Screening for Choice Sets*

Tan Gan[†] Yingkai Li[‡]

Abstract

We study a screening problem in which an agent privately knows which actions or technologies are feasible and can disclose only a subset to a principal. Once disclosed, feasible options are verifiable and their payoff consequences are publicly known, so private information concerns feasibility rather than payoffs, misreporting restricts the principal's choices directly rather than distorting her beliefs. Assuming feasible sets are ordered by inclusion, we establish a simple characterization of the optimal mechanism, where the principal either behaves as if there is no asymmetric information or locally provides no reward for better proposals. We derive comparative statics and illustrate the framework in applications to managing persuasion, action elicitation, and production-technology elicitation.

Keywords— screening, choice sets, set inclusion order
JEL— D81, D86

*Yingkai Li thanks the NUS Start-up Grant for financial support. We thank Joyee Deb, Piotr Dworczak, Nima Haghpanah, Marina Halac, Elliot Lipnowski, Harry Pei, Kaihao Yang, and Jidong Zhou for helpful comments and suggestions.

†Department of Managerial Economics and Strategy, London School of Economics. Email: t.gan2@lse.ac.uk

‡Department of Economics, National University of Singapore. Email: yk.li@nus.edu.sg

1 Introduction

Screening and mechanism design are central pillars of modern economic theory, traditionally modeling private information as payoff-relevant heterogeneity over a commonly known set of feasible actions. In such Myersonian environments, agents possess hidden information about preferences, costs, or states, and incentive constraints arise because different types value the same allocation differently. Yet in many economic situations, the principal understands the payoff consequences of feasible actions but does not know which actions are feasible, relying instead on the agent to identify or disclose them. Engineers propose designs/algorithms, managers propose business strategies, firms propose experimental tests to regulators, and civil servants propose policy reforms. In each case, proposed options are verifiable once revealed, but feasible alternatives can be concealed beforehand. Misreporting therefore affects the principal's opportunity set rather than her beliefs about payoffs, a strategic tension that standard screening models are not designed to capture. This paper develops a screening framework in which private information concerns feasibility sets rather than payoff parameters, and studies how a principal should commit to decision rules that discipline strategic disclosure of feasible choice sets.

We formulate a screening model in which the agent privately observes a set of feasible technologies and may report only a subset of these technologies to the principal. The principal always has access to a default technology, which induces a baseline choice set of payoff vectors—pairs (u, v) representing the agent's and principal's payoffs—from which she can choose even if the agent reports no additional technologies. Each reported technology set expands this baseline and induces a larger choice set of payoff vectors from which the principal selects upon receiving the report. Unlike classical Myersonian screening models, the agent's private information here does not concern payoff-relevant parameters but the set of feasible technologies itself. This distinction has two fundamental consequences. First, by concealing feasible technologies, the agent's report directly alters the principal's choice set and hence her attainable payoff, even when the principal perfectly knows the agent's type. Second, once a feasible option is disclosed, its payoff consequences are publicly known, so no inference problem remains. In particular, when an agent with a larger feasible technology set mimics an agent with a smaller feasible technology set, he obtains exactly the smaller-set agent's payoff rather than an information rent. The principal's design problem is therefore to commit ex ante to a decision rule mapping reported technology sets to payoff choices so as to maximize her expected payoff while inducing truthful revelation of the agent's full feasible set.

Modeling private information as knowledge of feasible choice sets rather than payoff parameters naturally entails screening over high-dimensional objects. Without additional structure, this creates severe tractability challenges. A key modeling assumption is that technology sets are ordered by inclusion: more capable agents have access to larger technology sets. This nested structure preserves tractability while capturing a wide range of environments, as we illustrate in the applications. Methodologically, our approach aligns with the recent trend in multidimensional screening such as Yang (2025), which imposes a complete order on the type space while allowing rich heterogeneity in other dimensions and permitting the principal to employ high-dimensional screening instruments.

After presenting the model, Section 2.1 discusses three applications to illustrate how the abstract notions of technologies and choice sets correspond to concrete economic environments. In problems of managing persuasion, a sender privately knows which experiments or information structures can be conducted, while the receiver can verify any disclosed experiment but relies on the sender to make feasible experiments available. In problems of action elicitation, an expert privately knows which detailed actions or reforms can be implemented, while the principal cannot identify or implement these actions without the expert proposing them. In problems of production-technology elicitation, a manager privately knows which production methods, projects, or business strategies are feasible, while the principal can evaluate and authorize any disclosed strategy but cannot access undisclosed ones. In each case, feasible options are verifiable once revealed, concealment restricts the principal's feasible choice set, and more capable agents naturally possess larger sets of feasible options. We will apply our main result to solve these examples later.

We begin our formal analysis by isolating the role of private access to technologies independently of private information. Even when the principal perfectly knows the agent's technology set, the agent can still conceal feasible technologies and thereby restrict the principal's choice set. This feature has no analogue in standard screening models, where under complete information, the principal can directly select her preferred allocation. We show that disciplining such concealment requires extreme off-path incentives: the optimal complete-information mechanism takes a "shoot-the-agent" form in which any report that withholds feasible technologies is punished by assigning the agent the lowest payoff attainable under the default technology. As a result, private access to technologies imposes a single restriction on the principal: the agent must be guaranteed a minimum utility determined by the worst punishment under the default technology.

To organize the analysis that follows, we summarize the complete-information benchmark by the complete information curve, defined as the agent's equilibrium payoff in the optimal mechanism under the complete information benchmark. For each technology set, this payoff equals the agent's utility under the principal's preferred feasible outcome whenever that utility exceeds the punishment guaranteed by the default technology, and it equals the punishment level otherwise. This curve is always pointwise attainable, but it need not be monotone in set inclusion.

We then introduce private information about technology sets. We show that the principal's mechanism design problem can be decomposed into two steps. First, the principal chooses a promised utility function, assigning each reported technology set a payoff for the agent. Second, given any such promised utility, the principal selects her payoff-maximizing point in the corresponding feasible choice set that delivers this promise. While the second step may be computationally non-trivial in some applications, it is conceptually straightforward. This decomposition reduces the screening problem to choosing a promised utility function subject to incentive compatibility and technological feasibility.

A key observation is that, under nested technology sets, incentive compatibility takes a particularly simple form: the promised utility function must be weakly increasing in set inclusion and must guarantee the agent at least his default-report payoff. Moreover, Lemma 4 shows that under nested technology sets it is without loss of generality to restrict attention to promised utility functions that lie within the monotone envelope of the complete information curve—that is, between its lower and upper monotone envelopes. This is illustrated in Fig. 1a. Crucially, once attention is restricted to this envelope, technical feasibility ceases to pose additional constraints: every weakly increasing promised utility function within the envelope is feasible. As a consequence, the screening problem collapses to choosing a one-dimensional monotone promised utility function within an explicitly constructed envelope determined by the complete-information benchmark.

We further characterize the structure of the optimal mechanism. The princi-

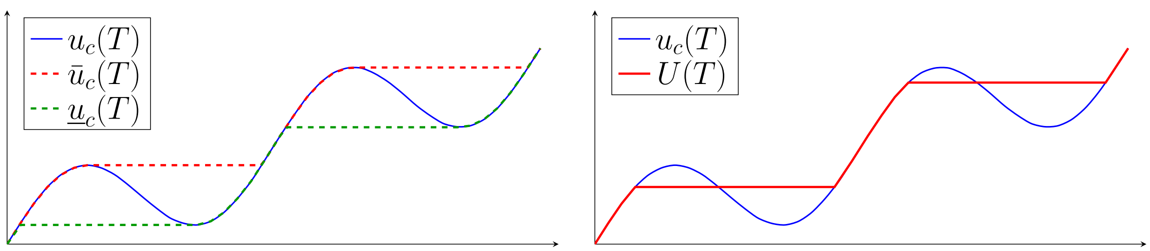

(a) Non-monotone function $u_c(T)$ with its upper monotone closure $\bar{u}_c(T)$ and lower monotone closure $\underline{u}_c(T)$.

(b) Optimal promised utility $U(T)$ given the complete information curve $u_c(T)$.

Figure 1: Illustration of monotone closures and optimal promised utility.

pal's optimal policy exhibits a bang–bang structure. It partitions the type space into a finite collection of intervals. On each interval, either the principal implements the complete-information allocation, so each type receives exactly its complete-information payoff; or the principal keeps the agent's promised utility constant throughout the interval. This is illustrated in Fig. 1b. Moreover, the number of flat segments is bounded by K, the number of decreasing segments of the complete-information curve. As a result, computing the optimal mechanism reduces to choosing at most K departure points from the complete-information curve. The entire mechanism-design problem therefore collapses to the selection of at most K scalar cutoff points and is easy to compute by dynamic programming. The resulting bang–bang structure resembles institutional environments in which principals differentiate rewards finely across disclosed capabilities in some ranges but deliberately compress rewards in other ranges.

We then derive general comparative statics. The principal's equilibrium payoff is increasing when the distribution of agent capabilities improves in the sense of first-order stochastic dominance under set inclusion. It is also increasing in the size of the principal's outside-option technology. The agent, by contrast, is always worse off when the principal's outside-option technology expands, while the effect of improved capabilities is ambiguous. These results formalize when improvements in agents' feasible technology sets benefit the principal, and clarify how enhanced agent capabilities can either benefit or harm agents, depending on how they reshape the optimal bunching regions.

Finally, we apply the framework to three economically relevant environments: managing Bayesian persuasion when senders control experiment design, eliciting expert policy actions from informed civil servants, and contracting with CEOs who privately control additional business strategies under moral hazard. In each application, we show how the complete-information benchmark, its monotone envelope, and the resulting bang–bang structure deliver sharp and sometimes counterintuitive predictions—for example, why richer feasible technology sets may fail to raise agents' equilibrium payoffs, and how even a completely flat promised utility curve may generate rich economic predictions due to the nature of the optimization problem that we conceptually neglect in the general characterization.

In summary, our analysis provides a tractable approach to screening problems in which private information concerns what can be done rather than what payoffs are. By shifting attention from hidden preferences to hidden feasibility and exploiting the nested-set structure, we obtain a unified theory of screening for choice sets that departs fundamentally from classical mechanism design.

1.1 Literature

This paper relates to the literature on project selection, where an agent privately observes a set of feasible projects and may propose one to a principal who can verify the proposed project's characteristics and can reject or implement the project. Armstrong and Vickers (2010) provide the canonical model: the principal chooses a delegation or acceptance set, and the agent selects a feasible project within it. This captures one of our key frictions, as concealment restricts the principal's feasible set. However, for tractability, their paper assumes that an agent can report only one project and that the principal's policy must be deterministic—either accepting or rejecting. Guo and Shmaya (2023) adopt a regret-minimization objective that helps them overcome some tractability issues and discuss the implications of allowing the agent to propose multiple projects simultaneously.

Our paper provides a unified yet tractable framework that embeds project proposals as a special case. In our model, the agent can report a subset of technologies, which impose structures over payoff pairs, instead of a single payoff pair or any arbitrary subset of payoff pairs. The principal can commit to taking any actions enabled by the reported technologies rather than simply accepting or rejecting a proposal. These modifications substantially expand the scope of the general model and enable us to study strategic proposals regarding signal structures, production technologies, and unknown actions.

This paper also relates to our earlier work Bergemann et al. (2025), which studies the specific example of managing persuasion, where the principal aims to discipline the strategic persuasion of an agent with private access to experiments/signals. To avoid tractability challenges, Bergemann et al. (2025) adopts a robustness approach in which the principal minimizes worst-case regret over uncertainty in both the feasibility signal and the sender's partially aligned preferences. The present paper uses a classical Bayesian objective and focuses only on uncertainty over feasibility, thereby isolating the novel effect of screening on feasibility and providing a clearer contrast with the classical literature. Curello and Sinander (2025) analyze a dynamic disclosure problem in which the agent privately observes a breakthrough that expands the feasible frontier. The uncertainty concerns the arrival time of the breakthrough, and optimal mechanisms exhibit deadline structures. Their result remains unchanged if the agent privately knows the arrival time at the beginning, and this modified static model is included in our framework. An agent with an earlier arrival time has a larger choice set; the corresponding complete-information curve is strictly decreasing; and the optimal promised-utility curve is flat, so the optimal mechanism exhibits the same deadline structures.

Another relevant strand is mechanism design with verifiable evidence or partial disclosure, where agents may possess hard information that cannot be fabricated but can be withheld. Classic and modern contributions study how disclosure constraints affect implementability and mechanism structure (e.g., Green and Laffont, 1986; Bull and Watson, 2007; Deneckere and Severinov, 2008; Ben-Porath et al., 2019). Our model shares one similarity in that the "message" the agent can send depends on his type. However, in that literature, the message is still "cheap" in the sense that it does not directly change payoffs. Upon observing the message, the principal must infer payoffs from the message structure, and this inference is often imperfect. In contrast, in our model, the agent can directly affect the principal's payoff by hiding technologies, and once technologies are reported (on or off path), the principal can always perfectly evaluate the consequences of her choices.

Finally, the paper is also related to multidimensional screening, which is generally intractable without strong structure. A recent approach regains tractability by imposing an order on types. For example, Yang (2025) focuses on consumers ranked by their willingness to pay and finds conditions under which nested bundling is optimal in a multidimensional pricing environment. Our paper adopts a similar methodological idea—imposing an inclusion order—but with private information being the feasible *set* rather than preferences.

2 Model

Primitives We consider a principal-agent model of screening for choice sets. There is a compact set $\mathcal{T}$ of all possible technologies t. For a set of technologies $T \subseteq \mathcal{T}$, we define

$$C(T) : 2^{\mathcal{T}} \to 2^{\mathbb{R}^2}$$

as the principal's choice set. That is, if the principal gets access to a subset of technologies $T \subseteq \mathcal{T}$, she will be able to choose a pair of payoffs $(u, v) \in C(T) \subseteq \mathbb{R}^2$, where (u, v) represents the payoffs to the agent and the principal, respectively. $C(T)$ is compact-valued and is monotone in set inclusion: $C(T') \subseteq C(T)$ for any $T' \subseteq T$.

The principal has a default technology $t_0 \in \mathcal{T}$ from which she can choose, and she relies on the agent's report on additional technologies to expand her feasible choice set. Each agent has private access to a subset of technologies $T \subset \mathcal{T}$, which is the agent's private information (type). Without loss of generality, we can always let $t_0 \subseteq T$ so $C(T)$ is the choice set that the principal faces if the type T agent truthfully reports all his technologies.

The principal does not observe the agent's private type T and holds a prior belief F. With a slight abuse of notation, we denote the support of F as $\mathcal{T}$. Due to the tractability challenge of high-dimensional screening problems, we primarily focus on the environment where the agent can be ranked by his capability: the more capable agent has a larger set of available technologies.

Assumption 1 (Nested Technology Sets)**.** *The set of available technologies is nested. That is, for any $T, T' \in \mathcal{T}$, either $T \subseteq T'$ or $T' \subseteq T$.*

As we will discuss through examples and analysis, this assumption is reasonable in many applications and creates tractability in high-dimensional screening problems while maintaining some degree of flexibility to capture the complex nature of choice sets.

Reports and Mechanisms An agent with type T can strategically report a subset T' such that $\{t_0\} \subseteq T' \subseteq T$ to the principal. That is, he can conceal available technologies but can never falsify technologies that do not exist. This assumption is appropriate for technologies, as they are hard to invent but easy to verify. If the agent reports nothing, then the principal is left with the default technology $\{t_0\}$.

To discipline the strategic report of the agent, the principal can ex-ante commit to a decision rule $a(T)$

$$a(T) : \mathcal{T} \to \Delta C(T)$$

that maps a reported set of technology to a distribution of payoff pairs within the feasible choice set induced by the technology. Because the principal can commit to arbitrary random actions, without loss of generality, we assume $a(T)$ is a deterministic mapping to the expected payoff pairs such that

$$a(T) : T \mapsto (u, v) \in \text{conv}(C(T)).$$

Although the setting is different from the standard Myersonian screening models, the idea of the revelation principle still holds, and it is without loss of generality to focus on direct mechanisms where the agent truthfully reports his technology set T. The objective of the mechanism design is:

$$\begin{aligned} \max_{a(\cdot) \in C(T)} \quad & \mathbf{E}_F[a_2(T)] \\ \text{s.t.} \quad & a_1(T) \geq a_1(T') \quad \forall T' \subset T. \end{aligned}$$

2.1 Applications

The language of our model is deliberately abstract and contains many layers to encompass various screening problems related to feasible choice sets. We start by discussing a few examples to illustrate how our general model relates to different applications. These applications are also stated in general terms. We will provide more detailed examples and characterize their solutions after the main result.

Manage Persuasion This problem concerns how a receiver with commitment power can discipline the strategic persuasion of a sender as in Bergemann et al. (2025). There is a state space Θ and an action space A. Both the sender's utility function $u(a,\theta)$ and the receiver's utility function $v(a,\theta)$ are publicly known. The receiver shares a common prior G about the state θ with the sender.

The receiver's (principal's) default choice set is to take a (possibly random) action $a \in \Delta A$ for all states based on the prior G, so the choice set under default technology is

$$C(\{t_0\}) = \{(u,v)|\ \exists a \in \Delta A,\ \text{s.t. } u = \mathbf{E}_{\theta\sim G}[u(a,\theta)]\,,\ v = \mathbf{E}_{\theta\sim G}[v(a,\theta)]\}.$$

The sender can conduct experiments to persuade the receiver, as in the literature on Bayesian persuasion (Kamenica and Gentzkow, 2011). An experiment consists of a signal space S and a signal mapping $\sigma : \Theta \to S$. Given any experiment (S,σ), the principal can adopt a (possibly random) strategy a that maps from S to ΔA. A technology t corresponds to an experiment (S,σ). The choice set with type T is

$$\begin{aligned} C(T) = \{(u,v)|\ &\exists (S,\sigma) \in T,\ a : S \to \Delta A \\ &\text{s.t. } u = \mathbf{E}_{\theta\sim G, s\sim\sigma(\theta)}[u(a(s),\theta)]\,,\ v = \mathbf{E}_{\theta\sim G, s\sim\sigma(\theta)}[v(a(s),\theta)]\}. \end{aligned}$$

However, the agent may not have the ability to perfectly fine-tune the experiments. Thus, there is only a set of experiments available, which is captured by the agent's private type T. The agent can strategically report a subset of available experiments. Anticipating this, the principal optimally commits to the choice of the experiment and her action upon the signal realization.

In this setting, Assumption 1 requires that the agent be ranked by how many experiments they can conduct: a more capable agent can conduct more experiments in terms of set inclusion. Assumption 1 does not impose any restrictions on what these experiments are. Note that our general framework can also accommodate the cost of the experiment $c(\sigma)$, which may vary with the choice of the experiment. This simply shifts the technology t downward by the cost of the experiment on the coordinate axis that represents the payoff of the agent.

Actions Elicitation This problem concerns how a principal who has a good understanding of the state but lacks the expertise or time to identify what detailed strategies are feasible. There is a state space $\theta \in \Theta$ drawn from common prior G, and a signal structure $\sigma : \Theta \to S$ that is privately observed by the principal. An action $a : S \to (u,v)$ is a mapping from states to payoff pairs. The principal has a known set of actions A_0 from which she can choose. Thus, her default choice set is

$$C(\{t_0\}) = \{(u, v) | \ \exists a : S \to \Delta A_0, \\ \text{s.t. } u = \mathbf{E}_{\theta \sim G, s \sim \sigma(\theta)}[u(a(s), \theta)] \,, \ v = \mathbf{E}_{\theta \sim G, s \sim \sigma(\theta)}[v(a(s), \theta)]\}.$$

The agent has private access to some additional actions, such as a detailed feasible schedule or a novel algorithm. He can choose to conceal information in his reports, but he cannot falsify an action, as the principal will be able to easily verify and understand the payoff consequences. A technology t corresponds to an additional action. The choice set with type T is

$$C(T) = \{(u, v) | \ \exists a : S \to \Delta(A_0 \cup T), \\ \text{s.t. } u = \mathbf{E}_{\theta \sim G, s \sim \sigma(\theta)}[u(a(s), \theta)] \,, \ v = \mathbf{E}_{\theta \sim G, s \sim \sigma(\theta)}[v(a(s), \theta)]\}.$$

In this setting, Assumption 1 requires that the agent is ranked by how many additional actions they have private access to: a more capable agent has more actions in terms of set inclusion.

Eliciting Production Technologies In this example, we highlight that the subsequent interaction between the principal and the agent, after the agent reports his set of available technologies, can be sophisticated. Consider the application where the board of a company (the principal) is contracting with a CEO (the agent). A business strategy (G, c) consists of a distribution $G \in \Delta\mathbb{R}$ over the company's profit and the cost of the agent. The company has an initial set of business strategies $\mathcal{G}_0$, from which the agent can privately choose. The principal can use a non-negative wage scheme $w(\cdot) : \mathbb{R} \to \mathbb{R}^+$ to provide incentives for the agent. Thus, the principal's default choice set is

$$C(\{t_0\}) = \{(u, v) | \ \exists w(\cdot), \ (G, c) \in \mathcal{G}_0, \ \text{s.t. } v = \int y - w(y) \, \mathrm{d}G(y), \\ u = \int w(y) - c \, \mathrm{d}G(y) \geq \max_{(G', c') \in \mathcal{G}_0} \int w(y) - c' \, \mathrm{d}G'(y)\}.$$

The agent may have some other business strategies, but unlocking them requires the principal's permission or even additional investment. A technology t corresponds to a set of business strategies $\mathcal{G}_t$, which may not necessarily be a singleton. For example, it is possible that a technology simultaneously introduces a productive business strategy and a shirking business strategy, and they cannot be separated.

After the agent reports the available set of additional business strategies (technology) $\cup_{t\in T}\mathcal{G}_t$, the principal can permit a subset T' of them: $\mathcal{G}(T') = \cup_{t\in T'}\mathcal{G}_t$. With such selective permission, the set of payoffs that she can induce via a wage scheme is

$$\tilde{C}(\mathcal{G}(T')) = \{(u,v)|\ \exists w(\cdot),\ (G,c) \in \mathcal{G}_0 \cup \mathcal{G}(T'),\ \text{s.t. } v = \int y - w(y)\,\mathrm{d}G(y),$$
$$u = \int w(y) - c\,\mathrm{d}G(y) \geq \max_{(G',c')\in\mathcal{G}_0\cup\mathcal{G}(T')} \int w(y) - c'\,\mathrm{d}G'(y)\}.$$

Note that due to the complex nature of incentive compatibility, $\tilde{C}(\mathcal{G}(T'))$ may not be monotonic in set inclusion in T'. However, because the principal has the freedom to choose what $\mathcal{G}'$ to permit and what wage scheme to provide. Her ultimate choice set is monotonic in set inclusion in T:

$$C(T) = \cup_{T'\subseteq T}\tilde{C}(\mathcal{G}(T')).$$

In this example, Assumption 1 holds as long as a more capable agent has more available business plans.

3 Analysis

3.1 Complete Information Benchmark

The agent in our model not only has private information about what technology is available, but also has private access to these technologies. To highlight the difference between these two elements, we start with the complete information benchmark where the principal's prior F has a degenerate support $\mathcal{T} = \{T_0\}$.

Ideally, the principal wants the agent to report T_0, after which the principal can select her favorite choice. However, because the agent has private access to the technologies, he can choose to report only a subset $T' \subseteq T$, which limits the principal's choice set. This highlights the first difference between our framework and the standard Myersonian framework: even in the complete information benchmark, the agent can and may directly affect the principal's payoff by changing the principal's available options.

To discipline the agent's strategic selection of technologies, it is without loss of generality for the principal to use the "shoot the agent" mechanism. That is, the principal commits to minimizing the agent's payoff whenever the agent reports $T \neq T_0$. The value of this minimum payoff is clearly decreasing in T in set inclusion. Thus, among all $T \neq T_0$, the best choice for the agent is to report nothing: $T' = \{t_0\}$. The corresponding value $\underline{u}$, where

$$\underline{u} = \min_{(u,v)\in C(\{t_0\})} u, \tag{1}$$

is the minimum value that the principal has to offer to the agent on the equilibrium path to ensure that he reports T truthfully. Therefore, we formally define the "shoot the agent" mechanism as

$$a(T) \in \begin{cases} \arg\min_{(u,v)\in C(T)} u & \text{if } T \neq T_0 \\ \arg\max_{(u,v)\in C(T), u\geq \underline{u}} v & \text{if } T = T_0. \end{cases}$$

Proposition 1. *In the complete information benchmark, "shoot the agent" is an optimal mechanism.*

For the purpose of subsequent analysis, we are interested in understanding the agent's payoffs under the optimal mechanism in the complete information benchmark. To simplify the exposition, we impose a generic assumption that the optimal choice of the principal with any technology $T \in \mathcal{T}$ is unique. All our results generalize without this assumption and we delay the discussion in the extension.

Assumption 2. *The optimal choice of the principal* $a^p(T) = \arg\max_{(u,v)\in \text{conv}(C(T))} v$ *is unique given any type* T.

With this assumption, $a_1^p(T)$ is the agent's payoff when the principal chooses her optimal choice within T, and we know

Corollary 1. *Under the optimal mechanism of the complete information benchmark, the agent's payoff is*

$$u_c(T) = \max\{\underline{u}, a_1^p(T)\}. \tag{2}$$

We denote $u_c(T)$ as the complete information curve. In the absence of private information about the technologies T, private access to the technology gives the agent very minimal power: he is willing to reveal all he has and let the principal choose her preferred option as long as it is weakly better than the worst punishment under the default technology.

3.2 Optimal Mechanism

After discussing the complete information benchmark and introducing the complete information curve, we are ready to discuss the optimal mechanism with private information.

Promised Utility Function The first step of the analysis is to point out that the design problem can be separated into two independent problems. Unless otherwise specified, we always impose Assumption 1. Therefore, to simplify notation, we also denote the agent's type T as a real number in $[0, 1]$, where a real number $T \leq T'$ implies the technology set $T \subseteq T'$.

For an (incentive-compatible) direct mechanism $a(T)$, $U(T) = a_1(T)$ is the agent's payoff when he reports T. We call this $U(T)$ function the promised utility function. Once the promised utility function is given, the rest of the design problem is non-strategic:

Lemma 1. *It is without loss of optimality to focus on direct mechanisms $a(T)$ where*

$$\begin{aligned} a(T) = \arg\max_{(u,v)} \quad & v, \\ \text{s.t.} \quad & (u, v) \in \text{conv}(C(T)), u = U(T) \end{aligned}$$

In some applications, e.g., business strategies with moral hazard, the analysis of the exact boundary of $C(T)$ can be challenging. Nevertheless, it is conceptually simple, and we will define

$$V(T, u) = \max_{(u',v') \in \text{conv}(C(T)), u'=u} v',$$

and proceed as if we fully understand $V(T, u)$. This decomposition helps us to smooth out the complexity of the problem and identify general properties independent of the detailed structure of the feasible choice set $C(T)$.

Not all promised utility functions can be achieved by some choice $(u, v) \in C(T)$. The second step is to characterize which promised utility functions are (technologically) feasible.

Lemma 2. *$U(T)$ is feasible if and only if $U(T) \in [\underline{u}_f(T), \bar{u}_f(T)]$, where*

$$\underline{u}_f(T) = \min_{(u,v)\in C(T)} u, \quad \bar{u}_f(T) = \max_{(u,v)\in C(T)} u.$$

$\underline{u}_f(T)$ is decreasing in T and $\bar{u}_f(T)$ is increasing in T.

The feasibility constraint is often quite messy in optimal control problems. However, as we will show later in Lemma 4, under our set inclusion type assumption (Assumption 1), the feasibility constraints will not bind in the optimal mechanism.

The third step is to characterize which promised utility functions $U(T)$ are incentive compatible. Clearly, if $T \geq T'$, then the type T agent always has the option to report his subset T' and gets $U(T')$. Thus, $U(T)$ must be weakly increasing in T. In fact, this is the only incentive constraint.

Lemma 3. *A promised utility function $U(T)$ is incentive compatible if and only if*

(i) $U(T)$ is weakly increasing;

(ii) $U(\{t_0\}) \geq \underline{u}$.

Lemma 3 highlights another difference between our screening model of feasibility and the Myersonian screening model of payoff-relevant information. In our model, high types only get a weakly higher payoff relative to low types, instead of obtaining a strictly higher payoff which depends on the mechanism's allocation rule. In fact, when high types misreport as a low type, they receive exactly the same payoff as the low type; whereas in the Myersonian screening model, high types obtain a strictly higher payoff even with the same allocation, which creates information rent. This difference comes from the fact that the principal does not need to make any inferences once the set of technologies is reported: she can perfectly evaluate the consequences of all her choices.

With the help of Lemmas 1 to 3, we can rewrite the principal's design problem as

$$\begin{aligned} \max_{U(\cdot)} \quad & \int V(T, U(T))\, \mathrm{d}F(T) \\ \text{s.t.} \quad & U(T) \text{ is weakly increasing,} \\ & U(T) \in [\underline{u}_f(T), \bar{u}_f(T)], \\ & U(\{t_0\}) \geq \underline{u}. \end{aligned} \tag{OPT}$$

Optimal Promised Utility Functions Perhaps surprisingly, the complete-information curve $u_c(T)$, defined in Equation (2), is crucial to the characterization of the optimal promised utility function. To see this, we start with two simple observations. First, by construction, $u_c(T)$ is always feasible. Second, if $u_c(T)$ is increasing in T, then it is the optimal promised utility function, and the principal can attain the complete-information payoff. The central question, therefore, is how to characterize the optimal promised utility function when the complete-information curve is not monotone. To this end, we define the upper monotone closure $\bar{u}_c(T)$ as the minimum monotone function such that $\bar{u}_c(T) \geq u_c(T)$ for all T. That is,

$$\bar{u}_c(T) = \max_{T' \leq T} u_c(T'), \quad \forall T \in \mathcal{T}. \tag{3}$$

See the red dashed line in Figure 1 for an illustration. Similarly, the lower monotone closure $\underline{u}_c(T)$ is the maximum monotone function such that $\underline{u}_c(T) \leq u_c(T)$ for all T, and

$$\underline{u}_c(T) = \min_{T' \geq T} u_c(T'), \quad \forall T \in \mathcal{T}. \tag{4}$$

A promised utility function $U(T)$ is in the monotone envelope (of the complete information curve) if

$$U(T) \in [\underline{u}_c(T), \bar{u}_c(T)], \quad \forall T \in \mathcal{T}. \tag{5}$$

It turns out that it is without loss of generality to focus on a monotone promised utility function in the monotone envelopes. The intuition is that we can always project the promised utility function into the monotone envelope while keeping the monotonicity, as illustrated in Figure 2, to improve the principal's payoff. Moreover, once

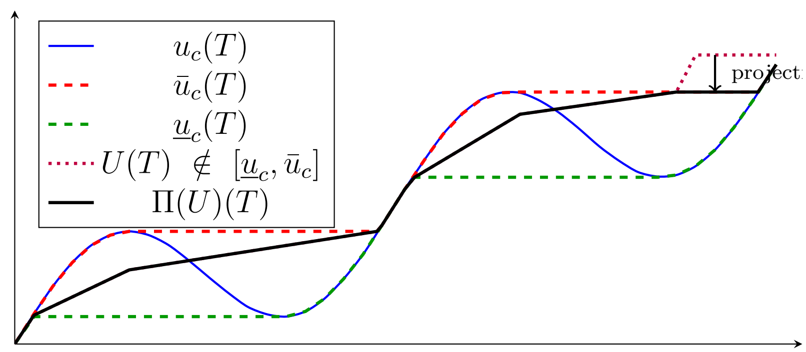

Figure 2: Illustration for Lemma 4: a monotone promised utility $U(T)$ that lies outside the monotone envelope can be projected into the envelope by clamping to $\bar{u}_c(T)$ where it violates feasibility/optimality; this projection weakly increases the principal's payoff.

we focus on the monotone promised utility function within the monotone envelope, the feasibility constraint can be neglected.

Lemma 4. *Any promised utility function in the monotone envelope is feasible. Moreover, there exists a monotone promised utility function in the monotone envelope, which solves the design problem* (OPT)*.*

Lemma 4 greatly simplifies the design problem: we just need to search over all monotone promised utility functions within the monotone envelope without worrying about any other constraints. To prepare for a sharper characterization of the optimal promised utility, we quantify the non-monotonicity of the complete-information curve. Let K denotes the number of maximal intervals on which the complete-information curve is strictly decreasing, or formally,

$$K \equiv \#\Big\{I \subseteq [0,1]: \ I \text{ is a maximal interval such that } u_c(\cdot) \text{ is strictly decreasing in I}\Big\}.$$

The principal's optimal policy features a bang–bang structure across intervals: on each interval, she either implements the complete-information allocation or keeps the agent's promised utility constant throughout. Moreover, the number of constant-promise intervals is bounded by K.

Theorem 1 (Optimal Mechanisms)**.** *There exists a partition of the type space into intervals $\{I_j\}_{j \in J_1 \cup J_2}$ such that $|J_1| - 1 \leq |J_2| \leq K$ and*

1. *If $j \in J_1$, $U(T) = u_c(T)$ for any type $T \in I_j$;*

2. If $j \in J_2$, $U(T)$ is a constant function in I_j.

Theorem 1 provides a rationale for two seemingly abnormal organizational behaviors: why institutions may deliberately limit incentives for marginally better proposals, and why they may fail to fully exploit newly disclosed feasible options. More strikingly, it shows that these two behaviors cannot appear in isolation in an optimal mechanism but instead may arise jointly. According to Theorem 1, either the principal fully exploits added feasibility—so richer feasible option sets are used in the best possible way for the principal—and continuous, strict incentives are provided, or she must both compress rewards and deliberately refrain from myopic exploitation to sustain truthful disclosure. One distortion cannot arise without the other.

This joint emergence of reward compression and deliberate constraints from the efficient exploitation of feasible options has clear organizational implications. In some capability ranges, better proposals lead to both better internal decisions and better treatment of the proposer: additional experiments, richer action plans, or more sophisticated business strategies are actively used to improve organizational outcomes, and the agent is rewarded accordingly. In other ranges, institutions deliberately maintain stable treatment even as proposals improve. Committees may follow standard approval rules, boards may adhere to fixed compensation packages, or regulators may apply uniform authorization pathways, despite being presented with richer feasible options. In these regions, additional disclosed capabilities expand the principal's internal choice set but are not fully exploited in the short run, because committing to stable treatment is necessary to sustain transparent disclosure of what is feasible.

4 Properties and Comparative Analysis

4.1 Rationalizable Mechanisms

We first show that the characterization of the optimal mechanism in Theorem 1 is essentially the strongest possible without imposing further parametric assumptions on the primitives, by demonstrating that any such promised utility function can arise as the unique optimal mechanism under some choice of primitives.

Define a candidate promised utility function $u(T)$ to be any function of the form described in Theorem 1. Such a function is fully determined by the complete-information curve $u_c(T)$ and at most K departure points from it. Because candidate promised utility functions are far lower-dimensional than the space of model primitives, varying only a subset of primitives suffices to rationalize every candidate promised utility function. To illustrate the high-dimensional nature of the primitives, we consider two classes of variations, each of which can rationalize all candidate promised utility functions. To simplify the exposition, we will assume that the complete information curve $u_c(T)$ is continuous in T, and hence the optimal promised utility $U(T)$ is also continuous in T.

Varying Distribution First, we fix the technology set $C(T)$, which determines the complete-information curve $u_c(T)$, and vary the distribution F over types $T \in [0, 1]$.

Proposition 2. *For any $C(T)$ and any candidate promised utility function $U(T)$, there exists a distribution F such that $U(T)$ is F-a.s. uniquely optimal.*

Varying Technology Set Second and, more importantly, we fix the complete-information curve $u_c(T)$ and the distribution F over types $T \in [0, 1]$ and vary the technology set by changing the choice set $C(T)$ associated with each type T. Here, the type $T \in [0, 1]$ represents the rank of the agent in terms of capability.

Proposition 3. *For any complete-information curve $u_c(T)$, any fully supported F, and any candidate promised utility function $U(T)$, there exists a monotone choice set function $C(T)$ that is consistent with $u_c(T)$, such that $U(T)$ is F-a.s. uniquely optimal.*

This exercise highlights that although types T are fully ranked, the model retains a high-dimensional character. The ranking orders agents by capability, but substantial heterogeneity remains in the internal geometry of choice sets and payoff possibilities. Consequently, many distinct underlying primitives are consistent with the same ranked type structure and complete-information curve.

4.2 Monotone Comparative Statics

Next, we derive monotone comparative statics for the principal and the agent's optimal payoff. Throughout, we maintain Assumption 1 and identify types with $T \in [0, 1]$ so that $T \leq T'$ if and only if the technology set of type T is contained in that of type T'.

For a type distribution F on $[0, 1]$, let $\text{OPT}(F)$ denote the value of (OPT) and let U^* be the corresponding optimal promised utility.

Technology Expansions We first consider a comparative static that improves the agent's feasible technologies *type-by-type.* In many applications, technological progress, improved institutional capacity, or regulatory changes expand the set of feasible technologies available to each agent. Importantly, such expansions may be heterogeneous across types and therefore may change the support of the type distribution; moreover, even if the baseline type space is totally ordered by set inclusion (Assumption 1), the *expanded* technology sets need not remain nested. Nonetheless, the principal's optimal expected payoff is monotone under such expansions.

Definition 1 (Technology expansion dominance)**.** *We say that* $\hat{F}$ dominates F *in the sense of technology expansion, denoted* $\hat{F} \succeq F$*, if there exists a measurable map* ϕ *such that*

$$\phi(T) \supseteq T \quad \textit{for all } T, \qquad \textit{and} \qquad \hat{T} = \phi(T) \ \ \textit{when} \ \ T \sim F,$$

so that $\hat{F}$ *is the distribution of* $\hat{T}$ *induced by* ϕ*.*

Proposition 4 (Principal's payoff under technology expansion)**.** *If* $\hat{F} \succeq F$ *and* F *satisfies Assumption 1, then* $\text{OPT}(\hat{F}) \geq \text{OPT}(F)$*.*

Definition 1 allows the post-expansion distribution $\hat{F}$ to violate the set-inclusion order: ϕ may expand technologies in a type-dependent way, so that $\phi(T)$ and $\phi(T')$ need not be comparable even when $T \subseteq T'$. Our argument only uses that the *baseline* distribution F satisfies Assumption 1.

The proof of Proposition 4 uses the following technical lemma, which shows that one can select an optimal mechanism (for the baseline problem) whose induced principal payoff is monotone along the baseline chain.

Lemma 5 (Monotone interim principal payoff)**.** *There exists an optimal promised utility* U^* *for* (OPT) *such that* $T \mapsto V(T, U^*(T))$ *is weakly increasing.*

To prove Proposition 4, intuitively, the principal can directly adopt the optimal promised utility U^* for the dominated distribution F, selected so that the induced principal payoff $V^*(T) \equiv V(T, U^*(T))$ is weakly increasing in the type of the agent

(Lemma 5). More specifically, let $\pi(S)$ be the largest type in support of F contained in S, and the mechanism promises $U^*(\pi(S))$ and then chooses the payoff-maximizing outcome in $C(S)$ delivering that promised utility. Feasibility holds because $\pi(S) \subseteq S$. Incentive compatibility holds because underreporting shrinks S, which can only (weakly) decrease $\pi(S)$ and hence the promised utility. Finally, at the true expanded set $\hat{T}$ we have $\pi(\hat{T}) \geq T$, so (by monotonicity of V^*) the principal's payoff is at least her payoff at T. Averaging over $T \sim F$ yields $\text{OPT}(\hat{F}) \geq \text{OPT}(F)$.

Technology expansion dominance is a *pointwise* improvement: each realized technology set is (weakly) enlarged. A different and standard comparison is a *distributional* improvement in the scalar capability index in the sense of first-order stochastic dominance (FOSD). The next corollary records that the principal's value is also monotone under FOSD.

Corollary 2 (First-order stochastic dominance)**.** *Suppose $\hat{F}$ first-order stochastically dominates F. Then* $\text{OPT}(\hat{F}) \geq \text{OPT}(F)$.

Under first-order stochastic dominance, there exists a standard monotone coupling ϕ such that $\hat{T}$ is a technological expansion of distribution T based on ϕ. Therefore, $\hat{F} \succeq F$ in the sense of Definition 1. Proposition 4 then immediately implies $\text{OPT}(\hat{F}) \geq \text{OPT}(F)$, and Corollary 2 holds.

When the technologies available to the agent expand, their effects on the agent's utility are ambiguous. Intuitively, whether the agent's utility improves depends on the alignment between the principal and the agent. To demonstrate this idea, we provide a simple illustration using binary types for the agent.

Specifically, consider the example with binary types $T_0 \leq T_1$. Let q be the probability of T_1 in the prior distribution F. Note that when q increases, the distribution increases in the FOSD sense. We also assume that $\underline{u} \leq u_c(T_0)$ and $\underline{u} \leq u_c(T_1)$, where $\underline{u} = \min_{(u,v)\in C(\{t_0\})} u$. We consider the impact of q on the agent's utility in two cases.

- $u_c(T_0) < u_c(T_1)$. In this case, the optimal promised utility aligns with the complete information curve. When q increases, the agent receives a higher promised utility with a higher probability. The ex ante payoff of the agent increases.

- $u_c(T_0) > u_c(T_1)$. In this case, there is a conflict of interests in the incentives. Due to the incentive compatibility constraint, Theorem 1 implies that the principal offers a constant promised utility to the agent in the optimal mechanism.

Let u_q be this constant utility. Note that in this example, u_q is decreasing in q. This is because the principal's payoff given any constant promised utility $u \in [u_c(T_1), u_c(T_0)]$ is

$$(1-q) \cdot V(T_0, u) + q \cdot V(T_1, u)$$

which is concave in u. Moreover, its derivative is $(1-q) \cdot V_u(T_0, u) + q \cdot V_u(T_1, u)$, where $V_u(T_0, u)$ is decreasing and weakly positive, and $V_u(T_1, u)$ is decreasing and weakly negative for $u \in [u_c(T_1), u_c(T_0)]$. Therefore, by increasing q, the maximizer u_q decreases. The ex ante payoff of the agent decreases.

Expansion of Default Technologies While technological expansion in general may have ambiguous effects on the agent, we show that expanding the *default* technology unambiguously improves the payoff of the principal while decreasing the payoff of the agent. To study comparative statics with respect to the default set, we allow the default set to be an arbitrary set $t_0 \subseteq \mathcal{T}$ and maintain that

$$t_0 \subseteq T \quad \text{for every type } T \text{ in the support.}$$

Thus, an agent of type T can report any set S such that $t_0 \subseteq S \subseteq T$, and if he discloses nothing beyond the default, the principal can still choose a payoff in $C(t_0)$.

Let

$$\underline{u}(t_0) \;\equiv\; \min_{(u,v) \in C(t_0)} u \tag{6}$$

denote the lowest payoff the principal can deliver to the agent using only default technologies. By set inclusion and monotonicity of $C(\cdot)$, if $t_0' \subseteq t_0$, then $C(t_0') \subseteq C(t_0)$ and hence

$$\underline{u}(t_0) \;\leq\; \underline{u}(t_0').$$

Proposition 5 (Principal's payoff under default expansion)**.** *For any $t_0' \subseteq t_0$, the principal's optimal expected payoff is weakly higher under default technology t_0 than under t_0'.*

Proof. For any $t_0' \subseteq t_0$, any promised utility $U(T)$ that is optimal under default technology t_0' remains feasible and incentive compatible for the principal under default technology t_0, as $\underline{u}(t_0) \;\leq\; \underline{u}(t_0')$. Therefore, the optimal payoff of the principal is weakly higher under default technology t_0. □

The following result illustrates the monotonicity of the optimal promised utility of the agent. Intuitively, a smaller set of default options increases the agent's utility pointwise for all types in the optimal mechanism in order to incentivize the truthful reporting of the agent, due to the fact that a smaller set of default options leads to a higher utility for the agent in the worst punishment for deviation.

Proposition 6 (Agent's payoff under default expansion)**.** *Fix* $t_0' \subseteq t_0$*. Let* $U^{t_0}(\cdot)$ *be any optimal promised utility function in the economy with default* t_0*. Then there exists an optimal promised utility function* $U^{t_0'}(\cdot)$ *in the economy with default* t_0' *such that*

$$U^{t_0'}(T) \;\geq\; U^{t_0}(T) \qquad \forall T.$$

5 Applications

5.1 Manage Persuasion

Consider the application of FDA drug approvals. In this application, the pharmaceutical company can design tests to reveal the effectiveness of the drugs. Each drug has two states $\{+, -\}$, indicating whether the drug is effective or not. The prior probability of $+$ is $q \in (0, 1)$. Each test has binary signals $\{s_+, s_-\}$. Let $\alpha \triangleq \Pr[s_+ \mid -]$ be the false positive rate and $\beta \triangleq \Pr[s_- \mid +]$ be the false negative rate. The pharmaceutical company has access to a set of tests that can trade off between the false positive and false negative rates. In particular, there exists a decreasing convex function ψ such that the feasible pair of α and β is given by $\beta = \psi(\alpha)$. Moreover, $\psi(0) = 1, \psi(1) = 0$, $\psi'(0) = -\infty, \psi'(1) = 0$. A type $T \in [\underline{T}, \bar{T}]$ pharmaceutical company has access to tests corresponding to a range of false positive rates $[\underline{\alpha}(T), \bar{\alpha}(T)]$. We assume that $\underline{\alpha}$ is decreasing in T and $\bar{\alpha}$ is increasing in T, so that a more capable company can access more tests with a wider range of tradeoffs.

The FDA can decide whether to approve the drug based on the test results. The FDA's utility of rejection is 0. If an effective drug is approved, the FDA has a utility of 1. If an ineffective drug is approved, the FDA has a utility of $-L$, where $L > 0$. We assume that the FDA prefers not to approve given the prior, i.e., $q - (1-q)L < 0$. The utility of the pharmaceutical company is state independent. They have a utility of 1 if and only if the drug is approved.

Optimal Mechanism Given any false positive rate α, the FDA's expected utility from approving the drug for a positive signal s_+ is $R(\alpha) \triangleq q(1-\psi(\alpha)) - (1-q)\alpha L$. Let $\hat{\alpha} \in (0,1)$ be the false positive rate such that $R(\hat{\alpha}) = 0$.[1] The corresponding approval probability is

$$P(\alpha) \triangleq \begin{cases} q(1-\psi(\alpha)) + (1-q)\alpha & \alpha \le \hat{\alpha}, \\ 0 & \alpha > \hat{\alpha}. \end{cases}$$

By the first order condition, the optimal payoff for the FDA is attained by α^* such that $\psi'(\alpha^*) = -\frac{(1-q)L}{q}$. The convexity of ψ implies that $0 < \alpha^* < \hat{\alpha} < 1$.

Let $\alpha(T)$ be the parameter in the optimal test when type T is publicly known to the principal. The first best curve is then defined as $u_c(T) = P(\alpha(T))$. The analysis of the optimal mechanism is divided into three cases under our assumption of set inclusion.

1. $\alpha^* \in [\underline{\alpha}(\underline{T}), \bar{\alpha}(\underline{T})]$. In this case, FDA benefits from choosing α^* for all types, and the first best curve is a constant function. The optimal promised utility satisfies $U(T) = u_c(T)$ for all types $T \in [\underline{T}, \bar{T}]$.

2. $\alpha^* > \bar{\alpha}(\underline{T})$. In this case, the optimal parameter for the FDA is $\alpha(T) = \min\{\bar{\alpha}(T), \alpha^*\}$. The approval probability, and hence the first best curve, is weakly increasing in T. By Theorem 1, the optimal promised utility to the agent coincides with the first best curve.

3. $\alpha^* < \underline{\alpha}(\underline{T})$. In this case, the optimal parameter for the FDA is $\alpha(T) = \max\{\underline{\alpha}(T), \alpha^*\}$. Based on the definition of approval probability P, the first best curve is constantly zero when $\underline{\alpha}(T) > \hat{\alpha}$, then jumps to $P(\hat{\alpha})$, and then gradually decreases in T. By Theorem 1, the optimal promised utility satisfies that there exists $u^* \in [0, P(\hat{\alpha})]$ such that

$$U(T) = \begin{cases} 0 & \underline{\alpha}(T) > \hat{\alpha}, \\ u^* & \underline{\alpha}(T) \le \hat{\alpha}. \end{cases}$$

In this context, the FDA always rejects the drug if the false positive rate the firm can access is always too large. Moreover, the FDA ensures a constant approval

[1]The existence of an interior $\hat{\alpha}$ is guaranteed by the fact that $\psi(0) = 1, \psi(1) = 0$.

probability u^* if the smallest possible false positive rate is below a certain cutoff. When $P(\underline{\alpha}(T)) > u^*$, the FDA adopts the test with parameter $\underline{\alpha}(T)$ and rejects drugs with positive signals with a positive probability to maintain the constant approval probability. When $P(\underline{\alpha}(T)) < u^*$, the FDA adopts a suboptimal test $\tilde{\alpha}$ such that $P(\tilde{\alpha}) = u^*$. The FDA approves the drug according to the test result in this case.

5.2 Action Elicitation

Consider the example of civil servants. In this application, there is a binary state $\theta \in \{-1, 1\}$, and the probability of the state being 1 is $q \in (0, 1)$. The politician receives a private signal s about the state θ such that the posterior is $\mathbb{P}(\theta = 1|s) = s$. Let G be the distribution over signal s. The politician always has two actions: the status quo a_0 and fighting with civil servants a_f. The politician's utility v and the civil servant's utility u under these two actions are

$$\begin{aligned} v(a_0, \theta) &= 0 & u(a_0, \theta) &= 0; \\ v(a_f, \theta) &= -m & u(a_f, \theta) &= -n. \end{aligned}$$

For the purpose of simplifying the exposition, we assume $n \geq 1$. Besides these two actions, there are potentially many actions for reforms a_d^i. A reform a_d^i is indexed by the direction $i \in \{-1, 1\}$ and the radical degree d. The politician prefers to conduct the correct reform but dislikes reforms in the wrong direction:

$$v(a_d^i, \theta) = d \text{ if } i = \theta, \quad v(a_d^i, \theta) = -d \text{ if } i \neq \theta.$$

The civil servants, on the other hand, are relatively biased against the politician in that they hate changes so

$$u(a_d^i, \theta) = d - \alpha d \text{ if } i = \theta, \quad u(a_d^i, \theta) = -d - \alpha d \text{ if } i \neq \theta.$$

Only the civil servant privately knows how to conduct reforms, and the civil servant has a private capability type $T \in [0, 1]$ drawn according to distribution F. A civil servant with type T knows how to conduct a_d^i for any $i = \{-1, 1\}$ and for any $d \in [0, T]$.

Optimal Mechanism The solution depends crucially on the capability of the politicians c_p, where

$$c_p =\in t_0^1 \max\{2s-1, 1-2s\}\,\mathrm{d}G(s)$$

This is the politician's payoff if she can optimally use the most radical reforms according to her own private signals. Since the politician always prefers the most radical reform, the first best curve is

$$u_c(T) = c_p \cdot T - \alpha \cdot T.$$

1. $c_p \geq \alpha$. In this case, the first best curve is $u_c(T)$ increasing in the agent's type by allowing the politician to optimally adopt the most radical reform for all types. By Theorem 1, the optimal promised utility coincides with the first best curve.

2. $c_p \leq \alpha$. In this case, the first best curve is $u_c(T)$ decreasing in the agent's type. By Theorem 1, the optimal promised utility is a constant function. To implement this constant promised utility while maximizing the politician's payoff, there exists $\hat{T} \in [0, 1]$ such that

   - If $T \leq \hat{T}$, the politician commits to adopting the reform a_T^i with the direction i that matches her posterior if and only if her posterior is sufficiently precise. Otherwise, the politician chooses to fight with the civil servants a_f. As T increases, the probability of a_f decreases.
   - If $T \geq \hat{T}$, the politician commits to adopting the reform a_T^i with direction i that matches her posterior if and only if her posterior is sufficiently precise. Otherwise, the politician chooses the status quo a_0. As T increases, the probability of a_0 increases.

   Note that in this case, the probability of reforms is maximized given the middle types of the civil servants.

5.3 Business Strategies with Moral Hazard

Consider the application of CEO performance pay. In this application, the CEO has three possible business strategies in ex ante:

$$\mathcal{G} = \{(G_0, 0), (G_1, c), (G_r, c)\}$$

where G_0 is the shirk option that leads to a reward of 1 with a probability of p_0 and 0 with a probability of $1 - p_0$, G_1 is the safe action that leads to a reward of 1 with a probability of p_1 and 0 with a probability of $1 - p_1$, and G_r is the risky action that leads to a reward of 1 and -4 with probabilities of p_r and a reward of 2 with a probability of $1 - 2p_r$.

The CEO with a private type $T \in [0, 0.5]$ can improve the outcome of the risky action to G_r^s for any $s \leq T$, where G_r^s leads to a reward of 1 and -4 with a probability of p_r^s and 2 with a probability of $1 - 2p_r^s$. To illustrate the interesting trade-offs, we focus on the parameters where $p_0 = 0.1, p_1 = 0.3$ and $p_r^s = 0.5 - s$. The cost of effort is $c = 0.1$.

Optimal Mechanism We first characterize the first best curve. When $T \in [0, 0.2]$, due to the existence of the risky action that leads to low expected rewards, in order to incentivize the CEO to choose the safe action, the firm needs to provide a positive payment to the CEO even when the reward is 0. Whether this is profitable compared to allowing the CEO to shirk depends on T. Specifically, in this example, when $T \in [0, \frac{3}{22}]$, it is too costly to incentivize the safe action, and the CEO receives zero payoff for shirking. When $T \in [\frac{3}{22}, 0.2]$, it is beneficial to incentivize the safe action. The CEO has a payoff jump at $\frac{3}{22}$, and then it gradually decreases as it becomes less costly to incentivize the CEO to choose the safe action.

When $T = 0.25$, the firm now becomes indifferent between incentivizing the safe action and the risky action. Therefore, the CEO's payoff is a constant function for $T \in [0.2, 0.25]$ by incentivizing him to choose the safe action. The payoff jumps down at 0.25, and the firm extracts all the surplus by incentivizing the CEO to choose the risky action. Combining all cases, we have

$$u_c(T) = \begin{cases} 0 & T \in [0, \frac{3}{22}), \\ \frac{5-18T}{20(1+2T)} & T \in [\frac{3}{22}, 0.2), \\ \frac{1}{20} & T \in [0.2, 0.25), \\ 0 & T \in [0.25, 0.5]. \end{cases}$$

By Theorem 1, the optimal promised utility satisfies that there exists $u^* \in [0, 0.1]$ such that

$$U(T) = \begin{cases} 0 & T \in [0, \frac{3}{22}), \\ u^* & T \in [\frac{3}{22}, 0.5]. \end{cases}$$

To implement this promised utility function, the CEO is incentivized to shirk if his type is below $\frac{3}{22}$. If his type is above $\frac{3}{22}$, there exists $\hat{T} \in [\frac{3}{22}, 0.25]$ such that

- If $T \in [\frac{3}{22}, \hat{T})$, the risky proposal is approved with an internal probability, and the approval probability increases in T.
- If $T \in [\hat{T}, 0.5]$, the risky proposal is approved with a probability of 1, and the CEO is overcompensated for exerting costly effort to incentivize them to reveal the type T truthfully.

6 Conclusion

This paper studies a screening problem in which private information concerns *what can be done* rather than *how payoffs vary across types*. An agent privately observes which technologies (or proposals, experiments, reforms, business strategies) are feasible and can disclose only a subset to a principal. Because disclosed options are verifiable and their payoff consequences are publicly known, strategic behavior operates through concealment that shrinks the principal's opportunity set, not through information rents in the Myersonian sense. This distinction is central for organizational and regulatory settings in which principals can evaluate proposals once presented but cannot access options that are withheld.

Our model differs substantially from the conventional screening model. First, even under complete information about feasibility, the agent's private access creates a commitment problem: the principal must make withholding unattractive. The optimal benchmark therefore features severe off-path discipline ("shoot the agent"), implying that private access primarily pins down a minimum utility that the agent must be guaranteed on-path. Second, once private information over feasible sets is introduced (under the nested-set order), incentive compatibility collapses to a simple monotonicity requirement on the agent's promised utility. The optimal mechanism then has a stark structure: the principal either lets the agent's promised utility track the complete information benchmark, or she deliberately compresses rewards by holding the agent's promised utility locally constant.

This bang–bang structure clarifies when and why institutions exhibit reward compression and the under-use of newly disclosed options. When the complete information benchmark is increasing, the principal can fully exploit additional feasibility while still preserving truthful disclosure; incentives are sharp, and differentiation is fine. When the complete information benchmark locally declines with capability, truthful disclosure requires the principal to flatten promised utilities over a region, which in turn forces her to forgo some myopic exploitation of the richer choice set in that region. Importantly, these distortions are not arbitrary: the number of flat regions is bounded by the number of downward-sloping regions of the complete information benchmark. Economically, this means that the complexity of optimal screening is governed by the extent of conflict revealed by the complete-information benchmark, and optimal "bunching" is limited and structured rather than pervasive.

Taken together, the analysis provides a tractable theory of screening for choice sets and a unified economic rationale for reward compression and selective under-exploitation of feasible options in optimal institutional design.

References

Armstrong, M. and Vickers, J. (2010). A model of delegated project choice. *Econometrica*, 78(1):213–244.

Ben-Porath, E., Dekel, E., and Lipman, B. L. (2019). Mechanisms with evidence: Commitment and robustness. *Econometrica*, 87(2):653–675.

Bergemann, D., Gan, T., and Li, Y. (2025). Managing persuasion robustly: The optimality of quota rules.

Bull, J. and Watson, J. (2007). Hard evidence and mechanism design. *Games and Economic Behavior*, 58(1):75–93.

Curello, G. and Sinander, L. (2025). Screening for breakthroughs. *Theoretical Economics*, 20(4):1323–1365.

Deneckere, R. and Severinov, S. (2008). Mechanism design with partial state verifiability. *Games and Economic Behavior*, 64(2):487–513.

Green, J. R. and Laffont, J.-J. (1986). Partially verifiable information and mechanism design. *Review of Economic Studies*, 53(3):447–456.

Guo, Y. and Shmaya, E. (2023). Regret-minimizing project choice. *Econometrica*, 91(5):1567–1593.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Yang, F. (2025). Nested bundling. *American Economic Review*, 115(9):2970–3013.

A Missing Proofs

A.1 Optimal Mechanisms

Proof of Proposition 1. Given any mechanism, the agent's utility is at least $\underline{u}$ by not reporting anything and letting the principal choose from the default technology t_0. Among all the choices that ensures the agent a utility at least $\underline{u}$, the principal's payoff is at most

$$\max_{(u,v)\in C(T), u\geq \underline{u}} v.$$

Note that the "shoot the agent" mechanism attains this upper bound while ensuring truth telling. Therefore, it must be the optimal mechanism. □

Proof of Corollary 1. If $a_1^p(T) \geq \underline{u}$, the "shoot the agent" mechanism optimally chooses $a^p(T)$ when receiving a truthful report of T. The agent's utility is $a_1^p(T)$ in this case. If $a_1^p(T) < \underline{u}$, since the feasible choice is convex, the principal's payoff when promising a utility $z \geq \underline{u} \geq a_1^p(T)$ to the agent is decreasing in z. Therefore, the principal's payoff is maximized by promising a utility $\underline{u}$ to the agent. This choice of promised utility is also feasible due to the set inclusion assumption. Combining the two cases, Corollary 1 holds. □

Proof of Lemma 1. Given any direct mechanism $a(T)$, consider another mechanism $\hat{a}(T)$ such that $\hat{a}_1(T) = a_1(T) = U(T)$, and

$$\begin{aligned} \hat{a}_2(T) = \max_{v} \quad & v \\ \text{s.t.} \quad & (U(T), v) \in \text{conv}(C(T)). \end{aligned}$$

Since $(U(T), a_2(T)) \in \text{conv}(C(T))$ as mechanism $a(T)$ must be feasible, we have $\hat{a}_2(T) = a_2(T)$ for any type T. Moreover, the construction that $\hat{a}_1(T) = a_1(T)$ ensures that mechanism $\hat{a}(T)$ also incentivize the agent to report T truthfully. Therefore, mechanism $\hat{a}(T)$ is an optimal mechanism that takes the form in the statement of Lemma 1. □

Proof of Lemma 2. By definition, $U(T)$ is feasible if and only if there exists v such that $(U(T), v) \in \text{conv}(C(T))$. Since the projection of $\text{conv}(C(T))$ on its first coordinate is a compact and convex interval, letting $\underline{u}_f(T) = \min_{(u,v)\in C(T)} u, \bar{u}_f(T) = \max_{(u,v)\in C(T)} u$, $U(T)$ is feasible if and only if $U(T) \in [\underline{u}_f(T), \bar{u}_f(T)]$. □

Proof of Lemma 3. We first prove the only if direction. Since the promised utility is incentive compatible, for any type $T \geq T'$, due to the set inclusion assumption (Assumption 1), an agent with type T can always misreport as T' to ensure a utility of $U(T')$. Therefore, $U(T)$ must be weakly increasing. The condition that $U(\{t_0\}) \geq \underline{u}$ is implied by the feasibility.

Now we prove the if direction. Since $U(T)$ is weakly increasing, and the agent can only misreport as a lower type. Therefore, misreporting weakly lower the agent's utility, and hence the mechanism is incentive compatible. □

Proof of Lemma 4. We first prove that $U(T) \leq \bar{u}_c(T)$ for all types T. The other direction, where $U(T) \geq \underline{u}_c(T)$, can be proved in the same manner.

Suppose that there exists $\hat{T}$ such that $U(\hat{T}) > \bar{u}_c(\hat{T})$. Consider another promised utility

$$\hat{U}(T) = \min\{U(T), \underline{u}_c(T)\}.$$

By construction, $\hat{U}(T)$ is still weakly increasing and hence can be implemented in an incentive compatible way. Moreover, for any type T, either $\hat{U}(T) = U(T)$, where the principal's payoff is the same in this case, or $U(T) > \hat{U}(T) = \bar{u}_c(T) \geq u_c(T)$. In the latter case, since the feasible choice of the principal given any type T is concave, the optimal payoff of the principal is decreasing in the promised utility of the agent once the promised utility exceeds $u_c(T)$. Moreover, the inequalities are strict under Assumption 2. Therefore, we have

$$V(T, U(T)) \leq V(T, \hat{U}(T))$$

for all types T. This implies that $\hat{U}(T)$ gives a higher expected payoff to the principal than $U(u)$. □

Proof of Theorem 1. Let U^* be an optimal promised utility. By Lemma 4, we may assume without loss that U^* is weakly increasing and lies in the monotone envelope, hence it is feasible for every T.

Fix any type T. Because $\text{conv}(C(T))$ is convex, the function $u \mapsto V(T, u)$ is concave on its feasible domain. Moreover, $u_c(T)$ is the maximizer of $V(T, u)$ among all feasible promised utilities for type T. Therefore,

$$V(T, u) \text{ is decreasing in } u \text{ on } [u_c(T), \infty), \text{and increasing in } u \text{ on } (-\infty, u_c(T)]. \quad (7)$$

Moreover, $u_c(T)$ is unique and the above monotonicity is strict under Assumption 2.

Step 1: Bang–bang form (either u_c or constant on each interval). Partition the type space into maximal (closed) intervals $\{\hat{I}_m\}_{m=1}^M$ such that on each $\hat{I}_m$ exactly one of the following holds:

$$U^*(T) = u_c(T) \ \forall T \in \hat{I}_m; \qquad U^*(T) > u_c(T) \ \forall T \in \hat{I}_m; \qquad U^*(T) < u_c(T) \ \forall T \in \hat{I}_m.$$

(These are maximal in the sense that two adjacent intervals never satisfy the same relation.)

On any interval $\hat{I}_m$ where $U^*(T) = u_c(T)$, nothing needs to be changed.

Case A: $U^* > u_c$ **on** $\hat{I} = [\underline{T}, \bar{T}]$**.** Define a new promise $\tilde{U}$ on $\hat{I}$ by

$$\tilde{U}(T) \;\equiv\; \max\Big\{U^*(\underline{T}),\; \max_{\underline{T}\le s\le T} u_c(s)\Big\}, \qquad T \in [\underline{T}, \bar{T}], \tag{8}$$

and set $\tilde{U}(T) = U^*(T)$ for $T \notin \hat{I}$. Then $\tilde{U}$ is weakly increasing on $\hat{I}$ (it is a running maximum) and matches U^* at the left endpoint. Also, since U^* is weakly increasing and $U^*(s) > u_c(s)$ on $\hat{I}$, we have

$$u_c(T) \;\le\; \tilde{U}(T) \;\le\; U^*(T) \qquad \forall T \in \hat{I}.$$

By (7), lowering the promised utility while staying weakly above $u_c(T)$ increases the principal's payoff pointwise, hence

$$V(T, \tilde{U}(T)) \;\ge\; V(T, U^*(T)) \qquad \forall T \in \hat{I}.$$

Thus, replacing U^* by $\tilde{U}$ on $\hat{I}$ weakly improves the objective and preserves feasibility and monotonicity. Finally, by construction of the running maximum on $\hat{I}$, the function $\tilde{U}$ is either constant or coincides with u_c.

Case B: $U^* < u_c$ **on** $\hat{I} = [\underline{T}, \bar{T}]$**.** Define a new promise $\tilde{U}$ on $\hat{I}$ by

$$\tilde{U}(T) \;\equiv\; \min\Big\{U^*(\bar{T}),\; \min_{T\le s\le \bar{T}} u_c(s)\Big\}, \qquad T \in [\underline{T}, \bar{T}], \tag{9}$$

and set $\tilde{U}(T) = U^*(T)$ for $T \notin \hat{I}$. The map $T \mapsto \min_{T\le s\le \bar{T}} u_c(s)$ is weakly increasing (a running minimum from the right), hence so is $\tilde{U}$. Moreover, since $U^*(s) < u_c(s)$ on $\hat{I}$ and U^* is weakly increasing,

$$U^*(T) \;\le\; \tilde{U}(T) \;\le\; u_c(T) \qquad \forall T \in \hat{I}.$$

By (7), raising the promised utility while staying weakly below $u_c(T)$ increases the principal's payoff pointwise, hence

$$V(T, \tilde{U}(T)) \;\ge\; V(T, U^*(T)) \qquad \forall T \in \hat{I}.$$

Thus replacing U^* by $\tilde{U}$ on $\hat{I}$ weakly improves the objective and preserves feasibility and monotonicity. By construction, on $\hat{I}$ the function $\tilde{U}$ is either constant or coincides with u_c.

Applying the above replacement separately on each strict-inequality interval $\hat{I}_m$ yields an optimal promised utility (still denoted U^*) such that on each maximal interval it is either equal to $u_c(T)$ or constant. Let $\{I_j\}_{j\in J_1\cup J_2}$ be the resulting maximal intervals, where J_1 indexes those with $U^*(T) = u_c(T)$ and J_2 indexes those on which U^* is constant. This proves items (1)–(2) in the theorem.

Moreover, by maximality, between any two distinct J_1-intervals there must be at least one J_2-interval, so $|J_1| - 1 \le |J_2|$.

Step 2: Bounding the number of constant intervals by K. Let $\{D_k\}_{k=1}^K$ be the collection of maximal intervals on which $u_c(\cdot)$ is strictly decreasing. The fact that $|J_2| \le K$ holds immediately by the following two observations.

1. *Every constant interval I_j with $j \in J_2$ intersects some D_k.* Let $I_j = [a, b]$ and $U(T) \equiv \bar{u}$ on I_j. Since I_j is in J_2, there exists $\hat{T} \in (a, b)$ with $\bar{u} \neq u_c(\hat{T})$. Suppose that u_c is weakly increasing. If $\bar{u} < u_c(\hat{T})$, then $u_c(T) > \bar{u}$ for all $T \in [\hat{T}, b]$. Increasing the promised utility in the range of $[\hat{T}, b]$ by a small $\varepsilon > 0$ improves the objective under Assumption 2 according to (7). Similarly, if $\bar{u} > u_c(\hat{T})$, an analogous small decrease in a left-neighborhood improves the principal's payoff. Therefore, u_c cannot be weakly increasing on I_j, so it must strictly decrease somewhere on I_j, and hence I_j intersects with at least one maximal strictly decreasing interval D_k.

2. *Each D_k intersects at most one constant interval I_j with $j \in J_2$.* Suppose instead that some D_k intersects two distinct constant intervals. Since D_k is connected and U is weakly increasing, there exists $\tau \in D_k$ at which U jumps from $\bar{u}_L$ on the left to $\bar{u}_R > \bar{u}_L$ on the right (both constants on neighborhoods contained in D_k). Because u_c is strictly decreasing on D_k, either $\bar{u}_L < u_c(\tau)$ or $\bar{u}_R > u_c(\tau)$ (indeed at least one must hold since $\bar{u}_L < \bar{u}_R$). If $\bar{u}_L < u_c(\tau)$, then for T just left of τ we have $\bar{u}_L < u_c(T)$, so increasing U slightly there (not exceeding $\bar{u}_R$) preserves monotonicity and, by (7), strictly raises the objective under Assumption 2. If $\bar{u}_R > u_c(\tau)$, then for T just right of τ we have $\bar{u}_R > u_c(T)$, so decreasing U slightly there (not below $\bar{u}_L$) preserves monotonicity and strictly raises the objective. Either way we contradict optimality. Hence D_k can intersect at most one constant interval.

Combining the bounds gives $|J_1| - 1 \le |J_2| \le K$, completing the proof. □

A.2 Properties and Comparative Analysis

Proof of Proposition 2. Given any promised utility $U(T)$, let $\{I_j\}_{j \in J_1}$ be the partitions of intervals such that $U(T) = u_c(T)$ for any $j \in J_1$ and $T \in I_j$. Let F be a distribution with support in $\{I_j\}_{j \in J_1}$. By definition, $U(T)$ implements the first best for distribution F. Therefore, $U(T)$ is optimal for distribution F. Moreover, under Assumption 2 where the optimal promised utility is unique for all types of the agent, any deviation of the promised utility for a set of types with positive measure in the support of F leads to a strict utility loss for the principal. Therefore, $U(T)$ is F-a.s. uniquely optimal. □

Proof of Proposition 3. We construct the choice set as

$$C(T) = \{(u,v) : u \in [\min\{\underline{u}, \min_T u_c(T)\}, \max_T u_c(T)], v \in [0, H(T) - \alpha(T) \cdot (u - u_c(T))^2]\}$$

where $\alpha(T)$ is a non-negative function that is to be determined later, and $H(T)$ is the parameter chosen to maintain the set inclusion assumption.

Given any promised utility $U(T)$, let $\{I_j\}_{j \in J_1}$ be the partitions of intervals such that $U(T) = u_c(T)$ for any $j \in J_1$ and $T \in I_j$, and let $\{I_j\}_{j \in J_2}$ be the partitions of intervals such that $U(T)$ is a constant for any $j \in J_2$ and $T \in I_j$ Let $\alpha(T) = 1$ for any type $T \in I_j$ with $j \in J_1$, and let $\alpha(T) = 0$ for any type $T \in I_j$ with $j \in J_2$. That is, the optimal utility of the principal is independent of the promised utility of the agent for any type $T \in I_j$ with $j \in J_2$.[2] Note that in our construction, $U(T)$ implements the first best given the distribution F. Therefore, $U(T)$ is the optimal promised utility. Moreover, in our construction, the promised utility for types $T \in I_j$ with $j \in J_1$ is uniquely pinned down, up to a set with zero measure. The continuity of the complete information curve $u_c(T)$ and the monotonicity of the promised utility function then imply that the optimal promised utilities for types $T \in I_j$ with $j \in J_2$ are pinned down as well. Therefore, $U(T)$ is F-a.s. uniquely optimal. □

Proof of Lemma 5. Let $U(\cdot)$ be any feasible and incentive-compatible promised utility. Define

$$\bar{T}(T) \in \arg\max_{T' \le T} V(T', U(T')),$$

breaking ties by selecting the *largest* maximizer, and define a new promised utility

$$\hat{U}(T) \equiv U(\bar{T}(T)).$$

Because the maximization set $\{T' \le T\}$ expands with T, the tie-breaking rule implies that $\bar{T}(T)$ is weakly increasing, and since $U(\cdot)$ is weakly increasing, $\hat{U}(\cdot)$ is weakly increasing as well.

Feasibility follows from set inclusion: $\bar{T}(T) \le T$ implies $C(\bar{T}(T)) \subseteq C(T)$, hence $[\underline{u}_f(\bar{T}(T)), \bar{u}_f(\bar{T}(T))] \subseteq [\underline{u}_f(T), \bar{u}_f(T)]$. Since $U(\bar{T}(T))$ is feasible for $\bar{T}(T)$, it is also feasible for T. Therefore $\hat{U}$ is feasible and incentive compatible.

Finally, for every T,

$$V(T, \hat{U}(T)) = V(T, U(\bar{T}(T))) \ge V(\bar{T}(T), U(\bar{T}(T))) = \max_{T' \le T} V(T', U(T')),$$

where the inequality uses that for fixed u, $V(T,u)$ is weakly increasing in T because $C(T') \subseteq C(T)$ whenever $T' \le T$. Thus $T \mapsto V(T, \hat{U}(T))$ is the running maximum of $T \mapsto V(T, U(T))$ and hence weakly increasing.

[2]It is also possible to construct parameters $\alpha(T) > 0$ for all types such that Assumption 2 is satisfied, and the candidate $U(T)$ remains F-a.s. uniquely optimal when $U(T)$ is an interior solution. The proof of that construction is more involved and hence omitted here to simplify the exposition.

Taking expectations under F gives $\int V(T, \hat{U}(T))\, dF(T) \geq \int V(T, U(T))\, dF(T)$. Applying this construction to an optimal U yields an optimal U^* with the desired monotonicity. □

Proof of Proposition 4. Let U^* be optimal for the baseline distribution F and satisfy Lemma 5. Define $V^*(T) \equiv V(T, U^*(T))$, which is weakly increasing in T.

Fix the expansion map ϕ in Definition 1. Although the support of $\hat{F}$ need not be nested, we will construct a feasible and incentive-compatible mechanism for $\hat{F}$ that achieves at least $\mathrm{OPT}(F)$.

Given any report S (a set of technologies containing the default), define its *projection onto the baseline chain* as

$$\pi(S) \;\equiv\; \sup\{T : \; T \subseteq S\}.$$

Because baseline types form a chain under Assumption 1, $\pi(\cdot)$ is well-defined and satisfies: (i) if $S \subseteq S'$, then $\pi(S) \leq \pi(S')$; and (ii) $\pi(T) = T$ for baseline types.

Now define a promise rule on *all* reports by

$$\hat{U}(S) \;\equiv\; U^*(\pi(S)).$$

This promise rule is weakly increasing in set inclusion because both $\pi(\cdot)$ and $U^*(\cdot)$ are weakly increasing. Moreover, it is feasible: since $\pi(S) \subseteq S$, any utility feasible for $\pi(S)$ is feasible for S (Lemma 2), and $U^*(\pi(S))$ is feasible for $\pi(S)$ by construction.

Let the principal, upon report S, choose the payoff-maximizing point delivering $\hat{U}(S)$, i.e., implement payoff $V(S, \hat{U}(S))$. This defines a feasible and incentive-compatible mechanism for $\hat{F}$.

Now evaluate the principal's payoff at an expanded type $\hat{T} = \phi(T)$. Since $\phi(T) \supseteq T$, we have $\pi(\hat{T}) \geq T$. Therefore,

$$\begin{aligned} V(\hat{T}, \hat{U}(\hat{T})) = V\big(\hat{T}, U^*(\pi(\hat{T}))\big) \;&\geq\; V\big(\pi(\hat{T}), U^*(\pi(\hat{T}))\big) \\ = V^*(\pi(\hat{T})) \;&\geq\; V^*(T), \end{aligned}$$

where the first inequality uses $\pi(\hat{T}) \subseteq \hat{T}$ and monotonicity of $V(\cdot, u)$ in set inclusion, and the second uses that V^* is weakly increasing and $\pi(\hat{T}) \geq T$.

Taking expectations with respect to $T \sim F$ (so that $\hat{T} = \phi(T) \sim \hat{F}$) yields that the constructed mechanism attains at least $\int V^*(T)\; \mathrm{d}F(T) = \mathrm{OPT}(F)$ under $\hat{F}$. Hence $\mathrm{OPT}(\hat{F}) \geq \mathrm{OPT}(F)$. □

Proof of Proposition 6. Fix any optimal promised utility $U^{t_0}(\cdot)$ and $U^{t_0'}(\cdot)$ for the default technologies t_0 and t_0' respectively. Let $\hat{T}$ be the smallest type such that $U^{t_0'}(\hat{T}) < U^{t_0}(\hat{T})$, and let $\hat{T}_-$ denote the largest type below $\hat{T}$ (so $U^{t_0'}(T) \geq U^{t_0}(T)$ for all $T \leq \hat{T}_-$).[3]

[3] If $\hat{T}$ is the lowest type in the support of F, we define $\hat{T}_-$ as the default technology t_0.

Consider the subproblem of (OPT) restricted to types $T \geq \hat{T}$, taking as given the boundary condition at $\hat{T}_{-}$: promised utilities must satisfy

$$U(T) \ \geq \ U^{t_0}(\hat{T}_{-}) \qquad \forall T \geq \hat{T},$$

and must be weakly increasing. Because the objective in (OPT) is additively separable across types and the only cross-type restriction is monotonicity, the restriction of an optimal solution to the tail $\{T \geq \hat{T}\}$ must solve this continuation problem given the boundary value.

In particular, the restriction of $U^{t_0}(\cdot)$ to $\{T \geq \hat{T}\}$ is optimal for the tail problem with the boundary constraint $U^{t_0}(\hat{T}_{-})$. Since $U^{t_0'}(\hat{T}_{-}) \geq U^{t_0}(\hat{T}_{-})$ and $U^{t_0'}(\cdot)$ is weakly increasing, the tail of $U^{t_0'}(\cdot)$ is feasible for this tail problem. Hence, the tail of $U^{t_0}(\cdot)$ yields a weakly higher principal payoff on $\{T \geq \hat{T}\}$ than the tail of $U^{t_0'}(\cdot)$.

Now define another promised utility

$$\widetilde{U}(T) \ = \ \begin{cases} U^{t_0'}(T), & T \leq \hat{T}_{-}, \\ U^{t_0}(T), & T \geq \hat{T}. \end{cases}$$

By construction and the definition of $\hat{T}$, we have $U^{t_0}(\hat{T}) \geq U^{t_0'}(\hat{T}) \geq U^{t_0'}(\hat{T}_{-})$, so $\widetilde{U}(\cdot)$ is weakly increasing. Moreover, feasibility and the lower-bound constraint are preserved because $\widetilde{U}(\cdot)$ coincides with feasible promised utilities in each region, and $t_0' \subseteq t_0$ implies $\underline{u}(t_0') \geq \underline{u}(t_0)$.

Therefore, $\widetilde{U}(\cdot)$ is feasible and incentive compatible under default t_0', and it yields a weakly higher principal payoff than $U^{t_0'}(\cdot)$. Therefore, $\widetilde{U}(\cdot)$ remains optimal when the default technology is t_0', and the constraint that $\widetilde{U}(T) \ \geq \ U^{t_0}(T)$ is satisfied for all type T. □