A mechanism-design approach to property rights*

Piotr Dworczak[†] Ellen V. Muir[‡]

December 31, 2025

Abstract

We propose a framework for studying the optimal design of *property rights*, broadly defined as rights relating to the control of an economic resource. An agent with private information participates in a trading mechanism that determines the final allocation of the resource. The mechanism may yield outcomes that are misaligned with social preferences and may fail to induce efficient investment by the agent. A designer, who does not directly control the mechanism, can influence its outcome by assigning the agent a menu of rights that determine the agent's *outside options* in the mechanism. We show that the optimal menu requires at most two types of rights. The first is an *option-to-own* that grants the agent control over the resource upon paying a pre-specified price. The second is needed only when the agent's hold-up problem is sufficiently severe, and its form depends on whether investment is contractible.

Keywords: Property rights, mechanism design, outside options, limited commitment, hold-up problem
JEL-Classification: D23, D86, D82

*We thank Sandeep Baliga, Dirk Bergemann, Alessandro Bonatti, Laura Doval, Mehmet Ekmekci, Marina Halac, Bruno Jullien, Elliot Lipnowski, Simon Loertscher, Christoph Luelfesmann, Paul Milgrom, Stephen Morris, Georg Nöldeke, Alessandro Pavan, Patrick Rey, Anna Sanktjohanser, Patrick Schmitz, Ilya Segal, Bruno Strulovici, Quitzé Valenzuela-Stookey, Michael Whinston, and seminar audiences at Berlin, Bonn, Boston College, Columbia, CY Cergy Paris, Harvard, Iowa, LUISS, Mannheim, MIT, PSE, Stanford, UChicago, VSET, Yale, the 2025 Columbia-Duke-MIT-Northwestern IO Theory Conference, EARIE 2025, the 2025 Chicago Booth Market Design Conference, EC 2024, the 2024 Canadian Economic Theory Conference, the 2024 Women in Economic Theory Conference, the 2024 ASSA Annual Meeting, the 20th Annual CSIO/TSE Conference on Industrial Organization, the 2023 Markets, Contracts and Organizations Conference at ANU, and the 2022 Swiss Theory Day for helpful comments and discussions. Julien Manili provided excellent research assistance. Dworczak gratefully acknowledges the support received under the ERC Starting grant IMD-101040122.

[†]Department of Economics, Northwestern University; Group for Research in Applied Economics. Email: piotr.dworczak@northwestern.edu.

[‡]Sloan School of Management, MIT. Email: emuir@mit.edu.

1 Introduction

The assignment of property rights has important implications for the distribution of surplus within society and—in the presence of transaction costs—economic efficiency (Coase, 1960, Williamson, 1979). Consequently, there are trade-offs associated with the design of these rights. Awarding a full property right over an economic resource incentivizes the owner to make efficient investment decisions relating to its use. However, when transaction costs (such as private information) are present, strong property rights may inhibit the future reallocation of economic resources to agents who can utilize them most efficiently. Moreover, the assignment of property rights may give rise to market power or conflict with society's distributive objectives.

One example that highlights the trade-offs involved in the design of property rights is the design of radio spectrum licenses. As the Federal Communications Commission (FCC)—which regulates the use of spectrum in the US—has noted, granting license holders full property rights encourages costly investment in the infrastructure needed to use spectrum efficiently. Yet, because the market-wide reassignment of spectrum involves significant complexity and transaction costs (as illustrated by the FCC's 2016-17 "Incentive Auction"), awarding such strong rights may impede the efficient reallocation of spectrum in response to technological progress.[1] This raises a natural question: How should spectrum licenses—and property rights more generally—be *optimally* designed?

The starting point of our analysis is the observation that property rights differ from other legal contracts in that they do not typically specify the contracting parties; rather, they give the holder unilateral authority over certain uses of the underlying resource. For instance, a full property right gives the owner the *option* to use the resource or generate income from it, regardless of what other opportunities might arise through interaction with other economic agents (such as selling the resource).[2] We therefore model property rights as determining the holder's *outside options* in economic interactions. This simplified perspective allows us to use mechanism-design techniques to characterize optimal property rights.

Our framework highlights the key trade-offs involved in the design of property rights by recasting this problem as a dynamic contracting problem between a *designer*, a *principal* and an *agent*. The designer first determines the agent's property rights: a flexible menu of

[1] The current design of spectrum licenses effectively grants holders full property rights for a fixed term. As Milgrom et al. (2017) observe: "Existing license designs present regulators with a stark choice between encouraging entry and innovation or ensuring that licensees' complementary, long-term investments are secure."

[2] In the legal nomenclature going back to Hohfeld (1917), property rights are *in rem* rights—"good against the world"—as opposed to *in personam* rights that typically arise from specific bilateral (or multilateral) agreements (see also Balbuzanov and Kotowski, 2019).

outside options relating to the control of an economic resource that are available to the agent in subsequent interactions. The agent then makes an investment decision that affects her valuation for the resource. We model investment as a binary choice for the agent: If she pays a cost, her value is drawn from a distribution that first-order stochastically dominates the default distribution. Following the agent's investment decision, a payoff-relevant public state is realized. The principal then chooses a trading mechanism (with transfers) that screens the agent's private information and determines the final allocation. The mechanism chosen by the principal must respect the agent's rights, i.e., it must ensure the agent's participation given the outside options created by her property rights.

A central assumption in our framework is that the designer does not directly control the trading mechanism. Conditional on the realization of the public state, the mechanism is chosen in a sequentially rational way by the principal who maximizes an objective function that need not align with the designer's preferences. We refer to this as the *non-commitment* friction. (In some applications, the principal could represent the "future self" of the designer, in which case non-commitment can be understood as a form of time inconsistency.) Non-commitment also gives rise to a *hold-up problem*: The agent may fail to undertake efficient investments if the principal's mechanism appropriates the resulting surplus.

In our framework, the designer indirectly influences the outcome of the trading mechanism by choosing the agent's property rights. For example, the designer may award the agent a conventional property right that gives her full control over the resource, thereby guaranteeing the agent the *option* to keep the resource regardless of the principal's objective and the realized state. Other designs might give the agent conditional rights, such as an option to demand a monetary payment from the principal in exchange for relinquishing control over the resource, or an option to acquire control over the resource upon paying a pre-specified price. By varying the strength and form of these rights, the designer can shape the agent's investment incentives and the principal's flexibility in choosing the mechanism, thereby mitigating the non-commitment friction and the agent's hold-up problem.

Before giving an overview of our results, we emphasize two simplifying assumptions implicitly made in our framework. First, by studying a setting with a single agent, we abstract away from the problem of who (among multiple agents) should be allocated the property right to the resource. In the language of mechanism design, we depart from the traditional focus on how to allocate a given good to agents differing in their values, and instead focus on the problem of *designing the good itself*—here, understood as designing the set of rights to the underlying economic resource. Second, we abstract from the potential impact of property rights on the distribution of bargaining power. That is, we assume that it is always the principal who chooses the trading mechanism, even if the agent holds full property rights. This is in line with our focus on modeling property rights as determining outside options of the holder. This assumption restricts the set of applications of our framework, but is appropriate in settings in which the principal represents a government or a market regulator.

To isolate the roles of the two main frictions in our framework, we first consider a simplified model without investment. In this case, the designer assigns the agent non-trivial property rights purely for the purpose of better aligning the outcome of the principal's chosen trading mechanism with the designer's preferences. The optimal property right always takes the form of an *option-to-own*. An option-to-own gives the agent full rights to the resource conditional on paying a pre-specified price to the principal. The designer can vary the strength of the option-to-own by adjusting this price. For example, setting the price to zero leads to a full property right, while setting a sufficiently high price is equivalent to giving no right to the agent. The optimal property right thus takes a simple form, while at the same time offering greater flexibility than a conventional property right.

More generally, for the full version of the model involving investment, we show that the optimal property right endows the agent with a menu of at most *two* types of rights. One of the rights takes the form of an option-to-own. The second right is introduced to strengthen investment incentives, and is required only if the hold-up problem is sufficiently severe. The form of this second right depends on whether the designer can make the agent's rights contingent on investment. If investment is observable (and contractible), the second right takes the form of a cash payment for undertaking the investment. If investment is not observable (and hence non-contractible), the second right takes the form of a partial property right that awards the agent control over a fraction of the resource (or, equivalently, gives the resource to the agent with some probability).

Methodologically, in our framework, property rights give rise to a set of outside options that are available to the agent when she interacts with the principal. The principal thus solves an instance of a mechanism design problem with type-dependent outside options, as in the work of Lewis and Sappington (1989), Maggi and Rodriguez-Clare (1995), and Jullien (2000). Assuming payoffs are linear in the allocations, we derive a novel solution technique for such problems based on an extension of the classical ironing technique due to Myerson (1981). In contrast to these earlier papers, the focus of our paper is then *designing* these outside options. Specifically, the designer's problem is to choose the optimal type-dependent reservation utility function for an agent who subsequently participates in a screening mechanism. We characterize solutions to this problem by exploiting the linear dependence of the principal's optimal mechanism on the agent's outside option function that the ironing procedure uncovers. These techniques are portable to other settings involving type-dependent outside options and may thereby be useful beyond the analysis of optimal property rights.[3]

For the main application of our results, we consider a dynamic resource allocation problem—such as the assignment of electromagnetic spectrum—in which a regulator cannot commit to future trading mechanisms (e.g., spectrum auctions) but can design the *resource use license* (e.g., spectrum license). The regulator trades off incentives for the license holder to undertake value-increasing investments against the ease with which control over the resource can be reassigned in the future if new efficient uses of the resource emerge. We find that the optimal license includes a provision that can be interpreted as a *renewable lease*, giving the license holder the opportunity to retain control conditional on paying a pre-specified price. We further illustrate the flexibility of our framework by briefly discussing the implications of our results in other applications, including regulating a private rental market, the design of patent policies, public procurement, and private contracting between firms.

The remainder of this paper is organized as follows. We provide an overview of the related literature in Section 1.1. Section 2 introduces the simplified version of the model without investment (and without uncertainty about the public state). Section 3 contains a derivation of the main result (Theorem 1) in the simple model. Then, in Section 4, we introduce the general model and extend the main result. Section 5 discusses applications. We conclude with a discussion of future research directions in Section 6.

1.1 Related literature

Building on the seminal contributions of Coase (1960) and Williamson (1979), the property-rights literature has largely focused on two types of transaction costs: private information and hold-up problems.[4] Option-to-own contracts have been proposed as a potential remedy for both frictions. However, to the best of our knowledge, we are the first to show that options-to-own arise as part of an *optimal* (second-best) solution when property rights can be chosen from a large non-parametric class. This flexible modeling approach resonates with the legal literature, which recognizes a wide range of property rights beyond the simple, unconditional ownership structures typically modeled in economics.[5]

Myerson and Satterthwaite (1983) first showed that there may be no bargaining procedure that results in efficient outcomes when contracting parties possess private information, and property rights are assigned exclusively to a single party. Cramton, Gibbons, and Klemperer (1987) further clarified that efficiency may be attainable if the parties have suffi-

[3] Baron et al. (2025) and Valenzuela-Stookey (2025) extend our ironing technique to analyze the optimal design of a task-allocation system in the presence of, respectively, a status-quo mechanism and workers' idiosyncratic preferences over tasks.

[4] For a comprehensive survey of this literature, see Segal and Whinston (2012).

[5] See Calabresi and Melamed (1972) and the related discussion in Segal and Whinston (2016).

ciently balanced ownership shares, and Schmitz (2002) demonstrated that an appropriately chosen default option (specified ex ante) can sometimes restore first-best efficiency under two-sided private information.[6] Most closely related to our work are papers that analyze the second-best design of property rights in this context. Che (2006) showed that options-to-own decrease the subsidy needed to implement the first-best outcome, while Segal and Whinston (2016) unified much of this literature by studying the subsidy-minimizing option-to-own that maximizes surplus subject to budget balance. Relative to this earlier work, in our setting the reasons for the non-implementability of the first-best outcome are fundamentally different. In the literature on efficient bargaining, impossibility results from the interaction of two-sided private information with the budget constraint. In our framework, the first best (which need not be allocative efficiency) is prevented by the non-commitment friction—the fact that the designer does not directly control the trading mechanism (and that the mechanism often features distortions due to agent's private information).[7] This approach allows us to fully characterize the second-best outcome (which turns out to be an option-to-own in the model without investment) and consider additional applications (such as resource use license design).

The incomplete-contracts literature—initiated by the seminal work of Grossman and Hart (1986) and Hart and Moore (1990)—focused instead on frictions due to relationship-specific investments that must be taken prior to trading, without the possibility of signing complete contracts. Several solutions to the resulting hold-up problem have been proposed in the literature. Aghion et al. (1994) argued that investment efficiency can be recovered by allowing for contracts that make appropriate provisions regarding renegotiation. The beneficial role of options-to-own have been investigated by Hart (1995) who showed that a price contract can improve upon a simple ownership structure, and by Nöldeke and Schmidt (1995, 1998) who identified settings in which options-to-own can restore first-best levels of investment. Relative to this literature, we shift the focus from the problem of optimal *reallocation* of residual control rights among multiple parties—which is central for classical questions such as the boundaries of the firm—to the problem of a designer who *designs* optimal property rights for a single agent—a case more relevant for market-design applications. This perspective enables us to characterize optimal rights even though the first best is typically not implementable in our setting (partly because we endow the agent with private information at the trading

[6]See also Jehiel and Pauzner (2006), who extend the analysis of Cramton, Gibbons, and Klemperer (1987) by allowing for interdependent values.

[7]In the organizational economics literature, closely related issues are discussed under the heading of *ex post governance problems*. Such problems arise in settings where ex ante contracting is possible but parties cannot commit to ex post actions, so outcomes are determined by authority, control, or governance structures rather than by an ex ante efficient plan (see, for example, Gibbons, 2005).

stage).[8] We find that options-to-own continue to play an important role even when the designer can choose property rights from a non-parametric set. At the same time, the optimal property right may become more complicated—additionally featuring a monetary transfer or a partial property right that grants control over a fraction of the resource—due to the interaction of the hold-up problem with private information at the trading stage. We further elaborate on the relationship between our results and the insights from the incomplete-contracts literature in Section 4.

The problem of efficient investment has also been studied within the more traditional mechanism-design literature. In particular, Rogerson (1992) showed that the Vickrey-Clarke-Groves (VCG) mechanism ensures efficient pre-mechanism investments because it makes participants internalize the social gains from changes in their private valuations. In contrast to these papers, our designer cannot directly control the mechanism—the mechanism is chosen by the principal (whose preferences may differ from those of designer) in a sequentially rational way. Instead, the designer in our model affects investment incentives indirectly by endowing the agent with property rights. Rogerson's insight arises as a special case of our framework: When both the principal and the designer want to maximize efficiency (and investment only affects the private value of the resource), it is optimal to allocate no rights to the agent. Moreover, in special cases of our model, the designer may be able to "force" the principal to implement a VCG mechanism by using an option-to-own with a price equal to the (social) opportunity cost of the resource.

One interpretation of our framework is that the designer is a regulator who—by awarding property rights to the agent—restricts the set of implementable mechanisms. This perspective relates our work to papers studying optimal design of a contractual environment between private parties. For instance, Hermalin and Katz (1993) asked whether courts can improve outcomes by modifying or refusing to enforce contracts between private parties and found that such interventions are rarely beneficial (since private contracting is typically efficient in their framework). More recently, Bhaskar et al. (2023), Hitzig and Niswonger (2023) and Thereze and Vaidya (2025) consider similar questions in delegation models, focusing on how a regulator can achieve various policy goals by restricting the menu of contracts available to private parties. More closely related to our paper is Lüelfesmann (2025), who considers a principal-agent model in which an agent with private information takes an action that affects the principal's payoff, and determines whether delegating full authority over the agent's

[8] Schmitz (2006) introduced private information into the canonical Grossman–Hart–Moore framework and showed that classical insights may fail in this case. Matouschek (2004) and Baliga and Sjöström (2018) also allow for private information at the contracting stage but abstract from investment. As in our model, property rights in Baliga and Sjöström (2018) generate type-dependent outside options, though they focus on parameter ranges where the first best is implementable.

action to the principal or the agent yields higher welfare.

Additional discussions of the related literature can be found in Section 5 (and Appendix B) in the context of applications. Most notably, we discuss how our conclusions concerning optimal license design relate to the policy proposals of Weyl and Zhang (2022).

2 Simple framework

We begin with a simplified framework that omits some of the economic forces emphasized in the introduction—most notably investment. We also abstract away from any uncertainty about the public state at the time of assigning property rights, which makes property rights more powerful than in the general framework that allows for contractual incompleteness. This stripped-down setting makes the fundamental economic drivers of our results and the main technical ideas more transparent (Section 3). We then extend the analysis to incorporate investment (Section 4), completing our description of the economic trade-offs involved in designing property rights.

Overview. We introduce a third party—a *designer*—into an otherwise standard contracting problem between a principal and an agent. The interaction between the principal and the agent determines who controls an underlying economic resource, and the designer has preferences over the outcome of this interaction. We assume that the designer cannot directly regulate the interaction between the principal and the agent. Instead, she can influence its outcome by assigning the agent a menu of rights that governs the agent's outside options with respect to control of the resource. After these property rights are assigned, the agent's private type is realized. The principal then selects a trading mechanism that is sequentially rational, given the rights conferred on the agent.

We now describe the components of the model in greater detail. We begin by introducing the principal-agent problem in the absence of any rights assigned by the designer.

Principal-agent problem. The agent's private type $\theta \in \Theta := [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$ is distributed according to a cumulative distribution function (CDF) F. We assume that F admits an absolutely continuous density f on Θ. The agent's utility is linear in the allocation x (interpreted as either a probability or a quantity) and the transfer t, with the type θ normalized to represent the agent's marginal (and possibly negative) value for the allocation. An agent of type θ who receives the allocation $x \in [0, 1]$ and makes a payment $t \in \mathbb{R}$ then obtains utility $\theta x - t$.

The principal chooses a trading mechanism, which—by the revelation principle—we can take to be a direct revelation mechanism satisfying appropriate incentive-compatibility and individual-rationality constraints. Formally, the principal chooses a mechanism $\langle x(\theta), t(\theta) \rangle$, where $x : \Theta \to [0, 1]$ denotes the allocation rule and $t : \Theta \to \mathbb{R}$ denotes the transfer rule.

Given a direct mechanism $\langle x(\theta), t(\theta)\rangle$, incentive-compatibility requires that,

$$U(\theta) \coloneqq \theta x(\theta) - t(\theta) \geq \theta x(\theta') - t(\theta'), \quad \forall \theta, \theta' \in \Theta. \tag{IC}$$

We normalize the agent's outside option (in the absence of any property rights) to 0. Individual rationality then requires that

$$U(\theta) \geq 0, \quad \forall \theta \in \Theta. \tag{IR$_\emptyset$}$$

The principal maximizes the objective function $\int_\Theta [V(\theta)x(\theta) + \alpha t(\theta)]\, dF(\theta)$, where $V : \Theta \to \mathbb{R}$ is upper semi-continuous in θ, and $\alpha > 0$ is the weight that the principal places on revenue.

We introduce a running example that we will use throughout for illustration.

Running Example. The principal is a seller of an indivisible good. The agent is a buyer, and θ represents the buyer's value for the good. The seller has an opportunity cost $c \in (\underline{\theta}, \bar{\theta})$ and maximizes revenue: $V(\theta) = -c$, for all $\theta \in \Theta$, and $\alpha = 1$.

Menu of rights. Before the principal and the agent interact (and before the agent's type is realized), the designer chooses a menu of rights M that the agent has access to when interacting with the principal. Specifically, we allow for any menu of the form

$$M = \{(x_i, t_i)\}_{i \in I},$$

where $x_i \in [0, 1]$ denotes an allocation, $t_i \in \mathbb{R}$ denotes a payment made by the agent to the principal, and I is an arbitrary index set. We assume that M is a compact subset of $[0, 1] \times \mathbb{R}$. Any right in the menu M can be executed by the agent in lieu of participating in the trading mechanism chosen by the principal. Consequently, each $(x_i, t_i) \in M$ constitutes an outside option available to the agent when contracting with the principal.

The principal is able to replicate any outcome in which the agent executes a right from M within her mechanism. Consequently, it is without loss of generality to restrict the principal to mechanisms that ensure full participation. When the agent holds a menu of rights M, the constraint (IR$_\emptyset$) then becomes

$$U(\theta) \geq \max\{0, \max_{i \in I} \{\theta x_i - t_i\}\}, \quad \forall \theta \in \Theta. \tag{IR$_M$}$$

Thus, given the set of rights M, the principal solves

$$\max_{\langle x(\theta),\, t(\theta)\rangle} \int_\Theta [V(\theta)x(\theta) + \alpha t(\theta)]\, dF(\theta) \quad \text{s.t.} \quad \text{(IC)}, \quad (\text{IR}_M). \tag{P}$$

We let $\langle x^\star(\theta\,;\, M), t^\star(\theta\,;\, M)\rangle$ denote the optimal mechanism for the principal when the participation constraint (IR_M) is induced by menu M.

Designer's problem. The designer's problem is then

$$\max_M \int_\Theta [V^\star(\theta)x^\star(\theta;\, M) + \alpha^\star t^\star(\theta;\, M)]\, dF(\theta), \tag{D}$$

where $V^\star : \Theta \to \mathbb{R}$ is continuous in θ, and $\alpha^\star \geq 0$ is the weight that the designer places on transferring a unit of money from the agent to the principal.[9] In the (non-generic) case where there are multiple optimal mechanisms for the principal, we take the standard approach of assuming that the tie is broken in the designer's favor. However, our results continue to hold under a large class of tie-breaking rules, including designer-inferior selection, as explained in Appendix A.4.

Running Example. The designer maximizes a convex combination of the seller's and the buyer's payoffs, where $\lambda \in [0, 1]$ is the weight on the buyer's payoff. That is, $V^\star(\theta) = \lambda\theta - c$ and $\alpha^\star = 1 - \lambda$. When $\lambda = 0$, the designer's and the principal's interests are fully aligned. When $\lambda = 1$, the designer would like to implement the efficient allocation.

2.1 Discussion

Property rights. Modeling the rights held by the agent in terms of a menu $M = \{(x_i, t_i)\}_{i \in I}$ yields a flexible framework that includes a rich set of possibilities:

- $M = \{(1, 0)\}$ captures a conventional (unconditional) property right: The agent holds residual rights of control over the resource and can select the $x = 1$ allocation at no cost (while being free to relinquish control if offered sufficient monetary compensation).
- $M = \{(0, -p)\}$ captures a right where the agent can demand a monetary transfer $p \in \mathbb{R}_+$ from the principal.

[9] We assume that the weight $\alpha^\star$ is non-negative. Otherwise, the designer would want to force the principal (through the menu M) to give an infinite transfer to the agent. One way to accommodate the case $\alpha^\star < 0$ is to impose a bound on transfers t in the menu M—see the supporting calculations in Online Appendix B.3 for details.

- $M = \{(1,0),(0,-p)\}$ captures a standard property right for the agent along with a *resale option* that allows the agent to sell the resource back to the principal at a price $p \in \mathbb{R}_+$.

- $M = \{(1, p)\}$ represents an option-to-own, giving the agent the right to acquire control over the resource conditional on paying the principal a fixed price $p \in \mathbb{R}_+$.

- $M = \{(y,0)\}$ with $y \in (0,1)$ captures a "partial property right." Depending on the application, y could represent the probability of allocation, a limited quantity covered by the right, or even—in a reduced-form way—geographic or temporal restrictions on the property right.[10]

- $M = \{(r, p(r))\}_{r\in[0,1]}$, for some function $p : (0,1) \to \mathbb{R}_+$, is a flexible menu that allows the agent to purchase their preferred partial property right $r \in [0,1]$ at a price $p(r)$.

Model Frictions. Our framework features two fundamental frictions that prevent the designer (who can be thought of as representing social preferences) from achieving the first-best outcome.

1. Non-commitment: The designer cannot fully commit the principal to her preferred mechanism, which matters whenever the designer's and principal's preferences are misaligned.[11]

2. Private information: Because of the agent's private information (and the resulting incentive-compatibility constraints), the first-best allocation may not be implementable at the contracting stage.

The central challenge that property rights address is the non-commitment friction: If the designer could commit the principal to her preferred mechanism, there would be no need to assign any property rights. The full commitment benchmark therefore provides an upper bound on the designer's payoff.[12] Private information introduces a gap between the full

[10] If y represents a probability, then the right can be implemented by conditioning ownership on some exogenous future event such as a court decision. For example, there is variation across jurisdictions in the degree of protection for intellectual property rights that determines the ex-ante probability of retaining *de facto* ownership.

[11] In the extended model with investment, this friction matters even with aligned preferences due to the agent's associated hold-up problem.

[12] Even in the simplified framework, the optimal assignment of property rights will generally fail to achieve the full commitment benchmark. To see this, set $V^{\star}(\theta) = -\infty$ in our running example—so that the agent's use of the good creates a massive negative externality—and notice that the principal—who does not internalize this externality—cannot be prevented from selling to the agent for some realizations of θ.

commitment benchmark and the first-best outcome. Absent private information, a familiar "Coasian irrelevance" result applies: Property rights can affect the distribution of surplus between the principal and agent but not the allocation of the good, making the design problem trivial (at least in the absence of investment). Private information also affects how close optimally assigned property rights can come to achieving the full commitment benchmark. Figure 1 summarizes the role of these two frictions in our simple framework.

Figure 1: The non-commitment friction separates the full commitment benchmark (C) from the no-rights outcome (A), while private information prevents optimal rights (B) from achieving the full commitment solution. Private information also ensures there is a gap between the full commitment (C) and first-best (D) benchmarks.

Running Example. When $\lambda = 0$, the designer wants to maximize the principal's payoff, so the non-commitment friction is turned off. It is optimal for the designer to offer no rights to the buyer in this case. When $\lambda = 1$, the designer would like to commit the principal to implementing the VCG mechanism, which achieves full efficiency. She can achieve that goal by assigning to the buyer an option-to-own with price p equal to the seller's opportunity cost c, so in this case the private-information friction is effectively turned off.[13] For $\lambda \in (0, 1)$, both frictions are active and solving the designer's problem and determining the optimal property right is non-trivial.

Our explicit modeling of the non-commitment friction distinguishes our framework both from the incomplete-contracts literature and the literature studying the role of property rights in addressing informational frictions. The non-commitment friction could arise in two types of environments. First, the principal and the agent could represent private parties acting in a market. The designer has the ability to regulate the market, assigning property rights in order to steer market outcomes toward the social optimum without having the ability to fully dictate the terms of every transaction. Second, the principal could represent a "future self" of the designer in a setting where the designer suffers from a time inconsistency problem. In this case, assigning property rights partially restores the designer's commitment to future trading mechanisms.

[13]If the buyer has an option-to-own at price c, then selling the good at a price equal to c is optimal for the seller.

Our modeling of the trading stage differs from the typical incomplete-contracts framework: We assume that the agent has private information and that there is a principal who chooses an incentive-compatible mechanism with transfers. This means that our framework assumes a separation between the notion of property rights and bargaining power: The principal enjoys full bargaining power—in the sense that she chooses the trading mechanism—regardless of the rights M held by the agent. However, as long as the principal attaches a positive weight α to revenue (which we have assumed), the choice of M does affect the eventual split of surplus between the agent and the principal. Property rights would be economically ineffective if $\alpha = 0$, as the principal would then simply "buy out" any rights in M with a sufficiently large cash payment.

3 Analysis and results

We can now state the main result of the paper that characterizes the structure of the optimal menu of rights $M^\star$ chosen by the designer.

Theorem 1. *There exists an optimal menu such that* $M^\star = \{(1, p)\}$ *for some* $p \in \mathbb{R}$.

The optimal menu $M^\star$ takes a simple and economically interpretable form. The menu consists of an *option-to-own* $(1, p)$ which gives the agent the right to control the resource by paying a pre-specified price p. By varying the price p, the designer can continuously adjust the strength of the agent's property right. When $p = 0$, the option-to-own becomes a conventional property right. When $p \geq \bar{\theta}$, the agent never finds it optimal to execute the right, so the option-to-own becomes equivalent to no right.

Running Example. By Theorem 1, the designer should assign to the buyer an option-to-own with some price p_λ, for any $\lambda \in [0, 1]$. It is straightforward to show that p_λ is decreasing in the weight λ attached to the agent's payoff, with $p_0 = \bar{\theta}$ (equivalent to no right) and $p_1 = c$ (equivalent to inducing a VCG mechanism), and $\bar{\theta} > p_\lambda > c$, for any $\lambda \in (0, 1)$. Giving a full property right is never optimal for the designer.[14]

In the remainder of this section, we sketch the proof of Theorem 1 (proofs of several technical steps are relegated to Appendix A). The proof overview casts some light on how the parameter p that characterizes the optimal menu is pinned down by the primitives of the model. We will further explore the economic implications of our characterization in

[14] As the proof of Theorem 1 reveals, if the agent held a full property right, the principal would offer to buy back the good from the agent at a price $\underline{p}$ lower than c, which would result in inefficient allocation whenever $\theta \in (\underline{p}, c)$.

Subsection 4.4 after introducing the general framework and in Section 5, where we study an application based on our running example.

3.1 Proof of Theorem 1

We proceed backwards, by first solving the principal's problem, and then solving the designer's problem.

Step 1: Formulating the principal's problem

Fix an arbitrary menu of rights M. We reformulate the principal's problem by expressing the consequences of menu M as a type-dependent outside option.

Lemma 1. *A choice of menu M by the designer is equivalent to choosing an outside option function $R : \Theta \to \mathbb{R}$ for the agent in the trading mechanism, where R is non-negative, non-decreasing and convex, with a right derivative that is bounded above by* 1*.*

Lemma 1 shows that the principal's problem reduces to optimization over the set of type-dependent outside option functions R. The proof follows from the observation that given a menu $M = \{(x_i, t_i)\}_{i\in I}$, we can set

$$R(\theta) = \max\{0, \max_{i\in I}\{x_i\theta - t_i\}\}.$$

Applying the envelope theorem shows that a direct mechanism $\langle x(\theta), t(\theta)\rangle$ chosen by the principal is incentive-compatible if and only if x is a non-decreasing function and, for any $\theta \in \Theta$, the agent's utility under truthful reporting is given by

$$U(\theta) = \underline{u} + \int_{\underline{\theta}}^{\theta} x(\tau)\, d\tau, \tag{1}$$

where $\underline{u} \in \mathbb{R}$ denotes the utility of the lowest type $\underline{\theta}$. This implies that U is a convex function with $U'(\theta) = x(\theta)$ almost everywhere. After standard transformations, this yields

$$\int_{\Theta} [V(\theta)x(\theta) + \alpha t(\theta)]\, dF(\theta) = \int_{\Theta} \left[V(\theta) + \alpha J^B(\theta)\right] x(\theta) dF(\theta) - \alpha\underline{u},$$

where $J^B(\theta) \coloneqq \theta - (1-F(\theta))/f(\theta)$ is the virtual value function. Combining this with Lemma

1, the principal's problem (P) can be rewritten as

$$\max_{x(\theta),\, \underline{u}\geq 0} \int_{\underline{\theta}}^{\overline{\theta}} W(\theta)x(\theta)d\theta - \alpha\underline{u} \tag{P$'$}$$
$$\text{s.t.} \quad x \text{ is non-decreasing, and } U(\theta) = \underline{u} + \int_{\underline{\theta}}^{\theta} x(\tau)\, d\tau \geq R(\theta), \quad \forall \theta \in \Theta,$$

where $W(\theta) := \big(V(\theta) + \alpha J^B(\theta)\big) f(\theta)$. We will refer to the constraint $U(\theta) \geq R(\theta)$ as the *outside option constraint*.

Step 2: Solving the principal's problem

Problems of the form (P$'$) have been analyzed in the literature, most notably by Jullien (2000), who uses weak duality to derive a solution under additional monotonicity assumptions. We develop a new method to solve problem (P$'$) that is based on an appropriate generalization of the ironing procedure of Myerson (1981). For the case of linear utilities that we study, our method is simpler, in that it does not require "guessing" the correct Lagrange multiplier, and more powerful, in that it does not require additional regularity assumptions. To emphasize the portability of the method to other applications involving type-dependent outside options, we solve problem (P$'$) for any outside option function R that satisfies the homogeneity assumption of Jullien (2000) and is such that $R(\theta) = u_0 + \int_{\underline{\theta}}^{\theta} x_0(\tau)d\tau$ for some $u_0 \geq 0$ and non-decreasing allocation rule $x_0 : \Theta \to [0, 1]$.

The following "ironing procedure" allows us to construct a solution to problem (P$'$). First, for all $\theta \in \Theta$, we define

$$\mathcal{W}(\theta) := \int_{\theta}^{\overline{\theta}} W(\tau)\, d\tau \quad \text{ and } \quad \overline{\mathcal{W}} := \text{co}(\mathcal{W}),$$

where co is an operator that returns the concave closure of a given function. Next, we define

$$\underline{\theta}^\star := \sup\{\{\theta \in \Theta : \overline{\mathcal{W}}'(\theta) \geq \alpha\} \cup \{\underline{\theta}\}\},$$
$$\overline{\theta}^\star := \inf\{\{\theta \in \Theta : \overline{\mathcal{W}}'(\theta) \leq 0\} \cup \{\overline{\theta}\}\}.$$

These definitions are illustrated in Figure 2. Informally, $\underline{\theta}^\star$ is the type at which the slope of $\overline{\mathcal{W}}$ is equal to α (or the lowest type $\underline{\theta}$ if the slope is always below α). Similarly, $\overline{\theta}^\star$ is the type at which the slope of $\overline{\mathcal{W}}$ is equal to 0 (or the highest type $\overline{\theta}$ if the slope is always above 0). Equivalently, $\overline{\theta}^\star$ is a global maximizer of $\overline{\mathcal{W}}$. The formal definitions handle the possibility that multiple types may satisfy these conditions and the fact that $\overline{\mathcal{W}}$ may be non-differentiable at some (countably many) points. Because $\overline{\mathcal{W}}$ is concave, we have $\underline{\theta}^\star \le \overline{\theta}^\star$.

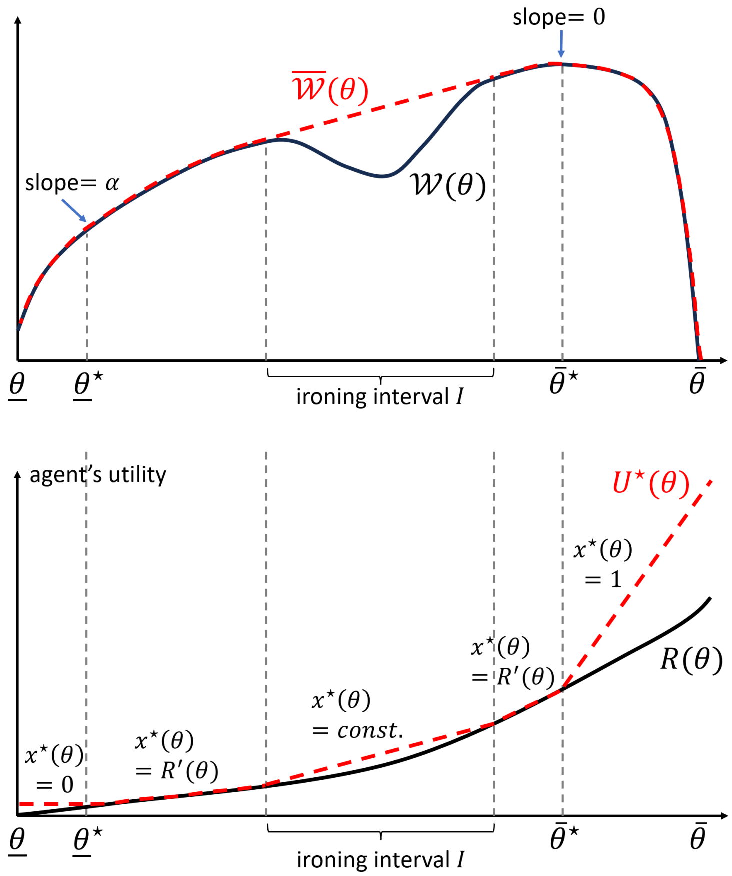

Figure 2: An illustration of the ironing procedure (top panel) and the mapping from the ironing procedure to the optimal indirect utility function U^* (bottom panel).

Let $\mathcal{I}$ be the (at most countable) collection of maximal open intervals (a, b) within $(\underline{\theta}^\star, \overline{\theta}^\star)$ with the property that $\mathcal{W}$ lies strictly below $\overline{\mathcal{W}}$ on (a, b).[15] (In Figure 2, there is a single such interval.) Let $\mathcal{I}^c$ be the complement collection of maximal (relatively) closed intervals $[a, b]$ within $(\underline{\theta}^\star, \overline{\theta}^\star)$ with the property that $\mathcal{W}$ coincides with $\overline{\mathcal{W}}$ on $[a, b]$. Intuitively, the allocation rule must be "ironed" on each $(a, b) \in \mathcal{I}$. Formally, we define

$$\underline{u}^\star = R(\underline{\theta}^\star) \text{ and } x^\star(\theta) = \begin{cases} 0 & \theta \le \underline{\theta}^\star, \\ \frac{\int_a^b R'(\tau)d\tau}{b-a} & \theta \in (a, b) \text{ for some } (a, b) \in \mathcal{I}, \\ R'(\theta) & \theta \in [a, b] \text{ for some } [a, b] \in \mathcal{I}^c, \\ 1 & \theta \ge \overline{\theta}^\star. \end{cases} \tag{2}$$

[15]By maximality we mean that such an interval (a, b) cannot be contained in another interval (a', b') with the same property.

The allocation rule $x^\star$ is equal to 0 below $\underline{\theta}^\star$ and 1 above $\overline{\theta}^\star$. By the choice of the payment $\underline{u}^\star$, the outside option constraint binds at $\theta = \underline{\theta}^\star$. Then, within the interval $[\underline{\theta}^\star, \overline{\theta}^\star]$, $x^\star$ coincides with $R'(\theta)$ on "non-ironing intervals" (the outside option constraint holds with equality everywhere in such intervals), and is constant on "ironing intervals" (the outside option constraint holds with equality only at the endpoints of such intervals). Figure 2 illustrates with an example.

The following lemma states that the ironing procedure defined above characterizes the solution to the principal's problem.

Lemma 2. *The pair* $(x^\star, \underline{u}^\star)$ *as defined in* (2) *solves problem* (P′).

For illustration and intuition, consider first the simplest case in which the objective function W is non-decreasing. In this case, $\mathcal{W}$ is concave, and hence $\mathcal{W} = \overline{\mathcal{W}}$. Thus, $\mathcal{I} = \emptyset$, and ironing is not needed. Furthermore, $\underline{\theta}^\star$ is defined by $W(\underline{\theta}^\star) = -\alpha$, and $\overline{\theta}^\star$ is defined by $W(\overline{\theta}^\star) = 0$ (assuming such solutions exist). For $\theta \geq \overline{\theta}^\star$, the principal's objective is positive, so she chooses the maximal allocation 1, and the outside option constraint is slack. For $\theta \leq \overline{\theta}^\star$, the principal's objective is negative, so she would like to choose the minimal allocation 0; however, that could be in conflict with the outside option constraint. The optimal solution in this region is thus the "cheapest" way for the principal to satisfy the constraint. Recall that α is the principal's value for money; if $W(\theta) < -\alpha$, it becomes "cheaper" for the principal to satisfy the outside option constraint with a monetary transfer than with a positive allocation. Thus, the principal optimally sets $x^\star(\theta) = 0$ for types below $\underline{\theta}^\star$, and she uses a monetary payment $\underline{u}^\star = R(\underline{\theta}^\star)$ to satisfy the outside option constraint for all these types. For the remaining types $\theta \in [\underline{\theta}^\star, \overline{\theta}^\star]$, the principal uses the outside option allocation $x_0 \equiv R'$ to satisfy the constraint; she sets $x^\star(\theta) = R'(\theta)$ which makes the outside option constraint hold with equality everywhere in that interval. The corresponding indirect utility function U of the agent is constant (equal to $\underline{u}^\star$) below $\underline{\theta}^\star$, coincides with $R(\theta)$ on $[\underline{\theta}^\star, \overline{\theta}^\star]$, and has slope 1 above $\overline{\theta}^\star$.

The case of a non-monotone W is analogous, except that we must first "iron" $W(\theta)$ into its monotone version $-\overline{\mathcal{W}}'(\theta)$. Ironing is accomplished by first concavifying the integral of W, and then differentiating it to identify the intervals I on which the ironed objective is constant. Intuitively, suppose that $U(\theta)$ is set to its lowest feasible level $R(\theta)$ in the interval $[\underline{\theta}^\star, \overline{\theta}^\star]$ (i.e., the outside option constraint holds with equality everywhere). This makes the corresponding allocation rule x strictly increasing as long as the outside option is strictly increasing. If the principal's objective function W is decreasing around some type within $[\underline{\theta}^\star, \overline{\theta}^\star]$, the principal can do better by making the allocation flat around that type. The new allocation should still be as low as possible, and thus the endpoints of the ironing interval will satisfy the outside option constraint with equality (while the constraint may be slack in the interior).

Mathematically, we rely on the observation that—if we view allocation rules as CDFs—the outside option constraint takes a form similar to second-order stochastic dominance of the candidate distribution x by the fixed distribution x_0 defining the outside option.[16] The ironing procedure makes the allocation rule x flat on "ironing intervals"—this operation corresponds to taking a mean-preserving spread of the distribution x_0, and thus preserves the constraint that x is second-order stochastically dominated by x_0.

While the solution to problem (P′) is of independent interest, the key observation that we use to prove Theorem 1 is as follows. (Here, we implicitly assume that the solution to problem (P′) is unique but we show how to relax this assumption in Appendix A.4.)

Corollary 1. *The optimal solution* $(x^\star, \underline{u}^\star)$ *to problem* (P′) *defined in* (2) *depends linearly on the outside option* R.[17]

Corollary 1 follows from direct inspection of the solution to the principal's problem, and uncovers a (somewhat surprising) property of the linear utility model.[18] The key observation is that there exists a subset of types $\Theta^\star$ such that, no matter what R is, the outside option constraint always holds with equality for $\theta \in \Theta^\star$ and does not bind for $\theta \notin \Theta^\star$.[19] Our ironing procedure constructs a set $\Theta^\star = \bigcup \mathcal{I}^c \cup \{\underline{\theta}^\star, \overline{\theta}^\star\}$ with exactly this property. Indeed, $\Theta^\star$ depends solely on the principal's objective function and the distribution of types (but not on R). The optimal mechanism $(x^\star, \underline{u}^\star)$ then depends on R only through a linear transformation applied within each of the intervals identified by the ironing procedure, as equation (2) makes explicit.

Corollary 1 also highlights a key difference between our setting and that of Segal and Whinston (2016). The key challenge in solving mechanism design problems involving type-dependent outside options is to identify the set of binding participation constraints. In their multi-agent setting, Segal and Whinston (2016) restrict attention to option-to-own contracts and problems satisfying the monotonicity condition of Jullien (2000), where it suffices to

[16]Our constraint differs from a standard second-order stochastic dominance constraint by the presence of the constants u_0 and $\underline{u}$—this complicates our proof but does not pose a substantial challenge. See Kleiner et al. (2021) for a general theory of optimization subject to second-order stochastic dominance constraints. Our approach to the ironing procedure resembles most closely the one described in Akbarpour ⓡ al. (2024).

[17]Formally, if $(x_i^\star, \underline{u}_i^\star)$ solves the problem (P′) under the outside option R_i, for $i \in \{1, 2\}$, then $(x^\star, \underline{u}^\star) = \lambda(x_1^\star, \underline{u}_1^\star) + (1-\lambda)(x_2^\star, \underline{u}_2^\star)$ is a solution to problem (P′) under $R = \lambda R_1 + (1-\lambda)R_2$, for any $\lambda \in (0, 1)$.

[18]This result is surprising because it is known from the literature on mechanism design with type-dependent outside options that the mechanism chosen by the principal generally depends on the outside option function R in a complicated way, and that such problem do not generally admit a constructive solution.

[19]The constraint may hold with equality for $\theta \notin \Theta^\star$ under some R, but all constraints corresponding to $\theta \notin \Theta^\star$ can be relaxed without affecting the solution.

impose the participation constraint for a single type for each agent.[20] Beyond this special case, imposing the participation constraint for a single type is not sufficient. For our single-agent setting, we identify a set of types $\Theta^\star$ with two properties: *(i)* the outside option constraint holds with equality for types in $\Theta^\star$, and *(ii)* the outside option constraint is slack for types outside $\Theta^\star$. Although the set of binding participation constraints could be a subset of $\Theta^\star$ for any given outside option function R, the set $\Theta^\star$ determines the structure of the optimal mechanism for *any* outside-option function—leading to Corollary 1.

Step 3: Solving the designer's problem

Given the solution to the principal's problem derived in the previous step, we can simplify the formulation of the designer's problem. Instead of optimizing over feasible functions R, the designer can optimize over $u \geq 0$ and a non-decreasing allocation rule x that together define $R(\theta) \equiv u + \int_{\underline{\theta}}^{\theta} x(\tau)d\tau$—this reparameterization preserves all conditions that a feasible function R must satisfy by Lemma 1. A consequence of Corollary 1 is that the designer's problem is linear in u and x:

Lemma 3. *The designer's problem of choosing the optimal menu M is equivalent to solving the problem, for some function $\Phi : \Theta \to \mathbb{R}$,*

$$\max_{\substack{x \text{ non-decreasing,} \\ u \geq 0}} \quad \int_{\underline{\theta}}^{\bar{\theta}} \Phi(\theta) dx(\theta) - \alpha^\star u. \tag{3}$$

By Corollary 1, the allocation rule selected by the principal is linear in the outside option R. Because the designer's payoff is linear in the final allocation, it follows that the designer's problem is also linear in R. Given this observation and the change of variables described previously, Lemma 3 is a matter of bookkeeping: The function Φ is explicitly derived in the proof of Lemma 3 by integrating the objective function by parts so that the allocation rule x enters the designer's objective as a measure against which Φ is integrated.

Since $\alpha^\star \geq 0$, it is optimal to set $u = 0$ (which means that the lowest type of the agent does not benefit from holding the designer's optimal right). Then, problem (3) reduces to maximizing a linear functional over the convex set of non-decreasing allocation rules. By Bauer's maximum principle, there exists an optimal allocation rule that is an extreme point of the feasible set; and extreme points of the set of non-decreasing allocation rules are exactly step functions. Thus, we obtain the following result:

[20] Figueroa and Skreta (2009), Loertscher and Wasser (2019) and Loertscher and Muir (2025) also solve mechanism design problems involving countervailing incentives by exploiting this approach and parameterizing the optimal mechanisms in terms of a single type (referred to as the "critical type") for each agent.

Corollary 2. *Problem* (3) *admits a solution of the form:* $x^\star(\theta) = \mathbf{1}_{\theta \geq \theta^\star}$*, for some* $\theta^\star \in \Theta$*, and* $u^\star = 0$*.*

Corollary 2 implies Theorem 1: The optimal menu corresponding to the solution described by the corollary is precisely a singleton menu consisting of an option-to-own with price $p = \theta^\star$.

Summarizing our argument, an option-to-own corresponds to an extreme point of the set of feasible outside option functions R that the designer can induce by assigning rights to the agent. The designer's problem is *linear* in the outside option R. Therefore, an extreme point of the feasible set is among the maximizers of the designer's objective function. An analogous argument will be used in the generalized framework, to which we turn next.

4 General framework

The general framework features a pre-trading investment stage. In the presence of the non-commitment friction—the designer's inability to commit the principal to the trading mechanism—this gives rise to a hold-up problem. Property rights serve a dual purpose: they provide incentives for efficient investment and align the principal's mechanism with the designer's preferences. In addition, we explicitly model contractual incompleteness by introducing a *public state* (realized and observed at the trading stage) and allowing property rights to depend on a *contractible state* (that is partially informative about the public state).

4.1 Model

Overview. There are three time periods ($T = 0, 1, 2$) and three players (designer, principal, and agent). At time $T = 0$, the designer chooses the menu of rights that determines the agent's outside options. At time $T = 1$, the agent decides whether to undertake a costly investment which determines the joint distribution of the agent's type and a public state. At time $T = 2$, the agent's private type and the state are realized, and the principal chooses a trading mechanism in a sequentially rational manner, respecting the rights that the designer endowed the agent with at time $T = 0$. An overview of the model is presented in Figure 3.

Menu of rights. At time $T = 0$, the designer chooses a menu of rights held by the agent. The menu of rights has the same structure as in the baseline framework except that we allow the menu to be contingent on the *contractible state* $s \in \mathcal{S}$ realized in period $T = 2$. We will denote the contingent menu of rights by $\mathcal{M} = \{M_s\}_{s \in \mathcal{S}}$.

Investment. At time $T = 1$, the agent makes a binary investment decision. Investing is associated with a (sunk) cost $c > 0$. The investment decision determines the joint distribution of the agent's type $\theta \in \Theta := [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$ and the public state $\omega \in \Omega$, where Ω is an arbitrary

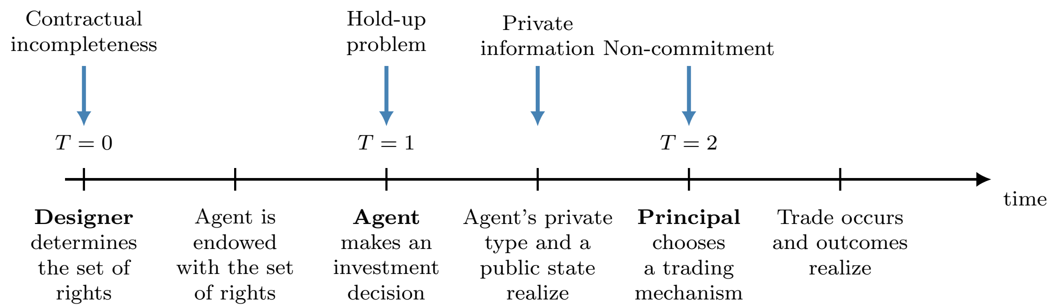

Figure 3: General framework overview and timeline.

finite set. If the agent invests, the public state is drawn from a distribution G; otherwise it is drawn from a distribution $\underline{G}$. In either case, the agent's type is drawn from a conditional distribution F_ω, for every $\omega \in \Omega$, with absolutely continuous density f_ω on Θ. We assume that investing increases the agent's type in the sense of first-order stochastic dominance:

$$\mathbb{E}_{\omega\sim(G-\underline{G})}[F_\omega(\theta)] \leq 0, \quad \forall \theta \in \Theta,$$

where for any (signed) measure μ on Ω and function $\phi : \Omega \to \mathbb{R}$, $\mathbb{E}_{\omega\sim\mu}[\phi(\omega)] := \int \phi(\omega) d\mu(\omega)$.

Trading Mechanisms. At time $T = 2$, the public state ω and the agent's type θ are realized. The agent's type θ is her private information, while the public state ω is observed by all parties. The contractible state s is given by $s = S(\omega)$, where $S : \Omega \to \mathcal{S}$ denotes a deterministic function that pins down the realized s as a function of ω.[21] The principal then chooses a trading mechanism, which we can again take to be a direct revelation mechanism satisfying appropriate incentive-compatibility and individual-rationality constraints.

Formally, for every realized ω, the principal chooses a mechanism $\langle x_\omega(\theta), t_\omega(\theta)\rangle$, where $x_\omega : \Theta \to [0, 1]$ denotes the allocation rule and $t_\omega : \Theta \to \mathbb{R}$ denotes the transfer rule. Given a direct mechanism $\langle x_\omega(\theta), t_\omega(\theta)\rangle$, incentive-compatibility requires that, for every $\omega \in \Omega$,

$$U_\omega(\theta) := \theta x_\omega(\theta) - t_\omega(\theta) \geq \theta x_\omega(\theta') - t_\omega(\theta'), \quad \theta, \theta' \in \Theta. \tag{IC$_\omega$}$$

Given a contingent menu of rights $\mathcal{M}$ and the realized contractible state s, we require agent's indirect utility to be at least at the level she would obtain by executing the best outside option from M_s: for every $\omega \in \Omega$,

$$U_\omega(\theta) \geq \max\{0, \max_{(x_i,t_i)\in M_{S(\omega)}} \{\theta x_i - t_i\}\}, \quad \forall \theta \in \Theta. \tag{IR$_\omega$}$$

[21] The assumption that the contractible state s is a deterministic function of ω is without loss of generality as long as both ω and s are publicly observed. This is because we can always redefine the public state as $\tilde{\omega} = (\omega, s)$.

Principal's problem. Given a realized state ω, the principal solves the problem

$$\max_{\langle x_\omega(\theta),\, t_\omega(\theta)\rangle} \int_\Theta [V_\omega(\theta)x_\omega(\theta) + \alpha t_\omega(\theta)]\, dF_\omega(\theta) \qquad (\text{P}_\omega)$$
$$\text{s.t.} \quad (\text{IC}_\omega), \quad (\text{IR}_\omega),$$

where $V_\omega : \Theta \to \mathbb{R}$ is upper semi-continuous in θ, and $\alpha > 0$ is the weight that the principal places on revenue. We denote by $\langle x^\star_\omega(\theta\,;\, M_{S(\omega)}), t^\star_\omega(\theta\,;\, M_{S(\omega)})\rangle$ the optimal mechanism for the principal in state ω.

Agent's problem. Anticipating the optimal mechanism used by the principal, the agent invests if and only if

$$\mathbb{E}_{\omega\sim(G-\underline{G})}\left[\int_\Theta \big(\theta x^\star_\omega(\theta;\, M_{S(\omega)}) - t^\star_\omega(\theta;\, M_{S(\omega)})\big)\, dF_\omega(\theta)\right] \geq c, \qquad (\text{I-OB})$$

where the left-hand side of equation (I-OB) is the differential payoff to the agent between the case when she invests ($\omega \sim G$) and when she does not invest ($\omega \sim \underline{G}$). Note that the agent's investment decision affects—through the distribution of the state ω—her type distribution, the mechanism chosen by the principal, and the set of rights the agent holds at the contracting stage.

Designer's problem. The designer's problem is then

$$\max_{\mathcal{M}} \quad \mathbb{E}_{\omega\sim G}\left[\int_\Theta \big(V^\star_\omega(\theta)x^\star_\omega(\theta;\, M_{S(\omega)}) + \alpha^\star t^\star_\omega(\theta;\, M_{S(\omega)})\big)\, dF_\omega(\theta)\right] \qquad (\text{D})$$
$$\text{s.t.} \quad (\text{I-OB}),$$

where $V^\star_\omega : \Theta \to \mathbb{R}$ is continuous in θ, and $\alpha^\star \geq 0$ is the weight that the designer places on transferring a unit of money from the agent to the principal. Unless stated otherwise, we assume: *(i)* the designer prefers to induce investment (which is why we included the investment-obedience constraint in the designer's problem), *(ii)* there exists some menu $\mathcal{M}$ that satisfies (I-OB), but *(iii)* the agent does not invest if she is not assigned any rights.

4.2 Discussion

The general framework introduces the *hold-up problem* (resulting from the agent's investment decision) and explicitly models the degree of *contractual incompleteness* through the interplay of the public and the contractible state.

The hold-up problem arises because the agent may decline to invest—even when investment is socially desirable—if private rents are insufficient to cover her cost. Since the principal chooses the mechanism in a sequentially rational manner, she may find it optimal to expropriate these rents ex post. Property rights must therefore be strong enough to protect the agent's incentives and overcome this friction.

The flexibility of property rights in the general framework is moderated by the degree of contractual incompleteness, captured by the coarseness of the contractible state $s = S(\omega)$ relative to the underlying state ω. Our framework flexibly accommodates both extremes: unconditional rights ($S(\omega) = s_0,\ \forall\omega$) and fully state-contingent rights ($S(\omega) = \omega,\ \forall\omega$). The incomplete-contracts literature has emphasized that, for practical purposes, the relevant set $\mathcal{S}$ is typically small, with most models taking $\mathcal{S}$ to be a singleton.

A special case of contractual incompleteness concerns the agent's investment decision. We implicitly assumed that investment is *de facto* observable to the principal, since its only payoff-relevant consequence is a change in the distribution of the state ω, whose realization the principal observes.[22] Whether property rights can be made contingent on investment, however, depends on the contractible state. If $S(\text{supp}(G)) \cap S(\text{supp}(\underline{G})) = \emptyset$, then *investment is contractible*: the designer can assign an entirely different set of rights conditional on the agent failing to invest. If $S(\omega)$ has the same distribution under G and $\underline{G}$, then *investment is non-contractible*: the designer must allocate the same rights regardless of whether the agent invested.

4.3 Main result

The following result is an extension of Theorem 1 to the general framework.

Theorem 1′. There exists an optimal menu $\mathcal{M}^\star$ that, for any $s \in \mathcal{S}$, takes the form $M_s^\star = \{(1, p_s), (y_s, p'_s)\}$ for some $p_s, p'_s \in \mathbb{R}$ and $y_s \in [0, 1)$.

Theorem 1′ shows that the optimal property right remains relatively simple in the general framework. Conditional on any contractible state s, it suffices to offer the agent a menu with at most two rights, one of which is an option-to-own. The second right takes the form of either a partial property right—granting the agent partial control over the resource at a reduced price—or a cash transfer to the agent (when $y_s = 0$ and $p'_s < 0$). As we demonstrate later, the form of this second right depends critically on whether investment is contractible. A further difference from Theorem 1 is that the terms of the rights (e.g., the price in the option-to-own) can now depend on the contractible state s.

Theorem 1′ does not rule out cases in which the optimal menu (conditional on some s) gives the agent no effective choice over which right to exercise, or even grants no rights at

[22] All our results extend to the case in which the principal does not observe the investment decision, at the cost of more cumbersome notation.

all. This can occur if one (or both) of the options in the menu are priced so highly that the agent never finds it optimal to use them.[23]

The proof of Theorem $1'$ is analogous to the proof of Theorem 1, with one key difference that explains the emergence of a second right in the optimal menu. The designer's problem remains linear in the control variables, which are now s-contingent outside option functions $R_s(\theta)$. However, an additional constraint appears: the investment-obedience constraint (I-OB). This constraint is linear in $\{R_s(\theta)\}_{s\in\mathcal{S}}$ because the optimal mechanism chosen by the principal is still linear in the outside option function $R_s(\theta)$. A linear problem with a single linear constraint admits a solution that is a convex combination of at most *two* extreme points.[24] A convex combination of two extremal outside-option functions (i.e., option-to-own contracts) corresponds precisely to the two-item menu in Theorem $1'$.

Economically, this means that the inclusion of a second option in the optimal menu is linked to the need to incentivize investment by the agent. This intuition is formalized in the following corollary.

Corollary 3. *If the investment-obedience constraint* (I-OB) *is slack at the optimal solution, then there exists an optimal solution to the designer's problem that takes the form of an s-contingent option-to-own: $M_s^\star = \{(1, p_s)\}$ for some $p_s \in \mathbb{R}$, for any $s \in \mathcal{S}$.*

Without the investment-obedience constraint, the designer chooses the price p_s in the option-to-own so as to align the principal's mechanism as closely as possible with her own preferences. Even when the investment-obedience constraint is present, it may be slack at the optimum, for example, when the designer chooses an option-to-own with a low price because she places a high weight on the agent's welfare.

In light of Corollary 3, it is precisely the presence of a binding investment-obedience constraint that can necessitate the inclusion of a second option in the optimal menu. Consequently, the form of this second option depends on whether investment is contractible.

Corollary 4. *Suppose that the investment cost c is sufficiently high.*

If investment is non-contractible (i.e., $S(\omega)$ has the same distribution under G and $\underline{G}$), there exists an optimal menu $M_s^\star = \{(1,\, p_s),\, (y_s,\, p_s')\}$ with $y_s > 0$, for each $s \in \mathcal{S}$.

If investment is contractible (i.e., $S(supp(G)) \cap S(supp(\underline{G})) = \emptyset$), there exists an optimal menu with $M_s^\star = \{(1,\, p_s),\, (0,\, -T)\}$ for $s \in S(supp(G))$ and $M_s^\star = \emptyset$ for $s \in S(supp(\underline{G}))$, for some $T > 0$.

[23] In applications, we show that a variety of configurations can be optimal: the menu may collapse to a singleton consisting of an option-to-own $(1,\, p_s)$, a pure transfer $(0,\, -p_s)$, or a costless partial right $(y_s,\, 0)$.

[24] This observation follows immediately from the results of Bauer (1958) and Szapiel (1975). Analogous results have been utilized in numerous recent papers in mechanism design (see, for example, Fuchs and Skrzypacz, 2015; Bergemann et al., 2018; Loertscher and Muir, 2024; Kang, 2023) and information design (see, for example, Le Treust and Tomala, 2019; Doval and Skreta, 2023).

Corollary 4 highlights a key distinction between the cases in which investment is contractible and those in which it is not. In the contractible case, the optimal menu—conditional on investment—consists of an option-to-own and a cash payment: The agent either keeps the good by paying a price p_s or relinquishes control in exchange for a monetary payment T. If the agent does not invest, she obtains no rights. To gain intuition, it is instructive to reinterpret how the agent chooses an outside option from her menu; so far, we have been assuming that the agent can only select one option. However, choosing one option from the menu $\{(1, p_s), (0, -T)\}$ is equivalent to paying the agent a cash transfer T *on top of* offering an option-to-own with price $p_s + T$ (conditional on investment). This alternative interpretation shows that when investment is contractible, it may be desirable to incentivize it with a direct monetary payment. In the remainder of the paper, we will use the term *lump-sum transfer* to refer to the alternative interpretation under which the cash payment T is given to the agent as long as she invests.

When investment is not contractible, lump-sum transfers are ineffective, since the agent would receive them regardless of her investment decision. Instead, the designer must rely on the fact that investment increases the agent's valuation θ. The optimal menu then incentivizes investment by raising the rents of higher types relative to lower types—by only offering options in which the agent obtains the good with strictly positive probability.

It is natural to conjecture that the simplicity of the optimal menu relies on our assumption that the agent makes a *binary* investment decision. Indeed, the proof of Theorem 1′ extends to cases in which additional (linear) constraints are introduced. If there are K such constraints—for example, because the agent has K alternative levels of investment to which she can deviate—then at most $K + 1$ options are required in the optimal menu (with an option-to-own always among them). This intuition, however, is only partially correct. What matters for the cardinality of the optimal menu is not the total number of constraints but the number of *binding* constraints. For instance, if investment is modeled as a continuous choice and the socially efficient level of investment is characterized by a first-order condition, then a single linear equation may be sufficient to capture the agent's obedience constraint. In that case, Theorem 1′ applies *verbatim*.[25]

Our methods did not rely on the fact that the linear constraint captured investment incentives. Any constraint that is linear in the allocation of the period-2 mechanism leads to the same mathematical conclusions. The constraint could capture other frictions, like the ones resulting from the agent's information acquisition as in Bergemann and Välimäki

[25] For a concrete example, suppose that the agent chooses a continuous effort level $e \in [0, 1]$ subject to a strictly convex cost $c(e)$. The agent's type θ is drawn from $\bar{F}$ with probability e, and from a lower distribution $\underline{F}$ with probability $1 - e$. Then, the first-order condition for some target level $e^\star$ is sufficient to ensure that the agent chooses $e^\star$, and the optimal menu contains at most two items.

(2002).

Relation to the incomplete-contracts literature. The hold-up problem resulting from relationship-specific non-contractible investments lies at the heart of the incomplete-contracts literature. Classical results (Grossman and Hart, 1986; Hart and Moore, 1990) imply that if there is no information asymmetry at the trading stage and only one agent invests, giving full property rights to that agent would result in an efficient outcome. This is consistent with a special case of Theorem $1'$. Suppose that the designer wants to induce the efficient outcome, the property right is unconditional, and consider the limiting case of our model in which the agent has no private information about her type.[26] Then, awarding the agent a full property right (an option-to-own with a zero price) is optimal.

However, Theorem $1'$ also shows what happens when these assumptions fail. First, consider the case when the designer maximizes an objective different from efficiency. With no information asymmetry, an ex-post efficient outcome will be implemented by the principal but property rights affect how much surplus the agent retains as a function of her realized type. Thus, the designer may strictly prefer a menu consisting of an option-to-own and a partial property right to a full property right. Second, suppose that the designer maximizes efficiency but the agent has private information at the trading stage. The classical conclusion may again fail:[27] The optimal property right is constructed to steer the outcome of the principal's mechanism toward that of a VCG mechanism—something that can be accomplished with an appropriate choice of the price in the option-to-own.

4.4 Optimal prices under regularity conditions

In this subsection, we analyze the structure of the optimal property rights—and provide a closed-form characterization of the optimal option-to-own price—under additional assumptions. Specifically, we assume that the principal's objective function $V_\omega(\theta)$ is non-decreasing, and that the (buyer's and seller's) virtual surplus functions

$$J^B_\omega(\theta) := \theta - \frac{1 - F_\omega(\theta)}{f_\omega(\theta)} \quad \text{and} \quad J^S_\omega(\theta) := \theta + \frac{F_\omega(\theta)}{f_\omega(\theta)}$$

are strictly increasing in θ, for every ω. To simplify notation, we also assume that the contractible state s is either singleton (in particular, investment is non-contractible) or binary

[26] The case with no private information can be obtained by taking the limit of our model in which the public state ω becomes arbitrarily informative about the realization of θ.

[27] Schmitz (2006) first observed this in a model where private information pertains only to the value of the outside option for the agent (but not the agent's value for trading).

in which case it reveals whether investment took place or not (investment is contractible).[28] This allows us to drop the subscript s in the notation (in the contractible case, the menu is empty conditional on no investment so we only need to characterize the menu conditional on investment). We refer to this specification of our model as the "monotone case."

We first characterize the optimal mechanism for the principal under these assumptions.

Proposition 1. *In the monotone case, for any outside option R, and conditional on any $\omega \in \Omega$, the principal chooses an optimal mechanism that induces an indirect utility*

$$U_\omega(\theta) = \begin{cases} R(\underline{\theta}^\star_\omega) & \theta < \underline{\theta}^\star_\omega, \\ R(\theta) & \theta \in [\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega], \\ R(\overline{\theta}^\star_\omega) + \theta - \overline{\theta}^\star_\omega & \theta > \overline{\theta}^\star_\omega, \end{cases}$$

where $\underline{\theta}^\star_\omega \leq \overline{\theta}^\star_\omega$ are defined by

$$V_\omega(\underline{\theta}^\star_\omega) + \alpha J^S_\omega(\underline{\theta}^\star_\omega) = 0 \quad \text{and} \quad V_\omega(\overline{\theta}^\star_\omega) + \alpha J^B_\omega(\overline{\theta}^\star_\omega) = 0,$$

whenever an interior solution $[\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega] \subset \Theta$ exists.

In the monotone case, the principal's problem admits an intuitive solution: The outside option constraint holds with equality at an "intermediate" interval of types $[\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega]$; the principal buys out rights using a cash payment for types $\theta \leq \underline{\theta}^\star_\omega$, and she allocates with probability one to types $\theta \geq \overline{\theta}^\star_\omega$. This is a direct consequence of the "ironing procedure" that we developed in Section 3.1.[29] In Figure 2, the monotone case corresponds to the special case where the objective function $\mathcal{W}$ is concave and therefore coincides with its concavification $\overline{\mathcal{W}}$. Intuitively, the principal wants to maximize the allocation for types higher than $\overline{\theta}^\star_\omega$ and minimize the allocation for types lower than $\overline{\theta}^\star_\omega$. Thus, the outside option constraint is slack for $\theta \geq \overline{\theta}^\star_\omega$. On the remainder of the type space, the principal uses the allocation rule to satisfy the outside option constraint for types above $\underline{\theta}^\star_\omega$, and the monetary payment to satisfy the outside option constraint for types below $\underline{\theta}^\star_\omega$. This intuition is embedded in the definitions of $\underline{\theta}^\star_\omega$ and $\overline{\theta}^\star_\omega$ from Proposition 1: The upper threshold $\overline{\theta}^\star_\omega$ is the cutoff type above which the principal would like to sell the resource to the agent, taking into account both the allocative effect and the revenue; the lower threshold $\underline{\theta}^\star_\omega$ is the cutoff type below which the principal would prefer to buy the resource from the agent.

[28] Formally, the public state is $\tilde{\omega} = (\mathbf{1}_i, \omega)$, where $\mathbf{1}_i$ is an indicator equal to 1 if and only if the agent invested, and $S(\tilde{\omega}) = \mathrm{id}(\mathbf{1}_i)$ where $\mathrm{id} : \{0, 1\} \to \{0, 1\}$ is the identity function.

[29] While Proposition 1 follows immediately from our ironing procedure, we could also derive it using weak duality based on Jullien (2000) because the dual variable takes a simple form in the monotone case.

Studying the monotone case allows us to provide a sharper characterization of the optimal property right. Suppose first that the investment-obedience constraint is slack and that the designer is indifferent over monetary transfers between the agent and the principal, i.e., $\alpha^\star = 0$. In this case, an option-to-own is optimal (Corollary 3), and its price p must satisfy

$$\mathbb{E}_{\omega\sim G}\Big[\, V^\star_\omega(p) f_\omega(p) \,\big|\, p \in [\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega] \Big] = 0, \tag{4}$$

whenever the optimal p is interior (see Appendix A.8 for supporting calculations). If the left-hand side of (4) is always positive, then $p = \underline{\theta}$ (equivalent to a full property right) is optimal; if it is always negative, then $p = \bar{\theta}$ (equivalent to no right) is optimal.

To interpret condition (4), note that the designer's first-best allocation in the second stage, assuming $V^\star_\omega(\theta)$ is non-decreasing, is to allocate the good to all types above the threshold $\theta^\star_\omega$ such that $V_\omega(\theta^\star_\omega) = 0$. The option-to-own price p is therefore chosen to implement this allocation *on average* across states ω. Crucially, the expectation in (4) is not taken with respect to the unconditional distribution of ω, but conditional on the event that p lies in the interval $[\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega]$ for the realized ω. This conditioning ensures that the option-to-own actually affects the final allocation of the resource. If $p < \underline{\theta}^\star_\omega$, the principal buys out the agent's option-to-own; if $p > \overline{\theta}^\star_\omega$, the principal offers the agent a strictly more attractive price within the mechanism. In both cases, the final allocation is locally constant in the price of the option-to-own, and thus the designer should ignore these contingencies when choosing the optimal p.

Next, suppose that the designer values revenue, i.e., $\alpha^\star > 0$. Then, (assuming an interior solution) the optimal price must satisfy

$$\mathbb{E}_{\omega\sim G}\left[\big(V^\star_\omega(p) + \alpha^\star J^B_\omega(p)\big)\, f_\omega(p) \mathbf{1}_{p\in[\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega]}\right] - \alpha^\star \mathbb{P}_{\omega\sim G}(p < \underline{\theta}^\star_\omega) = 0. \tag{5}$$

In this case, the designer would like to allocate to types for which $V^\star_\omega(\theta) + \alpha^\star J^B_\omega(\theta)$ is positive, where the virtual surplus term $J^B_\omega(\theta)$ measures how the allocation translates into revenue. However, equation (5) contains an additional effect captured by the second term. Whenever ω is such that $p < \underline{\theta}^\star_\omega$, the principal buys out the agent's right with a monetary payment. That payment is decreasing in p: the more attractive the option-to-own, the higher the compensation the principal must offer the agent to induce her to give up the right. Thus, a lower price p has no impact on the allocation in this region, but it reduces the principal's revenue—which the designer values at $\alpha^\star$. This additional effect pushes the designer to select a higher price p in the optimal menu. In particular, it implies that a full property right will be suboptimal whenever $\alpha^\star > 0$ and $\underline{\theta}^\star_\omega$ is bounded away from $\underline{\theta}$ for all ω (i.e., if the principal buys out the right with positive probability for every realization of ω).

Finally, we analyze how a binding investment-obedience constraint modifies the optimal price in the option-to-own, distinguishing between the cases where investment is contractible and where it is not. Throughout, we assume that the hold-up problem is sufficiently severe for the conclusion of Corollary 4 to apply.

The contractible-investment case. Suppose investment is contractible. By Corollary 4, the optimal menu can be implemented by awarding the agent a lump-sum transfer T and *additionally* letting her execute an option-to-own with price p, conditional on investment. Then, p must satisfy (assuming an interior solution)

$$\mathbb{E}_{\omega \sim G}\Big[\left(V_\omega^\star(p) + \alpha^\star p\right) f_\omega(p) \,\Big|\, p \in [\underline{\theta}_\omega^\star,\, \overline{\theta}_\omega^\star] \Big] = 0. \tag{6}$$

The formula for the optimal price is strikingly simple. Relative to (5) (the no-investment case), the second term disappears and the virtual surplus term $J_\omega^B(p)$ is replaced by p. In effect, the agent's information rents no longer enter the pricing condition. Consequently, when the investment-obedience constraint binds in the contractible case, the optimal option-to-own price will be *lower* than in the corresponding case without the hold-up problem.

For intuition, recall from Corollary 4 that the designer incentivizes investment through a combination of a lump-sum transfer and an option-to-own (offered conditional on investing). Optimality requires that the designer cannot profit by slightly lowering the price p in the option-to-own—thus relaxing the investment-obedience constraint—and then reducing the lump-sum transfer to make the constraint bind again. Lowering p has two effects: it changes the allocation, which the designer values locally at $V_\omega^\star(p)$, and it affects the revenue. Without adjusting the lump-sum transfer, the revenue effect would be captured by the virtual surplus term $J_\omega^B(p)$, as in (5). However, the binding investment-obedience constraint pins down agent's expected information rents conditional on investment, and hence the incremental net revenue—after adjusting the lump-sum transfer—excludes the information rent term. Similarly, the expected compensatory payment made by the principal buying out the agent's rights is absorbed by the lump-sum transfer. This means the designer no longer has an incentive to distort the allocation downward to reduce information rents, and so the optimal price p in the option-to-own is lower.

A noteworthy implication of (6) is that the optimal price p does *not* depend on the parameters of the agent's investment problem (such as the cost c), as long as the conclusion of Corollary 4 applies. Intuitively, the option-to-own is used to implement the desired allocation in the mechanism, while the lump-sum transfer is adjusted based on these parameters to ensure that the agent undertakes investment.

The non-contractible-investment case. When investment is non-contractible, the optimal menu may include two options, $M^\star = \{(1, p), (y, p')\}$, with two different prices p and p'. In contrast to the contractible case, it is no longer straightforward to disentangle the effects of the two options on investment incentives. Some intuition can be gained by introducing a Lagrange multiplier $\gamma \geq 0$ on the investment-obedience constraint. Although γ is determined endogenously, it must be non-decreasing in the cost of investment c, and can thus be interpreted as a measure of the severity of the hold-up problem. The optimal price p in the option-to-own $(1, p)$ must now satisfy (assuming an interior solution):[30]

$$\begin{aligned} &\mathbb{E}_{\omega\sim G}\left[\left(V^\star_\omega(p) + \alpha^\star J^B_\omega(p)\right) f_\omega(p)\mathbf{1}_{p\in[\underline{\theta}^\star_\omega,\, \overline{\theta}^\star_\omega]}\right] - \alpha^\star \mathbb{P}_{\omega\sim G}(p < \underline{\theta}^\star_\omega) \\ &+\gamma\left(\mathbb{E}_{\omega\sim G-\underline{G}}\left[(1 - F_\omega(p))\mathbf{1}_{p\in[\underline{\theta}^\star_\omega,\, \overline{\theta}^\star_\omega]}\right] + \mathbb{P}_{\omega\sim G}(p < \underline{\theta}^\star_\omega) - \mathbb{P}_{\omega\sim \underline{G}}(p < \underline{\theta}^\star_\omega)\right) = 0. \end{aligned} \tag{7}$$

The first line of equation (7) is the same as that of (5), while the second line describes two modifications due to a binding investment-obedience constraint. The first modification reflects the fact that the distribution of agent's values increases in the first-order stochastic dominance order after investing, i.e., $1 - F_\omega(p)$ increases on average across ω. This term will tend to make the optimal price p lower: The designer increases the incentives to invest by expanding the region in which the agent is the residual claimant. The second modification addresses the potential change in the probability that the principal will buy out the agent's rights. If that probability is higher conditional on investment, then strengthening the agent's right provides an additional incentive to invest (in anticipation of a higher expected compensatory payment from the principal), and hence also leads to a lower optimal price p.

5 Applications

In this section, we discuss an application of our framework to the problem of dynamic resource allocation. We then briefly mention several other applications, with details and supporting results relegated to Appendix B. Our goal is to provide an overview of how our framework can be mapped into various economic environments; a detailed analysis of policy implications in each environment is beyond the scope of this paper.

5.1 Dynamic resource allocation

A regulator (who is both the designer and the principal) allocates a scarce resource (e.g., electromagnetic spectrum or access to an oil tract) in a dynamic environment. The agent is

[30] In Appendix A.8, we show that equation (7) also captures the relevant trade-offs for the price p' of the partial right that could be included in the optimal menu.

assumed to control the resource at $T = 0$. (In Section 6, we comment on how our framework could be extended to model the problem of the initial allocation of the optimally-designed property rights.) At $T = 1$, the agent decides whether to invest in infrastructure that determines her value θ for keeping the resource in $T = 2$. The state ω is the value for the regulator of allocating the resource to an alternative use. The regulator is concerned with allocative efficiency, in that $V^\star_\omega(\theta) = V_\omega(\theta) = \theta - \omega$. Additionally, the regulator cares about revenue, and may attach a higher weight to revenue at $T = 2$, that is, $\alpha \geq \alpha^\star \geq 0$.

In this application, the agent is subject to a hold-up problem; additionally, the regulator suffers from time inconsistency (if $\alpha > \alpha^\star$). Time inconsistency could, for example, be the result of political pressure to raise a certain amount of revenue when reallocating scarce public resources.[31] The menu of rights selected by the regulator corresponds to the design of a license determining the agent's future rights to the resource.

Assuming regular distributions of types and a high enough cost of investment (making the investment-obedience constraint bind), we can apply the results from Section 4.4. We first assume that investment is contractible; resource use licenses sometimes include explicit clauses requiring proper maintenance or investment, such as "prudent operator standards" in oil and gas leases, or minimal coverage requirements in spectrum licenses. By Corollary 4, the optimal property right takes the form of a lump-sum transfer T for undertaking investment combined with an option-to-own with price p. In practice, the lump-sum transfer can be implemented by discounting the initial price for the license and imposing a punishment for failing to meet the required operating standards (monetary fine plus loss of the ability to execute the option-to-own). The option-to-own can be implemented as a renewable lease: As the lease termination date approaches, the current lessee may choose to pay the renewal fee p to keep the license for another term. By equation (6), price p satisfies

$$p = \mathbb{E}_{\omega \sim G}\left[\frac{\omega}{1+\alpha^\star} \; f_\omega(p) \,\middle|\, p \in [\underline{\theta}^\star_\omega, \, \overline{\theta}^\star_\omega] \right] = 0.$$

When the regulator is only concerned with efficiency ($\alpha^\star = 0$), she sets the price of the option-to-own to equal the expectation of her opportunity cost of the resource, with the expectation taken conditional on the option-to-own influencing the allocation. When the regulator is also concerned with revenue ($\alpha^\star > 0$), she sets the price p to be *lower*. To see why, recall that we assumed that the investment-obedience constraint binds: As $\alpha^\star$ grows larger, p goes down, but so does the lump-sum transfer T. Thus, in the contractible case,

[31] For example, the design goals for the "Incentive Auction" reallocating spectrum from TV broadcasters to mobile broadband operators included an explicit revenue target to cover FCC's costs and subsidize the federal budget; see Milgrom et al. (2012).

a regulator interested in raising revenue allocates stronger property rights to minimize the monetary transfer required to induce investment. In the limit as $\alpha^\star$ becomes large enough, the regulator chooses to allocate an unconditional property right (which maximizes the agent's willingness to pay for the license).

There is a high-level similarity between our optimal license and the "self-assessment mechanism" (and its variants) analyzed by Posner and Weyl (2017), Milgrom et al. (2017), and Weyl and Zhang (2022).[32] Both designs replace a rigid property right with a type of price mechanism that attempts to provide investment incentives for the current license holder conditional on a high value for keeping the resource. The right is less valuable to the license holder conditional on having a low value for the resource, which permits more efficient reallocation. The details of these designs, however, are different. In the case of the self-assessment mechanism, it is the license holder that names a price P; she then pays a fraction β of the price P to the regulator while committing to sell the license to anyone willing to offer P for it.[33] In our case, a price p is pre-specified, and it is the agent deciding whether to keep the license by paying p to the regulator (she may also be allowed to keep the license at a lower price if the state ω is low). In essence, our property right gives the agent an option to *guarantee* control over the resource but sacrifices some aspect of price discovery since the price p is fixed; in contrast, the self-assessment mechanism always exposes the current holder to some risk of losing control over the resource and uses that threat to extract more revenue from the holder conditional on having a high value. While our paper is the first to *derive* the optimal license design, the framework we propose does not include the self-assessment mechanism as a special case—leaving open the question of comparing the two designs more formally.[34]

In practice, it could be difficult to assess the extent to which an efficient level of investment is undertaken. If investment is not contractible, the optimal property right may become more complicated. By Corollary 4, the license may give two types of rights to the agent: an option-to-own with price p, and a partial right that gives access to the resource with probability y. In this case, as the lease termination date approaches, the current lessee either pays the renewal fee p to keep the license or submits a request for renewal at a lower fee p'; the request is then approved with probability y. While the regulator most likely could not commit to explicit randomization, she could instead commit to a review standard

[32] Early proponents of the self-assessment mechanism include Harberger (1965) and Tideman (1969).

[33] Weyl and Zhang (2022) propose a version of the self-assessment mechanism in which the price P is instead determined in a second price auction held between the incumbent and the entrants.

[34] Implementing the self-assessment mechanism requires a certain level of commitment to future trading mechanisms that we have ruled out by assumption. However, it is not clear how to formalize such partial commitment (full commitment makes any property right obsolete). We return to this issue in Section 6.

determining the average likelihood of a favorable decision.

The formula for optimal prices (given by equation 7) is more complicated than in the contractible case. In Appendix B, we derive explicit expressions in a numerical example. Prices tend to be lower when investment is more difficult to induce and higher when the regulator cares more about revenue (i.e., when $\alpha^\star$ is higher). A noteworthy special case—recovering a classical insight of Rogerson (1992)—arises when there is no time inconsistency, the regulator is only concerned with efficiency, and investment only affects the agent's private value for the resource. Then, it is optimal to allocate no rights to the agent because the VCG mechanism employed to reallocate resources at $T = 2$ already ensures efficient investment incentives.

5.2 Additional applications

5.2.1 Regulating a rental market

Our framework encompasses cases when the designer and the principal are separate entities with conflicting objectives. In Appendix B.2, we develop an application in which the designer is a policymaker, the principal is a rental company, and the agent is a business tenant (renting office space from the rental company). The rental company maximizes revenue while the policymaker is concerned with efficiency. Inefficiency arises because the tenant makes a sunk investment (e.g., installing specialized equipment) between two rental periods, which gives the rental company monopoly power over the tenant in the second period.

The menu of rights chosen by the policymaker captures regulation of a private rental market. A full property right corresponds to mandating a long-term lease. Other feasible regulations include giving the tenant the right to renew the lease at a pre-specified rent ("renewable-lease contract") or rent control.

Our framework predicts an important role for the *renewable-lease contract*. In Appendix B.2, we calculate the optimal price in the renewable-lease contract under parametric assumptions. The optimal price is equal to the conditional expectation of the future market rental rate minus a discount (which is strictly positive whenever the agent's investment constraint is binding). Similar regulation is often used in practice. For example, in the United Kingdom, the Landlord and Tenant Act 1954 provides commercial tenants with the right to renew any lease pertaining to a premises that it occupies for business purposes.

5.2.2 Patent policy

A classical economic question is how to reward and incentivize innovation and scientific discoveries. In Appendix B.3, we develop an application to patent policy design. The agent is a firm making a sunk investment in a new technology, and has private information about its marginal costs of production conditional on successful innovation. The principal is a patent office deciding whether the agent should have monopoly rights to the invention. The designer corresponds to a regulator designing patent policy. On top of the resulting hold-up problem, the patent office and the regulator might disagree on the welfare weight attached to consumer surplus.

We show that when the patent office has access to a reliable way of assessing the usefulness of the invention—corresponding to our assumption that investment is observable—the optimal patent policy may include a cash prize for the discovery. While cash prizes have been historically used to incentivize major discoveries,[35] in many cases regulators cannot verify whether an innovation is socially useful. Moreover, paying for discoveries could induce moral-hazard problems.[36]

When the investment is not observable and the patent office cannot pay the firm, we instead show that the optimal contract takes the form of allocating a *monopoly right free of charge with some fixed probability*. This result is driven by a conflict between the patent office's objective and the firm's willingness to pay to obtain monopoly rights. Consumer surplus from allowing free competition is particularly high when the costs of production are low (sales volume is large); but when the firm has low costs, it benefits more from obtaining monopoly rights. If the patent office attempted to screen firms using monetary payments, it would grant monopoly rights more often when the social harm from doing so is largest. As a result, the patent office optimally allocates the monopoly right with a probability that does not vary with the firm's production costs.[37] Optimal regulation takes the form of imposing a review standard that implements a probability of granting a patent that is sufficiently high to induce efficient investment.

5.2.3 Vaccine development

In Appendix B.4, we consider an application in which investment is observed and commissioned by a regulator who acts as both the designer and the principal. The agent is a pharmaceutical company developing a vaccine during a pandemic, with private information about marginal production costs. The regulator cares about patient welfare but faces uncertainty

[35] For example, The Longitude Act 1714 passed by the British Government offered a prize of 20,000 pounds (several million in purchasing power parity today) for invention of a clock that could operate with accuracy at sea.

[36] Kremer (1998) describes the possibility of bribery and rent-seeking, while Cohen et al. (2019) document the problem of "patent trolls" that would be exacerbated by offering additional financial incentives for "fake" discoveries.

[37] From a technical perspective, this application gives rise to an objective function for the principal that is *decreasing* in the agent's type. Consequently, ironing becomes a crucial step in characterizing the solution.

about the severity of the pandemic (or efficacy of the vaccine) at the time of contracting. The friction is that—in the absence of a contract—the regulator may not be interested in purchasing the vaccines after the investment costs have been sunk by the company.

The optimal contract between the regulator and the private producer takes the form of an *advanced market commitment*.[38] The regulator pays the company a lump-sum payment (for developing the vaccine) and guarantees a unit purchase price p (which we characterize in Appendix B.4). Ex post, when the pandemic turns out to be severe, the principal offers a unit price higher than p to increase the production of vaccines. When the pandemic is mild, the principal prefers to compensate the producer in cash, rather than buying the vaccines at the price p. In intermediate cases, the transaction price is equal to the contract price p.

5.2.4 Supply chain contracting

Our final application considers private contracting (see Appendix B.5 for details). A large firm (who is both the designer and the principal) wants to buy customized inputs from a small supplier (the agent). The supplier can invest in relationship-specific technology to produce the inputs at a privately observed marginal cost. Both firms maximize profits. Through the close interaction with the supplier, the large firm obtains a signal about the supplier's costs. The signal realization is non-contractible, leading to a ratchet effect. We investigate the form of the optimal contract offered by the large firm. On one extreme, without a long-term contract, firms can freely bargain ex post. On the other extreme, the large firm can commit to purchasing the entire future production of the supplier (which is effectively a merger in our setting). Possible intermediate arrangements include commitment by the large firm to buy a certain number of inputs at a pre-specified price.

Applying Theorem $1'$ shows that if investment by the supplier is not observable (e.g., the large firm cannot verify the quality of the inputs prior to assembling the final product), the large firm will optimally commit to a two-price scheme, committing to buy up to y units at some price p', or to buy any number of units at some lower price p. If investment by the supplier is observable, assuming the cost of investment is high enough, the large firm will offer an upfront payment for setting up production and then a guaranteed purchase price for any number of units.

The presence of private information at the trading stage (as well as the ratchet effect) make this application distinct from the typical setting in the incomplete-contracts literature. Without private information, Nöldeke and Schmidt (1995) find that the first-best outcome

[38] There has been a recent upsurge of interest in advanced market commitments among economists, particularly in relation to the use of these contracts as means to incentivize the production of vaccines (see, for example, Kremer et al., 2020a,b; Athey et al., 2020).

can be implemented (without relying on renegotiation, as in Aghion et al., 1994) by using an option contract that guarantees the seller a base price (lump-sum cash payment) plus an option price for delivery. Interestingly, if investment is observable and the conditions imposed in Section 4.4 hold, we arrive at the same conclusion, despite differences in the model and the fact that our optimal contract does not achieve the first best.

6 Concluding remarks

In this paper, we studied the design of property rights in an environment in which the designer cannot commit to a future trading mechanism. Instead, the designer can allocate property rights determining the holder's outside options in the mechanism. This perspective allowed us to employ a mechanism-design approach to characterize the optimal property right in a single-agent setup. The optimal right is more flexible than a full property right, and often allows the agent to retain control over the economic resource conditional on paying a pre-specified price. In this section, we briefly review extensions of our framework, and comment on future research directions.

Property rights as a form of partial commitment. In our framework the designer cannot commit to future trading mechanisms. From this perspective, property rights partly restore the designer's control over future allocations by specifying outside options that must be made available to the agent. There are other natural assumptions one could make about the degree of commitment. For instance, the designer might be able to *ban* certain outcomes (e.g., rent control restricts the set of prices a landlord can charge to a tenant). In the model, this would correspond to specifying a set of outcomes that cannot be offered in the mechanism run by the principal. If the designer can flexibly ban certain outcomes, mandate others, and condition these restrictions on the state, then she can effectively commit to the future mechanism. It is an interesting direction for future research to investigate how the strength of the designer's commitment power affects the form of optimal property rights.

State-contingent property rights. Even when ω is not contractible, the designer may be able to condition rights on ω indirectly by delegating the choice of the menu of rights to the principal. That is, the designer could design a *menu of (sub)menus*: At time $T = 2$, the principal first chooses a submenu from the menu, and then the agent can execute an outside option from the submenu. As long as the menu is constructed in such a way that the principal's relative preferences between submenus depend on the realized state ω, the designer can implement the dependence of the agent's outside option function on the state. Formally, the design problem is then one of choosing a state-contingent outside option for the agent but subject to an additional *incentive-compatibility constraint for the principal*. It is easy to show, by means of examples, that this extra flexibility may benefit the designer. This extension of our model would bear some similarity to the critique of the incomplete-contracts model offered by Maskin and Tirole (1999). Beyond theoretical curiosity, we find this research direction interesting because it could provide an optimality foundation for state-contingent property rights such as eminent domain.

Property rights versus bargaining power. In our framework, property rights affect the agent's outside options but not bargaining power. However, in the classical incomplete-contracts literature (Grossman and Hart, 1986; Hart and Moore, 1990), property rights were often associated with bargaining power. A natural extension of our framework is to symmetrize the positions of the principal and the agent by endowing both of them with private information and endogenizing their bargaining power. The trading stage could be modeled as a third party running an incentive-compatible mechanism as in Loertscher and Marx (2022) with type-dependent outside options and welfare weights reflecting the agents' relative bargaining power. The designer would then choose a menu of rights for both agents together with the bargaining weights. Our techniques, including the ironing approach to type-dependent outside options, could be helpful in analyzing this more general problem. On a conceptual level, this extension would allow a richer analysis of property rights, including the question of whose rights take precedence in case of conflict, as well as the role of abatement and easement.

Optimal allocation of property rights. We abstracted away from the problem of how to *allocate* optimally-designed rights by focusing on a single-agent setting and assuming that the agent simply holds the rights from the outset. This approach highlights the role of property rights in affecting future economic interactions. For example, it makes sense to think about the problem of designing a spectrum license separately from the problem of designing a spectrum *auction*. This is in part because—once the license is designed—allocating it to one of several agents is a standard mechanism design problem. Nevertheless, the two parts of the problem—the design of the right and its allocation—may interact in interesting ways. For example, when the designer is interested in maximizing revenue and the form of the right affects agents' willingness to pay.[39] We leave this direction for future research.

References

Aghion, P., M. Dewatripont, and P. Rey (1994): "Renegotiation Design with Unverifiable Information," *Econometrica*, 62, 257–282.

[39] Similar interactions have been analyzed in the literature on bidding with securities (see, for example, DeMarzo et al., 2005).

Akbarpour, M. ⓡ P. Dworczak ⓡ S. D. Kominers (2024): "Redistributive Allocation Mechanisms," *Journal of Political Economy*, 132, 1831–1875.

Athey, S., M. Kremer, C. Snyder, and A. Tabarrok (2020): "In the Race for a Coronavirus Vaccine, We Must Go Big. Really, Really Big." *New York Times*.

Balbuzanov, I. and M. H. Kotowski (2019): "Endowments, Exclusion, and Exchange," *Econometrica*, 87, 1663–1692.

Baliga, S. and T. Sjöström (2018): "A Theory of the Firm Based on Haggling, Coordination, and Rent-Seeking," Working paper.

Baron, J., R. Lombardo, J. Ryan, J. Suh, and Q. Valenzuela-Stookey (2025): "Mechanism Reform for Task Allocation," Working paper.

Bauer, H. (1958): "Minimalstellen von Funktionen und Extremalpunkte," *Archivder Mathematik*, 9, 389–393.

Bergemann, D., A. Bonatti, and A. Smolin (2018): "The Design and Price of Information," *American Economic Review*, 108, 1–48.

Bergemann, D. and J. Välimäki (2002): "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70, 1007–1033.

Bhaskar, D., A. McClellan, and E. Sadler (2023): "Regulation Design in Insurance Markets," *American Economic Review*, 113, 2546–80.

Bonnans, J. F. and A. Shapiro (2000): *Perturbation Analysis of Optimization Problems*, New York, NY: Springer.

Calabresi, G. and A. D. Melamed (1972): "Property Rules, Liability Rules, and Inalienability: One View of the Cathedral," *Harvard Law Review*, 85, 1089–1128.

Che, Y.-K. (2006): "Beyond the Coasian Irrelevance: Asymmetric Information," Unpublished notes.

Coase, R. H. (1960): "The Problem of Social Cost," *The Journal of Law & Economics*, 3, 1–44.

Cohen, L., U. G. Gurun, and S. D. Kominers (2019): "Patent Trolls: Evidence from Targeted Firms," *Management Science*, 65, 5461–5486.

Cramton, P., R. Gibbons, and P. Klemperer (1987): "Dissolving a Partnership Efficiently," *Econometrica*, 55, 615–632.

DeMarzo, P. M., I. Kremer, and A. Skrzypacz (2005): "Bidding with Securities: Auctions and Security Design," *American Economic Review*, 95, 936–959.

Doval, L. and V. Skreta (2023): "Constrained Information Design," *Mathematics of Operations Research*, 49, 78–106.

Dworczak, P. ⓡ S. D. Kominers ⓡ M. Akbarpour (2021): "Redistribution through

Markets," *Econometrica*, 89, 1665–1698.

Figueroa, N. and V. Skreta (2009): "The role of optimal threats in auction design," *Journal of Economic Theory*, 144, 884–897.

Fuchs, W. and A. Skrzypacz (2015): "Government Interventions in a Dynamic Market with Adverse Selection," *Journal of Economic Theory*, 158, 371–406.

Gallini, N. T. (1992): "Patent Policy and Costly Imitation," *The RAND Journal of Economics*, 23, 52–63.

Gibbons, R. (2005): "Four formal (izable) theories of the firm?" *Journal of Economic Behavior & Organization*, 58, 200–245.

Gilbert, R. and C. Shapiro (1990): "Optimal Patent Length and Breadth," *The RAND Journal of Economics*, 21, 106–112.

Grossman, S. J. and O. D. Hart (1986): "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 691–719.

Harberger, A. C. (1965): "Issues of Tax Reform in Latin America," in *Fiscal Policy for Economic Growth in Latin America: Papers and Proceedings of a Conference Held in Santiago, Chile, 1962. Baltimore, MD: Johns Hopkins University Press.*

Hart, O. (1995): *Firms, Contracts, and Financial Structure*, Oxford University Press.

Hart, O. and J. Moore (1990): "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98, 1119–1158.

Hermalin, B. E. and M. L. Katz (1993): "Judicial Modification of Contracts between Sophisticated Parties: A More Complete View of Incomplete Contracts and Their Breach," *Journal of Law, Economics, & Organization*, 9, 230–255.

Hitzig, Z. and B. Niswonger (2023): "Optimal Defaults, Limited Enforcement and the Regulation of Contracts," Working paper.

Hohfeld, W. N. (1917): "Fundamental Legal Conceptions as Applied in Judicial Reasoning," *The Yale Law Journal*, 26, 710–770.

Hopenhayn, H., G. Llobet, and M. Mitchell (2006): "Rewarding Sequential Innovators: Prizes, Patents, and Buyouts," *Journal of Political Economy*, 114, 1041–1068.

Jehiel, P. and A. Pauzner (2006): "Partnership Dissolution with Interdependent Values," *The RAND Journal of Economics*, 37, 1–22.

Jullien, B. (2000): "Participation Constraints in Adverse Selection Models," *Journal of Economic Theory*, 93, 1–47.

Kang, Z. Y. (2023): "The Public Option and Optimal Redistribution," Working paper.

Kleiner, A., B. Moldovanu, and P. Strack (2021): "Extreme Points and Majorization: Economic Applications," *Econometrica*, 89, 1557–1593.

Klemperer, P. (1990): "How Broad Should the Scope of Patent Protection Be?" *The RAND Journal of Economics*, 21, 113–130.

Kremer, M. (1998): "Patent Buyouts: A Mechanism for Encouraging Innovation," *The Quarterly Journal of Economics*, 113, 1137–1167.

Kremer, M., J. Levin, and C. M. Snyder (2020a): "Advance Market Commitments: Insights from Theory and Experience," *AEA Papers and Proceedings*, 110, 269–73.

——— (2020b): "Designing Advance Market Commitments for New Vaccines," Working paper.

Le Treust, M. and T. Tomala (2019): "Persuasion with Limited Communication Capacity," *Journal of Economic Theory*, 184, 104940.

Lewis, T. R. and D. E. M. Sappington (1989): "Countervailing Incentives in Agency Problems," *Journal of Economic Theory*, 49, 294–313.

Loertscher, S. and L. M. Marx (2022): "Incomplete Information Bargaining with Applications to Mergers, Investment, and Vertical Integration," *American Economic Review*, 112, 616–49.

Loertscher, S. and E. Muir (2024): "Optimal Labor Procurement under Minimum Wages and Monopsony Power," Working paper.

——— (2025): "Optimal Hotelling Auctions," Working paper.

Loertscher, S. and C. Wasser (2019): "Optimal structure and dissolution of partnerships," *Theoretical Economics*, 14, 1063–1114.

Lüelfesmann, C. (2025): "Authority and screening in a principal agent model," *Journal of Economic Behavior & Organization*, 231, 106862.

Maggi, G. and A. Rodriguez-Clare (1995): "On Countervailing Incentives," *Journal of Economic Theory*, 66, 238–263.

Maskin, E. and J. Tirole (1999): "Unforeseen Contingencies and Incomplete Contracts," *The Review of Economic Studies*, 66, 83–114.

Matouschek, N. (2004): "Ex Post Inefficiencies in a Property Rights Theory of the Firm," *The Journal of Law, Economics, and Organization*, 20, 125–147.

Milgrom, P., L. Ausubel, J. Levin, and I. Segal (2012): "Incentive Auction: Rules and Discussion," *Appendix C of FCC 12-118 (Notice of Proposed Rulemaking on the Incentive Auction).*

Milgrom, P., E. G. Weyl, and A. L. Zhang (2017): "Redesigning Spectrum Licenses," *Regulation*, 40(3), 22–26.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–78.

Myerson, R. and M. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265–281.

Nöldeke, G. and K. M. Schmidt (1995): "Option Contracts and Renegotiation: A Solution to the Hold-up Problem," *The RAND Journal of Economics*, 26, 163–179.

——— (1998): "Sequential Investments and Options to Own," *The RAND Journal of Economics*, 29, 633–653.

Posner, E. and E. G. Weyl (2017): "Property is Only Another Name for Monopoly," *Journal of Legal Analysis*, 9 (51).

Rogerson, W. P. (1992): "Contractual Solutions to the Hold-Up Problem," *The Review of Economic Studies*, 59, 777–793.

Schmitz, P. W. (2002): "Simple contracts, renegotiation under asymmetric information, and the hold-up problem," *European Economic Review*, 46, 169–188.

——— (2006): "Information Gathering, Transaction Costs, and the Property Rights Approach," *American Economic Review*, 96, 422–434.

Segal, I. and M. Whinston (2012): "Property Rights," in *Handbook of Organizational Economics*, ed. by R. Gibbons and J. Roberts, Princeton: Princeton University Press.

Segal, I. and M. D. Whinston (2016): "Property Rights and the Efficiency of Bargaining," *Journal of the European Economic Association*, 14, 1287–1328.

Szapiel, W. (1975): "Points Extrémaux dans les Ensembles Convexes(I). Théorie Générale," *Bulletin de l'Académie Polonaise des Sciences. Sériedes Sciences Mathématiques, Astronomiques et Physiques*, 23, 389–393.

Thereze, J. and U. Vaidya (2025): "Delegated Contracts," Working paper.

Tideman, N. (1969): "Three Approaches to Improving Urban Land Use," *PhD dissertation, University of Chicago.*

Valenzuela-Stookey, Q. (2025): "Automation and Task Allocation Under Asymmetric Information," Working paper.

Weyl, E. G. and J. Tirole (2012): "Market Power Screens Willingness-to-Pay," *The Quarterly Journal of Economics*, 127, 1971–2003.

Weyl, E. G. and A. L. Zhang (2022): "Depreciating Licenses," *American Economic Journal: Economic Policy*, 14, 422–48.

Williamson, O. E. (1979): "Transaction-Cost Economics: The Governance of Contractual Relations," *The Journal of Law & Economics*, 22, 233–261.

Wright, B. D. (1983): "The Economics of Invention Incentives: Patents, Prizes, and Research Contracts," *The American Economic Review*, 73, 691–707.

A Remaining Proofs

A.1 Proof of Lemma 1

Given a menu of rights $M = \{(x_i, t_i)\}_{i \in I}$, let $R(\theta) = \max\{0, \max_{i \in I}\{\theta x_i - t_i\}\}$. Since R is constructed by maximizing over a family of affine functions, this implies that R is convex and admits a right derivative. Moreover, since each affine function $\theta x_i - t_i$ has a non-negative gradient $x_i \in [0, 1]$, this implies that R is non-decreasing in θ and that $|\partial_+ R(\theta)| \in [0, 1]$, where $\partial_+ R$ denotes the right derivative of R. Conversely, suppose that we have a type-dependent outside option function R that is non-negative, non-decreasing and convex, and admits a right derivative that is bounded above by 1. Then, for all $\theta \in \Theta$, we can set $y(\theta) = \partial_+ R(\theta)$ and $S(\theta) = \theta \partial_+ R(\theta) - R(\theta)$. Since R is convex, the allocation rule y is non-decreasing. The envelope theorem then implies that the menu $M = \{(y(\theta), S(\theta))\}_{\theta \in \Theta}$ implements the reservation utility function R and is such that $R(\theta) = \max\{0, \max_{\theta' \in \Theta}\{\theta y(\theta') - S(\theta')\}\}$ as required.

A.2 Proof of Lemma 2

Consider first an auxiliary problem in which we fix $\underline{u}$ at some level weakly above u_0. Note that our assumption that the principal's objective function W is upper semi-continuous in θ implies that it is without loss of generality to restrict attention to right-continuous allocation rules. We will treat the allocation rule x as a CDF by extending it to the real line and assuming that $x(\theta) = 0$ for all $\theta < \underline{\theta}$, and $x(\theta) = 1$ for all $\theta \geq \bar{\theta}$.[40] Applying Leibniz's rule, integrating by parts, and using $\mathcal{W}(\bar{\theta}) = 0$ and $\lim_{\theta \nearrow \underline{\theta}} x(\theta) = 0$:

$$\int_{\underline{\theta}}^{\bar{\theta}} W(\theta) x(\theta) d\theta = -\int_{\underline{\theta}}^{\bar{\theta}} x(\theta) d\left(\int_{\theta}^{\bar{\theta}} W(\tau) d\tau\right) = \int_{\underline{\theta}}^{\bar{\theta}} \mathcal{W}(\theta) dx(\theta).$$

The problem is now to choose a CDF x to maximize

$$\int_{\underline{\theta}}^{\bar{\theta}} \mathcal{W}(\theta) dx(\theta) \text{ subject to } \int_{\underline{\theta}}^{\theta} x(\tau) d\tau \geq (u_0 - \underline{u}) + \int_{\underline{\theta}}^{\theta} x_0(\tau) d\tau, \, \forall \theta.$$

Up to the constant term $u_0 - \underline{u}$, the constraint states that x must be second-order stochastically dominated by x_0. In particular, if $\mathcal{W}$ is non-decreasing and concave, then the optimal x must satisfy the inequality as an equality (whenever this is feasible). Formally, define

[40] While the optimal mechanism might have $x(\bar{\theta}) < 1$, imposing $x(\bar{\theta}) = 1$ is without loss of generality since it is not affecting the principal's expected payoff and preserves all the constraints.

$\bar{x}(\theta) := x_0(\theta)\mathbf{1}_{\theta \geq \theta_0}$, where θ_0 is defined by

$$u_0 - \underline{u} + \int_{\underline{\theta}}^{\theta_0} x_0(\tau) d\tau = 0 \text{ (and } \theta_0 = \underline{\theta} \text{ if there is no solution).} \tag{8}$$

The allocation $\bar{x}$ is feasible by construction. If $\mathcal{W}$ is non-decreasing and concave, then any feasible x yields a lower objective than $\bar{x}$ because $\bar{x}$ second-order stochastically dominates any feasible x. Moreover, if a monotone x is second-order stochastically dominated by $\bar{x}$, then x is feasible.

The key idea of the proof (mimicking the logic behind classical "ironing") is to define a relaxed problem in which the objective is concave non-decreasing, and then show that the value of the relaxed problem can be achieved in the original problem.

Let $\overline{\mathcal{W}}$ be the concave closure of $\mathcal{W}$, and let $\overline{\mathcal{W}}_+$ be the non-decreasing concave closure of $\mathcal{W}$. Note that $\overline{\mathcal{W}}_+$ differs from $\overline{\mathcal{W}}$ only in that $\overline{\mathcal{W}}_+(\theta)$ is constant—equal to the global maximum $\mathcal{W}(\overline{\theta}^\star)$—for all $\theta \geq \overline{\theta}^\star$, where $\overline{\theta}^\star$ is defined as in the main text. Clearly, $\mathcal{W} \leq \overline{\mathcal{W}}_+$ and $\overline{\mathcal{W}}_+$ is non-decreasing and concave. By our previous argument, we have obtained an upper bound on the value of the problem equal to $\int_{\underline{\theta}}^{\bar{\theta}} \overline{\mathcal{W}}_+(\theta) d\bar{x}(\theta)$. We will now construct an allocation rule $x^\star$ that is feasible in the original problem and achieves this upper bound. Define $\mathcal{I}'$ to be the (at most countable) collection of maximal open intervals (a, b) within $(\underline{\theta}, \overline{\theta}^\star)$ with the property that $\mathcal{W}$ lies strictly below $\overline{\mathcal{W}}$ on (a, b). Note that the definition of $\mathcal{I}'$ differs from the definition of $\mathcal{I}$ in the main text only in that the former is defined on $(\underline{\theta}, \overline{\theta}^\star)$, and the latter on $(\underline{\theta}^\star, \overline{\theta}^\star)$. Define

$$x^\star(\theta) = \begin{cases} \frac{\int_a^b \bar{x}(\tau) d\theta}{b-a} & \theta \in (a, b) \text{ for some } (a, b) \in \mathcal{I}', \\ \bar{x}(\theta) & \theta \in (\underline{\theta}, \overline{\theta}^\star) \setminus \bigcup \mathcal{I}', \\ 1 & \theta \geq \overline{\theta}^\star. \end{cases}$$

Intuitively, $x^\star$ (viewed as a CDF) only attaches probability mass to types θ at which the objective $\mathcal{W}$ coincides with the concavified objective $\overline{\mathcal{W}}_+$. Note that $x^\star$ is feasible. It is non-decreasing because $\bar{x}$ is non-decreasing. Moreover, it is second-order stochastically dominated by $\bar{x}$ because it is obtained from $\bar{x}$ by a series of mean-preserving spreads within $(\underline{\theta}, \overline{\theta}^\star)$, and by a single first-order stochastic dominance shift above $\overline{\theta}^\star$—this suffices for feasibility, as noted previously.

We now argue that $x^\star$ achieves the upper bound of the value function. Let $x^\star_-(\overline{\theta}^\star)$ denote the left limit of $x^\star$ at $\overline{\theta}^\star$. Then,

$$\int_{\underline{\theta}}^{\bar{\theta}} \mathcal{W}(\theta)dx^\star(\theta) = \int_{(\underline{\theta},\overline{\theta}^\star)\setminus\bigcup\mathcal{I}'} \mathcal{W}(\theta)dx^\star(\theta) + \mathcal{W}(\overline{\theta}^\star)(1 - x^\star_-(\overline{\theta}^\star))$$

$$= \int_{(\underline{\theta},\overline{\theta}^\star)\setminus\bigcup\mathcal{I}'} \overline{\mathcal{W}}(\theta)dx^\star(\theta) + \overline{\mathcal{W}}(\overline{\theta}^\star)(1 - x^\star_-(\overline{\theta}^\star)) = \int_{\underline{\theta}}^{\bar{\theta}} \overline{\mathcal{W}}_+(\theta)d\bar{x}(\theta), \tag{9}$$

where the first equality follows from the fact that $x^\star$ puts no mass on types in the set $\bigcup\mathcal{I}'$ and types above $\overline{\theta}^\star$; the second equality follows from the fact that, by construction, $\mathcal{W} = \overline{\mathcal{W}}$ on the support of $x^\star$ within $(\underline{\theta},\, \overline{\theta}^\star)$, while the equality at $\overline{\theta}^\star$ follows because $\mathcal{W}$ and $\overline{\mathcal{W}}$ must coincide at the global maximum; and the third equality follows by linearity of $\overline{\mathcal{W}}_+$ in intervals $(a,\, b)$ belonging to $\mathcal{I}'$ and the fact that in such intervals $x^\star$ is a mean-preserving spread of $\bar{x}$, as well as from the fact that $\overline{\mathcal{W}}_+$ is constant above $\overline{\theta}^\star$. This proves that $x^\star$ is optimal.

In the last step of the proof, we maximize over $\underline{u}$. Note that—given the above derivation—the problem of choosing the optimal $\underline{u}$ can be written as

$$\max_{\underline{u}\geq u_0} \left\{ \overline{\mathcal{W}}_+(\theta_0(\underline{u}))x_0(\theta_0(\underline{u})) + \int_{(\theta_0(\underline{u}),\bar{\theta}]} \overline{\mathcal{W}}_+(\theta)dx_0(\theta) - \alpha\underline{u} \right\},$$

where $\theta_0(\underline{u})$ is defined as in (8), now with the dependence on $\underline{u}$ made explicit in the notation. Given that $\alpha > 0$, it is never optimal to choose $\underline{u}$ such that the equation $u_0 - \underline{u} + \int_{\underline{\theta}}^{\theta_0} x_0(\tau)d\tau = 0$ defining $\theta_0(\underline{u})$ does not have a solution, since this would make the outside option constraint slack everywhere. Given that $u_0 - \underline{u} + \int_{\underline{\theta}}^{\theta_0(\underline{u})} x_0(\tau)d\tau = 0$ must hold, we can maximize over the cutoff type θ_0 directly:

$$\max_{\theta_0} \left\{ \overline{\mathcal{W}}_+(\theta_0)x_0(\theta_0) + \int_{(\theta_0,\bar{\theta}]} \overline{\mathcal{W}}_+(\theta)dx_0(\theta) - \alpha \int_{\underline{\theta}}^{\theta_0} x_0(\tau)d\tau - \alpha u_0 \right\}.$$

Integration by parts yields $\int_{\underline{\theta}}^{\theta_0} x_0(\tau)d\tau = \theta_0 x_0(\theta_0) - \int_{\underline{\theta}}^{\theta_0} \tau dx_0(\tau)$. Additionally, we have

$$\int_{(\theta_0,\bar{\theta}]} \overline{\mathcal{W}}_+(\theta)dx_0(\theta) = \int_{[\underline{\theta},\bar{\theta}]} \overline{\mathcal{W}}_+(\theta)dx_0(\theta) - \int_{[\underline{\theta},\theta_0]} \overline{\mathcal{W}}_+(\theta)dx_0(\theta).$$

Omitting terms that do not depend on θ_0 and rearranging, we obtain an equivalent representation of the problem:

$$\max_{\theta_0\geq\underline{\theta}} \left\{ (\overline{\mathcal{W}}_+(\theta_0) - \alpha\theta_0)x_0(\theta_0) - \int_{[\underline{\theta},\theta_0]} (\overline{\mathcal{W}}_+(\theta) - \alpha\theta)dx_0(\theta) \right\}.$$

Integrating the second term by parts yields another equivalent representation:

$$\max_{\theta_0 \geq \underline{\theta}} \left\{ \int_{[\underline{\theta},\, \theta_0]} (\overline{\mathcal{W}}_+'(\theta) - \alpha) x_0(\theta) d\theta \right\}. \tag{10}$$

The function $\overline{\mathcal{W}}_+(\theta)$ is concave, and hence differentiable almost everywhere, with a decreasing derivative. Thus, the optimal θ_0 is the supremum over types θ such that $\overline{\mathcal{W}}_+'(\theta) \geq \alpha$ (with $\theta_0 = \underline{\theta}$ if the derivative is always below α). Note that $\overline{\mathcal{W}}(\theta) = \overline{\mathcal{W}}_+(\theta)$ for all θ such that $\overline{\mathcal{W}}_+'(\theta) \geq \alpha$, and hence the optimal θ_0 coincides with the definition of $\underline{\theta}^\star$ given in the main text.

Finally, we can plug the optimal $\theta_0 = \underline{\theta}^\star$ into the definition of $\bar{x}$ to obtain

$$x^\star(\theta) = \begin{cases} \frac{\int_a^b x_0(\tau) \mathbf{1}_{\tau \geq \underline{\theta}^\star} d\tau}{b-a} & \theta \in (a,\, b) \text{ for some } (a,\, b) \in \mathcal{I}', \\ x_0(\theta) \mathbf{1}_{\theta \geq \underline{\theta}^\star} & \theta \in (\underline{\theta},\, \overline{\theta}^\star) \setminus \bigcup \mathcal{I}', \\ 1 & \theta \geq \overline{\theta}^\star. \end{cases}$$

Notice that $\underline{\theta}^\star$ cannot belong to the interior of any interval $(a,\, b) \in \mathcal{I}'$ because, by definition, $\overline{\mathcal{W}}$ is linear on any such $(a,\, b)$. Thus, $x^\star(\theta)$ must be 0 for any $\theta \leq \underline{\theta}^\star$, and we can define $\mathcal{I}$ to be the intersection of $\mathcal{I}'$ with $(\underline{\theta}^\star,\, \overline{\theta}^\star)$—this gives us the definition of $\mathcal{I}$ from the main text. Finally, noting that $x_0(\theta) = R'(\theta)$ almost everywhere and that $\underline{u}^\star = R(\underline{\theta}^\star)$, we can verify that the optimal $(x^\star,\, \underline{u}^\star)$ defined above coincide with those defined by equation (2).

A.3 Proof of Lemma 3

Let $W^\star(\theta) := (V^\star(\theta) + \alpha^\star J^B(\theta)) f(\theta)$ denote the designer's objective multiplied by the density of types. Using the explicit solution to the principal's problem derived in Lemma 2, we can write the designer's expected payoff conditional on choosing an outside option function R as

$$-\alpha^\star R(\underline{\theta}^\star) + \sum_{(a,\, b) \in \mathcal{I}} \frac{\int_a^b R'(\tau) d\tau \int_a^b W^\star(\theta) d\theta}{b-a} + \int_{(\underline{\theta}^\star,\, \overline{\theta}^\star) \setminus \bigcup \mathcal{I}} R'(\theta) W^\star(\theta) d\theta,$$

where we have omitted the term $\int_{\overline{\theta}^\star}^{\bar{\theta}} W^\star(\theta) d\theta$ that does not depend on the chosen R.

We can now change variables (relying on the characterization from Lemma 1) by letting $R(\theta) = u + \int_{\underline{\theta}}^{\theta} x(\tau) d\tau$, for some $u \geq 0$, and non-decreasing allocation rule x. This gives rise to the following optimization problem for the designer: Maximize over non-decreasing allocation rules x and $u \geq 0$

$$-\alpha^\star \left(u + \int_{\underline{\theta}}^{\underline{\theta}^\star} x(\theta)\, d\theta \right) + \sum_{(a,\, b) \in \mathcal{I}} \frac{\int_a^b x(\tau)\, d\tau \int_a^b W^\star(\theta)\, d\theta}{b-a} + \int_{(\underline{\theta}^\star,\, \overline{\theta}^\star) \setminus \bigcup \mathcal{I}} x(\theta) W^\star(\theta)\, d\theta.$$

Since the objective is linear in $x(\theta)$, using integration by parts,[41] we can rewrite the problem as $\max_{x(\theta),\, u \geq 0} \int_{\underline{\theta}}^{\overline{\theta}} \Phi(\theta) dx(\theta) - \alpha^\star u$, where

$$\Phi(\theta) = -\alpha^\star \left(\underline{\theta}^\star - \theta\right)_+ + \sum_{(a,\, b) \in \mathcal{I}} \frac{(b - \max\{a, \theta\})_+ \int_a^b W^\star(\theta) d\theta}{b-a} + \sum_{[a,\, b] \in \mathcal{I}^c} \mathbf{1}_{\{\theta \leq b\}} \left(\int_{\max\{a,\, \theta\}}^b W^\star(\tau) d\tau \right).$$

A.4 Remark about tie-breaking rules

In the proof of Theorem 1, we have assumed a particular tie-breaking rule—not necessarily designer-preferred—in case of principal's indifference, implicit in how we defined the cutoffs $\underline{\theta}^\star$, $\overline{\theta}^\star$ as well as the ironing intervals $\mathcal{I}$ in the proof of Lemma 2. However, the proof of Lemma 2 allows us to characterize all solutions to the principal's problem. Indeed, any solution $x^\star$ must satisfy the string of equalities (A.2), and any optimal $\underline{\theta}^\star$ must solve problem (10). It follows that all solutions to problem (P′) can be obtained by modifying our baseline solution $(x^\star,\, \underline{u}^\star)$ in the following ways:

1. $\overline{\theta}^\star$ can be taken to be *any* global maximum of $\mathcal{W}$ (not necessarily the smallest one);

2. If $\mathcal{W} = \overline{\mathcal{W}}$ is *affine* on some interval $[a,\, b]$, then we can take any mean-preserving spread of $x^\star$ in that interval (in the baseline solution, $x^\star(\theta) = R'(\theta)$ on $[a,\, b]$);

3. $\underline{\theta}^\star$ can be taken to be *any* type θ with the property $\alpha = \overline{\mathcal{W}}'(\theta)$ if there are multiple such θ (not necessarily the largest one).

We will call a tie-breaking rule *consistent* if it breaks the principal's indifference by maximizing an auxiliary objective function $\int_{\underline{\theta}}^{\overline{\theta}} \phi(\theta) x^\star(\theta) d\theta - \beta \underline{u}^\star$, where $\phi:\, \Theta \to \mathbb{R}$ is continuous. Clearly, maximizing or minimizing the designer's payoffs are both consistent tie-breaking rules.

We claim that the solution picked by a consistent tie-breaking rule is linear in R, as in Corollary 1. The reason is that the optimal choice of $\overline{\theta}^\star$ and $\underline{\theta}^\star$ will be invariant to R; moreover, maximizing $\int_{\underline{\theta}}^{\overline{\theta}} \phi(\theta) x^\star(\theta) d\theta$ over mean-preserving spreads of $R'(\theta)$ in some interval $[a,\, b]$ can be solved by applying an ironing procedure analogous to the one that we used to

[41] In particular, we use the fact that $\int_a^b g(\theta) x(\theta) d\theta = \int_{\underline{\theta}}^{\overline{\theta}} \mathbf{1}_{\{\theta \leq b\}} \left(\int_{\max\{a,\, \theta\}}^b g(\tau) d\tau \right) dx(\theta)$.

solve the principal's problem. As we have shown, this procedure results in an $R-$invariant partition of $[a,\, b]$ into (at most countably many) subintervals on which either *(i)* the optimal $x^{\star}(\theta)$ is equal to $R'(\theta)$, in which case the subinterval can be included in the collection $\mathcal{I}^c$, or *(ii)* the optimal $x^{\star}(\theta)$ is constant, in which case the subinterval can be included in the collection $\mathcal{I}$. Overall, a consistent tie-breaking rule results in a solution whose structure is the same as in the proof of Lemma 2, except that the R-invariant cutoff types $\underline{\theta}^{\star}$ and $\overline{\theta}^{\star}$, as well as the R-invariant collection of ironing intervals, may be different. Thus, the solution is still linear in R.

A.5 Proof of Theorem 1′

Observe that *(i)* a choice of a contingent menu $\mathcal{M}$ by the designer corresponds to the choice of an $s-$contingent outside-option function R_s (as described by Lemma 1), and *(ii)* conditional on the realized state ω, the solution to the principal's problem is still described by Lemma 2, except that we need to keep track of the dependence of the principal's optimal mechanism on the public state (which we will do by indexing all relevant variables by ω).

Instead of optimizing over feasible outside-option functions R_s, the designer can optimize over $u_s \geq 0$ and non-decreasing allocation rule x_s that together define $R_s(\theta) \equiv u_s + \int_{\underline{\theta}}^{\theta} x_s(\tau)d\tau$—this reparameterization preserves all conditions that a feasible function R_s must satisfy by Lemma 1. We can then state an analog of Lemma 2 for the general framework. Let $H(s)$ (respectively, $\underline{H}$) denote the distribution of the contractible state conditional on the public state being distributed according to G (respectively, $\underline{G}$).

Lemma A.1. *The designer's problem of choosing the optimal contingent menu $\mathcal{M}$ is equivalent to solving the problem*

$$\max_{\substack{\{x_s \text{ non-decreasing,} \\ u_s \geq 0\}_{s\in\mathcal{S}}}} \quad \mathbb{E}_{s\sim H}\left[\int_{\underline{\theta}}^{\bar{\theta}} \Phi_s(\theta)dx_s(\theta)\right] - \alpha^{\star}\mathbb{E}_{s\sim H}\left[u_s\right] \tag{11}$$

$$\text{subject to} \quad \mathbb{E}_{s\sim H}\left[\int_{\underline{\theta}}^{\bar{\theta}} \Psi_s(\theta)dx_s(\theta)\right] - \mathbb{E}_{s\sim \underline{H}}\left[\int_{\underline{\theta}}^{\bar{\theta}} \underline{\Psi}_s(\theta)dx_s(\theta)\right] + \mathbb{E}_{s\sim (H-\underline{H})}\left[u_s\right] \geq \tilde{c},$$

for some constant $\tilde{c} \geq 0$, and functions Φ_s, Ψ_s, $\underline{\Psi}_s : \Theta \to \mathbb{R}$, for $s \in \mathcal{S}$.

We prove Lemma A.1 in Subsection A.5.1. Intuitively, by Corollary 1, the allocation rule selected by the principal is linear in the outside option R_s. Because both the designer's and the agent's payoffs are linear in the final allocation, the designer's problem is also linear in R_s (with a linear constraint corresponding to the agent's investment-obedience constraint). Given this observation and the change of variables described previously, Lemma A.1 is a matter of bookkeeping.

Problem (11) consists of maximizing a linear functional subject to a single linear constraint. It follows that there exists an optimal allocation rule that is a convex combination of at most two extreme points of the set of non-decreasing (s-contingent) allocation rules.

Lemma A.2. *Problem* (11) *admits a solution* $(x_s^\star,\, u_s^\star)$ *such that, either*

1. $u_s^\star = 0$ *and* $x_s^\star$ *takes on at most one value other than* 0 *or* 1*, for any* $s \in \mathcal{S}$*, or*

2. $u_s^\star > 0$ *for exactly one* $s \in \mathcal{S}$*, and* $x_s^\star$ *is a step function for all* $s \in \mathcal{S}$*.*

We prove Lemma A.2 in Subsection A.5.2. Lemma A.2 implies Theorem 1′: In either case, the menu corresponding to the optimal $\{(x_s^\star,\, u_s^\star\}_{s\in\mathcal{S}}$ offers at most two options. In fact, in case *2* of Lemma A.2, the optimal menu is singleton with an option-to-own for all but one realization of the contractible state s.

A.5.1 Proof of Lemma A.1

When analyzing the designer's problem, we must take into account that the solution to the principal's problem depends both on the induced outside option function R_s and on the public state ω. We will make that dependence explicit in our notation, indexing all the variables that appear in the proof of Lemma 2 by ω.

We begin with the agent's obedience constraint (I-OB). Using the envelope formula to pin down transfers used by the principal, we can write the agent's expected payoff from participating in the stage $T = 2$ mechanism as

$$\mathbb{E}_{\omega\sim\tilde{G}}\left[\underline{u}_\omega^\star(R_{S(\omega)}) + \int_{\underline{\theta}_\omega^\star}^{\overline{\theta}_\omega^\star} x_\omega^\star(\theta;\, R_{S(\omega)})(1 - F_\omega(\theta))d\theta + \int_{\overline{\theta}_\omega^\star}^{\bar{\theta}} (\theta - \overline{\theta}_\omega^\star)dF_\omega(\theta)\right],$$

where $\tilde{G} = G$ if the agent invested and $\tilde{G} = \underline{G}$ otherwise, and $\underline{u}_\omega^\star(R_{S(\omega)})$ and $x_\omega^\star(\theta;\, R_{S(\omega)})$ denote the optimal lump-sum payment and the optimal allocation rule, respectively, used by the principal in state ω (given that the outside option function in state ω is $R_{S(\omega)}$).

In particular, when the agent has no rights, the principal allocates the good with probability one to types $\theta \geq \overline{\theta}_\omega^\star$ (and with probability zero otherwise). Define

$$\tilde{c} := c - \mathbb{E}_{\omega\sim(G-\underline{G})}\left[\int_{\overline{\theta}_\omega^\star}^{\bar{\theta}} (\theta - \overline{\theta}_\omega^\star)dF_\omega(\theta)\right]$$

as the cost of investment net of the agent's benefit from investing in the absence of any rights. By the assumption that the agent does not invest if she is not allocated any rights,

$\tilde{c} > 0$. Using the explicit solution to the principal's problem from Lemma 2, we can write the agent's obedience constraint as

$$\mathbb{E}_{\omega\sim(G-\underline{G})}\left[R_{S(\omega)}(\underline{\theta}_\omega^\star) + \sum_{(a,\,b)\in\mathcal{I}_\omega} \frac{\int_a^b R'_{S(\omega)}(\tau)d\tau \int_a^b (1-F_\omega(\theta))d\theta}{b-a}\right.$$
$$\left. + \int_{(\underline{\theta}_\omega^\star,\,\overline{\theta}_\omega^\star)\setminus\bigcup\mathcal{I}_\omega} R'_{S(\omega)}(\theta)(1-F_\omega(\theta))d\theta\right] \geq \tilde{c}.$$

Next, denoting by $W_\omega^\star(\theta) \coloneqq (V_\omega^\star(\theta) + \alpha^\star J_\omega^B(\theta))f_\omega(\theta)$ the designer's objective multiplied by the density of types, we can write her expected payoff from choosing $\{R_s\}_{s\in\mathcal{S}}$ (and conditional on the agent investing) as

$$\mathbb{E}_{\omega\sim G}\left[-\alpha^\star R_{S(\omega)}(\underline{\theta}_\omega^\star) + \sum_{(a,\,b)\in\mathcal{I}_\omega} \frac{\int_a^b R'_{S(\omega)}(\tau)d\tau \int_a^b W_\omega^\star(\theta)d\theta}{b-a} + \int_{(\underline{\theta}_\omega^\star,\,\overline{\theta}_\omega^\star)\setminus\bigcup\mathcal{I}_\omega} R'_{S(\omega)}(\theta)W_\omega^\star(\theta)d\theta\right],$$

where we have omitted the term $\mathbb{E}_{\omega\sim G}\left[\int_{\overline{\theta}_\omega^\star}^{\bar{\theta}} W_\omega^\star(\theta)d\theta\right]$ that does not depend on the designer's choice variables.

We can now change variables by letting $R_s(\theta) = u_s + \int_{\underline{\theta}}^{\theta} x_s(\tau)d\tau$, for some $u_s \geq 0$, and non-decreasing allocation rule x_s. This gives rise to the following optimization problem for the designer: Maximize over allocation rules $\{x_s\}_{s\in\mathcal{S}}$ and non-negative $\{u_s\}_{s\in\mathcal{S}}$

$$\mathbb{E}_{\omega\sim G}\left[-\alpha^\star\left(u_{S(\omega)} + \int_{\underline{\theta}}^{\underline{\theta}_\omega^\star} x_{S(\omega)}(\theta)\,d\theta\right) + \sum_{(a,\,b)\in\mathcal{I}_\omega} \frac{\int_a^b x_{S(\omega)}(\tau)\,d\tau \int_a^b W_\omega^\star(\theta)\,d\theta}{b-a}\right.$$
$$\left. + \int_{(\underline{\theta}_\omega^\star,\,\overline{\theta}_\omega^\star)\setminus\bigcup\mathcal{I}_\omega} x_{S(\omega)}(\theta)W_\omega^\star(\theta)\,d\theta\right]$$

subject to

$$\mathbb{E}_{\omega\sim(G-\underline{G})}\left[u_{S(\omega)} + \int_{\underline{\theta}}^{\underline{\theta}_\omega^\star} x_{S(\omega)}(\theta)d\theta + \sum_{(a,\,b)\in\mathcal{I}_\omega} \frac{\int_a^b x_{S(\omega)}(\tau)d\tau \int_a^b (1-F_\omega(\theta))d\theta}{b-a}\right.$$
$$\left. + \int_{(\underline{\theta}_\omega^\star,\,\overline{\theta}_\omega^\star)\setminus\bigcup\mathcal{I}_\omega} x_{S(\omega)}(\theta)(1-F_\omega(\theta))d\theta\right] \geq \tilde{c}.$$

Let $G(\omega|s)$ (respectively, $\underline{G}(\omega|s)$) denote the distribution of ω conditional on $s \in \mathcal{S}$ if the agent invested (respectively, if the agent did not invest). Since both the objective and the constraints are linear in $x_s(\theta)$, changing the order of integration and using integration by parts, we can rewrite the problem as

$$\max_{\{x_s(\theta),\, u_s \geq 0\}_{s\in\mathcal{S}}} \quad \mathbb{E}_{s\sim H}\left[\int_{\underline{\theta}}^{\bar{\theta}} \mathbb{E}_{\omega\sim G(\cdot|s)}\left[\Phi_\omega(\theta)\right] dx_s(\theta)\right] - \alpha^\star \mathbb{E}_{s\sim H}\left[u_s\right]$$

$$\text{subject to} \quad \mathbb{E}_{s\sim H}\left[\int_{\underline{\theta}}^{\bar{\theta}} \mathbb{E}_{\omega\sim G(\cdot|s)}\left[\Psi_\omega(\theta)\right] dx_s(\theta)\right] - \mathbb{E}_{s\sim \underline{H}}\left[\int_{\underline{\theta}}^{\bar{\theta}} \mathbb{E}_{\omega\sim \underline{G}(\cdot|s)}\left[\Psi_\omega(\theta)\right] dx_s(\theta)\right] + \mathbb{E}_{s\sim (H-\underline{H})}\left[u_s\right] \geq \tilde{c},$$

where

$$\Phi_\omega(\theta) \coloneqq -\alpha^\star \left(\underline{\theta}^\star_\omega - \theta\right)_+ + \sum_{(a,\,b)\in\mathcal{I}_\omega} \frac{(b-\max\{a,\theta\})_+ \int_a^b W^\star_\omega(\theta) d\theta}{b-a} + \sum_{[a,\,b]\in\mathcal{I}^c_\omega} \mathbf{1}_{\{\theta\leq b\}} \left(\int_{\max\{a,\,\theta\}}^b W^\star_\omega(\tau) d\tau\right),$$

$$\Psi_\omega(\theta) \coloneqq \left(\underline{\theta}^\star_\omega - \theta\right)_+ + \sum_{(a,\,b)\in\mathcal{I}_\omega} \frac{(b-\max\{a,\theta\})_+ \int_a^b (1-F_\omega(\theta)) d\theta}{b-a} + \sum_{[a,\,b]\in\mathcal{I}^c_\omega} \mathbf{1}_{\{\theta\leq b\}} \left(\int_{\max\{a,\,\theta\}}^b (1-F_\omega(\tau)) d\tau\right).$$

We obtain the lemma by setting, for any $s \in \mathcal{S}$, $\Phi_s(\theta) \coloneqq \mathbb{E}_{\omega\sim G(\cdot|s)}\left[\Phi_\omega(\theta)\right]$, $\Psi_s(\theta) \coloneqq \mathbb{E}_{\omega\sim G(\cdot|s)}\left[\Psi_\omega(\theta)\right]$, and $\underline{\Psi}_s(\theta) \coloneqq \mathbb{E}_{\omega\sim \underline{G}(\cdot|s)}\left[\Psi_\omega(\theta)\right]$.

A.5.2 Proof of Lemma A.2

By Lemma A.1, the designer's problem is to maximize a linear functional subject to a single linear constraint. Thus, there exists a solution that is a convex combination of at most two extreme points.[42] Extreme points in the space of (non-decreasing) allocation rules are step functions of the form $\mathbf{1}_{\theta\geq\theta^\star}$. Moreover, the set of extreme points of a product of spaces is the product of extreme points of these spaces. Thus, for each s, the optimal x_s can be written as a two-step function, and in particular its image may contain at most one value other than 0 or 1.

The optimal solution to problem (11) must maximize the Lagrangian, with Lagrange

[42] Formally, and as summarized by Kang (2023), this follows from the results of Bauer (1958) and Szapiel (1975).

multiplier γ on the investment-obedience constraint:[43]

$$\mathbb{E}_{s\sim H}\left[\int_{\underline{\theta}}^{\bar{\theta}}\Phi_s(\theta)dx_s(\theta)\right]+\gamma\left(\mathbb{E}_{s\sim H}\left[\int_{\underline{\theta}}^{\bar{\theta}}\Psi_s(\theta)dx_s(\theta)\right]-\mathbb{E}_{s\sim\underline{H}}\left[\int_{\underline{\theta}}^{\bar{\theta}}\underline{\Psi}_s(\theta)dx_s(\theta)\right]\right)$$
$$-\alpha^{\star}\mathbb{E}_{s\sim H}[u_s]+\gamma\mathbb{E}_{s\sim(H-\underline{H})}[u_s].$$

A consequence of complementarity-slackness conditions (paired with the requirement that each u_s must maximize the Lagrangian) is that either $\gamma = 0$, or γ is pinned down by

$$\max_{s\in\mathcal{S}}\{(\gamma-\alpha^{\star})h(s)-\gamma\underline{h}(s)\}=0, \tag{12}$$

where $h(s)$ and $\underline{h}(s)$ are probability mass functions corresponding to H and $\underline{H}$, respectively (using the fact that Ω, and hence also $\mathcal{S}$, is finite). In the latter case, u_s can be any positive number for all $s \in \mathcal{S}^{\star}$, where $\mathcal{S}^{\star}$ is the set of maximizers of the above expression (with $u_s = 0$ for $s \notin \mathcal{S}^{\star}$).

There are two cases to consider. It could be that $u_s^{\star} = 0$ for all $s \in \mathcal{S}$—this corresponds to case *(i)* in Lemma A.2. In the opposite case, $u_s^{\star} > 0$ for some $s \in \mathcal{S}^{\star}$ and γ is pinned down by equation (12). Then, any $u_s \geq 0$, for $s \in \mathcal{S}^{\star}$, maximizes the Lagrangian. As long as $h \neq \underline{h}$, we can pick cutoff allocation rules $x_s(\theta)$ maximizing the Lagrangian that do not satisfy the obedience constraint when paired with $u_s = 0$,[44] and then satisfy the obedience constraint by picking a high enough $u_{s^{\star}}$ for some $s^{\star} \in \mathcal{S}^{\star}$ such that $h(s^{\star}) > \underline{h}(s^{\star})$ (with $u_s = 0$ for all $s \neq s^{\star}$). Indeed, $h(s^{\star}) > \underline{h}(s^{\star})$ (i.e., the contractible state $s^{\star}$ happens with higher probability when investment takes place) guarantees that the agent will have a strict incentive to invest for high enough monetary payment paid out in contractible state $s^{\star}$. This corresponds to case *(ii)* in Lemma A.2. Finally, consider the case in which $h = \underline{h}$, i.e., the contractible signal has the same distribution regardless of the investment decision. In this case, it is always optimal to set $u_s^{\star} = 0$ for all $s \in \mathcal{S}$, since the choice of u_s does not influence the agent's investment decision—corresponding again to case *(i)* of Lemma A.2.

A.6 Proofs of Corollaries 3 and 4

Corollary 3 follows from the proof of Lemma A.2 by observing that if the investment-obedience constraint is dropped from the designer's optimization problem ($\gamma = 0$ and the obedience constraint is satisfied at the unconstrained solution), then the optimum is attained by an extreme point $x_s(\theta) = \mathbf{1}_{\theta \geq p_s}$ and $u_s = 0$, for some $p_s \in \mathbb{R}$, for all $s \in \mathcal{S}$. This cor-

[43]Existence of a Lagrange multiplier follows from Theorem 2.165 in Bonnans and Shapiro (2000).
[44]Such x_s must exist, as otherwise we could not have a solution with some $u_s > 0$.

responds to offering the agent an option-to-own with a price indexed by the contractible state s.

Corollary 4 follows from the proof of Lemma A.2. If investment is non-contractible, i.e., when $h = \underline{h}$, the conclusion that $u_s^\star = 0$ for all $s \in \mathcal{S}$ implies that $y_s > 0$ in the corresponding optimal menu $M_s^\star$.

Suppose that investment is contractible. If c is high enough, it will not be possible to satisfy the investment-obedience constraint with $R_s(\underline{\theta}) = 0$ for all $s \in \mathcal{S}$. Thus, it must be that $u_s^\star > 0$ in the solution to the designer's problem, and hence equation (12) must hold. Contractibility of investment means that any $s \in S(\text{supp}(G))$ maximizes the expression in (12) and $\gamma^\star = \alpha^\star$. It then follows that two mechanisms with the same $\mathbb{E}_{s\sim H}[u_s^\star]$ are payoff-equivalent for the designer and lead to the same investment incentives for the agent (only the expected payment conditional on investment matters). Thus, we can construct a solution in which $u_s = T$ for some T and all $s \in S(\text{supp}(G))$.

A.7 Proof of Proposition 1

We begin with a technical lemma.

Lemma A.3. *Under the assumptions of Proposition 1, (i) $\mathcal{W}_\omega = \overline{\mathcal{W}}_\omega$ on $[\underline{\theta}_\omega^\star, \overline{\theta}_\omega^\star]$, (ii) $\overline{\mathcal{W}}_\omega'(\theta) \leq \alpha$ for all $\theta \leq \underline{\theta}_\omega^\star$, with equality at $\theta = \underline{\theta}_\omega^\star$, and (iii) $\overline{\theta}_\omega^\star$ is the global maximum of $\mathcal{W}_\omega$.*

Proof of Lemma A.3. We drop the subscript ω to simplify the exposition. We first prove that $W(\theta) = (V(\theta) + \alpha J^B(\theta))f(\theta)$ is non-decreasing on $[\underline{\theta}^\star, \overline{\theta}^\star]$. It suffices to show that, for all $\theta \in [\underline{\theta}^\star, \overline{\theta}^\star]$, $\frac{V'(\theta)+2\alpha}{\alpha}f(\theta) + \frac{V(\theta)+\alpha\theta}{\alpha}f'(\theta) \geq 0$. Using the definition of $\underline{\theta}^\star$ and $\overline{\theta}^\star$ given in Proposition 1, for all $\theta \in [\underline{\theta}^\star, \overline{\theta}^\star]$, we have $\frac{1-F(\theta)}{f(\theta)} \geq \frac{V(\theta)+\alpha\theta}{\alpha} \geq -\frac{F(\theta)}{f(\theta)}$. If $f'(\theta)$ is negative, we have $\frac{V'(\theta)+2\alpha}{\alpha}f(\theta) + \frac{V(\theta)+\alpha\theta}{\alpha}f'(\theta) \geq 2f(\theta) - \frac{F(\theta)}{f(\theta)}f'(\theta) \geq 0$, where the second inequality follows from the monotonicity of the seller virtual surplus. When $f'(\theta)$ is positive, we have $\frac{V'(\theta)+2\alpha}{\alpha}f(\theta) + \frac{V(\theta)+\alpha\theta}{\alpha}f'(\theta) \geq 2f(\theta) + \frac{1-F(\theta)}{f(\theta)}f'(\theta) \geq 0$, where the second inequality follows from the monotonicity of the buyer virtual surplus.

Next, we prove that $W(\theta) \leq -\alpha$ for $\theta \leq \underline{\theta}^\star$, that is, $\left[V(\theta) + \alpha\left(\theta - \frac{1-F(\theta)}{f(\theta)}\right)\right]f(\theta) \leq -\alpha$. Indeed, we have

$$\left[V(\theta) + \alpha\left(\theta - \frac{1-F(\theta)}{f(\theta)}\right)\right]f(\theta) = \underbrace{\left[V(\theta) + \alpha\left(\theta + \frac{F(\theta)}{f(\theta)}\right)\right]}_{\leq 0}f(\theta) - \alpha \leq -\alpha.$$

The same calculation shows that $W(\underline{\theta}^\star) = -\alpha$, and $W(\theta) \geq -\alpha$ for $\theta \geq \underline{\theta}^\star$.

Overall, we have shown that $\mathcal{W}(\theta) = \int_\theta^{\bar{\theta}} W(\tau)d\tau$ has a slope higher than α for $\theta \leq \underline{\theta}^\star$ and lower than α for $\theta \geq \underline{\theta}^\star$, is concave on $[\underline{\theta}^\star, \overline{\theta}^\star]$, and has a global maximum at $\overline{\theta}^\star$ (since

$W(\theta)$ crosses zero once from below at $\overline{\theta}^{\star}$). It follows that $\mathcal{W}$ is equal to its concave closure on $[\underline{\theta}^{\star}, \overline{\theta}^{\star}]$. Moreover, $\overline{\mathcal{W}}'(\theta) \le \alpha$ for $\theta \le \underline{\theta}^{\star}$, with equality at $\theta = \underline{\theta}^{\star}$. Finally, $\underline{\theta}^{\star}$ and $\overline{\theta}^{\star}$, as defined in Proposition 1, correspond to the $\underline{\theta}^{\star}$ and $\overline{\theta}^{\star}$ defined in Section 3.1. □

Given Lemma A.3, Proposition 1 follows directly from Lemma 2. The collection $\mathcal{I}_{\omega}$ is empty, so there are no ironing intervals: $U_{\omega}(\theta)$ coincides with $R(\theta)$ on $[\underline{\theta}^{\star}_{\omega}, \overline{\theta}^{\star}_{\omega}]$; below $\underline{\theta}^{\star}_{\omega}$, $x^{\star}_{\omega} = 0$, so U_{ω} is constant, equal to $R(\underline{\theta}^{\star}_{\omega})$; and above $\overline{\theta}^{\star}_{\omega}$, $x^{\star}_{\omega} = 1$, giving the expression for U_{ω} from Proposition 1. Finally, the formulas for $\underline{\theta}^{\star}_{\omega}$ and $\overline{\theta}^{\star}_{\omega}$ are special cases of the general definitions, simplified under the observations made in Lemma A.3.

A.8 Supplementary material for Section 4.4

Following the proof of Theorem 1′, relying on Proposition 1 and the simplifying assumptions in Section 4.4, we obtain

$$\Phi_{\omega}(\theta) = -\alpha^{\star}\left(\underline{\theta}^{\star}_{\omega} - \theta\right)_{+} + \mathbf{1}_{\{\theta \le \overline{\theta}^{\star}_{\omega}\}}\left(\int_{\max\{\underline{\theta}^{\star}_{\omega}, \theta\}}^{\overline{\theta}^{\star}_{\omega}} W^{\star}_{\omega}(\tau)d\tau\right),$$

$$\Psi_{\omega}(\theta) = \left(\underline{\theta}^{\star}_{\omega} - \theta\right)_{+} + \mathbf{1}_{\{\theta \le \overline{\theta}^{\star}_{\omega}\}}\left(\int_{\max\{\underline{\theta}^{\star}_{\omega}, \theta\}}^{\overline{\theta}^{\star}_{\omega}} (1 - F_{\omega}(\tau))d\tau\right).$$

We will consider separately the cases of contractible and non-contractible investment.

In the non-contractible case, we know the optimal x takes the form

$$x(\theta) = \begin{cases} 0 & \theta < \theta^{\star}_{1}, \\ y & \theta^{\star}_{1} \le \theta < \theta^{\star}_{2}, \\ 1 & \theta \ge \theta^{\star}_{2}. \end{cases}$$

As in the proof of Theorem 1′, we can study the behavior of the Lagrangian which now takes the form $\int_{\underline{\theta}}^{\bar{\theta}} \left(\Phi(\theta) + \gamma\Psi(\theta) - \gamma\underline{\Psi}(\theta)\right) dx(\theta)$, where we have used the fact that $u^{\star} = 0$ in the non-contractible-investment case. Since the optimal x must maximize the Lagrangian, both $\theta^{\star}_{1}$ and $\theta^{\star}_{2}$ must satisfy the first-order condition:

$$\Phi'(\theta^{\star}) + \gamma\Psi'(\theta^{\star}) - \gamma\underline{\Psi}'(\theta^{\star}) \overset{(FOC)}{=} 0.$$

Here, $\overset{(FOC)}{=}$ is short-hand notation for equality at interior points at which the left-hand side is differentiable, and for the appropriate weak inequalities at boundary points $\underline{\theta}$ and $\bar{\theta}$. At points of non-differentiability, with slight abuse of notation, we can interpret the condition as saying that the left derivative of the left-hand side must be non-negative while the right derivative of the left-hand side must be non-positive. Writing out the first-order condition more explicitly yields formula (7). If there is only a single point $\theta^\star$ satisfying the first-order condition, then offering a singleton menu with option-to-own with price $p = \theta^\star$ is optimal. If instead the optimal menu contains two options, then $\theta_1^\star < \theta_2^\star$ must both satisfy the first-order condition, while prices in the optimal menu $M^\star = \{(1, p), (y, p')\}$ are given by $p' = y\theta_1^\star$ and $p = \theta_2^\star - y(\theta_2^\star - \theta_1^\star)$. Equations (4) and (5) are special cases of (7), where additionally in case of (4) we have divided both sides of the equation by $\mathbb{P}(\theta^\star \in [\underline{\theta}_\omega^\star, \overline{\theta}_\omega^\star])$. We can do this because it is without loss of generality to restrict the set of candidate $\theta^\star$ to the closure of the set $\{\theta^\star \in \Theta : \mathbb{P}_{\omega\sim G}(\theta^\star \in [\underline{\theta}_\omega^\star, \overline{\theta}_\omega^\star]) > 0\}$ since the Lagrangian is constant in $\theta^\star$ outside this range.

In the contractible case, we know from the proof of Theorem 1 that—as long as the cost of investment c is sufficiently high—we must have $u^\star = 0$ and $\gamma = \alpha^\star$. Because the agent will hold no rights conditional on not investing, the Lagrangian in this case takes the simpler form $\int_{\underline{\theta}}^{\bar{\theta}} (\Phi(\theta) + \alpha^\star\Psi(\theta))\, dx(\theta)$. We also know that the optimal $x^\star(\theta) = \mathbf{1}_{\theta\geq\theta^\star}$, and since $x^\star$ maximizes the Lagrangian, $\theta^\star$ must satisfy the first-order condition $\Phi'(\theta^\star) + \alpha^\star\Psi'(\theta^\star) \overset{(FOC)}{=} 0$. A straightforward transformation of the first-order condition yields formula (6).

B Supplementary material for Section 5

B.1 Dynamic resource allocation

In this appendix, we consider a numerical example for Subsection 5.1. Abusing notation slightly, we assume that the state ω and the agent's type θ (conditional on investment) are independent random variables that are distributed uniformly on $[0, 1]$, and that the resource is useless to the agent ($\theta = 0$) absent investment.[45] Investment is not contractible.

By Proposition 1, at time $T = 2$, the regulator allocates the resource to all agent types above $\frac{\omega+\alpha}{1+2\alpha}$, buys back any rights from agent types below $\frac{\omega}{1+2\alpha}$ by offering them a cash payment, and lets the remaining types execute their optimal right. There exist cutoffs $\underline{c}$ and $\bar{c}$ satisfying $0 < \underline{c} < \bar{c}$ such that: If $c \leq \underline{c}$, then investment takes place even when the agent has no rights; and if $c = \bar{c}$, then only a full property right incentivizes investment. We assume that $c \in (\underline{c}, \bar{c})$ and analyze the optimal property rights in three cases.

Case $\alpha = \alpha^\star = 1$: When the regulator maximizes the sum of allocative efficiency and revenue in both periods, the optimal license is simply a renewable lease with a price p. The optimal price p makes the agent indifferent between investing or not investing.

[45] Formally, to make the example consistent with our general model, we could specify the public state as $\tilde{\omega} = (\omega, \mathbf{1}_i)$, where $\mathbf{1}_i = 1$ if and only if investment took place, and the distribution of θ is pinned down by $\mathbf{1}_i$.

Case $\alpha = 1$, $\alpha^\star = 0$: When the regulator is concerned with efficiency ex-ante but attaches a positive weight to $T = 2$ revenue, the optimal license is a partial property right $\{(y,\, 0)\}$, where the probability y makes the agent indifferent between investing or not investing.

Case $\alpha = \alpha^\star = 0$: In this case, the regulator maximizes efficiency in both periods.[46] The optimal mechanism at $T = 2$ takes the form of allocating the good to agent types above ω and buying out any rights for the remaining types—which reduces to a standard VCG mechanism when the agent has no rights. Thus, by Rogerson (1992), it is optimal for the regulator to assign no rights to the agent in this case.

Supporting calculations. We first determine the bounds $\bar{c}$ and $\underline{c}$. When no rights are assigned to the agent, investment is taken when $\int_0^1 \left(\int_{\frac{\omega+\alpha}{1+2\alpha}}^1 \left(\theta - \frac{\omega+\alpha}{1+2\alpha}\right) d\theta\right) d\omega \geq c$, or, equivalently, $\underline{c} := \frac{1}{6}(1+2\alpha)\left[\left(\frac{1+\alpha}{1+2\alpha}\right)^3 - \left(\frac{\alpha}{1+2\alpha}\right)^3\right] \geq c$. Under a full property right, investment is taken when $\int_0^1 \left(\int_{\frac{\omega}{1+2\alpha}}^1 \left(\theta - \frac{\omega}{1+2\alpha}\right) d\theta\right) d\omega \geq c$, or, equivalently, $\bar{c} := \frac{1}{6}(1+2\alpha)\left[1 - \left(\frac{2\alpha}{1+2\alpha}\right)^3\right] \geq c$.

Notice that in the case $\alpha = 0$, the principal uses a VCG mechanism (since $\underline{\theta}^\star_\omega = \overline{\theta}^\star_\omega = \omega$), which proves the claim for the case $\alpha = \alpha^\star = 0$. From now on, we assume that $\alpha = 1$.

Recall from Appendix A.8 that $\Phi_\omega(\theta)$ and $\Psi_\omega(\theta)$ are defined as the values of the functions $\Phi(\theta)$ and $\Psi(\theta)$, respectively, conditional on a given state ω. In the current application, we have

$$\Phi'_\omega(\theta) + \gamma\Psi'_\omega(\theta) = \begin{cases} \alpha^\star & \theta < \frac{\omega}{3}, \\ (\gamma - 1 - 2\alpha^\star)\,\theta + \omega + \alpha^\star - \gamma & \theta \in \left(\frac{\omega}{3},\, \frac{\omega+1}{3}\right), \\ 0 & \theta > \frac{\omega+1}{3}. \end{cases}$$

Therefore, using the assumption that G is uniform on $[0,\, 1]$,

$$\int_0^1 \left[\Phi'_\omega(\theta) + \gamma\Psi'_\omega(\theta)\right] dG(\omega) = \begin{cases} \int_0^{3\theta} \left[(\gamma - 1 - 2\alpha^\star)\,\theta + \omega + \alpha^\star - \gamma\right] d\omega + \alpha^\star(1 - 3\theta) & \theta < 1/3, \\ \int_{3\theta-1}^1 \left[(\gamma - 1 - 2\alpha^\star)\,\theta + \omega + \alpha^\star - \gamma\right] d\omega & \theta \in (1/3,\, 2/3), \\ 0 & \theta > 2/3, \end{cases}$$

and

$$\int_0^1 \left[\Phi''_\omega(\theta) + \gamma\Psi''_\omega(\theta)\right] dG(\omega) = \begin{cases} \theta\,(6\gamma - 12\alpha^\star + 3) - 3\gamma & \theta < 1/3, \\ -\theta\,(6\gamma - 12\alpha^\star + 3) + 1 - 7\alpha^\star + 5\gamma & \theta \in (1/3,\, 2/3), \\ 0 & \theta > 2/3. \end{cases}$$

From now on, we will take a look at the two cases, $\alpha^\star = 1$ and $\alpha^\star = 0$, separately.

Case $\alpha^\star = 1$. Specializing to the case $\alpha^\star = 1$, in the interval $[0, 1/3]$, the Lagrangian is

[46] Formally, since we assumed $\alpha > 0$, we consider the limit of solutions as $\alpha \to 0$.

concave and its derivative is strictly positive at 0. The derivative at $\theta = 1/3$ is $1/2 - (2/3)\gamma$. Then, on [1/3, 2/3], the second derivative changes from $3\gamma - 3$ to γ. The first derivative at 2/3 is 0. If γ is above 3/4, then the derivative at 1/3 is negative, and it must remain negative for all $\theta \geq 1/3$ because it must be 0 at 2/3. Thus, in this case, we have a global maximum that lies in $(0, 1/3]$. If γ is below 3/4, then since the function is concave in [0, 1/3], the first derivative must be positive on that interval. And since the derivative is positive at 1/3 but zero at 2/3, while the function changes from concave to convex, we must have now a unique global maximum that lies in [1/3, 2/3]. Thus, we have shown that, in all cases, an option-to-own is optimal. As γ changes from 0 to ∞, the optimal price takes all values between 0 and 2/3 (note also that if a price 2/3 is optimal, then any price between 2/3 and 1 is also optimal). Of course, the optimal price p must then satisfy the investment-obedience constraint with equality, that is, $(1 - G(3p))\,\bar{c} + \int_{3p-1}^{3p} \int_p^1 (\theta - p)_+ \, d\theta \, dG(\omega) + G(3p-1)\,\underline{c} = c$. As p varies from 0 to 2/3, the left-hand side takes on any value between $\underline{c}$ and $\bar{c}$.

Case $\alpha^\star = 0$. We now specialize to the case $\alpha^\star = 0$. If $\gamma \geq 1$, then on $[0, 1/3]$ the function is concave, and thus decreasing. On $[1/3, 2/3]$, the function is convex, and the first derivative is negative. Thus, the function is globally decreasing. It is thus optimal to give a full property right. However, except for the case $c = \bar{c}$, this would make the investment-obedience constraint slack, requiring γ to be 0.

If $\gamma \in [1/4, 1)$, then the first derivative at 1/3 is still negative. On [0, 1/3], the function is first concave and then convex, starting with a zero derivative, and ending with a negative derivative. Thus, the function is decreasing in this region. On [1/3, 2/3], the function is first convex and then concave, starting with a negative derivative, and ending with a zero derivative. We conclude that there are two local maxima: one at 0 and one at 2/3.

Finally, suppose that $\gamma < 1/4$, so that the first derivative at 1/3 is positive. Now, on $[0, 1/3]$, the function is first concave and then convex, starting with a zero derivative, and ending with a positive derivative. Thus, the function is first decreasing and then increasing in this region. On $[1/3, 2/3]$, the function is first convex and then concave, starting with a positive derivative, and ending with a zero derivative. Thus, we conclude again that there are two local maxima: one at 0 and one at 2/3.

Because the function is constant on [2/3, 1], whenever 2/3 is optimal, so is 1. We conclude that, regardless of the value of γ, the function is maximized either at 0 or at 1; however, this will not allow us to satisfy the investment-obedience constraint except for the boundary cases $c = \underline{c}$ and $c = \bar{c}$. Thus, in all other cases, it must be that γ takes a value that makes both 0 and 1 global maxima, in which case the designer can satisfy the investment-obedience constraint with equality by randomizing over full right and no property right with some probability y: $y\bar{c} + (1 - y)\underline{c} = c$. This concludes the proof for this case.

B.2 Regulating a rental market

The designer is a policymaker and the principal is a company leasing a rental unit to an agent (who could be a residential tenant or a business owner). The agent occupies the unit at time $T = 0$, and decides whether to invest (e.g., whether to take good care of the apartment or install specialized equipment in the office space). We assume that investment results in a higher value θ for staying in the unit for another lease term $T = 2$, but is not observable. The state ω is the price the rental company could receive by leasing to a new tenant (the market rental price). The rental company maximizes revenue: $V_\omega(\theta) = -\omega$ and $\alpha = 1$. The designer, on the other hand, is concerned with efficiency: $V^\star_\omega(\theta) = \theta - \omega$ and $\alpha^\star = 0$.

In this application, the menu of rights chosen by the policymaker captures regulation of a private rental market. The rental company has monopoly power over the tenant, since the tenant makes a sunk investment (and moving is implicitly assumed to be costly). This introduces a potential inefficiency, as the rental company might dictate prices above the market rate, which could further disincentivize investment.

Theorem 1′ and Proposition 1 predict an important role for the renewable-lease contract. To derive tighter predictions, let us further assume that, absent investment, the agent's value for staying in the rental unit is drawn from uniform distribution on $[0,\, 1]$, investing increases the value by a constant $\Delta > 0$, and that $\text{supp}(G) \subseteq [\Delta,\, 1 - \Delta]$.

Consider first the case when property rights are fully state-contingent, i.e., the regulation can condition on the market rental rate ω. Then, the optimal regulation takes the form of a renewable lease at a price $p = \omega - \gamma\Delta$, where γ is the Lagrange multiplier on the agent's investment-obedience constraint. This means that the agent is allowed to renew the lease at a price that is (potentially) discounted relative to the market price, and that the discount is larger when investment is more difficult to incentivize. If we restrict attention to parameters for which the agent's investment is socially efficient (i.e., if the cost c is sufficiently small), then $\gamma = 0$—the renewable-lease price is in fact *equal* to the market price which allows the designer to induce the VCG mechanism. The regulation has bite because the rental company would charge a higher price to the agent, exploiting its monopoly position.

If regulation cannot be conditional on ω, then the first-order condition determining the optimal price in the renewable-lease contract becomes $p = \mathbb{E}_{\omega \sim G}[\omega \mid p \in [\underline{\theta}^\star_\omega,\, \overline{\theta}^\star_\omega]] - \gamma\Delta$. Thus, the designer is trying to achieve a similar outcome but this time targeting the *expected* market rental rate, where the expectation is conditional on the market price ω being in a certain range that is endogenous to the choice of p. In particular, if the market rental rate ω is high, then $p < \underline{\theta}^\star_\omega$ and the rental company prefers to pay the agent to leave over forgoing the market rental rate. Similarly, if ω is low, then $p > \overline{\theta}^\star_\omega$ and the rental company offers to renew the agent's lease at a price strictly below p. Thus, the price p set by the designer only has bite when the market rental rate is in the intermediate region. In general, it is no longer the case that $\gamma = 0$, even when investment is socially efficient.

Supporting calculations. Using Proposition 1, we can pin down the interval $[\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega]$ on which the outside option constraint holds with equality:

$$\omega = \underline{\theta}^\star_\omega + \frac{F(\underline{\theta}^\star_\omega)}{f(\underline{\theta}^\star_\omega)} \text{ and } \omega = \overline{\theta}^\star_\omega - \frac{1 - F(\overline{\theta}^\star_\omega)}{f(\overline{\theta}^\star_\omega)}.$$

Due to our assumption that $\text{supp}(G) \subseteq [\Delta, 1-\Delta]$, we have that $\Delta \le \underline{\theta}^\star_\omega \le \overline{\theta}^\star_\omega \le 1$. This in turn implies that $\mathbb{E}_{\omega \sim (\underline{G} - G)}[F_\omega(\theta)] = \Delta$ for all $\theta \in [\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega]$. Then,

$$\Phi'_\omega(\theta) + \gamma \Psi'_\omega(\theta) = \begin{cases} 0 & \theta < \underline{\theta}^\star_\omega, \\ -[\theta - \omega + \gamma\Delta] & \theta \in (\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega), \\ 0 & \theta > \overline{\theta}^\star_\omega. \end{cases}$$

In case of state-contingent rights, we have a unique maximum at $\theta^\star = \omega - \gamma\Delta$. By Rogerson (1992), if investment is socially efficient, then setting $\gamma = 0$ incentivizes the agent to invest, and hence an option-to-own with price ω is optimal. In case of non-contingent rights, the optimal p must satisfy the first-order condition: $p = \mathbb{E}_{\omega \sim G}\left[\omega \mid p \in [\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega]\right] - \gamma\Delta$.

B.3 Patent policy

A classical economic question is how to reward and incentivize innovation. For example, Wright (1983) analyzed the choice between patents, prizes, and direct contracting, and showed that each of these alternatives can be an effective intervention depending on information available to a regulator. Other papers (see, for example, Klemperer, 1990; Gilbert and Shapiro, 1990; Gallini, 1992) studied the trade-off between the length and breadth of patents. While our baseline model cannot capture the notion of patent length, we can ask how the designer can optimally use patent breadth (allocation x in our model) and monetary payments (transfer t in our model) to induce socially efficient investment.

In this application, the agent is a firm making a costly investment at $T = 1$ in a new technology. The principal is a patent office deciding whether the agent should have monopoly rights to the invention. The designer corresponds to a regulator designing patent policy. Let k be the marginal cost of production for the firm conditional on investment, and—for simplicity—suppose that market demand for the product is given by $D(p) = 1 - p$. If the firm is a monopolist, it chooses to produce $(1-k)/2$, the price is $(1+k)/2$, and the profit is $(1-k)^2/4$. If the firm is not granted a monopoly, we assume there is perfect competition at the marginal cost k; the firm will not make profits, total production will be $1 - k$, and the price will be k. Thus, the utility of the firm from obtaining a monopoly at $T = 2$ is $\theta \equiv (1 - k)^2/4$. The designer attempts to maximize total surplus given by the sum of consumer surplus and firm profits, while the principal places a potentially higher weight $\omega \geq 1$ on consumer surplus.[47] A simple calculation shows that this scenario corresponds to $V_\omega(\theta) \equiv \theta(1 - (3/2)\omega)$ and $V^\star_\omega(\theta) \equiv V_1(\theta)$, for all $\omega \in \Omega$.

A full property right in this application gives the innovator a monopoly in the market for the invention. However, this hurts consumer surplus. In particular, the principal's objective $V_\omega(\theta)$ is *decreasing* in θ. This is because granting a monopoly right to the firm is particularly inefficient when the costs of production are low (θ is high). Our question in this context is whether investment can be incentivized by giving the innovator a partial right; an intermediate $x \in (0, 1)$ can be interpreted either as awarding the monopoly right with some probability (e.g., the regulator sets a review standard for patent applications) or as the patent breadth (e.g., the degree of protection against substitute products).[48] Additionally, if investment is observable, then the regulator can offer a direct cash prize for the innovation. To simplify our analysis, we assume that the distribution of costs (conditional on investment) does not depend on ω, and that the density of θ is non-decreasing (we later comment on how our results change without that last assumption).[49]

First, we suppose that the patent office has access to a transparent and credible way of assessing the usefulness of the invention—corresponding to our assumption that investment is observable and contractible. Then—as long as the weight on revenue is not too high—the optimal property right will include a cash prize for the discovery. Furthermore, if the support of ω is lower bounded by $(4/3)\alpha + (2/3)$—that is, if the principal puts sufficiently more weight on consumer surplus than on revenue—she will always prefer to buy out any rights of the innovator with cash.[50] In that case, the optimal property right is a cash payment conditional on investment.

We now suppose that investment is not observable and that the patent office cannot pay the firm. Under the same assumption that the support of ω is lower bounded by $(4/3)\alpha + (2/3)$, the optimal contract takes the form of allocating a monopoly right free

[47] For example, the principal could have redistributive preferences as in Dworczak ⓡ al. (2021).

[48] This also resonates with previous work demonstrating how the flexible allocation of market power and monopoly rights can improve innovation policy relative to simply awarding innovators full monopoly rights in the form of a patent (see, in particular, Hopenhayn et al., 2006 and Weyl and Tirole, 2012).

[49] For large enough ω, $V_\omega(\theta)$ is negative and decreasing; the assumption of non-decreasing density ensures that $V_\omega(\theta)f(\theta)$ preserves that property. As the solution to the principal's problem (P$'$) reveals, the structure of the ironing intervals depends on the monotonicity of the original objective *multiplied* by the density. This is a consequence of the fact that the outside option constraint does *not* depend on the distribution of types.

[50] See Kremer (1998) for historical cases of patent buyouts and a detailed analysis of how governments can determine the buyout price.

of charge with some fixed probability y (or with breadth y) that makes the investment-obedience constraint bind. Intuitively, when ω is high, conditional on the new technology being already developed, the patent office would prefer not to grant a monopoly right, and she is particularly reluctant to grant it when costs of production are low (because consumer surplus under perfect competition is particularly high in this case). However, it is firms with low production costs that have a higher willingness to pay for obtaining the monopoly right; hence, the best the patent office can do is allocate the monopoly right with a probability that does not depend on production costs.

When the principal puts a sufficiently high weight on revenue (relative to the realized ω) or when the density of θ is not non-decreasing, it might become optimal to "sell" the monopoly rights to firms with low costs. In that case, the optimal regulation takes a more complicated form, potentially specifying a fee that a firm applying for a patent may choose to pay to increase the probability of obtaining the patent (a type of "fast track" procedure). Allowing the firm to purchase a patent may be the cheapest way to incentivize investment because it promises the innovator a higher probability of obtaining monopoly rights precisely when these monopoly rights are most valuable (costs are low).

Supporting calculations. We start with two general observations. First, dropping the dependence on ω in the notation, suppose that $\mathcal{W}(\theta) \leq \mathcal{W}(\underline{\theta}) \cdot (\bar{\theta} - \theta)/(\bar{\theta} - \underline{\theta})$, that is, $\mathcal{W}$ lies everywhere below its concave closure. Following the proof of Lemma 2, we can then conclude that there are three cases:

1. If $\mathcal{W}(\underline{\theta}) = \int_{\underline{\theta}}^{\bar{\theta}} W(\tau)d\tau < -\alpha$, then $x^\star(\theta) \equiv 0$, and $u^\star = R(\bar{\theta})$, that is, the principal buys out all rights with money.

2. If $\mathcal{W}(\underline{\theta}) \in [-\alpha, 0]$, then $x^\star(\theta) = (R(\bar{\theta}) - R(\underline{\theta}))/(\bar{\theta} - \underline{\theta})$, and $u^\star = R(\underline{\theta})$, that is, the allocation rule is constant.

3. If $\mathcal{W}(\underline{\theta}) > 0$, then $x^\star(\theta) \equiv 1$, and $u^\star = R(\underline{\theta})$, that is, the agent always gets the good.

Second, it is easy to modify our solution to handle the case in which the principal is not allowed to pay the agent (assuming that the designer is then constrained to choose $R_s(\underline{\theta}) = 0$ for any s). We simply set $\underline{u}_s$ to 0, for any $s \in \mathcal{S}$, in the proof of Theorem 1′, which implies that the solution is modified by setting $\underline{\theta}^\star_\omega = \underline{\theta}$, for any ω—the principal never buys out the agent's rights. Then, the solution to case 1 above is the same as the solution to case 2.

Let us now apply these observations to the patent application. We have $V_\omega(\theta) = \theta(1-\frac{3}{2}\omega)$ and $V^\star_\omega(\theta) = -\frac{1}{2}\theta$. To simplify notation, let $\beta_\omega \coloneqq -(1 - \frac{3}{2}\omega)$. Recall also that $\theta \equiv \frac{1}{4}(1-k)^2$ so we can assume that θ is distributed on $[0, 1/4]$. To verify that $\mathcal{W}(\theta) \leq \mathcal{W}(\underline{\theta}) \cdot (\bar{\theta} - \theta)/(\bar{\theta} - \underline{\theta})$, as in the observation we made above, we have to check that, for all $\theta \in [0, 1/4]$,

$\int_{\theta}^{\frac{1}{4}} \left[-\beta_\omega \tau + \alpha J^B(\tau)\right] dF(\tau) \leq -(1-4\theta)\beta_\omega \mathbb{E}[\theta]$. Rewriting, we obtain, $\beta_\omega \int_{\theta}^{\frac{1}{4}} \tau dF(\tau) - (1-4\theta)\mathbb{E}[\theta] \geq \alpha\theta(1-F(\theta))$. The bound $\bar{\omega}$ can be defined by solving

$$\beta_{\bar{\omega}} \inf_{\theta \in [0,1/4]} \left\{ \frac{\int_{\theta}^{\frac{1}{4}} \tau dF(\tau) - (1-4\theta)\mathbb{E}[\theta]}{\theta(1-F(\theta))} \right\} = \alpha.$$

To obtain an explicit upper bound on $\bar{\omega}$, we observe that a sufficient condition is that $W(\theta) \equiv -\beta_\omega \theta f(\theta) + \alpha J^B(\theta) f(\theta)$ is decreasing. The derivative of this expression is $(2\alpha - \beta_\omega) f(\theta) + (\alpha - \beta_\omega)\theta f'(\theta) \leq (2\alpha - \beta_\omega) f(\theta)$, which is negative if $\beta_\omega \geq 2\alpha$ (where we used the fact that $f' \geq 0$). This means that $\bar{\omega} \leq (4/3)\alpha + (2/3)$.

Summarizing, if the lower bound of the support of ω is above $(4/3)\alpha + (2/3)$, whatever the outside option is, the principal will either offer a cash payment to buy out the rights (when this is allowed), or offer a constant probability y of allocating the monopoly right for free. In both cases, the principal will make sure that the highest type is getting exactly her outside option. This implies that the designer's problem reduces to choosing an outside option for the highest type that is just high enough to induce investment. In case monetary payments are allowed and investment is observable, the designer can achieve that via a cash payment; in case monetary payments are not allowed and the investment is not observable, the designer can achieve that by choosing a probability y of granting the monopoly right.

B.4 Vaccine development

The agent is a pharmaceutical company developing a vaccine at $T = 1$, during a pandemic. There is a unit mass of patients, and x represents the number of units purchased by the government at $T = 2$. Suppose that k is the marginal cost of production conditional on successful discovery of the vaccine. Let ω be the social value of vaccinating a single patient (which we assume is independent of k) that may depend, for example, on the severity of the pandemic. We set $\theta \equiv -k$. We assume the regulator cares exclusively about patient welfare, and let $V^\star_\omega(\theta) = V_\omega(\theta) = \omega$. Additionally, we let $1 = \alpha \geq \alpha^\star$.

In this application, our framework casts light on the optimal design of a contract between the government and a private producer. The friction is that—in the absence of a contract—the government may not be interested in purchasing the product after the investment costs have been sunk by the firm. However, the government can reward the investment with a cash transfer or a guaranteed sale price for all or some of the developed products. Note that it is natural to assume that these quantities should not depend on the state ω—while the severity of the pandemic may be publicly observed, it would be difficult to enforce such dependence in a legal contract. In this case, the optimal contract can essentially be thought of as an advanced market commitment.

We assume that investment is observable. By the analysis in Section 4.4 (and under the same regularity assumptions), as long as the cost of investment is sufficiently high, the optimal contract can be implemented as a lump-sum payment (for developing the vaccine) plus a guaranteed unit purchase price $p = \frac{\mathbb{E}_{\omega \sim G}[\omega | \omega \in [\underline{\omega}_p, \overline{\omega}_p]]}{\alpha^\star}$, for some functions $\underline{\omega}_p$, $\overline{\omega}_p$, assuming that p belongs to the support of the costs (otherwise, it coincides with one of the bounds). Intuitively, when $\omega < \underline{\omega}_p$ (the pandemic is not severe), the principal prefers to compensate the producer in cash, rather than buying the vaccines at the price p. When $\omega > \overline{\omega}_p$ (the pandemic is severe), the principal will offer a higher price than p to the producer to increase the production of vaccines. Thus, only in the intermediate range of ω can the price p set by the contract affect the $T = 2$ allocation.

Consistent with our discussion of the contractible case in Section 4.4, the optimal price does not depend on the exact cost of investment and the distribution of marginal costs; these factors only influence the size of the lump-sum payment. When $\alpha^\star = 0$, that is, when the government is not concerned about revenue at the stage of signing the contact, p will be equal to the upper bound of the distribution of costs; it is optimal to commit to purchasing all vaccines. When $\alpha^\star = 1$, so that the government has time-consistent preferences, the optimal price is the same as the regulator would choose if she wanted to implement the VCG mechanism. This is surprising, because the government was not assumed to maximize total surplus. The reason is related to the discussion of the optimal price in the contractible case given in Section 4.4. In the optimal contract, on the margin, the government must be indifferent between incentivizing investment using a slightly higher lump-sum payment or a slightly higher guaranteed purchase price—it thus behaves *as if* it was fully internalizing the producer's marginal costs (i.e., as if it was maximizing total surplus).

Supporting calculations. In this application, we have negative types: $\theta \equiv -k$. Moreover, $V_\omega(\theta) = \omega$, $V^\star_\omega(\theta) = \omega$, $\alpha = 1$, and $\alpha^\star \leq 1$. By Proposition 1, we have the thresholds $\omega + \underline{\theta}^\star_\omega + \frac{F(\underline{\theta}^\star_\omega)}{f(\underline{\theta}^\star_\omega)} = 0$ and $\omega + \overline{\theta}^\star_\omega - \frac{1-F(\overline{\theta}^\star_\omega)}{f(\overline{\theta}^\star_\omega)} = 0$, assuming they fall within $[\underline{\theta}, \bar{\theta}]$ (otherwise, they are equal to one of the bounds). Following the derivation in Appendix A.8, we have $\Phi'_\omega(\theta) + \alpha^\star \Psi'_\omega(\theta) = -[\omega + \alpha^\star\theta]\, f(\theta)\mathbf{1}\left\{\theta \in (\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega)\right\}$. Rewriting the first-order condition from Appendix A.8 yields that a necessary condition for optimality is $\theta^\star = \frac{\mathbb{E}_{\omega\sim G}[\omega | \omega \in [\underline{\omega}_{\theta^\star}, \bar{\omega}_{\theta^\star}]]}{\alpha^\star}$, with $\theta^\star = \bar{\theta}$ if the right-hand side expression is above $\bar{\theta}$, and $\theta^\star = \underline{\theta}$ if the right-hand side expression is below $\underline{\theta}$, where the bounds in the condition $\omega \in [\underline{\omega}_{\theta^\star}, \bar{\omega}_{\theta^\star}]$ are defined implicitly by $\theta^\star \in [\underline{\theta}^\star_\omega, \overline{\theta}^\star_\omega]$.

B.5 Supply chain contracting

There is a large firm (playing the role of the designer and the principal) buying x units of customized inputs from a small supplier (the agent). The supplier can invest at time $T = 1$ in relationship-specific technology to produce the inputs at marginal cost $k \equiv -\theta$. The firm maximizes profits and has a constant marginal value of 1 for each unit of the input. That is, we have $V_\omega(\theta) = 1$ and $\alpha = 1$. Through the close interaction with the supplier, we assume that the firm learns something about the supplier's costs: the state ω is a (non-contractible) signal of θ. Finally, we set $V^\star_\omega(\theta) = 1$ and $\alpha^\star = 1$, which corresponds to the firm proposing a contract to the supplier.

Theorem 1′ directly predicts the form of the optimal contract for the large firm. If investment by the small supplier is not observable, the large firm will in general choose a two-price scheme, committing to buy up to y units at some price p', or to buy any number of units at some lower price p. If investment by the small supplier is observable, assuming the cost of investment is high enough, the large firm will offer an upfront payment for setting up production and then a guaranteed purchase price for any number of units.