

A mechanism-design approach to property rights*

Piotr Dworczak[†] Ellen V. Muir[‡]

February 23, 2024

Abstract

We propose a framework for studying the optimal design of rights relating to the control of an economic resource, which we broadly refer to as *property rights*. An agent makes an investment decision affecting her valuation for the resource, and then participates in a trading mechanism chosen by a principal in a sequentially rational way, leading to a hold-up problem. A designer—who would like to incentivize efficient investment and whose preferences may differ from those of the principal—can endow the agent with a menu of rights that determine the agent’s outside options in the interaction with the principal. We characterize the optimal menu of rights as a function of the designer’s and the principal’s objectives, and the investment technology. The optimal menu requires at most two types of rights, including an *option-to-own* that grants the agent control over the resource upon paying a pre-specified price.

Keywords: Property rights, mechanism design, hold-up problem, option-to-own

JEL-Classification: D23, D86, D82

*We thank Dirk Bergemann, Alessandro Bonatti, Laura Doval, Marina Halac, Simon Loertscher, Elliot Lipnowski, Paul Milgrom, Stephen Morris, Georg Nöldeke, Bruno Jullien, Alessandro Pavan, Patrick Rey, Anna Sanktjohanser, Ilya Segal, Bruno Strulovici, and seminar audiences at Harvard/MIT, Yale, Columbia, Boston College, Bonn, Cergy Paris Université, MIT Sloan, UChicago, Stanford, University of Iowa, 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.

[‡]Joint Center for History and Economics, Harvard University. Email: emuir@fas.harvard.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 related to its use. However, when transaction costs (such as bargaining frictions) 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: If radio spectrum licenses grant holders full property rights over the underlying spectrum, then this incentivizes investment in the costly infrastructure that enables individual holders to efficiently utilize their spectrum. However, due to the complexity and high transaction costs associated with the market-wide reassignment of spectrum (see, for example, the FCC’s 2016-17 “Incentive Auction”), awarding spectrum licenses that endow holders with such strong property 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 explicitly specify the sides of the contract; rather, they give the holder the authority to unilaterally implement certain outcomes pertaining to the underlying economic resource. For example, a full property right always gives the owner an *option* to use or generate income from the resource regardless of what other options (such as selling the resource) might become available by interacting with other agents.² Consequently, we 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.

¹The current design of spectrum licenses effectively grants holders full property rights over the underlying spectrum 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.”

²In the legal nomenclature going back to Hohfeld (1917), property rights can be thought of as *in rem* rights that are “good against the world,” as opposed to *in personam* rights that typically arise out of bilateral (or multilateral) negotiations (see also the discussion in Balbuzanov and Kotowski, 2019).

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 outside options relating to the control of an economic resource 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 public state that pins down the principal’s opportunity cost for allocating the resource to the agent 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 the agent’s property rights).

The key assumption in our framework is that the designer does not directly control the trading mechanism. Conditional on the realization of the state, the mechanism is chosen in a sequentially rational way by the principal who maximizes an objective function that need not represent the designer’s ex-ante preferences. We refer to this friction as *ex-post inefficiency*. (In some applications, the principal could represent the “future self” of the designer, in which case ex-post inefficiency can be understood as a form of time inconsistency for the designer.) Lack of commitment also results in a *hold-up problem*: The agent may fail to undertake efficient investments if the subsequent trading mechanism extracts the resulting surplus.

The design of property rights can alleviate these frictions. For example, a conventional property right cedes full control over the resource to the agent, 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 designer in exchange for relinquishing control over the resource, or an option to acquire control over the resource conditional on paying a pre-specified price. By strengthening the agent’s rights, the designer affects the agent’s investment incentives as well as the principal’s flexibility at the stage of choosing a trading mechanism.

Before giving an overview of our results, we emphasize two (implicit) simplifying assumptions made within our framework. First, by studying a setting with a single agent, we abstract away from the problem of how to *select* the agent who should have property rights 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 away from the potential impact of prop-

erty 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. The assumption narrows down the set of applications of our framework but holds in environments in which the principal represents a government or a market regulator.

Our main finding is that the optimal property right is relatively simple but more flexible than a full property right. Specifically, in our framework, regardless of the designer’s preferences, optimality can be achieved by endowing the agent with a menu of at most two types of rights. One of the rights takes the form of an *option-to-own*. An option-to-own gives the agent the right to retain control over the resource conditional on paying a pre-specified price. The designer can vary the strength of the option-to-own by adjusting the price. For example, setting the price to zero is equivalent to a full property right, while setting a sufficiently high price is equivalent to giving no right to the agent. The second type of right in the optimal menu is only required if the agent’s cost of investing is sufficiently high and its form depends on whether the designer can make the agent’s rights contingent on investment. If investment is observable (and contractible), then the second right takes the form of a cash payment for undertaking the investment. If investment is not observable (and hence non-contractible), then 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).

From a methodological perspective, property rights in our framework give rise to a flexible set of outside options available to the agent in the interaction with the principal. Thus, the principal solves an instance of a mechanism-design problem with type-dependent outside options, as in the work of [Lewis and Sappington \(1989\)](#) and [Jullien \(2000\)](#). We derive a novel solution technique for such problems based on an extension of the classical ironing technique due to [Myerson \(1981\)](#). The designer’s problem is then 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 our 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.

We illustrate the usefulness and flexibility of our framework by considering five examples. First, motivated by applications such as the allocation of electromagnetic spectrum and mining rights, we study a dynamic resource allocation problem in which a regulator cannot commit to future trading mechanisms (e.g., spectrum auctions) but can design the *resource*

use license. When designing the 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 typically takes the form of a renewable lease that gives the license holder the opportunity to retain control conditional on paying a pre-specified price. Second, we consider the problem of how to optimally regulate a private rental market. In this application, we interpret the designer and the principal as separate economic agents: The designer is a market regulator concerned with efficiency, while the principal is a private rental company maximizing profits. We provide conditions under which optimal market regulation provides tenants with a right to renew their lease at a price tied to the market rental rate—a form of regulation that is frequently seen in practice. Third, inspired by a classic problem in economics, we discuss how a regulator might reward and incentivize innovation by committing to an appropriate patent policy. Specifically, we use a stylized model to illustrate why it may be optimal for the regulator to commit its patent office to a certain “review standard” that is independent of product profitability. We also discuss cases in which direct cash prizes or charging fees for granting patents may emerge as optimal tools. Our fourth example casts light on the optimal design of a contract between the government and a private producer. Inspired by applications such as vaccine development, we provide an optimality foundation for the practice of offering advanced market commitments. Finally, we investigate the classical ratchet effect by studying the optimal form of contractual rights between a large firm and a small supplier. In this context, the optimal menu of rights involves the large firm committing to a two-price purchase scheme.

The remainder of this paper is organized as follows. We provide an overview of the related literature in Section 1.1. Section 2 introduces and discusses the model. In Section 3 we state and prove our main result (Theorem 1), which characterizes the optimal menu of rights. The proofs of all auxiliary results can be found in Appendix A. Section 4 introduces and analyzes each of our five examples. We conclude with a discussion of future research directions in Section 5.

1.1 Related literature

Building on the seminal contribution of Coase (1960) and Williamson (1979), the economic literature concerning property rights has largely focused on two forms of transaction costs: private information and hold-up problems. Options-to-own have been proposed as potential solutions to both frictions. However, to the best of our knowledge, we are the first to demonstrate that options-to-own are part of an *optimal* solution when property rights can

be chosen from a large non-parametric class. The flexible approach to modeling property rights resonates with the legal literature, which considers other forms of property rights beyond the simple, unconditional property rights most commonly studied in the economics literature.³

In the context of private information as a type of transaction cost, [Myerson and Satterthwaite \(1983\)](#) first pointed out that there may be no bargaining procedure that results in efficient outcomes when contracting parties possess private information but property rights are assigned exclusively to one of the parties. [Cramton, Gibbons, and Klemperer \(1987\)](#) further clarified the importance of the initial allocation of property rights by showing that efficiency may be attainable if the involved parties have sufficiently balanced ownership shares. Most closely related to our work are papers analyzing the second-best design of property rights in this context. In particular, [Che \(2006\)](#) showed that using an option-to-own allows the designer to decrease the subsidy needed to implement the first-best outcome. [Segal and Whinston \(2016\)](#) unified much of this literature by studying the subsidy-minimizing choice of property rights from a relatively large parametric class; they also characterized the option-to-own that maximizes surplus subject to maintaining budget balance in the mechanism. Even without the hold-up problem, our framework and results would be different: The designer-preferred outcome in our setting (which need not be allocative efficiency) is prevented not by multi-sided private information but by the fact that the designer does not directly control the trading mechanism. This simplifies the analysis, and in particular allows us to characterize the optimal property right. Without the investment stage, the optimal right always takes the form of an option-to-own in our setting.

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 also been investigated. [Hart \(1995\)](#) showed that a price contract can improve upon a simple ownership structure, and [Nöldeke and Schmidt \(1995, 1998\)](#) identified settings in which options-to-own can restore first-best levels of investment. By studying a setting with a single agent, we shift focus away from the problem of optimal *reallocation* of residual rights of control among multiple parties, and towards the problem of the optimal *design* of these rights. This perspective allows us to characterize optimal rights even though the first best is typically not implementable in our setting, which features private

³See [Calabresi and Melamed \(1972\)](#) and the related discussion in [Segal and Whinston \(2016\)](#).

information at the trading stage.⁴ Despite these difference, we find that options-to-own play an important role even if the designer can choose from a non-parametric set of property rights. The optimal property right for addressing the hold-up problem may sometimes be more complicated: Depending on the observability (and contractibility) of investment, it may be necessary to complement an option-to-own with either a monetary transfer or a partial property right that grants control over a fraction of the resource (or the entire resource with some probability). In some of our applications, the property right chosen by the designer imposes restrictions on the private parties’ contracting space—this perspective was explored by [Hermalin and Katz \(1993\)](#) who asked whether courts could improve private contracting in this way; they find a mostly negative answer due to private contracting being efficient in their framework in most cases.⁵

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 valuations.⁶ 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. That being said, we recover a version of [Rogerson’s](#) insight by showing that if both the principal and the designer are interested in maximizing efficiency (and investment only affects the private value of the resource), then it is optimal to allocate no rights to the agent. Moreover, in the special case of our model with no uncertainty about the public state, the designer may sometimes 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.

A VCG mechanism may fail to induce efficient investments in the common value of the resource. [Weyl and Zhang \(2022\)](#) consider the trade-off between common-value investment incentives and allocative efficiency; they propose a new form of a partial property right—a depreciating license—that outperforms both a full property right and a short-term lease

⁴[Matouschek \(2004\)](#) and [Baliga and Sjöström \(2018\)](#) allow for private information at the contracting stage but do not consider the investment problem. As in our model, the property rights in [Baliga and Sjöström \(2018\)](#) lead to type-dependent outside options; however, [Baliga and Sjöström \(2018\)](#) focus on parameters for which the first best is implementable.

⁵Recently, [Hitzig and Niswonger \(2023\)](#) study a similar question in a different setting, with application to regulation of platform labor contracts.

⁶Several extensions of this result have been examined in the literature: [Bergemann and Välimäki \(2002\)](#) analyzed efficient information acquisition, [Hatfield et al. \(2019\)](#) clarified the link between ex-ante efficient investment, ex-post efficiency and strategy-proofness, and [Akbarpour et al. \(2023\)](#) studied investment incentives in a setting where the mechanism must be computationally tractable.

contract. Our model of investment is simpler but it can capture a common-value component in a reduced form way. Despite differences in modeling assumptions and our focus on *optimal* rights, we similarly find that a conditional property right—in particular one that involves a notion of a price—may be preferred to classical property rights. In Section 4, we review applications of our framework and comment on how the policy prescriptions we derive agree or differ from those formulated in more applied literatures on license design and optimal patent protection.

2 Model

Overview. We consider a model involving three time periods and three players: a designer, a principal, and an agent. At time $t = 0$, the designer chooses a *menu of rights* that determines the agent’s outside options (pertaining to the control of an economic resource being traded at $t = 2$). At time $t = 1$, the agent decides whether to undertake a costly investment. This investment decision 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 1.




Figure 1: Model overview and timeline.

Menu of rights. At time $t = 0$, the designer chooses a menu of rights M held by the agent

in subsequent periods. 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 in period $t = 2$, and the set I is arbitrary. 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 at $t = 2$, in the sense that any $(x_i, t_i) \in M$ constitutes an outside option available to the agent when contracting with the principal.

Investment. At time $t = 1$, the agent takes 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 $\theta \in \Theta := [\underline{\theta}, \bar{\theta}] \subset \mathbb{R}$ and the public state $\omega \in \Omega \subset \mathbb{R}$. If the agent invests, the public state is drawn from a distribution G , and the agent's type is drawn from a conditional distribution F_ω .⁷ If the agent does not invest, the respective distributions are denoted \underline{G} and \underline{F}_ω . We assume that, for every ω , F_ω and \underline{F}_ω admit absolutely continuous densities on Θ (denoted f_ω and \underline{f}_ω , respectively). For every ω , F_ω first-order stochastically dominates \underline{F}_ω , so that the primary role of investment is that it increases the agent's type.

Trading Mechanisms. At time $t = 2$, the agent's private type θ and the public state ω are realized (the state ω is observed by both the agent and the principal). The principal then 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, for every realized ω , the principal chooses a mechanism $\langle x_\omega(\theta), t_\omega(\theta) \rangle$, where $x_\omega : \Theta \rightarrow [0, 1]$ denotes the allocation rule and $t_\omega : \Theta \rightarrow \mathbb{R}$ denotes the transfer rule.

We assume the agent's utility is linear in the allocation x (interpreted as either a probability or quantity) and the transfer t , with the type θ normalized to equal the agent's marginal value for the allocation. An agent with type θ who receives an allocation $x \in [0, 1]$ and makes a payment $t \in \mathbb{R}$ then obtains utility $\theta x - t$. Given a direct mechanism $\langle x_\omega(\theta), t_\omega(\theta) \rangle$, incentive-compatibility requires that, for all $\theta, \theta' \in \Theta$, and $\omega \in \Omega$,

$$U_\omega(\theta) := \theta x_\omega(\theta) - t_\omega(\theta) \geq \theta x_\omega(\theta') - t_\omega(\theta'). \quad (\text{IC})$$

The agent's outside option in the absence of any rights is normalized to 0. However, the menu M chosen by the designer gives rise to an endogenous type-dependent outside option

⁷For all variables in our model that depend on ω , we assume that they are measurable functions of ω .

determined by the agent's optimal choice of a right from M at time $t = 2$. The principal is able to replicate all outcomes in which the agent executes some outside option from the menu M within her mechanism. Consequently, it is without loss of generality to assume that the mechanism chosen by the principal ensures participation; hence, for every $\theta \in \Theta$ and $\omega \in \Omega$, we require

$$U_\omega(\theta) \geq \max\{0, \max_{i \in I} \{\theta x_i - t_i\}\}. \quad (\text{IR})$$

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

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

where $V_\omega : \Theta \rightarrow \mathbb{R}$ is upper semi-continuous in θ , and $\alpha > 0$ is the weight that the principal places on revenue. We denote by $\langle x_\omega^*(\theta; M), t_\omega^*(\theta; M) \rangle$ the optimal mechanism for the principal when the participation constraint (IR) is induced by menu M . The optimal mechanism is generically unique; in case of indifference by the principal our proofs utilize a particular tie-breaking rule that simplifies the exposition. Our results continue to hold under a large class of tie-breaking rules, including designer-preferred selection, as explained in Appendix A.6.

Agent's problem. We can now formally state the agent's problem; the agent will invest if and only if

$$\int_{\Omega} \int_{\Theta} (\theta x_\omega^*(\theta; M) - t_\omega^*(\theta; M)) dF_\omega(\theta) dG(\omega) - c \geq \underline{U}, \quad (\text{I-OB})$$

where $\underline{U} \geq 0$ captures the agent's expected payoff from not investing. We will consider two cases of our model depending on whether the investment decision of the agent is observable (and contractible). In the *non-contractible case*, we set

$$\underline{U} = \int_{\Omega} \int_{\Theta} (\theta x_\omega^*(\theta; M) - t_\omega^*(\theta; M)) dF_\omega(\theta; M) dG(\omega),$$

capturing the idea that the agent enjoys her rights M and faces the same mechanism whether

or not she invested. In the *contractible case*, we set

$$\underline{U} = \int_{\Omega} \int_{\Theta} (\theta \underline{x}_{\omega}^*(\theta; \emptyset) - \underline{t}_{\omega}^*(\theta; \emptyset)) d\underline{F}_{\omega}(\theta) d\underline{G}(\omega),$$

where $\langle \underline{x}_{\omega}^*(\theta; \emptyset), \underline{t}_{\omega}^*(\theta; \emptyset) \rangle$ is the principal's optimal mechanism assuming that $M = \emptyset$ and the agent's type θ is drawn from \underline{F}_{ω} given the realized ω . That is, if the agent does not invest, she does not enjoy the rights assigned by the designer; moreover, the principal knows that the agent's type is drawn from a lower distribution.

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

$$\begin{aligned} \max_M \int_{\Omega} \int_{\Theta} [V_{\omega}^*(\theta) x_{\omega}^*(\theta; M) + \alpha^* t_{\omega}^*(\theta; M)] dF_{\omega}(\theta) dG(\omega) & \quad (\text{D}) \\ \text{s.t. (I-OB),} & \end{aligned}$$

where $V_{\omega}^* : \Theta \rightarrow \mathbb{R}$ is continuous in θ , and $\alpha^* \geq 0$ 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 M that satisfies (I-OB), but (iii) the agent does not invest if $M = \emptyset$.

2.1 Discussion

While somewhat abstract, our setting has the advantage of capturing a wide range of applications. We discuss some of our modeling assumptions and their interpretations below.

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 whereby the agent can demand a monetary transfer p from the principal (who then has full control over the period $t = 2$ allocation);
- $M = \{(1, 0), (0, -p)\}$ 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 ;

- $M = \{(1, p)\}$ represents a renewable lease or an *option-to-own*, giving the agent the right to acquire control over the resource conditional on paying the principal a fixed price p ;
- $M = \{(y, 0)\}$ with $y \in (0, 1)$ captures a “partial property right.” The interpretation of partial property rights will vary depending on the application. If y represents a probability, then the right can be implemented by conditioning ownership on some exogenous future event such as a court decision; the designer can adjust y by varying how difficult it is to contest the right in front of a court.⁸ If y represents a quantity traded, then a partial right applies only to some fraction of the total available volume. Finally, in a reduced-form way, y can capture geographic or temporal restrictions on the property right.
- $M = \{(s, p(s))_{s \in [0,1]}\}$, for some function $p : (0, 1) \rightarrow \mathbb{R}_+$, is a flexible menu then allows the agent to purchase their preferred partial property right $s \in [0, 1]$ at a price $p(s)$.

The investment stage. Our model is agnostic about the interpretation of the allocation, and whether or not the agent controls the underlying economic resource when making the investment decision. One possibility is that the agent makes a relationship-specific investment prior to trading with the principal at $t = 2$ and the property rights assignment represents an underlying legal framework; another is that the agent is allocated a good at $t = 0$ along with a legal contract specifying what rights the agents has with regard to extending her control over the good to the second period. Our applications explore both possibilities.

We modeled investment as a binary decision to highlight the key forces in our framework in the simplest possible way. However, our results extend—in an appropriate sense that we explain later—to richer environments in which the agent decides how much to invest.

The trading stage. Our modeling of the trading stage is different 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 has several implications. First, there exists an ex-post efficient mechanism but it need not be selected by the principal. The key assumption is that the principal selects the mechanism at time $t = 2$ in a sequentially rational way to maximize her payoff, which may in general differ from the social optimum, as represented by the designer’s preferences. The principal may

⁸For example, there is variation across jurisdictions in the degree of protection of intellectual property rights that determines the ex-ante probability of retaining *de facto* ownership.

also represent a third party or a “future self” of the designer (exhibiting a form of time-inconsistency). Second, 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.

Model Frictions. Our framework features two fundamental frictions. The first one is a *hold-up problem* created by the assumption that the principal cannot commit to her mechanism in order to incentivize the agent’s investment. The second one, which we refer to as *ex-post inefficiency*, is the possible divergence between the preferences of the designer and the principal, resulting in socially suboptimal allocation in the second-period mechanism.⁹ Property rights are a tool used by the designer to address both of these frictions, by shifting rents to the agent and affecting the mechanism selected by the principal. We will occasionally “turn off” one of the frictions (by either removing the investment stage from the game, or by aligning the designer’s and the principal’s preferences) to obtain sharper predictions.

Aside from some special cases, the effectiveness of property rights in our framework is limited. This is in part due to the fact that we have built in an *incomplete-contracts* friction by assuming that property rights cannot be conditioned on the realization of the state ω , even though the state is publicly observed. This assumption seems realistic for most applications and captures the idea that property rights endow the holder with robust guarantees that are not contingent on circumstances that would be difficult to verify in front of a court. That being said, our methods and results extend to the case of state-contingent rights, as we explain in Section 5.

3 Analysis and results

We begin by stating the main result of the paper that characterizes the structure of the optimal menu of rights M^* chosen by the designer.

Theorem 1. *There exists an optimal menu that takes the form $M^* = \{(1, p), (y, p')\}$ for some $p, p' \in \mathbb{R}$ and $y \in [0, 1)$.*

⁹Here, we use the word “efficiency” broadly to refer to the designer’s objective, representing a socially desirable outcome. We are not assuming that the designer maximizes allocative efficiency in the narrow sense, even though we will often study this case in applications.

As Theorem 1 shows, the optimal menu M takes a simple and economically interpretable form. The menu consists of at most two types of rights, including an *option-to-own* $(1, p)$ which gives the agent the right to control the resource by paying a pre-specified price p . The second item in the menu takes the form of either a partial property right—giving the agent partial control over the resource at a lower price, possibly for free—or a cash payment to the agent (when $y = 0$ and $p' < 0$). We later show that the form of the second item in the menu depends crucially on whether investment is observable or not.

Note that Theorem 1 does not preclude the possibility that the optimal menu gives the agent no choice over which right to execute (or even no rights whatsoever)—this is because one (or both) of the options in the menu could have a sufficiently high price that the agent never wants to execute it. We will show that a number of configurations can emerge as optimal in applications—the optimal menu could be a singleton containing an option-to-own $(1, p)$, a cash transfer $(0, -p)$, or a partial right allocated for free $(y, 0)$.

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 parameters p, p', y characterizing the optimal menu are pinned down by the primitives of the model. We will further explore the economic implications of our characterization in Subsection 3.3, where we derive tighter predictions under additional regularity conditions, and in Section 4, where we study applications.

3.1 Proof of Theorem 1

We proceed backwards, by first solving the principal’s problem in period $t = 2$, then considering the agent’s investment problem in period $t = 1$, and finally solving the designer’s problem in period $t = 0$.

Step 1: Formulating the principal’s problem

We first focus on solving the principal’s problem, given an arbitrary menu of rights M and a realization $\omega \in \Omega$ of the public state. For ease of exposition, we drop any explicit dependence of the principal’s objective function and the choice of mechanism on these variables. We reformulate the principal’s problem by expressing the consequences of any menu of rights M that the agent may hold as a type-dependent outside option.

Lemma 1. *A choice of menu M by the designer is equivalent to choosing an outside option function $R : \Theta \rightarrow \mathbb{R}$ for the agent in the second-period 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 maximizing 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

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

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. Moreover, for all $\theta \in \Theta$, we have

$$t(\theta) = \theta x(\theta) - \int_{\underline{\theta}}^{\theta} x(\tau) d\tau - \underline{u}.$$

After standard transformations, this yields

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

where $B(\theta) := \theta - (1 - F(\theta))/f(\theta)$ is the virtual value function. Combining this with Lemma 1, the principal's problem (P) can be rewritten as

$$\begin{aligned} & \max_{x: \Theta \rightarrow [0,1], \underline{u} \geq 0} \int_{\underline{\theta}}^{\bar{\theta}} W(\theta)x(\theta)d\theta - \alpha \underline{u} & (P') \\ \text{s.t. } & x \text{ is non-decreasing, and } U(\theta) = \underline{u} + \int_{\underline{\theta}}^{\theta} x(\tau) d\tau \geq R(\theta), \quad \forall \theta \in \Theta, \end{aligned}$$

where $W(\theta) := (V(\theta) + \alpha B(\theta)) f(\theta)$. 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 a generic upper semi-continuous objective $W(\theta)$, and an outside option function R 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 $x_0 : \Theta \rightarrow [0, 1]$.¹⁰

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

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

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

$$\begin{aligned} \underline{\theta}^* &:= \sup\{\{\theta \in \Theta : \bar{\mathcal{W}}'(\theta) \geq \alpha\} \cup \{\underline{\theta}\}\}, \\ \bar{\theta}^* &:= \inf\{\{\theta \in \Theta : \bar{\mathcal{W}}'(\theta) \leq 0\} \cup \{\bar{\theta}\}\}. \end{aligned}$$

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

Let \mathcal{I} be the (at most countable) collection of maximal open intervals (a, b) within $(\underline{\theta}^*, \bar{\theta}^*)$ with the property that \mathcal{W} lies strictly below $\bar{\mathcal{W}}$ on (a, b) .¹¹ (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 $(\underline{\theta}^*, \bar{\theta}^*)$ with the property that \mathcal{W} coincides with $\bar{\mathcal{W}}$ on $[a, b]$. Intuitively, the allocation rule must be “ironed” on each $(a, b) \in \mathcal{I}$. Formally, we define

$$\underline{x}^* = R(\underline{\theta}^*) \text{ and } x^*(\theta) = \begin{cases} 0 & \theta \leq \underline{\theta}^*, \\ \frac{\int_a^b R'(\tau) d\tau}{b-a} & \theta \in (a, b) \text{ for some } (a, b) \in \mathcal{I}, \\ R'(\theta) & \theta \in [a, b] \text{ for some } [a, b] \in \mathcal{I}^c, \\ 1 & \theta \geq \bar{\theta}^*. \end{cases} \quad (2)$$

¹⁰This last assumption is called *homogeneity* by Jullien (2000) and is crucial for our method to work. Lemma 1 guarantees that we can write the function R this way for any choice of M .

¹¹By maximality we mean that such an interval (a, b) cannot be strictly contained in another interval (a', b') with the same property.




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).

The allocation rule x^* is equal to 0 below $\underline{\theta}^*$ and 1 above $\bar{\theta}^*$. By the choice of the payment \underline{u}^* , the outside option constraint binds at $\theta = \underline{\theta}^*$. Then, within the interval $[\underline{\theta}^*, \bar{\theta}^*]$, x^* coincides with $R'(\theta)$ on “non-ironing intervals” (the outside option constraint binds everywhere in such intervals), and is constant on “ironing intervals” (the outside option constraint binds 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^*, \underline{u}^*) 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 $\mathcal{W} = \bar{\mathcal{W}}$. Thus, $\mathcal{I} = \emptyset$, and ironing is not needed. Furthermore, $\underline{\theta}^*$ is defined by $W(\underline{\theta}^*) = -\alpha$, and $\bar{\theta}^*$ is defined by $W(\bar{\theta}^*) = 0$ (assuming such solutions exist). For $\theta \geq \bar{\theta}^*$, the principal’s objective is positive, so she chooses the maximal allocation 1, and the outside option constraint is slack. For $\theta \leq \bar{\theta}^*$, 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 the allocation. Thus, the principal optimally sets $x^*(\theta) = 0$ for types below $\underline{\theta}^*$, and she uses the lump-sum payment $\underline{u}^* = R(\underline{\theta}^*)$ to satisfy the outside option constraint for all these types. For the remaining types $\theta \in [\underline{\theta}^*, \bar{\theta}^*]$, the principal uses the outside option allocation $x_0 \equiv R'$ to satisfy the constraint; she sets $x^*(\theta) = R'(\theta)$ 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}^*) below $\underline{\theta}^*$, coincides with $R(\theta)$ on $[\underline{\theta}^*, \bar{\theta}^*]$, and has slope 1 above $\bar{\theta}^*$.

The case of a non-monotone W is analogous, except that we must first “iron” $W(\theta)$ into its monotone version $-\overline{W}'(\theta)$. 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 $[\underline{\theta}^*, \bar{\theta}^*]$ (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 $[\underline{\theta}^*, \bar{\theta}^*]$, 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.¹² 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 .

Remark 1. The objective function W in problem (P') incorporates the density of types f . This implies that the properties of the solution—in particular 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 (unlike, for example, a supply constraint).

¹²Our 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 et al. \(2023\)](#).

While the solution to problem (P') is of independent interest, the key observation that we will need to prove Theorem 1 is as follows.

Corollary 1. *The optimal solution (x^*, \underline{u}^*) to problem (P') defined in (2) depends linearly on the outside option R .¹³*

Corollary 1 is a consequence of the ironing procedure: The collection of ironing intervals \mathcal{I} , and the cutoff types $\underline{\theta}^*$ and $\bar{\theta}^*$ depend only on the principal's objective function and the distribution of types; they do *not* depend on R . Intuitively, the principal can determine the set of types at which the outside option constraint binds before she knows what the outside option of each type is. Of course, the optimal mechanism (x^*, \underline{u}^*) ultimately depends on R but only through a linear transformation applied within each of the intervals identified by the ironing procedure.

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*

$$\max_{\substack{x \text{ non-decreasing,} \\ u \geq 0}} \int_{\underline{\theta}}^{\bar{\theta}} \Phi(\theta) dx(\theta) - \alpha^* u \quad \text{subject to} \quad \int_{\underline{\theta}}^{\bar{\theta}} \Psi(\theta) dx(\theta) + \mathbf{1}_{cont} \cdot u \geq \tilde{c}, \quad (3)$$

for some constant $\tilde{c} \geq 0$ and functions $\Phi, \Psi : \Theta \rightarrow \mathbb{R}$, where $\mathbf{1}_{cont}$ is an indicator function that is 1 in the contractible case, and 0 in the non-contractible case.

By Corollary 1, the allocation rule selected by the principal is linear in the outside option R . Because both the designer's and the agent's payoffs are linear in the final allocation, it follows that the designer's problem is also linear in R (with a linear constraint corresponding to the agent's investment-obedience constraint). Given this observation and the change of variables described previously, Lemma 3 is a matter of bookkeeping: The functions Φ and

¹³Formally, if $(x_i^*, \underline{u}_i^*)$ is the solution to problem (P') under outside option R_i , for $i \in \{1, 2\}$, then $(x^*, \underline{u}^*) = \lambda(x_1^*, \underline{u}_1^*) + (1-\lambda)(x_2^*, \underline{u}_2^*)$ is a solution to problem (P') under outside option $R = \lambda R_1 + (1-\lambda)R_2$, for any $\lambda \in (0, 1)$.

Ψ are derived by taking expectations over ω and integrating by parts so that the allocation x enters the designer’s objective as a measure against which Φ and Ψ are integrated.

Problem (3) consists of maximizing a linear functional subject to a single linear constraint over a non-negative number and a non-decreasing function. 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 functions.¹⁴

Lemma 4. *Problem (3) admits a solution (x^*, u^*) such that either (i) $u^* = 0$ and x^* takes on at most one value other than 0 or 1, or (ii) $u^* > 0$ and $x^*(\theta) \in \{0, 1\}$ for all $\theta \in \Theta$.*

Lemma 4 implies Theorem 1: In each case described in the lemma, the optimal outside option function R is spanned by a menu containing at most two elements. Whenever $u^* > 0$, one of the elements is a cash payment, $(0, -u^*)$. In remaining cases, the non-zero values taken by x^* and the cutoff types at which x^* jumps determine the options of the form (y, p') . In particular, one of the two elements in the optimal menu can be taken to be an option-to-own $(1, p)$ (possibly with a price p that makes it redundant).

3.2 Discussion

Theorem 1 predicts that the optimal menu for the designer takes a relatively simple form: It suffices to offer the agent two types of rights in the optimal menu, and one of these rights is an option-to-own. Economically, an option-to-own is attractive because it allows the designer to incentivize investment in a flexible way. If the price is set to be low, an option-to-own behaves almost like a full property right and provides high incentives to invest; it also forces the principal to either allocate the good to the agent with high probability or compensate her with monetary transfers. Thus, a low price will be used when the cost of investment is high, or when the designer has a stronger preference than the principal to allocate the good to the agent. If, instead, the price is set to be high, an option-to-own does not alter the allocation in the principal’s mechanism too much, and provides only a small “nudge” to invest. Thus, a high price might be used when investment is relatively easy to induce.

Mathematically, an option-to-own is special because it is an extreme point of the set of feasible outside option functions R that the designer can induce by assigning rights to the agent. As the proof of Theorem 1 demonstrates, the designer’s problem is linear in the outside option R . Because the problem features a single (linear) constraint, there exists an

¹⁴This result follows from an infinite-dimensional extension of Carathéodory’s theorem found in Kang (2023) and has many analogs in recent papers in mechanism design (see, for example, Fuchs and Skrzypacz, 2015; Bergemann et al., 2018; Loertscher and Muir, 2023) and information design (see, for example, Le Treust and Tomala, 2019; Doval and Skreta, 2022).

optimal solution that is a convex combination of at most two extreme points. A simple corollary of the proof of Theorem 1 is that a single option-to-own would be optimal absent the investment-obedience constraint:

Corollary 2. *If the investment-obedience constraint is slack at the optimal solution, then there exists an optimal solution to the designer’s problem that takes the form of an option-to-own: $M^* = \{(1, p)\}$ for some $p \in \mathbb{R}$.*

In a version of our model without an investment-obedience constraint, the designer sets the price p in the option-to-own in a way that maximally aligns the principal’s mechanism with the designer’s preferences. Even if the investment-obedience constraint is present, it may be slack at the optimal solution if the misalignment of the designer’s and principal’s preferences is sufficiently large. For example, if the principal maximizes revenue while the designer puts sufficient weight on the agent’s welfare, she may choose to offer an option-to-own with a low price to shift more rents to the agent. As a by-product, the agent may have a strict incentive to invest.

In light of Corollary 2, it is the presence of a binding investment-obedience constraint that can lead to the necessity of including a second option in the optimal menu. Unsurprisingly then, the form of the second option in the optimal menu depends on whether investment is observable.

Corollary 3. *Suppose that the investment cost c is sufficiently high. In the non-contractible case, there exists an optimal menu $M^* = \{(1, p), (y, p')\}$ with $y > 0$ and $p'/y \in [\underline{\theta}, \bar{\theta}]$. In the contractible case, there exists an optimal menu $M^* = \{(1, p), (0, -T)\}$ with $T > 0$.*

Corollary 3 reveals a key difference between the cases when investment is observable (and contractible) and when it is not. In the contractible case, the optimal menu consists of an option-to-own and an option-to-sell—the agent either keeps the good by paying a price p or relinquishes control in exchange for a monetary payment T . For intuition, it is useful to observe that offering the menu $\{(1, p), (0, -T)\}$ (from which the agent selects a single option) is equivalent to paying the agent a lump-sum payment T and offering an option-to-own with price $p + T$, conditional on investment. (In the remainder of the paper, we will use the term “lump-sum payment” to refer to this alternative interpretation under which the cash payment is always given to the agent, regardless of other options selected from the menu.) Thus, in the contractible case, the designer can incentivize investment with cash. In contrast, when investment is not observable (or not contractible), offering a lump-sum payment to the agent is ineffective for incentivizing investment because the agent collects the lump-sum payment regardless of the investment decision. Instead, the designer incentivizes

investment by leveraging the fact that investment increases the agent’s value—the optimal menu increases the rents of higher types relative to lower types by only including options in which the agent obtains the good with strictly positive probability.

The simplicity of the optimal menu relies on our simplifying assumption that the agent takes a binary investment decision. However, the proof of Theorem 1 easily extends to the case when more (linear) constraints are added. If there are K constraints—for example because the agent has K alternative levels of investment to which she can deviate—at most $K + 1$ options are needed in the optimal menu offered by the designer (and an option-to-own is one of them). However, such a bound will typically not be tight. What matters is the number of *binding* constraints. For example, if investment is modeled as a continuous choice and the socially-efficient level of investment is pinned down by a first-order condition, then a single linear equation may be sufficient to capture the agent’s obedience constraint, and Theorem 1 applies *verbatim*.¹⁵

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 (2002).

3.3 The monotone case

In this subsection, we analyze the structure of the optimal property right under additional regularity conditions. Suppose that, for any $\omega \in \Omega$, the principal’s objective function $V_\omega(\theta)$ is non-decreasing in θ . Furthermore, suppose that the virtual surplus functions $B_\omega(\theta) := -(1 - F_\omega(\theta))/f_\omega(\theta)$ and $S_\omega(\theta) := \theta + F_\omega(\theta)/f_\omega(\theta)$ (which are usually associated with buyers and sellers, respectively, in mechanism design problems) are strictly increasing, and that the density $f_\omega(\theta)$ is continuously differentiable.

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*

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

¹⁵For a concrete example, suppose that the agent makes a continuous choice of effort $e \in [0, 1]$ subject to a strictly convex cost $c(e)$; the agent’s type θ is drawn from F_ω with probability e , and from \underline{F}_ω with probability $1 - e$. Then, the first-order condition for some target level of investment e^* is sufficient to guarantee that e^* is chosen by the agent, and the optimal menu has at most two items.

where $\underline{\theta}_\omega^* \leq \bar{\theta}_\omega^*$ are defined by

$$V_\omega(\underline{\theta}_\omega^*) + \alpha S_\omega(\underline{\theta}_\omega^*) = 0 \quad \text{and} \quad V_\omega(\bar{\theta}_\omega^*) + \alpha B_\omega(\bar{\theta}_\omega^*) = 0,$$

whenever an interior solution $[\underline{\theta}_\omega^*, \bar{\theta}_\omega^*] \subset \Theta$ exists.

In the monotone case, the principal's problem admits a simple and intuitive solution: The outside option constraint binds at an “intermediate” interval of types $[\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]$; the principal buys out rights using a cash payment for types $\theta \leq \underline{\theta}_\omega^*$, and she allocates with probability one to types $\theta \geq \bar{\theta}_\omega^*$. This is a direct consequence of the “ironing procedure” that we developed in Section 3.1.¹⁶ Intuitively, the principal wants to maximize the allocation for types higher than $\bar{\theta}_\omega^*$ and minimize the allocation for types lower than $\bar{\theta}_\omega^*$. Thus, the outside option constraint is slack for $\theta \geq \bar{\theta}_\omega^*$. On the remainder of the type space, the principal uses the allocation rule to satisfy the outside option constraint for types above $\underline{\theta}_\omega^*$, and the monetary payment to satisfy the outside option constraint for types below $\underline{\theta}_\omega^*$. This intuition is embedded in the definitions of $\underline{\theta}_\omega^*$ and $\bar{\theta}_\omega^*$ from Proposition 1: The upper threshold $\bar{\theta}_\omega^*$ 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 $\underline{\theta}_\omega^*$ is the cutoff type below which the principal would prefer to buy the resource from the agent.

Studying the monotone case allows us to provide more intuitions regarding the price used in the option-to-own included in the optimal menu (see Appendix A.8 for supporting calculations). For the rest of this section, we assume that the distribution of the public state ω does not depend on the agent's investment decision ($G = \underline{G}$).

The contractible case. In case investment is observable, as long as the hold-up problem is sufficiently severe, we know from Corollary 3 that the optimal menu can be implemented by awarding the agent a lump-sum payment T conditional on investment and letting her execute an option-to-own with some price p . Under standard regularity conditions guaranteeing the validity of first-order conditions, the price p must satisfy

$$\mathbb{E}_{\omega \sim G} \left[(V_\omega^*(p) + \alpha^* p) f_\omega(p) \mid p \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*] \right] = 0, \quad (4)$$

as long as p is interior; $p = \underline{\theta}$ is optimal if the left-hand side of equation (4) is always positive, and $p = \bar{\theta}$ is optimal if the left-hand side of equation (4) is always negative. The expectation in expression (4) is taken conditional on the event that p lies between $\underline{\theta}_\omega^*$ and $\bar{\theta}_\omega^*$ given the realization of ω , which guarantees that the option-to-own affects the final allocation of the

¹⁶While 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.

resource: If p is below $\underline{\theta}_\omega^*$, the principal buys out the agent's option-to-own, and if p is above $\bar{\theta}_\omega^*$, the principal offers a better price to the agent within the mechanism.

Intuitively, recall that in the contractible case, the designer incentivizes investment with a combination of a lump-sum payment and an option-to-own (that is offered conditional on investing). If the designer uses both tools, optimality requires that she cannot benefit by slightly lowering the price p in the option-to-own—thus relaxing the investment-obedience constraint—and then slightly lowering the lump-sum payment to make it bind again. Lowering the price p in the option-to-own increases the allocation in the mechanism for nearby types—the designer values this change at $V_\omega^*(p)$ —but it also affects the revenue. Normally, the effect on revenue would be captured by the virtual surplus term, $B_\omega(p)$. However, the binding investment-obedience constraint pins down the agent's expected information rents conditional on investment, and hence the incremental net revenue—after adjusting the lump-sum payment—excludes the information rent term: The net revenue from selling to type p is simply p . Because the designer values revenue at α^* , the net effect is given by $V_\omega^*(p) + \alpha^*p$. The optimal price p in the option-to-own makes the net effect zero in expectation. An interesting corollary is that the optimal price in formula (4) does not depend on the parameters of the agent's investment problem (such as the cost c). In a sense, the designer uses the option-to-own to achieve the desired physical allocation in the second stage, and then adjusts the lump-sum payment to make sure that the agent undertakes investment.

The non-contractible case. In the non-contractible case, by Corollary 3, the optimal menu may include two options, $M^* = \{(1, p), (y, p')\}$, with two different prices p and p' . Unlike in the case of observable investment, it is difficult to separate the effects of the two options on the investment incentives. However, we can provide some intuition if we explicitly introduce a Lagrange multiplier $\gamma \geq 0$ on the investment-obedience constraint in problem (3). While the multiplier γ is endogenous, it must be non-decreasing in the cost of investment c , and is hence intuitively related to the severity of the hold-up problem.

Suppose first that it is optimal to offer a singleton menu with an option-to-own $(1, p)$. Then, p must satisfy the following first-order condition (assuming an interior solution):

$$\begin{aligned} \alpha^* \frac{\mathbb{P}_{\omega \sim G}(p < \underline{\theta}_\omega^*)}{\mathbb{P}_{\omega \sim G}(p \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*])} - \mathbb{E}_{\omega \sim G} \left[(V_\omega^*(p) + \alpha^* B_\omega(p)) f_\omega(p) \mid p \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*] \right] \\ - \gamma \mathbb{E}_{\omega \sim G} \left[\underline{F}_\omega(p) - F_\omega(p) \mid p \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*] \right] = 0. \end{aligned} \quad (5)$$

To understand the expression, consider first the case when the designer does not care about revenue, $\alpha^* = 0$, and there is no hold-up problem, $\gamma = 0$. In this case, the designer's

first-best allocation in the second stage, assuming monotonicity of $V_\omega^*(\theta)$, is to allocate to all agent's types above the threshold θ_ω^* such that $V_\omega(\theta_\omega^*) = 0$. Thus, the designer sets the price p to achieve her optimal allocation *on average* across ω , conditional on ω falling within the range in which the option-to-own has bite. Suppose now that $\alpha^* > 0$. In this case, the designer would like to allocate to types for which $V_\omega^*(\theta) + \alpha^*B_\omega(\theta)$ is positive. However, there is a second, more subtle, effect captured by the first term in equation (5). Whenever ω is such that $p < \theta_\omega^*$, the principal will buy out the agent's right with a monetary payment that is decreasing in p (the more attractive the option-to-own, the higher the compensation the principal must offer to the agent). Thus, in this region, a lower price p has no effect on the allocation but it decreases the principal's revenue, which the designer values at α^* . This effect will push the designer to choose a higher price p in the optimal menu, in particular implying that a full property right will be suboptimal if θ_ω^* is bounded away from $\underline{\theta}$ (across ω). Finally, suppose that $\gamma > 0$, implying that the investment-obedience constraint binds. By assumption, the distribution of agent's values increases in the first-order stochastic dominance order after investing: $\underline{F}_\omega(p) \geq F_\omega(p)$. Thus, the last term in expression (5) 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. Once again, however, providing incentives to invest through an option-to-own is only effective when the price p falls in the region $[\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]$ where it affects the final allocation of the resource.

In Appendix A.8, we show that equation (5) captures the relevant trade-offs also when it is optimal to offer two options in the menu. Mathematically, two options are offered when the Lagrangian is non-monotone and has more than one global maximum, implying multiple solutions to the first-order condition (5). In that case, the solutions to equation (5) pin down the optimal cutoff types at which the agent switches between executing different outside options—prices p and p' can then be calculated from these cutoff types. Economically, non-monotonicity of the Lagrangian can arise due to the conflict between incentivizing investment (which pushes prices p and p' to be lower) and giving the principal more flexibility at the trading stage (which pushes prices p and p' to be higher).

4 Applications

In this section, we discuss several application of our framework, and illustrate the results with simple numeral examples.¹⁷ Our goal is to provide an overview of how our framework could be mapped into various economic environments and how our results relate to previous analyses of these environments; a detailed analysis of policy implications in each environment

¹⁷Supporting calculations for these applications can be found in Appendix B.

is beyond the scope of this paper.

4.1 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 assumed to control the resource initially. (In Section 5, we comment on how our framework could be extended to model the problem of the initial allocation of the optimally-designed property rights.) The agent decides in $t = 1$ 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 some alternative use in $t = 2$. The regulator is concerned with allocative efficiency, in that $V_\omega^*(\theta) = V_\omega(\theta) = \theta - \omega$. Additionally, the regulator cares about revenue, and may attach a higher weight to revenue at $t = 2$, that is, $\alpha \geq \alpha^* \geq 0$.

In this application, the agent is subject to a hold-up problem; additionally, the regulator suffers from time-inconsistency (if $\alpha > \alpha^*$). Time inconsistency could be, for example, the result of political pressure to raise a certain amount of revenue when reallocating scarce public resources.¹⁸ 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 relatively high cost of investment, we can apply the analysis from Section 3.3. 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. In this case, the optimal property right takes the form of an option-to-own combined with an option-to-sell. 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. The option-to-sell additionally gives the lessee the right to require monetary compensation for relinquishing control. Both rights are conditional on meeting the required investment level.

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).¹⁹ 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

¹⁸For 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).

¹⁹Early proponents of the self-assessment mechanism include Harberger (1965) and Tideman (1969).

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.²⁰ 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 (if she doesn't pay the price p , she may still keep the license but only if the state ω is low enough). 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.²¹

In practice, it could be difficult to assess the extent to which an efficient level of investment is undertaken. Thus, we next turn to the case when investment is not contractible. The optimal property right may become more complicated. By Corollary 3, the license may give two types of rights to the agent: an option-to-own with some price p , and a partial right that results in a full property right 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 determining the average likelihood of a favorable decision.

It turns out, however, that under additional assumptions the optimal license may be simpler—we illustrate this with a numerical example that further sheds light on the trade-offs involved in the optimal license design.

Numerical example. Assume that both 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.

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

²⁰Weyl 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.

²¹Implementing 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 5.

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 $\alpha = \alpha^ = 1$:* 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.

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

Case $\alpha = \alpha^ = 0$:* In this case, the regulator maximizes efficiency in both periods.²² 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.

The numerical example illustrates how the form of the optimal property right varies with the regulator’s ex-ante and ex-post preferences over revenue. If there’s no time inconsistency and the regulator is only concerned with efficiency, it is optimal to allocate no rights to the agent because the VCG mechanism employed to reallocate resources at $t = 2$ ensures efficient investment incentives.²³ When the regulator cares about revenue at the ex-post stage but not at the ex-ante stage, it is optimal to provide investment incentives through a partial property right. Intuitively, a full property right is the most effective way of inducing investment when investment is not observable, as it makes the agent fully internalize its benefits. However, a full property right would typically make the investment-obedience constraint slack while distorting the efficient ex-post reallocation of the resource. Indeed, when the agent holds a full property right, the regulator chooses to sacrifice allocative efficiency at the reallocation stage for some realizations of ω in order to decrease the monetary compensation paid to the agent. Thus, at the ex-ante stage, the regulator specifies a contract that awards the agent a full property right with just enough probability to make the investment-obedience

²²Formally, since we assumed $\alpha > 0$, we consider the limit of solutions as $\alpha \rightarrow 0$.

²³This case does not arise in the analysis of Weyl and Zhang (2022) because investment in their model creates a common value. We can capture investments in the common value of the resource by assuming that the distribution G of the regulator’s opportunity cost first-order stochastically dominates the corresponding distribution \underline{G} conditional on no investment. Under that assumption, there exists a region in the parameter space in which the VCG mechanism would not lead to the efficient investment level, and the agent would be optimally assigned a non-trivial property right.

constraint bind. Finally, when the regulator cares about revenue at the ex-ante stage as well, she wants to design property rights so as to increase the revenue from the reallocation mechanism. She thus optimally switches from incentivizing investment via a partial property right (that yields no revenue at $t = 2$) to a renewable lease which generates revenue through the renewal fee p . The renewable lease is still effective at inducing investment (provided that p is sufficiently low) because it makes the agent internalize the benefits from investment conditional on realizing a high value for the resource.

4.2 Regulating a rental market

Next, we introduce an application in which the designer and the principal are two separate entities with conflicting objectives. 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 in it (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: $V_\omega^*(\theta) = \theta - \omega$ and $\alpha^* = 0$.

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. A full property right is interpreted in this context as mandating a long-term lease; other feasible regulations take the form of rent control or giving the tenant the right to stay by paying a pre-specified rent to the rental company.

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 $\text{supp}(G) \subseteq [\Delta, 1 - \Delta]$.

As a benchmark, consider first the case when ω is known ex-ante. 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, 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 ω is initially unknown (and possibly correlated with θ), then the first-order condition determining the optimal price in the renewable-lease contract is

$$p = \mathbb{E}[\omega \mid p \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]] - \gamma\Delta.$$

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 $p < \underline{\theta}_\omega^*$ holds and the rental company will prefer to pay the agent to leave, rather than forgoing the market rental rate. Similarly, if ω is low, then $p > \bar{\theta}_\omega^*$ holds and the company will offer 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.

It is worth noting that regulation similar to the one described here 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. In terms of residential leases, rent-control policies that are common in large cities around the world impose bounds on how much rent can increase from period to period, although they do not typically give the tenant the *right* to stay. However, in many countries rent control is combined with some degree of protection against eviction, which to some extent approximates a renewable-lease contract.

4.3 Patent policy

A classical economic question is how to reward and incentivize innovation and scientific discoveries. 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 agent 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.²⁴ A simple calculation shows that this scenario corresponds to $V_\omega(\theta) \equiv \theta(1 - (3/2)\omega)$ and $V_\omega^*(\theta) \equiv V_1(\theta)$, for all $\omega \in \Omega$.

A property right in this application gives the innovator full monopoly power 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).²⁵ 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 is uncorrelated with ω , and that the density of θ is differentiable and non-decreasing.²⁶

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

²⁴For example, the principal could have redistributive preferences as in Dworzak [©] al. (2021).

²⁵This 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).

²⁶For large enough ω , $V_\omega(\theta)$ is negative and decreasing; in light of Remark 1, we need the assumption of non-decreasing density to ensure that $V_\omega(\theta)f(\theta)$ preserves that property. Economically, this means that the principal is mostly concerned about granting monopoly rights to a firm with low costs (which would not be the case if having low costs is statistically unlikely).

innovator with cash.²⁷ In that case, the optimal property right is a cash payment conditional on investment.

While cash prizes have been historically used to incentivize major discoveries,²⁸ in many cases regulators cannot verify whether an innovation is socially useful. Moreover, paying for discoveries could induce moral-hazard problems.²⁹ From now on we 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 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 decreasing, it might become optimal to “sell” the monopoly rights to firms with low costs. In that case, the optimal regulation may take 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).

4.4 Vaccine development

Next, 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 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

²⁷See [Kremer \(1998\)](#) for historical cases of patent buyouts and a detailed analysis of how governments can determine the buyout price.

²⁸For 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. The Millennium Prize Problems selected by the Clay Institute serve as a modern-day example.

²⁹[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 an additional financial incentives for “fake” discoveries.

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 $\theta \equiv -k$. If the regulator cares exclusively about patient welfare, then $V_\omega^*(\theta) = V_\omega(\theta) = \omega$. Additionally, we let $1 = \alpha \geq \alpha^*$.

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.³⁰

We assume that investment is observable (the government can verify that the vaccine is effective). By the analysis in Section 3.3 (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

$$p = \frac{\mathbb{E}[\omega \mid \omega \in [\underline{\omega}_p, \bar{\omega}_p]]}{\alpha^*},$$

for some functions $\underline{\omega}_p, \bar{\omega}_p$, 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 $\omega > \bar{\omega}_p$ (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 3.3, 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 $\alpha^* = 0$, that is, when the government is not concerned about revenue at the stage of signing the contract, p will be equal to the upper bound of the distribution of costs; it is optimal to commit to purchasing all vaccines. When $\alpha^* = 1$, 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

³⁰There 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](#)).

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 3.3: 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).

4.5 Supply chain contracting

Finally, we exploit the possible correlation between θ and ω to capture an application with the classical ratchet effect. There is a large firm (playing both the role of the designer and the principal) buying some amount x 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, the firm can learn the supplier’s costs; the state ω is a noisy signal of θ . Setting $V_\omega^*(\theta) = 1$ and $\alpha^* = 1$ corresponds to the firm proposing a contract to the supplier.

In this application, we can investigate the optimal form of contractual rights between two firms, similar to the problem considered by the incomplete-contracts literature. Firms can freely bargain given the realized information in the future, or effectively merge by having the large firm purchase the entire future production of the supplier. Intermediate arrangements are also possible, such as the commitment by the large firm to buy a certain number of inputs at a pre-specified price.

Theorem 1 predicts the form of the optimal contract for the large firm. If investment by the small 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 in general commit to a two-price scheme, committing to buy up to y units at some price p' , and 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.

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 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 3.3 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.³¹ However, the role of prices is different across these two results. In Nöldeke and Schmidt (1995), the option price is pinned down by a condition ensuring efficient investment by the seller, while the base price can be freely adjusted to affect the split of surplus between the two parties. In our framework, with observable investment, both the cash payment and the option price affect the seller’s incentive to invest—the option price is used to lower the cost of incentivizing investment by making sure that the seller captures some of the benefits from increasing her type (lowering her costs). With unobservable investment (which is in fact closer to the setting of Nöldeke and Schmidt, 1995), our optimal contract is potentially more complicated and features an additional price for delivering a fraction of the seller’s production capacity.

5 Concluding remarks

In this paper, we studied the design of property rights in an environment in which the designer cannot commit to future trading mechanisms, giving rise to ex-post inefficiency and a hold-up problem. We modeled property rights as a set of outside options available to the agent. This perspective allowed us to employ a mechanism-design approach to characterize the optimal property right. 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. We investigated several applications of our results, including the design of spectrum licenses, the regulation of private rental markets, patent policy and procurement contracting by governments and large firms. In this section, we briefly review extensions of our framework, and comment on future research directions.

Property rights as a form of partial commitment. The frictions in our framework result from the inability of the designer to commit to future trading mechanisms. From this perspective, property rights are partly restoring 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 example, we could allow agent’s rights to be state-contingent—this would not affect our theoretical results significantly but would make property rights a more powerful tool for the designer, and could lead to new insights in applications. Another possibility, commonly encountered in

³¹The comparison would not be affected if we assumed that the seller maximizes total surplus—rather than her own profit—when choosing the initial contract.

practice, is that the designer might be able to *ban* certain outcomes (for example, 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 outcome, 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. We assumed throughout that the state ω is publicly observable but not contractible. As discussed in the preceding paragraph, a model in which rights can be made contingent on ω would give more power to the designer. 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*.³² It is easy to show, by means of examples, that this extra flexibility may benefit the designer. 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. One of the key assumptions of our framework was that property rights affect outside options but not bargaining power. It is then natural to model the trading stage as the problem of optimal mechanism design by a principal. 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 “one-sided” framework is to symmetrize the positions of the principal and the agent by endowing both of them with private information and endogenizing the bargaining power. The trading stage could be modeled as a third party running an incentive-compatible mechanism á la 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

³²This extension of our model would bear some similarity to the critique of the incomplete-contracts model offered by Maskin and Tirole (1999).

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 optimal property rights. In this paper, 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.

A direct extension of our framework would feature N agents with private signals, realized at $t = 0$, about their values conditional on having control over the resource in the future and undertaking investment. The designer would then have to take into account how the design of property rights affects the outcomes of the mechanism run at time $t = 0$ to allocate these rights to one of the agents. If the designer were only concerned with the efficiency of the allocation, we conjecture that under appropriate single-crossing assumptions (higher signal realization are associated with a higher distribution of future values), our characterization of optimal property rights would apply with minimal modifications. The reason is that a standard second price auction would allocate the rights to the agent with the highest signal realization regardless of the exact design of these rights. However, if the designer were additionally concerned with the revenue raised at $t = 0$, the optimal design of property rights would interact non-trivially with the optimal design of the mechanism to allocate them.³³ The primary link would be willingness to pay; for example, by designing stronger spectrum licenses, the designer could increase the bidders' values in the auction allocating them, but at the cost of lowering the revenue from future auctions. 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.

³³Similar interactions have been analyzed in the literature on bidding with securities (see, for example, DeMarzo et al., 2005).

- AKBARPOUR, M., K. LI, S. LI, S. D. KOMINERS, AND P. MILGROM (2023): “Algorithmic Mechanism Design with Investment,” *Econometrica* (forthcoming).
- AKBARPOUR, M. [Ⓔ] P. DWORCZAK [Ⓔ] S. D. KOMINERS (2023): “Redistributive Allocation Mechanisms,” *Journal of Political Economy* (forthcoming).
- 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 Hagglng, Coordination, and Rent-Seeking,” Working paper.
- BAUER, H. (1958): “Minimalstellen von Funktionen und Extrempunkte,” *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.
- 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 (2022): “Constrained Information Design,” *Mathematics of Operations Research* (forthcoming).
- DWORCZAK, P. [Ⓔ] S. D. KOMINERS [Ⓔ] M. AKBARPOUR (2021): “Redistribution through Markets,” *Econometrica*, 89, 1665–1698.
- FUCHS, W. AND A. SKRZYPACZ (2015): “Government Interventions in a Dynamic Market

- with Adverse S,” *Journal of Economic Theory*, 158, 371–406.
- GALLINI, N. T. (1992): “Patent Policy and Costly Imitation,” *The RAND Journal of Economics*, 23, 52–63.
- 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.
- HATFIELD, J. W., F. KOJIMA, AND S. D. KOMINERS (2019): “Strategy-Proofness, Investment Efficiency, and Marginal Returns: An Equivalence,” *Becker Friedman Institute for Research in Economics Working Paper*.
- 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.
- 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. 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 (2023): “Optimal Labor Procurement under Minimum Wages and Monopsony Power,” Working paper.
- 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., 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. SATTERTHWAITTE (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 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.
- 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. Séries des Sciences Mathématiques, Astronomiques et Physiques*, 23, 389–393.
- TIDEMAN, N. (1969): “Three Approaches to Improving Urban Land Use,” *PhD dissertation, University of Chicago*.
- 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 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 $\theta \geq \bar{\theta}$.³⁴ Applying Leibniz's rule, integrating by parts, and using $\mathcal{W}(\bar{\theta}) = 0$ and $\lim_{\theta \nearrow \underline{\theta}} x(\theta) = 0$:

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

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

³⁴While 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.

x must satisfy the inequality as an equality (whenever this is feasible). Formally, define

$$\bar{x}(\theta) := x_0(\theta)\mathbf{1}_{\theta \geq \theta_0},$$

where θ_0 is defined by

$$u_0 - \underline{u} + \int_{\underline{\theta}}^{\theta_0} x_0(\tau) d\tau = 0 \text{ (and } \theta_0 = \underline{\theta} \text{ if there is no solution).} \quad (6)$$

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 $\mathcal{W}(\bar{\theta}^*)$ —for all $\theta \geq \bar{\theta}^*$, where $\bar{\theta}^*$ 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^* 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 $(\underline{\theta}, \bar{\theta}^*)$ 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 $(\underline{\theta}, \bar{\theta}^*)$, and the latter on $(\underline{\theta}^*, \bar{\theta}^*)$. Define

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

Intuitively, x^* (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^* 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 $(\underline{\theta}, \bar{\theta}^*)$, and by a single first-order stochastic dominance shift above $\bar{\theta}^*$ —this suffices for feasibility, as noted previously.

We now argue that x^* achieves the upper bound of the value function. Let $x_-^*(\bar{\theta}_\omega^*)$ denote the left limit of x^* at $\bar{\theta}_\omega^*$. Then,

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \mathcal{W}(\theta) dx^*(\theta) &= \int_{(\underline{\theta}, \bar{\theta}^*) \cup \mathcal{I}'} \mathcal{W}(\theta) dx^*(\theta) + \mathcal{W}(\bar{\theta}^*)(1 - x_-^*(\bar{\theta}^*)) \\ &= \int_{(\underline{\theta}, \bar{\theta}^*) \cup \mathcal{I}'} \bar{\mathcal{W}}(\theta) dx^*(\theta) + \bar{\mathcal{W}}(\bar{\theta}^*)(1 - x_-^*(\bar{\theta}^*)) = \int_{\underline{\theta}}^{\bar{\theta}} \bar{\mathcal{W}}_+(\theta) d\bar{x}(\theta), \end{aligned} \quad (7)$$

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

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

where $\theta_0(\underline{u})$ is defined as in (6), 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:

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

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

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

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

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

Integrating the second term by parts yields another equivalent representation:

$$\max_{\theta_0 \geq \underline{\theta}} \left\{ \int_{[\underline{\theta}, \theta_0]} (\bar{\mathcal{W}}'_+(\theta) - \alpha)x_0(\theta) d\theta \right\}. \quad (8)$$

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

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

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

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

A.3 Proof of Lemma 3

In the proof, we separately address the contractible and the non-contractible case (with respect to the investment decision of the agent). To streamline exposition, we first cover the non-contractible case, and then explain how to modify the proof to cover the contractible case.

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 and on the public state ω —we will make that dependence explicit in our notation. In particular, let $\underline{u}_\omega^*(R) := R(\underline{\theta}_\omega^*)$, and let $\langle x_\omega^*(\theta; R), t_\omega^*(\theta; R) \rangle$ denote the mechanism chosen by the principal conditional on ω and R .

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

$$\int_{\Omega} \left(\underline{u}_\omega^*(R) + \int_{\underline{\theta}_\omega^*}^{\bar{\theta}_\omega^*} x_\omega^*(\theta; R)(1 - \tilde{F}_\omega(\theta))d\theta + \int_{\bar{\theta}_\omega^*}^{\bar{\theta}} (\theta - \bar{\theta}_\omega^*)d\tilde{F}_\omega(\theta) \right) \tilde{G}(\omega),$$

where $\tilde{F}_\omega = F_\omega$ and $\tilde{G} = G$ if the agent invested, and $\tilde{F}_\omega = \underline{F}_\omega$ and $\tilde{G} = \underline{G}$ otherwise. In particular, when the agent has no rights ($R \equiv 0$), the principal allocates the good with probability one to types $\theta \geq \bar{\theta}_\omega^*$ (and with probability zero otherwise). Define

$$\tilde{c} := c - \left(\int_{\Omega} \int_{\bar{\theta}_\omega^*}^{\bar{\theta}} (\theta - \bar{\theta}_\omega^*)dF_\omega(\theta)dG(\omega) - \int_{\Omega} \int_{\bar{\theta}_\omega^*}^{\bar{\theta}} (\theta - \bar{\theta}_\omega^*)d\underline{F}_\omega(\theta)d\underline{G}(\omega) \right)$$

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, $\tilde{c} > 0$. We can now write the agent's obedience constraint as

$$\begin{aligned} & \int_{\Omega} \left(R(\underline{\theta}_\omega^*) + \sum_{(a,b) \in \mathcal{I}_\omega} \frac{\int_a^b R'(\tau)d\tau}{b-a} \int_a^b (1 - F_\omega(\theta))d\theta + \int_{\substack{(\underline{\theta}_\omega^*, \bar{\theta}_\omega^*) \\ \setminus \cup \mathcal{I}_\omega}} R'(\theta)(1 - F_\omega(\theta))d\theta \right) dG(\omega) - \tilde{c} \\ & \geq \int_{\Omega} \left(R(\underline{\theta}_\omega^*) + \sum_{(a,b) \in \mathcal{I}_\omega} \frac{\int_a^b R'(\tau)d\tau}{b-a} \int_a^b (1 - \underline{F}_\omega(\theta))d\theta + \int_{\substack{(\underline{\theta}_\omega^*, \bar{\theta}_\omega^*) \\ \setminus \cup \mathcal{I}_\omega}} R'(\theta)(1 - \underline{F}_\omega(\theta))d\theta \right) d\underline{G}(\omega). \end{aligned}$$

Next, denoting by $W_\omega^*(\theta) := (V_\omega^*(\theta) + \alpha^*B_\omega(\theta))f_\omega(\theta)$ the designer's objective multiplied by the density of types, we can write her expected payoff conditional on choosing an outside option function R (and conditional on the agent investing) as

$$\int_{\Omega} \left(-\alpha^*R(\underline{\theta}_\omega^*) + \sum_{(a,b) \in \mathcal{I}_\omega} \frac{\int_a^b R'(\tau)d\tau}{b-a} \int_a^b W_\omega^*(\theta)d\theta + \int_{\substack{(\underline{\theta}_\omega^*, \bar{\theta}_\omega^*) \\ \setminus \cup \mathcal{I}_\omega}} R'(\theta)W_\omega^*(\theta)d\theta \right) dG(\omega),$$

where we have omitted the term $\int_{\Omega} (\int_{\bar{\theta}_\omega^*}^{\bar{\theta}} W_\omega^*(\theta)d\theta)dG(\omega)$ that does not depend on the chosen R .

We can now change variables 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 x and $u \geq 0$

$$-\alpha^*u + \int_{\Omega} \left(-\alpha^* \int_{\underline{\theta}}^{\underline{\theta}_{\omega}^*} x(\theta) d\theta + \sum_{(a,b) \in \mathcal{I}_{\omega}} \frac{\int_a^b x(\tau) d\tau}{b-a} \int_a^b W_{\omega}^*(\theta) d\theta + \int_{\bigcup_{\omega \in \mathcal{I}_{\omega}} (\underline{\theta}_{\omega}^*, \bar{\theta}_{\omega}^*)} x(\theta) W_{\omega}^*(\theta) d\theta \right) dG(\omega)$$

subject to

$$\begin{aligned} & \int_{\Omega} \left(\int_{\underline{\theta}}^{\underline{\theta}_{\omega}^*} x(\theta) d\theta + \sum_{(a,b) \in \mathcal{I}_{\omega}} \frac{\int_a^b x(\tau) d\tau}{b-a} \int_a^b (1 - F_{\omega}(\theta)) d\theta + \int_{\bigcup_{\omega \in \mathcal{I}_{\omega}} (\underline{\theta}_{\omega}^*, \bar{\theta}_{\omega}^*)} x(\theta) (1 - F_{\omega}(\theta)) d\theta \right) dG(\omega) - \tilde{c} \\ & \geq \int_{\Omega} \left(\int_{\underline{\theta}}^{\underline{\theta}_{\omega}^*} x(\theta) d\theta + \sum_{(a,b) \in \mathcal{I}_{\omega}} \frac{\int_a^b x(\tau) d\tau}{b-a} \int_a^b (1 - \underline{F}_{\omega}(\theta)) d\theta + \int_{\bigcup_{\omega \in \mathcal{I}_{\omega}} (\underline{\theta}_{\omega}^*, \bar{\theta}_{\omega}^*)} x(\theta) (1 - \underline{F}_{\omega}(\theta)) d\theta \right) d\underline{G}(\omega). \end{aligned}$$

Since both the objective and the constraints are linear in $x(\theta)$, using integration by parts,³⁵ we can rewrite the problem as

$$\max_{x(\theta), u \geq 0} \int_{\underline{\theta}}^{\bar{\theta}} \Phi(\theta) dx(\theta) - \alpha^*u \quad \text{subject to} \quad \int_{\underline{\theta}}^{\bar{\theta}} \Psi(\theta) dx(\theta) \geq \tilde{c}, \quad (9)$$

where

$$\begin{aligned} \Phi(\theta) = & \int_{\Omega} \left(-\alpha^* (\underline{\theta}_{\omega}^* - \theta)_+ + \sum_{(a,b) \in \mathcal{I}_{\omega}} (b - \max\{a, \theta\})_+ \frac{\int_a^b W_{\omega}^*(\theta) d\theta}{b-a} \right. \\ & \left. + \sum_{[a,b] \in \mathcal{I}_{\omega}^c} \mathbf{1}_{\{\theta \leq b\}} \left(\int_{\max\{a, \theta\}}^b W_{\omega}^*(\tau) d\tau \right) \right) dG(\omega), \end{aligned}$$

³⁵In particular, we use the fact that $\int_a^b g(\theta)x(\theta)d\theta = \int_{\underline{\theta}}^{\bar{\theta}} \mathbf{1}_{\{\theta \leq b\}} \left(\int_{\max\{a, \theta\}}^b g(\tau)d\tau \right) dx(\theta)$.

and

$$\begin{aligned}
\Psi(\theta) = & \int_{\Omega} \left((\underline{\theta}_{\omega}^* - \theta)_+ + \sum_{(a,b) \in \mathcal{I}_{\omega}} (b - \max\{a, \theta\})_+ \frac{\int_a^b (1 - F_{\omega}(\theta)) d\theta}{b - a} \right. \\
& \left. + \sum_{[a,b] \in \mathcal{I}_{\omega}^c} \mathbf{1}_{\{\theta \leq b\}} \left(\int_{\max\{a, \theta\}}^b (1 - F_{\omega}(\tau)) d\tau \right) \right) dG(\omega) \\
- & \int_{\Omega} \left((\underline{\theta}_{\omega}^* - \theta)_+ + \sum_{(a,b) \in \mathcal{I}_{\omega}} (b - \max\{a, \theta\})_+ \frac{\int_a^b (1 - \underline{F}_{\omega}(\theta)) d\theta}{b - a} \right. \\
& \left. + \sum_{[a,b] \in \mathcal{I}_{\omega}^c} \mathbf{1}_{\{\theta \leq b\}} \left(\int_{\max\{a, \theta\}}^b (1 - \underline{F}_{\omega}(\tau)) d\tau \right) \right) d\underline{G}(\omega).
\end{aligned}$$

Thus, we have represented the designer's problem as maximizing a linear functional subject to a single linear constraint.

The contractible case. In the contractible case, the transformations are analogous but notation is further complicated by the fact that, conditional on no investment, the principal's mechanism is designed optimally for the distribution \underline{F}_{ω} of the agent's type. We define the cost of investment net of the agent's benefit from investing in the absence of any rights as

$$\tilde{c} := c - \left(\int_{\Omega} \int_{\bar{\theta}_{\omega}^*}^{\bar{\theta}} (\theta - \bar{\theta}_{\omega}^*) dF_{\omega}(\theta) dG(\omega) - \int_{\Omega} \int_{\underline{\bar{\theta}}_{\omega}^*}^{\bar{\theta}} (\theta - \underline{\bar{\theta}}_{\omega}^*) d\underline{F}_{\omega}(\theta) d\underline{G}(\omega) \right),$$

where $\underline{\bar{\theta}}_{\omega}^*$ denotes the analog of $\bar{\theta}_{\omega}^*$ obtained by replacing F_{ω} with \underline{F}_{ω} in its definition. The agent's obedience constraint in the designer's problem becomes

$$\int_{\Omega} \left(u + \int_{\underline{\theta}}^{\underline{\theta}_{\omega}^*} x(\theta) d\theta + \sum_{(a,b) \in \mathcal{I}_{\omega}} \frac{\int_a^b x(\tau) d\tau}{b-a} \int_a^b (1 - F_{\omega}(\theta)) d\theta + \int_{\substack{(\underline{\theta}_{\omega}^*, \bar{\theta}_{\omega}^*) \\ \setminus \bigcup \mathcal{I}_{\omega}}} x(\theta) (1 - F_{\omega}(\theta)) d\theta \right) dG(\omega) \geq \tilde{c}.$$

Finally, we can write the designer's problem as

$$\max_{x(\theta), u \geq 0} \int_{\underline{\theta}}^{\bar{\theta}} \Phi(\theta) dx(\theta) - \alpha^* u \quad \text{subject to} \quad \int_{\underline{\theta}}^{\bar{\theta}} \Psi(\theta) dx(\theta) + u \geq \tilde{c}, \quad (10)$$

where $\Phi(\theta)$ is defined as in the non-contractible case, and

$$\begin{aligned} \Psi(\theta) = \int_{\Omega} & \left((\underline{\theta}_{\omega}^* - \theta)_+ + \sum_{(a,b) \in \mathcal{I}_{\omega}} (b - \max\{a, \theta\})_+ \frac{\int_a^b (1 - F_{\omega}(\theta)) d\theta}{b - a} \right. \\ & \left. + \sum_{[a,b] \in \mathcal{I}_{\omega}^c} \mathbf{1}_{\{\theta \leq b\}} \left(\int_{\max\{a, \theta\}}^b (1 - F_{\omega}(\tau)) d\tau \right) \right) dG(\omega). \end{aligned}$$

By using the indicator function $\mathbf{1}_{\text{cont}}$, we obtain the unified statement of Lemma 3 covering both the contractible and non-contractible case.

A.4 Proof of Lemma 4

By Lemma 3, 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.³⁶ Extreme points in the space of (non-decreasing) allocation rules are cutoff functions of the form $\mathbf{1}_{\theta \geq \theta^*}$. Thus, the optimal x can be written as a two-step function, and in particular its image may contain at most one value other than 0 or 1.

In the non-contractible case, it is clear that it is optimal to set $u = 0$. This corresponds to case (i) in Lemma 4. The same conclusion is true in the contractible case when $u = 0$ in the optimal solution.

Suppose that $u > 0$ in the optimal solution in the contractible case. Observe that the optimal solution must maximize the Lagrangian, with Lagrange multiplier γ ,³⁷

$$\int_{\underline{\theta}}^{\bar{\theta}} (\Phi(\theta) + \gamma \Psi(\theta)) dx(\theta) + (\gamma - \alpha)u,$$

and that in case $u > 0$ is optimal, we must have $\gamma = \alpha^*$. Indeed, $\alpha^* \leq \gamma$ as otherwise the unique optimal choice would be $u = 0$, and $\alpha^* \geq \gamma$ as otherwise the Lagrangian would not have a maximum. But if $\gamma = \alpha^*$, then any $u \geq 0$ maximizes the Lagrangian. Thus, we can pick a cutoff allocation rule $x(\theta)$ maximizing $\int_{\underline{\theta}}^{\bar{\theta}} (\Phi(\theta) + \alpha^* \Psi(\theta)) dx(\theta)$ that does not satisfy the obedience constraint when paired with $u = 0$,³⁸ and then satisfy the obedience constraint by picking $u \geq 0$, so that $\int_{\underline{\theta}}^{\bar{\theta}} \Psi(\theta) dx(\theta) + u = \tilde{c}$. This corresponds to case (ii) in Lemma 4.

³⁶Formally, this follows from the results of Bauer (1958) and Szapiel (1975), as summarized by Kang (2023), which can be seen as a version of Carathéodory's theorem for an infinite-dimensional linear space.

³⁷Existence of a Lagrange multiplier follows from Theorem 2.165 in Bonnans and Shapiro (2000).

³⁸Such an x must exist, as otherwise we could not have a solution with $u > 0$.

A.5 Proofs of Corollaries 1, 2, and 3

Corollary 1 follows by direct inspection of the solution derived in Lemma 2. Corollary 2 follows from the proof of Lemma 4 by observing that if the investment-obedience constraint is dropped from the designer's optimization problem, then the optimum is attained by an extreme point $x(\theta) = \mathbf{1}_{\theta \geq \theta^*}$ and $u = 0$. This corresponds to offering the agent an option-to-own. Finally, Corollary 3 follows from the proof of Lemma 4 and two observations. In the non-contractible case, the conclusion that $u = 0$ implies that $y > 0$ in the corresponding optimal menu M^* . Moreover, if p'/y did not belong to $[\underline{\theta}, \bar{\theta}]$, it would never be chosen by any type of the agent (hence, it could be replaced by $(y, y\bar{\theta})$ without affecting the designer's payoff). In the contractible case, if c is high enough, no R with $R(\underline{\theta}) = 0$ satisfies the investment-obedience constraint. Thus, it must be that $u > 0$ in the solution to the designer's problem, which corresponds to including the option $(0, -T)$, with $T > 0$, in the optimal menu.

A.6 Remark about tie-breaking rules

In the proof of Theorem 1, we have assumed a particular tie-breaking rule in case of principal's indifference, implicit in how we defined the cutoffs $\underline{\theta}^*$, $\bar{\theta}^*$ 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^* must satisfy the string of equalities (A.2), and any optimal $\underline{\theta}^*$ must solve problem (8). It follows that all solutions to problem (P') can be obtained by modifying our baseline solution (x^*, \underline{u}^*) in the following ways:

1. $\bar{\theta}^*$ can be taken to be *any* global maximum of \mathcal{W} (not necessarily the smallest one);
2. If $\mathcal{W} = \bar{\mathcal{W}}$ is *affine* on some interval $[a, b]$, then we can take any mean-preserving spread of x^* in that interval (in the baseline solution, $x^*(\theta) = R'(\theta)$ on $[a, b]$);
3. $\underline{\theta}^*$ can be taken to be *any* type θ with the property $\alpha = \bar{\mathcal{W}}'(\theta)$ 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 $\int_{\underline{\theta}}^{\bar{\theta}} \phi(\theta)x^*(\theta)d\theta - \beta \underline{u}^*$, where $\phi : \Theta \rightarrow \mathbb{R}$ 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 $\bar{\theta}^*$ and $\underline{\theta}^*$ will be invariant to R ; moreover, maximizing $\int_{\underline{\theta}}^{\bar{\theta}} \phi(\theta)x^*(\theta)d\theta$ 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 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 $x^*(\theta)$ is equal to $R'(\theta)$, in which case the subinterval can be included in the collection \mathcal{I}^c , or (ii) the optimal $x^*(\theta)$ 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 $\underline{\theta}^*$ and $\bar{\theta}^*$, as well as the R -invariant collection of ironing intervals, may be different. Thus, the solution is still linear in R .

A.7 Proof of Proposition 1

We begin with a technical lemma.

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

Proof of Lemma A.1. We drop the subscript ω to simplify the exposition. We first prove that $W(\theta) = (V(\theta) + \alpha B(\theta))f(\theta)$ is non-decreasing on $[\underline{\theta}^*, \bar{\theta}^*]$. It suffices to show that, for all $\theta \in [\underline{\theta}^*, \bar{\theta}^*]$,

$$\frac{V'(\theta) + 2\alpha}{\alpha} f(\theta) + \frac{V(\theta) + \alpha\theta}{\alpha} f'(\theta) \geq 0.$$

Using the definition of $\underline{\theta}^*$ and $\bar{\theta}^*$ given in Proposition 1, for all $\theta \in [\underline{\theta}^*, \bar{\theta}^*]$, 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 $\theta \leq \underline{\theta}^*$, that is,

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

Indeed, we have

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

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

Overall, we have shown that $\mathcal{W}(\theta) = \int_{\theta}^{\bar{\theta}} W(\tau) d\tau$ has a slope higher than α for $\theta \leq \underline{\theta}^*$ and lower than α for $\theta \geq \underline{\theta}^*$, is concave on $[\underline{\theta}^*, \bar{\theta}^*]$, and has a global maximum at $\bar{\theta}^*$ (since $W(\theta)$ crosses zero once from below at $\bar{\theta}^*$). It follows that \mathcal{W} is equal to its concave closure on $[\underline{\theta}^*, \bar{\theta}^*]$. Moreover, $\bar{\mathcal{W}}'(\theta) \leq \alpha$ for $\theta \leq \underline{\theta}^*$, with equality at $\theta = \underline{\theta}^*$. Finally, $\underline{\theta}^*$ and $\bar{\theta}^*$, as defined in Proposition 1, correspond to the $\underline{\theta}^*$ and $\bar{\theta}^*$ defined in Section 3.1. \square

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

A.8 Supplementary material for Section 3.3

Following the proof of Theorem 1, and relying on Proposition 1 to simplify the designer's problem, we obtain

$$\max_{x(\theta), u} \int_{\underline{\theta}}^{\bar{\theta}} \Phi(\theta) dx(\theta) - \alpha^* u \quad \text{subject to} \quad \int_{\underline{\theta}}^{\bar{\theta}} \Psi(\theta) dx(\theta) + \mathbf{1}_{\text{cont}} \cdot u \geq \tilde{c},$$

where

$$\Phi(\theta) := \int_{\Omega} \left(-\alpha^* (\underline{\theta}_\omega^* - \theta)_+ + \mathbf{1}_{\{\theta \leq \bar{\theta}_\omega^*\}} \int_{\max\{\underline{\theta}_\omega^*, \theta\}}^{\bar{\theta}_\omega^*} [V_\omega^*(\tau) + \alpha^* B_\omega(\tau)] dF_\omega(\tau) \right) dG(\omega) \equiv \mathbb{E}_{\omega \sim G}[\Phi_\omega(\theta)],$$

and, given the assumption $G = \underline{G}$,

$$\Psi(\theta) := \int_{\Omega} \left(\mathbf{1}_{\{\theta \leq \bar{\theta}_\omega^*\}} \left(\int_{\max\{\underline{\theta}_\omega^*, \theta\}}^{\bar{\theta}_\omega^*} (F_\omega(\tau) - F_\omega(\tau)) d\tau \right) \right) dG(\omega) \equiv \mathbb{E}_{\omega \sim G}[\Psi_\omega(\theta)]$$

in the non-contractible case, while

$$\Psi(\theta) := \int_{\Omega} \left((\underline{\theta}_{\omega}^* - \theta)_+ + \mathbf{1}_{\{\theta \leq \bar{\theta}_{\omega}^*\}} \int_{\max\{\underline{\theta}_{\omega}^*, \theta\}}^{\bar{\theta}_{\omega}^*} (1 - F_{\omega}(\tau)) d\tau \right) dG(\omega) \equiv \mathbb{E}_{\omega \sim G}[\Psi_{\omega}(\theta)]$$

in the contractible case.

As in the proof of Theorem 1, we can study the behavior of the Lagrangian

$$\int_{\underline{\theta}}^{\bar{\theta}} (\Phi(\theta) + \gamma \Psi(\theta)) dx(\theta) + (\mathbf{1}_{\text{cont}} \cdot \gamma - \alpha)u,$$

where γ is the Lagrange multiplier on the investment-obedience constraint.

In the contractible case, we know from the proof of Theorem 1 that when $u > 0$ (which is necessarily true in the optimal mechanism when the cost of investment is high enough), we must have $\gamma = \alpha^*$. We also know that $x(\theta) = \mathbf{1}_{\theta \geq \theta^*}$, and since the optimal x maximizes the Lagrangian that has one-sided derivatives everywhere, the optimal θ^* must satisfy the generalized first-order condition

$$\Phi'(\theta^*) + \alpha^* \Psi'(\theta^*) \stackrel{(FOC)}{=} 0, \quad (11)$$

where $\stackrel{(FOC)}{=}$ 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. Since we have

$$\Phi'_{\omega}(\theta) + \alpha^* \Psi'_{\omega}(\theta) = \begin{cases} 0 & \theta < \underline{\theta}_{\omega}^*, \\ -[V_{\omega}^*(\theta, \omega) + \alpha^* \theta] f_{\omega}(\theta) & \theta \in (\underline{\theta}_{\omega}^*, \bar{\theta}_{\omega}^*), \\ 0 & \theta > \bar{\theta}_{\omega}^*. \end{cases}$$

the first-order condition becomes

$$-\mathbb{E}_{\omega \sim G} \left[\mathbf{1}_{\{\theta^* \in [\underline{\theta}_{\omega}^*, \bar{\theta}_{\omega}^*]\}} (V_{\omega}^*(\theta^*) + \alpha^* \theta^*) f_{\omega}(\theta^*) \right] \stackrel{(FOC)}{=} 0.$$

It is without loss of generality to restrict the set of candidate θ^* to the closure of the set $\{\theta^* \in \Theta : \mathbb{P}_{\omega \sim G}(\theta^* \in [\underline{\theta}_{\omega}^*, \bar{\theta}_{\omega}^*]) > 0\}$ since the Lagrangian is constant in θ^* outside this range. If $u > 0$ is interpreted as a lump-sum payment, then setting the price $p = \theta^*$ in the option-to-own implements the outside option function $R(\theta) = u + \int_{\underline{\theta}}^{\theta} x(\tau) d\tau$. A straightforward

transformation of the first-order condition yields formula (4).

In the non-contractible case, we know that $u = 0$ and the optimal x takes the form

$$x(\theta) = \begin{cases} 0 & \theta < \theta_1^*, \\ y & \theta_1^* \leq \theta < \theta_2^*, \\ 1 & \theta \geq \theta_2^*. \end{cases}$$

Since the optimal x must maximize the Lagrangian, both θ_1^* and θ_2^* must satisfy the first-order condition (11), with α^* replaced by a generic Lagrange multiplier γ . Since we have

$$\Phi'_\omega(\theta) + \gamma\Psi'_\omega(\theta) = \begin{cases} \alpha^* & \theta < \underline{\theta}_\omega^*, \\ - \left[V_\omega^*(\theta) + \alpha^* B_\omega(\theta) + \gamma \frac{F_\omega(\theta) - F_\omega(\theta^*)}{f_\omega(\theta)} \right] f_\omega(\theta) & \theta \in (\underline{\theta}_\omega^*, \bar{\theta}_\omega^*), \\ 0 & \theta > \bar{\theta}_\omega^*, \end{cases}$$

the first-order condition is

$$\alpha^* \mathbb{P}_{\omega \sim G}(\theta^* < \underline{\theta}_\omega^*) - \mathbb{E}_{\omega \sim G} \left[\left(V_\omega^*(\theta^*) + \alpha^* B_\omega(\theta^*) + \gamma \frac{F_\omega(\theta^*) - F_\omega(\theta^*)}{f_\omega(\theta^*)} \right) f_\omega(\theta) \mathbf{1}_{\theta^* \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]} \right] \stackrel{(FOC)}{=} 0.$$

It is again without loss of generality to restrict the set of candidate θ^* to the closure of the set $\{\theta^* \in \Theta : \mathbb{P}_{\omega \sim G}(\theta^* \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]) > 0\}$; in particular, if $\mathbb{P}_{\omega \sim G}(\theta^* \geq \underline{\theta}_\omega^*) = 0$ for θ^* below some threshold, then θ^* below that threshold cannot be optimal (since the Lagrangian is increasing in that region). Thus, we can rewrite the condition as

$$\alpha^* \frac{\mathbb{P}_{\omega \sim G}(\theta^* < \underline{\theta}_\omega^*)}{\mathbb{P}_{\omega \sim G}(\theta^* \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*])} - \mathbb{E}_{\omega \sim G} \left[(V_\omega^*(\theta^*) + \alpha^* B_\omega(\theta^*)) f_\omega(\theta^*) \mid \theta^* \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*] \right] - \gamma \mathbb{E}_{\omega \sim G} \left[F_\omega(\theta^*) - F_\omega(\theta^*) \mid \theta^* \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*] \right] \stackrel{(FOC)}{=} 0.$$

If there is only a single point θ^* satisfying the first-order condition, then offering a singleton menu with option-to-own with price $p = \theta^*$ is optimal. This gives us condition (5). If instead the optimal menu contains two options, then $\theta_1^* < \theta_2^*$ must both satisfy the first-order condition, while prices in the optimal menu $M^* = \{(1, p), (y, p')\}$ are given by $p' = y\theta_1^*$ and $p = \theta_2^* - y(\theta_2^* - \theta_1^*)$.

B Supporting calculations for Section 4

B.1 Calculations for Subsection 4.1

Using the result derived in Proposition 1, we can explicitly calculate the interval $[\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]$ on which the outside option constraint binds:

$$\underline{\theta}_\omega^* = \frac{\omega}{1+2\alpha} \quad \text{and} \quad \bar{\theta}_\omega^* = \frac{\omega + \alpha}{1+2\alpha}.$$

Let us 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 $\underline{\theta}_\omega^* = \bar{\theta}_\omega^* = \omega$), which proves the claim for the case $\alpha = \alpha^* = 0$. From now on, we assume that $\alpha = 1$.

Using the notation from Appendix A.8, we have

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

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

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

and

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

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

Case $\alpha^* = 1$. In this case, we have

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

and

$$\int_0^1 [\Phi''_{\omega}(\theta) + \gamma \Psi''_{\omega}(\theta)] dG(\omega) = \begin{cases} \theta(6\gamma - 9) - 3\gamma & \theta < 1/3, \\ -\theta(6\gamma - 9) - 6 + 5\gamma & \theta \in (1/3, 2/3), \\ 0 & \theta > 2/3. \end{cases}$$

In the interval $[0, 1/3]$, the function is 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 $\alpha^* = 0$. In this case, the derivatives are given by

$$\int_0^1 [\Phi'_\omega(\theta) + \gamma\Psi'_\omega(\theta)] dG(\omega) = \begin{cases} (3\gamma + \frac{3}{2})\theta^2 - 3\theta\gamma & \theta < 1/3, \\ -(3\gamma + \frac{3}{2})\theta^2 + (5\gamma + 1)\theta - 2\gamma & \theta \in (1/3, 2/3), \\ 0 & \theta > 2/3, \end{cases}$$

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

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 Calculations for Subsection 4.2

Using Proposition 1, we can pin down the interval $[\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]$ on which the outside option constraint binds:

$$\omega = \underline{\theta}_\omega^* + \frac{F(\underline{\theta}_\omega^*)}{f(\underline{\theta}_\omega^*)} \quad \text{and} \quad \omega = \bar{\theta}_\omega^* - \frac{1 - F(\bar{\theta}_\omega^*)}{f(\bar{\theta}_\omega^*)}.$$

Due to our assumption that $\text{supp}(G) \subseteq [\Delta, 1 - \Delta]$, we have that $\Delta \leq \underline{\theta}_\omega^* \leq \bar{\theta}_\omega^* \leq 1$. This in turn implies that $\underline{F}(\theta) - F(\theta) = \Delta$ for all $\theta \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]$.

Suppose first that ω is known. Then,

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

We thus have a unique maximum at

$$\theta^* = \omega - \gamma\Delta.$$

By Rogerson (1992), if investment is socially efficient, then setting $\gamma = 0$ will incentivize the agent to invest, and hence an option-to-own with price ω must be optimal.

Now let us suppose that $\omega \sim G$. Then,

$$\Phi'_\omega(\theta) + \gamma\Psi'_\omega(\theta) = -\mathbf{1}_{\{\underline{\theta}_\omega^* \leq \theta \leq \bar{\theta}_\omega^*\}} [\theta - \omega + \gamma\Delta].$$

Thus, the optimal p must satisfy the first-order condition:

$$p = \mathbb{E}_\omega [\omega \mid p \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]] - \gamma\Delta.$$

B.3 Calculations for Subsection 4.3

First, we make a general observation. Dropping the dependence on ω in the notation, suppose that

$$\mathcal{W}(\theta) \leq \frac{\bar{\theta} - \theta}{\bar{\theta} - \underline{\theta}} \mathcal{W}(0).$$

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

2. If $\mathcal{W}(0) \in [-\alpha, 0]$, then $x^*(\theta) = \frac{R(\underline{\theta}) - R(\theta)}{\theta - \underline{\theta}}$, and $u^* = R(\underline{\theta})$, that is, the allocation rule is constant.
3. If $\mathcal{W}(0) > 0$, then $x^*(\theta) \equiv 1$, and $u^* = R(\underline{\theta})$, that is, the agent always gets the good.

Note also that it is easy to modify our methods to handle the case in which the principal is not allowed to pay the agent (assuming that the designer is then constrained to choose $R(\underline{\theta}) = 0$). We simply set \underline{u} to 0 in the proof of Theorem 1, which means that $\underline{\theta}^* = \underline{\theta}$, for any ω . Then, case 1 above becomes case 2.

Let us now apply this observation to the application from Section 4.3. We have $V_\omega(\theta) = \theta(1 - \frac{3}{2}\omega)$, $V_\omega^*(\theta) = -\frac{1}{2}\theta$. To simplify notation, let $\beta_\omega := -(1 - \frac{3}{2}\omega)$. 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 $\mathcal{W}(\theta) \leq \frac{\bar{\theta} - \theta}{\bar{\theta} - \underline{\theta}}\mathcal{W}(0)$, as in the observation we made above, we have to check that, for all $\theta \in [0, 1/4]$,

$$\int_{\theta}^{\frac{1}{4}} [-\beta_\omega \tau + \alpha B(\tau)] dF(\tau) \leq -(1 - 4\theta)\beta_\omega \mathbb{E}[\theta].$$

Rewriting, we obtain,

$$\beta_\omega \frac{\int_{\theta}^{\frac{1}{4}} \tau dF(\tau) - (1 - 4\theta)\mathbb{E}[\theta]}{\theta(1 - F(\theta))} \geq \alpha.$$

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

$$W(\theta) \equiv -\beta_\omega \theta f(\theta) + \alpha B(\theta) f(\theta)$$

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 lies 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 Calculations for Subsection 4.4

In this application, we have negative types: $\theta \equiv -k$. Moreover, $V_\omega(\theta) = \omega$, $V_\omega^*(\theta) = \omega$, $\alpha = 1$, and $\alpha^* \leq 1$.

By Proposition 1, we have the thresholds

$$\omega + \underline{\theta}_\omega^* + \frac{F(\underline{\theta}_\omega^*)}{f(\underline{\theta}_\omega^*)} = 0 \quad \text{and} \quad \omega + \bar{\theta}_\omega^* - \frac{1 - F(\bar{\theta}_\omega^*)}{f(\bar{\theta}_\omega^*)} = 0,$$

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

$$\Phi'_\omega(\theta) + \alpha^* \Psi'_\omega(\theta) = \begin{cases} 0 & \theta < \underline{\theta}_\omega^*, \\ -[\omega + \alpha^* \theta] f(\theta) & \theta \in (\underline{\theta}_\omega^*, \bar{\theta}_\omega^*), \\ 0 & \theta > \bar{\theta}_\omega^*. \end{cases}$$

Rewriting the first-order condition from Appendix A.8 yields that a necessary condition for optimality is

$$\theta^* = \frac{\mathbb{E}[\omega | \omega \in [\underline{\omega}_{\theta^*}, \bar{\omega}_{\theta^*}]]}{\alpha^*},$$

with $\theta^* = \bar{\theta}$ if the right-hand side expression is above $\bar{\theta}$, and $\theta^* = \underline{\theta}$ if the right-hand side expression is below $\underline{\theta}$, where the bounds in the condition $\omega \in [\underline{\omega}_{\theta^*}, \bar{\omega}_{\theta^*}]$ are defined implicitly by $\theta^* \in [\underline{\theta}_\omega^*, \bar{\theta}_\omega^*]$.

B.5 Calculations for Subsection 4.5

The conclusions follow directly from Theorem 1.