q}])$ such that $\theta \geq \theta'$ implies $\xi(\theta) \succeq \xi(\theta')$, which is again independent of the induced belief.

State-dependent type ranking Example 2 illustrates that when the ranking of types is not uniform across states, full transparency may not be optimal.²⁵ We now provide a more systematic analysis of this phenomenon, using the previous results. To provide the starker contrast with Proposition 2, we consider a setting that shares a key feature of Example 2: we can partition $\Delta(\Omega)$ into two regions such that within each region the ranking of types—as determined by $u(q, \theta|\mu)$ —is the same, but it differs across regions.

Concretely, suppose that $\Omega = \{\omega_1, \omega_2\}$. Furthermore, assume

$$u(q, \theta, \omega) = \begin{cases} q\theta & \text{if } \omega = \omega_2 \\ q(c - b\theta) & \text{otherwise,} \end{cases},$$

where $c \in \mathbb{R}$, and $b > 0$. Identify beliefs with the probability that the state is ω_2 and define

$$\hat{\mu} = \frac{b}{1 + b}.$$

For $\mu < \hat{\mu}$, we have that $u(q, \theta|\mu)$ is decreasing in θ , whereas if $\mu > \hat{\mu}$, then $u(q, \theta|\mu)$ is increasing in θ .

Consider now the designer's optimal payoff $W : \Delta(\Omega) \rightarrow \mathbb{R}$ as a function of the different beliefs he may induce. When $\mu \geq \hat{\mu}$, the designer's payoff can be obtained by solving the program in Equation 9 as before. Instead, when $\mu < \hat{\mu}$, the designer's payoff can be obtained by solving a problem analogous to Equation 9, but where the space of implementable allocations is the set of decreasing q , Q_{\downarrow} , and the participation constraint of $\bar{\theta}$ binds. It follows that W is convex on $[0, \hat{\mu}]$ and $(\hat{\mu}, 1]$. Thus, the support of designer's optimal experiment is included in $\{0, \hat{\mu}, 1\}$.

²⁴Restricting attention to deterministic mechanisms, Ottaviani and Prat (2001) obtain the optimality of full disclosure without such a linearity assumption. Their model, however, is different from ours and that of the aforementioned papers: the agent's type is not payoff relevant and the agent's type and the state are affiliated. Thus, while related in spirit, Proposition 2 is distinct from their result.

²⁵See Szabadi (2018) for a similar observation.

Proposition 3 (State-dependent type ranking). *Suppose the agent's payoff satisfies the assumptions above. If $(1 - \hat{\mu})W(0) + \hat{\mu}W(1) \geq W(\hat{\mu})$, full transparency is optimal. Otherwise, full transparency is not optimal: if $\mu_0 < \hat{\mu}$, it is optimal to split μ_0 to 0 and $\hat{\mu}$; if $\mu_0 > \hat{\mu}$, it is optimal to split μ_0 to $\hat{\mu}$ and 1. In particular, if $W(\hat{\mu}) > \max\{W(0), W(1)\}$, then full transparency is not optimal.*

By Proposition 3, whether full transparency is optimal depends on the designer and agent's payoffs and the type distribution, but only through their impact on the value the function W takes at points $\{0, \hat{\mu}, 1\}$. At $\hat{\mu}$, the agent earns no rents—as $u(\cdot | \hat{\mu})$ is constant across types—which pushes against full transparency. At the same time, efficiency may dictate the designer to condition the allocation rule on the state, which favors information disclosure. The piecewise convexity of W implies that if $W(\hat{\mu})$ dominates W at the extreme beliefs, the rent extraction motive dominates and the designer does not engage in full disclosure.

4.2 Comparison with Myersonian Mechanism Design

We now compare optimal calibrated mechanism design and the Myersonian benchmark. In the Myersonian benchmark, the designer is not concerned with the information the mechanism reveals about the state, and hence provides a natural upper bound on the designer's payoffs in calibrated mechanism design. The gap between the designer's optimal payoff across both benchmarks quantifies the loss from the calibration constraint. If no gap exists, the calibration constraint is non-binding and an optimal calibrated mechanism can be found solving the Myersonian benchmark. Instead, if a gap exists, the optimal mechanism in the Myersonian benchmark reveals information about the state in a way that it fails to be incentive compatible or individually rational under calibration.

Myersonian benchmark In the Myersonian benchmark, the designer chooses a direct mechanism $(\xi, t) : \Theta \times \Omega \rightarrow \Delta([0, \bar{q}]) \times \mathbb{R}$ subject to incentive and participation constraints that must hold on average across states under the prior μ_0 .²⁶ Formally,

$$W_{\text{My}} \equiv \max_{(q, t) : \Theta \times \Omega \rightarrow [0, \bar{q}] \times \mathbb{R}} \int_{\Theta} \mathbb{E}_{\mu_0} [w(\xi(\theta, \omega), \theta, \omega) + t(\theta, \omega)] F(d\theta) \quad (\text{OPT}_{\text{My}})$$

$$\text{s.t. } \begin{cases} (\forall \theta \in \Theta) \mathbb{E}_{\mu_0} [u(\xi(\theta, \omega), \theta, \omega) - t(\theta, \omega)] \geq 0 \\ (\forall \theta, \theta' \in \Theta) \mathbb{E}_{\mu_0} [u(\xi(\theta, \omega), \theta, \omega) - t(\theta, \omega)] \geq \mathbb{E}_{\mu_0} [u(\xi(\theta', \omega), \theta, \omega) - t(\theta', \omega)] \end{cases},$$

where $w(\xi, \theta, \omega)$ and $u(\xi, \theta, \omega)$ are the linear extensions of $w(\cdot, \theta, \omega)$ and $u(\cdot, \theta, \omega)$, respectively.

Program OPT_{My} is a mechanism design problem with a multidimensional allocation, corresponding to assigning (a distribution over) q in each state. As a result, the distinction between the Myersonian benchmark and optimal calibrated design shows in the monotonicity requirements the allocation $\xi(\theta, \omega)$ must satisfy for a transfer $t : \Theta \times \Omega \rightarrow \mathbb{R}$ to exist that implements $\xi(\theta, \omega)$. Indeed, implementability of $\xi : \Theta \times \Omega \rightarrow \Delta([0, \bar{q}])$ is equivalent to *integral monotonicity* (Rochet, 1987; Pavan et al., 2014):²⁷

$$(\forall \theta, \theta' \in \Theta) \int_{\theta'}^{\theta} \int_{\Omega} [u_2(\xi(s, \omega), s, \omega) - u_2(\xi(\theta', \omega), s, \omega)] d\mu_0 ds \geq 0, \quad (\text{IM})$$

²⁶Because payoffs are quasilinear, considering mechanisms that do not randomize on transfers is without loss of generality.

²⁷The equi-Lipschitz condition on u ensures we can take the derivative inside the integral.

where recall u_2 is the derivative of u in its second coordinate.

Comparison with calibrated mechanism design To facilitate the comparison with [Proposition 2](#), we focus on deterministic mechanisms $(q, t) : \Theta \times \Omega \rightarrow [0, \bar{q}] \times \mathbb{R}$. Remarkably, even if the agent's payoff net of transfers, $u(q, \theta, \omega)$, is supermodular in (q, θ) for all $\omega \in \Omega$, the characterization of the set of implementable $q(\cdot)$ cannot be simplified beyond integral monotonicity without further assumptions. Because integral monotonicity is a global, implicitly defined constraint, verifying implementability and computing the optimal mechanism is more computationally involved in the Myersonian benchmark than in calibrated mechanism design. Indeed, [Proposition 2](#) implies the optimal deterministic calibrated mechanism coincides with the state-by-state optimal deterministic mechanism under this assumptions. In other words, the optimal deterministic calibrated mechanism can be obtained by selecting allocations $q(\cdot)$ that satisfy

$$Q_{\text{cal}} = \{q : \Theta \times \Omega \rightarrow [0, \bar{q}] : (\forall \omega \in \Omega) q(\cdot, \omega) \text{ is increasing}\}.$$

That is, the allocation in the optimal calibrated mechanism must satisfy monotonicity state-by-state. Instead, the optimal deterministic Myersonian mechanism can be obtained by selecting allocations $q(\cdot)$ that satisfy [Equation IM](#), which we denote by Q_{My} .

When $u(q, \theta, \omega)$ is supermodular in (q, θ) for all $\omega \in \Omega$, the above discussion implies the designer's optimal payoff in the Myersonian and calibration settings can be written as follows:

$$\begin{aligned} W_{\text{My}}^D &= \max_{q \in Q_{\text{My}}} \int_{\Theta} \mathbb{E}_{\mu_0} [J(q(\theta, \omega), \theta, \omega; F)] F(d\theta), \\ W_{\text{cal}}^D &= \max_{q \in Q_{\text{cal}}} \int_{\Theta} \mathbb{E}_{\mu_0} [J(q(\theta, \omega), \theta, \omega; F)] F(d\theta), \end{aligned} \quad (10)$$

where the superscript D in the objective is a reminder that we restrict attention to mechanisms that are deterministic conditional on the state, or the induced belief.

By reducing the comparison across settings to monotonicity requirements on the space of allocations, the above expressions provide us with an immediate way of comparing the designer's payoffs across settings. In particular, when the optimal Myersonian mechanism satisfies the state-by-state monotonicity constraints, we have that the calibration constraint entails no loss to the designer. We record this observation for future use:

Observation 1. *Suppose $u(q, \theta, \omega)$ is supermodular in (q, θ) for all $\omega \in \Omega$. Then, if the allocation rule in the Myersonian benchmark satisfies monotonicity state-by-state, $W_{\text{cal}}^D = W_{\text{My}}^D$.*

Two natural questions are under what conditions the solution to OPT_{My} is deterministic and satisfies state-by-state monotonicity. We answer them simultaneously by studying the relaxed program. Inspection of [Equation 10](#) reveals that if the virtual surplus is supermodular in (q, θ) for every ω , then the solution q_{rel} to the relaxed problem

$$W_{\text{rel}} = \max_{q : \Theta \times \Omega \rightarrow [0, \bar{q}]} \int_{\Theta} \mathbb{E}_{\mu_0} [J(q(\theta, \omega), \theta, \omega; F)] F(d\theta), \quad (11)$$

satisfies monotonicity state-by-state by Topkis' theorem. Moreover, a stochastic mechanism is

equivalent to a deterministic mechanism which depends on the random reports of a fictitious agent (Pavan et al., 2014). The virtual surplus in this fictitious setting coincides with that in the integrand on the right-hand side of [Equation 11](#)—the type reports of the fictitious agent are payoff irrelevant—and is maximized by q_{rel} .

Proposition 4 (Sufficient condition for no gap). *Suppose the virtual surplus $J((q, \theta, \omega); F)$ is supermodular in (q, θ) for all ω . Then, the designer's payoffs under the optimal Myersonian and calibrated mechanisms coincide.*

By contrast to [Proposition 2](#), [Proposition 4](#) relies on assumptions on the type distribution and the designer's payoff. As [Example 2](#) illustrates, the supermodularity of the virtual surplus can fail when the type distribution is not regular, creating a gap between the designer's payoff at the optimal Myersonian and calibrated mechanisms. [Example 3](#) illustrates such a gap can also arise when the designer's payoff is not supermodular:

Example 3 (Payoff gap when w is not supermodular). *Suppose states are binary, $\Omega = \{\omega_L, \omega_H\} = \{1, 3\}$, and equally likely. Suppose types are uniformly distributed, $\theta \sim U[0, 1]$. Finally, let $q \in [0, 1]$ denote the probability the seller's good is allocated. Payoffs are given by:*

$$u(q, \theta, \omega) = q\theta\omega$$

$$w(q, \theta, \omega) = 2(1 - 2\theta)q.$$

Note that w is increasing in q when $\theta < 1/2$ and decreasing in q when $\theta > 1/2$.²⁸ In this case, the virtual surplus evaluated at different states is:

$$J((q, \theta, \omega); F) = \begin{cases} (1 - 2\theta)q & \text{if } \omega = \omega_L \\ (2\theta - 1)q & \text{otherwise} \end{cases}.$$

In the Myersonian benchmark, implementable allocations are elements of Q_{My} , which in this case is equivalent to requiring that $\mathbb{E}_{\mu_0}[q(\cdot, \omega)\omega]$ is increasing. The optimal Myersonian allocation obtains from pointwise maximizing the virtual surplus, and is given by:

$$q_{\text{My}}(\theta, \omega) = \begin{cases} 1 & \text{if } \omega = \omega_L \text{ and } \theta < 1/2 \\ 1 & \text{if } \omega = \omega_H \text{ and } \theta > 1/2 \\ 0 & \text{otherwise} \end{cases}.$$

The designer's payoff under the Myersonian mechanism is 1/4.

By [Proposition 2](#), q_{My} cannot be implemented by a calibrated mechanism as it is not increasing state-by-state. Intuitively, when $\omega = \omega_L$, types above 1/2 would learn from the calibrated information structure that they do not obtain the good, whereas types below 1/2 do, and would misreport their types.

Instead, in the optimal calibrated mechanism, the designer sets $q_{\text{cal}}(\theta, \omega_H) = \mathbb{1}[\theta \geq 1/2]$ and sets

²⁸The designer can be viewed as an online advertising platform and the agent as an advertiser. State ω represents the click-through rate of an ad slot, and higher θ corresponds to a larger advertiser willing to pay more for exposure. The designer's payoff captures both the value created by advertising and the disutility from showing ads of large advertisers, e.g., due to user brand fatigue.

$q(\theta, \omega_L)$ to be constant in θ . The designer's payoff under calibration is $W_{cal} = 1/8 < W_{My}$.

4.3 Optimal Calibrated Auction

In this section, we consider a multiple agent application and study the design of the optimal calibrated auction. [Proposition 1](#) implies the optimal calibrated auction induces a generalized two-stage mechanism, and hence the optimal generalized two-stage mechanism provides an upper bound on the designer's optimal payoff under calibration. However, computing the optimal generalized two-stage mechanism is complicated because (i) no tractable characterization of joint distributions over posterior beliefs is available, and (ii) the allocation rule may condition on the state and not only the agents' beliefs. For that reason, our analysis below relies on [Observation 1](#): We show the optimal Myersonian auction can be implemented by fully revealing the state, and hence, remains incentive compatible and individually rational when the agents have access to the calibrated information structure. Below, we first specialize our multi-agent model and notation to the auction application and then link our assumptions to online advertising.

Suppose there is a single good for sale and the state is multidimensional, $\omega = (\omega_i, \omega_{0i})_{i \in [N]} \in \mathbb{R}_+^{2N}$, and distributed according to prior distribution μ_0 . Suppose that for all $i \in [N]$, $\Theta_i = [0, 1]$, with $\theta_i \sim F_i$ with full-support density f_i . That is, we are assuming agents' types are independent of the state, and hence, independent across each other. Denote by $q_i \in [0, 1]$ the probability agent i is allocated the good, and note that feasibility implies that $0 \leq \sum_{i=1}^N q_i \leq 1$.

We assume the agents' and the designer's utilities are quasilinear in transfers. Agent i 's payoff net of transfers is $u_i(q_i, \theta_i, \omega) = q_i(\omega_i \theta_i + \omega_{0i})$. Thus, state components ω_i capture the value responsiveness to agent's private information, whereas state components ω_{0i} capture the overall shift. The state components can be correlated (and asymmetric) across agents, allowing for interdependent values. The designer's payoff net of transfers is $w(q, \theta, \omega) = \sum_i q_i w_i(\theta, \omega)$ for some functions $(w_i)_{i \in [N]}$. Below, we study the designer-optimal calibrated mechanism.

To fix ideas, consider the following mapping to an online advertising environment. The designer is an advertising platform, and the good is an advertising slot on a given webpage targeted to a selected category of users in a given week. Agents are firms that wish to display their ads, and their private types represent the expected revenue from a click on their ad. State components ω_i could capture individual click-through rates or match values, while state components ω_{0i} could capture individual display values, that is, the expected revenue from an ad being displayed irrespective of whether it is clicked (for instance, due to brand-building effects). The state is observed through proprietary data available to the platform and can be used in the design of the auction. The platform values the resulting revenue but may also have additional efficiency considerations, summarized by w_i .

As anticipated, we characterize the optimal calibrated mechanism by showing that it coincides with the Myersonian optimal one. To this end, consider the Myersonian problem, in which the designer chooses $(q(\theta, \omega), t(\theta, \omega)) \in [0, 1]^N \times \mathbb{R}^N$. Because agent i 's payoff is linear in θ_i , arguments analogous to those in [Section 4.2](#) imply a feasible $q(\theta, \omega)$ is implementable if and only if for all i , $\mathbb{E}_{F_{-i}, \mu_0}[q_i(\theta_i, \theta_{-i}, \omega) \omega_i]$ is increasing in θ_i . In a slight abuse of notation, denote by Q_{My} the set of all such functions and define

the virtual surplus as

$$J((q, \theta, \omega); F) = \sum_{i=1}^N q_i(\theta, \omega) \left(w_i(\theta, \omega) + \left(\theta_i - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \right) \omega_i + \omega_{0i} \right). \quad (12)$$

Standard arguments imply the individual rationality constraint of $\theta_i = 0$ binds for all i , and an optimal mechanism solves

$$W_{\text{My}} = \max_{q \in Q_{\text{My}}} \int_{[0,1]^N} \mathbb{E}_{\mu_0} [J(q(\theta, \omega), \theta, \omega; F)] f(\theta) d\theta.$$

Proposition 5 (No gap in regular auctions). *Suppose that (i) for all $i \in \{1, \dots, N\}$, F_i is Myerson regular, and for all i, j, θ , and ω , $w_{i\theta_i}(\theta, \omega) \geq 0$, $w_{i\theta_i}(\theta, \omega) \geq w_{j\theta_i}(\theta, \omega)$.²⁹ Then, $W_{\text{cal}} = W_{\text{My}}$.*

The proof of Proposition 5 in Appendix D.2 shows that under our assumptions the optimal Myersonian mechanism can be obtained by solving the relaxed program. Importantly, the assumption that $w_{i\theta_i}(\theta, \omega) \geq w_{j\theta_i}(\theta, \omega)$ ensures that an increase in agent i 's type increases the designer's payoff of giving the object to agent i by more than the value of giving it to other agents. This, in turn, ensures agent i 's allocation probability is increasing in her type.

Viewed through the lens of the online advertising example, Proposition 5 implies that in regular environments, while the advertising platform benefits from having the data on click-through rates and display values, it does not benefit from the informational advantage over bidders that such data entails. Its objective is maximized by making the click-through rates and display values readily available to bidders and running optimal auctions in all instances.

5 Microfoundation

In this section, we provide a microfoundation for calibrated mechanism design by analyzing the outcome distributions that can arise when an agent repeatedly engages with the same mechanism (Section 5.1) and contrast this to what can be implemented when the designer can offer the agent a *fully dynamic* mechanism (Section 5.2). To keep the presentation simple, we present the results with minimal notation, and refer the reader to Appendix C for details.

Throughout, we consider the case of a single agent, whose type (i) is redrawn each period from the same distribution and (ii) is independent of the state. The reason for (i) is as follows. When the designer offers the agent a fully dynamic mechanism, the revelation principle implies that it is without loss of generality for the designer to ask the agent for type reports. Moreover, logic similar to that in Myerson (1986) implies that the designer only elicits one type report when the agent's type is persistent, and hence, the agent has no possibility of experimenting with the mechanism. Hence, to put repeated and dynamic mechanisms on a more similar footing, assuming the agent's type is redrawn each period is necessary. However, when the agent's type is repeatedly drawn from a distribution that depends on the state, the agent learns about the state both through her own type and her allocations in the mechanism.³⁰ Thus, we assume (ii) so that the agent learns about the state only through her

²⁹In the statement, $w_{k\theta_i}$ denotes the derivative of w_k with respect to θ_i , for $k \in \{1, \dots, N\}$.

³⁰To be sure, Theorem 3 extends to the case in which the agent's type is fully persistent and correlated with the state.

interaction with the mechanism. Lastly, we consider the single-agent case as extending the results in this section to multiple agents requires addressing subtle issues in strategic experimentation, which we plan to pursue in future work.

5.1 Repeated Interactions with a Mechanism

We consider first the case in which the agent interacts repeatedly with the same mechanism ϕ in each period of an infinite horizon interaction. In line with [Section 2](#), a repeated mechanism is a mapping

$$\phi : M \times \Omega \times \mathcal{E} \rightarrow \Delta(A),$$

where M is a finite set of messages and \mathcal{E} is a finite set endowed with some measure, denoted η . The results in [Section 3](#) imply that assuming \mathcal{E} is finite is without loss of generality and it simplifies the proofs. In contrast to [Section 2](#), we allow the mechanism to have an arbitrary message space. The reason is that we cannot invoke the revelation principle when the designer offers the same mechanism repeatedly: unless the agent's best response is the same across periods, the composition of the mechanism with the agent's reporting strategy yields a time-dependent, direct mechanism. To avoid keeping track of participation and reporting strategies separately in what follows, we assume a message $m_\phi \in M$ exists such that for all $(\omega, \varepsilon) \in \Omega \times \mathcal{E}$, $\phi(m_\phi, \omega, \varepsilon) = \delta_{a_\phi}$.

Timing Given ϕ , the agent faces the following extensive form. Nature draws (ω, ε) once at the beginning, unobserved to the agent. In each period, nature first draws the agent's type, which the agent observes. The agent then sends a message m into the mechanism. The mechanism then draws the allocation from $\phi(\cdot | m, \omega, \varepsilon)$, which the agent observes.

Given the mechanism ϕ and the extensive form game it induces, the agent's strategy specifies for each period t and each period- t type $\theta \in \Theta$, a distribution over M , as a function of the agent's past observations, which include her past types, messages, and allocations. Importantly, we assume the agent does not observe her payoffs to focus on the agent learning through the mechanism.

We assume the agent is infinitely patient, that is, she has limit-of-means preferences. Her average payoff through period T when the realization is (ω, ε) and the type-message-allocation sequence is $(\theta_t, m_t, a_t)_{t=1}^T$ is given by:

$$U_T((\theta_t, m_t, a_t)_{t=1}^T, \omega, \varepsilon) = \frac{1}{T} \sum_{t=1}^T u(a_t, \theta_t, \omega).$$

A strategy σ is a best response for the agent if for all alternative strategies σ' , we have that

$$\liminf_{T \rightarrow \infty} \mathbb{E}_\sigma [U_T] \geq \limsup_{T \rightarrow \infty} \mathbb{E}_{\sigma'} [U_T], \quad (13)$$

where \mathbb{E}_σ is the expectation relative to the measure induced over the terminal histories by the prior on Ω , the distribution on \mathcal{E} , the agent's type distribution f , the mechanism ϕ , and the agent's reporting strategy σ .³¹

³¹The Ionescu-Tulcea extension theorem implies this measure is always well-defined for any mechanism and any agent's

Implementation Our notion of implementation is based on the induced *occupation measure* on $A \times \Theta \times \Omega$, that is, the (limit) expected frequency of tuples (a, θ, ω) when the agent best responds to the mechanism. For this reason, we restrict attention to mechanisms ϕ for which (i) a best-response σ exists, and (ii) its induced occupation measure v_σ over $A \times \Theta \times M \times \Omega \times \mathcal{E}$ exists, defined as follows³²

$$v_\sigma(a, \theta, m, \omega, \varepsilon) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_\sigma \left[\sum_{t=1}^T \mathbb{1}[(a_t, \theta_t, m_t, \omega', \varepsilon') = (a, \theta, m, \omega, \varepsilon)] \right] = \lim_{T \rightarrow \infty} v_\sigma^T(a, \theta, m, \omega, \varepsilon), \quad (14)$$

where the last identity defines v_σ as the limit of the up to period T occupation measures v_σ^T , which are always well-defined.

Under our definition of best response, which is the same as in Hart (1985), existence of a best response implies the agent's payoff at the best-response strategy is well-defined.³³ Even if the occupation measure in Equation 14 is enough to calculate the agent's payoffs, that the agent's payoffs are well-defined does not mean the occupation measure is well-defined. Because outcome distributions—and not payoffs—are usually the focus of mechanism design, we require that both the mechanism has a best response *and* it induces a well-defined occupation measure.

Definition 5 (Implementation). *Outcome distribution $\vartheta \in \Delta(A \times \Theta \times \Omega)$ can be implemented by a repeated mechanism if a mechanism ϕ and a best-response strategy σ exist such that*

$$\vartheta(a, \theta, \omega) = \sum_{\varepsilon \in \mathcal{E}, m \in M} v_\sigma(a, \theta, m, \omega, \varepsilon).$$

We are now ready to state the main result of this section. Theorem 3 shows that the outcome distributions implemented by repeated mechanisms can be implemented by two-stage mechanisms, and hence by calibrated mechanisms:

Theorem 3 (Microfoundation of Calibrated Mechanism Design). *Outcome distribution $\vartheta \in \Delta(A \times \Theta \times \Omega)$ is implementable by a repeated mechanism if and only if ϑ can be implemented by an incentive compatible and individually rational two-stage mechanism, that is for all $(a, \theta, \omega) \in A \times \Theta \times \Omega$,*

$$\vartheta(a, \theta, \omega) = \mu_0(\omega) f(\theta) \int_{\Delta(\Omega)} \alpha(a|\theta, \mu) \beta(d\mu|\omega), \quad (15)$$

where $\beta : \Omega \rightarrow \Delta(\Delta(\Omega))$ is Bayes plausible and $\alpha(\cdot|\cdot, \mu) : \Theta \rightarrow \Delta(A)$ is incentive compatible and individually rational on the support of $\mu_0 \otimes \beta$.

The proof of this and all results in this section can be found in Appendix C.

Theorem 3 provides a microfoundation for calibrated mechanism design. Whenever the designer is concerned with agents learning from the outcome of the mechanism and cares only about the long-run outcome distribution, it is as if he is designing a two-stage mechanism.

We now provide a proof sketch for Theorem 3, which is also useful to understand the proof of the result

strategy. See Appendix C.1 for details.

³²Throughout this section, limits of measures should always be understood in the weak* sense.

³³Note that in mechanism design one always focuses on mechanisms that have well-defined best responses in single-agent settings, and equilibria in multi-agent ones.

in the next section. For simplicity, let $\tilde{\Omega} = \Omega \times \mathcal{E}$ with elements $\tilde{\omega}$. Suppose repeated mechanism ϕ implements ϑ , and let $\nu_\sigma \in \Delta(A \times \Theta \times M \times \tilde{\Omega})$ denote the induced occupation measure. As the analysis so far illustrates, tracking the joint distribution over allocations, types, states, *and* beliefs is important to show that ϑ can be implemented via a two-stage mechanism. To this end, we extend the up to period T occupation measures, $\nu_\sigma^T \in \Delta(A \times \Theta \times M \times \tilde{\Omega})$, to account for the frequency of beliefs through period T . In fact, we define two sequences of extended occupation measures over $A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega})$: the first, $\bar{\nu}_\sigma^{T,1}$, calculates the frequency of a tuple $(a, \theta, m, \tilde{\omega}, \mu)$ by counting the beliefs at the beginning of period t and the second, $\bar{\nu}_\sigma^{T,2}$, by counting the beliefs at the end of period t . Whereas the martingale property of beliefs implies these two sequences have the same (subsequential) limits, they have different conditional independence properties, which we use to derive the representation of ϑ via a two-stage mechanism. Suppose for simplicity that $\bar{\nu}_\sigma^{T,1}$ (and hence, $\bar{\nu}_\sigma^{T,2}$), have limit $\bar{\nu}_\sigma$, though this assumption is not needed for the proof.³⁴ A consequence of the martingale property of beliefs is that only the long-run beliefs of the agent are in the support of $\bar{\nu}_\sigma$.

The proof consists of three steps. In the first step, we show that $\bar{\nu}_\sigma$ admits the following decomposition:

$$\bar{\nu}_\sigma(\{(a, \theta, m, \tilde{\omega})\} \times \tilde{\Delta}) = \int_{\tilde{\Delta}} \mu(\tilde{\omega}) f(\theta) \rho(m|\theta, \mu) \alpha'(a|m, \mu) \tau(d\mu),$$

where (i) $\tau \in \Delta(\Delta(\tilde{\Omega}))$ has mean $\mu_0 \otimes \eta$, where recall η is the measure on \mathcal{E} , and (ii) $\rho : \Theta \times \Delta(\tilde{\Omega}) \rightarrow \Delta(M)$ is a “Markovian reporting strategy”, and (iii) α' is *almost* the allocation rule in the two-stage mechanism, and hence the prime notation. Moreover, on the support of μ , $\alpha'(\cdot|\cdot, \mu)$ coincides with $\phi(\cdot|\cdot, \tilde{\omega})$, implying that $\phi(\cdot|\cdot, \tilde{\omega})$ is constant in $\tilde{\omega}$ on the support of μ . This is the step which exploits the different conditional independence properties of $\bar{\nu}_\sigma^{T,1}$ and $\bar{\nu}_\sigma^{T,2}$. We use $\bar{\nu}_\sigma^{T,1}$ to show the conditional independence of types and beliefs—all agent types in period t have the same belief at the beginning of period t —and $\bar{\nu}_\sigma^{T,2}$ to show the conditional independence of the allocation and the state—the belief at the end of period t contains all the information about the state contained in the allocation.

In the second step, we show that $\rho : \Theta \times \Delta(\tilde{\Omega}) \rightarrow \Delta(M)$ is indeed a best response for the agent when her type is θ and her belief is μ . In other words, the support of $\rho(\cdot|\theta, \mu)$ is contained in

$$\arg \max_{m \in M} \sum_{\tilde{\omega} \in \tilde{\Omega}} \mu(\tilde{\omega}) \sum_{a \in A} \phi(a|m, \tilde{\omega}) u(a, \theta, \omega),$$

for beliefs on the support of τ . Hence, we can use ρ and ϕ to define a direct mechanism $\alpha(\cdot|\cdot, \mu) : \Theta \rightarrow \Delta(A)$ that satisfies the agent’s participation and incentive constraints when her belief is μ . Together, Steps 1 and 2 allow us to obtain the representation of ϑ as in Equation 15.³⁵

Whereas the above steps are enough to show that ϑ is implementable by *some* individually rational and incentive compatible two-stage mechanism, they do not necessarily imply that the distribution over posteriors τ is the one induced by the information structure calibrated to ϕ, π_ϕ . The last step of the proof shows that even if this is not the case, the agent *adequately learns* the information contained in π_ϕ in the sense of Aghion et al. (1991). Indeed, Lemma C.4 shows a strategy exists that approximately

³⁴By assumption, the marginal of $\bar{\nu}_\sigma^T$ on $A \times \Theta \times M \times \tilde{\Omega}$ converges to ν_σ . Moreover, we show that the marginal of $\bar{\nu}_\sigma^T$ on $\Delta(\tilde{\Omega})$ also converges (Lemma C.3). However, this is not enough to ensure the convergence of $\bar{\nu}_\sigma^T$.

³⁵Whereas the above two-stage mechanism is described in terms of beliefs over $\Omega \times \mathcal{E}$, we show in the appendix how to derive from it a two-stage mechanism in terms of beliefs over Ω .

delivers the payoff from learning π_ϕ , so that the agent's payoff under σ is at least the payoff she would obtain if she had learned π_ϕ . Because the payoff from learning π_ϕ is the maximal payoff the agent can possibly attain, we conclude that the payoff under σ is the payoff the agent would attain when facing the calibrated information structure π_ϕ (and best responding to it).

5.2 Dynamic Mechanisms

In this section, we consider the case in which the designer can offer the agent a *dynamic* mechanism, that is, one that conditions the allocation in each period on the history of past allocations and reports. The analysis herein allows us to describe the limits implied by calibration on the set of implementable outcomes.

Dynamic mechanisms A dynamic mechanism $\varphi = (\varphi_t)_{t \in \mathbb{N}}$ is a sequence of mappings that condition on the state, the history of participation decisions, type reports and allocations, and today's report, and output an allocation. Formally, expand the set of type reports and allocations by a non-participation message and the outside option, which we denote by $\Theta A_\emptyset = \Theta \times A \cup \{(\emptyset, a_\emptyset)\}$.³⁶ For each $t \in \mathbb{N}$, define the mechanism in period t , $\varphi_t : \Omega \times (\Theta A_\emptyset)^{t-1} \times \Theta \rightarrow \Delta(A)$. Because the designer can flexibly design the mechanism in each period we no longer rely on the randomization device.

A dynamic mechanism induces an extensive-form game for the agent, in which in each period, the agent decides whether to participate, and conditional on participation what type to report. Whenever the agent chooses not to participate, she obtains her outside option a_\emptyset . We denote by p the agent's participation strategy and by σ the agent's reporting strategy.

Implementation Our notion of implementation continues to be based on the occupation measure over the set of allocations, types, participation decisions, type reports, and states, induced by the distributions μ_0 and f , the mechanism φ , and the agent's participation and reporting strategy. However, as we show in [Appendix D.3.1](#), it is without loss to focus on mechanisms such that (i) participation with probability 1 and truthtelling is a best response for the agent, and (ii) the mechanism implements the outside option with probability 1 in all future periods following a non-participation decision by the agent.³⁷ Thus, we focus on dynamic mechanisms φ such that (i) a best response exists, and (ii) the occupation measure over $A \times \Theta \times \Omega$ is well-defined.

Incentives in dynamic mechanisms Dynamic mechanisms allow the designer to condition the agent's allocation on the history of past participation decisions and reports (and allocations), and hence allow the designer to implement outcomes that satisfy weaker notions of truthtelling and participation, which we explain next.

Because the designer can condition the mechanism on the history of past reports, he can compare the frequency of type reports against the type distribution. So long as the agent is telling the truth, the

³⁶This notation allows us to keep the definitions of the histories when the agent participates and does not participate symmetric, and saves us on including the agent's participation strategy in the histories.

³⁷That is, starting from a dynamic mechanism φ and a best response strategy (p, σ) , one can construct an alternative mechanism φ' such that participation and truthtelling are a best response for the agent and preserves the distribution over $(\Omega \times \Theta \times A)^\infty$ induced by (p, σ) and φ .

frequency of reports will match the type distribution f over large blocks of time. In fact, any reporting strategy whose expected frequency of reports matches the type distribution will be indistinguishable from truthtelling.

Definition 6 (Undetectable deviations). *An undetectable deviation is a reporting strategy $\sigma : \Theta \rightarrow \Delta(\Theta)$ such that for all $\theta' \in \Theta$*

$$\sum_{\theta \in \Theta} f(\theta) \sigma(\theta' | \theta) = f(\theta').$$

By tracking the empirical distribution of type reports, the designer can dissuade the agent from employing detectable deviations. Thus, in a dynamic mechanism, the designer should be concerned with only discouraging undetectable deviations. This leads to a weaker notion of incentive compatibility for allocation rules:

Definition 7 (Unprofitable undetectable deviations). *The allocation rule $\alpha : \Theta \times \Delta(\Omega) \rightarrow \Delta(A)$ lacks profitable undetectable deviations at belief $\mu \in \Delta(\Omega)$ if for all undetectable deviations σ ,*

$$\sum_{\theta \in \Theta} f(\theta) \sum_{a \in A} \alpha(a | \theta, \mu) \sum_{\omega \in \Omega} \mu(\omega) u(a, \theta, \omega) \geq \sum_{\theta \in \Theta} f(\theta) \sum_{\theta' \in \Theta} \sigma(\theta' | \theta) \sum_{a \in A} \alpha(a | \theta', \mu) \sum_{\omega \in \Omega} \mu(\omega) u(a, \theta, \omega).$$

A two-stage mechanism ψ with allocation rule α lacks profitable undetectable deviations if $\alpha(\cdot | \cdot, \mu)$ lacks profitable undetectable deviations for all beliefs in the support of the mechanism.

To illustrate the difference between the lack of profitable undetectable deviations and incentive compatibility, consider the following example from Ball and Kattwinkel (2023). Suppose the agent types are binary, $\{\theta_1, \theta_2\}$, and equally likely. The set of allocations, $q \in \{0, 1\}$, describes whether the agent receives a good. Finally, suppose the agent's payoff is $u(q, \theta) = q\theta$ and $\theta_1 < \theta_2$. Consider the mechanism that allocates the good to θ_2 : While it is not incentive compatible, it lacks profitable undetectable deviations. The constraint that the deviation must be undetectable implies the gains from θ_1 obtaining the good come at the expense of θ_2 getting the good.

Consider now the agent's *participation* incentives in the dynamic mechanism: once the agent rejects the mechanism once, the agent obtains her outside option in all continuation histories independent of her participation decision and her types. In other words, whereas the agent can always ensure her outside option by rejecting the mechanism in a given period, she is effectively quitting the mechanism forever for all her types. The following definition introduces the notion of individual rationality satisfied by the mechanism in the long run.

Definition 8 (Ex ante individual rationality). *The allocation rule $\alpha : \Theta \times \Delta(\Omega) \rightarrow \Delta(A)$ is ex ante individually rational at belief $\mu \in \Delta(\Omega)$ if*

$$\sum_{\theta \in \Theta} f(\theta) \sum_{a \in A} \alpha(a | \theta, \mu) \sum_{\omega \in \Omega} \mu(\omega) u(a, \theta, \omega) \geq \sum_{\theta \in \Theta} f(\theta) \sum_{\omega \in \Omega} \mu(\omega) u(a_\phi, \theta, \omega).$$

A two-stage mechanism ψ with allocation rule α is ex ante individually rational if $\alpha(\cdot | \cdot, \mu)$ is ex ante individually rational for all beliefs in the support of the mechanism.

We are now ready to state the main result of this section:

Theorem 4 (Implementable Outcomes via Dynamic Mechanisms). *A dynamic mechanism exists that implements outcome $\vartheta \in \Delta(A \times \Theta \times \Omega)$ if and only if ϑ can be implemented by an ex ante individually rational two-stage mechanism which lacks profitable undetectable deviations. That is, if and only if for all $(a, \theta, \omega) \in A \times \Theta \times \Omega$*

$$\vartheta(a, \theta, \omega) = \mu_0(\omega) f(\theta) \int_{\Delta(\Omega)} \alpha(a|\theta, \mu) \beta(d\mu|\omega), \quad (16)$$

where $\beta : \Omega \rightarrow \Delta(\Delta(\Omega))$ is Bayes plausible and $\alpha(\cdot|\cdot, \mu) : \Theta \rightarrow \Delta(A)$ lacks profitable undetectable deviations and is ex ante individually rational on the support of $\mu_0 \otimes \beta$.

Theorem 4 characterizes the outcome distributions implementable by dynamic mechanisms as those implemented by two-stage mechanisms that satisfy the incentive constraints: unprofitability of undetectable deviations and ex ante individual rationality. Notably, both notions of incentive constraints apply in the aggregate over the type distribution, which reflects the transient nature of the agent's private information.

Comparing Theorem 3 and Theorem 4, we see that dynamic mechanisms allow the designer to weaken the incentive constraints of the agent, but do not allow him to engage in richer—i.e., type-dependent—disclosures. Despite dynamic mechanisms implying weaker incentive constraints, we can build on the results of Rochet (1987) and Rahman (2024) to show that in settings with transferable utility, where $a = (q, t)$, dynamic and repeated mechanisms implement the same set of *physical* allocations $q : \Theta \times \Omega \rightarrow \mathbb{R}$.³⁸ Indeed, Rahman (2024) shows the lack of profitable undetectable deviations is equivalent to cyclical monotonicity in Rochet (1987). Thus, in settings with transferable utility, Theorems 3 and 4 imply that dynamic mechanisms do not allow the designer to expand on the set of implementable distributions over (q, θ, ω) .

The proof of the only if direction is similar to that of Theorem 3, in that we similarly extend the occupation measure to account for the agent's beliefs and show it satisfies the conditional independence properties implied by a two-stage mechanism. In a dynamic mechanism, however, the agent can ensure the payoff of some, but not all deviations. The latter property is what delivers that the two-stage mechanism must lack profitable undetectable deviations.

The proof of the if direction, instead, harnesses a construction in Margaria and Smolin (2018). The proof proceeds in two steps. In the first step, we analyze a fictitious model without state uncertainty in which a designer faces a privately informed agent, so that implementable outcomes are elements of $\Delta(A \times \Theta)$. We show that if $\vartheta'(a, \theta) = f(\theta) \alpha'(a|\theta) \in \Delta(A \times \Theta)$ is such that α' lacks profitable undetectable deviations and is ex ante individually rational, then a dynamic mechanism exists that implements ϑ' .³⁹ This is the step that relies on Margaria and Smolin (2018). We construct a dynamic mechanism, which can be split into blocks of random length. Each block consists of two phases: a reporting

³⁸This is easily seen in the example after Definition 7. Note the mechanism that allocates the good to the agent if and only if her type is θ_2 satisfies monotonicity. Hence, a transfer scheme exists that implements this allocation rule with transfers.

³⁹In a repeated principal-agent game with communication, Meng (2021) shows that the principal can guarantee in the patient limit his complete information payoff subject to the constraint that her actions satisfy the cyclic monotonicity condition in Rochet (1987). We view the results as complementary: We focus on implementable outcome distributions, instead of payoffs, when the agent has limit of the means preferences, which makes our notion of implementation exact.

phase and an adjustment phase. In the reporting phase, the mechanism uses the agent's reports to determine the allocation. Instead, in the adjustment phase, the mechanism simulates type reports so that the frequency of type reports matches the type distribution (in expectation) over the length of the block, whenever this is not the case at the end of the reporting phase. These two steps ensure that the expected frequency of type reports and allocations matches ϑ' . We then leverage that $\alpha'(\cdot|\theta)$ lacks profitable undetectable deviations to show the agent cannot do better than by telling the truth. Hence, the induced frequency of types and allocations also matches ϑ' . Moreover, the construction ensures that after any history, truthtelling delivers a continuation payoff equal to the ex ante payoff. Because α' is ex ante individually rational, we conclude the participation constraints are satisfied.

The second step uses the above result and the representation of the outcome distribution via a two-stage mechanism to construct a dynamic mechanism that implements any outcome distribution that satisfies the properties in [Theorem 4](#). Indeed, one can construct a dynamic mechanism which uses a finite number of steps to disclose information to the agent via the realized allocations,⁴⁰ and then continues as in the above construction to implement the allocation rule $\alpha(\cdot|\cdot, \mu)$.

6 Conclusions

Many economic institutions—online platforms, lenders, regulators—rely on mechanisms that remain fixed while agents interact with them repeatedly. When the mechanism's operation depends on a state known only to the designer, agents can learn this state from their outcomes, constraining what the mechanism can implement. We introduce calibrated mechanism design, a static solution concept that requires mechanisms to remain incentive compatible given the information they endogenously reveal about the designer's private state through repeated use. In private value environments, the calibration constraint pushes the designer toward full transparency, precluding Crémer-McLean-style schemes under transferable utility. In single agent-settings, calibrated mechanisms are equivalent to two-stage mechanisms. This equivalence yields a practical algorithm for finding optimal calibrated mechanisms, combining tools from information design and mechanism design. We provide a microfoundation by showing calibrated mechanisms characterize exactly what is implementable when an infinitely patient agent repeatedly interacts with the same mechanism, and study the implications on implementable outcomes of allowing the designer to offer fully dynamic mechanisms.

The most important direction for future work is deepening the analysis of multi-agent settings. On the one hand, understanding when generalized two-stage mechanisms coincide with calibrated mechanisms would enable the study of multi-agent applications, while abstracting from the dynamics of experimentation. On the other hand, extending our microfoundation to the multi-agent case would further ground calibrated mechanism design. More broadly, our framework suggests that any institution whose repeated operation leaks information about its designer's knowledge faces a fundamental tradeoff between conditioning the mechanism on this information and the information this leaks to participants, and calibrated mechanism design offers a disciplined way to analyze it.

⁴⁰A standard argument implies that if ϑ satisfies [Equation 16](#), then a finite support β' exists such that ϑ satisfies [Equation 16](#) with β' in place of β .

References

AGHION, P., P. BOLTON, C. HARRIS, AND B. JULLIEN (1991): “Optimal Learning by Experimentation,” *Review of Economic Studies*, 58, 621–654.

ALIPRANTIS, C. D. AND K. C. BORDER (2006): *Infinite Dimensional Analysis: a Hitchhiker’s Guide*, Springer.

ALONSO, R. AND O. CÂMARA (2016): “Bayesian Persuasion with Heterogeneous Priors,” *Journal of Economic Theory*, 165, 672–706.

ARIELI, I., Y. BABICHENKO, AND F. SANDOMIRSKIY (2024): “Feasible Conditional Belief Distributions.”

ATTAR, A., E. CAMPIONI, T. MARIOTTI, AND A. PAVAN (2025): “Keeping the Agents in the Dark: Competing Mechanisms, Private Disclosures, and the Revelation Principle,” *Working Paper*.

BALL, I. AND D. KATTWINKEL (2023): “Quota Mechanisms: Finite-Sample Optimality and Robustness,” *Working Paper*.

BERGEMANN, D., P. DUETTING, R. PAES LEME, AND S. ZUO (2022a): “Calibrated Click-Through Auctions,” in *Proceedings of the ACM Web Conference 2022*, 47–57.

BERGEMANN, D., T. HEUMANN, AND S. MORRIS (2022b): “Screening with Persuasion,” *Working Paper*.

BERGEMANN, D. AND J. HÖRNER (2018): “Should First-Price Auctions Be Transparent?” *American Economic Journal: Microeconomics*, 10, 177–218.

BERGEMANN, D. AND M. PESENDORFER (2007): “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137, 580–609.

BLUME, L. E., M. M. BRAY, AND D. EASLEY (1982): “Introduction to the Stability of Rational Expectations Equilibrium,” *Journal of Economic Theory*, 26, 313–317.

BOGACHEV, V. I. (2007): *Measure Theory*, Springer.

CALZOLARI, G. AND A. PAVAN (2006): “On the Optimality of Privacy in Sequential Contracting,” *Journal of Economic Theory*, 130, 168–204.

CESA-BIANCHI, N., T. CESARI, R. COLOMBONI, F. FUSCO, AND S. LEONARDI (2024): “The Role of Transparency in Repeated First-Price Auctions with Unknown Valuations,” in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, 225–236.

CRÉMER, J. AND R. P. MCLEAN (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 1247–1257.

DASKALAKIS, C., C. PAPADIMITRIOU, AND C. TZAMOS (2016): “Does Information Revelation Improve Revenue?” in *Proceedings of the 2016 ACM Conference on Economics and Computation*, 233–250.

DOVAL, L. AND V. SKRETA (2022): “Mechanism Design with Limited Commitment,” *Econometrica*, 90, 1463–1500.

DWORCZAK, P. (2020): “Mechanism Design with Aftermarkets: Cutoff Mechanisms,” *Econometrica*, 88, 2629–2661.

ESŐ, P. AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74, 705–731.

ESPONDA, I. (2008): “Information Feedback in First Price Auctions,” *RAND Journal of Economics*, 39, 491–508.

FOSTER, D. P. AND R. V. VOHRA (1997): “Calibrated Learning and Correlated Equilibrium,” *Games and Economic Behavior*, 21, 40–55.

FU, H., P. JORDAN, M. MAHDIAN, U. NADAV, I. TALGAM-COHEN, AND S. VASSILVITSKII (2012): “Ad Auctions with Data,” in *International Symposium on Algorithmic Game Theory*, Springer, 168–179.

GENTZKOW, M. AND E. KAMENICA (2017): “Bayesian Persuasion with Multiple Senders and Rich Signal Spaces,” *Games and Economic Behavior*, 104, 411–429.

GOLREZAEI, N., A. JAVANMARD, AND V. MIRROKNI (2019): “Dynamic Incentive-Aware Learning: Robust Pricing in Contextual Auctions,” *Advances in Neural Information Processing Systems*, 32.

GREEN, J. (1977): “The Non-Existence of Informational Equilibria,” *Review of Economic Studies*, 44, 451–463.

GREEN, J. R. AND J.-J. LAFFONT (1987): “Posterior Implementability in a Two-Person Decision Problem,” *Econometrica*, 69–94.

GREEN, J. R. AND N. L. STOKEY (2022): “Two Representations of Information Structures and Their Comparisons,” *Decisions in Economics and Finance*, 45, 541–547.

GUESNERIE, R. AND J.-J. LAFFONT (1984): “A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm,” *Journal of Public Economics*, 25, 329–369.

HABERMAN, A. AND R. JAGADEESAN (2025): “Auctions with Withdrawal Rights: A Foundation for Uniform Price,” in *Proceedings of the 26th ACM Conference on Economics and Computation*, 35–35.

HÄFNER, S., M. PYCIA, AND H. ZENG (2025): “Mechanism Design with Information Leakage,” *Working Paper*.

HART, S. (1985): “Nonzero-Sum Two-Person Repeated Games with Incomplete Information,” *Mathematics of Operations Research*, 10, 117–153.

JACKSON, M. O. AND H. F. SONNENSCHEIN (2007): “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, 75, 241–257.

KALLENBERG, O. (2017): *Random Measures, Theory and Applications*, vol. 1, Springer.

KANORIA, Y. AND H. NAZERZADEH (2020): “Dynamic Reserve Prices for Repeated Auctions: Learning from Bids,” *Working Paper*.

KARTIK, N., S. LEE, AND D. RAPPOPORT (2024): “Single-Crossing Differences in Convex Environments,” *Review of Economic Studies*, 91, 2981–3012.

KRÄHMER, D. (2020): “Information Disclosure and Full Surplus Extraction in Mechanism Design,” *Journal of Economic Theory*, 187, 105020.

KREPS, D. M. (1977): “A Note on “Fulfilled Expectations” Equilibria,” *Journal of Economic Theory*, 14, 32–43.

LACLAU, M. AND L. RENOU (2017): “Public Persuasion,” *Working Paper*.

LI, H. AND X. SHI (2017): “Discriminatory Information Disclosure,” *American Economic Review*, 107, 3363–3385.

MARGARIA, C. AND A. SMOLIN (2018): “Dynamic Communication with Biased Senders,” *Games and Economic Behavior*, 110, 330–339.

MASKIN, E. AND J. TIROLE (1990): “The Principal-Agent Relationship with an Informed Principal: The Case of Private Values,” *Econometrica*, 379–409.

MENG, D. (2021): “On the Value of Repetition for Communication Games,” *Games and Economic Behavior*, 127, 227–246.

MILGROM, P. R. (1981): “Rational Expectations, Information Acquisition, and Competitive Bidding,” *Econometrica*, 921–943.

MILGROM, P. R. AND R. J. WEBER (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 1089–1122.

MYERSON, R. B. (1983): “Mechanism Design by an Informed Principal,” *Econometrica*, 1767–1797.

——— (1986): “Multistage Games with Communication,” *Econometrica*, 323–358.

NEDELEC, T., C. CALAUZÈNES, N. EL KAROUI, AND V. PERCHET (2022): “Learning in Repeated Auctions,” *Found. Trends Mach. Learn.*, 15, 176–334.

NEDELEC, T., N. E. KAROUI, AND V. PERCHET (2019): “Learning to Bid in Revenue-Maximizing Auctions,” in *Proceedings of the 36th International Conference on Machine Learning*, vol. 97, 4781–4789.

NIEMEYER, A. (2022): “Posterior Implementability in an N-Person Decision Problem,” *Working Paper*.

OTTAVIANI, M. AND A. PRAT (2001): “The Value of Public Information in Monopoly,” *Econometrica*, 69, 1673–1683.

PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601–653.

RADNER, R. (1979): “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices,” *Econometrica*, 655–678.

RAHMAN, D. M. (2024): “Detecting Profitable Deviations,” *Journal of Mathematical Economics*, 111, 102946.

RENOU, L. AND T. TOMALA (2015): “Approximate Implementation in Markovian Environments,” *Journal of Economic Theory*, 159, 401–442.

ROCHET, J.-C. (1987): “A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context,” *Journal of Mathematical Economics*, 16, 191–200.

RUBIN, H. AND O. WESLER (1958): “A Note on Convexity in Euclidean N-Space,” in *Proceedings of the American Mathematical Society*, vol. 9, 522–523.

SMOLIN, A. (2023): “Disclosure and Pricing of Attributes,” *RAND Journal of Economics*, 54, 570–597.

SZABADI, B. (2018): “Essays in Microeconomic Theory: Information Design, Partnerships, and Matching,” Ph.D. thesis, Northwestern University.

YAMASHITA, T. (2018): “Optimal Public Information Disclosure by Mechanism Designer,” *Working Paper*.

Mathematical conventions

Throughout the appendix, we take all sets to be Polish spaces, that is, completely metrizable, separable, topological spaces, and endow them with their Borel σ -algebra. We endow product spaces with their product σ -algebra. For a Polish space X , we let \mathcal{B}_X denote its Borel σ -algebra and $\Delta(X)$ the set of all Borel probability measures on X , endowed with the weak* topology. Thus, $\Delta(X)$ is also a Polish space (Aliprantis and Border, 2006), and it is compact, whenever X is compact (Aliprantis and Border, 2006, Theorem 15.11 and Theorem 15.15).

Notational conventions If X is a Polish space, \tilde{X} denotes a measurable subset of X , i.e., an element of the Borel σ -algebra on X , and $C_b(X)$ denotes the set of continuous and bounded functions on X . Given a measure $\nu \in \Delta(\times_{i=1}^N Y_i)$, we denote by $\nu_{Y_j Y_k \dots Y_l}$ the marginal of ν on $Y_j Y_k \dots Y_l$. When one of the $Y_i = \Delta(X_i)$, we write Δ instead of Y_i in the subscript, when it is unlikely to generate confusion.

Throughout the appendix, we define different distributions that arise in our proofs. Because we endow product spaces with their product topology and their product Borel σ -algebra, it is enough to define these new measures on the measurable rectangles and we follow this convention throughout.

Disintegration We rely on the notion of disintegration in many of our proofs (Bogachev, 2007, Chapter 10.6). We define disintegration in the context of product sets $X \times Y$, as this is the one that shows up in the proof, but it is more general than this. Given a measure $\nu \in \Delta(X \times Y)$, $\lambda : X \times \mathcal{B}_Y \rightarrow [0, 1]$ is the disintegration of ν along X if the following holds

1. For all $\tilde{Y} \in \mathcal{B}_Y$, $x \mapsto \lambda_x(\tilde{Y})$ is measurable,
2. For ν_X -almost everywhere $x \in X$, $\tilde{Y} \mapsto \lambda_x(\tilde{Y})$ is a probability measure, and
3. For every bounded measurable function $g : X \times Y \rightarrow \mathbb{R}$,

$$\int_{X \times Y} g(x, y) \nu(d(x, y)) = \int_X \int_Y g(x, y) \lambda_x(dy) \nu_X(dx).$$

Kallenberg (2017, Theorem 1.23) ensures that $\{\lambda_x : x \in X\}$ exists and is unique ν_X -almost everywhere.

A Omitted proofs from Section 2

Proof of Theorem 1. Suppose the agents' payoffs are state independent and in a slight abuse of notation let $u_i(a_i, \theta_i)$ denote agent i 's utility.

The calibrated mechanism design problem is

$$\max_{\phi: \Theta \times \Omega \times [0,1] \rightarrow \Delta(A)} \sum_{\omega \in \Omega} \mu_0(\omega) \sum_{\theta \in \Theta} f(\theta|\omega) \int_0^1 w(\phi(\theta, \omega, \varepsilon), \theta, \omega) \lambda(d\varepsilon),$$

subject to the following constraints holding for all $(\omega, \varepsilon) \in \Omega \times [0, 1]$, $i \in [N]$, $\theta_i \in \Theta$ and $\theta'_i \in \Theta$:

$$\begin{aligned} \sum_{a_i \in A_i} \mathbb{E}_{f_{-i}(\cdot|\omega)} \left[\sum_{a_{-i} \in A_{-i}} \phi(\theta_i, \theta_{-i}, \omega, \varepsilon)(a_i, a_{-i}) u_i(a_i, \theta_i) \right] &\geq \sum_{(a_i \in A_i)} \mathbb{E}_{f_{-i}(\cdot|\omega)} \left[\sum_{a_{-i} \in A_{-i}} \phi(\theta'_i, \theta_{-i}, \omega, \varepsilon)(a_i, a_{-i}) u_i(a_i, \theta_i) \right] u_i(a_i, \theta_i) \\ \sum_{a_i \in A_i} \mathbb{E}_{f_{-i}(\cdot|\omega)} \left[\sum_{a_{-i} \in A_{-i}} \phi(\theta_i, \theta_{-i}, \omega, \varepsilon)(a_i, a_{-i}) u_i(a_i, \theta_i) \right] u_i(a_i, \theta_i) &\geq u_i(a_{i\emptyset}, \theta_i). \end{aligned}$$

In other words, for each agent i , her interim allocation rule $\pi_{\phi,i}(\omega, \varepsilon)$ must be an element of $S_{IC/IR,i}^*$, where the latter is the set of interim allocation rules $S_i^* : \Theta_i \rightarrow \Delta(A_i)$ that satisfy the following incentive compatibility and individual rationality constraints:

$$\begin{aligned} (\forall \theta_i, \theta'_i \in \Theta_i) \sum_{a_i \in A_i} s_i^*(a_i|\theta_i) u_i(a_i, \theta_i) &\geq \sum_{a_i \in A_i} s_i^*(a_i|\theta'_i) u_i(a_i, \theta_i) \\ (\forall \theta_i \in \Theta_i) \sum_{a_i \in A_i} s_i^*(a_i|\theta_i) u_i(a_i, \theta_i) &\geq u_i(a_{i\emptyset}, \theta_i). \end{aligned}$$

Because the individual rationality and incentive constraints must hold for each pair (ω, ε) , the designer's problem is separable across variables for different ω, ε : the sets of variables $\phi(\cdot, \omega, \varepsilon)$ appear in different sets of constraints and the objective function is additively separable across those variables. Consequently, the designer's problem can be solved as a collection of independent problems, one for each ω, ε . \square

B Omitted proofs from Section 3

In this section, we present the proofs of [Theorem 2](#) and [Proposition 1](#). We proceed as follows: We first prove [Proposition 1](#), as when $N = 1$ its proof implies the “if” direction of [Theorem 2](#). We then prove the “only if” direction of [Theorem 2](#).

Proof of [Proposition 1](#). We focus on the case in which types and states are independently distributed, and explain how to extend the proof when they are not.

Let $\vartheta \in \Delta(A \times \Theta \times \Omega)$ denote the outcome distribution implemented by an incentive compatible and individually rational calibrated mechanism. We show transition probabilities $\beta : \Omega \rightarrow \Delta(\Delta(\Omega)^N)$, $\bar{\alpha} : \Theta \times \Omega \times \Delta(\Omega)^N \rightarrow \Delta(A)$, and $\alpha_i : \Theta_i \times \Delta(\Omega) \rightarrow \Delta(A_i)$ for $i \in \{1, \dots, N\}$ exist such that

$$\vartheta(a, \theta, \omega) = \mu_0(\omega) f(\theta) \int_{\Delta(\Omega)^N} \bar{\alpha}(a|\theta, \omega, \mu_1, \dots, \mu_N) \beta(d(\mu_1, \dots, \mu_N)|\omega), \quad (\text{B.1})$$

and for all $i \in \{1, \dots, N\}$, (i) α_i satisfies [item 3 of Definition 4](#), and (ii) on the support of $\mu_0 \otimes \beta$, $\alpha_i(\cdot|\cdot, \mu_i)$ is incentive compatible and individually rational when agent i holds belief μ_i .

Let $\pi_{\omega,i} : [0, 1] \rightarrow S_i^*$ denote the mapping $\varepsilon \mapsto \pi_i(\omega, \cdot)$. For $\tilde{S}_i^* \in \Delta(A_i)^{\Theta_i}$, define

$$\Pr_i(\{\omega\} \times \tilde{S}_i^*) = \mu_0(\omega) \lambda(\pi_{\omega,i}^{-1}(\tilde{S}_i^*)) = \int_{\tilde{S}_i^*} \mu_i(\omega|s_i^*) \tau_{\phi,i}(ds_i^*),$$

where the second equality follows from disintegration of $\Pr_i \in \Delta(\Omega \times S_i^*)$ along S_i^* , and corresponds to the definition of Bayes rule for agent i . Define the measurable mappings, $T_i : S_i^* \rightarrow \Delta(\Omega)$ and

$T : S^* \rightarrow \Delta(\Omega)^N$ as follows: $T_i(s_i^*) = \mu_i(\cdot | s_i^*)$ and $T(s^*) = (T_1(s_1^*), \dots, T_N(s_N^*))$.

Define a joint distribution $Q \in \Delta(A \times \Theta \times \Omega \times \Delta(\Omega)^N)$ as follows:

$$Q(\{(a, \theta, \omega)\} \times \times_{i=1}^N \tilde{\Delta}_i) = \mu_0(\omega) f(\theta) \int_{\pi_{\omega}^{-1}(T^{-1}(\times \tilde{\Delta}_i))} \phi(a | \theta, \omega, \varepsilon) \lambda(d\varepsilon),$$

where $\pi_{\omega}^{-1}(T^{-1}(\times \tilde{\Delta}_i)) = \cap_{i=1}^N \{\varepsilon : T_i(\pi_{\omega, i}(\varepsilon)) \in \tilde{\Delta}_i\}$.

We note the following properties of Q . First, consider its marginal over $\Theta \times \Omega \times \Delta(\Omega)^N$,

$$Q_{\Theta \Omega \Delta^N}(\{(\theta, \omega)\} \times \times_{i=1}^N \tilde{\Delta}_i) = \mu_0(\omega) f(\theta) \lambda(\cap_{i=1}^N \{\varepsilon : T_i(\pi_{\omega, i}(\varepsilon)) \in \tilde{\Delta}_i\}),$$

which implies that the disintegration of $Q_{\Theta \Omega \Delta^N}$ along $\Theta \times \Omega$, $\beta : \Theta \times \Omega \rightarrow \Delta(\Delta(\Omega)^N)$ does not depend on θ . This automatically implies that Q admits the following disintegration:

$$Q(\{(a, \theta, \omega)\} \times \times_{i=1}^N \tilde{\Delta}_i) = \mu_0(\omega) f(\theta) \int_{\times_{i=1}^N \tilde{\Delta}_i} \bar{\alpha}(a | \theta, \omega, \mu_1, \dots, \mu_N) \beta(d(\mu_1, \dots, \mu_N) | \omega), \quad (\text{B.2})$$

which, in turn, delivers [Equation B.1](#). Moreover, note that the marginal of β on the beliefs of agent i , $\beta_i : \Omega \rightarrow \Delta(\Delta(\Omega))$, satisfies

$$\beta_i(\tilde{\Delta}_i | \omega) = \lambda(\pi_{\omega, i}^{-1}(T_i^{-1}(\tilde{\Delta}_i))).$$

Consider now the marginal on $A_i \times \Theta_i \times \Omega \times \Delta(\Omega)$ of Q , $Q_{A_i \Theta_i \Omega \Delta_i}$, which satisfies:

$$\begin{aligned} Q_{A_i \Theta_i \Omega \Delta_i}(\{(a_i, \theta_i, \omega)\} \times \tilde{\Delta}_i) &= \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a_{-i} \in A_{-i}} Q(\{(a, \theta, \omega)\} \times \tilde{\Delta}_i \times \Delta(\Omega)^{N-1}) = \\ &= \mu_0(\omega) f_i(\theta_i) \int_{\pi_{\omega, i}^{-1}(T_i^{-1}(\tilde{\Delta}_i))} \left(\sum_{\theta_{-i} \in \Theta_{-i}} f_{-i}(\theta_{-i}) \sum_{a_{-i} \in A_{-i}} \phi(a_i, a_{-i} | \theta_i, \theta_{-i}, \omega, \varepsilon) \right) \lambda(d\varepsilon) \\ &= \mu_0(\omega) f_i(\theta_i) \int_{\pi_{\omega, i}^{-1}(T_i^{-1}(\tilde{\Delta}_i))} \pi_{\omega, i}(\varepsilon) (a_i | \theta_i) \lambda(d\varepsilon) = \mu_0(\omega) f_i(\theta_i) \int_{T_i^{-1}(\tilde{\Delta}_i)} s_i^*(a_i | \theta_i) (\lambda \circ \pi_{\omega, i}^{-1})(ds_i^*). \end{aligned} \quad (\text{B.3})$$

Lastly, $Q_{A_i \Theta_i \Omega \Delta_i}$ admits the following representation via disintegration:

$$\begin{aligned} Q_{A_i \Theta_i \Omega \Delta_i}(\{(a_i, \theta_i, \omega)\} \times \tilde{\Delta}_i) &= \mu_0(\omega) f_i(\theta_i) \int_{\tilde{\Delta}_i} \alpha_i(a_i | \theta_i, \omega, \mu_i) \beta_i(d\mu_i | \omega) \\ &= \mu_0(\omega) f_i(\theta_i) \int_{\tilde{\Delta}_i} \alpha_i(a_i | \theta_i, \omega, \mu_i) (\lambda \circ \pi_{\omega, i}^{-1} \circ T_i^{-1})(d\mu_i) \end{aligned} \quad (\text{B.4})$$

Together with the uniqueness of disintegration and the sufficiency property of beliefs, [Equations B.3](#) and [B.4](#) imply that α_i does not depend on ω . The incentive compatibility and individual rationality of α_i follows from that of the calibrated mechanism.

Finally, consider the case in which θ and ω are not independent. Then, the experiment β in [Equation B.2](#) induces a joint distribution over the beliefs of N fictitious agents whose prior over the state is given by μ_0 . Agent i 's updated beliefs when her type is θ_i obtain from a transformation of μ (Alonso and Câmara, 2016; Laclau and Renou, 2017).⁴¹ Thus, up to changing $f(\theta)$ by $f(\theta | \omega)$, and interpreting

⁴¹Let $\mu_0(\cdot | \theta_i) \in \Delta(\Omega)$ denote the prior of the agent with type θ_i and $\mu(\cdot | s_i)$ denote the updated belief of an agent with prior

the draw from the Blackwell experiment as the posterior of an agent with prior belief μ_0 , the result follows. \square

Proof of Theorem 2.

“Only if” direction Similar to the proof of [Proposition 1](#), we focus on the case in which θ and ω are independent. Suppose $\vartheta \in \Delta(A \times \Theta \times \Omega)$ is implemented by an incentive compatible and individually rational two-stage mechanism. That is,

$$\vartheta(a, \theta, \omega) = \mu_0(\omega) f(\theta) \int_{\Delta(\Omega)} \alpha(a|\theta, \mu) \beta(d\mu|\omega), \quad (\text{B.5})$$

and α is incentive compatible and individually rational on the support of $\mu_0 \otimes \beta$. We construct an incentive compatible and individually rational calibrated mechanism that implements ϑ .

First, if ϑ satisfies [Equation B.5](#), [Rubin and Wesler \(1958\)](#) and [Carathéodory’s theorem \(Aliprantis and Border, 2006, Theorem 5.32\)](#) imply that a finite support $\beta' : \Omega \rightarrow \Delta(\{\mu_1, \dots, \mu_K\})$ exists such that for all $(a, \theta, \omega) \in A \times \Theta \times \Omega$ ⁴²

$$\vartheta(a, \theta, \omega) = \mu_0(\omega) f(\theta) \sum_{k=1}^K \alpha(a|\theta, \mu_k) \beta'(\{\mu_k\}|\omega).$$

For each $\omega \in \Omega$, partition $[0, 1] = \bigcup_{k=1}^{K-2} [b_k^\omega, b_{k+1}^\omega] \cup [b_{K-1}^\omega, 1]$, where $b_1 = 0$, and for all $k \in \{1, \dots, K-2\}$, $b_{k+1}^\omega = \sum_{l=1}^k \beta'(\{\mu_l\}|\omega)$. Define for $\varepsilon \in [b_k^\omega, b_{k+1}^\omega]$

$$\phi(a|\theta, \omega, \varepsilon) = \alpha(a|\theta, \mu_k).$$

The calibrated information structure is $\pi_\phi(\omega, \varepsilon) = \phi(\cdot|\cdot, \omega, \varepsilon) = \alpha(\cdot|\theta, \mu_k)$ for $\varepsilon \in [b_k^\omega, b_{k+1}^\omega]$ if $m \leq K-2$ or $\varepsilon \in [b_{K-1}^\omega, 1]$.

We now show that for all θ and all $s \in \text{supp } \pi_\phi$, the mechanism ϕ is incentive compatible and individually rational. Note that π_ϕ has finite support, and let $s \in \text{supp } \pi_\phi$ and let $\mu(\cdot|s)$ denote the updated posterior. Then, k exists such that the following holds:

$$\mu(\omega|s) = \frac{\mu_0(\omega) \lambda(\{\varepsilon : \pi(\omega, \varepsilon) = s\})}{\sum_{\omega' \in \Omega} \mu_0(\omega') \lambda(\{\varepsilon : \pi(\omega', \varepsilon) = s\})} = \frac{\mu_0(\omega) (b_{k+1}^\omega - b_k^\omega)}{\sum_{\omega' \in \Omega} \mu_0(\omega') (b_{k+1}^{\omega'} - b_k^{\omega'})} = \mu_k(\omega).$$

Moreover, because $\phi(\cdot|\cdot, \omega, \varepsilon) = \alpha(\cdot|\cdot, \mu_k)$, then it satisfies the agent’s incentive compatibility and individual rationality constraints when she holds belief μ_k .⁴³ \square

μ_0 upon observing signal s_i^* . When the signal is s_i^* , the agent with type θ_i updates her beliefs to:

$$\mu_i(\cdot|\theta_i, s_i^*) = \frac{\mu_0(\cdot|\theta_i) \cdot \frac{\mu(\cdot|s_i^*)}{\mu_0(\cdot)}}{\left\| \mu_0(\cdot|\theta_i) \cdot \frac{\mu(\cdot|s_i^*)}{\mu_0(\cdot)} \right\|},$$

where the \cdot and $/$ operations are meant componentwise, and $\|\cdot\|$ is the l^1 -norm.

⁴²Namely, [Equation B.5](#) implies that for all (θ, ω) , $\vartheta(\cdot|\theta, \omega) \in \text{cl co}\{\alpha(\cdot|\theta, \mu) : \mu \in \Delta(\Omega)\}$. [Rubin and Wesler \(1958\)](#) implies that our under assumptions $\text{cl co}\{\alpha(\cdot|\theta, \mu) : \mu \in \Delta(\Omega)\} = \text{co}\{\alpha(\cdot|\theta, \mu) : \mu \in \Delta(\Omega)\}$, and the rest of the claim follows from [Carathéodory’s theorem](#).

⁴³To extend the result to the case in which the agent’s type is correlated with the state, note the following. Knowing μ_0

Supplementary Appendix

C Omitted proofs from Section 5

C.1 Repeated Mechanisms

In this section, we present the proof of [Theorem 3](#). To do so, we first complete the formal definition of the game induced by repeating mechanism $\phi: M \times \Omega \times \mathcal{E} \rightarrow \Delta(A)$, by specifying the histories, strategy space, and the distribution over terminal histories induced by the agent's strategy and the mechanism. Having laid this groundwork, we describe the proof strategy, and then provide the formal details of the proof. Throughout this section, we use the shorthand $\tilde{\Omega} = \Omega \times \mathcal{E}$, and denote its elements by $\tilde{\omega}$.

Histories and strategies Histories through period $t \in \mathbb{N}$ are defined as $H^t \equiv (\Theta \times M \times A)^{t-1}$. The set of infinite histories from the agent's point of view is H^∞ . The set of terminal histories is $\mathcal{H}^\infty \equiv \tilde{\Omega} \times H^\infty$, where recall \mathcal{E} is finite and endowed with some measure η .

The agent's behavioral strategy is defined as a collection $\sigma \equiv (\sigma_t)_{t \in \mathbb{N}}$ such that for all $t \geq 1$

$$\sigma_t: H^t \times \Theta \rightarrow \Delta(M).$$

The tuple of distributions (μ_0, η, f) together with the mechanism ϕ and the agent's strategy σ determine a joint distribution over \mathcal{H}^∞ by the Ionescu-Tulcea theorem (Bogachev, 2007, Theorem 10.7.3). We provide more details on this probability distribution below. Denote by $\mathbb{P}_{(\mu_0, \eta, f, \phi, \sigma)}$ and $\mathbb{E}_{(\mu_0, \eta, f, \phi, \sigma)}$ the probability distribution over the terminal histories and the expectation with respect to this distribution, respectively. Whenever it is not likely to lead to confusion, we drop the dependence on $(\mu_0, \eta, f, \phi, \sigma)$, and whenever we want to emphasize the dependence on the agent's strategy we note the dependence on σ .

The distribution over terminal histories \mathcal{H}^∞ For future use, we review the construction of \mathbb{P}_σ . For each t , the distributions (μ_0, η, f) together with the mechanism ϕ and the agent's strategy σ determine a distribution over $\tilde{\Omega} \times H^t$, which we denote by $\mathbb{P}_\sigma^t \in \Delta(\tilde{\Omega} \times H^t)$. Note that for any subset $\tilde{\mathcal{H}}^t \subset \tilde{\Omega} \times H^t$,

$$\mathbb{P}_\sigma^t(\tilde{\mathcal{H}}^t) = \mathbb{P}_\sigma^{t+1}(\tilde{\mathcal{H}}^t \times (\Theta \times M \times A)). \quad (\text{C.1})$$

Moreover,

$$\mathbb{P}_\sigma^{t+1}(\tilde{\omega}, h^t, \theta, m, a) = \mathbb{P}_\sigma^t(\tilde{\omega}, h^t) f(\theta) \sigma_t(h^t, \theta)(m) \phi(a|m, \tilde{\omega}). \quad (\text{C.2})$$

By the Ionescu-Tulcea theorem, the distribution $\mathbb{P}_\sigma \in \Delta(\tilde{\Omega} \times H^\infty)$ is the unique distribution that satisfies that for all $t \in \mathbb{N}$, $\tilde{\mathcal{H}}^t \subset \tilde{\Omega} \times H^t$,

$$\mathbb{P}_\sigma(\tilde{\mathcal{H}}^t \times \prod_{s=t+1}^{\infty} (\Theta \times M \times A)) = \mathbb{P}_\sigma^t(\tilde{\mathcal{H}}^t). \quad (\text{C.3})$$

updates to μ_k conditional on s is enough to pin down the agent's belief $\mu(\cdot|\theta, s)$, with respect to which the agent's incentive compatibility and individual rationality constraints are defined (see [footnote 41](#)).

Belief system The agent's beliefs over $\tilde{\Omega}$ at the beginning of each t are determined by the belief system, which in a slight abuse of notation we denote by $\mu_t : H^t \rightarrow \Delta(\tilde{\Omega})$. The belief system satisfies

$$\mathbb{P}_\sigma^t(h^t) \mu_t(\tilde{\omega}|h^t) = \mathbb{P}_\sigma^t(\tilde{\omega}, h^t).$$

That is, whenever h^t is such that $\mathbb{P}_\sigma(\{\tilde{h} \in H^\infty : \tilde{h}^t = h^t\}) > 0$,

$$\mu_t(\tilde{\omega}|h^t) = \frac{\mathbb{P}_\sigma^t(\tilde{\omega}, h^t)}{\mathbb{P}_\sigma^t(h^t)} = \mathbb{P}_\sigma^t(\tilde{\omega}|h^t).$$

Given $\mathbb{P}_\sigma \in \Delta(\mathcal{H}^\infty)$, define $\mu_\infty(\tilde{\omega}|h^\infty) \equiv \mathbb{P}_\sigma(\tilde{\omega}|h^\infty)$ to be the belief system conditional on the whole terminal history h^∞ .

Remark C.1 (Belief system and strategies as functions on \mathcal{H}^∞). *Whereas the beliefs and strategies are defined on the finite histories, it is sometimes convenient to write them as functions on \mathcal{H}^∞ that are adapted to H^t .*

A property of the belief system We collect here a property of the belief system which we use in our proofs below.

Lemma C.1 (Martingale property under weak* convergence). $\mu_t(h^\infty) \xrightarrow{w^*} \mu_\infty(h^\infty)$ \mathbb{P}_σ -almost surely.

This and the proof of other technical results are in [Appendix D](#).

C.1.1 Proof of [Theorem 3 \(necessity\)](#)

We are now ready to present the proof of [Theorem 3](#), starting by the “only if” direction. Let $\vartheta \in \Delta(A \times \Theta \times \Omega)$ denote the outcome distribution implemented by repeated mechanism ϕ under best response σ , and let ν_σ denote the associated occupation measure, the definition of which we reproduce below for ease of reference:

$$\nu_\sigma(a, \theta, m, \tilde{\omega}) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_\sigma \left[\sum_{t=1}^T \mathbb{1}[(a_t, \theta_t, m_t, \tilde{\omega}') = (a, \theta, m, \tilde{\omega})] \right] = \lim_{T \rightarrow \infty} \nu_\sigma^T(a, \theta, m, \tilde{\omega}), \quad (\text{C.4})$$

where recall limits are in the weak* sense. We show that ϑ can be implemented by an incentive compatible and individually rational two-stage mechanism.

To this end, we consider two sequences of *extended* occupation measures on $A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega})$:

$$\bar{\nu}_\sigma^{T,1}(\{(a, \theta, m, \tilde{\omega})\} \times \tilde{\Delta}) = \frac{1}{T} \mathbb{E}_\sigma \left[\sum_{t=1}^T \mathbb{1}[(a_t, \theta_t, m_t, \tilde{\omega}') = (a, \theta, m, \tilde{\omega})] \mathbb{1}[\mu_t \in \tilde{\Delta}] \right], \quad (\text{C.5})$$

$$\bar{\nu}_\sigma^{T,2}(\{(a, \theta, m, \tilde{\omega})\} \times \tilde{\Delta}) = \frac{1}{T} \mathbb{E}_\sigma \left[\sum_{t=1}^T \mathbb{1}[(a_t, \theta_t, m_t, \tilde{\omega}') = (a, \theta, m, \tilde{\omega})] \mathbb{1}[\mu_{t+1} \in \tilde{\Delta}] \right]. \quad (\text{C.6})$$

We note the following. First, [Equation C.5](#) counts the beliefs at the beginning of period t , while [Equation C.6](#) counts the beliefs at the end of period t (after the realization of θ , m , and a .) [Equation C.5](#) is key to obtain the (limit) independence of the belief and type distributions, while [Equation C.6](#) allows us to obtain the (limit) independence of the allocation and the state, conditional on the induced belief.

Second, v_σ^T is the marginal of both $\bar{v}_\sigma^{T,1}$ and $\bar{v}_\sigma^{T,2}$. Third, by [Lemma C.1](#), $\mu_t \xrightarrow{w^*} \mu_\infty$, and hence both $\bar{v}_\sigma^{T,1}$ and $\bar{v}_\sigma^{T,2}$ have the same set of subsequential limits, which we record for future reference below (see [Appendix D](#) for the proof):

Lemma C.2. *The occupation measures $\bar{v}_\sigma^{T,1}$ and $\bar{v}_\sigma^{T,2}$ have the same set of subsequential limits.*

The proof of necessity of [Theorem 3](#) proceeds in five steps. First, we show that the marginal of $\bar{v}_\sigma^{T,1}$ on $\Delta(\tilde{\Omega})$, which we denote by τ_σ^T weak*-converges to $\mathbb{P}_\sigma \circ \mu_\infty^{-1}$. We denote this limit measure by τ_σ . By [Lemma C.2](#), τ_σ is also the (limit) marginal of $\bar{v}_\sigma^{T,2}$ on $\Delta(\tilde{\Omega})$.

Second, we show that up to a subsequence $\bar{v}_\sigma^{T,1}, \bar{v}_\sigma^{T,2} \xrightarrow{w^*} \bar{v}_\sigma$. Furthermore, transition probabilities $\tau_\sigma \in \Delta(\Delta(\tilde{\Omega})), \rho : \Theta \times \Delta(\tilde{\Omega}) \rightarrow \Delta(M), \alpha' : M \times \Delta(\tilde{\Omega}) \rightarrow \Delta(A)$ exist such that

$$v_\sigma(a, \theta, m, \tilde{\omega}) = \int_{\Delta(\tilde{\Omega})} \mu(\tilde{\omega}) f(\theta) \rho(m|\theta, \mu) \alpha'(a|m, \mu) \tau_\sigma(d\mu). \quad (\text{C.7})$$

Hence, the agent's payoff when faced with mechanism ϕ and playing strategy σ can be written as:

$$\mathbb{E}_{\bar{v}_\sigma} [u(a, \theta, \omega)] = \int_{\Delta(\tilde{\Omega})} \sum_{\theta \in \Theta} f(\theta) \sum_{m \in M} \rho(m|\theta, \mu) \mathbb{E}_{\tilde{\omega} \sim \mu} \left[\sum_{a \in A} \alpha'(a|\mu, m) u(a, \theta, \omega) \right] \tau_\sigma(d\mu). \quad (\text{C.8})$$

Third, we show that for all $\theta \in \Theta$

$$\mathbb{E}_{\tau_\sigma} \left\{ \sum_{m \in M} \rho(m|\theta, \mu) \mathbb{E}_{\tilde{\omega} \sim \mu} \left[\sum_{a \in A} \alpha'(a|\mu, m) u(a, \theta, \omega) \right] - \max_{m \in M} \mathbb{E}_{\tilde{\omega} \sim \mu} \left[\sum_{a \in A} \alpha'(a|\mu, m) u(a, \theta, \omega) \right] \right\} = 0. \quad (\text{C.9})$$

[Equations C.8](#) and [C.9](#) allow us to identify the incentive compatible and individually rational allocation rule of the two-stage mechanism that implements ϑ .

Fourth, whereas the previous steps identify a two-stage mechanism expressed in terms of posterior beliefs over $\tilde{\Omega}$, we show how to obtain a two-stage mechanism expressed in terms of posterior beliefs over Ω . Finally, we show that the agent's payoff in [Equation C.8](#) coincides with the payoff she would get when best responding to the information structure calibrated to ϕ .

Step 1 Having defined the extended occupation measure in [Equation C.5](#), we present here a property we use in our proof. Let τ_σ^T denote the marginal of $\bar{v}_\sigma^{T,1}$ on $\Delta(\tilde{\Omega})$. That is, for any measurable subset $\tilde{\Delta} \subset \Delta(\tilde{\Omega})$, define

$$\tau_\sigma^T(\tilde{\Delta}) = \frac{1}{T} \mathbb{E}_\sigma \left[\sum_{t=1}^T \mathbb{1}[\mu_t \in \tilde{\Delta}] \right]. \quad (\text{C.10})$$

In [Appendix D](#), we prove the following:

Lemma C.3. *The sequence of measures $(\tau_\sigma^T)_{T \in \mathbb{N}}$ defined by [Equation C.10](#) converges in the weak* sense to the push-forward measure $\tau_\sigma \equiv \mathbb{P}_\sigma \circ \mu_\infty^{-1}$, where $\mu_\infty(h^\infty) = \mathbb{P}_\sigma(\cdot | h^\infty)$.*

Step 2 To show that Equation C.7 holds, we show the following properties of $\bar{v}_\sigma^{T,1}$ and $\bar{v}_\sigma^{T,2}$. On the one hand, $\bar{v}_\sigma^{T,1}$ satisfies that for all $g \in C_b(A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega}))$,

$$\int_{A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega})} g(a, \theta, m, \tilde{\omega}, \mu) d\bar{v}_\sigma^{T,1} = \int_{\Theta \times M \times \Delta(\tilde{\Omega})} \mathbb{E}_\mu [\mathbb{E}_{\phi(\cdot|m, \tilde{\omega})} [g(a, \theta, m, \tilde{\omega}, \mu)]] d\bar{v}_{\sigma, \Theta M \Delta}^{T,1}, \quad (\text{C.11})$$

and for all $q \in C_b(\Theta \times \Delta(\Omega))$,

$$\int_{\Theta \times \Delta(\tilde{\Omega})} q(\theta, \mu) d\bar{v}_{\sigma, \Theta \Delta}^{T,1} = \int_{\Delta(\tilde{\Omega})} \int_\Theta f(\theta) q(\theta, \mu) d\tau_\sigma^T. \quad (\text{C.12})$$

On the other hand, $\bar{v}_\sigma^{T,2}$ satisfies that for all $g \in C_b(A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega}))$,

$$\int_{A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega})} g(a, \theta, m, \tilde{\omega}, \mu) d\bar{v}_\sigma^{T,2} = \int_{\Theta \times M \times A \times \Delta(\tilde{\Omega})} \mathbb{E}_{\tilde{\omega} \sim \mu} [g(a, \theta, m, \tilde{\omega}, \mu)] d\bar{v}_{\sigma, \Theta M A \Delta}^{T,2}. \quad (\text{C.13})$$

In the expressions above, the subscripts on $\bar{v}_\sigma^{T,k}$ next to σ are the spaces over which we take the marginals, and Δ is shorthand notation for $\Delta(\tilde{\Omega})$. Because $\Delta(A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega}))$ is compact (Aliprantis and Border, 2006, Theorem 15.11), $\bar{v}_\sigma^{T,1}$ has a convergent subsequence $(\bar{v}_\sigma^{T,n,1})_{n \in \mathbb{N}}$, which by Lemma C.2 is also a convergent subsequence of $\bar{v}_\sigma^{T,2}$. Let \bar{v}_σ denote the weak* limit along T_n . The continuity of the projection implies that ν_σ is the marginal of \bar{v}_σ on $A \times \Theta \times M \times \tilde{\Omega}$, and $\tau_\sigma \equiv \mathbb{P}_\sigma \circ \mu_\infty^{-1}$ is the marginal on $\Delta(\Omega)$. Equations C.12 and C.13 together imply that \bar{v}_σ admits the decomposition in the right hand side of Equation C.7, and the result follows.

To show Equation C.11 holds, use that $\bar{v}_\sigma^{T,1}$ has finite support to write it as follows:

$$\begin{aligned} \bar{v}_\sigma^{T,1}(a, \theta, m, \tilde{\omega}, \mu) &= \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t} \mathbb{P}_\sigma^t(\tilde{\omega}, h^t) f(\theta) \sigma_t(h^t, \theta)(m) \phi(a|m, \tilde{\omega}) \mathbb{1}[\mu_t = \mu] \\ &= \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t} \mathbb{P}_\sigma^t(h^t) \mu(\tilde{\omega}) f(\theta) \sigma_t(h^t, \theta)(m) \phi(a|m, \tilde{\omega}) \mathbb{1}[\mu_t = \mu], \end{aligned} \quad (\text{C.14})$$

where the second equality uses that $\mu_t(h^t) = \mathbb{P}_\sigma^t(\cdot|h^t)$.

Equation C.14 implies the following holds for every bounded continuous function $g \in C_b(A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega}))$:

$$\mathbb{E}_{\bar{v}_\sigma^{T,1}} [g(a, \theta, m, \tilde{\omega}, \mu)] = \mathbb{E}_{\bar{v}_{\sigma, \Theta M \Delta}^{T,1}} [\mathbb{E}_\mu [\mathbb{E}_{\phi(\cdot|m, \tilde{\omega})} [g(a, \theta, m, \tilde{\omega}, \mu)]]],$$

This completes the proof that Equation C.11 holds. Letting $V_g(\theta, M, \mu) = \mathbb{E}_\mu [\mathbb{E}_{\phi(\cdot|m, \tilde{\omega})} [g(a, \theta, m, \tilde{\omega}, \mu)]]$, we have that

$$\int_{A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega})} g d\bar{v}_\sigma^{T,1} = \int_{\Theta \times M \times \Delta(\tilde{\Omega})} V_g d\bar{v}_{\sigma, \Theta M \Delta}^{T,1} \Leftrightarrow \int_{A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega})} (g - V_g) d\bar{v}_\sigma^{T,1} = 0.$$

Because $g - V_g \in C_b(A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega}))$ and $\bar{v}_\sigma^{T,n,1} \xrightarrow{w^*} \bar{v}_\sigma$, we conclude that

$$\mathbb{E}_{\bar{v}_\sigma} [g(a, \theta, m, \tilde{\omega}, \mu)] = \int_{\Theta \times M \times \Delta(\tilde{\Omega})} \mathbb{E}_\mu [\mathbb{E}_{\phi(\cdot|m, \tilde{\omega})} [g(a, \theta, m, \tilde{\omega}, \mu)]] d\bar{v}_{\sigma, \Theta M \Delta}. \quad (\text{C.15})$$

To show that Equation C.12 holds, note that the marginal of $\bar{v}_\sigma^{T,1}$ on $\Theta \times \Delta(\tilde{\Omega})$ equals $\tau_\sigma^T \otimes f$. Indeed, fix

any continuous function $q \in C_b(\Theta \times \Delta(\tilde{\Omega}))$ and note that for all T

$$\mathbb{E}_{\bar{v}_{\sigma, \Theta \Delta}^{T,1}} [q(\theta, \mu)] = \mathbb{E}_{\tau_{\sigma}^T} \left[\sum_{\theta \in \Theta} f(\theta) q(\theta, \mu) \right].$$

Letting $V_q(\mu) = \sum_{\theta \in \Theta} f(\theta) q(\theta, \mu)$, we have that for all T

$$\int_{\Delta(\tilde{\Omega}) \times \Theta} (q(\theta, \mu) - V_q(\mu)) d\bar{v}_{\sigma, \Theta \Delta}^{T,1} = 0.$$

Because $q - V_q \in C_b(\Theta \times \Delta(\tilde{\Omega}))$ and $\bar{v}_{\sigma}^{T_n,1} \xrightarrow{w^*} \bar{v}_{\sigma}$, we conclude that

$$\int_{\Theta \times \Delta(\tilde{\Omega})} q(\theta, \mu) d\bar{v}_{\sigma, \Theta \Delta} = \int_{\Delta(\tilde{\Omega})} \int_{\Theta} f(\theta) q(\theta, \mu) d\tau_{\sigma}. \quad (\text{C.16})$$

Lastly, to show that Equation C.13 holds, note that we can write $\bar{v}_{\sigma}^{T,2}$ as follows (once again, we use that for finite T , it has finite support):

$$\begin{aligned} \bar{v}_{\sigma}^{T,2}(a, \theta, m, \tilde{\omega}, \mu) &= \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t} \mathbb{P}_{\sigma}^t(h^t) \mathbb{P}_{\sigma}^t(\tilde{\omega} | h^t) f(\theta) \sigma_t(h^t, \theta)(m) \phi(a | m, \tilde{\omega}) \mathbb{1}[\mu_{t+1}(h^t, \theta, m, a) = \mu] = \\ &= \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t: \mu_{t+1}(h^t, \theta, m, a) = \mu} \mathbb{P}_{\sigma}^{t+1}(\tilde{\omega}, h^t, \theta, m, a) = \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t: \mu_{t+1}(h^t, \theta, m, a) = \mu} \mu(\tilde{\omega}) \mathbb{P}_{\sigma}^{t+1}(h^t, \theta, m, a) \\ &= \mu(\tilde{\omega}) \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t: \mu_{t+1}(h^t, \theta, m, a) = \mu} \mathbb{P}_{\sigma}^{t+1}(h^t, \theta, m, a) = \mu(\tilde{\omega}) \bar{v}_{\sigma, A \Theta M \Delta}^{T,2}(a, \theta, m, \mu), \end{aligned}$$

where the last expression follows from noting that the term multiplying $\mu(\omega)$ in the first expression in the third line is $\sum_{\tilde{\omega}} v_{\sigma}^{T,2}(a, \theta, m, \tilde{\omega}, \mu)$.

Then, for every $g \in C_b(A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega}))$, we have that

$$\mathbb{E}_{\bar{v}_{\sigma}^{T,2}} [g(a, \theta, m, \tilde{\omega}, \mu)] = \mathbb{E}_{\bar{v}_{\sigma, A \Theta M \Delta(\tilde{\Omega})}^{T,2}} [\mathbb{E}_{\tilde{\omega} \sim \mu} [g(a, \theta, m, \tilde{\omega}, \mu)]],$$

which completes the proof that Equation C.13 holds. Letting $V_g(a, \theta, m, \mu) = \mathbb{E}_{\tilde{\omega} \sim \mu} [g(a, \theta, m, \tilde{\omega}, \mu)]$, we have that

$$\int (g - V_g) d\bar{v}_{\sigma}^{T,2} = 0.$$

Because $g - V_g \in C_b(A \times \Theta \times M \times \tilde{\Omega} \times \Delta(\tilde{\Omega}))$ and $\bar{v}_{\sigma}^{T_n,2} \xrightarrow{w^*} \bar{v}_{\sigma}$, we conclude that

$$\mathbb{E}_{\bar{v}_{\sigma}} [g(a, \theta, m, \tilde{\omega}, \mu)] = \int_{A \times \Theta \times M \times \Delta(\tilde{\Omega})} \mathbb{E}_{\tilde{\omega} \sim \mu} [g(a, \theta, m, \tilde{\omega}, \mu)] \bar{v}_{\sigma, A \Theta M \Delta}(d(a, \theta, m, \mu)). \quad (\text{C.17})$$

Equations C.16 and C.17 imply \bar{v}_{σ} admits the following disintegration:

$$\bar{v}_{\sigma}(\{(a, \theta, m, \tilde{\omega})\} \times \tilde{\Delta}) = \int_{\tilde{\Delta}} \mu(\tilde{\omega}) f(\theta) \alpha'(a | \theta, m, \mu) \rho(m | \theta, \mu) \tau_{\sigma}(d\mu), \quad (\text{C.18})$$

where we disintegrated $\bar{v}_{\sigma, A \Theta M \Delta}$ first along $\Theta \times \Delta(\tilde{\Omega})$ —and used Equation C.16 to obtain the independence

of Θ and $\Delta(\tilde{\Omega})$ —and then further disintegrated the distribution of $A \times M$ conditional on $\Theta \times \Delta(\tilde{\Omega})$. Now, Equation C.15 implies that the following also holds

$$\bar{v}_\sigma(\{(a, \theta, m, \tilde{\omega})\} \times \tilde{\Delta}) = \int_{\tilde{\Delta}} \mu(\tilde{\omega}) f(\theta) \phi(a|m, \tilde{\omega}) \rho(m|\theta, \mu) \tau_\sigma(d\mu), \quad (\text{C.19})$$

where once again we use the uniqueness of disintegration. Because Equations C.18 and C.19 hold for any tuple $(a, \theta, m, \tilde{\omega})$ and measurable subset $\tilde{\Delta}$ of $\Delta(\tilde{\Omega})$, we conclude that (i) $\alpha'(a|\theta, m, \mu)$ does not depend on θ τ_σ -almost everywhere, and (ii) $\phi(\cdot|m, \tilde{\omega})$ is constant on $\tilde{\omega}$ in the support of μ τ_σ -almost everywhere. This concludes the proof of Step 2.

Step 3 We now argue that the agent achieves

$$u^*(\mu) \equiv \sum_{\theta \in \Theta} f(\theta) \max_{m \in M} \sum_{\tilde{\omega} \in \tilde{\Omega}} \mu(\tilde{\omega}) \sum_{a \in A} \phi(a|m, \tilde{\omega}) u(a, \theta, \omega) = \sum_{\theta \in \Theta} f(\theta) \max_{m \in M} \sum_{\tilde{\omega} \in \tilde{\Omega}} \mu(\tilde{\omega}) \sum_{a \in A} \alpha'(a|m, \mu) u(a, \theta, \omega), \quad (\text{C.20})$$

on the support of τ_σ , where the second equality follows from Step 2. Toward a contradiction, suppose this is not the case; that is,

$$\mathbb{E}_{\tau_\sigma} \left[\sum_{\theta \in \Theta} f(\theta) \sum_{m \in M} \rho(m|\theta, \mu) \sum_{\tilde{\omega} \in \tilde{\Omega}} \mu(\tilde{\omega}) \sum_a \phi(a|m, \tilde{\omega}) u(a, \theta, \omega) \right] < \mathbb{E}_{\tau_\sigma} \left[\sum_{\theta \in \Theta} u^*(\mu) \right] = U^*.$$

We show that the agent can achieve a payoff arbitrarily close to U^* by playing according to σ until some finite T and then best-responding to her beliefs at time T in every period thereafter; a contradiction.

Consider a strategy σ' which until some period T plays according to σ and after period T best responds to $\mu_T(h^T) \in \Delta(\tilde{\Omega})$. Because payoffs accumulated on a finite number of periods are irrelevant to long-run payoffs, this strategy results in a payoff:

$$\begin{aligned} & \sum_{h^T \in H^T} \mathbb{P}_\sigma^T(h^T) \sum_{\theta \in \Theta} f(\theta) \max_{m \in M} \left[\sum_{\tilde{\omega} \in \tilde{\Omega}} \mu_T(\tilde{\omega}) \sum_{a \in A} \phi(a|m, \tilde{\omega}) u(a, \theta, \omega) \right] = \sum_{h^T \in H^T} \mathbb{P}_\sigma^T(h^T) u^*(\mu_T(h^T)) \\ &= \mathbb{E}_{\mathbb{P}_\sigma^T \circ \mu_T^{-1}} [u^*(\mu)] = \mathbb{E}_{\mathbb{P}_\sigma \circ \mu_T^{-1}} [u^*(\mu)], \end{aligned}$$

where the last equality follows as μ_T is adapted to the histories through T . Similar arguments to Lemma C.3 imply that $\mathbb{P}_\sigma \circ \mu_T^{-1} \xrightarrow{w^*} \mathbb{P}_\sigma \circ \mu_\infty^{-1} \equiv \tau_\sigma$. Noting that $u^* : \Delta(\tilde{\Omega}) \rightarrow \mathbb{R}$ is continuous and bounded (as it is the maximum of linear functions in beliefs), we obtain that as $T \rightarrow \infty$,

$$\mathbb{E}_{\mathbb{P}_\sigma \circ \mu_T^{-1}} [u^*(\mu)] \rightarrow \mathbb{E}_{\tau_\sigma} [u^*(\mu)].$$

It follows that for every $\delta > 0$, we can find T large enough so that $|\mathbb{E}_{\mathbb{P}_\sigma \circ \mu_T^{-1}} [u^*(\mu)] - U^*| < \delta$, contradicting the optimality of σ .

We conclude that $\alpha : \Theta \times \Delta(\tilde{\Omega}) \rightarrow \Delta(A)$ defined as follows:

$$\alpha(a|\theta, \mu) = \sum_{m \in M} \rho(m|\theta, \mu) \alpha'(a|m, \mu),$$

is incentive compatible and individually rational τ_σ -almost everywhere.

Step 4 We now show how to derive a two-stage mechanism $\beta^* : \Omega \rightarrow \Delta(\Delta(\Omega))$ and an allocation rule $\alpha^* : \Theta \times \Delta(\Omega) \rightarrow \Delta(A)$ that implement ϑ . First, note that the agent's payoff when her type is θ and the induced belief is $\mu \in \Delta(\tilde{\Omega})$, can be written as

$$\sum_{\tilde{\omega} \in \tilde{\Omega}} \mu(\tilde{\omega}) \sum_{a \in A} \alpha(a|\theta, \mu) u(a, \theta, \tilde{\omega}) = \sum_{\omega \in \Omega} \mu_\Omega(\omega) \sum_{a \in A} \alpha(a|\theta, \mu) u(a, \theta, \omega),$$

where μ_Ω is the marginal of μ on Ω and the equality follows because the realization of ε is payoff-irrelevant. By Step 3, $\alpha(\cdot|\theta, \mu)$ is individually rational and incentive compatible when the agent holds belief μ_Ω .

Furthermore, for each $(a, \theta, \omega) \in A \times \Theta \times \Omega$, we have

$$\vartheta(a, \theta, \omega) = \sum_{\varepsilon \in \mathcal{E}} \int_{\Delta(\tilde{\Omega})} \mu(\omega, \varepsilon) f(\theta) \alpha(a|\theta, \mu) \tau_\sigma(d\mu) = \int_{\Delta(\tilde{\Omega})} \mu_\Omega(\omega) f(\theta) \alpha(a|\theta, \mu) \tau_\sigma(d\mu).$$

For each $\theta \in \Theta$, consider the joint distribution $Q_\theta \in \Delta(A \times \Delta(\Omega))$ defined as follows:

$$Q_\theta(\{a\} \times \tilde{\Delta}) = \int_{\Delta(\tilde{\Omega})} \mathbb{1}[\mu_\Omega \in \tilde{\Delta}] \alpha(a|\theta, \mu) \tau_\sigma(d\mu) = \int_{\tilde{\Delta}} \alpha^*(a|\theta, \mu_\Omega) \tau^*(d\mu_\Omega),$$

where the third equality follows from disintegration (note $\alpha^*(\cdot|\cdot, \mu_\Omega) = \mathbb{E}[\alpha(\cdot|\cdot, \tilde{\mu}) | \tilde{\mu}_\Omega = \mu_\Omega]$). By the first argument in Step 4, $\alpha^*(\cdot|\cdot, \mu_\Omega)$ is individually rational and incentive compatible when the agent holds μ_Ω . We obtain that

$$\vartheta(a, \theta, \omega) = \int_{\Delta(\Omega)} \mu_\Omega(\omega) f(\theta) \alpha^*(a|\theta, \mu_\Omega) \tau^*(d\mu_\Omega).$$

Defining for all $\omega \in \Omega$ and measurable subsets $\tilde{\Delta} \in \Delta(\Omega)$,

$$\beta^*(\tilde{\Delta}|\omega) = \int_{\tilde{\Delta}} \frac{\mu(\omega)}{\mu_0(\omega)} \tau^*(d\mu),$$

Steps 2-4 together imply that ϑ can be implemented by the incentive compatible and individually rational two-stage mechanism (β^*, α^*) .

Step 5: The agent adequately learns Finally, we argue that the agent earns the same payoff as if she had access to the information structure calibrated to ϕ, π_ϕ .

Lemma C.4. *Let τ_ϕ denote the belief distribution induced by the calibrated information structure π_ϕ . Then, the agent's payoff under σ equals*

$$U(\tau_\phi) \equiv \mathbb{E}_{\tau_\phi} \left[\sum_{\theta \in \Theta} f(\theta) \max_{m \in M} \sum_{\tilde{\omega} \in \tilde{\Omega}} \mu(\tilde{\omega}) \sum_{a \in A} \phi(a|m, \tilde{\omega}) u(a, \theta, \tilde{\omega}) \right]. \quad (\text{C.21})$$

The proof of this is standard, and hence we defer it to [Appendix D](#).

C.1.2 Proof of Theorem 3 (sufficiency)

Suppose $\vartheta \in \Delta(A \times \Theta \times \Omega)$ is implemented by an incentive compatible and individually rational two-stage mechanism. That is,

$$\vartheta(a, \theta, \omega) = \mu_0(\omega) f(\theta) \int_{\Delta(\Omega)} \alpha(a|\theta, \mu) \beta(d\mu|\omega), \quad (\text{C.22})$$

and α is incentive compatible and individually rational on the support of $\mu_0 \otimes \beta$. As in the proof of Theorem 2, a finite support $\beta' : \Omega \rightarrow \Delta(\{\mu_1, \dots, \mu_K\})$ exists such that (β', α) implement ϑ . As in Green and Stokey (2022), the experiment β' can be generated by a finite information structure $\pi : \Omega \times \mathcal{E} \rightarrow \Delta(\Omega)$, where (i) \mathcal{E} is finite, (ii) \mathcal{E} is independent of Ω , and (iii) $\mu = \pi(\omega, \varepsilon)$.

Construct a mechanism $\phi : \Theta \times \Omega \times \mathcal{E} \rightarrow \Delta(A)$ such that $\phi(\cdot|\theta, \omega, \varepsilon) = \alpha(\cdot|\theta, \pi(\omega, \varepsilon))$. (Note that π is information structure calibrated to ϕ , but expressed in beliefs.) Consider now the extensive form game induced by such a mechanism.⁴⁴

If the agent truthfully reports her type, then the occupation measure induces outcome distribution ϑ . Hence, under truthtelling, the agent's payoff is:

$$\begin{aligned} U(\sigma_{\text{truth}}) &= \sum_{(a, \theta, \omega) \in A \times \Theta \times \Omega} \vartheta(a, \theta, \omega) u(a, \theta, \omega) = \mathbb{E}_{\tau_\phi} \left[\sum_{\theta \in \Theta} f(\theta) \sum_{a \in A} \alpha(a|\theta, \mu) \sum_{\omega \in \Omega} \mu(\omega) u(a, \theta, \omega) \right] \\ &= \mathbb{E}_{\tau_\phi} \left[\sum_{\theta \in \Theta} f(\theta) \max \left\{ \max_{\theta' \in \Theta} \sum_{a \in A} \alpha(a|\theta', \mu) \sum_{\omega \in \Omega} \mu(\omega) u(a, \theta, \omega), \sum_{\omega \in \Omega} \mu(\omega) u(a_\phi, \theta, \omega) \right\} \right] = \\ &= \mathbb{E}_{\mathcal{E}} \left[\sum_{\theta \in \Theta} f(\theta) \max \left\{ \max_{\theta' \in \Theta} \sum_{\omega \in \Omega} \mu(\omega) \sum_{a \in A} \phi(a|\theta', \omega, \varepsilon) u(a, \theta, \omega), \sum_{\omega \in \Omega} \mu(\omega) u(a_\phi, \theta, \omega) \right\} \right], \end{aligned} \quad (\text{C.23})$$

where (i) τ_ϕ is the belief distribution induced by the information structure π , and (ii) the first equality is by definition of the occupation measure, the second is the definition that ϑ is implemented by the two-stage mechanism, the third follows from incentive compatibility and individual rationality of α , and the fourth is definitional.

Moreover, the payoff in the last line of Equation C.23 is the payoff the agent obtains by using the “learning” strategy in Lemma C.4, which first extracts all the mechanism can teach her about the state and then uses that information to optimize over her participation and reporting strategies. It follows that truthtelling (and participation) are optimal and ϑ is implemented by repeated mechanism ϕ .

C.2 Dynamic Mechanisms

In this section, we present the proof of Theorem 4. To do so, we first complete the formal definition of the game, by specifying the histories, strategy space, and the distribution over terminal histories induced by the agent's strategy and the mechanism. Having laid this groundwork, we describe the proof strategy, and then provide the formal details of the proof.

⁴⁴We could expand the mechanism by allowing the agent to have a message which triggers the outside option, but this is not necessary as α is individually rational.

Mechanisms, histories, and strategies A dynamic mechanism $(\varphi_t)_{t \in \mathbb{N}}$ is a sequence of mappings that condition on the state, the agent's report history, the allocation history, and today's report and output an allocation. By the revelation principle, it is without loss of generality to restrict attention to mechanisms that solicit type reports.

As in the main text, we expand the set of type reports and allocations by the non-participation decision and the outside option, which we denote by $\Theta A_\emptyset \equiv A \times \Theta \cup \{(\emptyset, a_\emptyset)\}$. Then, $\hat{H}^t = (\Theta A_\emptyset)^{t-1}$ denotes the histories of reports (inclusive of the non-participation decision) and allocations at the beginning of time $t \in \mathbb{N}$, and let $\hat{\mathcal{H}}^t = \Omega \times \hat{H}^t$. Similarly, let $\hat{H}^\infty = \times_{t \in \mathbb{N}} (\Theta A_\emptyset)$ denote the set of all possible report-allocation outcome paths, and let $\hat{\mathcal{H}}^\infty = \Omega \times \hat{H}^\infty$. A dynamic mechanism is then a collection of mappings $(\varphi_t)_{t \in \mathbb{N}}$ such that $\varphi_t : \hat{\mathcal{H}}^t \times \Theta \rightarrow \Delta(A)$.

To define the agent's strategy, let $H^t = \Theta^{t-1} \times \hat{H}^{t-1}$, where the coordinates denote the sequence of realized types, reports (inclusive of participation decisions), and allocations through period $t-1$. A behavioral strategy is a mapping $(p_t, \sigma_t) : H^t \times \Theta \rightarrow [0, 1] \times \Delta(\Theta)$.

The distribution over terminal histories \mathcal{H}^∞ To obtain the complete description of the paths on the tree we need to append Ω to H^t ; hence the paths through period $t-1$ are $\Omega \times H^t \equiv \mathcal{H}^t$. The distributions over states, agent's types, the agent's strategy, and the mechanism induce a distribution over the terminal histories $\mathcal{H}^\infty \equiv \Omega \times H^\infty$, which we denote by $\mathbb{P}_{(p, \sigma)} \in \Delta(\Omega \times H^\infty)$. We denote by $\mathbb{E}_{(p, \sigma)}$ the expectation under this measure. The distribution $\mathbb{P}_{(p, \sigma)} \in \Delta(\Omega \times H^\infty)$ is the unique distribution that satisfies that for all $t \in \mathbb{N}$, $\tilde{\mathcal{H}}^t \subset \Omega \times \mathcal{E} \times H^t$,

$$\mathbb{P}_{(p, \sigma)}(\tilde{\mathcal{H}}^t \times \prod_{s=t+1}^{\infty} (\Theta \times \Theta A_\emptyset)) = \mathbb{P}_{\sigma}^t(\tilde{\mathcal{H}}^t),$$

where the distributions $(\mathbb{P}_{(p, \sigma)}^t)_{t \in \mathbb{N}}$ satisfy (under participation and truthtelling)

$$\mathbb{P}_{(p, \sigma)}^{t+1}(\omega, h^t, \theta, \theta', a) = \mathbb{P}_{(p, \sigma)}^t(\omega, h^t) f(\theta) \mathbb{1}[\theta' = \theta] \varphi_t(a | \omega, \hat{h}^t, \theta').$$

Implementation We focus on incentive-compatible mechanisms φ for which (i) a best response, (p, σ) , exists, and (ii) the occupation measure $\nu_\sigma \in \Delta(A \times \Theta \times \Omega)$ exists, where

$$\nu_{(p, \sigma)}(a, \theta, \omega) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{(p, \sigma)} \left[\sum_{t=1}^T \mathbb{1}[(a_t, \theta_t, \omega') = (a, \theta, \omega)] \right], \quad (\text{C.24})$$

where the limit is in the weak* sense. In contrast to [Appendix C.1](#), we do not keep track of the agent's type reports in the occupation measure, only the agent's types. Under (p, σ) only truthtelling histories have positive probability.

C.2.1 Proof of [Theorem 4](#) (necessity)

Let $\vartheta \in \Delta(A \times \Theta \times \Omega)$ denote the outcome distribution implemented by an incentive compatible dynamic mechanism φ , and let $\nu_{(p, \sigma)}$ denote the occupation measure under the agent's truthtelling strategy. Below, we show that $\nu_{(p, \sigma)}$, and hence ϑ , can be implemented by a two-stage mechanism which lacks profitable undetectable deviations and is ex ante individually rational.

Analogously to the proof of [Theorem 3](#), we define two sequences of extended occupation measures on $A \times \Theta \times \Omega \times \Delta(\Omega)$ defined as follows. Letting $\tilde{\Delta}$ denote a measurable subset of $\Delta(\Omega)$, define

$$\bar{v}_{(p,\sigma)}^{T,1}(\{(a, \theta, \omega)\} \times \tilde{\Delta}) = \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t} \mathbb{P}_{(p,\sigma)}^t(\omega, h^t) f(\theta) \sigma_t(h^t, \theta) \varphi(\omega, \hat{h}^t, \theta)(a) \mathbb{1}[\mu_t(h^t) \in \tilde{\Delta}], \quad (\text{C.25})$$

$$\bar{v}_{(p,\sigma)}^{T,2}(\{(a, \theta, \omega)\} \times \tilde{\Delta}) = \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t} \mathbb{P}_{(p,\sigma)}^t(\omega, h^t) f(\theta) \sigma_t(h^t, \theta) \varphi(\omega, \hat{h}^t, \theta)(a) \mathbb{1}[\mu_{t+1}(h^t, \theta, \theta, a) \in \tilde{\Delta}]. \quad (\text{C.26})$$

The proof proceeds similarly to that in [Appendix C.1](#). First, we show that the occupation measure $\nu_{(p,\sigma)} \in \Delta(A \times \Theta \times \Omega)$ admits the following decomposition

$$\nu_{(p,\sigma)}(a, \theta, \omega) = \int_{\Delta(\Omega)} f(\theta) \mu(\omega) \alpha(a|\theta, \mu) \tau_{(p,\sigma)}(d\mu), \quad (\text{C.27})$$

where $\tau_{(p,\sigma)}$ is the distribution over terminal beliefs (cf. [Lemma C.3](#)) and the transition probability $\alpha : \Theta \times \Delta(\Omega) \rightarrow \Delta(A)$ is our candidate allocation rule. Consequently, the agent's equilibrium payoff can be written as follows:

$$\sum_{(a, \theta, \omega) \in A \times \Theta \times \Omega} \nu_{(p,\sigma)}(a, \theta, \omega) u(a, \theta, \omega) = \int_{\Delta(\Omega)} \left[\sum_{\theta \in \Theta} f(\theta) \sum_{\omega \in \Omega} \mu(\omega) \sum_{a \in A} \alpha(a|\theta, \mu) u(a, \theta, \omega) \right] \tau_{(p,\sigma)}(d\mu). \quad (\text{C.28})$$

Second, we show that the allocation rule lacks profitable undetectable deviations and is ex ante individually rational.

The occupation measure satisfies Equation C.27 To prove that Equation C.27 holds, we first show that for all $g \in C_b(A \times \Theta \times \Omega \times \Delta(\Omega))$ and all $T \in \mathbb{N}$,

$$\int_{A \times \Theta \times \Omega \times \Delta(\Omega)} g(a, \theta, \omega, \mu) d\bar{v}_{(p,\sigma)}^{T,2} = \int_{A \times \Theta \times \Delta(\Omega)} \mathbb{E}_\mu[g(a, \theta, \omega, \mu)] d\bar{v}_{(p,\sigma), A \Theta \Delta(\Omega)}^{T,2} \quad (\text{C.29})$$

and for all $q \in C_b(\Theta \times \Delta(\Omega))$ and all $T \in \mathbb{N}$,

$$\int_{\Theta \times \Delta(\Omega)} q(\theta, \mu) d\bar{v}_{(p,\sigma), \Theta \Delta}^{T,1} = \int_{\Delta(\Omega)} \sum_{\theta \in \Theta} f(\theta) q(\theta, \mu) d\bar{v}_{(p,\sigma), \Delta}^{T,1}, \quad (\text{C.30})$$

where the subscripts on \bar{v} next to (p, σ) are the spaces over which we take the marginals, and Δ is shorthand notation for $\Delta(\Omega)$. We skip the proof of this step as it basically repeats the proof of the analogous step in [Appendix C.1](#).

Because $\Delta(A \times \Theta \times \Omega \times \Delta(\Omega))$ is compact (Aliprantis and Border, 2006, Theorem 15.11), $\bar{v}_{(p,\sigma)}^{T,1}$ has a convergent subsequence $(\bar{v}_{(p,\sigma)}^{T_n,1})_{n \in \mathbb{N}}$, which by [Lemma C.2](#) is also a convergent subsequence of $\bar{v}_{(p,\sigma)}^{T,2}$. Let $\bar{v}_{(p,\sigma)}$ denote the weak* limit along T_n . The continuity of the projection implies that $\nu_{(p,\sigma)}$ is the marginal of $\bar{v}_{(p,\sigma)}$ on $A \times \Theta \times \Omega$, and $\tau_{(p,\sigma)} \equiv \mathbb{P}_\sigma \circ \mu_\infty^{-1}$ is the marginal on $\Delta(\Omega)$. Moreover, [Equation C.29](#) and [Equation C.30](#) together imply that $\nu_{(p,\sigma)}$ admits the decomposition on the right hand side of [Equation C.27](#), and the result follows.

The allocation rule lacks profitable undetectable deviations We now show the allocation rule α admits no profitable undetectable deviations. An undetectable deviation is a transition probability σ' from $\Theta \times \Delta(\Omega)$ to $\Delta(\Theta)$ such that for all $\mu \in \Delta(\Omega)$ and $\theta' \in \Theta$

$$\sum_{\theta \in \Theta} f(\theta) \sigma'(\theta' | \theta, \mu) = f(\theta'). \quad (\text{C.31})$$

Consider a deviation by the agent to (p, σ') instead of (p, σ) . That is, when his type is θ and belief is μ , the agent chooses type θ' with probability $\sigma'(\theta' | \theta, \mu)$. In what follows, we index the induced distributions over histories only by σ and σ' as we are only changing the agent's reporting strategy. In particular, denote by $\mathbb{P}_{\sigma'}$ the induced probability distribution over terminal histories when the agent uses (p, σ') instead of (p, σ) .

We first claim that for every t the marginal of $\mathbb{P}_{\sigma'}^t$ over $\Omega \times \hat{H}^t$ coincides with that of \mathbb{P}_{σ}^t . Recall that for every t we have that

$$\mathbb{P}_{\sigma'}^{t+1}(\omega, h^t, \theta, \theta', a) = \mathbb{P}_{\sigma'}^t(\omega, h^t) f(\theta) \sigma'(\theta' | \mu_t(h^t), \theta) \varphi_t(a | \omega, \hat{h}^t, \theta').$$

Adding up over θ on both sides and using Equation C.31, we get:

$$\sum_{\theta \in \Theta} \mathbb{P}_{\sigma'}^{t+1}(\omega, h^t, \theta, \theta', a) = \mathbb{P}_{\sigma'}^t(\omega, h^t) f(\theta') \varphi_t(a | \omega, \hat{h}^t, \theta').$$

Now, note that $h^t = (\hat{h}^t, \tilde{\theta}^{t-1})$ for some sequence $\tilde{\theta}^{t-1} \in \Theta^{t-1}$. If we add up on both sides over all such sequences we get

$$\sum_{\theta \in \Theta, \tilde{\theta}^{t-1} \in \Theta^{t-1}} \mathbb{P}_{\sigma'}^{t+1}(\omega, \hat{h}^t, \tilde{\theta}^{t-1}, \theta, \theta', a) = \sum_{\tilde{\theta}^{t-1} \in \Theta^{t-1}} \mathbb{P}_{\sigma'}^t(\omega, \hat{h}^t, \tilde{\theta}^{t-1}) f(\theta') \varphi_t(a | \omega, \hat{h}^t, \theta').$$

Note that if the distribution over $\Omega \times \hat{H}^t$ induced by σ' up to period t is the same as that induced by σ , we get that the right-hand side equals:

$$\mathbb{P}_{\sigma, \hat{\mathcal{H}}^t}^t(\omega, \hat{h}^t) f(\theta') \varphi_t(a | \omega, \hat{h}^t, \theta'),$$

and hence $\mathbb{P}_{\sigma', \hat{\mathcal{H}}^{t+1}}^{t+1}(\omega, \hat{h}^t, \theta', a) = \mathbb{P}_{\sigma, \hat{\mathcal{H}}^t}^t(\omega, \hat{h}^t) f(\theta') \varphi_t(a | \omega, \hat{h}^t, \theta') = \mathbb{P}_{\sigma, \hat{\mathcal{H}}^{t+1}}^{t+1}(\omega, \hat{h}^t, \theta', a)$. By definition of $\mathbb{P}_{\sigma'}$, we conclude that $\mathbb{P}_{\sigma', \hat{\mathcal{H}}^\infty} = \mathbb{P}_{\sigma, \hat{\mathcal{H}}^\infty}$. Hence, the joint distribution over states, reports, and allocations is the same under σ and σ' .

Let $\bar{v}_{\sigma'}^{T,1}, \bar{v}_{\sigma'}^{T,2} \in \Delta(A \times \Theta \times \hat{\Theta} \times \Omega \times \Delta(\Omega))$ denote the analogue of the occupation measures in Equations C.25 and C.26 corresponding to σ' , extended to account for the agent's reports. Below, the notation $\hat{\Theta}$ signifies those are the agent's reports. In what follows, recalling that the belief system depends only on the reported history and not the type history is useful. Equation C.31 implies that for all measurable

subsets $\tilde{\Delta}$ of $\Delta(\Omega)$,

$$\begin{aligned}
\sum_{\theta \in \Theta} \bar{v}_{\sigma'}^{T,1}(\{(a, \theta, \theta', \omega)\} \times \tilde{\Delta}) &= \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t} \mathbb{P}_{\sigma'}^t(\omega, h^t) f(\theta') \varphi_t(\omega, \hat{h}^t, \theta')(a) \mathbb{1}[\mu_t(h^t) \in \tilde{\Delta}] \\
&= \frac{1}{T} \sum_{t=1}^T \sum_{\hat{h}^t \in \hat{H}^t} \mathbb{P}_{\sigma'}^t(\omega, \hat{h}^t) f(\theta') \varphi_t(\omega, \hat{h}^t, \theta')(a) \mathbb{1}[\mu_t(\hat{h}^t) \in \tilde{\Delta}] \\
&= \frac{1}{T} \sum_{t=1}^T \sum_{\hat{h}^t \in \hat{H}^t} \mathbb{P}_{\sigma'}^t(\omega, \hat{h}^t) f(\theta') \varphi_t(\omega, \hat{h}^t, \theta')(a) \mathbb{1}[\mu_t(\hat{h}^t) \in \tilde{\Delta}] = \bar{v}_{\sigma'}^{T,1}(\{(a, \theta', \omega)\} \times \tilde{\Delta}).
\end{aligned} \tag{C.32}$$

The first equality uses the definition of undetectability, the second uses that all the terms depend only on the reported history, the third uses that σ and σ' induce the same distribution over states, reports, and allocations, and the last is the definition of the occupation measure induced by σ . In words, the marginal of $\bar{v}_{\sigma'}^{T,1}$ over allocations, reports, states, and beliefs, $\bar{v}_{\sigma', A\hat{\Theta}\Omega\Delta}^{T,1}$ coincides with $\bar{v}_{\sigma}^{T,1}$.

We now show that $\bar{v}_{\sigma'}^{T,1}$ and $\bar{v}_{\sigma'}^{T,2}$ have a convergent subsequence with limit $\bar{v}_{\sigma'} \in \Delta(A \times \Theta \times \hat{\Theta} \times \Omega \times \Delta(\Omega))$ that admits the following decomposition:

$$\mathbb{E}_{\bar{v}_{\sigma'}}[u(a, \theta, \omega)] = \int_{\Delta(\Omega)} \left[\sum_{\theta \in \Theta} f(\theta) \sum_{\theta' \in \Theta} \sigma'(\theta, \mu)(\theta') \sum_a \alpha(a|\theta', \mu) u(a, \theta, \mu) \right] d\tau_{(p, \sigma)},$$

where $u(a, \theta, \mu)$ is the linear extension of $u(a, \theta, \cdot)$.

We proceed as follows: First, we show that for each T , under $v_{\sigma'}^{T,1}$, the allocation is independent of the true type conditional on the period- t belief and the reported type. Indeed,

$$\begin{aligned}
\sum_{\omega \in \Omega} v_{\sigma'}^{T,1}(a, \theta, \theta', \omega, \mu) &= \frac{1}{T} \sum_{t=1}^T \sum_{h^t \in H^t} \mathbb{P}_{\sigma'}^t(h^t) \left(\sum_{\omega \in \Omega} \mathbb{P}_{\sigma'}(\omega|h^t) \varphi_t(\omega, \hat{h}^t, \theta')(a) \right) f(\theta) \sigma'(\mu_t(h^t), \theta)(\theta') \mathbb{1}[\mu_t(h^t) = \mu] \\
&= \frac{f(\theta) \sigma'(\theta, \mu)(\theta')}{f(\theta')} \left(\frac{1}{T} \sum_{t=1}^T \sum_{\hat{h}^t: \mu_t(\hat{h}^t) = \mu} \mathbb{P}_{\sigma'}^t(\hat{h}^t) \sum_{\omega \in \Omega} \mathbb{P}_{\sigma'}(\omega|\hat{h}^t) f(\theta') \varphi_t(\hat{h}^t, \theta')(a) \right) \\
&= \frac{f(\theta) \sigma'(\theta, \mu)(\theta')}{f(\theta')} v_{\sigma', A\hat{\Theta}\Delta}^{T,1}(a, \theta', \mu) = \frac{f(\theta) \sigma'(\theta, \mu)(\theta')}{f(\theta')} v_{\sigma, A\Theta\Delta}^{T,1}(a, \theta', \mu),
\end{aligned}$$

where the third and fourth equalities use [Equation C.25](#). Moreover, the same analysis as that under σ implies the agent's true type is independent of the belief.

Therefore, $\bar{v}_{\sigma', A\hat{\Theta}\Delta}^{T,1}$ admits decomposition:

$$\bar{v}_{\sigma', A\hat{\Theta}\Delta}^{T,1}(a, \theta, \theta', \mu) = \frac{f(\theta) \sigma'(\theta'|\theta, \mu)}{f(\theta')} \bar{v}_{\sigma', A\hat{\Theta}\Delta}^{T,1}(a, \theta', \mu) = \frac{f(\theta) \sigma'(\theta'|\theta, \mu)}{f(\theta')} \bar{v}_{\sigma, A\Theta\Delta}^{T,1}(a, \theta', \mu), \tag{C.33}$$

where the second equality follows from [Equation C.32](#).

Second, by the same arguments as in [Appendix C.1](#), $\bar{v}_{\sigma'}^{T,2}(a, \theta, \theta', \omega, \mu)$ admits decomposition $\mu(\omega) \bar{v}_{\sigma', A\hat{\Theta}\Delta}^{T,2}(a, \theta, \theta', \mu)$ for each T .

Third, convergent subsequences $v_{\sigma'}^{T_{nm}, 1}$ and $v_{\sigma'}^{T_{nm}, 2}$ exist with limit $\bar{v}_{\sigma'}$ (cf. [Lemma C.2](#)).⁴⁵ We note two

⁴⁵Lemma C.2 implies that $\bar{v}_{\sigma'}^{T,1}$ and $\bar{v}_{\sigma'}^{T,2}$ have the same set of subsequential limits. Indeed, let g

things. On the one hand, because our previous arguments show that the set of measures admitting the above decompositions is closed, the limit $\bar{v}_{\sigma'}$ admits the decomposition. That is,

$$\int_{A \times \Theta \times \hat{\Theta} \times \Omega \times \Delta(\Omega)} u(a, \theta, \omega) \bar{v}_{\sigma'}(d(a, \theta, \theta', \omega, \mu)) = \int_{A \times \hat{\Theta} \times \Delta(\Omega)} \sum_{\theta \in \Theta} \frac{f(\theta) \sigma'(\theta' | \theta, \mu)}{f(\theta')} \left(\sum_{\omega \in \Omega} \mu(\omega) u(a, \theta, \omega) \right) d\bar{v}_{\sigma', A \hat{\Theta} \Delta}.$$

On the other hand, because T_{n_m} is a subsequence of T_n and $\bar{v}_{\sigma', A \hat{\Theta} \Delta}^{T,1} = \bar{v}_{\sigma, A \Theta \Omega \Delta}^{T,1}$ and $\bar{v}_{\sigma, A \Theta \Omega \Delta}^{T,1} \xrightarrow{w^*} \bar{v}_{\sigma}$, we can conclude that $\bar{v}_{\sigma', A \hat{\Theta} \Delta} = \bar{v}_{\sigma, A \Theta \Delta}$ and admits the same decomposition as \bar{v}_{σ} . We conclude that

$$\mathbb{E}_{\bar{v}_{\sigma'}}[u(a, \theta, \omega)] = \int_{\Delta(\Omega)} \left[\sum_{\theta, \theta' \in \Theta} f(\theta) \sigma'(\theta' | \theta, \mu) \sum_a \alpha(a | \theta', \mu) \left(\sum_{\omega \in \Omega} \mu(\omega) u(a, \theta, \omega) \right) \right] d\tau_{(p, \sigma)}.$$

Consequently,

$$\begin{aligned} \limsup_{T \rightarrow \infty} \mathbb{E}_{\sigma'}[U_T] &\geq \lim_{m \rightarrow \infty} \mathbb{E}_{\sigma'}[U_{T_{n_m}}] = \mathbb{E}_{\bar{v}_{\sigma'}}[u(a, \theta, \omega)] \\ &= \mathbb{E}_{\tau_{(p, \sigma)}} \left[\sum_{\theta, \theta', a} f(\theta) \sigma'(\theta, \mu) (\theta') \alpha(a | \theta', \mu) u(a, \theta, \mu) \right], \end{aligned}$$

where $u(a, \theta, \mu)$ is the linear extension of $u(a, \theta, \cdot)$. Because σ is a best response, we have that

$$\mathbb{E}_{\tau_{(p, \sigma)}} \left[\sum_{\theta, a} f(\theta) \alpha(a | \theta, \mu) u(a, \theta, \mu) \right] \geq \mathbb{E}_{\tau_{(p, \sigma)}} \left[\sum_{\theta, \theta', a} f(\theta) \sigma'(\theta, \mu) (\theta') \alpha(a | \theta', \mu) u(a, \theta, \mu) \right],$$

which implies the two-stage mechanism lacks profitable undetectable deviations.

The allocation rule is ex ante individually rational Define

$$U_{\text{net}}(\mu) = \sum_{\theta \in \Theta} f(\theta) \sum_{\omega \in \Omega} \mu(\omega) \left[\sum_{a \in A} \alpha(a | \theta, \mu) u(a, \theta, \omega) - u(a_{\emptyset}, \theta, \omega) \right],$$

to be the agent's (ex ante) payoff net of the outside option at belief μ . Ex ante individual rationality of α is equivalent to $U_{\text{net}}(\mu) \geq 0$ for all μ in the support of $\tau_{(p, \sigma)}$.

Toward a contradiction, assume that $U_{\text{net}}(\mu) < 0$ with positive probability under $\tau_{(p, \sigma)}$. By [Lemma D.1](#)

denote any continuous bounded function on $\Omega \times \Delta(\Omega) \times \Theta \times \Theta \times A$. Let $D_T(g) = \mathbb{E}_{\bar{v}_{\sigma'}^{T,2}}[g] - \mathbb{E}_{\bar{v}_{\sigma'}^{T,1}}[g] = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{\sigma'}[g(a_t, \theta_t, \theta'_t, \omega, \mu_{t+1}) - g(a_t, \theta_t, \theta'_t, \omega, \mu_t)]$. The argument in [Lemma C.2](#) implies that $D_T(g) \rightarrow 0$ as $T \rightarrow \infty$ (this does not rely on the existence of a limit, just the convergence of beliefs and the continuity of g). Now, let T_n be such that $\bar{v}_{\sigma'}^{T_n,1} \xrightarrow{w^*} \bar{v}$. Note that

$$\int g d\bar{v}_{\sigma'}^{T_n,2} = \int g d\bar{v}_{\sigma'}^{T_n,1} + D_{T_n}(g) \rightarrow \int g d\bar{v} + 0,$$

so a subsequential limit of $\bar{v}_{\sigma'}^{T,1}$ is a subsequential limit of $\bar{v}_{\sigma'}^{T,2}$. Switching the role of 1 and 2, we obtain the opposite set inclusion.

in Appendix D, a set $B \subset \Delta(\Omega)$ open relative to $\Delta(\Omega)$ exists such that⁴⁶

$$\int_B U_{\text{net}}(\mu) \tau_{(p,\sigma)}(d\mu) < 0.$$

Moreover, we can pick B such that $\tau_{(p,\sigma)}(\partial B) = 0$, where ∂B denotes the boundary of B relative to $\Delta(\Omega)$.⁴⁷ Lastly, let $\delta > 0$ be such that

$$\int_B U_{\text{net}}(\mu) \tau_{(p,\sigma)}(d\mu) \leq -2\delta. \quad (\text{C.34})$$

Let $(\mu_t(h^t))_{t \in \mathbb{N}, h^t \in H^t}$ denote the belief process under (p, σ) . For $L \in \mathbb{N}$, define a strategy (p^L, σ^L) as follows:

1. $(p_t^L(h^t, \cdot), \sigma_t^L(h^t, \cdot)) = (p_t(h^t, \cdot), \sigma_t(h^t, \cdot))$ if either $t < L$ OR $(t \geq L \text{ and } \mu_L(h^L) \notin B)$, where h^L precedes h^t ,
2. Otherwise, $(p_t^L(h^t, \cdot), \sigma_t^L(h^t, \cdot)) = (0, \sigma_t(h^t, \cdot))$ (note that when the agent quits the strategy can be specified arbitrarily.)

Note the agent's average payoff through period T under (p^L, σ^L) can be written as follows:

$$\mathbb{E}_{(p^L, \sigma^L)} [U_T] = \mathbb{E}_{(p, \sigma)} [U_T] - \mathbb{E}_{(p, \sigma)} \left[\frac{1}{T} \sum_{t=1}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[t \geq L \text{ and } \mu_L \in B] \right].$$

We show that for sufficiently large L , (p^L, σ^L) is a profitable deviation. For $T \geq L$, write

$$\begin{aligned} & \mathbb{E}_{(p, \sigma)} \left[\frac{1}{T} \sum_{t=L}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[\mu_L \in B] \right] - \int_B U_{\text{net}}(\mu) \tau_{(p, \sigma)}(d\mu) = \\ &= \mathbb{E}_{(p, \sigma)} \left[\frac{1}{T} \sum_{t=1}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[\mu_t \in B] \right] - \int_B U_{\text{net}}(\mu) \tau_{(p, \sigma)}(d\mu) \\ & - \mathbb{E}_{(p, \sigma)} \left[\frac{1}{T} \sum_{t=1}^{L-1} (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[\mu_t \in B] \right] \\ & + \mathbb{E}_{(p, \sigma)} \left[\frac{1}{T} \sum_{t=L}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) (\mathbb{1}[\mu_L \in B] - \mathbb{1}[\mu_t \in B]) \right]. \end{aligned} \quad (\text{C.35})$$

Let $K = \max_{\theta, \omega, a} |(u(a, \theta, \omega) - u(a_\phi, \theta, \omega))|$, and note that we can bound the term in the last line of Equation C.35 as follows:

$$\left| \mathbb{E}_{(p, \sigma)} \left[\frac{1}{T} \sum_{t=L}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) (\mathbb{1}[\mu_L \in B] - \mathbb{1}[\mu_t \in B]) \right] \right| \leq K \mathbb{E}_{(p, \sigma)} \left[\frac{1}{T} \sum_{t=L}^T |\mathbb{1}[\mu_L \in B] - \mathbb{1}[\mu_t \in B]| \right].$$

⁴⁶A set $X \subseteq \Delta(\Omega)$ is open relative to $\Delta(\Omega)$ if an open set $Y \subseteq \mathbb{R}^{|\Omega|}$ exists such that $X = Y \cap \Delta(\Omega)$. The boundary relative to $\Delta(\Omega)$ is analogously defined via open sets relative to $\Delta(\Omega)$.

⁴⁷Lemma D.1 provides an interval of radii $r \in (0, r_0)$ such that $\int_{B(\hat{\mu}, r)} U_{\text{net}}(\mu) \tau_{(p, \sigma)}(d\mu) < 0$. Note that only countable many such r can have $\tau_{(p, \sigma)}(\partial B(\hat{\mu}, r)) > 0$ (the boundaries for different radii are disjoint), so we can always pick r such that $\tau_{(p, \sigma)}(\partial B(\hat{\mu}, r)) = 0$ and preserve the negative sign.

Because $\mu_t \xrightarrow{w^*} \mu_\infty$ $\mathbb{P}_{(p,\sigma)}$ -a.s. (Lemma C.1) and $\tau_{(p,\sigma)}(\partial B) = 0$, we conclude:⁴⁸

$$\lim_{L \rightarrow \infty} \sup_{t \geq L} |\mathbb{1}[\mu_L \in B] - \mathbb{1}[\mu_t \in B]| = 0 \quad \mathbb{P}_{(p,\sigma)}\text{-a.s.}$$

Then,

$$\mathbb{E}_{(p,\sigma)} \left[\frac{1}{T} \sum_{t=L}^T |\mathbb{1}[\mu_L \in B] - \mathbb{1}[\mu_t \in B]| \right] \leq \mathbb{E}_{(p,\sigma)} \left[\sup_{t \geq L} |\mathbb{1}[\mu_L \in B] - \mathbb{1}[\mu_t \in B]| \right]$$

and choose \bar{L} large enough so that for all $L \geq \bar{L}$, we have that:

$$\mathbb{E}_{(p,\sigma)} \left[\sup_{t \geq L} |\mathbb{1}[\mu_L \in B] - \mathbb{1}[\mu_t \in B]| \right] \leq \delta/K. \quad (\text{C.36})$$

Consider now the term in the third line of Equation C.35 and note that it is bounded in absolute value by $K(L-1)/T$, which tends to 0 as $T \rightarrow \infty$. Similarly, the term in the second line of Equation C.35 vanishes as $T \rightarrow \infty$.⁴⁹ Thus, for $L \geq \bar{L}$, we can find \bar{T} such that for all $T \geq \bar{T}$ ⁵⁰

$$|\mathbb{E}_{(p,\sigma)} \left[\frac{1}{T} \sum_{t=L}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[\mu_L \in B] \right] - \int_B U_{\text{net}}(\mu) \tau_{(p,\sigma)}(d\mu)| \leq \frac{3}{2} \delta, \quad (\text{C.37})$$

and hence

$$\mathbb{E}_{(p,\sigma)} \left[\frac{1}{T} \sum_{t=L}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[\mu_L \in B] \right] \leq -\frac{1}{2} \delta. \quad (\text{C.38})$$

We conclude that

$$\limsup_{T \rightarrow \infty} \mathbb{E}_{(p^L, \sigma^L)}[U_T] \geq \lim_{T \rightarrow \infty} \mathbb{E}_{(p,\sigma)}[U_T] + \frac{1}{2} \delta,$$

a contradiction.

C.2.2 Proof of Theorem 4 (sufficiency)

We now show that all outcome distributions $\vartheta \in \Delta(A \times \Theta \times \Omega)$ that admit the decomposition in Theorem 4 can be implemented via a dynamic mechanism. To this end, let τ and $\alpha : \Theta \times \Delta(\Omega) \rightarrow \Delta(A)$

⁴⁸The property that $\tau_{(p,\sigma)}(\partial B) = 0$ ensures that $\mathbb{1}[\mu_t(h^\infty) \in B]$ is eventually constant almost surely. Let $E = \{h^\infty : \mu_\infty(h^\infty) \notin \partial B\}$. On E , either μ_∞ is in the interior of B (relative to $\Delta(\Omega)$) or in the interior of B^C (relative to $\Delta(\Omega)$), that is, an $\epsilon > 0$ exists such that $(B(\mu_\infty, \epsilon) \cap \Delta(\Omega)) \subset B$ or B^C . In either case, for each h^∞ , there exists $N(h^\infty)$ such that for all $t \geq N(h^\infty)$, $\mu_t(h^\infty) \in (B(\mu_\infty, \epsilon) \cap \Delta(\Omega))$ and hence $\mathbb{1}[\mu_t(h^\infty) \in B]$ is eventually constant. When $\tau_{(p,\sigma)}(\partial B) = 0$, we have that E has probability 1 under $\mathbb{P}_{(p,\sigma)}$.

⁴⁹Indeed, $\mathbb{E}_{(p,\sigma)} \left[\frac{1}{T} \sum_{t=1}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[\mu_t \in B] \right] = \mathbb{E}_{\bar{V}_{(p,\sigma)}^{T,1}} \left[(u(a, \theta, \omega) - u(a_\phi, \theta, \omega)) \mathbb{1}[\mu \in B] \right]$ and our previous analysis implies it converges to $\int_B U_{\text{net}}(\mu) \tau_{(p,\sigma)}(d\mu)$.

⁵⁰Choose \bar{T} so that for all $T \geq \bar{T}$:

$$\begin{aligned} & |\mathbb{E}_{(p,\sigma)} \left[\frac{1}{T} \sum_{t=1}^T (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[\mu_t \in B] \right] - \int_B U_{\text{net}}(\mu) \tau_{(p,\sigma)}(d\mu)| \\ & + |\mathbb{E} \left[-\frac{1}{T} \sum_{t=1}^{L-1} (u(a_t, \theta_t, \omega) - u(a_\phi, \theta_t, \omega)) \mathbb{1}[\mu_t \in B] \right] | \leq \delta/2. \end{aligned}$$

denote the belief distribution and the ex ante individually rational allocation rule without profitable undetectable deviations corresponding to ϑ . That is,

$$\vartheta(a, \theta, \omega) = \int_{\Delta(\Omega)} \mu(\omega) f(\theta) \alpha(a|\theta, \mu) \tau(d\mu). \quad (\text{C.39})$$

The proof proceeds as follows:

1. We first consider a fictitious setting in which there is no state uncertainty and we are given an allocation rule $\alpha' : \Theta \rightarrow \Delta(A)$ that is ex ante individually rational and lacks profitable undetectable deviations for some utility function $u' : A \times \Theta \rightarrow \mathbb{R}$. [Proposition C.1](#) shows that a dynamic mechanism exists that implements α' .
2. We then show that if ϑ satisfies [Equation C.39](#), then a finite support belief distribution τ' exists such that ϑ and α satisfies [Equation C.39](#) with τ' instead of τ .
3. Lastly, we use this result to construct a dynamic game that implements ϑ .

Step 1 For this step, we consider a fictitious setting in which there is no state uncertainty and the designer faces a privately informed agent with payoffs $u' : A \times \Theta \rightarrow \mathbb{R}$, where $\theta \sim f \in \Delta(\Theta)$.⁵¹

Suppose we are given an allocation rule $\alpha' : \Theta \rightarrow \Delta(A)$ that admits no profitable undetectable deviations relative to u' as in [Definition 7](#) and is individually rational as in [Definition 8](#). We have the following result:

Proposition C.1. *Let $\vartheta' = f(\theta) \alpha'(a|\theta) \in \Delta(A \times \Theta)$ such that α' lacks profitable undetectable deviations and is ex ante individually rational. Then, a dynamic mechanism exists that implements ϑ' .*

Proof of [Proposition C.1](#). The proof is constructive. We build on the analysis of Margaria and Smolin (2018) and present a dynamic mechanism that alternates between communication and adjustment phases. In all phases, the mechanism selects allocations using reports θ' according to α' . In a communication phase, the reports are those sent by the agent. In an adjustment phase, the agent's reports are disregarded; instead, the mechanism simulates reports to guarantee that the occupation measure over reports coincides with f and these simulated reports are used to determine the allocation. The mechanism ensures that under any agent's strategy, the occupation measure over reports and allocations exists and equals ϑ' ; thus, any strategy corresponds to an undetectable deviation. The length of communication phases grows in time. Thus, under truthtelling the relative length of adjustment phases vanishes in time, and the expected occupation measure over types and allocations exists and equals ϑ' . Because α' lacks profitable undetectable deviations, it follows that truthtelling is optimal for the agent. Because α' is ex ante individually rational, it follows that the participation constraints are satisfied.

Formally, the mechanism consists of sequential blocks, each block starting with a communication phase followed by an adjustment phase. The lengths of communication phases are fixed at L_1, L_2, \dots

⁵¹Anticipating our construction in [item 3](#), for each belief μ in the support of τ , the allocation rule $\alpha'(\cdot|\cdot, \mu)$ has no profitable undetectable deviations relative to payoff function $u'(a, \theta) = \sum_{\omega \in \Omega} \mu(\omega) u(a, \theta, \omega)$.

such that $L_n \rightarrow \infty$ and $L_n / \sum_{k \leq n} L_k \rightarrow 0$, e.g., $L_n = n$. The length of adjustment phase N_n depends on the agent's reports in the communication phase in block n . Denote by T_n the first period of block n , which is the first period of the corresponding communication phase. The first period of the corresponding adjustment phase is $T_n + L_n + 1$. Denote by freq_n^1 the average report frequencies in this block at the beginning of the adjustment stage:

$$\text{freq}_n^1(\hat{\theta}) \triangleq \frac{1}{L_n} \sum_{t=T_n}^{T_n+L_n-1} \mathbf{1}(\hat{\theta}_t = \hat{\theta}). \quad (\text{C.40})$$

If $\text{freq}_n^1 = f$, then the adjustment phase is empty, and the mechanism proceeds to the next block. Otherwise, in the adjustment phase, the mechanism generates reports over N_n periods to guarantee that at the end of the adjustment phase the *expected* frequency of reports in this block equals f that is,

$$\mathbb{E}[\text{freq}_n^2 | \text{freq}_n^1] = f \quad (\text{C.41})$$

where

$$\text{freq}_n^2(\hat{\theta}) \triangleq \frac{1}{L_n + N_n} \sum_{t=T_n}^{T_n+L_n+N_n-1} \mathbf{1}(\hat{\theta}_t = \hat{\theta}). \quad (\text{C.42})$$

To do so, denote by $\eta \triangleq \min_{\theta} f(\theta)$ and observe that $f \in \Delta(\Theta)$ can be surrounded by a ball of radius η within the simplex $\Delta(\Theta)$. The adjustment phase lasts for N_n periods where:⁵²

$$N_n = \left\lceil L_n \frac{\|\text{freq}_n^1 - f\|_{\infty}}{\eta} \right\rceil, \quad (\text{C.43})$$

and in each period of the adjustment phase the mechanism generates the reports i.i.d. according to \tilde{f}_n^a :

$$f_n^a = f - (\text{freq}_n^1 - f) \frac{L_n}{N_n}. \quad (\text{C.44})$$

The construction ensures that $f_n^a \in \Delta(\Theta)$, because $\|f_n^a - f\|_{\infty} \leq \eta$, and that (C.41) holds, because

$$\mathbb{E}[\text{freq}_n^2 | \text{freq}_n^1] = \frac{1}{L_n + N_n} (L_n \text{freq}_n^1 + N_n f_n^a) = f.$$

This in turn guarantees that the long-run distribution of reports (generated jointly by the agent and the mechanism) exists and equals f irrespectively of the agent's strategy. Intuitively, the fact that each block becomes negligible relative to past history over time ensures the agent's reports in each block have less and less effect on the long run frequency of reports, whereas the adjustment phase ensures that the frequency of reports converges to f . Formally, for any history and T denote by $n^{\text{last}}(T)$ the number of the block to which T belongs and by $T^{\text{last}}(T)$ the first period of that block. Observe that for

⁵²Any $\|\cdot\|_p$ would work, but larger p results in weakly shorter adjustment phases.

any agent's strategy:

$$N_n \leq L_n \left(\max_{f'} \frac{\|f' - f\|_\infty}{\eta} + 1 \right) \triangleq L_n \bar{\rho}. \quad (\text{C.45})$$

Therefore,

$$\frac{|T - T^{\text{last}}(T)|}{T^{\text{last}}(T)} \leq \frac{L_{n^{\text{last}}(T)}(1 + \bar{\rho})}{\sum_{k < n^{\text{last}}(T)} L_k} \xrightarrow[T \rightarrow \infty]{a.s.} 0, \quad (\text{C.46})$$

where the limit result holds because $n^{\text{last}}(T) \xrightarrow[T \rightarrow \infty]{a.s.} \infty$ and $L_n / \sum_{k \leq n} L_k \xrightarrow[n \rightarrow \infty]{} 0$.

Then, for any agent's strategy, for any $\hat{\theta} \in \Theta$,

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \Pr(\hat{\theta}_t = \hat{\theta}) &= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{f(\hat{\theta}) T^{\text{last}}(T) + f^{\text{last}}(T)(T - T^{\text{last}}(T))}{T^{\text{last}}(T) + T - T^{\text{last}}(T)} \right] \\ &= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{f(\hat{\theta}) + f^{\text{last}}(T)(T - T^{\text{last}}(T)) / T^{\text{last}}(T)}{1 + (T - T^{\text{last}}(T)) / T^{\text{last}}(T)} \right] = f(\hat{\theta}), \end{aligned} \quad (\text{C.47})$$

where $f^{\text{last}}(T) \in \Delta(\Theta)$ is the report frequency in the last block up to period T , and the last line follows from Equation C.46.

Since the mechanism chooses allocations in all periods according to α' , it follows that for any agent's strategy σ' the induced occupation measure over allocations and type reports satisfies:

$$\lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}_{\sigma'} \left[\sum_{t=1}^T \mathbb{1}[(a_t, \hat{\theta}_t) = (a, \hat{\theta})] \right] = f(\hat{\theta}) \alpha'(a | \hat{\theta}) = \vartheta'(a, \hat{\theta}). \quad (\text{C.48})$$

In other words, for any reporting strategy the occupation measure over allocations and reports exists.

We now show that under truthtelling the occupation measure over types and allocations exists and equals $f(\theta) \alpha'(a | \theta) = \vartheta(a, \theta)$. To this end, assume that the agent always reports her true type. For any T , denote by $\tilde{L}^{\text{total}}(T)$ the total number of periods spent in communication phases before T and by $\tilde{N}^{\text{total}}(T)$ the total number of periods spent in adjustment phases before T . Observe that by the strong law of large numbers, because $L_n \rightarrow \infty$,

$$\frac{N_n}{L_n} \leq \frac{\|\text{freq}_n^1 - f\|_\infty}{\eta} + \frac{1}{L_n} \xrightarrow[n \rightarrow \infty]{a.s.} 0.$$

Therefore,

$$\frac{\tilde{N}^{\text{total}}(T)}{\tilde{N}^{\text{total}}(T) + \tilde{L}^{\text{total}}(T)} \xrightarrow[T \rightarrow \infty]{a.s.} 0,$$

because whenever $N_n / L_n \rightarrow 0$, $\lim_{T \rightarrow \infty} N^{\text{total}}(T) / (N^{\text{total}}(T) + L^{\text{total}}(T)) = \lim_{n \rightarrow \infty} N_n / (L_n + N_n) = 0$.

It follows that

$$\begin{aligned}
& \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \Pr((\theta_t, \hat{\theta}_t, a_t) = (\theta, \hat{\theta}, a)) \\
&= \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{\tilde{L}^{\text{total}}(T) \mathbf{1}(\theta = \hat{\theta}) f(\hat{\theta}) \alpha'(a|\hat{\theta}) + \tilde{N}^{\text{total}}(T) f^{\text{adj}}(T)(\theta, \hat{\theta}, a)}{\tilde{N}^{\text{total}}(T) + \tilde{L}^{\text{total}}(T)} \right] \\
&= \mathbf{1}(\theta = \hat{\theta}) f(\hat{\theta}) \alpha'(a|\hat{\theta}),
\end{aligned}$$

where $f^{\text{adj}}(T) \in \Delta(\Theta \times \Theta \times A)$ is the average frequency of types, reports, and allocations in the adjustment phases before T . Therefore, under truthtelling, the occupation measure over allocations and types equals

$$\vartheta'(a, \theta) = f(\theta) \alpha'(a|\theta). \quad (\text{C.49})$$

Hence, the agent's payoff in the dynamic mechanism under truthtelling is:

$$U^{\text{truth}} = \sum_{(a, \theta)} f(\theta) \alpha'(a|\theta) u'(a, \theta). \quad (\text{C.50})$$

It remains to show that the agent cannot achieve more than U^{truth} under any other strategy. To this end, fix an alternative strategy σ , and denote by $U(\sigma) = \limsup_{T \rightarrow \infty} U_T(\sigma)$ where:

$$U_T(\sigma) = \frac{1}{T} \sum_{t=1}^T \sum_{a, \theta} \Pr((a_t, \theta_t) = (a, \theta)) u'(a, \theta).$$

Consider any convergent subsequence $(U_{T_n})_{n=1}^\infty$ along times $\{T_n\}_{n=1}^\infty$. Because $\Delta(A \times \Theta \times \Theta)$ is compact (Aliprantis and Border, 2006, Theorem 15.11), a convergent (sub)subsequence at times $\{T_k\}_{k=1}^\infty \subseteq \{T_n\}_{n=1}^\infty$ exists along which the occupation measure induced by σ

$$\nu_\sigma^{T_k} \xrightarrow{w^*} \nu_\sigma,$$

for some $\nu_\sigma \in \Delta(A \times \Theta \times \Theta)$, which by (C.48) satisfies $\nu_\sigma(a, \hat{\theta}) = f(\hat{\theta}) \alpha'(a|\hat{\theta})$. It follows that for some undetectable deviation $\nu_\sigma(\hat{\theta}|\theta)$:

$$\lim_{n \rightarrow \infty} U_{T_n} = \lim_{k \rightarrow \infty} U_{T_k} = \sum_{\theta, \hat{\theta}, a} f(\theta) \nu_\sigma(\hat{\theta}|\theta) \alpha'(a|\hat{\theta}) u'(a, \theta) \leq U^{\text{truth}},$$

where the inequality follows because $\alpha'(a|\hat{\theta})$ lacks profitable undetectable deviations. Because this inequality holds for any convergent subsequence $(U_{T_n})_{n=1}^\infty$,

$$U(\sigma) = \limsup_{T \rightarrow \infty} U_T(\sigma) \leq U^{\text{truth}}.$$

Finally, observe that the construction ensures that after every history, truthtelling from there on delivers the continuation payoff U^{truth} . Since α' is ex ante individually rational, $U^{\text{truth}} \geq \sum_\theta f(\theta) u'(a_\phi, \theta)$, and thus the participation constraints are satisfied. This concludes the proof. \square

Step 2 Consider now the outcome distribution $\vartheta \in \Delta(A \times \Theta \times \Omega)$ satisfying Equation C.39. As we argue in the proof of Theorem 2, a finite $K \leq |A||\Theta||\Omega|$, $\{\mu_1, \dots, \mu_K\} \in \Delta(\Omega)$, and $\tau' \in \Delta(\Delta(\Omega))$ exists such that

$$\vartheta(a, \theta, \omega) = f(\theta) \sum_{k=1}^K \tau'(\mu_k) \mu_k(\omega) \alpha(a|\theta, \mu_k).$$

Step 3 We now use steps 1 and 2 to complete the proof of Theorem 4, so in what follows we use the finite support representation of ϑ in the previous step. By Bayes plausibility, a dynamic mechanism can generate the belief split τ' in T periods with $T \leq \lceil \log_{|A|}(|\Omega||\Theta||A|) \rceil$, by treating each sequence of allocations of length T as a message. This can be achieved by making the mechanism constant on the agent's type reports during the first T periods. Since each $\alpha(\cdot|\cdot, \mu_k)$ for $k \in \{1, \dots, K\}$ lacks profitable undetectable deviations and is individually rational, Proposition C.1 implies that a dynamic mechanism exists that implements ϑ by first generating the belief split τ' and then implementing $\alpha(\cdot|\cdot, \mu_k)$ in the corresponding continuation play.

D Proof of auxiliary results

D.1 Revelation principle for calibrated mechanism design

In the main text, we restricted attention to incentive compatible and individually rational calibrated mechanisms. We show in this appendix that this restriction is without loss of generality by considering mechanisms with arbitrary message spaces and participation and reporting decisions by the agents that constitute an equilibrium of the game induced by the mechanism and its calibrated information structure.

Mechanisms Let $2^{[N]} \setminus \emptyset$ denote the nonempty subsets of agents. Then, we can define a mechanism as a collection $\{(M_J, \phi_J) : J \in 2^{[N]} \setminus \emptyset\}$, where

$$\phi_J : M_J \times \Omega \times [0, 1] \rightarrow \Delta(A_J),$$

is the mechanism when agents in J participate, where $M_J = \times_{i \in J} M_i$ and $A_J = \times_{i \in J} A_i$.

Information Structure Let $\hat{S}_i = \Delta(A_i)^{M_i}$ denote the collection of menus of lotteries with labels M_i , and let $\hat{S} = \times_{i \in [N]} \hat{S}_i$. An information structure is (π, \hat{S}) , where $\pi : \Omega \times [0, 1] \rightarrow \hat{S}$.

Participation and reporting strategies It is notationally convenient to allow each agent to have her own randomization device $\varepsilon_i \sim U[0, 1]$ and write agents' strategies as mappings $(p_i, \sigma_i) : \Theta_i \times \hat{S}_i \times [0, 1] \rightarrow \{0, 1\} \times M_i$, where p_i denotes agent i 's participation decision, and σ_i her reporting strategy, conditional on participating. To distinguish the agents' randomization from that of the original mechanism, we reserve ε_0 for the realization of the mechanism's randomization device.

Given $(p_i, \sigma_i)_{i \in [N]}$ and a mechanism (ϕ, M) , fix a profile $(\theta, \hat{s}, \bar{\varepsilon}) \equiv (\theta_i, \hat{s}_i, \varepsilon_i)_{i \in [N]}$. This determines a set of agents that participate,

$$J(\theta, \hat{s}, \bar{\varepsilon}) = \{j \in [N] : p_j(\theta_j, \hat{s}_j, \varepsilon_j) = 1\},$$

and let $J_{-i}(\theta, \hat{s}, \bar{\varepsilon})$ denote the projection of $J(\theta, \hat{s}, \bar{\varepsilon})$ on $J \setminus \{i\}$. Note that J_{-i} only depends on $(\theta_{-i}, \hat{s}_{-i}, \bar{\varepsilon}_{-i})$. Lastly, write $\phi_{J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i}) \cup \{i\}}(m_i, \sigma_{J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i})}, \omega, \varepsilon_0) \in \Delta(A_{J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i}) \cup \{i\}})$ for

$$\sum_{m_{J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i})}} \left(\prod_{j \in J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i})} \sigma_j(\theta_j, \hat{s}_j)(m_j) \right) \phi_{J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i}) \cup \{i\}}(m_i, m_{J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i})}, \omega, \varepsilon_0)$$

Calibrated information structures Given $(p_i, \sigma_i)_{i \in [N]}$ and a mechanism (ϕ, M) , the information structure (π, \hat{S}) is calibrated with the mechanism and the agents' strategies if whenever $\pi(\omega, \varepsilon_0) = (\hat{s}_1, \dots, \hat{s}_N)$, then for all i, m_i

$$\hat{s}_i(\cdot | m_i) = \mathbb{E}_{\tilde{\theta}_{-i} \sim f_{-i}(\cdot | \omega), \varepsilon_{-i}} \left[\sum_{a_{-i} \in A_{-i}} \phi_{J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i}) \cup \{i\}}(m_i, \sigma_{J_{-i}(\theta_{-i}, \hat{s}_{-i}, \varepsilon_{-i})}, \omega, \varepsilon_0)(\cdot, a_{-i}) \right].$$

Below, to keep the presentation simple, we focus on the case in which the calibrated information structure has finite support.

Equilibrium Given $(p_i, \sigma_i)_{i \in [N]}$, a mechanism (ϕ, M) and an information structure (π, \hat{S}) calibrated with the mechanism and the agents' strategies, $(p_i, \sigma_i)_{i \in [N]}$ is an equilibrium if for all $i \in [N]$, all $\theta_i \in \Theta_i$, all $\hat{s}_i \in \hat{S}_i$, and $\varepsilon_i \in [0, 1]$, the following hold:

$$\begin{aligned} \sigma_i(\theta_i, \hat{s}_i, \varepsilon_i) &\in \arg \max_{m_i \in M_i} \sum_{a_i \in A_i} \hat{s}_i(a_i | m_i) \mathbb{E}_{\omega \sim \mu_i(\cdot | \theta_i, \hat{s}_i)} [u_i(a_i, \theta_i, \omega)], \\ p_i(\theta_i, \hat{s}_i, \varepsilon_i) &\in \arg \max_{p \in \{0, 1\}} p \sum_{a_i \in A_i} \hat{s}_i(a_i | \sigma_i(\theta_i, \hat{s}_i, \varepsilon_i)) \mathbb{E}_{\omega \sim \mu_i(\cdot | \theta_i, \hat{s}_i)} [u_i(a_i, \theta_i, \omega)] + (1-p) \mathbb{E}_{\omega \sim \mu_i(\cdot | \theta_i, \hat{s}_i)} [u_i(a_{i\emptyset}, \theta_i, \omega)], \end{aligned}$$

where $\mu_i(\theta_i, \hat{s}_i) \in \Delta(\Omega)$ denotes agent i 's updated beliefs about the state when her type is θ_i conditional on receiving signal \hat{s}_i .

Revelation Principle Fix $(p_i, \sigma_i)_{i \in [N]}$, a mechanism (ϕ, M) and an information structure (π, \hat{S}) calibrated with the mechanism such that $(p_i, \sigma_i)_{i \in [N]}$ is an equilibrium. We construct a direct mechanism (ϕ^*, Θ) and a calibrated information structure (π^*, S^*) calibrated with the mechanism under truthtelling and full participation such that truthtelling and full participation is an equilibrium.

First, note that we can extend each $\phi_J(\cdot) \in \Delta(A_J)$ to a mechanism $\bar{\phi}_J(\cdot) \in \Delta(A)$ as follows: for all $m \in M_J, \omega \in \Omega, \varepsilon_0 \in [0, 1]$, and $a_J \in A_J$,

$$\bar{\phi}_J(m_J, \omega, \varepsilon_0)(a) = \phi_J(m_J, \omega, \varepsilon_0)(a_J) \times \delta_{a_{-J, \emptyset}}.$$

Define a “pseudo”-mechanism as follows:

$$\hat{\phi}_N(\theta, \omega, \varepsilon_0, \bar{\varepsilon}) = \bar{\phi}_{J(\theta, \pi(\omega, \varepsilon_0), \bar{\varepsilon})}(\sigma_{J(\theta, \hat{s}, \bar{\varepsilon})}, \omega, \varepsilon_0).$$

where $\sigma_{J(\theta, \hat{s}, \bar{\varepsilon})}$ is the message vector generated by the strategies. Define the full participation mechanism $\phi_N^* : \Theta \times \Omega \times [0, 1] \mapsto \Delta(A)$ to be

$$\phi_N^*(\theta, \omega, \varepsilon_0)(a) = \int_{[0, 1]^N} \hat{\phi}_N(\theta, \omega, \varepsilon_0, \bar{\varepsilon})(a) \lambda^N(d\bar{\varepsilon})$$

Let $S_i^* = \Delta(A_i)^{\Theta_i}$ and define $\pi^*(\omega, \varepsilon_0) = (s_1^*, \dots, s_N^*) \in \times_{i \in [N]} S_i^*$, where

$$s_i^*(\cdot | \hat{\theta}_i) = \mathbb{E}_{\theta_{-i} \sim f_{-i}(\cdot | \omega)} \left[\sum_{a_{-i}} \phi_N^*(\hat{\theta}_i, \theta_{-i}, \omega, \varepsilon_0)(\cdot, a_{-i}) \right].$$

By definition, the information structure is calibrated relative to full participation and truthful reporting.

We now show that full participation and truthful reporting is a best response to others participating and truthfully reporting into the mechanism. To this end, consider agent i 's payoff from submitting report θ'_i when observing s_i^* . Denoting by $\Sigma(\omega, s_i^*)$ the set of ε_0 such that $\pi_i^* = s_i^*$, this payoff is given by.⁵³

$$\begin{aligned} & \sum_{\omega \in \Omega} \frac{\mu_0(\omega) f_i(\theta_i | \omega)}{Pr(s_i^* | \theta_i)} \sum_{\theta_{-i}} f_{-i}(\theta_{-i} | \omega) \int_{\Sigma(\omega, s_i^*)} \sum_a \phi_N^*(\theta'_i, \theta_{-i}, \omega, \varepsilon_0)(a_i, a_{-i}) \lambda(d\varepsilon_0) u_i(a_i, \theta_i, \omega) = \\ & \sum_{a_i \in A_i} \sum_{\omega \in \Omega} \frac{\mu_0(\omega) f_i(\theta_i | \omega)}{Pr(s_i^* | \theta_i)} u_i(a_i, \theta_i, \omega) \sum_{\theta_{-i}} f_{-i}(\theta_{-i} | \omega) \int_{\Sigma(\omega, s_i^*)} \sum_{a_{-i}} \phi_N^*(\theta'_i, \theta_{-i}, \omega, \varepsilon_0)(a_i, a_{-i}) \lambda(d\varepsilon_0) \\ & = \int_0^1 \left[\sum_{a_i} \sum_{\omega} \frac{\mu_0(\omega) f_i(\theta_i | \omega)}{Pr(s_i^* | \theta_i)} u_i(a_i, \theta_i, \omega) \int_{\Sigma(\omega, s_i^*)} (\star) \lambda(d\varepsilon_0) \right] \lambda(d\varepsilon_i), \end{aligned}$$

where

$$\begin{aligned} \star &= \mathbb{E}_{\theta_{-i} | \omega, \bar{\varepsilon}_{-i}} \left[\sum_{a_{-i}} \hat{\phi}_N(\theta'_i, \theta_{-i}, \varepsilon_0, \bar{\varepsilon}_{-i})(a_i, a_{-i}) \right] \\ &= \mathbb{E}_{\theta_{-i} | \omega, \bar{\varepsilon}_{-i}} \left[\sum_{a_{-i}} \bar{\phi}_{J(\theta'_i, \theta_{-i}, \hat{s}(\omega, \varepsilon_0), \bar{\varepsilon})}(\sigma_{J(\theta'_i, \theta_{-i}, \hat{s}(\omega, \varepsilon_0), \bar{\varepsilon})}, \omega, \varepsilon_0)(a_i, a_{-i}) \right] \\ &= \mathbb{1}[p_i(\theta'_i, \hat{s}_i(\omega, \varepsilon_0), \varepsilon_i) = 1] \hat{s}_i(a_i | \sigma_i(\theta'_i, \hat{s}_i, \varepsilon_i)) + (1 - \mathbb{1}[p_i(\theta'_i, \hat{s}_i(\omega, \varepsilon_0), \varepsilon_i) = 1]) \delta_{a_{i,\phi}}(a_i). \end{aligned}$$

Because agent i of type θ_i could have imitated type θ'_i , reporting θ_i dominates. By the same logic, when the agent reports θ_i , she obtains at least the payoff from participating in the mechanism.

D.2 Optimal Calibrated Auction

Proof of Proposition 5. The pointwise solution to the Myersonian problem allocates the good to agents in $N^*(\theta, \omega) = \arg\max_{i \in [N] \cup \{0\}} [w_i(\theta, \omega) + J_i(\theta_i) \omega_i + \omega_{0i}]$ where $i = 0$ corresponds to an outside option with $w_0 \equiv J_0 \equiv \omega_{00} \equiv 0$. The conditions of the proposition ensure that for all $i \in N$ and $j \in [N] \cup \{0\}$,

$$\frac{d}{d\theta_i} (w_i(\theta, \omega) + J_i(\theta_i, F_i) \omega_i + \omega_{0i}) \geq \frac{d}{d\theta_i} (w_j(\theta, \omega) + J_j(\theta_j, F_j) \omega_j + \omega_{0j}).$$

Thus, an optimal selection $q^*(\theta, \omega)$ exists such that for each i , θ_{-i} , and ω , $q^*(\theta_i, \theta_{-i}, \omega)$ is non-decreasing in θ_i (e.g., one that uniformly randomizes over $N^*(\theta, \omega)$).

Denote by Q_{full} the set of allocation rules implementable under full state disclosure. These are the rules such that for all i and ω , $\mathbb{E}_{F_{-i}}[q_i(\theta_i, \theta_{-i}, \omega)]$ is non-decreasing in θ_i . It follows that $q^* \in Q_{\text{full}}$, and hence $q^* \in Q_{\text{My}}$. Thus, q^* solves the Myersonian problem and can also be implemented by fully disclosing the state to the agents and conducting an optimal mechanism state-by-state. By revenue

⁵³In the expressions that follow, recall the full participation mechanism ϕ_N^* already averages over the agents' own randomization devices.

equivalence, the expected revenue of such implementation is the same as under no disclosure, and thus the designer obtains payoff W_{My} . \square

D.3 Technical results from Appendix C

Proof of Lemma C.1. The set of continuous bounded functions on $\tilde{\Omega}$ is separable and hence it has a countable dense subset $\{g_k\}_{k \in \mathbb{N}} \subset C_b(\tilde{\Omega})$. It is immediate to see that $\mu_n \xrightarrow{w^*} \mu$ if and only if for all $k \in \mathbb{N}$ $\int g_k d\mu_n \rightarrow \int g_k d\mu$.

For each $k \in \mathbb{N}$ define a real-valued, bounded, martingale on $(\mathcal{H}^\infty, \mathcal{B}_{\mathcal{H}^\infty}, \mathbb{P}_\sigma)$ as follows:

$$M_t^k(\tilde{\omega}, h^\infty) = \int_{\tilde{\Omega}} g_k(\omega', \varepsilon) d\mu_t(\tilde{\omega}, h^\infty)(\omega', \varepsilon).$$

Doob's martingale convergence theorem implies that $M_t^k(\tilde{\omega}, h^\infty) = \mathbb{E}[g_k | h^t] \rightarrow M_\infty^k(\tilde{\omega}, h^\infty) = \mathbb{E}[g_k | h^\infty]$ \mathbb{P}_σ -a.s. Let E_k denote the subset of \mathcal{H}^∞ where convergence happens, and note that $\mathbb{P}_\sigma(E_k) = 1$.

Let $E = \cap_k E_k$ and note that $\mathbb{P}_\sigma(E) = 1$. Then, on E , we have that for all $k \in \mathbb{N}$,

$$\int_{\tilde{\Omega}} g_k(\omega', \varepsilon) d\mu_t(\tilde{\omega}, h^\infty)(\omega', \varepsilon) \rightarrow M_\infty^k(\tilde{\omega}, h^\infty). \quad (\text{D.1})$$

Fix now a terminal history $(\tilde{\omega}, h^\infty)$. Because $\Delta(\tilde{\Omega})$ is compact (Aliprantis and Border, 2006, Theorem 15.11), the sequence $(\mu_t(\tilde{\omega}, h^\infty))_{t \in \mathbb{N}}$ has a convergent subsequence $\mu_{t_j}(\tilde{\omega}, h^\infty) \xrightarrow{w^*} \tilde{\mu}$. Passing the limit along t_j in Equation D.1 we have that for all $k \in \mathbb{N}$

$$\int_{\tilde{\Omega}} g_k(\omega', \varepsilon) d\mu_{t_j}(\tilde{\omega}, h^\infty)(\omega', \varepsilon) \rightarrow \int_{\tilde{\Omega}} g_k(\omega', \varepsilon) d\tilde{\mu}.$$

Because the set $\{g_k : k \in \mathbb{N}\}$ determines the convergent subsequences, any subsequential limits must be equal. Hence $\mu_t(\cdot)$ converges on E and call this limit $\tilde{\mu}_\infty$. Hence on E we have that $\mu_t(h^\infty) \xrightarrow{w^*} \tilde{\mu}_\infty(h^\infty)$.

Now, for each k , $\int g_k d\mu_\infty = \mathbb{E}[g_k | h^\infty]$ almost surely. Hence, $\tilde{\mu}_\infty$ is a version of the law of $\tilde{\Omega}$ conditional on h^∞ . This is $\mathbb{P}_\sigma(\cdot | h^\infty)$, completing the proof. \square

Proof of Lemma C.2. Fix a continuous and bounded function $g \in C_b(A \times \Theta \times \Omega \times \Delta(\Omega))$ and define for each terminal history (ω, h^∞)

$$\Delta_t(\omega, h^\infty) = g(a_t(h^\infty), \theta_t(h^\infty), \omega, \mu_t(h^\infty)) - g(a_t(h^\infty), \theta_t(h^\infty), \omega, \mu_{t+1}(h^\infty)).$$

As in Lemma C.1, let E denote the probability-1 subset of \mathcal{H}^∞ on which $\mu_t \xrightarrow{w^*} \mu_\infty$.⁵⁴ Then, on E , $\Delta_t(\omega, h^\infty) \rightarrow 0$, and hence,

$$\mathbb{E}_\sigma [\Delta_t(\omega, h^\infty)] \rightarrow 0,$$

⁵⁴To be sure, the proof of Lemma C.1 is written in the context of repeated mechanisms but it extends verbatim to dynamic mechanisms with simple notational adjustments.

as $t \rightarrow \infty$. Now, for every T ,

$$D_T(g) \equiv \mathbb{E}_{\bar{v}_\sigma^{T,1}}[g] - \mathbb{E}_{\bar{v}_\sigma^{T,2}}[g] = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_\sigma[\Delta_t],$$

and hence the left-hand side goes to 0 as $T \rightarrow \infty$.

Now, let T_n be such that $\bar{v}_\sigma^{T_n,2} \xrightarrow{w^*} \bar{v}$ for some $\bar{v} \in \Delta(A \times \Theta \times \Omega \times \Delta(\Omega))$. Note that

$$\mathbb{E}_{\bar{v}_\sigma^{T_n,1}}[g] = \mathbb{E}_{\bar{v}_\sigma^{T_n,2}}[g] + D_{T_n}[g] \rightarrow \mathbb{E}_{\bar{v}}[g] + 0,$$

so a subsequential limit of $\bar{v}_\sigma^{T_n,2}$ is a subsequential limit of $\bar{v}_\sigma^{T_n,1}$. Switching the role of 1 and 2, we obtain the opposite set inclusion and the result follows. \square

Proof of Lemma C.3. Fix a continuous function g on $\Delta(\tilde{\Omega})$. Then, we want to show that

$$\mathbb{E}_{\tau_T}[g] \rightarrow \mathbb{E}_{\mathbb{P}_\sigma \circ \mu_\infty^{-1}}[g].$$

By Lemma C.1, $\mu_t \xrightarrow{w^*} \mu_\infty$ \mathbb{P}_σ -almost surely and g is continuous, we have that $\mathbb{E}_\sigma[g(\mu_t)] \rightarrow \mathbb{E}_\sigma[g(\mu_\infty)]$ by dominated convergence theorem.⁵⁵ Because eventually constant sequences have Cesàro limits, we have that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}_\sigma[g(\mu_t)] \rightarrow \mathbb{E}_\sigma[g(\mu_\infty)].$$

And now we are basically done, because

$$\mathbb{E}_{\tau_T}[g] = \int_{H^\infty} \frac{1}{T} \sum_{t=1}^T g(\mu_t(h^\infty)) \mathbb{P}_\sigma(dh^\infty) = \frac{1}{T} \sum_{t=1}^T \mathbb{E}_\sigma[g(\mu_t)] \rightarrow \mathbb{E}_\sigma[g(\mu_\infty)] = \int g d(\mathbb{P}_\sigma \circ \mu_\infty^{-1}).$$

In other words, the occupation measure on beliefs induced by the strategy (and the prior, the type distribution, and the mechanism) is the push-forward measure $(\mathbb{P}_\sigma \circ \mu_\infty^{-1})$. In particular, that \mathbb{P}_σ is a measure implies that $(\mathbb{P}_\sigma \circ \mu_\infty^{-1})$ is a measure itself (Bogachev, 2007, Chapter 3.6). \square

Proof of Lemma C.4. We now show the agent can ensure the payoff $U(\tau_\phi)$ in Equation C.21, which corresponds to the agent's maximum payoff under the calibrated information structure π_ϕ . Recall that this information structure is the one that corresponds to the partition of $\tilde{\Omega}$, \mathcal{P} , defined as follows: $\tilde{\omega}, \tilde{\omega}'$ in the same cell P of \mathcal{P} if for all (a, m) , $\phi(a|m, \tilde{\omega}) = \phi(a|m, \tilde{\omega}')$. Conditional on cell P , the associated posterior is $\mu(|P) \in \Delta(\tilde{\Omega})$. Let $\tau_\phi \in \Delta(\Delta(\tilde{\Omega}))$ denote the induced belief distribution (with mean $\mu_0 \otimes \eta$).

To prove the result, we consider the strategy σ'_N parameterized by a number N and defined as follows. In the exploration phase of σ'_N , the agent plays each message m for N rounds. Let $\mu_{N|M}(h^{N|M})$ denote the agent's beliefs as a function of the realized sequence of allocations implied by $h^{N|M}$. For each

⁵⁵To be sure,

$$\mathbb{E}_\sigma[g(\mu_t)] = \int_{H^\infty} g(\mu_t(h^\infty)) \mathbb{P}_\sigma(dh^\infty).$$

history that succeeds $h^{N|M|}$, the agent of type θ plays the message m that solves

$$\max_{m \in M} \sum_{\tilde{\omega}} \mu_{N|M|}(h^{N|M|})(\tilde{\omega}) \sum_{a \in A} \phi(a|m, \tilde{\omega}) u(a, \theta, \omega).$$

It is immediate to verify that the agent's (limit) average payoff under σ'_N is given by:

$$\begin{aligned} U(\sigma'_N) &= \mathbb{E}_{\sigma'_N} \left[\sum_{\theta \in \Theta} f(\theta) \max_{m \in M} \sum_{\tilde{\omega}} \mu_{N|M|}(h^{N|M|})(\tilde{\omega}) \sum_{a \in A} \phi(a|m, \tilde{\omega}) u(a, \theta, \omega) \right] = \\ &= \sum_{h^{N|M|} \in H^{N|M|}} \mathbb{P}_{\sigma'_N}^{N|M|}(h^{N|M|}) u^*(\mu_{N|M|}(h^{N|M|})), \end{aligned} \quad (\text{D.2})$$

where u^* is as in [Equation C.20](#).

Below, we show that the distribution of beliefs under σ'_N , $\mathbb{P}_{\sigma'_N} \circ \mu_{N|M|}^{-1} \xrightarrow{w^*} \tau_\phi$. Consequently, as u^* is continuous and bounded on $\Delta(\tilde{\Omega})$, for any $\delta > 0$, we can choose N_δ so that for all $N \geq N_\delta$,

$$|U(\sigma'_N) - U(\tau_\phi)| \leq \delta.$$

Consequently, the agent's payoff under σ must be $U(\tau_\phi)$ because by definition for all $\delta > 0$ ⁵⁶

$$\liminf_{T \rightarrow \infty} \mathbb{E}_\sigma [U_T] \geq \limsup_{T \rightarrow \infty} \mathbb{E}_{\sigma'_{N_\delta}} [U_T] = U(\sigma'_{N_\delta}) \geq U(\tau_\phi) - \delta,$$

and hence,

$$\liminf_{T \rightarrow \infty} \mathbb{E}_\sigma [U_T] \geq \lim_{\delta \rightarrow 0} U(\tau_\phi) - \delta = U(\tau_\phi).$$

which completes the proof.

We now complete the missing step:

The law of $\mu_{N|M|}$ converges to τ_ϕ It is useful to write the bottom line of [Equation D.2](#) as follows:

$$\sum_{\tilde{\omega} \in \tilde{\Omega}} (\mu_0 \otimes \eta)(\tilde{\omega}) \mathbb{E}_{\mathbb{P}_{\sigma'_N}^{N|M|}(\cdot|\tilde{\omega})} [u^*(\mu_{N|M|})].$$

We show that the conditional law of $\mu_{N|M|}$, $\mathbb{P}_{\sigma'_N}(\cdot|\tilde{\omega})$, converges to the Dirac measure on $\mu(\cdot|P(\tilde{\omega}))$. Noting that $\tau_\phi = \sum_{\tilde{\omega} \in \tilde{\Omega}} (\mu_0 \otimes \eta)(\tilde{\omega}) \delta_{\mu(\cdot|P(\tilde{\omega}))}$ completes the proof.

For the exploration block, define for each $m \in M$, the empirical frequency $\hat{\phi}_{N,m} : (\Theta \times M \times A)^{N|M|} \rightarrow \Delta(A)$, as follows

$$\hat{\phi}_{N,m}(h^{N|M|})(a) = \frac{1}{N} \sum_{n=1}^N \mathbb{1}[A_{m,n} = a],$$

where $A_{m,n}$ is the n^{th} draw from A when the message is m . Let $\hat{\phi}_N : H^{N|M|} \rightarrow \Delta(A)^M$ denote the vector

⁵⁶The argument shows that the agent's equilibrium payoff is at least $U(\tau_\phi)$. However, it is immediate that $U(\tau_\phi)$ is the most the agent can make in the game as τ_ϕ extracts all information from the mechanism.

of empirical frequencies. The agent's belief at history $h^{N|M|}$ is given by:

$$\mu_{N|M}(h^{N|M|})(\tilde{\omega}) = \frac{(\mu_0 \otimes \eta)(\tilde{\omega}) \prod_{m \in M} \prod_{a \in A} \phi(a|m, \tilde{\omega})^{N\hat{\phi}_{N,m}(h^{N|M|})(a)}}{\sum_{\tilde{\omega}'} (\mu_0 \otimes \eta)(\tilde{\omega}') \prod_{m \in M} \prod_{a \in A} \phi(a|m, \tilde{\omega}')^{N\hat{\phi}_{N,m}(h^{N|M|})(a)}}. \quad (\text{D.3})$$

Denote by $\tau_{N|M}, \tilde{\omega} \in \Delta(\Delta(\tilde{\Omega}))$ the law of $\mu_{N|M}$ conditional on $\tilde{\omega}$, i.e., $\tau_{N|M}, \tilde{\omega} = \mathbb{P}_{\sigma'_N}(\cdot|\tilde{\omega}) \circ \mu_{N|M}^{-1}$. Below, we show that $\tau_{N|M}, \tilde{\omega}$ converges weakly to $\delta_{\mu(\cdot|P(\tilde{\omega}))}$.

Suppose the true state is $\tilde{\omega}^*$. Then, $(A_{m,1}, \dots, A_{m,N})$ are drawn i.i.d. from distribution $\phi(\cdot|m, \tilde{\omega}^*)$. Fix a continuous and bounded function g on A . Then, almost surely⁵⁷

$$\int_A g d\hat{\phi}_N(\cdot|m) = \frac{1}{N} \sum_{n=1}^N g(A_{m,n}) \rightarrow \mathbb{E}_\phi[g(A_{m,1})] = \sum_{a \in A} g(a) \phi(a|m, \tilde{\omega}^*),$$

by the strong law of large numbers applied to the i.i.d random variables $(g(A_{m,1}), \dots, g(A_{m,N}))$. Because this holds for all g , then $\hat{\phi}_N(\cdot|m) \xrightarrow{w^*} \phi(\cdot|m, \tilde{\omega}^*)$ almost surely when the true state is $\tilde{\omega}^*$.

Fix an arbitrary state $\tilde{\omega}$ and consider the ratio of the right-hand side of Equation D.3 at $\tilde{\omega}$ and $\tilde{\omega}^*$:

$$\frac{\mu_{N|M}(h^{N|M|})(\tilde{\omega})}{\mu_{N|M}(h^{N|M|})(\tilde{\omega}^*)} = \frac{(\mu_0 \otimes \eta)(\tilde{\omega})}{(\mu_0 \otimes \eta)(\tilde{\omega}^*)} \prod_{a \in A} \prod_{m \in M} \left(\frac{\phi(a|m, \tilde{\omega})}{\phi(a|m, \tilde{\omega}^*)} \right)^{N\hat{\phi}_{N,m}(h^{N|M|})(a)}. \quad (\text{D.4})$$

Suppose $\tilde{\omega} \notin P(\tilde{\omega}^*)$. By definition of the partition \mathcal{P} , a message $m \in M$ and allocation $a \in A$ exist such that $\phi(a|m, \tilde{\omega}) \neq \phi(a|m, \tilde{\omega}^*)$. Taking logarithm on both sides of Equation D.4 and dividing by N ,

$$\frac{1}{N} \log \left(\frac{\mu_{N|M}(h^{N|M|})(\tilde{\omega})}{\mu_{N|M}(h^{N|M|})(\tilde{\omega}^*)} \right) = \frac{1}{N} \log \left(\frac{(\mu_0 \otimes \eta)(\tilde{\omega})}{(\mu_0 \otimes \eta)(\tilde{\omega}^*)} \right) + \sum_{a' \in A} \sum_{m' \in M} \hat{\phi}_{N,m'}(h^{N|M|})(a') \log \left(\frac{\phi(a'|m', \tilde{\omega})}{\phi(a'|m', \tilde{\omega}^*)} \right). \quad (\text{D.5})$$

Because $\hat{\phi}_N(\cdot|m) \xrightarrow{w^*} \phi(\cdot|m, \tilde{\omega}^*)$ almost surely when the true state is $\tilde{\omega}^*$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \left(\frac{\mu_{N|M}(h^{N|M|})(\tilde{\omega})}{\mu_{N|M}(h^{N|M|})(\tilde{\omega}^*)} \right) = - \sum_{m \in M} D_{\text{KL}}(\phi(\cdot|m, \tilde{\omega}^*) || \phi(\cdot|m, \tilde{\omega})), \quad (\text{D.6})$$

where D_{KL} is the Kullback-Leibler divergence. Note that at least one of the terms in the KL-divergence is positive as $\phi(\cdot|m, \tilde{\omega}) \neq \phi(\cdot|m, \tilde{\omega}^*)$. Hence,

$$\lim_{N \rightarrow \infty} \log \left(\frac{\mu_{N|M}(h^{N|M|})(\tilde{\omega})}{\mu_{N|M}(h^{N|M|})(\tilde{\omega}^*)} \right) = -\infty,$$

meaning that $\mu_{N|M}(h^{N|M|})(\tilde{\omega}) / \mu_{N|M}(h^{N|M|})(\tilde{\omega}^*) \rightarrow 0$.

Suppose now that $\tilde{\omega} \in P(\tilde{\omega}^*)$. Then, Equation D.4 reduces to

$$\frac{\mu_{N|M}(h^{N|M|})(\tilde{\omega})}{\mu_{N|M}(h^{N|M|})(\tilde{\omega}^*)} = \frac{(\mu_0 \otimes \eta)(\tilde{\omega})}{(\mu_0 \otimes \eta)(\tilde{\omega}^*)}, \quad (\text{D.7})$$

for all N .

⁵⁷This almost surely is under the law of A under $\phi(\cdot|m, \tilde{\omega}^*)$.

Collecting both cases, we conclude that conditional on the true state being $\tilde{\omega}^*$,

$$\sum_{\tilde{\omega} \in P(\tilde{\omega}^*)} \mu_{N|M}(\tilde{\omega}) \rightarrow_{N \rightarrow \infty} 1,$$

and moreover, within the cell, the fixed-ratio property implies the law $\tau_{N|M, \tilde{\omega}^*} \xrightarrow{w^*} \delta_{\mu(\cdot|P(\tilde{\omega}^*))}$. We conclude that the unconditional belief distribution, $\sum_{\tilde{\omega} \in \bar{\Omega}} (\mu_0 \otimes \eta)(\tilde{\omega}) \tau_{N|M, \tilde{\omega}} \xrightarrow{w^*} \sum_{\tilde{\omega} \in \bar{\Omega}} (\mu_0 \otimes \eta)(\tilde{\omega}) \delta_{\mu(P(\tilde{\omega}))} = \tau_\phi$. In particular,

$$U(\sigma'_N) = \mathbb{E}_{\mathbb{P}_{\sigma'_N} \circ \mu_{N|M}^{-1}} [u^*(\mu)] \rightarrow \mathbb{E}_{\tau_\phi} [u^*(\mu)],$$

completing the proof. \square

Lemma D.1. *Suppose $u : \Delta(\Omega) \rightarrow \mathbb{R}$ satisfies that*

$$\int_{\Delta(\Omega)} u(\mu) \tau(d\mu) < \int_{\Delta(\Omega)} \max\{u(\mu), 0\} \tau(d\mu),$$

then a set $B \subseteq \Delta(\Omega)$ open relative to $\Delta(\Omega)$ exists such that

$$\int_B u(\mu) \tau(d\mu) < 0.$$

Proof. Define the positive and negative parts of u :

$$u_+(\mu) := \max\{u(\mu), 0\}, \quad u_-(\mu) := \max\{-u(\mu), 0\}.$$

Then $u = u_+ - u_-$ pointwise. Integrating and using the assumed strict inequality,

$$\int u d\tau = \int u_+ d\tau - \int u_- d\tau < \int u_+ d\tau \implies \int u_- d\tau > 0.$$

Hence the set $N := \{\mu \in \Delta(\Omega) : u(\mu) < 0\}$ has strictly positive mass under τ .

Embed $\Delta(\Omega) \subset \mathbb{R}^{|\Omega|}$. Extend τ to a finite Borel measure $\tilde{\tau}$ on \mathbb{R}^d by

$$\tilde{\tau}(A) := \tau(A \cap \Delta(\Omega)) \quad (A \subseteq \mathbb{R}^d \text{ Borel}),$$

and extend u to $\tilde{u} : \mathbb{R}^d \rightarrow [-1, 1]$ by $\tilde{u} = u$ on $\Delta(\Omega)$ and $\tilde{u} = 0$ on $\mathbb{R}^d \setminus \Delta(\Omega)$. Then $\tilde{u} \in L^1(\tilde{\tau})$ and $\tilde{\tau}(N) = \tau(N) > 0$.

Let $B_{\mathbb{R}^d}(x, r)$ denote the ball in $\mathbb{R}^{|\Omega|}$ with center x and radius r . By the Lebesgue differentiation theorem for finite Borel measures on \mathbb{R}^d , there is a $\tilde{\tau}$ -full-measure set $D \subseteq \mathbb{R}^d$ such that for every $x \in D$,

$$\lim_{r \downarrow 0} \frac{1}{\tilde{\tau}(B_{\mathbb{R}^d}(x, r))} \int_{B_{\mathbb{R}^d}(x, r)} \tilde{u} d\tilde{\tau} = \tilde{u}(x),$$

whenever $\tilde{\tau}(B_{\mathbb{R}^d}(x, r)) > 0$ (and this positivity holds for all sufficiently small r for $\tilde{\tau}$ -a.e. x). Since $\tilde{\tau}(N \cap D) > 0$, choose $\mu_0 \in N \cap D$. Then $\tilde{u}(\mu_0) = u(\mu_0) < 0$. Therefore the above limit is strictly negative,

so there exists $r_0 > 0$ such that for all $0 < r < r_0$,

$$\int_{B_{\mathbb{R}^d}(\mu_0, r)} \tilde{u} d\tilde{\tau} < 0.$$

For such an r , let $B := B_\Delta(\mu_0, r) = \Delta(\Omega) \cap B_{\mathbb{R}^d}(\mu_0, r)$, which is an open ball in $\Delta(\Omega)$. Using the definitions of $\tilde{\tau}$ and \tilde{u} ,

$$\int_B u d\tau = \int_{B_{\mathbb{R}^d}(\mu_0, r)} \tilde{u} d\tilde{\tau} < 0.$$

This proves the claim. \square

D.3.1 Revelation principle for limit of means preferences

We show in this section that when the designer uses dynamic mechanisms, it is without loss of generality for the designer to employ direct dynamic mechanisms that (i) implement the outside option at all histories after the agent first exercises her option not to participate in the mechanism, and (ii) for which the agent's best response is to always participate and truthfully report her type. This justifies the class of mechanisms we employ in the analysis of [Section 5.2](#).

Histories, mechanisms, and strategies As in the main text, to simplify notation, we do not include the agent's decision to participate in the mechanism in the histories of the game. Instead, we follow the convention that if the agent does not participate, it is as if she reported \emptyset and the allocation is a_\emptyset . Formally, let $MA_\emptyset = (M \times A) \cup \{(\emptyset, a_\emptyset)\}$. With this notation, a history through period t is an element of $\hat{H}_M^t \equiv (MA_\emptyset)^{t-1}$ and let $\hat{\mathcal{H}}_M^t = \Omega \times \hat{H}_M^t$.⁵⁸

A mechanism is a collection $\varphi \equiv (\varphi_t)_{t=1}^\infty$ such that the mechanism in period t is a mapping $\varphi_t : \hat{\mathcal{H}}_M^t \times M \rightarrow \Delta(A)$.

Let $H_M^t = (\Theta \times MA_\emptyset)^{t-1} = \Theta^{t-1} \times \hat{H}_M^t$. The agent's strategy, (p, σ) , is given by her participation strategy $p_t : H_M^t \times \Theta \rightarrow [0, 1]$, and conditional on participating, her reporting strategy $\sigma_t : H_M^t \times \Theta \rightarrow \Delta(M)$.

The distribution over terminal histories To obtain the complete description of the paths on the tree we need to append Ω to H_M^t ; hence the paths through period $t-1$ are $\Omega \times H_M^t \equiv \mathcal{H}_M^t$. The distributions over states and agent's types, the agent's strategy, and the mechanism induce a distribution over the terminal histories $\mathcal{H}_M^\infty \equiv \Omega \times H_M^\infty$, which we denote by $\mathbb{P}_{\varphi, (p, \sigma)} \in \Delta(\Omega \times H_M^\infty)$, as it is now useful to keep track of the mechanism. We denote by $\mathbb{E}_{(p, \sigma)}$ the expectation under this measure. The distribution $\mathbb{P}_{\varphi, (p, \sigma)} \in \Delta(\Omega \times H_M^\infty)$ is the unique distribution that satisfies that for all $t \in \mathbb{N}$, $\tilde{\mathcal{H}}_M^t \subset \Omega \times H_M^t$,

$$\mathbb{P}_{\varphi, (p, \sigma)}(\tilde{\mathcal{H}}_M^t \times \prod_{s=t+1}^\infty (\Theta \times MA_\emptyset)) = \mathbb{P}_{\varphi, (p, \sigma)}^t(\tilde{\mathcal{H}}_M^t),$$

⁵⁸We index histories by the messages to distinguish these histories from those when the designer uses direct mechanisms.

where the distributions $(\mathbb{P}_{\varphi, (p, \sigma)}^t)_{t \in \mathbb{N}}$ satisfy

$$\begin{aligned}\mathbb{P}_{\varphi, (p, \sigma)}^{t+1}(\omega, h_M^t, \theta, m, a) &= \mathbb{P}_{\varphi, (p, \sigma)}^t(\omega, h_M^t) f(\theta) p_t(h_M^t, \theta) \sigma_t(h_M^t, \theta)(m) \varphi_t(\omega, \hat{h}_M^t, m)(a), \\ \mathbb{P}_{\varphi, (p, \sigma)}^{t+1}(\omega, h_M^t, \theta, \emptyset, a) &= \mathbb{P}_{\varphi, (p, \sigma)}^t(\omega, h_M^t) f(\theta) (1 - p_t(h_M^t, \theta)) \mathbb{1}[a = a_\emptyset].\end{aligned}$$

Outcome distribution Our interest is in the distribution over payoff-relevant outcomes, $\Omega \times (\Theta \times A)^\infty$, and hence on the marginal of $\mathbb{P}_{\varphi, (p, \sigma)}$ on $\Omega \times (\Theta \times A)^\infty$, which we denote by $\bar{\mathbb{P}}_{\varphi, (p, \sigma)}$.

Best response We say that strategy (p, σ) is a best response for the agent if for all alternative strategies (p', σ') , we have that

$$\liminf_{T \rightarrow \infty} \mathbb{E}_{(p, \sigma)}[U_T] \geq \limsup_{T \rightarrow \infty} \mathbb{E}_{(p', \sigma')}[U_T], \quad (\text{D.8})$$

where recall U_T is the agent's average payoff until period T .

Direct and full participation mechanisms A special case of the above game is that in which $M = \Theta$, and whenever the agent does not participate, the mechanism chooses a_\emptyset with probability 1 for any message in all continuation histories. We call these mechanisms direct and full participation mechanisms. Below, when $M = \Theta$, we drop the dependence of the set histories on M .

Formally, let $\hat{\mathcal{H}}_\emptyset^t$ denote the subset of $\hat{\mathcal{H}}^t$ such that at some point the sequence (\emptyset, a_\emptyset) appears. We define mechanisms

$$\tilde{\varphi}_t : \hat{\mathcal{H}}^t \times \Theta \rightarrow \Delta(A),$$

such that $\tilde{\varphi}_t(\omega, \hat{h}^t, \cdot) = \mathbb{1}[a = a_\emptyset]$ whenever $(\omega, \hat{h}^t) \in \hat{\mathcal{H}}_\emptyset^t$.

Theorem D.1. *Suppose that (p, σ) is a best response to mechanism φ . Then, a direct and full participation mechanism $\tilde{\varphi}$ exists such that*

1. *Participation with probability 1 and truthtelling are a best response for the agent,*
2. *The distribution over $\Omega \times (\Theta \times A)^\infty$ induced by $(\varphi, (p, \sigma))$ is the same as that induced by $\tilde{\varphi}$ under participation and truthtelling.*

Proof. Fix a mechanism $\varphi = (\varphi^t)_{t \geq 1}$ and a best response (p, σ) for the agent in the sense of Equation 13. Let $\mathbb{P}_{\varphi, (p, \sigma)}$ denote the induced distribution over \mathcal{H}^∞ . We write $(\omega, (\theta_t, m_t, a_t)_{t \geq 1})$ for a generic realization, where $m_t = \emptyset \Rightarrow a_t = a_\emptyset$.

The proof proceeds in three steps. In the first step, we construct a direct (but not full participation) mechanism $\tilde{\varphi}$, which under participation and truthtelling after every history on path implements the same outcome distribution as $(\varphi, (p, \sigma))$. In the second step, we verify that participation and truthtelling after every history on path is a best response to $\tilde{\varphi}$. In the third step, we construct a direct and full participation mechanism from $\tilde{\varphi}$. That the agent can always quit the mechanism at each step and obtain a_\emptyset and Step 3 implies that participation and truthtelling after every history is also a best response to the full participation mechanism obtained from $\tilde{\varphi}$.

Step 1: We first construct the direct mechanism $\tilde{\varphi}_t : \hat{\mathcal{H}}^t \times \Theta \rightarrow \Delta(A)$. Define a collection of transition probabilities $\kappa_t : \mathcal{H}_M^t \times \Theta \rightarrow \Delta(M \cup \{\emptyset\})$ as follows:

$$\kappa_t(m_t | h_M^t, \theta_t) = (1 - p_t(h_M^t, \theta_t)) \mathbb{1}[m_t = \emptyset] + p_t(h_M^t, \theta_t) \sigma_t(h_M^t, \theta_t)(m_t).$$

We construct $\tilde{\varphi}$ recursively. In period 1, if the state is ω and the report is θ_1 , the designer draws fictitious m_1 from $\kappa_1(\cdot | \theta_1)$, and implements $a_1 = a_\emptyset$ if $m_1 = \emptyset$, and otherwise draws $a_1 \sim \varphi_1(\omega, m_1)$.

Recursively, for $t \geq 2$, if the sequence of reports, fictitious messages, and allocations is $(\theta'^{t-1}, m^{t-1}, a^{t-1}) = (\theta'_s, m_s, a_s)_{s=1}^{t-1}$ and the agent reports θ_t , the designer draws m_t from $\kappa_t(\cdot | \theta'^{t-1}, m^{t-1}, a^{t-1}, \theta_t)$ and implements $a_t = a_\emptyset$ if $m_t = \emptyset$, and otherwise draws $a_t \sim \varphi_t(\omega, m^{t-1}, a^{t-1}, m_t)$.⁵⁹

It is immediate that under truthtelling and participation the mechanism $\tilde{\varphi}$ implements the same distribution over $\Omega \times (A \times \Theta)^\infty$.⁶⁰

Step 2: Let (p^*, σ^*) denote the agent's strategy that participates and truthfully reports after every history. We now show that (p^*, σ^*) is a best response to $\tilde{\varphi}$ in the sense of Equation D.8.

To do so, we show that for any strategy $(\tilde{p}, \tilde{\sigma})$ in the game induced by the direct mechanism $\tilde{\varphi}$, a strategy (p', σ') exists such that

$$\mathbb{E}_{\tilde{\varphi}, (\tilde{p}, \tilde{\sigma})}[U_T] = \mathbb{E}_{\varphi, (p', \sigma')}[U_T] \quad \text{for all } T, \quad (\text{D.9})$$

where U_T is the average payoff through period T . Given Equation D.9 and the best-response property of (p, σ) to φ ,

$$\liminf_{T \rightarrow \infty} \mathbb{E}_{\varphi, (p, \sigma)}[U_T] \geq \limsup_{T \rightarrow \infty} \mathbb{E}_{\varphi, (p', \sigma')}[U_T] \quad \text{for all } (p', \sigma').$$

Using Step 1, we have $\mathbb{E}_{\varphi, (p, \sigma)}[U_T] = \mathbb{E}_{\tilde{\varphi}, (p^*, \sigma^*)}[U_T]$ for all T , and by Equation D.9 we have $\mathbb{E}_{\varphi, (p', \sigma')}[U_T] = \mathbb{E}_{\tilde{\varphi}, (\tilde{p}, \tilde{\sigma})}[U_T]$ for all T . Hence

$$\liminf_{T \rightarrow \infty} \mathbb{E}_{\tilde{\varphi}, (p^*, \sigma^*)}[U_T] \geq \limsup_{T \rightarrow \infty} \mathbb{E}_{\tilde{\varphi}, (\tilde{p}, \tilde{\sigma})}[U_T] \quad \text{for all } (\tilde{p}, \tilde{\sigma}),$$

which is exactly the definition of (p^*, σ^*) being a best response to $\tilde{\varphi}$. It remains to construct (p', σ') and verify (D.9).

We show how the agent can emulate the strategy $(\tilde{p}, \tilde{\sigma})$ in the game induced by the indirect mechanism via strategy (p', σ') . Define $\tilde{\kappa}_t : H^t \times \Theta \rightarrow \Delta(\Theta \cup \{\emptyset\})$ as follows

$$\tilde{\kappa}_t(\theta' | h^t, \theta_t) = (1 - \tilde{p}(h^t, \theta_t)) \mathbb{1}[\theta' = \emptyset] + \tilde{p}(h^t, \theta_t) \tilde{\sigma}(h^t, \theta_t)(\theta').$$

The strategy (p', σ') privately simulates the report process induced by $(\tilde{p}, \tilde{\sigma})$ in the direct mechanism, and conditional on the fictitious reports, generates the actual message in the indirect mechanism φ using the kernel κ_t in Step 1, evaluated at the fictitious type history.

⁵⁹Recall the fictitious reports encode the agent's participation decisions in the original mechanism.

⁶⁰In fact, if we kept track of the designer's draws of fictitious messages, the new mechanism implements the same distribution over terminal histories \mathcal{H}_M^∞ , and a fortiori, its marginal over $\Omega \times (A \times \Theta)^\infty$ is the same.

Formally, for $t \geq 1$ given the history of types, messages, and allocations through period t , $(\theta^{t-1}, m^{t-1}, a^{t-1})$, and the privately tracked fictitious type reports θ'^{t-1} ,⁶¹ the agent of type θ_t draws a fictitious report $\theta'_t \sim \tilde{\kappa}_t(\cdot | \theta^{t-1}, \theta'^{t-1}, a^{t-1}, \theta_t)$. If $\theta'_t = \emptyset$, then $m_t = \emptyset$ (the agent does not participate in period t). Otherwise, $\theta'_t \in \Theta$ and m_t is drawn from $\kappa_t(\cdot | \theta'^{t-1}, m^{t-1}, a^{t-1}, \theta'_t)$. The allocation is a_\emptyset upon rejection, and $a_t \sim \varphi_t(\omega, m^{t-1}, a^{t-1}, m_t)$, otherwise.

It is immediate that Equation D.9 holds and hence (p^*, σ^*) is a best response to $\tilde{\varphi}$.⁶² In the extensive form game induced by $\tilde{\varphi}$, the designer simulates the agent's participation and reporting strategies using (p, σ) based on the agent's type reports and determines allocations in the mechanism. When the agent's strategy is given by $(\tilde{p}, \tilde{\sigma})$, the process described in the above paragraph correspond to the designer's simulated participation and reporting strategies, and allocations continued to be determined by φ . Hence, in the extensive form game induced by $\tilde{\varphi}$, when the agent plays $(\tilde{p}, \tilde{\sigma})$, it is as if she faces mechanism φ and plays strategy (p', σ') . The best response property of (p, σ) implies that playing (p', σ') yields a weakly worse payoff, and hence $(\tilde{p}, \tilde{\sigma})$ is not a profitable deviation from (p^*, σ^*) in the direct mechanism $\tilde{\varphi}$.

Step 3: Modify $\tilde{\varphi}$ at all histories that include at least one non-participation decision, so that the mechanism implements the outside option a_\emptyset . With this modification, the mechanism satisfies the full participation property. It is immediate that participation and truthtelling after every history remains a best response. \square

⁶¹Recall that to minimize notation and make history lengths symmetric across participation and nonparticipation, we record the agent's rejection of the mechanism as the empty message \emptyset .

⁶²In fact, the construction ensures the stronger property that $\mathbb{P}_{\tilde{\varphi}, (\tilde{p}, \tilde{\sigma})}^T = \mathbb{P}_{\varphi, (p', \sigma')}^T$, when in a slight abuse of notation we keep track of the fictitious messages in $\mathbb{P}_{\tilde{\varphi}, (\tilde{p}, \tilde{\sigma})}^T$.