

Contractibility Design

Roberto Corrao*
Stanford

Joel P. Flynn†
Yale

Karthik A. Sastry‡
Princeton

September 11, 2026

Abstract

We introduce a model of contracting in which the principal, in addition to writing contracts, must engage in contractibility design: creating an evidence structure that allows them to prove when the agent has breached the contract. Designing an evidence structure entails both (i) front-end costs borne *ex ante*, such as those of drafting contracts, and (ii) back-end costs borne *ex post*, such as those of generating evidence. We find that, under even small front-end costs, optimal contracts are coarse, specifying finitely many contingencies out of a continuum of possibilities. In contrast, under even large back-end costs, optimal contracts are complete. Applied to the design of procurement contracts, our results rationalize: (i) the discreteness of contracts, (ii) the presence of similarly vague contracts in low-stakes and high-stakes settings, and (iii) the abrupt adjustment of contracts to changes in the economic environment.

*Stanford University Department of Economics, 579 Jane Stanford Way, Stanford, CA, 94305. Email: rcorrao@stanford.edu.

†Yale University Department of Economics, 30 Hillhouse Avenue, New Haven, CT, 06511. Email: joel.flynn@yale.edu.

‡Princeton University Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ, 08544. Email: ksastry@princeton.edu.

This paper was formerly circulated under the title: “Optimally Coarse Contracts.” We are grateful to Dirk Bergemann, Alessandro Bonatti, Oğuzhan Çelebi, Simone Cerreia-Vioglio, Tommaso Denti, Laura Doval, Piotr Dworzak (Discussant), Drew Fudenberg, Robert Gibbons, Marina Halac, Navin Kartik, Giacomo Lanzani, Elliot Lipnowski, Meg Meyer, Stephen Morris, Ellen Muir, Marco Ottaviani, Jacopo Perego, Andrea Prat, Ilya Segal, Ludvig Sinander, Francesco Squintani, Philipp Strack, Roland Strausz, Alexander Teytelboym, Jean Tirole, Rakesh Vohra, Alexander Wolitzky, Kai Hao Yang, and seminar participants at Columbia, Chicago, Collegio Carlo Alberto, PSE, Sciences Po, Oxford, Warwick, Nottingham, Yale, Arizona State, UBC, the North American Winter Meeting of the Econometric Society, SITE, Washington University in St. Louis, the NYU IO Day, UC Berkeley, UT Austin, Universitat Autònoma de Barcelona, and the University of Southern California for helpful comments. First posted version: November 2, 2023.

1 Introduction

Contracts are not enforceable by *fiat*, requiring careful drafting and the generation of evidence that can be used to prove or disprove the legality of a given action. Accordingly, the literature on contract law emphasizes two classes of contracting costs. The first class comprises *front-end* costs of “foreseeing possible future contingencies, determining the efficient obligations that should be enforced in each contingency, [...] and drafting the contract language that communicates their intent to courts” (Scott and Triantis, 2005). The second class comprises *back-end* costs of “observing and proving the existence (or nonexistence) of any relevant fact after uncertainty has been resolved” (Scott and Triantis, 2005). The fact that lawyers spend up to 60% of their time drafting and reviewing documents (Thomson Reuters, 2024) underscores the importance of these costs for contract design.

In practice, contracting parties weigh such costs of complex contracts against their incentive gains. The contracts that result from this balancing act are often vague: even high-stakes contracts can be perplexingly imprecise (see *e.g.*, Asanuma, 1989; Bewley, 1999).

The literature in microeconomic theory has approached this issue in two ways. On the one hand, the textbook mechanism design approach to optimal contracting (see *e.g.*, Bolton and Dewatripont, 2004; Laffont and Martimort, 2009) abstracts away from costs of writing contracts and allows all contracts to be as precise as possible. On the other hand, a large and influential literature on incomplete contracts studies the optimal design of contracts given the constraint that they are vague, often with the justification that complete contracts are prohibitively costly (see *e.g.*, Anderlini and Felli, 1994; Battigalli and Maggi, 2002).

In this paper, we propose a new framework for studying contractibility design in principal-agent settings in which information is incomplete and evidence is costly. Our goal is to understand how the nature of contracting costs affects the properties of optimal contracts. When are vaguely specified contracts optimal? What economic principles guide their design?

Formally, we consider a setting in which a principal contracts with an agent who has private information about a continuous type and may take a continuous action. The principal designs both a menu of contracts, which specify promised actions and payments, and an *evidence structure* that maps the action the agent takes to a set of evidence. The principal can use this evidence to prove whether the agent deviated from the promised action, imposing a fine upon the agent when they can prove such a breach of contract.¹ We assume only that the feasible evidence structures for the principal satisfy natural continuity and monotonicity properties and always allow the principal to prove that they excluded the agent.

Evidence structures are costly. We consider two classes of costs: front-end and back-end costs. Both share the same basic form and capture the idea that the principal bears a cost whenever the evidence structure allows the distinction of one action from another. Front-

¹We further prove that such *breach mechanisms* are without loss of optimality for the principal.

end costs are borne *ex ante* and depend on the entire evidence structure, independently of the actions taken and the eventual generation and use of evidence. By contrast, back-end costs are borne *ex post* and the principal internalizes their expected value when designing the evidence structure. Thus, back-end costs depend on the likelihood that different sets of evidence are generated. Importantly, evidence structures that perfectly distinguish all pairs of actions have a finite and potentially arbitrarily small cost. The joint design of the evidence structure and menu involves a trade-off between the benefits of sharper incentives and the costs of creating and employing an evidence structure that guarantees these incentives.

In our main results, we establish that front-end and back-end costs generate qualitatively distinct optimal contracts. When the principal faces only front-end costs, every optimal evidence structure is *coarse*: it allows for the distinction of *finitely many* equivalence classes of actions. This, in turn, implies that there exists a finite optimal menu of contracts. Thus, front-end costs always lead to incomplete contracts: there are finitely many contingencies, but within each contingency many actions are legally permissible.² By contrast, when the principal faces only back-end costs, all optimal evidence structures and menus have the cardinality of the continuum. That is, despite transaction costs, it is optimal for the principal to specify a *complete* contract. Taken together, our results show that large transaction costs are neither necessary nor sufficient for optimal contracts to be coarse. Instead, coarseness hinges on the *nature* of these costs and, more specifically, their timing.

The intuition behind the optimality of coarseness under front-end costs is that, when contracts are sufficiently precise, the marginal benefits of more precise contracting are an order of magnitude smaller than the corresponding marginal costs. More specifically, the principal allocates to each type of the agent the action that maximizes their virtual surplus among those actions that are distinguishable given the evidence structure. When there are already distinguishable actions that are close enough to the principal's second best, we show that adding another one even closer can only yield modest gains of *third-order*. By contrast, the front-end cost of adding that action in the middle of an interval of indistinguishable actions leads to a *second-order* loss, as the number of costly comparisons that must be made is quadratic in the number of actions being compared.

The contrasting result that back-end costs do not yield coarseness is subtle: for a fixed distribution of the agents' actions, the back-end cost is mathematically equivalent to a front-end cost. The difference arises because front-end costs are paid before the principal interacts with the agent, and thus do not depend on the agent's actions. This is akin to a fixed cost.

²We adopt the definition of contractual incompleteness used, among others, by [Spier \(1992\)](#) and [Scott and Triantis \(2005\)](#). That is, while the contract specifies an obligation for each outcome, it *does not* specify a different obligation for each outcome, even when it is efficient to do so.

Hence, under front-end costs, the principal may economize on specifying a detailed evidence structure by rendering only a few equivalence classes of actions contractible. By contrast, back-end costs are paid after the agent acts. Thus, there is a variable cost of generating evidence that depends on the action taken by the agent. Moreover, the principal can always design incentives to avoid incurring high variable costs of evidence. Because of this, there is no gain in restricting what evidence can be generated up-front, and it is optimal to leave open the possibility of perfect evidence generation.

When both costs are present, contracts are incomplete, but front-end and back-end costs interact to determine the optimal structure of evidence and the menu. Our approach allows us to derive an upper bound on the optimal number of contractible actions. This bound is invariant to the slope of the virtual surplus in the agent’s action, an object that captures the (first-order) sensitivity of payoffs. Thus, our analysis not only predicts the coarseness of contracts, but also that high- and low-stakes contracts could be equally coarse. This rationalizes the perhaps puzzling phenomenon that even billion-dollar contracts are vague. Moreover, the optimality of coarse contracts allows us to reduce the original infinite-dimensional problem to a finite-dimensional one. With this, under standard regularity conditions, simple first-order conditions describe optimal contracts.

In an application, we use the model and our theoretical results to study contractibility design in optimal procurement. Agents are contractors who differ in their privately known productivity and can exert effort to produce output for a principal. The principal is a purchaser who uses contracts to induce effort, but also bears front-end and back-end costs.

The purchaser optimally offers discrete levels of payment corresponding to discrete tiers of effort. Such arrangements are common in practice, even in settings where piece rates or more complicated nonlinear contracts are feasible (see *e.g.*, [Bewley, 1999](#)). Applying our necessary optimality conditions, we solve analytically for the optimal menu. Front-end costs affect only the number of elements of an optimal menu, while back-end costs only distort the payments received for each menu item. Turning to comparative statics, we show that: (i) the vagueness of optimal contracts is invariant to type-neutral changes in productivity and (ii) contracts are rigid in response to small changes to type-augmenting productivity, while large changes can induce a complete restructuring. Thus, our results rationalize the possibility that procurement contracts are vague in a way that is insensitive to economic value and adjust discontinuously to changes in the economic environment.

Related Literature. We build our theory on the formalization of evidence developed by [Green and Laffont \(1986\)](#) and [Hart, Kremer, and Perry \(2017\)](#). Two novel features of our evidentiary model are: (i) evidence is endogenously generated by the agent’s action, rather than exogenously by their type (see [Ball and Kattwinkel, 2019](#); [Krähmer and Strausz, 2025](#),

for recent contributions on mechanism design with evidence about exogenous types) and (ii) the evidence structure is optimally designed. We also formalize two general classes of costs of evidence structures (front-end and back-end) and show that their analysis is tractable, despite the richness of the space of evidence structures over which they are defined.

Endogenous evidence about the agent’s actions leads to endogenous contractibility of the agent’s actions. This links our analysis to the literature studying imperfect contractibility and the dichotomy between the letter of the contract and the spirit of the contract emphasized by Williamson (1975), Hart and Moore (2008), and Squintani (2019). We formalize this distinction via evidence structures, which nest as a particular and special case the type of imperfect contractibility (free disposal) studied by Grubb (2009) and Corrao, Flynn, and Sastry (2023). Differently from these papers, we study the *optimal* extent of contractibility. In addition, we derive the properties of contractibility from an evidentiary foundation.

Our work fits into a larger literature that provides foundations for incomplete contracts based on transaction costs (Simon, 1951; Coase, 1960). With respect to the Tirole (1999) classification of approaches, our analysis shows how costs of writing contracts (*i.e.*, front-end costs) lead to incompleteness while costs of enforcing contracts (*i.e.*, back-end costs) do not.³

Existing work on contracting with costly writing and/or enforcement studies problems with computational constraints on what events can be described (Anderlini and Felli, 1994, 1999; Al-Najjar, Anderlini, and Felli, 2006) and costs of writing that scale linearly with the number of clauses in a contract (Dye, 1985; Bajari and Tadelis, 2001; Battigalli and Maggi, 2002). Relative to this work, our analysis is different in two ways. First, we study an infinite (rather than finite) state and action space in which continuum contracts are feasible at finite cost. Thus, our analysis can directly speak to whether it is desirable to implement infinite contracts, like piece rates for workers (Holmström, 1979; Holmström and Milgrom, 1987) and nonlinear pricing with smooth quantity discounts (Wilson, 1993). This allows us to overcome the Tirole (1999) critique of the incomplete contracts literature that implicitly or explicitly assumes that such contracts are infinitely costly to write. Second, as mentioned above, our analysis proposes a new framework for justifying costly contractibility based on costly evidence. This microfoundation for the form of contractibility costs is especially important in light of our result that some, but not all, costs lead to optimal coarseness.

Finally, our methods are related to a literature on the finiteness of optimal contracts and menus. In Section 5.6, we elaborate on this connection in detail.

³There are other important perspectives on contractual incompleteness, including *ex post* renegotiation (Segal, 1999; Hart and Moore, 1999; Che and Hausch, 1999) and *ex ante* signaling (Spier, 1992). Our goal is not to subsume all such perspectives but instead to study how contracting costs affect incompleteness.

2 Model

2.1 The Agent and the Principal

The core of our model is a standard principal-agent setting. The agent's type $\theta \in \Theta = [0, 1]$ is drawn from a distribution F with strictly positive continuous density f . The agent privately knows their type and can take an action x in the interval $X = [0, \bar{x}] \subset \mathbb{R}$. The agent has a twice continuously differentiable utility function $u : X \times \Theta \rightarrow \mathbb{R}$. We assume that higher types value higher actions more and that preferences are monotone increasing in the action: (i) u is strictly supermodular in (x, θ) and (ii) for each $\theta \in \Theta$, $u(\cdot, \theta)$ is strictly monotone increasing over X . The case with strictly decreasing preferences over X is analogous. All agent types value the zero action the same as their outside option payoff, which we normalize to zero, or $u(0, \theta) = 0$ for all $\theta \in \Theta$. The agent has quasilinear preferences over actions and money $\tau \in \mathbb{R}$, so their transfer-inclusive payoff is $u(x, \theta) - \tau$.

The principal's payoff derives from three sources. The first is a payoff that depends on the agent's action and type given by $\pi : X \times \Theta \rightarrow \mathbb{R}$, a twice continuously differentiable function such that $\pi(0, \theta) = 0$ for all $\theta \in \Theta$. The second is the monetary payment $\tau \in \mathbb{R}$ from the agent. The third is the cost of contractibility, which we will introduce in due course. We define the virtual surplus function $J : X \times \Theta \rightarrow \mathbb{R}$ as:

$$J(x, \theta) = \pi(x, \theta) + u(x, \theta) - \frac{1 - F(\theta)}{f(\theta)} u_\theta(x, \theta). \quad (1)$$

This is the total surplus generated when agent θ takes action x , net of payments that are necessary to ensure local incentive compatibility. We assume that J is twice continuously differentiable, strictly supermodular in (x, θ) with $J_{x\theta} > 0$, and strictly quasiconcave in x .⁴

2.2 Evidence

In standard mechanism design, actions, or noisy signals thereof (*e.g.*, [Holmström, 1977](#)), are assumed to be perfectly contractible. Instead, we start from the premise that actions are contractible only to the extent that the principal can prove, with hard and legally verifiable evidence, that the agent breached the contract if they did indeed violate its terms. We formalize this by building a model of evidence regarding the agent's actions.

⁴Sufficient primitive conditions are that u , π , and F are three times continuously differentiable with $u_{x\theta}, u_{xx\theta}, \pi_{x\theta} > 0$, $u_{xx}, u_{x\theta\theta}, \pi_{xx} < 0$, and that F has the increasing hazard rate property. We assume strict supermodularity of J for clarity of exposition, as this ensures there is no mechanical incentive for the principal to bunch agents at finitely many allocations. As we later discuss, our analysis extends to settings with a strictly concave (but not supermodular) virtual surplus by leveraging the results of [Toikka \(2011\)](#).

Our approach is motivated by existing frameworks for modeling evidence about *exogenous* types that we adapt to modeling *endogenous* actions (Green and Laffont, 1986; Bull and Watson, 2007; Hart, Kremer, and Perry, 2017). Formally, we posit that the space of evidence is an interval $\Omega \subseteq \mathbb{R}$. An evidence structure is a correspondence $\mathcal{E} : X \rightrightarrows \Omega$ that returns a set of evidence $\mathcal{E}(x)$ when the agent takes action $x \in X$. The principal can exhibit any subset of $\mathcal{E}(x)$ to an external arbitrator. We assume that the agent is “innocent until proven guilty.” That is, the agent’s action is found consistent with the contract unless the principal can prove that the agent did not act in accordance with the contract. Formally, the principal can prove that the agent did not do some promised action y if there exists a piece of evidence generated by their action $e \in \mathcal{E}(x)$ that could not have been generated by doing what they promised to do, *i.e.*, $e \notin \mathcal{E}(y)$. The principal can *distinguish* action x from action y if and only if $\mathcal{E}(y)^c \cap \mathcal{E}(x) \neq \emptyset$.

We restrict attention to evidence structures with the following *regularity* properties:

Definition 1 (Regularity). *An evidence structure $\mathcal{E} : X \rightrightarrows \Omega$ is regular if it satisfies:*

1. (Continuity) Fix $x, y \in X$: (i) For every sequence $x_n \rightarrow x$ with $\mathcal{E}(x_n) \subseteq \mathcal{E}(y)$ for all $n \in \mathbb{N}$, it holds that $\mathcal{E}(x) \subseteq \mathcal{E}(y)$, (ii) Conversely, if $\mathcal{E}(x) \subseteq \mathcal{E}(y)$, then for every sequence $y_n \rightarrow y$, there exists a subsequence $y_{n_k} \rightarrow y$ and a sequence $x_k \rightarrow x$ such that $\mathcal{E}(x_k) \subseteq \mathcal{E}(y_{n_k})$ for all $k \in \mathbb{N}$.
2. (Definitive evidence of exclusion) For all $x \in X$, if $\mathcal{E}(x) \subseteq \mathcal{E}(0)$, then $x = 0$.
3. (Monotonicity) $\mathcal{E}$ is monotone in the strong set order.


Continuity is a technical assumption, necessary to derive the existence of optimal contracts given an evidence structure. Definitive evidence of exclusion asserts that the principal can always prove that the agent was excluded if that is indeed the case. Monotonicity rules out situations in which high actions generate lower evidence than low actions and corresponds to a Boolean version of the monotone likelihood-ratio property.⁵

Figure 1 illustrates evidence structures and the content of regularity. Examples A, B, and C each correspond to different regular evidence structures. Example D corresponds to an evidence structure that fails monotonicity. In particular, taking $X = \Omega = [0, 1]$, we can observe that $x' = 1$ generates the piece of evidence $e = 0$ and $x = 1/2$ generates the piece of evidence $e' = 1$. However, it is *not* the case that the higher action generates the higher piece of evidence: $e' = 1 \notin \mathcal{E}(x') = [0, 1/4]$.

An assumption implicit in our model is that the principal’s payoff is not evidence *to an external arbitrator* in and of itself. This is motivated by practice in the contracting settings

⁵To see this, observe that monotonicity of $\mathcal{E}$ in the strong set order is equivalent to the following property. For all $x, x' \in X$ and $e, e' \in \Omega$, if $e' \geq e$ and $x' \geq x$, then $\mathbb{I}[e' \in \mathcal{E}(x')]\mathbb{I}[e \in \mathcal{E}(x)] \geq \mathbb{I}[e' \in \mathcal{E}(x)]\mathbb{I}[e \in \mathcal{E}(x')]$.

Figure 1: Examples of Evidence Structures



Note: This figure shows examples of evidence structures. Examples A, B, and C are regular evidence structures (see Definition 1). Example D is an irregular evidence structure because it fails monotonicity.

in which we are interested. Typically, one cannot establish breach of contract by merely claiming that a loss was suffered; one must provide concrete evidence of the underlying conduct and prove its inconsistency with a codified contract (see, *e.g.*, [New York City Bar, 2014](#)). In such settings, [Maskin and Tirole's \(1999\)](#) suggested approach to specify contracts directly over payoffs would not *by itself* overcome the necessity of evidence generation.

2.3 Timeline


We now describe the full timeline of events, as summarized in Figure 2. First, the principal chooses a regular evidence structure $\mathcal{E}$. Second, the principal commits to a menu $\mathcal{M} = \{(y_i, T(y_i))\}_{i \in I} \subseteq X \times \mathbb{R}$ of promised actions $y_i \in X$ for some index set I and corresponding payments $T : X \rightarrow \mathbb{R}$. Third, the agent's type $\theta \in \Theta$ is drawn from the distribution F , the agent observes it, and then either selects an element $(y, T(y)) \in \mathcal{M}$ or takes their outside option.⁶ Fourth, if the agent participates in the mechanism, they choose an action $x \in X$ and this generates evidence $\mathcal{E}(x)$. Fifth, if x can be distinguished from y , then the agent is punished by an arbitrarily large fine, and otherwise pays $T(y)$ to the principal.

2.4 Costs of Evidence: Front-End *vs.* Back-End

We finally describe the costs of evidence. Legal scholarship ([Scott and Triantis, 2005](#)) identifies two distinct stages of a contractual relationship in which the principal bears costs: before the agent takes an action and after the agent takes an action. We formalize this idea by considering two separate components for the cost of evidence structures. First, there are front-end costs borne *ex ante* when the principal chooses an evidence structure (Stage 1 in Figure 2). These depend on the evidence structure but not the subsequently generated evidence. They include, for example, the costs of “foreseeing possible future contingencies,

⁶As is standard, we assume that the agent, if indifferent, takes the principal's preferred action.

Figure 2: Model Timeline



determining the efficient obligations that should be enforced in each contingency, [...] and drafting the contract language that communicates their intent to courts” (Scott and Triantis, 2005). Second, there are back-end costs borne *ex post* when the agent’s action generates evidence (Stage 4 of Figure 2). These depend on producing or interpreting evidence specific to the agent’s action: that is, “observing and proving the existence (or nonexistence) of any relevant fact after uncertainty has been resolved” (Scott and Triantis, 2005).

To model these costs in a parallel way, we first introduce a basic notion of the cost of distinguishing any two actions under an evidence structure $\mathcal{E}$. Specifically, we assume that whenever $z \in X$ can be distinguished from $x \in X$ (i.e., $\mathcal{E}(x)^c \cap \mathcal{E}(z) \neq \emptyset$), the evidence structure generates a pairwise distinction cost $g(x, z)$, where $g : X \times X \rightarrow \mathbb{R}_{++}$ is a twice continuously differentiable function and we let $\mathcal{G}$ denote the set of such functions.⁷ Thus, given an evidence structure $\mathcal{E}$, the total cost of distinguishing a hypothetical action x from those actions from which it can be separated is

$$\gamma_g(\mathcal{E}, x) = \int_{\{z \in X : \mathcal{E}(x)^c \cap \mathcal{E}(z) \neq \emptyset\}} g(x, z) \, dz. \quad (2)$$

The function g captures how hard it is to distinguish two actions. This may depend on the identities of those actions, how far apart those actions are, and/or whether those actions are distinguished *before* the agent acts or *after* the agent acts. The simplest case corresponds to paying a uniform cost of $\kappa > 0$ every time z is distinguished from x , or $g(x, z) = \kappa$.

Under *front-end costs*, the principal must pay to distinguish all actions that the agent may hypothetically take, regardless of what the agent ultimately does. The front-end cost is then:

$$\Gamma_f(\mathcal{E}) = \int_X \gamma_{g_f}(\mathcal{E}, x) \, dx, \quad (3)$$

where $g_f \in \mathcal{G}$ measures the cost of distinguishing pairs of actions *ex ante*. Intuitively, front-

⁷Observe that this implies that $g(x, x) > 0$. This has no practical implications. As the diagonal is measure zero, nothing in our analysis changes if we consider $\hat{g}(x, y) = g(x, y)\mathbb{I}[x \neq y]$.

end costs act as *fixed costs* of establishing the evidence structure: once the evidence structure has been designed, it can be subsequently used and reused without additional front-end cost.

Under *back-end costs*, the principal pays only to generate the evidence corresponding to the agent's ultimate action, and not any other hypothetical actions. That is, every time the agent takes action x , the principal incurs the cost $\gamma_{g_b}(\mathcal{E}, x)$, where $g_b \in \mathcal{G}$ measures the cost of distinguishing pairs of actions *ex post*. If the principal designs evidence structure $\mathcal{E}$ and the agent eventually takes action x , then their total *realized* cost is $\Gamma_f(\mathcal{E}) + \gamma_{g_b}(\mathcal{E}, x)$. When the principal chooses the evidence structure at the first stage, given a belief $Q \in \Delta(X)$ about the action that the agent will take,⁸ their *expected* back-end costs are:

$$\Gamma_b(\mathcal{E}, Q) = \int_X \gamma_{g_b}(\mathcal{E}, x) dQ(x). \quad (4)$$

Intuitively, back-end costs act as *variable costs* of generating evidence.⁹ They encode the costs of generating specific pieces of evidence, rather than the costs of setting up a structure that allows for evidence generation.

The key economic difference between front-end costs (Equation 3) and expected back-end costs (Equation 4) is that the uniform weighting in the front-end cost does not represent the (perceived) likelihood of actions, as does Q in the back-end cost, but rather the *fixed* relative contribution of each action to the total front-end costs. In fact, when Q is a fixed distribution with a density dQ , the back-end cost is *payoff* equivalent to a front-end cost with $g_f = g_b \times dQ$. Yet, because Q is endogenous to the choice of the evidence structure and menu, the analysis of the two costs will be significantly different.

2.5 The Principal's Problem

We next describe the principal's problem of contractibility design via evidence. We define $\phi : \Theta \rightarrow X$ and $t : \Theta \rightarrow \mathbb{R}$ as mappings from types to actions and transfers, respectively. We let $\mathcal{I}(\mathcal{M}, \mathcal{E})$ denote the pairs of measurable (ϕ, t) that arise in some pure-strategy subgame-perfect Nash equilibrium of the game between the principal and agent given $(\mathcal{M}, \mathcal{E})$. The principal jointly designs the evidence structure $\mathcal{E}$ and the menu $\mathcal{M}$ to maximize its expected payoffs net of front-end and back-end costs of evidence. We say that an evidence structure $\mathcal{E}$ and a menu $\mathcal{M}$ are *minimal* if there is no proper subset $\mathcal{M}' \subset \mathcal{M}$ such that the principal attains the same maximum payoff under both $(\mathcal{M}, \mathcal{E})$ and $(\mathcal{M}', \mathcal{E})$. In what follows, given

⁸More specifically, given any equilibrium, letting $\phi : \Theta \rightarrow X$ denote the corresponding equilibrium mapping from types to final actions, we have that $Q = F \circ \phi^{-1}$.

⁹This formulation assumes that the principal must pay back-end costs to prove that it excluded agents. If this is not the case, our analysis differs only in determining a threshold type that participates in the mechanism, and our main results hold as stated (see Remark 3 in Appendix E).

an evidence structure, we always restrict to menus that satisfy minimality. This refinement is without loss of optimality and removes artificial sources of multiplicity of solutions.

For convenience, we break the principal's problem into two parts. First, we maximize over the menu and the allocation to define the value of an evidence structure, including back-end costs but not front-end costs:

$$\mathcal{R}(\mathcal{E}) = \sup_{\mathcal{M}, (\phi, t) \in \mathcal{I}(\mathcal{M}, \mathcal{E})} \int_{\Theta} (\pi(\phi(\theta), \theta) + t(\theta)) dF(\theta) - \Gamma_b(\mathcal{E}, Q_\phi), \quad (5)$$

where $Q_\phi \in \Delta(X)$ is given by $Q_\phi(x) = \mathbb{P}_F[\phi(\theta) \leq x]$ and is the equilibrium distribution over actions induced by ϕ . We observe that the expected back-end cost, for a given $\mathcal{E}$, depends on the final outcomes implemented by the principal. An optimal evidence structure maximizes the total value over the space of all regular evidence structures:

$$\sup_{\mathcal{E}} \mathcal{R}(\mathcal{E}) - \Gamma_f(\mathcal{E}). \quad (6)$$

2.6 Extensions of the Framework

Beyond Binary Enforcement. In our baseline model, evidence is used only to determine whether or not the agent breached the contract. More generally, the principal could condition payments flexibly on the promised action and the evidence that they submit to the arbitrator. We describe this via a flexible mechanism $P : X \times 2^\Omega \rightarrow \mathbb{R}$, with $P(y, E)$ denoting the payment made by the agent to the principal when the agent promises action $y \in X$ and the principal exhibits $E \subseteq \Omega$ to the arbitrator. We do not introduce any new cost of exhibiting evidence conditional on its (costly) generation. Of course, the principal can only submit evidence that actually exists, that is, any subset $E \subseteq \mathcal{E}(x)$ given the agent's final action $x \in X$. Any such flexible mechanism respects our premise that final actions can be contracted upon only insofar as they generate different evidence.

Seen this way, our baseline model restricts to a subclass of mechanisms with binary enforcement. Specifically, we say P is a *breach mechanism* given $\mathcal{E}$ if there exists a tariff $T : X \rightarrow \mathbb{R}$ and a punishment $M \in \mathbb{R}$ such that:

$$P(y, E) = \begin{cases} T(y), & E \cap \mathcal{E}(y)^c = \emptyset, \\ M, & \text{otherwise.} \end{cases} \quad (7)$$

While this may seem restrictive, we show in Appendix D that considering only breach mechanisms is without loss of optimality for the principal. That is, for any evidence structure $\mathcal{E}$ and any equilibrium that the principal can implement via a flexible mechanism given that

evidence structure, there exists a breach mechanism that yields an identical expected payoff for the principal (Proposition 4). This holds even conditional on any maximum punishment $\bar{M} > 0$ representing limited liability. Intuitively, akin to the revelation principle, breach mechanisms can be designed to replicate the incentives provided by flexible mechanisms. A key observation underlying this is that additional evidence only increases the transfer paid by the agent, since the principal chooses which evidence to exhibit. Under the replicating breach mechanism, agents promise to take actions they ultimately perform, and they pay the transfer that would have occurred in equilibrium under the flexible mechanism if they took that final action.

Probabilistic Arbitration. In our baseline model, evidence and the determination of breach are deterministic. However, it is natural to imagine that arbitration of breaches is probabilistic. There are two possible types of arbitration errors that may result: type I errors, where the agent satisfied the contract but is nevertheless found to be in violation, and type II errors, where the agent violated the contract but is not found to be in violation. In Appendix E we extend our model to allow for exogenously probabilistic arbitration and we show that our main results generalize to such a setting (see Proposition 5).

More General Costs. Parameterizing front-end and back-end costs by $g_f, g_b \in \mathcal{G}$ imposes a specific structure. In Appendix F, we consider a more general, non-parametric family of front-end and back-end costs and provide general sufficient conditions under which our main results hold (see Propositions 7 and 8).

3 Main Results

We now study optimal evidence structures and menus. Here, we present the main results and the economic intuition behind them; we defer the main formal arguments to Section 5.


3.1 Front-End Costs Imply Coarse Contracts

We first define *coarseness*. For any non-empty set A , we denote its cardinality by $|A| \in \mathbb{N} \cup \{\infty\}$. We let $\mathcal{E}(X) = \{\mathcal{E}(x)\}_{x \in X} \subseteq 2^\Omega$ denote the image of the evidence correspondence $\mathcal{E}$, and we recall that $\mathcal{M} = \{(y_i, T(y_i))\}_{i \in I} \subseteq X \times \mathbb{R}$ denotes a menu of actions and payments.

Definition 2 (Coarseness). *An evidence structure $\mathcal{E}$ is coarse if $|\mathcal{E}(X)| \in \mathbb{N}$. A menu $\mathcal{M}$ is coarse if $|\mathcal{M}| \in \mathbb{N}$.*

A coarse evidence structure partitions the actions into finitely many classes, within which actions cannot be distinguished. In Figure 3, we illustrate a coarse evidence structure and

Figure 3: Coarse and Non-Coarse Evidence Structures



Note: Each panel shows an example evidence correspondence, with an action space $X = [0, 1]$ and an evidence space $\Omega = [0, 1]$. Panel A has the cardinality of the continuum and Panel B has the cardinality of the set of natural numbers, both of which are infinite. Panel C has a finite cardinality, and is therefore coarse by Definition 2.

two non-coarse evidence structures. A coarse menu specifies only a finite number of contingencies: promised actions and corresponding payments if the agent cannot be shown to have deviated from the promise. Therefore, the coarseness of the evidence structure has a natural interpretation as representing the *vagueness* of a contract: many possible actions generate the same evidence, and so many actions are legally consistent with the contract.

Toward stating our first main result, we define the maximum concavity of the virtual surplus function $\bar{J}_{xx} = \max_{x \in X, \theta \in \Theta} |J_{xx}(x, \theta)|$, the minimum supermodularity of the virtual surplus function $\underline{J}_{x\theta} = \min_{x \in X, \theta \in \Theta} J_{x\theta}(x, \theta)$, the maximum density of types $\bar{f} = \max_{\theta \in \Theta} f(\theta)$, and the minimal front-end cost of distinguishing any two actions $\underline{g}_f = \min_{x, z \in X} g_f(x, z)$. Under our maintained assumptions, we have that $\underline{J}_{x\theta}, \underline{g}_f > 0$. We first consider the case with only front-end costs and no back-end costs, *i.e.*, $g_f \in \mathcal{G}$ and $g_b \equiv 0$.

Theorem 1 (Front-end Costs Imply Coarseness). *If the principal only faces front-end costs, then an optimal evidence structure exists and every optimal evidence structure $\mathcal{E}^*$ is coarse:*

$$|\mathcal{E}^*(X)| \leq 6 + \left\lceil 3 \frac{\bar{x} \bar{J}_{xx} \bar{f}}{\underline{g}_f \underline{J}_{x\theta}} \right\rceil. \quad (8)$$

Moreover, every optimal menu $\mathcal{M}^$ is coarse and has the same bound on its cardinality.*

Under front-end costs, contracts are coarse. Underlying these coarse contracts are evidence structures that can only differentiate among a finite number of equivalence classes of actions. This optimal imprecision in the contract structure is reminiscent of the vague clauses that Scott and Triantis (2005) argue are “commonplace in commercial contracts.”

Our bound on the extent of richness in contracting is invariant to the slope of the virtual surplus function in the agent’s action, or the *sensitivity* of payoffs. In this way, the bound

rationalizes why high-stakes settings where billions of dollars are on the line may feature contracts that are comparably vague to settings in which the value is trivial.¹⁰ In Section 5, we provide intuition for the exact form of the bound. In our application, we moreover study an explicit class of problems and find that this bound is tight up to an affine transformation.¹¹

We also highlight two immediate remarks regarding Theorem 1. First, even arbitrarily small front-end costs can induce the qualitative property of coarseness. That is, if Theorem 1 holds given a particular front-end cost described by a cost of distinguishing g_f , then it also holds under the cost κg_f for any arbitrarily small $\kappa \in \mathbb{R}_{++}$. Second, the coarseness result holds despite the fact that the front-end cost of perfectly discerning evidence structures is finite. Thus, the result does *not* rely on the infeasibility of continuum evidence structures: instead, it shows their suboptimality under even arbitrarily small front-end costs.

3.2 Contractual Completeness Under Pure Back-End Costs

We next study the case with only back-end costs ($g_b \in \mathcal{G}$ and $g_f \equiv 0$). To do so, we define the modified virtual surplus, which includes the realized back-end costs for the principal:

$$H(x, \theta) = J(x, \theta) - \int_x^{\bar{x}} g_b(x, z) dz. \quad (9)$$

We say that H is non-trivial if there exists $\theta \in \Theta$ such that $H(\cdot, \theta)$ has a unique non-degenerate maximizer that is, moreover, in the interior of X . This condition rules out cases in which the principal prefers to allocate all agents either to 0 or $\bar{x}$.

Theorem 2 (Back-End Costs Do Not Imply Coarseness). *If the principal faces only back-end costs, then an optimal evidence structure $\mathcal{E}^*$ with cardinality $|\mathcal{E}^*(X)| = \infty$ exists. Moreover, if H is non-trivial, then every optimal evidence structure $\mathcal{E}^*$ and every optimal menu $\mathcal{M}^*$ have cardinality $|\mathcal{E}^*(X)| = |\mathcal{M}^*| = \infty$.*

Under back-end costs, contracts specify a continuum of contingencies that are tailored to different types of agents. This is despite the presence of transaction costs that, at first impression, may appear similar to the front-end costs which we considered in Theorem 1.

3.3 The Difference Between Front-End and Back-End Costs

How can the optimal evidence structure and menu be so different across cost regimes when the two costs are so similar? Indeed for every evidence structure $\mathcal{E}$, when $g_f = g_b$ and

¹⁰Ours is not the only natural definition of the “stakes” of contracting. For example, multiplicatively scaling the payoff by some fixed number $\lambda > 1$ would increase the squared curvature term by λ^2 and the supermodularity by λ , thus scaling the bound by a factor of λ and pushing toward richer contracts.

¹¹That is, the numbers 6 and 3 in the bound may not be sharp, but the remaining term is.

$X = [0, 1]$, there exists an action function ϕ such that the front-end and expected back-end costs impose the same total cost on the principal, *i.e.*, $\Gamma_b(\mathcal{E}, Q_\phi) = \Gamma_f(\mathcal{E})$. However, the presence of each cost creates very different marginal costs of changing the evidence structure. This distinction is crucial.

We provide a detailed economic argument for Theorems 1 and 2 in Section 5. But to illustrate the main intuition, we consider here a simple example based on the [Mussa and Rosen \(1978\)](#) model with $X = \Omega = [0, 1]$, $u(x, \theta) = \theta x$, $\pi(x) = -x^2/2$, and $\theta \sim U[0, 1]$.¹² We take symmetric front-end and back-end costs, $g_f \equiv g_b \equiv \kappa$ for some constant $\kappa \in [0, 1]$. For simplicity, we compare the following three evidence structures:

$$\mathcal{E}_a(x) = \begin{cases} \{0\} & \text{if } x = 0, \\ [0, 1] & \text{otherwise.} \end{cases} \quad \mathcal{E}_b(x) = \begin{cases} \{0\} & \text{if } x = 0, \\ [0, 1/2] & \text{if } x \in (0, 1/2], \\ [0, 1] & \text{if } x \in (1/2, 1]. \end{cases} \quad \mathcal{E}_c(x) = [0, x]. \quad (10)$$

Structure a makes exclusion contractible but otherwise has no power to distinguish the agent's actions. Structure b can distinguish exclusion from low actions $x \in (0, 1/2]$ and low actions from high actions $x \in (1/2, 1]$. Structure c can perfectly distinguish all actions.

We first consider the case with only front-end costs, as in Theorem 1. The front-end cost of each evidence structure can be computed using the definition $\Gamma_f(\mathcal{E}_i) = \int_X \gamma_\kappa(\mathcal{E}_i, x) dx$. These calculations, visualized in the top row of Figure 4, yield $\Gamma_f(\mathcal{E}_a) = 0$, $\Gamma_f(\mathcal{E}_b) = \kappa/4$, and $\Gamma_f(\mathcal{E}_c) = \kappa/2$. These are *fixed costs* of evidence, which can be computed before knowing the specific contract or what the agent will do.


The principal trades off these fixed costs against the benefits of a more precise contract. Under structure a , any included agent can choose the highest action without punishment, namely $x = 1$. Hence, by a revelation-principle logic, the principal can without loss restrict attention to assigning agents to actions in $\overline{D}_a = \{0, 1\}$. Similarly, under structure b , the principal can treat the feasible assignment set as $\overline{D}_b = \{0, 1/2, 1\}$: each agent takes the highest action in the class to which they are assigned. By contrast, under structure c , the principal can assign agents to any action in $\overline{D}_c = [0, 1]$, because any profitable breach by the agent is punished. In each case $i \in \{a, b, c\}$, the principal's value is thus:

$$\mathcal{R}(\mathcal{E}_i) = \int_0^1 \max_{x \in \overline{D}_i} \left\{ (2\theta - 1)x - \frac{1}{2}x^2 \right\} d\theta, \quad (11)$$

where the term in brackets is the virtual surplus $J(x, \theta)$. It is simple to compute $\mathcal{R}(\mathcal{E}_a) =$

¹²Strictly speaking, type $\theta = 0$ has utility $u(x, 0) = 0$ for all $x \in X$, which is increasing but not strictly increasing. As this type has zero probability, our results also apply to these payoffs.

Figure 4: Front-End and Back-End Costs in an Example



Note: This figure illustrates the calculation of front-end and back-end costs in the example described in Section 3.3. The columns correspond to the three evidence structures defined in Equation 10. The rows correspond to the calculation of front-end costs and back-end costs. We set $\kappa = 1$, but give the cost formula for any κ .

$12/192$, $\mathcal{R}(\mathcal{E}_b) = 15/192$, $\mathcal{R}(\mathcal{E}_c) = 16/192$. When $\kappa > 1/24$, evidence structure c yields lower overall value to the principal than the two coarse structures a and b . The gains from contractibility diminish quickly: simply adding the ability to separate low from high actions captures three-quarters of the gains from going from no contractibility to full contractibility. The fact that these gains diminish much more quickly than the front-end costs is the general force underlying the conclusion that contracts are coarse.

We next study the case with only back-end costs, as in Theorem 2. Back-end costs arise after actions are taken and evidence is generated. Thus, when the principal assigns type θ to an action $\phi(\theta)$, they must consider the cost of generating the corresponding evidence. For any type θ and feasible action x , the principal's payoffs are summarized by the modified virtual surplus function H (Equation 9) that augments the standard virtual surplus with back-end costs. Here, the cost term is $-\kappa(1 - x)$: if the agent takes action x , this must be distinguished from the actions in $(x, 1]$ at cost κ per action. Putting this together, the principal's assignment for each type θ solves

$$\phi_i(\theta) \in \arg \max_{x \in \bar{D}_i} \left\{ (2\theta - 1)x - \frac{1}{2}x^2 - \kappa(1 - x) \right\}. \quad (12)$$

Different evidence structures affect the assignment by constraining the domain $\bar{D}_i$. Under $\mathcal{E}_a$, $\bar{D}_a = \{0, 1\}$ and $\phi_a(\theta) = \mathbb{I}[\theta \geq 3/4 - \kappa/2]$. Under $\mathcal{E}_b$, $\bar{D}_b = \{0, 1/2, 1\}$ and $\phi_b(\theta) = 1/2 \mathbb{I}[\theta \in [5/8 - \kappa/2, 7/8 - \kappa/2]] + \mathbb{I}[\theta \in [7/8 - \kappa/2, 1]]$. Finally, under $\mathcal{E}_c$, all actions are contractible and $\phi_c(\theta) = \max \{0, \min \{1, 2\theta - 1 + \kappa\}\}$, the usual truncated linear assignment.

Depending on how the assignment of actions to agents changes, adding contractible actions—*e.g.*, going from a to b and adding the contractible action $1/2$ —may either increase or decrease the total expected back-end costs. However, the principal’s total payoff can never decline when more actions are contractible under the *optimal* assignment. This is because the principal can always secure the same payoff as they did previously by simply implementing the same assignment of agents to actions. Indeed, in this example, by calculating the expected back-end costs paid in each case, using the formula $\Gamma_b(\mathcal{E}_i, Q_{\phi_i}) = \int_{\Theta} \gamma_{\kappa}(\mathcal{E}_i, \phi_i(\theta)) d\theta$, the back-end costs under the optimal mechanism are *equal* in all cases (see the bottom row of Figure 4). As the principal also achieves a greater value from better matching types of the agent to actions, greater contractibility is unambiguously favorable.

To summarize, under front-end costs, the principal pays for an evidence structure before interacting with the agent. Thus, the principal can economize on costs by establishing a coarse structure and then using it to constrain all possible types. By contrast, paying back-end costs to generate evidence for any one type’s action is of no help in constraining the behavior of any other type. Thus, the principal designs the optimal assignment type-by-type, internalizing the cost of evidence generation as if it were an additional variable cost of production. As a result, it is of no benefit to the principal to rule out the possibility of generating evidence for certain contingencies, and the perfectly distinguishing evidence structure is always optimal.

3.4 The General Case: Coarseness and the Optimal Contract

We next show that, in the general case with *both* front-end and back-end costs, the qualitative properties of optimal contracts and evidence are as in Theorem 1.

Corollary 1. *If H is strictly quasi-concave or strictly quasi-convex, then every optimal evidence structure $\mathcal{E}^*$ and menu $\mathcal{M}^*$ are coarse.*

When H is strictly quasi-concave in x for all $\theta \in \Theta$, the result follows from the same arguments as Theorem 1, replacing J with H . When H is strictly quasi-convex, the principal assigns every type to actions 0 and $\bar{x}$, so all optimal menus have at most two elements.

Remark 1 (A General Richness Bound). When H is strictly quasi-concave, adapting the proof of Theorem 1 immediately yields the same bound with $\bar{H}_{xx} = \max_{x \in X, \theta \in \Theta} |H_{xx}(x, \theta)|$ replacing $\bar{J}_{xx}$ but with all other terms unchanged. Moreover, if g_b is symmetric—or distinguishing x from y is exactly as costly as it is to distinguish y from x —and twice continuously differentiable, then:

$$|\mathcal{E}^*(X)| \leq 6 + \left\lceil \frac{3\bar{x} (\bar{J}_{xx} + 3\bar{g}_{b,x} + \bar{x}\bar{g}_{b,xx})^2 \bar{f}}{g_f J_{x\theta}} \right\rceil, \quad (13)$$

where $\bar{g}_{b,x} = \max_{x \in X} |g_{b,x}(x, x)|$ is the maximum slope of g_b on the diagonal and $\bar{g}_{b,xx} = \max_{x,z \in X} |g_{b,xx}(x, z)|$ is the maximum curvature of g_b . An immediate consequence is that, differently from front-end costs, when g_b is a constant function, $g_b(x, z) \equiv \kappa$, the presence of back-end costs does not even affect the general bound on richness. $\triangle$

We finally establish the structure of optimal contracts with front-end and back-end costs.

Proposition 1. *Suppose that H is strictly quasi-concave or strictly quasi-convex, and fix an optimal menu $\mathcal{M}^* = \{(x_k, T(x_k))\}_{k=1}^{K^*}$ where $0 = x_1 < \dots < x_k < \dots < x_{K^*} = \bar{x}$.¹³ Then, for all $k \in \{2, \dots, K^* - 1\}$:*

$$\begin{aligned} & \int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} J_x(x_k, \theta) dF(\theta) + \left(F(\hat{\theta}_{k+1}) - F(\hat{\theta}_k)\right) \left(g_b(x_k, x_k) - \int_{x_k}^{\bar{x}} g_{b,x}(x_k, z) dz\right) \\ &= G_f(x_{k+1}, x_k) - G_f(x_k, x_k) - \int_{x_{k-1}}^{x_k} g_f(x, x_k) dx, \end{aligned} \quad (14)$$

where $G_f(w, x) = \int_0^w g_f(x, z) dz$ and $\hat{\theta}_k$ is given by: (i) $\hat{\theta}_1 = 0$ and $\hat{\theta}_{K^*+1} = 1$, (ii) $\hat{\theta}_k$ is the unique solution to $H(x_k, \hat{\theta}_k) - H(x_{k-1}, \hat{\theta}_k) = 0$ for $k \in \{3, \dots, K^* - 1\}$, and (iii) $\hat{\theta}_2$ and $\hat{\theta}_{K^*}$ are the unique solution to the same equality if one exists and are otherwise 0 and 1, respectively. The corresponding transfer payments, for all $k \in \{1, \dots, K^*\}$, are given by:

$$T(x_k) = \mathbb{I}[k \geq 2] \sum_{j=2}^k \left[u(x_j, \hat{\theta}_j) - u(x_{j-1}, \hat{\theta}_j) \right]. \quad (15)$$

Proof. See Appendix B.3. $\square$

This result provides a second-order difference equation that must be satisfied by the optimal coarse menu. The first term on the left-hand side measures the marginal increase in the principal's payoff and revenue from transfers. The second captures the marginal reduction in back-end costs, which scales with the measure of agents who are allocated to action x_k . The term on the right-hand side measures the marginal change in front-end costs. While in principle this difference equation is quite complicated, we show in Section 4 that it takes a tractable form in a canonical application and enables computation of and comparative statics for the optimal coarse contract.

¹³Here, we augment any optimal (minimal) menu with actions 0 and $\bar{x}$. This is innocuous and allows us to simplify the description of the second-order difference equation characterizing optimal menus.

4 Application to Procurement

We now apply our results to study procurement with endogenous evidence. The agent is a contractor with private information about their productivity and the principal is a purchaser that bears both front-end and back-end transaction costs. We describe in closed form the optimal evidence structure and menu, which features a discrete number of payments corresponding to discrete tiers of effort. Front-end and back-end costs play distinct roles: the former determines the coarseness of the menu, while the latter distorts down equilibrium effort and payment. Moreover, due to its coarseness, the optimal payment structure can be rigid in the face of small changes to productivity, while changing discontinuously in the face of large ones. Finally, we briefly describe additional applications to screening problems.

4.1 Set-up: Procurement with Evidence of Effort

A contractor (the agent) supplies a service to a purchaser (the principal). The contractor's payoff from providing effort level $e \in E = [0, 1]$ is

$$\tilde{u}(e, \vartheta) = -\alpha(1 - \vartheta)e - \beta \frac{e^2}{2}, \quad (16)$$

where $\vartheta \sim \tilde{F} = U[0, 1]$ is the contractor's privately observed productivity, $\alpha > 0$ is a parameter that decreases productivity, especially so for lower types, and $\beta > 0$ is a parameter that scales curvature in effort costs. High effort leads to a higher payoff for the purchaser:

$$\tilde{\pi}(e) = \eta e, \quad (17)$$

where $\eta > 0$ represents the return to effort. Both the purchaser and contractor have quasi-linear utility, and w represents a payment from the former to the latter.¹⁴

The purchaser faces uniform front-end and back-end costs: $g_f(x, y) \equiv \kappa_f$ and $g_b(x, y) \equiv \kappa_b$, where $\kappa_f, \kappa_b > 0$. The former represent the *ex ante* costs of foreseeing contingencies and codifying their difference from one another: for example, describing what it means for the contractor to exert high effort to remodel a bathroom, sell a house, or serve the company as CEO. The latter represents *ex post* costs of generating evidence of what action was taken and perhaps certifying this to an external arbitrator: that is, making sure high effort was indeed exerted, as per the agreed-upon definition. Finally, to simplify the closed-form characterization of optimal contracts, we assume $\beta \leq \eta - \kappa_b \leq 2\alpha$. This ensures that the purchaser desires high effort from the highest productivity types and wishes to exclude

¹⁴With respect to the general theory, the order of actions is flipped. We proceed to apply our results by rewriting the model under the ordering consistent with our general analysis.

the lowest productivity types from undertaking the project.

The purchaser designs a menu of contracts, pairs of effort levels e and payments $w(e)$, as well as a regular evidence structure, $\mathcal{E} : E \rightrightarrows \Omega$ with $\Omega = [0, 1]$, to maximize its expected payoff net of front-end and back-end costs.

4.2 Optimal Contracts Specify Finitely Many Effort Levels

We now apply our results to solve the purchaser's problem. We use the change of variables $x = 1 - e$ ("shirking") and $\theta = 1 - \vartheta$ ("unproductiveness") to bring the model under the assumptions of Section 2 (see Appendix B.4). The purchaser's modified virtual surplus function is

$$H(x, \theta) = (2\alpha\theta + \beta - \eta + \kappa_b)x - \beta\frac{x^2}{2}, \quad (18)$$

which is strictly concave and strictly supermodular. Hence, Corollary 1 implies that any optimal evidence structure and menu are coarse. We can therefore optimize over a number $K \in \mathbb{N}$ of distinct shirking levels $0 = x_1 < \dots < x_K = 1$.

Proposition 1 implies that optimal shirking levels solve a first-order condition that trades off the *ex post* benefits from precise screening against the *ex ante* costs of specifying the contract. Plugging the form of H and of the front-end cost into Equation 14, we have that each element of the menu x_k for $k \in \{2, \dots, K^* - 1\}$ satisfies:

$$\int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} (2\alpha\theta + \beta(1 - x_k) - \eta + \kappa_b) d\theta = \kappa_f(x_{k-1} + x_{k+1} - 2x_k), \quad (19)$$

where $\hat{\theta}_k = \frac{\beta}{4\alpha}(x_k + x_{k-1}) + \frac{\eta - \kappa_b - \beta}{2\alpha}$ are the types between which the purchaser is indifferent to assign x_k or x_{k-1} . This reduces to a second-order difference equation:


$$(x_{k+1} + x_{k-1} - 2x_k) \left[\frac{\beta^2}{16\alpha}(x_{k+1} - x_{k-1}) - \kappa_f \right] = 0, \quad (20)$$

with boundary conditions $x_1 = 0$ and $x_K = 1$. To solve for the optimal menu, we compute the solutions of this difference equation for each candidate K . We then solve for the optimal K^* . To do so, we observe that the difference equation has two families of solutions: a uniform grid over the points and an "alternating grid" with even spacings between odd points and even points. The second-order conditions of this problem imply that the alternating grid is not a maximum. The optimal menu is then given by:

Proposition 2 (Optimal Menu). *The purchaser optimally offers the following menu:*

$$e_k = \frac{k-1}{K^*-1} \quad w(e_k) = \frac{1}{2} \frac{k-1}{K^*-1} \left(\frac{\beta}{2} \frac{k-1}{K^*-1} + \eta - \kappa_b \right) \quad k \in \{1, \dots, K^*\}, \quad (21)$$

Figure 5: Optimal Procurement Contracts with Costly Evidence



Note: This figure illustrates the optimal contract from Proposition 2 with $\alpha = \beta = 1$, $\eta = 3/2$, $\kappa_b = 1/2$, and $\kappa_f = 1/32$. The optimal contract has $K^* = 4$. The first panel shows the assignment ϕ^* , the second panel shows the payment w^* (with some abuse of notation, extended to all of E), and the third shows the evidence structure $\mathcal{E}^*$.


where K^* satisfies $|K^* - \tilde{K}| < 1$ and $\tilde{K} = 1 + \frac{\beta^2}{12\alpha\kappa_f}$. Moreover, K^* decreases in α , increases in β , decreases in κ_f , and does not depend on η or κ_b .

Proof. See Appendix B.4. □

In the optimal menu, an example of which we illustrate in Figure 5, only a finite number of effort levels and payments are explicitly specified. After promising any one of these effort levels (except the highest), the contractor can, in principle, put in any level of effort within a range without having provably breached the contract (Panel C). But, given any of these efforts, the contractor would receive the same compensation (Panel B); as a result, they keep their promise by putting in the minimal effort consistent with the terms (Panel A). This aligns with real-world contracts that specify obligations under vague contingencies.

The optimality of uniformly spaced effort levels arises due to symmetries in the benefits and costs of more precise contracting. To understand symmetric benefits, we observe that the second derivative $H_{xx} = -\beta$ is constant as a function of (x, θ) . Starting from perfect contractibility, the opportunity cost of removing contractibility in some interval of the action space is the same *regardless* of where that interval is located. This is for two reasons. First, as virtual surplus is quadratic in this model, the purchaser has an equal opportunity cost of forgoing differentiation for high-output versus low-output contractors. Second, because the optimal assignment function is linear, the same measure of types is affected. The corresponding symmetry in costs arises because, for constant costs of distinguishing, producing evidence to distinguish any pair of effort levels has the same cost.

Figure 6: Comparing Front-end and Back-end Costs



Note: This figure illustrates the different implications of front-end and back-end costs for the optimal menu. It shows the optimal payment schedule $w^*(e)$ (with some abuse of notation, extended to all of E) in a parametrization with $\alpha = \beta = 1$, $\eta = 3/2$, and three different calibrations for the costs of evidence: “No costs” with $\kappa_f = \kappa_b = 0$ (gray dashed line), “Front-end only” with $\kappa_f = 1/32$ and $\kappa_b = 0$ (blue solid line), and “Back-end only” with $\kappa_f = 0$ and $\kappa_b = 1/2$.

4.3 The Distinct Roles of Front-End and Back-End Costs

We next contrast the economic role of front-end and back-end costs. To bring this into sharpest relief, Figure 6 shows the optimal payment schedule under three different assumptions for costs: no costs ($\kappa_f = \kappa_b = 0$), front-end costs only ($\kappa_f > 0$, $\kappa_b = 0$), and back-end costs only ($\kappa_f = 0$, $\kappa_b > 0$). As observed in our general analysis (Theorems 1 and 2), the front-end cost is essential for the qualitative prediction of a coarse menu. Under both the “no cost” and “back-end only” scenarios, the optimal contract specifies a continuum of effort levels and payments, although the exact schedule differs between the cases.

Economically, the front-end and back-end costs play separate roles in the comparative statics of the optimal contract. The number of contracted effort levels depends on the front-end cost parameter κ_f but not the back-end cost parameter κ_b , even in the case with both costs (Equation 21). The latter property arises because the back-end cost enters as if it is a linear production cost (see Equation 18). Such a linear cost scales the total surplus of the contract, but does not affect the cost or benefit of more *precise* contracting, which is instead governed by the curvature of the disutility of effort. In our general analysis, the same logic underlies the comparative statics of our richness bound, which depends on the back-end cost only through derivatives of g_b (see Remark 1), which are zero owing to uniformity.

The payments themselves, conditional on the contracted effort levels, depend on the back-end cost parameter κ_b and not the front-end cost parameter κ_f . This is the reason why, in Figure 6, the “front-end only” and “no costs” contracts offer the same payments for the four effort levels that arise in equilibrium under front-end costs, while the “back-end only” contract specifies a lower payment for each effort level. The back-end cost, like a production cost, distorts downward the effort level assigned to each productivity type; because more productive contractors are delivering each effort level when back-end costs are introduced,

the payment schedule shifts downward. The front-end cost, by contrast, affects payments only if it affects information rents and with quadratic utility, this effect is zero.

These comparative statics are consistent with the general intuition that front-end costs act as fixed costs and determine the coarseness of the contract by inducing the principal to economize on the specificity of the contract. The principal economizes as such because they can “re-use” the evidence structure regardless of the realized type of the agent. At the same time, back-end costs act as variable costs and distort the allocation of types to actions but do not affect the coarseness of the contract. Once again, this is because there is no gain to the principal in tying their hands up front to economize on variable costs at the back end, which must be paid type-by-type.

4.4 Comparative Statics: Rigidity in Contracts

We now study how the optimal menu $\{e_k, w(e_k)\}_{k=1}^{K^*}$ responds to changes in productivity. The parameter η can be described as a *neutral* shifter of productivity. The parameter α can be described as a *type-augmenting* shifter of productivity, as it decreases productivity more for agents with higher θ . For this analysis, we assume that parameters $(\alpha, \beta, \eta, \kappa_b, \kappa_f)$ are such that there is a unique K^* . This is true for almost all such vectors of parameters.¹⁵

Corollary 2. *The following statements are true about the optimal menu:*

1. *The optimal effort levels are invariant to η , and payments continuously increase in η .*
2. *For every α , there exists a neighborhood $\mathcal{A}$ of α such that the optimal effort levels and payments are invariant to α' for $\alpha' \in \mathcal{A}$, and there exist different optimal effort levels and payments for $\alpha' \notin \mathcal{A}$.*


Proof. See Appendix B.5 □

Neutral productivity shocks affect total surplus, but they do not affect the purchaser’s *marginal* incentives to write a more precise contract. This property was also echoed in our more general analysis through the lack of dependence of the richness bound on *first derivatives* of the surplus function and, therefore, the first derivatives of u or π (Theorem 1 and Corollary 1). In practice, this is consistent with the motivating observation that vague contracts are ubiquitous even in high-stakes settings, like billion-dollar commercial contracts.

Type-augmenting productivity shocks have a more subtle effect, which depends discontinuously on their size. In response to small shocks, both the number of effort levels and

¹⁵This follows because if there are multiple optimal values for K^* , then they must differ by one. Given this, it is simple to show that, for a situation in which K^* is not unique, an indifference equation of the following form must hold: $\frac{\beta^2}{24\alpha\kappa_f}\Lambda(K^*) = 1$, where Λ is a rational function.

Figure 7: How Payment Schedules Respond to Productivity Changes



Note: This figure illustrates the comparative statics of the optimal contract in response to type-augmenting productivity (α), as per Corollary 2. We plot the payments ($w^*(e)$) and the compensation for each agent ($w^*(\phi^*(\vartheta))$) for a baseline case ($\alpha = 1, K^* = 4$; black solid line), a “small shock” with slightly higher productivity ($\alpha = 0.8, K^* = 4$; gray dotted line), and a “large shock” with much higher productivity ($\alpha = 0.7, K^* = 5$; gray dashed line). We fix $\beta = 1$, $\eta = 3/2$, $\kappa_b = 1/2$, and $\kappa_f = 1/32$. The graph of compensation is truncated at $\vartheta = 0.3$, as lower types are excluded and receive no wages.

the corresponding payments are unaffected. When type-augmenting productivity increases by sufficiently small amounts, agents’ wages may increase, but only through discrete jumps across fixed steps. We illustrate such a scenario in the first row (Panel A) of Figure 7.

By contrast, in response to larger changes, the entire payment schedule can change: that is, the purchaser offers a different and potentially non-nested set of effort levels and corresponding payments. That is, in response to large increases in type-augmenting productivity, effort levels are destroyed and created, and some individual contractors may even be reassigned to lower effort levels. We illustrate such a scenario in the second row (Panel B) of Figure 7: in particular, in the left panel, note that the new payment schedule (gray dashed line) is sometimes above and sometimes below the old schedule (black solid line).

In this way, the model generates size-dependent rigidity in contracts: contracts are perfectly rigid in response to small shocks but feature discrete transitions in response to large shocks. This is true even though our analysis imposes no cost of adjusting the menu.

4.5 Additional Applications

Pay Grades in Labor Procurement. Our findings can help rationalize the ubiquity of pay grades that group workers to have common salaries. In his survey of wage-setting practices for US manufacturing and services firms, [Bewley \(1999\)](#) observes that piece rates that continuously vary compensation with output are relatively rare. A commonly mentioned problem with piece rates is the “cost of establishing the rates.” An alternative arrangement

is a *grade and step system* whereby the full set of tasks is segmented into discrete grades (job titles) and, within each grade, discrete steps with different salaries. Our analysis rationalizes such a system as natural under even small front-end costs of “establishing the rates.”

Quality Segmentation in Input Procurement. We can re-interpret our model such that the principal is a purchasing firm, the agent is a supplier whose costs are given by Equation 16, and the costly action is to produce an input of a given quality. Our result implies that a supplier contract specifies a coarse menu of quality levels and corresponding payments. Our result is consistent with Asanuma’s (1989) description of input sourcing by automobile and machinery manufacturers in Japan. Asanuma (1989) describes how purchasing firms segment suppliers into three categories differentiated by the quality of their inputs and contract in sharply different ways with suppliers in different categories.

Grading and Enforcement in Regulation. As in Baron and Myerson (1982), the agent is a firm that is a monopolist in its market and the principal is a regulator. The regulator imposes subsidies or fines that depend on the quality of the monopolist’s output, but the monopolist has private information about demand. The government designs an evidence structure that legally formalizes a definition of quality. In this setting, the model predicts that the government specifies coarse quality grades. These are common in public health regulations. For example, U.S. regulations governing *Salmonella* for raw poultry products coarsen the results of tests into three categories: much better than standards, just above standards, or below standards (USDA Food Safety and Inspection Service, 2021).

Penalties and Quality Conditions for Rental Goods. A monopolist sells a service like a car or vacation rental that can be utilized to different extents. They choose a menu of utilization levels and prices and face a cost of higher utilization, akin to the production cost in Mussa and Rosen (1978). Moreover, they must also explicitly codify utilization—for example, what constitutes a unit in “good” versus “bad” condition—given front-end and back-end costs. Our results imply that utilization is contracted upon in tiers, such as discrete grades of condition for a car or vacation rental. This lines up with common practice.¹⁶

5 Proofs of Main Results

We now sketch the proof for the main results, up to some technical steps that are relegated to the Appendix. Our argument moves backward through the model timeline (Figure 2).

¹⁶For example, the Europcar terms of service for the United Kingdom specify discrete condition levels for car returns and corresponding fees: if the front bumper of a Mini/Economy rental has a dent of less than 2cm, between 2cm and 5cm, between 5cm and 15cm, or a dent larger than 15cm, then the corresponding fees are £0, £542, £694, £738 (Europcar, 2024).

5.1 Contractibility in the Shadow of the Law

We first make a simple observation about how the evidence structure and potential for punishment shape the agent's action. After selecting (y, τ) from a menu $\mathcal{M}$, the agent can take all of the following actions without being proved to have breached the contract:

$$C_{\mathcal{E}}(y) = \{x \in X : \mathcal{E}(x) \subseteq \mathcal{E}(y)\}. \quad (22)$$

We call $C_{\mathcal{E}} : X \rightrightarrows X$ the *contractibility correspondence* induced by the evidence structure $\mathcal{E}$. The contractibility correspondence represents the set of actions to which the agent optimally restricts themselves in the so-called *shadow of the law*: in anticipation of a large punishment for legally provable deviations, the agent will avoid them. We call a contractibility correspondence C regular if there exists a regular evidence structure $\mathcal{E}$ such that $C = C_{\mathcal{E}}$. The next lemma characterizes regular contractibility correspondences in terms of *self-enforcing recommendations*, i.e., actions from which agents do not deviate:

Lemma 1. *A contractibility correspondence C is regular if and only if there exist two closed sets $\underline{D} \subseteq X$ and $\overline{D} \subseteq X$ such that $0 \in \underline{D}$, $0, \bar{x} \in \overline{D}$ and for all $y \in X$:*

$$C(y) = \left[\max_{z \leq y: z \in \underline{D}} z, \min_{z \geq y: z \in \overline{D}} z \right]. \quad (23)$$


Moreover, given C , $\underline{D}$ and $\overline{D}$ are unique and given by $\underline{D} = \{\min C(y)\}_{y \in X}$ and $\overline{D} = \{\max C(y)\}_{y \in X}$.

Proof. See Appendix C.1. □

For any regular evidence structure $\mathcal{E}$, we can uniquely define its sets of upper and lower self-enforcing recommendations $\overline{D}_{\mathcal{E}}$ and $\underline{D}_{\mathcal{E}}$. We call these the sets of self-enforcing recommendations because they are those that an agent with monotone increasing and decreasing preferences over actions, respectively, would follow. The result implies that regular contractibility correspondences are intervals bounded by their upper and lower envelopes, $\bar{\delta}(y) = \max C(y)$ and $\underline{\delta}(y) = \min C(y)$. The monotonicity of these functions is necessary by monotonicity of the evidence structure. In turn, their images uniquely recover the sets of self-enforcing recommendations, $\overline{D} = \bar{\delta}(X)$ and $\underline{D} = \underline{\delta}(X)$.

Geometrically, Lemma 1 implies that, at any point, $\bar{\delta}$ and $\underline{\delta}$ either coincide with the identity line, or they are flat. To illustrate this, Figure 8 shows the three contractibility correspondences that are generated by the example regular evidence structures shown earlier in Figure 1. Visually, contractibility correspondences can be composed from three basic struc-

Figure 8: Examples of Contractibility Correspondences



Note: This Figure illustrates the contractibility correspondences generated by the three regular evidence structures illustrated earlier in Figure 1 (Panels A-C). The action space is $X = [0, 1]$. Under each example, we give the corresponding set of self-enforcing recommendations.

tures: “disposal” (“lower triangles”), “creation” (“upper triangles”), or indistinguishability (“boxes”). Focusing on the sets of *upper* self-enforcing recommendations is sufficient:

Lemma 2. *For any regular evidence structure $\mathcal{E}$ such that $\underline{D}_{\mathcal{E}} \not\subseteq \{0, \bar{x}\}$, there exists a regular evidence structure $\mathcal{E}'$ with $\underline{D}_{\mathcal{E}'} \subseteq \{0, \bar{x}\}$ and such that $\mathcal{R}(\mathcal{E}') - \Gamma_f(\mathcal{E}') \geq \mathcal{R}(\mathcal{E}) - \Gamma_f(\mathcal{E})$, with the previous inequality being strict unless the principal faces only back-end costs.*

Proof. See Appendix B.1. □

The simple reason is that agents with monotone increasing preferences never prefer to deviate to lower actions, and building an evidence structure that distinguishes higher actions from lower actions is costly. Thus, for the remainder of the analysis, it is without loss of optimality to set $\underline{D}_{\mathcal{E}} = \{0\}$ and focus on $\overline{D}_{\mathcal{E}}$.¹⁷ With this, it is convenient to define $\overline{D} \subseteq 2^X$ as the set of closed subsets of X that contain 0 and $\bar{x}$.

5.2 Optimal Contracts and the Value of Evidence Structures

We now turn to “Contract Design” (Step 2 in Figure 2), the choice of the menu $\mathcal{M} = \{(y_i, T(y_i))\}_{i \in I}$. The following lemma shows that the solution to this problem can be described as a pointwise maximization of the modified virtual surplus function, subject to the constraint that actions are self-enforcing:

Lemma 3. *For every regular evidence structure $\mathcal{E}$ with $\underline{D}_{\mathcal{E}} = \{0\}$, we have $\mathcal{R}(\mathcal{E}) = \mathcal{H}(\overline{D}_{\mathcal{E}})$, where:*

$$\mathcal{H}(\overline{D}) = \int_{\Theta} \max_{x \in \overline{D}} H(x, \theta) dF(\theta) \quad \forall \overline{D} \in \overline{D}. \quad (24)$$

¹⁷This follows from the fact that, for all regular evidence structures $\mathcal{E}$ and $\mathcal{E}'$ satisfying $\overline{D}_{\mathcal{E}} = \overline{D}_{\mathcal{E}'}$, $\underline{D}_{\mathcal{E}'} = \{0, \bar{x}\}$, and $\underline{D}_{\mathcal{E}} = \{0\}$, the principal can implement under $\mathcal{E}$ the same equilibrium outcomes and front-end costs that are attainable under $\mathcal{E}'$, while incurring a weakly lower back-end cost under the former because a positive mass of agents can all take final action $\bar{x}$ in the selected equilibrium.

Proof. See Appendix A.1. □

We prove this result in three steps. In the first step, we characterize implementable *final action* functions $\phi : \Theta \rightarrow X$ given $\mathcal{E}$. Each ϕ is implementable given $\mathcal{E}$ if and only if it is monotone increasing and maps into the set of self-enforcing recommendations: $\phi(\Theta) \subseteq \overline{D}_{\mathcal{E}}$. These properties are intuitively *necessary* because of incentive compatibility and the fact that agents will always take the highest action that is feasible under the shadow of the law. A more technical argument is required to show they are *sufficient*, given that agents can contemplate global “double deviations”: type θ could imitate type θ' , receive their recommendation $\phi(\theta')$, and then take a different final action $x \in C_{\mathcal{E}}(\phi(\theta'))$.

In the second step, we show that, for every $\mathcal{E}$ and every ϕ that is implementable given $\mathcal{E}$:

$$\Gamma_b(\mathcal{E}, Q_{\phi}) = \int_{\Theta} \int_{\phi(\theta)}^{\bar{x}} g_b(\phi(\theta), z) dz dF(\theta), \quad (25)$$

where $Q_{\phi} = F \circ \phi^{-1}$ is the distribution over actions induced by ϕ . This equality follows from a change of variable and the fact that $\phi(\Theta) \subseteq \overline{D}_{\mathcal{E}}$. Moreover, it establishes a key fact: the back-end cost does not depend on the structure of $\mathcal{E}$ when evaluated at one of the *equilibrium* distributions of actions that it induces. This is the formalization of the variable costs logic that we highlighted in Section 3.3: as costs are paid *ex post* type-by-type, the evidence structure is itself irrelevant for costs conditional on the costs of distinguishing that are ultimately paid in equilibrium. A direct implication of this is that the back-end cost induces a cost $\int_x^{\bar{x}} g_b(x, z) dz$ for every equilibrium final action x .

In the third step, we express the transfer function $t : \Theta \rightarrow \mathbb{R}$ in terms of ϕ using local incentive compatibility and individual rationality. This allows us to define the modified virtual surplus function $H(x, \theta) = J(x, \theta) - \int_x^{\bar{x}} g_b(x, z) dz$, inclusive of the “as-if production cost” induced by the back-end cost of evidence. Finally, we derive the pointwise formulation of Equation 24 by observing that the pointwise maximization has a monotone solution, due to the supermodularity of J , and thus H , in (x, θ) .¹⁸

The Solution Under Back-end Costs. Theorem 2 follows as a corollary of Lemma 3. To see why, it is immediate that $\mathcal{H}(X) \geq \mathcal{H}(\overline{D})$ for all $\overline{D} \in \overline{\mathcal{D}}$. In the absence of front-end costs, any $\mathcal{E}$ such that $\overline{D}_{\mathcal{E}} = X$ is optimal. Intuitively, because the evidence structure matters only through the realized back-end costs that are incurred *ex post*, there is no gain to the principal in tying their hands and designing an imperfect evidence structure *ex ante*.

¹⁸If H were strictly concave (but not necessarily supermodular), then at this point one could apply the techniques of Toikka (2011) to characterize the value. While it would be relatively simple to adapt our subsequent arguments to the non-supermodular case, we do not do so in the interests of space.

The remainder of the proof in Appendix A.2 shows that it is *strictly* optimal for the principal to choose an evidence structure and menu with the cardinality of the continuum when H is non-trivial. Intuitively, the principal wishes to allocate a continuum of different actions to a positive measure subset of agents.

5.3 The Existence of an Optimal Evidence Structure

We next establish the existence of an optimal evidence structure in the presence of both front-end and back-end costs. Because of Lemmas 1 and 2, it suffices to show the existence of an optimal set of upper self-enforcing recommendations $\overline{D}$. We have shown in Lemma 3 that we can express the value of an evidence structure via the function $\mathcal{H} : \overline{D} \rightarrow \mathbb{R}$. By Berge's maximum theorem, this function is continuous in the Hausdorff topology. We can also express the front-end cost associated with a regular evidence structure $\mathcal{E}$ with $\underline{D}_{\mathcal{E}} = \{0\}$ as the function $\tilde{\Gamma}_f : \overline{D} \rightarrow \mathbb{R}$ defined by $\tilde{\Gamma}_f(\overline{D}_{\mathcal{E}}) = \Gamma_f(\mathcal{E})$. We show that this function is continuous in the Hausdorff topology. The existence of an optimal evidence structure finally follows from Weierstrass' theorem because $\overline{D}$ is compact in the Hausdorff topology.

Lemma 4. $\mathcal{H} : \overline{D} \rightarrow \mathbb{R}$ and $\tilde{\Gamma}_f : \overline{D} \rightarrow \mathbb{R}$ are continuous in the Hausdorff topology. Therefore, an optimal evidence structure $\mathcal{E}^*$ exists.

Proof. See Appendix A.3. □


5.4 The Costs and Benefits of Precision

We next use a variational approach to derive the optimality of coarse evidence structures and menus as well as a bound on their cardinality. The key step is to show that an evidence structure with self-enforcing recommendations $a < b$, but no other $x \in (a, b)$, is strictly better than one that adds a self-enforcing recommendation $a < z < b$, provided that a and b are sufficiently close. We recall that $H(x, \theta) = J(x, \theta) - \int_x^{\bar{x}} g_b(x, z) dz$ and $\bar{H}_{xx} = \max_{x, \theta} |H_{xx}(x, \theta)|$, $\bar{H}_{x\theta} = \min_{x, \theta} H_{x\theta}(x, \theta)$, and $\bar{f} = \max_{\theta} f(\theta)$. We focus on the case of Corollary 1 in which there are both front-end and back-end costs, and H is strictly quasi-concave and not strictly quasi-convex.¹⁹ Of course, if $g_b \equiv 0$ as in Theorem 1, then $H = J$ is strictly quasi-concave by assumption.

Benefits. We define the function that maximizes modified virtual surplus as $\phi^P(\theta) = \arg \max_{x \in X} H(x, \theta)$, which is unique (by strict quasi-concavity), increasing, and moreover strictly increasing on the interior of X (by strict supermodularity).

¹⁹Under strict quasi-convexity of H (or of J in the case of Theorem 1), the conclusion of coarseness is trivial as the principal assigns all agents to either 0 or $\bar{x}$.

Figure 9: Illustrating the Payoff Gains of Precise Contractibility (Lemma 5)



Note: This figure illustrates one step in deriving the bound in Lemma 5. The black curves illustrate the virtual surplus function for three types $\theta_0 < \theta_1 < \theta_2$, and type θ_1 is such that the principal's preferred allocation is $\phi^P(\theta_1) \in (a, b)$. We annotate the relevant calculations for the first-order and second-order terms of the bound in blue and red, respectively. The orange line is tangent to $H(\cdot, \theta_1)$ at $x = z$, representing the first-order term in a Taylor expansion of $H(\cdot, \theta_1)$ around z . The action $w \in (z, b)$ used in the bound of the second-order term comes from Taylor's remainder theorem.

Lemma 5. Fix any $\bar{D} \in \bar{\mathcal{D}}$ and any $a, b \in \bar{D} \cap \phi^P(\Theta)$ such that $a < b$ and $\bar{D} \cap (a, b) = \emptyset$. For all $z \in (a, b)$, we have that:

$$\mathcal{H}(\bar{D} \cup \{z\}) - \mathcal{H}(\bar{D}) \leq \frac{3}{2} \frac{\bar{H}_{xx}^2 \bar{f}}{\bar{H}_{x\theta}} (b-a)(b-z)(z-a). \quad (26)$$


Proof. See Appendix A.4. □

We sketch the derivation of this inequality below. When the action z is added between a and b , agents fall into three groups: agents who continue to take their previous action, agents who previously took action a and now take action z , and agents who previously took action b and now take action z . To describe which types fall into the latter two cases, we define the type that the principal is indifferent to assign to $x, y \in \phi^P(\Theta)$ with the function $\hat{\theta} : \phi^P(\Theta)^2 \rightarrow \Theta$. Agents $\theta \in [\hat{\theta}(a, z), \hat{\theta}(a, b)]$ are now assigned to z rather than a and agents $\theta \in [\hat{\theta}(a, b), \hat{\theta}(z, b)]$ are now assigned to z rather than b . Hence, the total change in value is:

$$\mathcal{H}(\bar{D} \cup \{z\}) - \mathcal{H}(\bar{D}) = \int_{\hat{\theta}(a, z)}^{\hat{\theta}(a, b)} (H(z, \theta) - H(a, \theta)) dF(\theta) + \int_{\hat{\theta}(a, b)}^{\hat{\theta}(z, b)} (H(z, \theta) - H(b, \theta)) dF(\theta). \quad (27)$$

To bound this object, we first bound the integrands, focusing for simplicity on $H(z, \theta) - H(b, \theta)$. We illustrate this procedure in Figure 9. Specifically, because $\phi^P(\theta) \in [a, b]$ for all $\theta \in [\hat{\theta}(a, b), \hat{\theta}(z, b)]$ due to strict quasi-concavity, we can apply Taylor's remainder theorem twice to conclude that $|H(z, \theta) - H(b, \theta)| \leq \bar{H}_{xx} \left((b-a)(b-z) + \frac{1}{2}(b-z)^2 \right)$. Intuitively, when concavity $\bar{H}_{xx}$ is higher, the principal gains more by tailoring actions to types.

Figure 10: Illustrating the Cost of Precise Contractibility (Lemma 6)



Note: This figure illustrates the cost-saving bound of Lemma 6 in the case of constant costs of distinguishing, or $g_f \equiv \kappa > 0$. The blue solid line illustrates $\bar{\delta}$ for a $\bar{D}$ which contains the isolated points $a < b$, and the orange dashed line illustrates $\bar{\delta}$ for $\bar{D} \cup \{z\}$ for some $z \in (a, b)$. The colored hashed areas illustrate the front-end costs under each, and the gray rectangle is proportional to the difference in front-end costs.

We then bound the limits of integration. By applying the implicit function theorem to $\hat{\theta}$, we can show that $\hat{\theta}(z, b) - \hat{\theta}(a, b) \leq \frac{1}{2} \frac{\bar{H}_{xx}}{H_{x\theta}}(z - a)$. If H is either very concave or not very supermodular, then many types of agents are allocated to the new contractible action z .

We can then combine all of the information to place the following upper bound on the second integral: $\frac{1}{2} \frac{\bar{H}_{xx}^2 \bar{f}}{H_{x\theta}}(z - a) ((b - a)(b - z) + \frac{1}{2}(b - z)^2)$. Performing the same steps for the first integral and factorizing yields the bound.

Costs. We next bound from below the marginal front-end cost of adding an isolated point. We remind that $\underline{g}_f = \min_{x, z \in X} g_f(x, z) > 0$.

Lemma 6. Fix any $\bar{D} \in \bar{\mathcal{D}}$ and any $a, b \in \bar{D}$ such that $a < b$ and $\bar{D} \cap (a, b) = \emptyset$. For all $z \in (a, b)$, we have that:

$$\tilde{\Gamma}_f(\bar{D} \cup \{z\}) - \tilde{\Gamma}_f(\bar{D}) \geq \underline{g}_f(b - z)(z - a). \quad (28)$$

Proof. See Appendix A.5. □

To understand this bound, we first re-express the front-end cost as the sum of costs across actions. We show that:

$$\Gamma_f(\mathcal{E}) = \int_X \left(\int_0^{\underline{\delta}_{\mathcal{E}}(x)} g_f(x, q) dq + \int_{\bar{\delta}_{\mathcal{E}}(x)}^{\bar{x}} g_f(x, q) dq \right) dx = \int_X \int_{\bar{\delta}_{\mathcal{E}}(x)}^{\bar{x}} g_f(x, q) dq dx, \quad (29)$$

where $\underline{\delta}_{\mathcal{E}} = \min C_{\mathcal{E}}$, $\bar{\delta}_{\mathcal{E}} = \max C_{\mathcal{E}}$. Intuitively, the principal pays to distinguish each hypothetical action x from potential deviations $q < \underline{\delta}_{\mathcal{E}}(x)$ and $q > \bar{\delta}_{\mathcal{E}}(x)$. As the principal sets

$\delta_{\mathcal{E}}(x) = 0$, only the second term remains. Graphically, the front-end cost corresponds to the area above the contractibility correspondence (*i.e.*, above $\bar{\delta}_{\mathcal{E}} = \max C_{\mathcal{E}}$) weighted by g_f .

To build intuition, we derive the bound in the simple case of constant costs of distinguishing ($g_f \equiv \kappa > 0$) and we illustrate this in Figure 10. In this case, the front-end cost is proportional to the area above $\bar{\delta}$. For each desired action x , the principal pays a cost proportional to the measure of potential deviations q from which x must be distinguished. The additional cost from adding self-enforcing recommendation z comes from distinguishing potential deviations $q \in (z, b]$ from desired actions $x \in (a, z]$. Thus, the cost is equal to $\kappa(b - z)(z - a)$, illustrated by the gray shading in Figure 10. To derive the bound under non-constant costs, we simply replace κ with the minimum value of g_f .

5.5 Establishing Coarseness

We finally combine these bounds to establish the optimality of coarseness:

Lemma 7. *If $\bar{D} \in \bar{\mathcal{D}}$ contains two points $a, b \in \bar{D} \cap \phi^P(\Theta)$ such that $0 < b - a < \frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 f}$ and $\bar{D} \cap (a, b) \neq \emptyset$, then $\bar{D}$ is not optimal. Moreover, any optimal $\bar{D}^* \in \bar{\mathcal{D}}$ satisfies:*

$$|\bar{D}^*| \leq 6 + \left\lceil 3 \frac{\bar{x} \bar{H}_{xx}^2 \bar{f}}{g_f H_{x\theta}} \right\rceil. \quad (30)$$

Proof. See Appendix A.6. □

To build intuition for the first part of the result, let us first assume that $\bar{D} \cap (a, b)$ contains a single point. Because a and b are close enough by assumption, Lemmas 5 and 6 yield that $\bar{D}$ cannot be optimal. Economically, this follows from the observation that the revenue benefits of adding an additional self-enforcing recommendation are an order of magnitude smaller than the front-end costs of doing so. The proof of the result then extends this logic to when $\bar{D} \cap (a, b)$ is an arbitrary set, possibly countably infinite, uncountably infinite and dense, or even uncountably infinite and nowhere dense (*e.g.*, a fractal set). The key idea is that we can remove all points between a and b and show that every time we re-introduce a point, the total value goes down. As the principal's objective is continuous in the Hausdorff topology, this guarantees that removing all such points is an improvement. The bound follows from the requirement that no two points in $\bar{D} \cap \phi^P(\Theta)$ can be too close.

5.6 Relationship To Other Work on Finite Menus

Finally, we comment on the technical relationship between our results and those in other analyses of design problems with finite menus.

Costs of Contracts and Finite Contracts. An influential literature in contract theory has studied how costs of writing contracts can give rise to optimally incomplete contracts (Dye, 1985; Bajari and Tadelis, 2001; Battigalli and Maggi, 2002). In these papers, the cost of a contract depends on the number of elements on the contract, a form of a front-end cost in the language of Scott and Triantis (2005). In our setting, we would write these costs as $\tilde{\Gamma}_f(\overline{D}) = \kappa_f |\overline{D}|$. These costs make infinite contracts infinitely costly and thus led to a critique by Tirole (1999): while writing linear contracts may be costly, it is not plausible to believe that such costs are infinite. Under our approach, incomplete contracts may be optimal even when perfect contractibility is arbitrarily cheap.

Remark 2 (Coarseness Obtains Despite Vanishing Marginal Costs). Consider a constant cost of distinguishing actions, $g_f \equiv \kappa_f > 0$. For any finite number of action recommendations $K \in \mathbb{N}$, the cost of having a uniform grid of K self-enforcing recommendations is:

$$\tilde{\Gamma}_f \left(\left\{ \frac{k-1}{K-1} \bar{x} \right\}_{k=1}^K \right) = \frac{\kappa_f \bar{x}^2}{2} \frac{K-2}{K-1}, \quad (31)$$

which is a *strictly concave* function of the number of self-enforcing recommendations K . As $K \rightarrow \infty$, we have $\tilde{\Gamma}_f \rightarrow \kappa_f \bar{x}^2 / 2$. Thus, even when the marginal cost saving of removing contractibility at infinity is zero, contracts can nevertheless be optimally coarse. $\triangle$

Moreover, our results demonstrate a subtlety in arguing that costs of contracts will necessarily yield incompleteness; this conclusion hinges on the nature of costs and whether they are incurred at the *ex ante* or the *ex post* stage.

Technical Approaches to Prove Finiteness of Menus. Here, we connect Theorem 1 to other results in the literature of mechanism and information design that derive the finiteness of optimal contracts, menus, or experiments. As explained, our main technical argument can be summarized in three steps: (i) deriving a lower bound on the cost of adding contractibility points; (ii) deriving an upper bound on the benefit of doing so; and (iii) applying a finite approximation result. The first step is the least related to prior work, since it pertains to the variational analysis of the front-end cost of distinguishing introduced in our model.

The second step is related to a literature in information theory that studies the optimal discretization of a continuous signal (see *e.g.*, Lloyd, 1982). Wilson (1989) used a similar approach to bound the loss of a principal screening an agent by using only finite menus with a fixed number of items. These results are refined and extended to the multidimensional case in Bergemann, Yeh, and Zhang (2021). We contribute by providing a bound that depends on the *structure* of the menu and the distances between adjacent triples of points, as opposed to

just the menu’s cardinality. This is important because any cardinality-based bound would be insufficient to derive our results due to the richness of our costs.

The third step is similar to the approximation arguments used by Bergemann and Pendorfer (2007), Lou (2022), and Bergemann, Heumann, and Morris (2025) to derive the optimality of finite information structures in information and mechanism design problems. Unlike these papers, we also provide an explicit bound on the extent of coarseness.²⁰ Finally, the tools developed in this paper could be applied to costly design in alternative economic settings, such as delegation and information disclosure.

6 Conclusion

In this paper, we introduce a model of contractibility design. Our analysis has two premises. The first is that contracts are only enforceable to the extent that the principal can prove that the agent deviated from the terms of the contract. The second is that the codification and generation of evidence that can be used to provide such proof potentially entail front-end or back-end costs. Front-end costs yield coarse contracts while back-end costs yield complete contracts. In the presence of both costs, we show how the agent’s and principal’s preferences affect the structure and potential coarseness of contracts. We argue that our model generates insights into the nature of procurement contracts, worker hiring, government regulation, and monopoly pricing. In particular, we rationalize why even very small front-end costs of writing contracts may give rise to incomplete contracts.

A Additional Proofs of the Main Results

The proofs of Lemmas 1 and 2 are mechanical and we defer them to Appendices C.1 and B.1. However, in what follows, we exploit that every regular evidence structure $\mathcal{E}$ affects the contracting problem solely via the induced regular contractibility correspondence $C = C_{\mathcal{E}}$ as explained in Section 5.1. Therefore, we often start with a regular C as the primitive object.

A.1 Proof of Lemma 3

We prove the result in three parts. First, we present a characterization of implementable allocations. Second, we use this characterization to derive the principal’s control problem. Third, we solve this control problem for the optimal contract.

²⁰Our approach is also different from alternative approaches based on duality (*e.g.*, Jung, Kim, Matějka, and Sims, 2019; Amador, Bagwell, and Carpizo, 2025) or variational arguments used in the optimal-delegation literature (*e.g.*, Alonso and Matouschek, 2008; Dovis, Kirpalani, and Sublet, 2024).

Part 1: Implementation. We begin by establishing a general taxation principle with partial contractibility induced by an evidence structure. To do this, we recast the inner problem of designing the menu (Equation 5) as the following mechanism design problem. Given a fixed contractibility correspondence C , the revelation principle allows us to restrict to direct and truthful mechanisms.²¹ Thus, a mechanism is a triple (ϕ, ξ, T) comprising a recommendation $\xi : \Theta \rightarrow X$, a final action or outcome $\phi : \Theta \rightarrow X$, and a tariff $T : X \rightarrow \bar{\mathbb{R}}$. The tariff and the recommendation jointly determine the transfer between the principal and the agent $T(\xi(\theta))$. The recommendation and the final action must be consistent with the contractibility correspondence, that is, $\phi(\theta) \in C(\xi(\theta))$. This, together with incentive constraints, determines the set of implementable mechanisms.

Definition 3 (Implementable Mechanism). *A mechanism (ϕ, ξ, T) is implementable given a contractibility correspondence C if the following three conditions are satisfied:*

1. *Obedience:*

$$\phi(\theta) \in \arg \max_{x \in C(\xi(\theta))} u(x, \theta) \quad \text{for all } \theta \in \Theta. \quad (\text{O})$$

2. *Incentive Compatibility:*

$$\xi(\theta) \in \arg \max_{y \in X} \left\{ \max_{x \in C(y)} u(x, \theta) - T(y) \right\} \quad \text{for all } \theta \in \Theta. \quad (\text{IC})$$

3. *Individual Rationality:*

$$u(\phi(\theta), \theta) - T(\xi(\theta)) \geq 0 \quad \text{for all } \theta \in \Theta. \quad (\text{IR})$$

We let $\mathcal{I}(C)$ denote the set of implementable mechanisms given C .

Given a regular contractibility correspondence C , we say that $T : X \rightarrow \bar{\mathbb{R}}$ is monotone with respect to C if $T(y) \geq T(x)$ for all $x, y \in X$ such that $x \in C(y)$. Next, we establish that monotonicity of the tariff with respect to C is necessary and sufficient for implementability.

Lemma 8 (C -Monotone Taxation Principle). *Fix a regular contractibility correspondence C . A final action function ϕ is implementable given C if and only if there exists a tariff $T : X \rightarrow \bar{\mathbb{R}}$ that is monotone with respect to C and such that:*

$$\phi(\theta) \in \arg \max_{x \in X} \{u(x, \theta) - T(x)\} \quad (32)$$

²¹Following the standard approach in mechanism design, we select the principal's preferred equilibrium and restrict to deterministic mechanisms. We omit the formal proof of the revelation principle in this case as it is standard and closely follows the steps in [Myerson \(1982\)](#).

and $u(\phi(\theta), \theta) - T(\phi(\theta)) \geq 0$ for all $\theta \in \Theta$. In this case, ϕ is supported by $\xi = \phi$ and T .

Proof. See Appendix B.2. □

With this result in hand, we now show:

Lemma 9 (Implementation). *A final action function ϕ is implementable under $C = [\underline{\delta}, \bar{\delta}]$, with self-enforcing recommendation sets $\underline{D} = \underline{\delta}(X)$ and $\bar{D} = \bar{\delta}(X)$, if and only if it is monotone increasing and such that $\phi(\Theta) \subseteq \bar{D}$. Moreover, ϕ is supported by $\xi = \phi$ and tariff:*

$$T(x) = T(0) + u(x, \phi^{-1}(x)) - \int_0^{\phi^{-1}(x)} u_\theta(\phi(s), s) ds, \quad \forall x \in \phi(\Theta), \quad (33)$$

where $\phi^{-1}(s) = \inf\{\theta \in \Theta : \phi(\theta) \geq s\}$.

Proof. (Only If for First Part) If ϕ is implementable, then there exists (ξ, T) that support ϕ . By Lemma 8, we may take that $\xi = \phi$. By (IC) and Lemma 8, there exists a transfer function $t : \Theta \rightarrow \mathbb{R}$ given by $t(\theta) = T(\phi(\theta))$ such that $u(\phi(\theta), \theta) - t(\theta) \geq u(\phi(\theta'), \theta) - t(\theta')$ for all $\theta, \theta' \in \Theta$. As u is strictly single-crossing, Proposition 1 in Rochet (1987) then implies that ϕ is monotone. Toward a contradiction suppose that $\phi(\theta) \notin \bar{D}$. Deviating to $\bar{\delta}(\phi(\theta)) > \phi(\theta)$ is a strict improvement for the agent. Thus, if ϕ is implementable, then $\phi(\Theta) \subseteq \bar{D}$ holds.

(If For First Part) Without loss of generality, we again prove this part for the case with monotone increasing preferences. Now suppose that $\phi(\theta) \in \bar{D}$ holds for all $\theta \in \Theta$ and ϕ is monotone increasing. Define the function $t : \Theta \rightarrow \mathbb{R}$ as

$$t(\theta) = Z + u(\phi(\theta), \theta) - \int_0^\theta u_\theta(\phi(s), s) ds, \quad (34)$$

for some $Z \leq 0$, and the tariff $T : X \rightarrow \bar{\mathbb{R}}$ as $T(x) = \inf_{\theta' \in \Theta} \{t(\theta') : x \in C(\phi(\theta'))\}$. Fix $x, y \in X$ such that $x \in C(y)$. By transitivity of C , for all $\theta \in \Theta$, if $y \in C(\phi(\theta))$, then $x \in C(\phi(\theta))$. This shows that $\{\theta \in \Theta : y \in C(\phi(\theta))\} \subseteq \{\theta \in \Theta : x \in C(\phi(\theta))\}$. Therefore, applying the construction of T , $T(y) \geq T(x)$. Thus, T is monotone with respect to C .

As T is monotone with respect to C , if we can show that $\phi(\theta) \in \arg \max_{x \in X} \{u(x, \theta) - T(x)\}$ and $u(\phi(\theta), \theta) - T(\phi(\theta)) \geq 0$, then we have shown by Lemma 8 that ϕ is implementable.

We start with the second condition. For every $\theta \in \Theta$, we have

$$u(\phi(\theta), \theta) - T(\phi(\theta)) \geq u(\phi(\theta), \theta) - t(\theta) = \int_0^\theta u_\theta(\phi(s), s) ds - Z. \quad (35)$$

The right-hand side of this last equation is monotone increasing in θ since it is absolutely continuous with derivative $u_\theta(\phi(\theta), \theta) = \int_0^{\phi(\theta)} u_{x\theta}(z, \theta) dz \geq 0$ for almost all $\theta \in \Theta$, because u is supermodular. Given that $Z \leq 0$, we have that $u(\phi(\theta), \theta) - T(\phi(\theta)) \geq 0$ for all $\theta \in \Theta$.

We are left to prove that (ϕ, T) satisfies Equation 32. We first prove that, for all $\theta, \theta' \in \Theta$:

$$u(\phi(\theta), \theta) - t(\theta) \geq \max_{x \in C(\phi(\theta'))} u(x, \theta) - t(\theta'). \quad (36)$$

This is a variation of the standard reporting problem under the final action function ϕ and transfers t , where each agent, on top of misreporting their type, can also consume everything allowed by C . Violations of this condition can take two forms. First, an agent of type θ could report type θ' and consume $x = \phi(\theta')$. We call this a single deviation. Second, an agent of type θ could report type θ' and consume $x \in C(\phi(\theta')) \setminus \{\phi(\theta')\}$. We call this a double deviation. Under our construction of transfers t and monotonicity of ϕ , by a standard mechanism-design argument, there is no strict gain to any agent of reporting θ' and consuming $x = \phi(\theta')$. Thus, there are no profitable single deviations under (ϕ, t) .

Suppose now that θ imitates θ' and takes final action $x \neq \phi(\theta')$. As $\phi(\theta') \in \overline{D}$, $x < \phi(\theta')$. But in that case, simply taking action $\phi(\theta')$ is better. But then this is a single deviation, which we have ruled out. To derive the tariff, for each $x \in \phi(\Theta)$ choose $\theta_x \in \Theta$ such that $\phi(\theta_x) = x$ and set $T(x) = t(\theta_x)$. $\square$

Part 2: Control Problem. We now turn the problem into a control problem:

Lemma 10. *Fix a regular $\mathcal{E}$ with corresponding regular contractibility correspondence C and induced $(\underline{D}, \overline{D})$ such that $\underline{D} = \{0\}$. Any optimal final action function solves:*

$$\begin{aligned} \mathcal{H}(\overline{D}) &= \max_{\phi} \int_{\Theta} H(\phi(\theta), \theta) dF(\theta) \\ \text{s.t. } &\phi(\theta') \geq \phi(\theta), \phi(\theta) \in \overline{D}, \quad \theta, \theta' \in \Theta : \theta' \geq \theta. \end{aligned} \quad (37)$$

Proof. We begin by eliminating the proposed allocation and transfers from the objective. From the proof of Lemma 9, transfers for any incentive compatible (ξ, ϕ, t) are given by:

$$t(\theta) = Z + u(\phi(\theta), \theta) - \int_0^{\theta} u_{\theta}(\phi(s), s) ds, \quad (38)$$

for some constant $Z \in \mathbb{R}$. Thus, any ξ that supports ϕ leads to the same principal payoff and can therefore be made equal to ϕ without loss of optimality. Moreover, we know that ϕ being incentive compatible is equivalent to ϕ being monotone increasing and $\phi(\theta) \in \overline{D}$.

We next re-express the back-end costs. We start from the definition,

$$\Gamma_b(\mathcal{E}, Q) = \int_X \gamma_{g_b}(\mathcal{E}, x) dQ(x) = \int_X \int_{\mathcal{E}(x)^c \cap \mathcal{E}(z) \neq \emptyset} g_b(x, z) dz dQ(x). \quad (39)$$

By definition of C , we can write the domain of the inner integral as $\{z : z \notin C(x)\}$. Moreover, from Lemma 1, $C(x) = [\underline{\delta}(x), \bar{\delta}(x)]$, for some envelope functions $\bar{\delta}, \underline{\delta}$. Therefore,

$$\Gamma_b(\mathcal{E}, Q) = \int_X \left(\int_0^{\bar{\delta}(x)} g_b(x, z) dz + \int_{\bar{\delta}(x)}^{\bar{x}} g_b(x, z) dz \right) dQ(x). \quad (40)$$

Moreover, $\underline{\delta} = 0$ by assumption that $\underline{D} = \{0\}$. We further observe, in light of the implementation argument above, that $x = \bar{\delta}(x)$ for any action x that is chosen in equilibrium. Using these simplifications, as well as a change of variables from $x, Q_\phi(x)$ to $\theta, F(\theta)$, we write

$$\Gamma_b(\mathcal{E}, Q_\phi) = \int_\Theta \int_{\phi(\theta)}^{\bar{x}} g_b(\phi(\theta), z) dz dF(\theta). \quad (41)$$

By these arguments and Lemma 9, we write the principal's problem as

$$\begin{aligned} & \max_{\phi, Z} \int_\Theta \left(\pi(\phi(\theta), \theta) + Z + u(\phi(\theta), \theta) - \int_0^\theta u_\theta(\phi(s), s) ds - \int_{\phi(\theta)}^{\bar{x}} g_b(\phi(\theta), z) dz \right) dF(\theta) \\ \text{s.t. } & \phi(\theta') \geq \phi(\theta), \quad \phi(\theta) \in \bar{D} \quad \forall \theta, \theta' \in \Theta : \theta' \geq \theta, \\ & u(\phi(\theta), \theta) - \left(Z + u(\phi(\theta), \theta) - \int_0^\theta u_\theta(\phi(s), s) ds \right) \geq 0 \quad \forall \theta \in \Theta. \end{aligned} \quad (42)$$

We further simplify this by applying integration by parts on the double integral of $u_\theta(\phi(s), s)$ over θ and s . Plugging into the principal's objective, we find that the principal solves:

$$\begin{aligned} & \max_{\phi, Z} \int_\Theta (H(\phi(\theta), \theta) + Z) dF(\theta) \quad \text{s.t.} \quad \phi(\theta') \geq \phi(\theta), \quad \phi(\theta) \in \bar{D} \quad \forall \theta, \theta' \in \Theta : \theta' \geq \theta \\ & u(\phi(\theta), \theta) - \left(Z + u(\phi(\theta), \theta) - \int_0^\theta u_\theta(\phi(s), s) ds \right) \geq 0 \quad \forall \theta \in \Theta, \end{aligned} \quad (43)$$

where we define $H(x, \theta) = \pi(x, \theta) + u(x, \theta) - \frac{1-F(\theta)}{f(\theta)} u_\theta(x, \theta) - \int_x^{\bar{x}} g_b(x, z) dz$. Thus, it is optimal to set $Z \in \mathbb{R}$ as large as possible such that:

$$V(\theta) = u(\phi(\theta), \theta) - \left(Z + u(\phi(\theta), \theta) - \int_0^\theta u_\theta(\phi(s), s) ds \right) \geq 0 \quad \forall \theta \in \Theta. \quad (44)$$

We know that $V'(\theta) = u_\theta(\phi(\theta), \theta) \geq 0$ almost everywhere as we have shown that $u(x, \cdot)$ is monotone. Thus, the tightest such constraint occurs when $\theta = 0$. Hence, the maximal Z must satisfy $V(0) = -Z \geq 0$. This implies that Z is optimally 0. $\square$

Part 3: The Optimal Contract. We first solve the pointwise problem in the control problem from Lemma 10 and verify that this solution is monotone. The pointwise problem

is $\max_{x \in \overline{D}} H(x, \theta)$. The maximum exists as H is continuous and $\overline{D}$ is compact. Moreover, by Theorem 4' in [Milgrom and Shannon \(1994\)](#), every selection from the $\arg \max$ is monotone increasing and hence implementable. For any such selection ϕ^* , the claim that we may take $\xi^* = \phi^*$ and the formula for the optimal tariff follow from applying Lemma 9.

A.2 Remainder of the Proof of Theorem 2

Under only back-end costs, $\mathcal{R}(\mathcal{E}) = \mathcal{H}(\overline{D}_{\mathcal{E}})$ and $\Gamma_f(\mathcal{E}) = 0$. Thus, the program that defines an optimal evidence structure (if it exists), Equation 6, is $\sup_{\mathcal{E}} \mathcal{H}(\overline{D}_{\mathcal{E}})$. From Lemma 10, it is immediate that $\mathcal{H}(X) \geq \mathcal{H}(\overline{D})$ for any $\overline{D} \in \overline{\mathcal{D}}$. Moreover, because there exists a subset $\tilde{\Omega} \subseteq \Omega$ and an order preserving bijection $\omega : X \rightarrow \tilde{\Omega}$, it follows that the regular evidence structure defined by $\mathcal{E}^*(x) \equiv \{\omega(z) : 0 \leq z \leq x\}$ is such that $\overline{D}_{\mathcal{E}^*} = X$.

We next show that, if H is non-trivial, then every optimal evidence structure $\mathcal{E}^*$ and every optimal menu $\mathcal{M}^*$ have cardinality $|\mathcal{E}^*(X)| = |\mathcal{M}^*| = \infty$. We first observe that, if $|\mathcal{E}^*(X)| < \infty$, then necessarily $|\mathcal{M}^*| < \infty$. Moreover, letting $\overline{D}^* = \cup_{i \in I} \{y_i\}$, where we write $\mathcal{M}^* = \{(y_i, T(y_i))\}_{i \in I}$, we observe that $|\overline{D}^*| < \infty$. It is therefore necessary only to contradict the optimality of an $\overline{D}$ such that $|\overline{D}| < \infty$. Suppose such an $\overline{D}$ maximized $\mathcal{H}$. Because H is non-trivial, there exists some $\theta^* \in \Theta$ such that $\arg \max_{x \in X} H(x, \theta^*)$ exists, is a singleton, and is on the interior of X . Letting the unique maximizer be $x^*(\theta^*)$, we have that $H_x(x^*(\theta^*), \theta^*) = 0$. As the maximizer is unique and non-degenerate and H is C^2 and strictly supermodular, by application of the implicit function theorem we have that $x^{*'}(\theta^*) > 0$, and x^* is strictly increasing on a neighborhood $I_{\theta^*} \subset \Theta$. Hence, we have that $|x^*(I_{\theta^*})| = \infty$. We can then set $\overline{D}' = \overline{D} \cup x^*(I_{\theta^*})$. Thus, using the fact that the density of types is bounded away from zero and that, for all $\theta \in I_{\theta^*}$, $H(x^*(\theta), \theta) > H(x, \theta)$ for all $x \neq x^*(\theta)$:

$$\mathcal{H}(\overline{D}') - \mathcal{H}(\overline{D}) \geq \int_{I_{\theta^*}} \left(\max_{x \in \overline{D}'} H(x, \theta) - \max_{x \in \overline{D}} H(x, \theta) \right) dF(\theta) > 0. \quad (45)$$

A.3 Proof of Lemma 4

By Lemma 1, we can uniquely represent each regular $\mathcal{E}$ by $(\underline{D}, \overline{D}) \in \underline{\mathcal{D}} \times \overline{\mathcal{D}}$, where $\underline{\mathcal{D}}$ is the collection of closed subsets of X that contain 0. Moreover, by Lemma 2, it is without loss of optimality to set $\underline{D} = \{0\}$. Finally, in Lemma 10, we defined the value function $\mathcal{H} : \overline{\mathcal{D}} \rightarrow \mathbb{R}$ such that $\mathcal{R}(\mathcal{E}) = \mathcal{H}(\overline{D}_{\mathcal{E}})$. Thus, it suffices to show that $\sup_{\overline{D} \in \overline{\mathcal{D}}} \mathcal{H}(\overline{D}) - \tilde{\Gamma}_f(\overline{D})$ admits a solution. Compactness of $\overline{\mathcal{D}}$ follows from the stronger statement in Lemma 13 that $\underline{\mathcal{D}} \times \overline{\mathcal{D}}$ is compact in the Hausdorff topology. We next show continuity of $\mathcal{H} : \overline{\mathcal{D}} \rightarrow \mathbb{R}$. By applying Berge's maximum theorem to the program $\tilde{\mathcal{H}}(\overline{D}, \theta) = \max_{x \in \overline{D}} H(x, \theta)$, the map $(\overline{D}, \theta) \mapsto \tilde{\mathcal{H}}(\overline{D}, \theta)$ is continuous in the appropriate product topology. Moreover, $\tilde{\mathcal{H}}$ is bounded, as u

(and its first derivatives), π , and g_b are bounded. Thus, the dominated convergence theorem implies that $\overline{D} \mapsto \mathcal{H}(\overline{D})$ is continuous in the Hausdorff topology. Finally, continuity of $\tilde{\Gamma}_f$ follows from Proposition 3. Thus, by Weierstrass' Theorem, the problem admits a solution.

A.4 Proof of Lemma 5

Let ϕ_z^* and ϕ^* denote the optimal allocations under $\overline{D} \cup \{z\}$ and $\overline{D}$. By Lemma 10:

$$\mathcal{H}(\overline{D} \cup \{z\}) - \mathcal{H}(\overline{D}) = \int_{\Theta} (H(\phi_z^*(\theta), \theta) - H(\phi^*(\theta), \theta)) dF(\theta). \quad (46)$$

We define $\overline{\phi}(\theta) = \min\{x \in \overline{D} : x \geq \phi^P(\theta)\}$ and $\underline{\phi}(\theta) = \max\{x \in \overline{D} : x \leq \phi^P(\theta)\}$. As H is strictly quasi-concave, the solution of $\max_{x \in \overline{D}} H(x, \theta)$ is either $\overline{\phi}(\theta)$ or $\underline{\phi}(\theta)$. Define $\Delta H(\theta) = H(\overline{\phi}(\theta), \theta) - H(\underline{\phi}(\theta), \theta)$ for all $\theta \in \Theta$. When $\Delta H(\theta) > 0$, it is $\overline{\phi}(\theta)$. When $\Delta H(\theta) < 0$, it is $\underline{\phi}(\theta)$. When $\Delta H(\theta) = 0$, either is optimal. Further, define $\hat{\theta} : \phi^P(\Theta)^2 \rightarrow \Theta$ as the type for which the principal is indifferent between assigning x or y , *i.e.*, the unique solution to $\frac{H(x, \theta) - H(y, \theta)}{x - y} = 0$ when $x \neq y$ and as the unique solution to $H_x(x, \theta) = 0$ when $x = y$. This function is well defined (because $x, y \in \phi^P(\Theta)$), symmetric, and jointly continuous. Moreover, by the implicit function theorem and the fact that H is twice continuously differentiable and strictly supermodular, it follows that for all $x \in \phi^P(\Theta)$, $\hat{\theta}(x, \cdot)$ is continuously differentiable over $[\phi^P(0), x]$. Hence, the following assignment function is optimal:

$$\phi_z^*(\theta) = \begin{cases} z & , \theta \in (\hat{\theta}(a, z), \hat{\theta}(z, b)), \\ \phi^*(\theta) & , \text{otherwise.} \end{cases} \quad (47)$$

Further noting that $\phi^*(\theta) = a$ for $\theta \in (\hat{\theta}(a, z), \hat{\theta}(a, b))$ and $\phi^*(\theta) = b$ for $\theta \in (\hat{\theta}(a, b), \hat{\theta}(z, b))$,

$$\mathcal{H}(\overline{D} \cup \{z\}) - \mathcal{H}(\overline{D}) = \int_{\hat{\theta}(a, z)}^{\hat{\theta}(a, b)} (H(z, \theta) - H(a, \theta)) dF(\theta) + \int_{\hat{\theta}(a, b)}^{\hat{\theta}(z, b)} (H(z, \theta) - H(b, \theta)) dF(\theta). \quad (48)$$

We now proceed in three steps to bound this difference.

Step 1: Bounding the Integrand. We first derive an upper bound for $|H(x, \theta) - H(y, \theta)|$, for any $x, y \in [a, b]$ and $\theta \in (\hat{\theta}(a, z), \hat{\theta}(z, b))$. Using Taylor's remainder theorem:

$$H(x, \theta) = H(y, \theta) + H_x(y, \theta)(x - y) + \frac{1}{2} H_{xx}(w(\theta), \theta)(x - y)^2, \quad (49)$$

for some $w(\theta) \in [x, y] \cup [y, x]$. We further apply Taylor's remainder theorem to $H_x(y, \theta)$:

$$H_x(y, \theta) = H_x(\phi^P(\theta), \theta) + H_{xx}(v(\theta), \theta)(y - \phi^P(\theta)) = H_{xx}(v(\theta), \theta)(y - \phi^P(\theta)), \quad (50)$$

where the first equality defines the point $v(\theta) \in [y, \phi^P(\theta)] \cup [\phi^P(\theta), y]$ and the second uses $H_x(\phi^P(\theta), \theta) = 0$, since ϕ^P is an interior maximizer of H . Combining these expansions:

$$\begin{aligned} |H(x, \theta) - H(y, \theta)| &\leq |H_x(y, \theta)||x - y| + \frac{1}{2}|H_{xx}(w(\theta), \theta)|(x - y)^2 \\ &\leq |H_{xx}(v(\theta), \theta)||y - \phi^P(\theta)||x - y| + \frac{1}{2}|H_{xx}(w(\theta), \theta)|(x - y)^2 \\ &\leq \bar{H}_{xx} \left(|y - \phi^P(\theta)||x - y| + \frac{1}{2}(x - y)^2 \right). \end{aligned} \quad (51)$$

We now apply this result to bound the integrands:

$$\begin{aligned} H(z, \theta) - H(a, \theta) &\leq |H(z, \theta) - H(a, \theta)| \leq \bar{H}_{xx} \left(|z - \phi^P(\theta)||a - z| + \frac{1}{2}(a - z)^2 \right) \\ &\leq \bar{H}_{xx} \left((b - a)(z - a) + \frac{1}{2}(z - a)^2 \right), \end{aligned} \quad (52)$$

where in the final line we use the inequality $|z - \phi^P(\theta)| \leq b - a$, which in turn follows from the fact that $z, \phi^P(\theta) \in [a, b]$. By applying the same arguments to the second integrand:

$$H(z, \theta) - H(b, \theta) \leq \bar{H}_{xx} \left((b - a)(b - z) + \frac{1}{2}(b - z)^2 \right). \quad (53)$$

Step 2: Bounding the Limits of Integration. By the implicit function theorem:

$$\hat{\theta}_y(x, y) = \frac{H_x(y, \hat{\theta}(x, y))}{H_\theta(x, \hat{\theta}(x, y)) - H_\theta(y, \hat{\theta}(x, y))}, \quad (54)$$

for all $x, y \in \phi^P(\Theta)$ such that $y < x$. For the denominator, we observe that:

$$H_\theta(x, \hat{\theta}(x, y)) - H_\theta(y, \hat{\theta}(x, y)) = \int_y^x H_{x\theta}(t, \hat{\theta}(x, y)) dt \geq \underline{H}_{x\theta}(x - y). \quad (55)$$

For the numerator, we apply Taylor's remainder theorem:

$$H(x, \hat{\theta}(x, y)) = H(y, \hat{\theta}(x, y)) + H_x(y, \hat{\theta}(x, y))(x - y) + \frac{1}{2}H_{xx}(w, \hat{\theta}(x, y))(x - y)^2, \quad (56)$$

for some $w \in [y, x]$, and recall that $H(x, \hat{\theta}(x, y)) = H(y, \hat{\theta}(x, y))$. From this, we obtain that:

$$H_x(y, \hat{\theta}(x, y)) = -\frac{1}{2}H_{xx}(w, \hat{\theta}(x, y))(x - y) \leq \frac{1}{2}\bar{H}_{xx}(x - y). \quad (57)$$

Combining these facts, we obtain that:

$$\hat{\theta}_y(x, y) \leq \frac{\frac{1}{2} \bar{H}_{xx}(x - y)}{\bar{H}_{x\theta}(x - y)} = \frac{1}{2} \frac{\bar{H}_{xx}}{\bar{H}_{x\theta}}. \quad (58)$$

A symmetric argument yields the same bound when $x < y$. By the mean-value theorem:

$$\hat{\theta}(a, b) - \hat{\theta}(a, z) \leq \frac{1}{2} \frac{\bar{H}_{xx}}{\bar{H}_{x\theta}}(b - z). \quad (59)$$

By symmetry of $\hat{\theta}$, we have $\hat{\theta}(z, b) - \hat{\theta}(a, b) = \hat{\theta}(b, z) - \hat{\theta}(b, a)$, obtaining the analogous bound.

Step 3: Bounding the Value. Combining steps 1 and 2, we conclude that

$$\begin{aligned} \mathcal{H}(\bar{D} \cup \{z\}) - \mathcal{H}(\bar{D}) &= \int_{\hat{\theta}(a, z)}^{\hat{\theta}(a, b)} (H(z, \theta) - H(a, \theta)) dF(\theta) + \int_{\hat{\theta}(a, b)}^{\hat{\theta}(z, b)} (H(z, \theta) - H(b, \theta)) dF(\theta) \\ &\leq \bar{H}_{xx} \left(\int_{\hat{\theta}(a, z)}^{\hat{\theta}(a, b)} \left((b - a)(z - a) + \frac{1}{2}(z - a)^2 \right) dF(\theta) + \int_{\hat{\theta}(a, b)}^{\hat{\theta}(z, b)} \left((b - a)(b - z) + \frac{1}{2}(b - z)^2 \right) dF(\theta) \right) \\ &\leq \bar{H}_{xx} \bar{f} \left(\left(\hat{\theta}(a, b) - \hat{\theta}(a, z) \right) \left((b - a)(z - a) + \frac{1}{2}(z - a)^2 \right) + \left(\hat{\theta}(z, b) - \hat{\theta}(a, b) \right) \left((b - a)(b - z) + \frac{1}{2}(b - z)^2 \right) \right) \\ &\leq \frac{1}{2} \frac{\bar{H}_{xx}^2 \bar{f}}{\bar{H}_{x\theta}} \left((b - z) \left((b - a)(z - a) + \frac{1}{2}(z - a)^2 \right) + (z - a) \left((b - a)(b - z) + \frac{1}{2}(b - z)^2 \right) \right) \\ &= \frac{1}{2} \frac{\bar{H}_{xx}^2 \bar{f}}{\bar{H}_{x\theta}} (b - z)(z - a) \left((b - a) + \frac{1}{2}(z - a) + (b - a) + \frac{1}{2}(b - z) \right) \leq \frac{3}{2} \frac{\bar{H}_{xx}^2 \bar{f}}{\bar{H}_{x\theta}} (b - a)(b - z)(z - a). \end{aligned} \quad (60)$$

A.5 Proof of Lemma 6

From Lemma 1, $C_{\mathcal{E}}(x) = [\underline{\delta}_{\mathcal{E}}(x), \bar{\delta}_{\mathcal{E}}(x)]$. Thus,

$$\Gamma_f(\mathcal{E}) = \int_X \left(\int_0^{\bar{\delta}_{\mathcal{E}}(x)} g_f(x, q) dq + \int_{\bar{\delta}_{\mathcal{E}}(x)}^{\bar{x}} g_f(x, q) dq \right) dx. \quad (61)$$

Moreover, by Lemma 2, we can restrict attention to the case in which $\underline{\delta}_{\mathcal{E}} = 0$. We therefore define the induced cost over self-enforcing recommendations, $\tilde{\Gamma}_f : \bar{\mathcal{D}} \rightarrow \mathbb{R}$, as

$$\tilde{\Gamma}_f(\bar{D}) = \int_X \int_{\bar{\delta}_{\bar{D}}(x)}^{\bar{x}} g_f(x, q) dq dx, \quad (62)$$

where we define $\bar{\delta}_{\bar{D}}(x) = \min_{q \geq x: q \in \bar{D}} q$. Recall that $\bar{D} \cap (a, b) = \emptyset$ and fix $z \in (a, b)$. We have

$$\bar{\delta}_{\bar{D} \cup \{z\}}(x) = z \mathbb{I}[x \in (a, z)] + \bar{\delta}_{\bar{D}}(x) (1 - \mathbb{I}[x \in (a, z)]) \quad (63)$$

and, moreover, $\bar{\delta}_{\bar{D}}(x) = b$ for all $x \in (a, b]$. Thus,

$$\tilde{\Gamma}_f(\bar{D} \cup \{z\}) - \tilde{\Gamma}_f(\bar{D}) = \int_a^z \left(\int_z^{\bar{x}} g_f(x, q) dq - \int_b^{\bar{x}} g_f(x, q) dq \right) dx = \int_a^z \int_z^b g_f(x, q) dq dx. \quad (64)$$

The inequality follows by observing that $g_f(x, q) \geq \underline{g}_f := \min_{x, q \in X} g_f(x, q) > 0$, and hence

$$\tilde{\Gamma}_f(\bar{D} \cup \{z\}) - \tilde{\Gamma}_f(\bar{D}) \geq \int_a^z \int_z^b \underline{g}_f dq dx = \underline{g}_f(b - z)(z - a). \quad (65)$$

A.6 Proof of Lemma 7

We first establish the suboptimality of any $\bar{D}$ with the properties described in the statement. Define $S = \bar{D} \cap [a, b]$. Because S is closed, by Corollary 3.90 in Aliprantis and Border (2006) there exists a sequence $\{S_n\}_{n \in \mathbb{N}}$ of finite subsets of S converging to S in the Hausdorff topology. Because $\bar{D} \cap (a, b) \neq \emptyset$, without loss of generality, one can take this sequence such that $S_1 = \{a, b, x_1\}$ and $S_{n+1} = S_n \cup \{x_{n+1}\}$ for some sequence of points $\{x_n\}_{n \in \mathbb{N}}$ in $S \setminus \{a, b\}$. For all $n \in \mathbb{N}$, define $\bar{D}_n = (\bar{D} \setminus (a, b)) \cup S_n$ and observe that each of these sets is closed and that $\bar{D}_n$ converges to $(\bar{D} \setminus (a, b)) \cup S = \bar{D}$ in the Hausdorff topology. Moreover, define $V_n = \mathcal{H}(\bar{D}_n) - \tilde{\Gamma}_f(\bar{D}_n)$ for all n . Moreover, $\mathcal{H}$ and $\tilde{\Gamma}_f$ are continuous in the Hausdorff topology (see Lemma 4 and Proposition 3, respectively), and so $V_n \rightarrow \mathcal{H}(\bar{D}) - \tilde{\Gamma}_f(\bar{D})$.

We next prove that, for all $n \in \mathbb{N}$, we have $V_{n-1} > V_n$, where $S_0 = \{a, b\}$, $\bar{D}_0 = \bar{D} \setminus (a, b)$, and $V_0 = \mathcal{H}(\bar{D}_0) - \tilde{\Gamma}_f(\bar{D}_0)$. Fix any $n \in \mathbb{N}$ and observe that there exist two points $a_n, b_n \in S_{n-1}$ such that $0 < b_n - a_n < \frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 \bar{f}}$, and $S_n \cap (a_n, b_n) = \{x_n\}$. In turn, this implies that $a_n, b_n \in \bar{D}_{n-1} \cap \phi^P(\Theta)$ and that $\bar{D}_{n-1} \cap (a_n, b_n) = \emptyset$. Combining Lemmas 5 and 6, with $a = a_n$, $b = b_n$, and $z = x_n$, we have

$$\begin{aligned} V_n - V_{n-1} &= \left(\mathcal{H}(\bar{D}_{n-1} \cup \{x_n\}) - \tilde{\Gamma}_f(\bar{D}_{n-1} \cup \{x_n\}) \right) - \left(\mathcal{H}(\bar{D}_{n-1}) - \tilde{\Gamma}_f(\bar{D}_{n-1}) \right) \\ &\leq \frac{3}{2} \frac{\bar{H}_{xx}^2 \bar{f}}{H_{x\theta}} (b_n - a_n)(b_n - x_n)(x_n - a_n) - \underline{g}_f(b_n - x_n)(x_n - a_n) \\ &= \underline{g}_f(b_n - x_n)(x_n - a_n) \left(\frac{3}{2} \frac{\bar{H}_{xx}^2 \bar{f}}{g_f H_{x\theta}} (b_n - a_n) - 1 \right) < 0, \end{aligned} \quad (66)$$

where the very last inequality observes that $b_n - a_n \leq b - a < \frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 \bar{f}}$. By passing to the limit, we thus have $\mathcal{H}(\bar{D}_0) - \tilde{\Gamma}_f(\bar{D}_0) = V_0 > V_1 > \dots > V_n > \dots \geq \lim_n V_n = \mathcal{H}(\bar{D}) - \tilde{\Gamma}_f(\bar{D})$. Thus, as desired, a $\bar{D}$ with the maintained property is not optimal.

We finally prove the finiteness of any optimal $\bar{D}$ as well as the bound on its cardinality. Fix an optimal set of self-enforcing recommendations $\bar{D}^* \in \bar{\mathcal{D}}$; at least one such $\bar{D}^*$ exists according to Lemma 4. All $\bar{D} \in \bar{\mathcal{D}}$ contain 0 and $\bar{x}$. We now establish that $\bar{D}^*$ can contain

at most one point in $(0, \phi^P(0)]$ and at most one point in $[\phi^P(1), \bar{x}]$. Toward a contradiction, suppose that there are two points $a, b \in \bar{D}^* \cap (0, \phi^P(0)]$ such that $a < b$. As H is strictly quasi-concave and strictly supermodular, we know that $H(x, \theta)$ is increasing in x for all $x \in (0, \phi^P(0)]$ and $\theta \in \Theta$. Hence, $H(b, \theta) > H(a, \theta)$ for all $\theta \in \Theta$. It follows that $\mathcal{H}(\bar{D}^*) - \mathcal{H}(\bar{D}^* \setminus (0, b)) = 0$. At the same time, by Lemma 6, we have that $\tilde{\Gamma}_f(\bar{D}^*) - \tilde{\Gamma}_f(\bar{D}^* \setminus (0, b)) \geq g_f(b - a)a$. Thus, it is a strict improvement to modify $\bar{D}^*$ to $\bar{D}^* \setminus (0, b)$, a contradiction. An entirely symmetric argument for points above $\phi^P(1)$ establishes that there can be at most one element $a \in \bar{D}^*$ such that $a \in [\phi^P(1), \bar{x}]$. Thus, $|\bar{D}^* \cap ([0, \phi^P(0)] \cup [\phi^P(1), \bar{x}])| \in \{2, 3, 4\}$.

By Lemma 7, we know that $\bar{D}^* \cap \phi^P(\Theta)$ cannot contain two points a and b that are closer than $\frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 f}$ and such that there exists a $z \in \bar{D}^* \cap (a, b)$. From this, it is immediate that $\bar{D}^* \cap \phi^P(\Theta)$ is a finite set. We let $L = |\bar{D}^* \cap \phi^P(\Theta)|$ and enumerate the L points as $\{x_1, \dots, x_L\}$ with $\phi^P(0) \leq x_1 < x_2 < \dots < x_k < \dots < x_L \leq \phi^P(1)$. Applying Lemma 7, we have that $x_{k+1} - x_{k-1} \geq \frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 f}$. Under this constraint, it is possible to fit at most $2 + 2 \left\lfloor \bar{x} / \left(\frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 f} \right) \right\rfloor$ points between any subset of $[0, \bar{x}]$. To see this constructively, for odd k set $x_1 = 0$ and then set $x_{2n+1} = x_{2n-1} + \frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 f}$ until $x_{2n+1} > \bar{x}$. For even k , start by setting $x_2 = \eta$, for some small $\eta > 0$ and then set $x_{2n} = x_{2n-2} + \frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 f}$ until $x_{2n} > \bar{x}$. This construction ensures that the feasibility constraint binds for all k , both even and odd. The even indexed points hit $\bar{x}$ first (unless $x_{2k+1} = \bar{x}$ for some k , in which you hit $\bar{x}$ one index sooner). This construction allows for at most $2 \left(1 + \left\lfloor \bar{x} / \left(\frac{2}{3} \frac{g_f H_{x\theta}}{H_{xx}^2 f} \right) \right\rfloor \right) \leq 2 + \left\lfloor 3 \frac{\bar{x} H_{xx}^2 f}{g_f H_{x\theta}} \right\rfloor$ points.

References

- AL-NAJJAR, N. I., L. ANDERLINI, AND L. FELLI (2006): “Undescribable events,” *The Review of Economic Studies*, 73(4), 849–868.
- ALIPRANTIS, C. D., AND K. BORDER (2006): *Infinite dimensional analysis: A Hitchhiker’s Guide*. Springer.
- ALONSO, R., AND N. MATOUSCHEK (2008): “Optimal delegation,” *The Review of Economic Studies*, 75(1), 259–293.
- AMADOR, M., K. BAGWELL, AND L. CARPIZO (2025): “Generalized Delegation,” Working paper.
- ANDERLINI, L., AND L. FELLI (1994): “Incomplete written contracts: Undescribable states of nature,” *The Quarterly Journal of Economics*, 109(4), 1085–1124.
- (1999): “Incomplete contracts and complexity costs,” *Theory and decision*, 46(1), 23–50.
- ASANUMA, B. (1989): “Manufacturer-supplier relationships in Japan and the concept of relation-specific skill,” *Journal of the Japanese and International Economies*, 3(1), 1–30.
- BAJARI, P., AND S. TADELIS (2001): “Incentives versus transaction costs: A theory of procurement contracts,” *RAND Journal of Economics*, pp. 387–407.
- BALL, I., AND D. KATTWINKEL (2019): “Probabilistic verification in mechanism design,” *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 389–390.
- BARON, D. P., AND R. B. MYERSON (1982): “Regulating a monopolist with unknown costs,” *Econometrica*, pp. 911–930.

- BATTIGALLI, P., AND G. MAGGI (2002): “Rigidity, discretion, and the costs of writing contracts,” *American Economic Review*, 92(4), 798–817.
- BERGEMANN, D., T. HEUMANN, AND S. MORRIS (2025): “Screening with Persuasion,” Working paper.
- BERGEMANN, D., AND M. PESENDORFER (2007): “Information structures in optimal auctions,” *Journal of Economic Theory*, 137(1), 580–609.
- BERGEMANN, D., E. YEH, AND J. ZHANG (2021): “Nonlinear pricing with finite information,” *Games and Economic Behavior*, 130, 62–84.
- BEWLEY, T. F. (1999): *Why Wages Don’t Fall During a Recession*. Harvard University Press.
- BOLTON, P., AND M. DEWATRIPONT (2004): *Contract Theory*. MIT Press, Cambridge, MA.
- BULL, J., AND J. WATSON (2007): “Hard evidence and mechanism design,” *Games and Economic Behavior*, 58(1), 75–93.
- CHE, Y.-K., AND D. B. HAUSCH (1999): “Cooperative investments and the value of contracting,” *American Economic Review*, 89(1), 125–147.
- COASE, R. H. (1960): “The problem of social cost,” *The Journal of Law and Economics*, 3, 1–44.
- CORRAO, R., J. P. FLYNN, AND K. A. SASTRY (2023): “Nonlinear Pricing with Underutilization: A Theory of Multi-part Tariffs,” *American Economic Review*, 113(3), 836–860.
- DOVIS, A., R. KIRPALANI, AND G. M. SUBLET (2024): “On the Optimal Allocation of Policy-Making,” Discussion paper, National Bureau of Economic Research.
- DYE, R. A. (1985): “Costly contract contingencies,” *International Economic Review*, pp. 233–250.
- EUROPCAR (2024): “Terms and Conditions of Hire,” Accessed January 25, 2025, from https://assets.ctfassets.net/wmdwnw615vg5/26z40cvms7JVQf0c0HgCmn/5876a9c86c9b46cc1af13571f53c1670/CGL_EN_GB.pdf.
- GREEN, J. R., AND J.-J. LAFFONT (1986): “Partially verifiable information and mechanism design,” *The Review of Economic Studies*, 53(3), 447–456.
- GRUBB, M. D. (2009): “Selling to overconfident consumers,” *American Economic Review*, 99(5), 1770–1807.
- HART, O., AND J. MOORE (1999): “Foundations of incomplete contracts,” *The Review of Economic Studies*, 66(1), 115–138.
- (2008): “Contracts as reference points,” *The Quarterly Journal of Economics*, 123(1), 1–48.
- HART, S., I. KREMER, AND M. PERRY (2017): “Evidence games: Truth and commitment,” *American Economic Review*, 107(3), 690–713.
- HOLMSTRÖM, B. (1977): “On incentives and control in organizations,” Ph.D. thesis, Stanford University.
- (1979): “Moral hazard and observability,” *The Bell Journal of Economics*, pp. 74–91.
- HOLMSTRÖM, B., AND P. MILGROM (1987): “Aggregation and linearity in the provision of intertemporal incentives,” *Econometrica*, pp. 303–328.
- JUNG, J., J. H. KIM, F. MATĚJKA, AND C. A. SIMS (2019): “Discrete actions in information-constrained decision problems,” *The Review of Economic Studies*, 86(6), 2643–2667.
- KRÄHMER, D., AND R. STRAUZ (2025): “Unidirectional incentive compatibility,” *Journal of Economic Theory*, p. 106051.
- LAFFONT, J.-J., AND D. MARTIMORT (2009): *The Theory of Incentives: the Principal-Agent Model*. Princeton University Press, Princeton, NJ.
- LLOYD, S. (1982): “Least squares quantization in PCM,” *IEEE transactions on information theory*, 28(2), 129–137.

- LOU, Y. (2022): “Optimal Delegation with Information Manipulation,” Working Paper 4298900, SSRN.
- MASKIN, E., AND J. TIROLE (1999): “Unforeseen contingencies and incomplete contracts,” *The Review of Economic Studies*, 66(1), 83–114.
- MILGROM, P., AND C. SHANNON (1994): “Monotone comparative statics,” *Econometrica*, pp. 157–180.
- MUSSA, M., AND S. ROSEN (1978): “Monopoly and product quality,” *Journal of Economic Theory*, 18(2), 301–317.
- MYERSON, R. B. (1982): “Optimal coordination mechanisms in generalized principal–agent problems,” *Journal of Mathematical Economics*, 10(1), 67–81.
- NEW YORK CITY BAR (2014): “Contract Litigation: Breach of Contract,” NYC Bar Association website, Legal Editor; retrieved from <https://www.nycbar.org/get-legal-help/article/business-and-corporate-law/contract-litigation/>.
- ROCHET, J.-C. (1987): “A necessary and sufficient condition for rationalizability in a quasi-linear context,” *Journal of Mathematical Economics*, 16(2), 191–200.
- SCOTT, R. E., AND G. G. TRIANTIS (2005): “Incomplete contracts and the theory of contract design,” *Case Western Reserve Law Review*, 56, 187.
- SEGAL, I. (1999): “Complexity and renegotiation: A foundation for incomplete contracts,” *The Review of Economic Studies*, 66(1), 57–82.
- SIMON, H. A. (1951): “A formal theory of the employment relationship,” *Econometrica: Journal of the Econometric Society*, pp. 293–305.
- SPIER, K. E. (1992): “Incomplete contracts and signalling,” *The RAND Journal of Economics*, pp. 432–443.
- SQUINTANI, F. (2019): “Optimal Overspecified Contracts,” Working paper.
- THOMSON REUTERS (2024): “Legal drafting challenges, risks and opportunities in a year of transformation,” July 11. Retrieved from: <https://legal.thomsonreuters.com/blog/legal-drafting-challenges-risks-and-opportunities/>.
- TIROLE, J. (1999): “Incomplete contracts: Where do we stand? (Walras-Bowley Lecture),” *Econometrica*, 67(4), 741–781.
- TOIKKA, J. (2011): “Ironing without control,” *Journal of Economic Theory*, 146(6), 2510–2526.
- USDA FOOD SAFETY AND INSPECTION SERVICE (2021): “FSIS Directive 10250.2: Performance Standards – Salmonella Verification Program for Raw Poultry Products,” Technical directive, U.S. Department of Agriculture, Food Safety and Inspection Service, Accessed: September 11, 2026.
- WILLIAMSON, O. E. (1975): “Markets and hierarchies: analysis and antitrust implications: a study in the economics of internal organization,” *University of Illinois at Urbana-Champaign’s Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship*.
- WILSON, R. B. (1989): “Efficient and competitive rationing,” *Econometrica*, pp. 1–40.
- (1993): *Nonlinear Pricing*. Oxford University Press, New York.

Supplemental Appendix to *Contractibility Design* by Corrao, Flynn, and Sastry

B Omitted Proofs

B.1 Proof of Lemma 2

In Lemma 12 in Appendix C.1, we show that a regular contractibility correspondence C can be equivalently represented via (i) two sets of self-enforcing recommendations $\underline{D}, \overline{D}$ (see Lemma 1 in the main text); or via (ii) $C(y) = [\underline{\delta}(y), \overline{\delta}(y)]$ for all $y \in X$, where $\underline{\delta} : X \rightarrow X$ is an upper semi-continuous increasing function and $\overline{\delta} : X \rightarrow X$ is a lower semi-continuous increasing function such that: (i) $y \in [\underline{\delta}(y), \overline{\delta}(y)]$, (ii) $\underline{\delta}(x) = \underline{\delta}(y)$ for all $x \in [\underline{\delta}(y), y)$, (iii) $\overline{\delta}(x) = \overline{\delta}(y)$ for all $x \in (y, \overline{\delta}(y)]$, and (iv) $\overline{\delta}(0) = 0$.

There are only two $\underline{\delta}$ functions satisfying the properties of Lemma 1 and such that $\underline{\delta}(X) \subseteq \{0, \bar{x}\}$: (i) $\underline{\delta}(x) = 0$ and (ii) $\underline{\delta}(x) = \bar{x}\mathbb{I}[x = \bar{x}]$. Assume by contradiction that $\underline{D} \not\subseteq \{0, \bar{x}\}$, and therefore $\underline{\delta}$ does not correspond to one of these two functions. Thus, there exists an $x_0 \in (0, \bar{x}]$ such that $\underline{\delta}(x_0) \in (0, \bar{x})$. Consider replacing this $\underline{\delta}$ with $\underline{\delta}' = 0$, defining a corresponding new evidence structure $\mathcal{E}'$. In Lemma 9 in Appendix A.1, we show that implementable actions and conditionally optimal tariffs depend only on the upper self-enforcing recommendations, summarized by the set $\overline{D}$ and function $\overline{\delta}$. Therefore, it suffices to consider the effect on front-end and back-end costs. For front-end costs, we observe (see the proof of Lemma 6 in Appendix A.5) that

$$\Gamma_f(\mathcal{E}) = \int_X \left[\int_0^{\underline{\delta}_{\mathcal{E}}(x)} g_f(x, z) dz + \int_{\overline{\delta}_{\mathcal{E}}(x)}^{\bar{x}} g_f(x, z) dz \right] dx. \quad (67)$$

Considered as a functional of $\underline{\delta}, \overline{\delta}$, Γ_f is monotone. Moreover, $\underline{\delta}_{\mathcal{E}} \geq \underline{\delta}_{\mathcal{E}'}$. Thus, $\Gamma_f(\mathcal{E}) \geq \Gamma_f(\mathcal{E}')$. Moreover, this inequality is clearly strict if $g_f > 0$, as $\underline{\delta}_{\mathcal{E}} > \underline{\delta}_{\mathcal{E}'}$ for a positive measure of recommendations.

Similarly, for the back-end costs, we observe (see the proof of Lemma 10 in Appendix A.1) that

$$\Gamma_b(\mathcal{E}, Q) = \int_X \left[\int_0^{\underline{\delta}_{\mathcal{E}}(x)} g_b(x, z) dz + \int_{\overline{\delta}_{\mathcal{E}}(x)}^{\bar{x}} g_b(x, z) dz \right] dQ(x) \quad (68)$$

and an analogous argument applies to show that $\Gamma_b(\mathcal{E}, Q) \geq \Gamma_b(\mathcal{E}', Q)$, for all Q . This suffices to show that the variant evidence structure $\mathcal{E}'$ implements the same outcomes and payments at lower front-end and back-end costs, completing the proof.

B.2 Proof of Lemma 8

(Only if) We begin by proving the necessity of the existence of a monotone tariff with respect to C . Suppose that ϕ is implementable. It follows that there exists (ξ, T) that supports ϕ . In particular, observe that (O) implies that $\phi(\theta) \in C(\xi(\theta))$ for all $\theta \in \Theta$. Next define $\hat{T} : X \rightarrow \bar{\mathbb{R}}$ as:

$$\hat{T}(x) = \inf_{y \in X} \{T(y) : x \in C(y)\}. \quad (69)$$

We next show that ϕ is also supported by $(\phi, \hat{T})$. By (O) of (ϕ, ξ, T) , we have $u(\phi(\theta), \theta) \geq u(x, \theta)$ for all $x \in C(\phi(\theta)) \subseteq C(\xi(\theta))$ (by transitivity) and for all $\theta \in \Theta$, yielding (O) of $(\phi, \phi, \hat{T})$. By (IR) of (ϕ, ξ, T) and the definition of $\hat{T}$, we have

$$u(\phi(\theta), \theta) - \hat{T}(\phi(\theta)) \geq u(\phi(\theta), \theta) - T(\xi(\theta)) \geq 0, \quad (70)$$

for all $\theta \in \Theta$, yielding (IR) of $(\phi, \phi, \hat{T})$. Next, assume toward a contradiction that $(\phi, \phi, \hat{T})$ does not satisfy (IC), that is, there exist $\theta \in \Theta$ and $y \in X$ such that

$$\max_{x \in C(y)} u(x, \theta) - \hat{T}(y) > u(\phi(\theta), \theta) - \hat{T}(\phi(\theta)). \quad (71)$$

By the definition of $\hat{T}$, there exists a sequence $z_n \in X$ such that $y \in C(z_n)$ for all n and $T(z_n) \downarrow \hat{T}(y)$. Thus, there exists n large enough such that

$$\begin{aligned} \max_{x \in C(z_n)} u(x, \theta) - T(z_n) &\geq \max_{x \in C(y)} u(x, \theta) - T(z_n) > u(\phi(\theta), \theta) - \hat{T}(\phi(\theta)) \\ &\geq u(\phi(\theta), \theta) - T(\xi(\theta)) = \max_{x \in C(\xi(\theta))} u(x, \theta) - T(\xi(\theta)). \end{aligned} \quad (72)$$

The first inequality follows from $C(y) \subseteq C(z_n)$ since $y \in C(z_n)$. The second strict inequality follows from Equation 71 and the fact that $T(z_n) \downarrow \hat{T}(y)$. The third inequality follows from the construction of $\hat{T}$. The final equality follows as (ϕ, ξ, T) satisfies (O). However, the previous inequality yields a contradiction of (IC) of (ϕ, ξ, T) , proving that $(\phi, \phi, \hat{T})$ satisfies (IC). This shows that $(\phi, \phi, \hat{T})$ is implementable, hence that Equation 32 holds and that $u(\phi(\theta), \theta) - \hat{T}(\phi(\theta)) \geq 0$ for all $\theta \in \Theta$.

Finally, we argue that $\hat{T}$ is monotone with respect to C . Fix $x, y \in X$ such that $y \in C(x)$. By Transitivity of C we have $\{\hat{x} \in X : x \in C(\hat{x})\} \subseteq \{\hat{x} \in X : y \in C(\hat{x})\}$, yielding that $\hat{T}(y) \leq \hat{T}(x)$, as desired.

(If) We now establish sufficiency. Suppose that there exists a tariff $T : X \rightarrow \bar{\mathbb{R}}$ that is monotone with respect to C and such that Equation 32 holds and $u(\phi(\theta), \theta) - T(\phi(\theta)) \geq 0$ for all $\theta \in \Theta$. We will show that (ϕ, ϕ, T) is implementable. (IR) is immediately satisfied.

Next, we show that (IC) is satisfied. Suppose, toward a contradiction, that it were not. That is, there exist $\theta \in \Theta$, $y \in X$, and $x \in C(y)$ such that

$$u(x, \theta) - T(y) > \max_{\hat{x} \in C(\phi(\theta))} u(\hat{x}, \theta) - T(\phi(\theta)) \geq u(\phi(\theta), \theta) - T(\phi(\theta)). \quad (73)$$

But then, we have the following contradiction of monotonicity of T in C :

$$u(x, \theta) - T(y) > u(\phi(\theta), \theta) - T(\phi(\theta)) \geq u(x, \theta) - T(x) \implies T(x) > T(y), \quad (74)$$

where the second inequality uses the fact that $\phi(\theta)$ solves the program in Equation 32. Finally, we show that (O) is satisfied. Toward a contradiction, assume that there exists $\theta \in \Theta$ and $x \in C(\phi(\theta))$ such that $u(x, \theta) > u(\phi(\theta), \theta)$. However, by monotonicity of T in C , we know that $T(\phi(\theta)) \geq T(x)$. Thus, $u(x, \theta) - T(x) > u(\phi(\theta), \theta) - T(\phi(\theta))$, yielding a contradiction to IC, which we just showed. This proves sufficiency.

Finally, the fact that any implementable final action function can be implemented as part of an allocation (ϕ, ϕ, T) follows by the construction in the necessity part of our proof.

B.3 Proof of Proposition 1

By Corollary 1, there always exists an optimal menu with finite actions, supported by a set of self-enforcing recommendations $\bar{D} = \{x_1, \dots, x_K\}$ for some $K \in \mathbb{N}$ and such that $x_1 < x_2 < \dots < x_K$. Throughout, we adopt the notational convention that $x_0 = 0$ and $x_{K+1} = \bar{x}$.

We next derive the optimal allocation and payments. As shown in the last step of the proof of Lemma 5, for any given $\bar{D}$,

$$\phi^*(\theta) = \begin{cases} \bar{\phi}(\theta) = \min\{x \in \bar{D} : x \geq \phi^P(\theta)\}, & \text{if } H(\bar{\phi}(\theta), \theta) > H(\underline{\phi}(\theta), \theta), \\ \underline{\phi}(\theta) = \max\{x \in \bar{D} : x \leq \phi^P(\theta)\}, & \text{otherwise.} \end{cases} \quad (75)$$

We now specialize this using the structure of $\bar{D}$. As H is strictly single-crossing, $H(x_k, \theta) - H(x_{k-1}, \theta) = 0$ has no solution if and only if (i) $H(x_k, 0) - H(x_{k-1}, 0) > 0$ or (ii) $H(x_k, 1) - H(x_{k-1}, 1) < 0$. As H is strictly quasi-concave, if $H(x_k, 0) - H(x_{k-1}, 0) > 0$, then $H(\cdot, 0)$ is strictly increasing at x_{k-1} , and therefore at all x_j for $j \leq k-1$. Thus, if $H(x_k, 0) - H(x_{k-1}, 0) > 0$ holds for k , it holds for all $j \leq k$. Define $\underline{k} = \max\{k \in \{1, \dots, K\} : H(x_k, 0) - H(x_{k-1}, 0) > 0\}$, with the convention that $\underline{k} = 1$ if this set is empty. Similarly, if $H(x_k, 1) - H(x_{k-1}, 1) < 0$, then $H(\cdot, 1)$ is strictly decreasing at x_k . Thus, if $H(x_k, 1) - H(x_{k-1}, 1) < 0$ holds for k , it holds for all $j \geq k$. Define $\bar{k} = \min\{k \in \{1, \dots, K\} : H(x_k, 1) - H(x_{k-1}, 1) < 0\}$, with the convention that $\bar{k} = K+1$ if this set is empty. As H

is strictly single crossing, $\bar{k} > \underline{k}$. We now have that $H(x_k, \theta) - H(x_{k-1}, \theta) = 0$ has a solution if and only if $k \in \{\underline{k} + 1, \dots, \bar{k} - 1\}$ (if $\bar{k} = \underline{k} + 1$, then this set is empty). For all $k \geq \bar{k}$, we have that $\hat{\theta}_k = 1$. For all $k \leq \underline{k}$, we have that $\hat{\theta}_k = 0$. For all $k \in \{\underline{k} + 1, \dots, \bar{k} - 1\}$, we have that $\hat{\theta}_k$ is the unique solution to $H(x_k, \hat{\theta}_k) = H(x_{k-1}, \hat{\theta}_k)$. As H is strictly quasi-concave, we know that $\phi^P(\hat{\theta}_k) \in (x_{k-1}, x_k)$, which implies that $\underline{\phi}(\hat{\theta}_k) = x_{k-1}$ and $\bar{\phi}(\hat{\theta}_k) = x_k$. Thus, we have that $\phi^*(\theta) = x_k$ for all $\theta \in (\hat{\theta}_k, \hat{\theta}_{k+1}]$.

We now derive the tariff. In the strictly quasi-concave case, applying Equation 33 from Lemma 9, we have that:

$$\begin{aligned} T(x_k) &= u(x_k, \hat{\theta}_k) - \mathbb{I}[k \geq 2] \sum_{j=1}^{k-1} \int_{\hat{\theta}_j}^{\hat{\theta}_{j+1}} u_\theta(x_j, s) \, ds \\ &= u(x_1, 0) + \mathbb{I}[k \geq 2] \sum_{j=2}^k \left[u(x_j, \hat{\theta}_j) - u(x_{j-1}, \hat{\theta}_j) \right], \end{aligned} \quad (76)$$

where the equality computes the integral and telescopes the summation. Observing that $x_1 = 0$ and $u(0, 0) = 0$ gives the desired formula. The same formula trivially holds in the strictly quasi-convex case.

We next derive the necessary first-order condition. We first observe that, given a set of self-enforcing recommendations $\{x_1, \dots, x_K\}$, the principal's payoff inclusive of back-end costs can be written as

$$\mathcal{H}(\{x_1, \dots, x_K\}) = \sum_{k=1}^K \int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} H(x_k, \theta) \, dF(\theta) \quad (77)$$

and, similarly, the front-end cost can be written as

$$\begin{aligned} \tilde{\Gamma}_f(\{x_1, \dots, x_K\}) &= \int_X \int_{\bar{\delta}_{\bar{D}}(x)}^{\bar{x}} g_f(x, z) \, dz \, dx = \sum_{k=1}^{K-1} \int_{x_k}^{x_{k+1}} \int_{x_{k+1}}^{\bar{x}} g_f(x, z) \, dz \, dx \\ &= \sum_{k=1}^{K-1} \int_{x_k}^{x_{k+1}} (G_f(\bar{x}, x) - G_f(x_{k+1}, x)) \, dx. \end{aligned} \quad (78)$$

A necessary condition for optimality is that payoffs cannot be locally improved by perturbing any action x_k , for $k \in \{2, \dots, K-1\}$. That is,

$$\frac{d}{d\varepsilon} \mathcal{H}(\{x_1, \dots, x_k + \varepsilon, \dots, x_K\})|_{\varepsilon=0} = \frac{d}{d\varepsilon} \tilde{\Gamma}_f(\{x_1, \dots, x_k + \varepsilon, \dots, x_K\})|_{\varepsilon=0}, \quad (79)$$

where we observe that both derivatives are well-defined. The left-hand side of this is

$$\begin{aligned} \int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} H_x(x_k^*, \theta) dF(\theta) &= \int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} H_x(x_k^*, \theta) dF(\theta) + \frac{\partial}{\partial x_k^*} \hat{\theta}_k \left(H(x_k^*, \hat{\theta}_k) - H(x_{k-1}^*, \hat{\theta}_k) \right) f(\hat{\theta}_k) \\ &\quad + \frac{\partial}{\partial x_k^*} \hat{\theta}_{k+1} \left(H(x_{k+1}^*, \hat{\theta}_{k+1}) - H(x_k^*, \hat{\theta}_{k+1}) \right) f(\hat{\theta}_{k+1}), \end{aligned} \quad (80)$$

where, in the second equality, we use the fact that $H(x_k^*, \hat{\theta}_k) = H(x_{k-1}^*, \hat{\theta}_k)$ when $\hat{\theta}_k \in (0, 1)$ and $\frac{\partial}{\partial x_k^*} \hat{\theta}_k = 0$ when $\hat{\theta}_k \in \{0, 1\}$. Moreover, using the definition $H(x, \theta) = J(x, \theta) - \int_x^{\bar{x}} g_b(x, z) dz$, we observe that:

$$H_x(x, \theta) = J_x(x, \theta) - \int_x^{\bar{x}} g_{b,x}(x, z) dz + g_b(x, x), \quad (81)$$

where the latter two terms do not depend on θ . Thus,

$$\int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} H_x(x_k^*, \theta) dF(\theta) = \int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} J_x(x_k^*, \theta) dF(\theta) + \left(F(\hat{\theta}_{k+1}) - F(\hat{\theta}_k) \right) \left(g_b(x_k, x_k) - \int_{x_k}^{\bar{x}} g_{b,x}(x_k, z) dz \right). \quad (82)$$

The right-hand side of Equation 79 can be written as

$$\begin{aligned} &(G_f(\bar{x}, x_k) - G_f(x_k, x_k)) - \int_{x_{k-1}}^{x_k} g_f(x, x_k) dx - (G_f(\bar{x}, x_k) - G_f(x_{k+1}, x_k)) \\ &= G_f(x_{k+1}, x_k) - G_f(x_k, x_k) - \int_{x_{k-1}}^{x_k} g_f(x, x_k) dx \end{aligned} \quad (83)$$

We obtain the desired Equation 14 by plugging Equations 82 and 83 into Equation 79.

B.4 Proof of Proposition 2

We split the argument into three parts. We first map the problem to one that satisfies the general assumptions of Section 2. We next calculate the optimal contract in the transformed problem. We finally map the allocation and tariff back to the original problem.

Step 1: Transformation of the Problem. We define $x = 1 - e \in X = [0, 1]$ as the agent's shirking and $\theta = 1 - \vartheta \sim F = U[0, 1]$ as their unproductivity. We define transformed preferences for the agent, which satisfy all of our maintained assumptions:

$$u(x, \theta) = \tilde{u}(1 - x, 1 - \theta) - \tilde{u}(1, 1 - \theta) = (\alpha\theta + \beta)x - \beta \frac{x^2}{2}. \quad (84)$$

We define the transformed payoff for the principal,

$$\pi(x) = \tilde{\pi}(1-x) - \tilde{\pi}(1) = \eta(1-x) - \eta = -\eta x. \quad (85)$$

We next rewrite the front-end and back-end costs in terms of shirking. We first observe that an evidence structure over effort induces an evidence structure over shirking defined by $\mathcal{E}^x : X \rightarrow [0, 1]$, where $X = [0, 1]$ and, for each $x \in [0, 1]$, $\mathcal{E}^x(x) = \mathcal{E}(1-x)$. Moreover, $\bar{\delta}_{\mathcal{E}^x}(x) = 1 - \underline{\delta}_{\mathcal{E}}(1-x)$ and $\underline{\delta}_{\mathcal{E}^x}(x) = 1 - \bar{\delta}_{\mathcal{E}}(1-x)$, for each $x \in [0, 1]$. Thus, in some abuse of notation, we can write

$$\Gamma_f(\mathcal{E}) = \tilde{\Gamma}_f(\mathcal{E}^x) = \kappa_f \int_E (\underline{\delta}_{\mathcal{E}^x}(x) + (1 - \bar{\delta}_{\mathcal{E}^x}(x))) dx. \quad (86)$$

An analogous argument applies to the back-end cost.

The (e, ϑ) problem satisfies the assumptions of Lemma 10 in the case of monotone *decreasing* preferences. In this case, the result implies that the principal problem given a set of self-enforcing efforts $\underline{D} \subseteq E$ is

$$\tilde{\mathcal{H}}(\underline{D}) := \max_{\tilde{\phi}} \int_0^1 \tilde{H}(\tilde{\phi}(\vartheta), \vartheta) d\tilde{F}(\vartheta) \text{ s.t. } \tilde{\phi}(\vartheta') \geq \tilde{\phi}(\vartheta), \tilde{\phi}(\vartheta) \in \underline{D}, \vartheta, \vartheta' \in [0, 1] : \vartheta' \geq \vartheta \quad (87)$$

where the modified virtual surplus function is

$$\begin{aligned} \tilde{H}(e, \vartheta) &= \tilde{u}(e, \vartheta) + \tilde{\pi}(e) - \frac{1 - \tilde{F}(\vartheta)}{\tilde{f}(\vartheta)} \tilde{u}_{\vartheta}(e, \vartheta) - \kappa_b e \\ &= (2\alpha\theta + \beta - \eta + \kappa_b)x - \beta \frac{x^2}{2} + \left(\eta - \kappa_b - \frac{\beta}{2} - 2\alpha\theta \right). \end{aligned} \quad (88)$$

Using this substitution, we observe the following mapping to the (x, θ) problem:

$$\begin{aligned} \tilde{\mathcal{H}}(\underline{D}) - \left(\eta - \alpha - \kappa_b - \frac{\beta}{2} \right) &= \mathcal{H}(\overline{D}) = \max_{\phi} \int_0^1 \left((2\alpha\theta + \beta - \eta + \kappa_b)\phi(\theta) - \beta \frac{\phi(\theta)^2}{2} \right) dF(\theta) \\ \text{s.t. } \phi(\theta') &\geq \phi(\theta), \phi(\theta) \in \overline{D}, \quad \theta, \theta' \in [0, 1] : \theta' \geq \theta, \end{aligned} \quad (89)$$

where $\overline{D} = \{(1-e) \in X : e \in \underline{D}\}$. Combining this with the earlier observation about front-end costs (Equation 86), we obtain that the original problem is equivalent to solving the transformed problem with modified surplus $H(x, \theta) = (2\alpha\theta + \beta - \eta + \kappa_b)x - \beta \frac{x^2}{2}$. Moreover, this function is strictly quasi-concave.

Corollary 1 implies that any optimal contractibility correspondence can be represented

by a coarse set of self-enforcing shirking recommendations, $\bar{D} = \{x_1, \dots, x_K\}$. In the transformed problem, these map to a set of self-enforcing effort recommendations $\{1 - x_1, \dots, 1 - x_K\}$. To apply Proposition 1, we must also establish that (without loss of optimality) $x_1 = 0$ and $x_K = \bar{x} = 1$. *Excludability* guarantees that $e_1 = 0$ (zero effort, or full shirking) is perfectly contractible. An additional simple argument is required to establish that (without loss of optimality) we can restrict attention to the case in which $e_K = 1$ (full effort or zero shirking) is also contractible. Imagine it were not, and optimal contractibility were represented by an evidence structure $\mathcal{E}$. Then, construct a variant evidence structure $\mathcal{E}'$ that perturbs $\mathcal{E}$ such that $e = 1$ is self-enforcing. Since the evidence space $\Omega = [0, 1]$ has the cardinality of the continuum, this is always possible by having $e = 1$ generate a unique (singleton) piece of evidence. We observe that $\Gamma_f(\mathcal{E}') = \Gamma_f(\mathcal{E})$ and $\Gamma_b(\mathcal{E}', Q) = \Gamma_b(\mathcal{E}, Q)$ for all Q , as both costs, when re-expressed with $(\underline{\delta}, \bar{\delta})$ as their arguments, are L_1 -continuous. But under $\mathcal{E}'$ there are weakly more self-enforcing recommendations. Thus, $\mathcal{E}'$ obtains higher value. Hence, it is without loss of optimality to restrict attention to contractibility containing $e = 1$ ($x = 0$) as a self-enforcing recommendation. Together, this argument shows that we can apply Proposition 1 to the transformed problem.

Step 2: Optimal Contract in the Transformed Problem. We next apply Proposition 1 to the transformed problem. The modified virtual surplus function follows Equation 18.

We first solve for the candidate optimal contract that solves the variational first-order condition in Proposition 1 for each K . This first-order condition for $k \in \{2, \dots, K - 1\}$ is

$$\int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} (2\alpha\theta + \beta(1 - x_k) - \eta + \kappa_b) d\theta - \kappa_f(-2x_k + x_{k-1} + x_{k+1}) = 0, \quad (90)$$

where $\hat{\theta}_k = \frac{\beta}{4\alpha}(x_k + x_{k-1}) + \frac{\eta - \kappa_b - \beta}{2\alpha}$. Calculating the integral and evaluating at $\hat{\theta}_k$, we write

$$\kappa_f(-2x_k + x_{k-1} + x_{k+1}) = \frac{\beta^2}{16\alpha}(x_{k+1} - x_{k-1})(x_{k+1} + x_{k-1} - 2x_k). \quad (91)$$

This condition can be rearranged to obtain

$$(x_{k+1} + x_{k-1} - 2x_k) \left[\frac{\beta^2}{16\alpha}(x_{k+1} - x_{k-1}) - \kappa_f \right] = 0. \quad (92)$$

This equation has two solutions,

$$x_k = \frac{x_{k+1} + x_{k-1}}{2}, \quad x_{k+1} = x_{k-1} + \Delta, \quad (93)$$

where $\Delta = \frac{16\alpha\kappa_f}{\beta^2}$. We now separately consider three cases.

We first consider case 1, in which all points are on the “uniform grid.” From the boundary conditions, $x_1 = 0$ and $x_K = 1$. Thus, $x_k = \frac{x_{k+1} + x_{k-1}}{2} \implies x_k = \frac{k-1}{K-1}$. We can verify that this is a local maximum by checking that the Hessian is negative definite at this solution. We calculate that:

$$\begin{aligned} \frac{\partial^2(\mathcal{H} - \tilde{\Gamma}_f)}{\partial x_k^2} &= \text{Hess}_{k-1,k-1}^{\mathcal{H}} = -\frac{\beta^2}{4\alpha(K-1)} + 2\kappa_f =: \Lambda, \\ \frac{\partial^2(\mathcal{H} - \tilde{\Gamma}_f)}{\partial x_k \partial x_{k+1}} &= \text{Hess}_{k,k-1}^{\mathcal{H}} = \text{Hess}_{k-1,k}^{\mathcal{H}} = \frac{\beta^2}{8\alpha(K-1)} - \kappa_f = -\frac{1}{2}\Lambda, \end{aligned} \quad (94)$$

where we note that row and column $k-1$ of $\text{Hess}^{\mathcal{H}}$ corresponds to the variable x_k and in the first equality we define Λ . Thus, the Hessian is a tridiagonal Toeplitz matrix, which implies that the Eigenvalues are, by Theorem 2.2 of [Kulkarni, Schmidt, and Tsui \(1999\)](#), given by:

$$\lambda_k = \Lambda \left(1 - \cos \left(\frac{k-1}{K-1} \pi \right) \right) \quad (95)$$

for $k \in \{2, \dots, K-1\}$. As $\cos \left(\frac{k-1}{K-1} \pi \right) \in (-1, 1)$ for all such k , we have that $\text{sgn}(\lambda_k) = \text{sgn}(\Lambda)$. Thus, the Hessian is negative definite if and only if $K < \bar{K} = 1 + \frac{\beta^2}{8\alpha\kappa_f}$. We will later verify that this holds whenever K is set optimally, confirming the optimality of the uniform grid.

We next consider case 2, in which all points are on the “alternating grid.” In this case, even points form a uniform grid with spacing $\Delta = \frac{16\alpha\kappa_f}{\beta^2}$ and the odd points form a uniform grid with spacing $\Delta = \frac{16\alpha\kappa_f}{\beta^2}$. When K is odd, given the boundary conditions that $x_1 = 0$ and $x_K = 1$, we have that this is possible only when $K = 1 + \frac{2}{\Delta}$, which is itself only possible when $\frac{\beta^2}{8\alpha\kappa_f}$ is an even integer. When K is even, the solution must be $x_k = \frac{k-1}{2}\Delta$ for k odd, and $x_k = 1 - \frac{K-k}{2}\Delta$ for k even. This is possible for any even K such that $\frac{2}{\Delta} < K < 2 + \frac{2}{\Delta}$.

We next show that the alternating grid (for odd or even K) is *not* a local maximum of the objective function. For a local maximum, a necessary condition is that the Hessian is negative semidefinite. We will show the existence of a vector $v \in \mathbb{R}^{K-2}$ such that $v \neq 0$ and $v' \text{Hess}^{\mathcal{H}} v > 0$, which implies that $\text{Hess}^{\mathcal{H}}$ is not negative semidefinite. To do this, we first calculate and simplify the second derivatives:

$$\begin{aligned} \frac{\partial^2(\mathcal{H} - \tilde{\Gamma}_f)}{\partial x_k^2} &= \text{Hess}_{k-1,k-1}^{\mathcal{H}} = -\frac{\beta^2}{8\alpha}\Delta + 2\kappa_f = 0, \\ \frac{\partial^2(\mathcal{H} - \tilde{\Gamma}_f)}{\partial x_k \partial x_{k+1}} &= \text{Hess}_{k,k-1}^{\mathcal{H}} = \text{Hess}_{k-1,k}^{\mathcal{H}} = \frac{\beta^2}{8\alpha}(x_{k+1} - x_k) - \kappa_f. \end{aligned} \quad (96)$$

Using this, we define $v_k = z_{k-1} - z_k$, where z_k denotes the unit vector in some dimension

k . This direction corresponds to increasing x_k and decreasing x_{k+1} . We calculate

$$v'_k \text{Hess}^{\mathcal{H}} v_k = 2 \left(\kappa_f - \frac{\beta^2}{8\alpha} (x_{k+1} - x_k) \right). \quad (97)$$

We now split the proof into two cases. First, consider the case in which $K > 4$. In this case, there must exist some x_k, x_{k+1} such that $x_{k+1} - x_k < \frac{\Delta}{2}$, since the grid is not uniform. Then $v'_k \text{Hess}^{\mathcal{H}} v_k > 2 \left(\kappa_f - \frac{\Delta \beta^2}{16\alpha} \right) = 0$ and, as desired, we have shown that the Hessian is not negative definite. Next, we consider the case in which $K = 4$. In this case, we take two candidate vectors. The first is $u = z_1 + z_2$, and we observe $u' \text{Hess}^{\mathcal{H}} u = 2 \left(\frac{\beta^2}{8\alpha} (x_3 - x_2) - \kappa_f \right)$. The second is $v_1 = z_1 - z_2$, and we observe $v'_1 \text{Hess}^{\mathcal{H}} v_1 = 2 \left(\kappa_f - \frac{\beta^2}{8\alpha} (x_3 - x_2) \right) = -u' \text{Hess}^{\mathcal{H}} u$. We finally note that, if $u' \text{Hess}^{\mathcal{H}} u = v'_1 \text{Hess}^{\mathcal{H}} v_1 = 0$, then $x_3 - x_2 = \frac{8\alpha\kappa_f}{\beta^2} = \frac{\Delta}{2}$ and the alternating grid collapses to the uniform grid.

The final case to consider is one in which some points are separated as per the alternating grid and others are separated as per the uniform grid. However, in this case, when points are separated according to the alternating grid, the previous argument on the Hessian applies, contradicting optimality.

We next derive the profit and optimal tariff evaluated at the uniform-grid solution:

Lemma 11. *The value to the monopolist of a K -item contract supported on a uniform grid can be written as $V(K) = \hat{\Pi}(K) - \hat{\Gamma}_f(K)$ where*

$$\hat{\Pi}(K) = \frac{\beta - \eta + \kappa_b + 2\alpha}{4\alpha} (2\alpha - \eta + \kappa_b) + \frac{\beta^2}{48\alpha} \frac{(2K-3)(2K-1)}{(K-1)^2} \quad \text{and} \quad \hat{\Gamma}_f(K) = \frac{\kappa_f}{2} \frac{K-2}{K-1}. \quad (98)$$

Moreover, the optimal allocation is supported by the tariff

$$T^*(x_k) = \frac{1}{2} \frac{k-1}{K^*-1} \left(\beta + \eta - \kappa_b - \frac{\beta}{2} \frac{k-1}{K^*-1} \right). \quad (99)$$

Proof. We observe that the principal's payoff, inclusive of back-end costs (but exclusive of front-end costs), can be written as

$$\mathcal{H}((x_k)_{k=1}^K) = \sum_{k=1}^K \int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} H(x_k, \theta) dF(\theta). \quad (100)$$

Using this, we calculate

$$\begin{aligned}\hat{\Pi}(K) &= \sum_{k=1}^K \int_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} \left((2\alpha\theta + \beta - \eta + \kappa_b)x_k - \beta \frac{x_k^2}{2} \right) d\theta = \sum_{k=1}^K \left[\alpha x_k \theta^2 - x_k \left(\eta - \kappa_b - \beta + \frac{\beta}{2} x_k \right) \theta \right]_{\hat{\theta}_k}^{\hat{\theta}_{k+1}} \\ &= \frac{\beta}{2\alpha(K-1)} \sum_{k=2}^{K-1} \left(\alpha x_k (\hat{\theta}_{k+1} + \hat{\theta}_k) - x_k \left(\eta - \kappa_b - \beta + \frac{\beta}{2} x_k \right) \right) + (1 - \hat{\theta}_K) \left(\alpha(1 + \hat{\theta}_K) + \frac{\beta}{2} - \eta + \kappa_b \right),\end{aligned}\quad (101)$$

where, in the second line, we use that $\hat{\theta}_{k+1} - \hat{\theta}_k = \frac{\beta}{2\alpha(K-1)}$ for $k < K$ and that $\hat{\theta}_{K+1} = 1$ and $x_K = 1$. We simplify the summation term as

$$\sum_{k=2}^{K-1} \left(\alpha x_k (\hat{\theta}_{k+1} + \hat{\theta}_k) - x_k \left(\eta - \kappa_b - \beta + \frac{\beta}{2} x_k \right) \right) = \frac{\beta}{12(K-1)} (K-2)(2K-3), \quad (102)$$

where we use that $\hat{\theta}_k + \hat{\theta}_{k+1} = -\frac{\beta - \eta + \kappa_b}{\alpha} + \frac{\beta}{\alpha} x_k$. To simplify the second term, we observe that $\hat{\theta}_K = \left(\frac{\eta - \kappa_b - \beta}{2\alpha} \right) + \frac{\beta}{4\alpha} \left(1 + \frac{K-2}{K-1} \right)$ and therefore

$$(1 - \hat{\theta}_K) \left(\alpha(\hat{\theta}_K + 1) + \frac{\beta}{2} - \eta + \kappa_b \right) = \frac{\beta - \eta + \kappa_b + 2\alpha}{4\alpha} (2\alpha - \eta + \kappa_b) + \frac{\beta^2}{16\alpha(K-1)^2} (2K-3). \quad (103)$$

Putting this together and simplifying, we write:

$$\hat{\Pi}(K) = \frac{\beta - \eta + \kappa_b + 2\alpha}{4\alpha} (2\alpha - \eta + \kappa_b) + \frac{\beta^2}{48\alpha} \frac{(2K-3)(2K-1)}{(K-1)^2}. \quad (104)$$

We next show the desired representation of $\hat{\Gamma}$ by direct calculation:

$$\hat{\Gamma}(K) = \kappa_f \sum_{k=1}^{K-1} \frac{1}{K-1} \frac{k-1}{K-1} = \frac{\kappa_f}{(K-1)^2} \sum_{k=1}^{K-1} (k-1) = \frac{\kappa_f}{2} \frac{K-2}{K-1}. \quad (105)$$

We finally compute the tariff. Starting from Equation 15,

$$\begin{aligned}T^*(x_k) &= \mathbb{I}[k \geq 2] \sum_{j=2}^k \left[u(x_j, \hat{\theta}_j) - u(x_{j-1}, \hat{\theta}_j) \right] \\ &= \mathbb{I}[k \geq 2] \sum_{j=2}^k \left[(\alpha \hat{\theta}_j + \beta)(x_j - x_{j-1}) + \frac{\beta}{2} (x_{j-1}^2 - x_j^2) \right] \\ &= \frac{1}{2} \frac{k-1}{K^* - 1} \left(\beta + \eta - \kappa_b - \frac{\beta}{2} \frac{k-1}{K^* - 1} \right),\end{aligned}\quad (106)$$

where we substitute in the expressions for $\hat{\theta}_j$ and x_j , simplify at each step, and evaluate at $K = K^*$. $\square$

To derive $\tilde{K}$, we take the first derivative of V :

$$V'(K) = \frac{\beta^2}{24\alpha(K-1)^3} - \frac{\kappa_f}{2(K-1)^2}. \quad (107)$$

We observe that $V'(K) > 0$ if and only if $K < \tilde{K} := \frac{\beta^2}{12\alpha\kappa_f} + 1$. We now prove that $|K^* - \tilde{K}| < 1$. If $K^* - \tilde{K} > 1$, then we know that $V(K^* - 1) > V(K^*)$ as $V' < 0$ for all $K^* - 1 < K < K^*$; this contradicts optimality. Similarly, if $\tilde{K} - K^* > 1$, we know that $V(K^* + 1) > V(K^*)$ as $V' > 0$ for all $K^* < K < K^* + 1$; this contradicts optimality. Recall that we needed to check if the Hessian was negative definite. This is true so long as $K^* \leq \bar{K}$. This holds whenever $\tilde{K} + 1 \leq \bar{K}$, which reduces to $\tilde{K} \geq 3$. It remains to check when $\tilde{K} \in (2, 3)$ and $K^* = 3$. Direct calculation shows that indifference between $K = 2$ and $K = 3$ occurs when $\kappa_f = \frac{\beta^2}{16\alpha}$. At this point, $\tilde{K} = 7/3$. Thus, whenever $K^* > 2$ is strictly optimal (which is when $\kappa_f < \frac{\beta^2}{16\alpha}$), we have that $K^* < \bar{K}$. The comparative statics follow from standard monotone comparative statics arguments, after the observations that $V_{K\alpha} < 0$, $V_{K\beta} > 0$, $V_{K\eta} = 0$, and $V_{K\kappa_b} = 0$.

Step 3: Mapping Back to the Original Problem. Since $e = 1 - x$, we observe that the optimal contract can be supported on $e_k = 1 - x_{K^*-k+1} = \frac{k-1}{K^*-1}$ with K^* self-enforcing recommendations. We next observe that, in the original problem, the IR constraint binds for type $\vartheta = 0$, who always (for any K^*) takes action $x = 1$ or $e = 0$ and receives direct payoff $\tilde{u}(0, 0) = 0$. Therefore, $\tilde{T}(e_k) = T(x_{K^*-k+1}) - C$ where C solves $T(1) - C = 0$, and hence $C = \frac{1}{2} \left(\frac{\beta}{2} + \eta - \kappa_b \right)$. Therefore, we calculate

$$\begin{aligned} \tilde{T}(e_k) &= \frac{1}{2} \left(1 - \frac{k-1}{K^*-1} \right) \left(\beta + \eta - \kappa_b - \frac{\beta}{2} \left(1 - \frac{k-1}{K^*-1} \right) \right) - \frac{1}{2} \left(\frac{\beta}{2} + \eta - \kappa_b \right) \\ &= -\frac{1}{2} \frac{k-1}{K^*-1} \left(\eta - \kappa_b + \frac{\beta}{2} \frac{k-1}{K^*-1} \right). \end{aligned} \quad (108)$$

We finally translate the transfers into wages by reversing the sign of the payments: $w(e_k) = -\tilde{T}(e_k)$ for all k . That is, a positive wage corresponds to a negative transfer from the agent (worker) to the principal (firm). This completes the proof.

B.5 Proof of Corollary 2

The first part of the statement follows immediately from the characterization of the optimal menu in Proposition 2. We now prove the second part. By assumption on $(\alpha, \beta, \eta, \kappa_b, \kappa_f)$, there exists a maximal neighborhood $\mathcal{B}$ of α such that, for all $\alpha' \in \mathcal{B}$, there exists a unique optimal K^* at $(\alpha', \beta, \eta, \kappa_b, \kappa_f)$. Let $\mathcal{A}$ be the set of $\alpha' \in \mathcal{B}$ such that $K^*(\alpha') = K^*(\alpha)$, where $K^*(\alpha')$ is the unique optimal K^* at parameter vector $(\alpha', \beta, \eta, \kappa_b, \kappa_f)$. We observe

from Lemma 11 that $V = \hat{\Pi} - \hat{\Gamma}$ is submodular in α . This implies that K^* is monotone in α and therefore that $\mathcal{A} = \mathcal{B}$. By Proposition 2, we have that $\{w^*(e_k)\}_{k=1}^{K^*}$ is invariant to α conditional on K^* . Moreover, again by Proposition 2, $\{w^*(e_k)\}_{k=1}^{K^*}$ is different from $\{w^*(e_k)\}_{k=1}^K$ for any $K \neq K^*$ and, as $\mathcal{B}$ is the maximal neighborhood of α such that K^* is unique, whenever $\alpha' \notin \mathcal{B}$, there exists an optimal $K^{**} \neq K^*$.

C Representations of Evidence and Their Properties

In this appendix, we derive several equivalent representations of evidence and establish their topological properties.

C.1 Representing Evidence

In this appendix, we prove Lemma 1. It follows as a consequence of the following, more general equivalence that we exploit in some of our proofs.

Lemma 12. *The following four properties of a contractibility correspondence $C : X \rightrightarrows X$ are equivalent:*

1. *C is a regular contractibility correspondence, that is, $C = C_{\mathcal{E}}$ for some regular evidence structure $\mathcal{E} : X \rightrightarrows \Omega$ (see definition of $C_{\mathcal{E}}$ in Equation 22).*
2. *C satisfies (i) Reflexivity: for every $y \in X$, $y \in C(y)$; (ii) Excludability: $C(0) = \{0\}$; (iii) Transitivity: for every $x, y, z \in X$, if $x \in C(y)$ and $z \in C(x)$, then $z \in C(y)$; (iv) Monotonicity: for every $x, y \in X$, if $x \leq y$, then $C(x) \leq_{SSO} C(y)$, where $\leq_{SSO}$ denotes the strong set order; (v) Continuity: C is closed valued and lower hemicontinuous.*
3. *We have $C(y) = [\underline{\delta}(y), \bar{\delta}(y)]$ for all $y \in X$, where $\underline{\delta} : X \rightarrow X$ is an upper semi-continuous increasing function and $\bar{\delta} : X \rightarrow X$ is a lower semi-continuous increasing function such that: (i) $y \in [\underline{\delta}(y), \bar{\delta}(y)]$, (ii) $\underline{\delta}(x) = \underline{\delta}(y)$ for all $x \in [\underline{\delta}(y), y)$, (iii) $\bar{\delta}(x) = \bar{\delta}(y)$ for all $x \in (y, \bar{\delta}(y)]$, and (iv) $\bar{\delta}(0) = 0$.*
4. *There exist two closed sets $\underline{D} \subseteq X$ and $\bar{D} \subseteq X$ such that $0 \in \underline{D}$, $0, \bar{x} \in \bar{D}$ and for all $y \in X$:*

$$C(y) = \left[\max_{z \leq y: z \in \underline{D}} z, \min_{z \geq y: z \in \bar{D}} z \right]. \quad (109)$$

Moreover, given C , $\underline{D}$ and $\bar{D}$ are unique and given by $\underline{D} = \underline{\delta}(X)$ and $\bar{D} = \bar{\delta}(X)$.

Proof. To prove this, we show that: 1 implies 2, 2 implies 3, 3 implies 4, and 4 implies 1.

1 implies 2. By Equation 22, it is immediate that any contractibility correspondence C that is induced by some evidentiary correspondence $\mathcal{E}$ is transitive, reflexive, closed-valued, and lower hemicontinuous. Observe that definitive evidence of exclusion implies

that $C(0) = \{0\}$, yielding excludability of C . It remains only to show monotonicity of C . Fix $y' \geq y$, $x \in C(y)$, and $x' \in C(y')$. If $x' \geq x$, then $\max\{x, x'\} = x' \in C(y')$ and $\min\{x, x'\} = x \in C(y)$. Suppose now that $x' < x$. We now show that $\mathcal{E}(x) \subseteq \mathcal{E}(y')$ and $\mathcal{E}(x') \subseteq \mathcal{E}(y)$, which yields monotonicity of C by the fact that these claims imply that $\max\{x, x'\} = x \in C(y')$ and $\min\{x, x'\} = x' \in C(y)$. By definition we have that $\mathcal{E}(x) \subseteq \mathcal{E}(y)$, $\mathcal{E}(x') \subseteq \mathcal{E}(y')$, $\mathcal{E}(y') \geq_{SSO} \mathcal{E}(y)$, and $\mathcal{E}(x) \geq_{SSO} \mathcal{E}(x')$. Fix $\omega \in \mathcal{E}(x)$ and $\omega' \in \mathcal{E}(x')$. If $\omega \leq \omega'$, as $\mathcal{E}(x) \geq_{SSO} \mathcal{E}(x')$, we have that $\omega = \min\{\omega, \omega'\} \in \mathcal{E}(x')$. As $\mathcal{E}(x') \subseteq \mathcal{E}(y')$, this implies that $\omega \in \mathcal{E}(y')$. If $\omega > \omega'$, as $\omega \in \mathcal{E}(x) \subseteq \mathcal{E}(y)$ and $\omega' \in \mathcal{E}(x') \subseteq \mathcal{E}(y')$, we know that $\omega = \max\{\omega, \omega'\} \in \mathcal{E}(y')$ by the fact that $\mathcal{E}(y') \geq_{SSO} \mathcal{E}(y)$. These steps imply that $\mathcal{E}(x) \subseteq \mathcal{E}(y')$. Now fix $\omega \in \mathcal{E}(x')$ and $\omega' \in \mathcal{E}(x)$. If $\omega > \omega'$, we have that $\omega = \max\{\omega, \omega'\} \in \mathcal{E}(x)$ as $\mathcal{E}(x) \geq_{SSO} \mathcal{E}(x')$. As $\mathcal{E}(x) \subseteq \mathcal{E}(y)$, this implies that $\omega \in \mathcal{E}(y)$. If $\omega \leq \omega'$, as $\omega \in \mathcal{E}(x') \subseteq \mathcal{E}(y')$ and $\omega' \in \mathcal{E}(x) \subseteq \mathcal{E}(y)$, we have that $\omega = \min\{\omega, \omega'\} \in \mathcal{E}(y)$ as $\mathcal{E}(y') \geq_{SSO} \mathcal{E}(y)$. These steps establish that $\mathcal{E}(x') \subseteq \mathcal{E}(y)$.

2 implies 3. Let $C : X \rightrightarrows X$ be a regular contractibility correspondence and define $\underline{\delta}(y) = \min C(y)$ and $\bar{\delta}(y) = \max C(y)$ for all $y \in X$. By the fact that C is closed-valued, $\bar{\delta}$ and $\underline{\delta}$ exist. By monotonicity, we have that $\underline{\delta}$ and $\bar{\delta}$ are increasing functions. By reflexivity, we know that $y \geq \underline{\delta}(y)$ and $y \leq \bar{\delta}(y)$ for all y . Moreover, by Lemma 17.29 in [Aliprantis and Border \(2006\)](#), $\bar{\delta}$ is lower semi-continuous and $\underline{\delta}$ is upper semi-continuous.

We now show that $C(y) = [\underline{\delta}(y), \bar{\delta}(y)]$. Assume by contradiction there exists some $y \in X$ and $x \in [\underline{\delta}(y), \bar{\delta}(y)]$ such that $x \notin C(y)$. Consider first the case where $x < y$. By the definition of $\underline{\delta}$, $\underline{\delta}(y) \in C(y)$ and $\underline{\delta}(y) < x$. As $x < y$, by monotonicity, we have that $C(x) \leq_{SSO} C(y)$. Thus, as $x \in C(x)$ and $\underline{\delta}(y) \in C(y)$, we know that $\max\{x, \underline{\delta}(y)\} = x \in C(y)$. This is a contradiction. Consider now the case where $y < x$. Again, $\bar{\delta}(y) \in C(y)$ and $x < \bar{\delta}(y)$. By monotonicity, we have that $\min\{x, \bar{\delta}(y)\} = x \in C(y)$. This is a contradiction.

We next show parts (ii), (iii), and (iv). Fix $x, y \in X$ and assume that $x \in [\underline{\delta}(y), \bar{\delta}(y))$, which implies $x \in C(y)$. We start with part (ii), and mirror the argument for part (iii). Suppose $x < y$. As C is monotone, we know that $\underline{\delta}(x) \leq \underline{\delta}(y)$. Suppose by contradiction that $\underline{\delta}(x) < \underline{\delta}(y)$. But then, given the other properties of δ , for all $z \in (\underline{\delta}(x), \underline{\delta}(y))$ we would have that $z \in C(x)$ but $z \notin C(y)$, which contradicts transitivity. For part (iii), consider the same scenario but reversed. Suppose $x > y$. As C is monotone, we know that $\bar{\delta}(x) \geq \bar{\delta}(y)$. Imagine this held with a strict inequality. Then there would exist $z \in (\bar{\delta}(y), \bar{\delta}(x))$ such that $z \in C(x)$ and $z \notin C(y)$, while $x \in C(y)$. This violates transitivity. It is immediate that $\bar{\delta}(0) = 0$ by excludability as $C(0) = \{0\}$.

3 implies 4. We first observe the following for $\underline{\delta}, \bar{\delta}$ with the relevant properties: for all $\underline{z} \in \underline{\delta}(X)$ and $\bar{z} \in \bar{\delta}(X)$, it holds $\underline{\delta}(\underline{z}) = \underline{z}$ and $\bar{\delta}(\bar{z}) = \bar{z}$. To show this, let $\underline{z} = \underline{\delta}(x)$ for some

$x \in X$. We have that $\underline{z} \in [\underline{\delta}(x), x]$. If $\underline{z} = x$, then we have that $\underline{\delta}(\underline{z}) = \underline{\delta}(x) = \underline{z}$. If $\underline{z} < x$, given property (ii), we must have $\underline{\delta}(\underline{z}) = \underline{\delta}(x) = \underline{z}$. The proof for $\bar{z} \in \bar{\delta}(X)$ is symmetric, using property (iii).

Define $\underline{D} = \underline{\delta}(X)$ and $\bar{D} = \bar{\delta}(X)$. First, observe that

$$\max_{z \leq x: z \in \underline{D}} z = \max_{z \leq x: z \in \underline{\delta}(X)} z \geq \underline{\delta}(x), \quad (110)$$

by construction. Let $\underline{z} = \max_{z \leq x: z \in \underline{D}} z$ and assume by contradiction that $\underline{z} > \underline{\delta}(x)$. If $\underline{z} = x$, then $x \in \underline{\delta}(X)$ and by the earlier argument we have that $x = \underline{\delta}(x) < \underline{z}$, yielding a contradiction. If instead $\underline{z} < x$, then by the earlier argument and the property (ii) of $\underline{\delta}$, we have $\underline{z} = \underline{\delta}(\underline{z}) = \underline{\delta}(x)$, yielding a contradiction. With this, we conclude that $\underline{z} = \underline{\delta}(x)$. With symmetric steps, we can show that $\min_{z \geq x: z \in \bar{D}} z = \bar{\delta}(x)$. Next, observe that necessarily we have $\underline{\delta}(0) = 0$, $\bar{\delta}(\bar{x}) = \bar{x}$, and $\bar{\delta}(0) = 0$ proving that $0 \in \underline{D}$ and $0, \bar{x} \in \bar{D}$. Finally, we need to show that $\underline{D}$ and $\bar{D}$ are closed. Take a sequence $z_n \in \underline{D}$ such that $z_n \rightarrow z$. Given that X is closed, we have that $z \in X$ and therefore $\underline{\delta}(z) \leq z$. Given that every z_n is in $\underline{D}$, the property defined above implies that $\underline{\delta}(z_n) = z_n$ for all n . Given that $\underline{\delta}$ is upper semi-continuous, it follows that $z = \lim_{n \rightarrow \infty} z_n = \lim_{n \rightarrow \infty} \underline{\delta}(z_n) \leq \underline{\delta}(z)$. This implies that $z = \underline{\delta}(z)$ (as $z \geq \underline{\delta}(z)$) and therefore that $z \in \underline{D}$. This shows that $\underline{D}$ is closed. A symmetric argument shows that $\bar{D}$ is closed.

4 implies 1. We show this by construction. For the argument, we take $\Omega = X$; as Ω has the cardinality of the continuum, this is without loss of generality, as the construction below can be composed with an isomorphism from X to Ω . We define

$$\mathcal{E}(x) = \left[\max_{z \leq x: z \in \underline{D}} z, \min_{z \geq x: z \in \bar{D}} z \right] =: [\underline{\delta}(x), \bar{\delta}(x)], \quad (111)$$

where, for convenience, we define $\underline{\delta}, \bar{\delta}$ in the second equality. We observe immediately that these functions have the three properties from part 3 of the statement. It is immediate to see that both these functions are monotone increasing, such that $\underline{\delta}(x) \leq x \leq \bar{\delta}(x)$, and respectively upper semi-continuous and lower semi-continuous by Lemma 17.30 in [Aliprantis and Border \(2006\)](#). To see the remaining properties, observe that the correspondences $x \rightrightarrows \{z \in \underline{D} : z \leq x\}$ and $x \rightrightarrows \{z \in \bar{D} : z \geq x\}$ are both upper hemicontinuous. Next, assume that $y \in [\underline{\delta}(x), x]$ and let $z = \underline{\delta}(x)$. We have $\underline{\delta}(y) \leq z$ by monotonicity. Moreover, by assumption $z \leq y$ and $z \in \underline{D}$, so that $z \leq \underline{\delta}(y)$ by definition. We then must have $z = \underline{\delta}(y)$. Symmetrically, assume that $y \in (x, \bar{\delta}(x)]$ and let $z = \bar{\delta}(x)$. We have $\bar{\delta}(y) \geq z$ by monotonicity. Moreover, by assumption $z \geq y$ and $z \in \bar{D}$, so that $z \geq \bar{\delta}(y)$ by definition. We then must have $z = \bar{\delta}(y)$. Finally, as $0 \in \bar{D}$, we have that $\bar{\delta}(0) = 0$.

We now show that the induced evidence structure satisfies definitive evidence of exclusion, monotonicity, and the continuity property. Definitive evidence of exclusion follows from $\bar{\delta}(0) = 0$, and $\bar{\delta}(x) \geq x > 0$ for any $x \neq 0$. Thus, $\mathcal{E}(0) = \{0\}$ and $\mathcal{E}(x) \not\subseteq \mathcal{E}(0)$ for any $x \neq 0$.

We next show monotonicity. To remind, this property implies that, for all $x, x' \in X$ and $e, e' \in \Omega$, if $e' \geq e$ and $x' \geq x$, then

$$\mathbb{I}[e' \in \mathcal{E}(x')]\mathbb{I}[e \in \mathcal{E}(x)] \geq \mathbb{I}[e' \in \mathcal{E}(x)]\mathbb{I}[e \in \mathcal{E}(x')]. \quad (112)$$

This amounts to showing that the intersection of conditions on the right-hand side implies the intersection of conditions on the left-hand side. If both right-hand side conditions hold, then $\underline{\delta}(x) \leq e' \leq \bar{\delta}(x)$ and $\underline{\delta}(x') \leq e \leq \bar{\delta}(x')$. Using the monotonicity of $\underline{\delta}, \bar{\delta}$, we have that $e \leq e' \leq \bar{\delta}(x) \leq \bar{\delta}(x')$ and $\underline{\delta}(x) \leq \underline{\delta}(x') \leq e \leq e'$. Thus, $\underline{\delta}(x) \leq e \leq \bar{\delta}(x)$ and $\underline{\delta}(x') \leq e' \leq \bar{\delta}(x')$. By implication, as desired, $e' \in \mathcal{E}(x')$ and $e \in \mathcal{E}(x)$.

We next show continuity of $\mathcal{E}$. Fix $x, y \in X$. First, consider a sequence $x_n \rightarrow x$ with $\mathcal{E}(x_n) \subseteq \mathcal{E}(y)$ for all n , that is, $\underline{\delta}(y) \leq \underline{\delta}(x_n) \leq \bar{\delta}(x_n) \leq \bar{\delta}(y)$ for all n . The continuity properties of $\underline{\delta}$ and $\bar{\delta}$ imply that $\mathcal{E}(x) \subseteq \mathcal{E}(y)$. Second, assume $\mathcal{E}(x) \subseteq \mathcal{E}(y)$ and fix a sequence $y_n \rightarrow y$. Define $x_n = \min \{ \max \{ x, \underline{\delta}(y_n) \}, \bar{\delta}(y_n) \}$. By construction, we have $\mathcal{E}(x_n) \subseteq \mathcal{E}(y_n)$ for all n . By the semicontinuity properties of $\underline{\delta}$ and $\bar{\delta}$, we have $\limsup_n \underline{\delta}(y_n) \leq \underline{\delta}(y) \leq x$ and $\liminf_n \bar{\delta}(y_n) \geq \bar{\delta}(y) \geq x$. With this, we have that $x_n \rightarrow x$ as desired. $\square$

C.2 The Topology of Sets of Self-Enforcing Recommendations

Let $\underline{\Delta} \times \bar{\Delta}$ denote the space of pairs of functions $(\underline{\delta}, \bar{\delta})$ that have all the properties in Lemma 1. Recall that we endow this set with the relative topology induced by the product L_1 topology over pairs of integrable functions. Also, recall that $\underline{\mathcal{D}}$ denotes the collection of closed subsets of X that contain 0 and $\bar{\mathcal{D}}$ denotes the collection of closed subsets of X that contain 0 and $\bar{x}$. Let $\underline{\mathcal{D}} \times \bar{\mathcal{D}}$ denote the set of pairs of self-enforcing recommendations sets, $(\underline{\mathcal{D}}, \bar{\mathcal{D}})$. Recall that we have endowed this space with the product topology induced by the Hausdorff topology on each collection of sets.

Lemma 13. *The set $\underline{\mathcal{D}} \times \bar{\mathcal{D}}$ is compact.*

Proof. We show that each of $\underline{\mathcal{D}}$ and $\bar{\mathcal{D}}$ is compact in the Hausdorff topology. Specifically, we explicitly establish the compactness of $\bar{\mathcal{D}}$ and observe that an entirely symmetric argument applies to establish the compactness of $\underline{\mathcal{D}}$. With this, the compactness of the product space $\underline{\mathcal{D}} \times \bar{\mathcal{D}}$ follows from Tychonoff's theorem.

Observe that $\bar{\mathcal{D}} \subset \mathfrak{F}$, where $\mathfrak{F}$ is the collection of nonempty closed sets of X . Theorem 3.85 in Aliprantis and Border (2006) establishes that $\mathfrak{F}$ is compact in the Hausdorff topology. Fix

a sequence $\{\overline{D}_n\}_{n \in \mathbb{N}} \subseteq \overline{\mathcal{D}}$. By the compactness of $\mathfrak{F}$, this sequence must have a subsequence $\overline{D}_{n_k}$ converging to some $\overline{D} \in \mathfrak{F}$ in the Hausdorff topology. Furthermore, we have that $0, \bar{x} \in \overline{D}_{n_k}$ for all k . Hence, by Hausdorff convergence of $\{\overline{D}_{n_k}\}_{k \in \mathbb{N}}$, we have that $0, \bar{x} \in \overline{D}$. Since $\overline{D}$ is a closed subset of X that contains 0 and $\bar{x}$, we have $\overline{D} \in \overline{\mathcal{D}}$. Finally, because the initial sequence was arbitrarily chosen, this implies that $\overline{\mathcal{D}}$ is compact. $\square$

Lemma 1 implies that $\underline{\Delta} \times \overline{\Delta}$ and $\underline{\mathcal{D}} \times \overline{\mathcal{D}}$ are isomorphic via the maps:

$$\begin{aligned} (\underline{\delta}, \overline{\delta}) &\mapsto (\underline{D}_{\underline{\delta}}, \overline{D}_{\overline{\delta}}) := (\underline{\delta}(X), \overline{\delta}(X)), \\ (\underline{D}, \overline{D}) &\mapsto (\underline{\delta}_{\underline{D}}(x), \overline{\delta}_{\overline{D}}(x)) := (\max\{z \in \underline{D} : z \leq x\}, \min\{z \in \overline{D} : z \geq x\}). \end{aligned} \quad (113)$$

We next prove that Hausdorff's convergence of sets of self-enforcing recommendations implies L_1 -convergence of the corresponding envelope functions to the envelope functions induced by the limit sets.

Lemma 14. *Fix two sequences $\{(\underline{D}_n, \overline{D}_n)\}_{n \in \mathbb{N}} \subseteq \underline{\mathcal{D}} \times \overline{\mathcal{D}}$ and $\{(\underline{\delta}_n, \overline{\delta}_n)\}_{n \in \mathbb{N}} \subseteq \underline{\Delta} \times \overline{\Delta}$ such that $(\underline{\delta}_n, \overline{\delta}_n) = (\underline{\delta}_{\underline{D}_n}, \overline{\delta}_{\overline{D}_n})$ for all $n \in \mathbb{N}$. If $(\underline{D}_n, \overline{D}_n) \rightarrow (\underline{D}, \overline{D})$, then $(\underline{\delta}_n, \overline{\delta}_n) \rightarrow (\underline{\delta}_{\underline{D}}, \overline{\delta}_{\overline{D}})$.*

Proof. Fix two sequences as in the statement and assume that $(\underline{D}_n, \overline{D}_n) \rightarrow (\underline{D}, \overline{D})$ in the product Hausdorff topology. We only show that $\overline{\delta}_n \rightarrow \overline{\delta}_{\overline{D}}$ in L_1 , as $\underline{\delta}_n \rightarrow \underline{\delta}_{\underline{D}}$ follows from an entirely symmetric argument. Together, these two facts imply convergence in the L_1 product topology.

For notational simplicity, denote $\overline{\delta}_{\overline{D}} = \overline{\delta}$ and let $X_{\overline{\delta}} \subseteq X$ denote the collection of points at which $\overline{\delta}$ is continuous. Because $\overline{\delta}$ is non-decreasing, it follows that $X \setminus X_{\overline{\delta}}$ is at most countable. We next show that $\overline{\delta}_n(x) \rightarrow \overline{\delta}(x)$ for all $x \in X_{\overline{\delta}}$. Before showing this, we observe that the previous claim concludes the argument because

$$\lim_{n \rightarrow \infty} \int_X |\overline{\delta}_n(x) - \overline{\delta}(x)| dx = \lim_{n \rightarrow \infty} \int_{X_{\overline{\delta}}} |\overline{\delta}_n(x) - \overline{\delta}(x)| dx = \int_{X_{\overline{\delta}}} \lim_{n \rightarrow \infty} |\overline{\delta}_n(x) - \overline{\delta}(x)| dx = 0, \quad (114)$$

where the first equality follows from the fact that $X_{\overline{\delta}}$ has full measure, the second follows from the dominated convergence theorem, and the last follows from the claim.

Next, fix $x \in X_{\overline{\delta}}$. Clearly, if $x = \bar{x}$, then $\overline{\delta}_n(x) \rightarrow \overline{\delta}(x)$ since they are all equal to $\bar{x}$. Thus, we next always assume that $x < \bar{x}$. We split the rest of the proof of the claim into two parts depending on whether x is inside or outside $\overline{D}$.

1. Assume that $x \in X_{\overline{\delta}} \setminus \overline{D}$ and define $\overline{D}_n(x) = \overline{D}_n \cap [x, \bar{x}]$ for all n as well as $\overline{D}(x) = \overline{D} \cap [x, \bar{x}]$. We first show that $\overline{D}_n(x) \rightarrow \overline{D}(x)$ in the Hausdorff topology. Let $Li(\overline{D}_n(x))$ and $Ls(\overline{D}_n(x))$ respectively denote the topological limit inferior and the topological limit superior

of the sequence $\{\overline{D}_n(x)\}_{n \in \mathbb{N}}$.²² Next, fix $z \in \overline{D}(x)$ and observe that $z > x$ because $x \in X_{\overline{\delta}} \setminus \overline{D}$. Because $z \in \overline{D}$ and $\overline{D}_n \rightarrow \overline{D}$, it follows that there exists a sequence $z_n \in \overline{D}_n$ such that $z_n \rightarrow z$. In particular, we must have $z_n > x$ for all n large enough, yielding that $z_n \in \overline{D}_n(x)$ for all n large enough, hence that $z \in Li(\overline{D}_n(x))$. Because z was arbitrarily chosen, this implies that $\overline{D}(x) \subseteq Li(\overline{D}_n(x))$. Next, fix $z \in Ls(\overline{D}_n(x))$. By definition of Ls , there exists a sequence $\{z_k\}_{k \in \mathbb{N}}$ such that $z_k \rightarrow z$ and $z_k \in \overline{D}_{n_k}(x)$ along a subsequence parametrized by k . Because $\overline{D}_{n_k} \rightarrow \overline{D}$ in the Hausdorff topology by assumption and $z_k \in \overline{D}_{n_k}$ for all k , it follows that $z \in \overline{D}$. Similarly, because $z_k \in [x, \overline{x}]$ for all k , it follows that $z \in [x, \overline{x}]$, yielding that $z \in \overline{D}(x)$. Because z was arbitrarily chosen, this implies that $Ls(\overline{D}_n(x)) \subseteq \overline{D}(x)$ and overall that $\overline{D}(x) = Li(\overline{D}_n(x)) = Ls(\overline{D}_n(x))$. Theorem 3.93 in Aliprantis and Border (2006) then yields that $\overline{D}_n(x) \rightarrow \overline{D}(x)$ in the Hausdorff topology. Finally, because $\overline{\delta}_n(x) = \min \overline{D}_n(x)$ for all n and $\overline{\delta}(x) = \min \overline{D}(x)$, Theorem 17.31 in Aliprantis and Border (2006) implies that $\overline{\delta}_n(x) \rightarrow \overline{\delta}(x)$.

2. Assume that $x \in X_{\overline{\delta}} \cap \overline{D}$ and define $\overline{D}_n(x)$ for all n and $\overline{D}(x)$ as above. Also, let $\text{int } \overline{D}$ and $\partial \overline{D}$ respectively denote the interior and the boundary points of $\overline{D}$. First, observe that $\overline{\delta}(x) = x$ because $x \in \overline{D}$. We next show that, for every $\varepsilon > 0$, there exists a point $x^\varepsilon \in \overline{D}(x)$ such that $0 < |x^\varepsilon - x| \leq \varepsilon$. We split the proof of this claim into two cases:

- a. If $x \in \text{int } \overline{D}$, then the claim is immediately true because $\text{int } \overline{D}$ is open.
- b. Assume now that $x \in \partial \overline{D}$ and, by contradiction, that there exists $\varepsilon > 0$ such that for all $z \in \overline{D}(x)$ we have either $|z - x| = 0$ or $|z - x| > \varepsilon$. Because $x < \overline{x}$, there must be $z \in \overline{D}(x)$ with $|z - x| > 0$, hence for all such z we must have $|z - x| > \varepsilon$. In turn, this implies that $\overline{D} \cap [x, x + \varepsilon] = \{x\}$, and hence that $\overline{\delta}(y) > x + \varepsilon$ for all $y \in (x, x + \varepsilon)$. With this, we have that $\overline{\delta}$ is discontinuous at x , yielding a contradiction.

Next, fix $\varepsilon > 0$ and $x^\varepsilon \in \overline{D}(x)$ as above. Because $x^\varepsilon \in \overline{D}$, Hausdorff convergence implies that there exists a sequence $x_n^\varepsilon \in \overline{D}_n$ such that $x_n^\varepsilon \rightarrow x^\varepsilon$. Moreover, because $x^\varepsilon > x$, for n large enough we must have $x_n^\varepsilon > x$, hence that $x_n^\varepsilon \in \overline{D}_n(x)$. Therefore, for all n large enough, we have

$$x \leq \overline{\delta}_n(x) = \min \overline{D}_n(x) \leq x_n^\varepsilon. \quad (115)$$

Passing to the limits, we have

$$x \leq \liminf_n \overline{\delta}_n(x) \leq \limsup_n \overline{\delta}_n(x) \leq x^\varepsilon. \quad (116)$$

By taking $\varepsilon \rightarrow 0$ we conclude that $\lim_n \overline{\delta}_n(x) = x = \overline{\delta}(x)$, as desired. $\square$

²²See for example Definition 3.80 in Aliprantis and Border (2006).

Before stating the main result of this section, observe that each cost function $\Gamma_{\mathcal{E}}$ induces cost functions in the spaces of envelope functions and sets of self-enforcing recommendations (by Lemma 1). In particular, we write these induced costs as $\Gamma_{\Delta} : \underline{\Delta} \times \overline{\Delta} \rightarrow [0, \infty]$ defined over the envelope functions of regular contractibility correspondences and $\Gamma_{\mathcal{D}} : \underline{\mathcal{D}} \times \overline{\mathcal{D}} \rightarrow [0, \infty]$ defined over the sets of self-enforcing recommendations of regular contractibility correspondences. Of course these induced costs are related by the identity $\Gamma_{\mathcal{D}}(\underline{D}, \overline{D}) = \Gamma_{\Delta}(\underline{\delta}_D, \overline{\delta}_D)$.

Proposition 3. *If a cost function Γ_{Δ} is lower semi-continuous (resp. continuous) in the product L_1 topology, then $\Gamma_{\mathcal{D}}$ is lower semi-continuous (resp. continuous) in the product Hausdorff topology.*

Proof. The result immediately follows from Lemma 14. □

This proposition is useful in our analysis as it is simple to verify that front-end costs are continuous in the product L_1 topology when written in terms of envelopes of contractibility correspondences (as they are simply integrals of these functions). However, it is not simple to directly verify Hausdorff continuity of the cost in the space of self-enforcing recommendations. This result provides that conclusion, which is an essential step in our existence proof.

D Flexible Use of Evidence

Our baseline analysis imposes a seeming restriction on the principal's use of evidence: payments depend only on the promised action and the binary event of whether they can prove that the agent violated the contract. In this appendix, we allow the principal to freely contract on the promised action and the evidence that they present to the arbitrator. We establish that, for any equilibrium under a mechanism that uses evidence flexibly, the principal can design a breach mechanism under which an equilibrium exists that yields the same expected payoff. Thus, it is without loss of optimality for the principal to employ breach mechanisms, justifying our focus on them.

A flexible mechanism $P : X \times 2^{\Omega} \rightarrow \mathbb{R}$ specifies payments from the agent to the principal $P(y, E)$ when the agent promises to do $y \in X$ and the principal presents evidence $E \subseteq \Omega$ to the arbitrator. The principal can present any subset E of the evidence generated by the agent's action $\mathcal{E}(x)$ to the arbitrator. Such a mechanism allows, for example, the principal to punish different violations of the contract in different ways (or not at all). The timing is as in our baseline model, but here we are explicit that after evidence has been generated, the principal elects which subset to exhibit to the arbitrator, who then executes payments in accordance with the flexible mechanism.

We also allow for limited liability in the form of a cap $\bar{M} \in \bar{\mathbb{R}}_{++}$ on the possible transfer from the agent to the principal: $P(y, E) \leq \bar{M}$ for all $y \in X$ and $E \subseteq \Omega$. If a mechanism P respects this constraint, we say that it respects limited liability; in the limit case where $\bar{M} = \infty$, this is trivially satisfied by any mechanism. We describe an equilibrium (a pure SPNE) as a joint distribution $W \in \Delta(\Theta \times X \times X)$ over tuples (θ, y, x) of types, promised actions, and final actions.

Our analysis in the main text implicitly restricts to a simpler class of mechanisms:

Definition 4 (Breach Mechanism). *Given a regular evidence structure $\mathcal{E}$, a flexible mechanism $P : X \times 2^\Omega \rightarrow \mathbb{R}$ is a breach mechanism for $\mathcal{E}$ if there exists a tariff $T : X \rightarrow \mathbb{R}$ and a punishment $M \in \mathbb{R}_{++}$ such that:*

$$P(y, E) = \begin{cases} T(y), & E \cap \mathcal{E}(y)^c = \emptyset, \\ M, & \text{otherwise.} \end{cases} \quad (117)$$

Under a breach mechanism, payments $T(y)$ are provided when the agent promises to do $y \in X$ so long as the principal cannot prove that the agent did not do y , *i.e.*, $E \cap \mathcal{E}(y)^c = \emptyset$. When the principal can prove that the agent did not do y , they levy a punishment M . A breach mechanism respects limited liability if $\max \{ \sup_{y \in X} T(y), M \} \leq \bar{M}$.

The following result establishes that breach mechanisms are without loss of optimality, both for our baseline analysis and if limited liability is imposed:

Proposition 4. *Fix some limited liability cap $\bar{M}$. For every regular evidence structure $\mathcal{E}$, contract P that respects limited liability $\bar{M}$, and equilibrium action distribution W given $(\mathcal{E}, P)$, there exists a breach mechanism $\hat{P}$ for $\mathcal{E}$ that respects limited liability $\bar{M}$ and an equilibrium action distribution $\hat{W}$ given $(\mathcal{E}, \hat{P})$ such that the principal's expected payoff from $(\mathcal{E}, P, W)$ equals that of $(\mathcal{E}, \hat{P}, \hat{W})$.*

Proof. Fix an equilibrium W under evidence structure $\mathcal{E}$ and contract P . Let $\phi : \Theta \rightarrow X$ and $\xi : \Theta \rightarrow X$ denote the equilibrium mappings from types to final actions and promises, respectively. We also define the value of the principal optimally exhibiting evidence as:

$$\mathcal{T}(x, y) = \max_{E \subseteq \mathcal{E}(x)} P(y, E), \quad (118)$$

which must have a solution by the hypothesis that we are in an equilibrium. Moreover, let $V(\theta) = u(\phi(\theta), \theta) - \mathcal{T}(\phi(\theta), \xi(\theta))$ denote type θ 's equilibrium payoff under P .

We first set the punishment for the breach mechanism M . If $\bar{M} < \infty$, set $M = \bar{M}$. If

$\bar{M} = \infty$, choose any finite $M > \max_{x \in X, \theta \in \Theta} u(x, \theta)$. In either case,

$$u(x, \theta) - M \leq V(\theta) \quad \text{for all } x \in X, \theta \in \Theta. \quad (119)$$

When $\bar{M} < \infty$, this follows because $\mathcal{T}(x, y) \leq \bar{M} = M$ for all (x, y) and equilibrium under P implies $V(\theta) \geq u(x, \theta) - \mathcal{T}(x, y)$ for every deviation (y, x) .

We fix a map $\zeta : X \rightarrow X$ such that:

$$\zeta(z) \in \arg \min_{y \in X} \mathcal{T}(z, y) \quad \forall z \in \phi(\Theta). \quad (120)$$

To see that such a map exists, fix $z \in \phi(\Theta)$ and choose θ_z such that $\phi(\theta_z) = z$. If there existed $y \in X$ such that $\mathcal{T}(z, y) < \mathcal{T}(z, \xi(\theta_z))$, then type θ_z could promise y and still take action z , contradicting equilibrium under P . Hence $\xi(\theta_z) \in \arg \min_{y \in X} \mathcal{T}(z, y)$, and the argmin is nonempty. In particular, for every type θ ,

$$\mathcal{T}(\phi(\theta), \zeta(\phi(\theta))) = \mathcal{T}(\phi(\theta), \xi(\theta)). \quad (121)$$

Putting this together, we define the map $T : X \rightarrow \mathbb{R}$ such that:

$$T(z) = \begin{cases} \mathcal{T}(z, \zeta(z)), & z \in \phi(\Theta), \\ M, & \text{else.} \end{cases} \quad (122)$$

By construction, $T(z) \leq M$ for all $z \in X$. We take $\hat{P}$ as the breach mechanism induced by tariff T and penalty M .

Any agent can secure their equilibrium payoff $V(\theta)$ under P also under $\hat{P}$ by taking the same final action and promising to take this action. To see this, note that if any agent $\theta \in \Theta$ promises $\phi(\theta)$ and takes action $\phi(\theta)$ under $\hat{P}$, then they pay the principal

$$T(\phi(\theta)) = \mathcal{T}(\phi(\theta), \zeta(\phi(\theta))) = \mathcal{T}(\phi(\theta), \xi(\theta)), \quad (123)$$

where the last equality follows from Equation 121. Thus, it remains to show that the agent can do no better under $\hat{P}$ than under P . We now define:

$$\hat{\mathcal{T}}(x, y) = \max_{E \subseteq \mathcal{E}(x)} \hat{P}(y, E), \quad (124)$$

for the induced transfer under $\hat{P}$. Toward a contradiction, suppose that there exists $x' \in X$,

$y' \in X$ and a type $\theta \in \Theta$ such that:

$$u(x', \theta) - \widehat{T}(x', y') > V(\theta). \quad (125)$$

There are two possible cases. First, suppose that $y' \in \phi(\Theta)$ and y' is such that $\mathcal{E}(x') \cap \mathcal{E}(y')^c = \emptyset$. In this case, we have that $\mathcal{E}(x') \subseteq \mathcal{E}(y')$. Thus, we have that:

$$u(x', \theta) - \widehat{T}(x', y') = u(x', \theta) - T(y') = u(x', \theta) - \mathcal{T}(y', \zeta(y')). \quad (126)$$

Because $\mathcal{E}(x') \subseteq \mathcal{E}(y')$, for all $z \in X$,

$$\mathcal{T}(x', z) = \max_{E \subseteq \mathcal{E}(x')} P(z, E) \leq \max_{E \subseteq \mathcal{E}(y')} P(z, E) = \mathcal{T}(y', z). \quad (127)$$

In particular, taking $z = \zeta(y')$ and combining the previous two equations, we have that:

$$u(x', \theta) - \widehat{T}(x', y') \leq u(x', \theta) - \mathcal{T}(x', \zeta(y')) \leq V(\theta), \quad (128)$$

where the last inequality follows from equilibrium under P , since under the original flexible mechanism type θ could have promised $\zeta(y')$ and taken action x' . This yields a contradiction.

Second, suppose that either $y' \notin \phi(\Theta)$ or y' is such that $\mathcal{E}(x') \cap \mathcal{E}(y')^c \neq \emptyset$. In this case, the transfer between the principal and the agent is M and the hypothesis (to be contradicted) becomes that:

$$u(x', \theta) - M > V(\theta). \quad (129)$$

But this contradicts Equation 119.

Thus, there is an equilibrium $\widehat{W}$ under $\widehat{P}$ such that agents' actions follow ϕ , yielding the same final action distribution as W . Moreover, all transfers between the agent and the principal are the same. The evidence structure is unchanged, so front-end costs are unchanged; and since the final action distribution is unchanged, back-end costs are unchanged. Because of this, the principal's expected payoffs are equal under $(\mathcal{E}, P, W)$ and $(\mathcal{E}, \widehat{P}, \widehat{W})$ and the statement is proven. $\square$

For intuition, consider a general mechanism that is *not* a breach mechanism, and under which some agent takes an action z by promising y , deviating from this promise, and making the payment $P(y, E(z, y))$, where $E(z, y) \subseteq \mathcal{E}(z)$ is the equilibrium evidence that the principal shows to the arbitrator. If $P(y, E(z, y))$ is the cheapest way to take action z , then it is implicitly the price that any agent who desires to do z would pay (Equation 120). The principal could instead construct a tariff in which z is directly offered at this price (see

Equation 122), and then punish breaches with the maximum punishment $\bar{M}$. This gives equally powerful incentives for any agent who previously “cheated” to take action z to now do so “honestly.” Thus, the class of breach mechanisms is sufficient for the principal’s design problem, as they can do no better under general mechanisms. This immediately recovers our baseline analysis when limited liability does not bind ($\bar{M} = \infty$).

A full analysis of our model under nontrivial limited liability ($\bar{M} < \infty$) is beyond the scope of this paper. But we sketch here how our results would extend in that environment. By Proposition 4, it is without loss of optimality for the principal to focus on breach mechanisms. With limited liability, the principal cannot levy a punishment exceeding $\bar{M}$. Moreover, the agent can always choose $\bar{x}$ and accept the punishment $\bar{M}$. Because all agents prefer higher actions, this is the best among all ways to *violate* one’s promise. Thus, it is as if the outside option of the agent is now $\bar{u}(\theta) = \max\{u(\bar{x}, \theta) - \bar{M}, 0\}$. If participation is ensured under this outside option, then the constraint that $T(y) \leq \bar{M}$ is slack. We further note that if agent $\theta \in \Theta$ prefers $(\bar{x}, \bar{M})$ over any other element of the menu, then so too does every agent $\theta' \geq \theta$. This implies that there exists a threshold type $\bar{\theta}$ such that (i) $\bar{\theta}$ is indifferent between the proposed action and payment and $(\bar{x}, \bar{M})$ and (ii) every $\theta > \bar{\theta}$ weakly prefers $(\bar{x}, \bar{M})$ to their best option in the menu. Single-crossing implies that, if (i) holds, then (ii) must hold. The relevant participation constraints for all types $\theta < \bar{\theta}$ are twofold. First, they prefer their menu selection to not participating. Second, they prefer their menu selection to $(\bar{x}, \bar{M})$: but this is guaranteed by monotonicity, single-crossing, and the definition of $\bar{\theta}$.

However, unlike the case without limited liability, we cannot easily get rid of the first participation constraint by making it bind only for the lowest type. Instead, the designer’s choice of the transfer and the final action function $\phi(\theta)$ also affects the indifference condition for type $\bar{\theta}$. This does *not* yield a type-dependent IR constraint, which would require a substantially different analysis, but rather the addition of a single constraint to capture indifference of $\bar{\theta}$. Thus, we can use a Lagrangian approach to take into account this additional constraint and maximize a modified version of the principal’s virtual surplus that takes into account the corresponding Lagrange multiplier. Up to this multiplier, the problem is then still equivalent to the pointwise maximization problem we have exploited in our arguments. This in turn implies that, in the relevant space of types $\theta \in [0, \bar{\theta}]$ that do not choose to take the largest action and the penalty $\bar{M}$, the optimal final action function will have the same properties we have derived in our results.²³

²³We omit further details in the interests of space, but such details are available on request.

E Probabilistic Arbitration

In our main analysis, we assumed that the process of determining breach of contract was deterministic and error-free. However, it is natural to imagine that an arbitrator may make mistakes in judging the evidence. These mistakes may further be of two varieties: type I errors, where the agent is adjudged to have violated the contract when the true evidence does not support that conclusion; and type II errors, where the agent is adjudged to have not violated the contract when the evidence in fact proves that they did. In this extension, we show that our main analysis and results extend to a setting in which punishments are bounded and an arbitrator is subject to both such type I and type II errors.

The model and timeline are modified in two ways. First, in the arbitration phase, after the agent has acted, the arbitrator can make errors. When $\mathcal{E}(x) \subseteq \mathcal{E}(y)$, *i.e.*, there is no evidence to support that the agent deviated, there is a probability $\alpha \in [0, 1)$ that the arbitrator wrongly concludes that the agent shall be punished (type I). Instead, when $\mathcal{E}(x) \cap \mathcal{E}(y)^c \neq \emptyset$, *i.e.*, there is evidence that proves the agent deviated, there is a probability $1 - \beta \in [0, 1]$ that the arbitrator wrongly concludes that the agent shall not be punished (type II). Second, the financial punishment that the arbitrator levies is given by some $\bar{P} < \infty$. This ensures that the mere possibility of type I errors does not dissuade all agents from participating.

An agent who acts in accordance with the contract expects to lose $\alpha\bar{P}$ in utility from being wrongfully punished, in addition to forfeiting their promised payment. An agent who fails to act in accordance with the contract expects to lose $\beta\bar{P}$ in utility. Moreover, the best final action conditional on committing a breach is the highest-payoff action. As we show below, a sufficiently large punishment and probability of detection suffices to ensure that no agent wishes to do this, rendering type II errors irrelevant to our analysis.

The more delicate issue is the possibility of type I errors. In their presence, it is as if the agent (conditional on participation) has the payoff $u(x, \theta) - (1 - \alpha)t - \alpha\bar{P}$ and the principal has the payoff $\pi(x, \theta) + (1 - \alpha)t + \alpha\bar{P}$. It may now be optimal to exclude some agents from participation in the mechanism. However, conditional on this optimal choice of exclusion, our remaining analysis is unchanged.

These claims are summarized in the following Proposition, which makes use of the definition $\bar{u} = \max_{x \in X, \theta \in \Theta} u(x, \theta)$:

Proposition 5. *If $(\beta - \alpha)\bar{P} > (1 - \alpha)\bar{u}$, then the claims of Theorems 1 and 2,²⁴ and Corollary 1 hold.*

Proof. We first show that no agent will participate in the mechanism and then breach the

²⁴Here, the definition of non-triviality must be strengthened to include the hypothesis that $H(\phi^P(\theta), \theta) > 0$ at some θ for which there is a unique, non-degenerate, and interior maximum.

contract under the assumption that $(\beta - \alpha)\bar{P} > (1 - \alpha)\bar{u}$. For an agent who selects $(y, T(y))$, compliance entails an expected payment of $\alpha\bar{P} + (1 - \alpha)T(y)$, whereas a detectable breach entails an expected payment of $\beta\bar{P} + (1 - \beta)T(y)$. Suppose, toward a contradiction, that an agent participates and chooses a detectable breach $x \notin C(y)$. Let $u_B = u(x, \theta) \leq \bar{u}$. Participation requires $U_B := u_B - \beta\bar{P} - (1 - \beta)T(y) \geq 0$. If $\beta < 1$, this implies $(1 - \beta)(\bar{P} - T(y)) \geq \bar{P} - u_B$. The agent could instead comply by choosing $x' \in C(y)$, yielding $U_C := u_C - \alpha\bar{P} - (1 - \alpha)T(y)$, where $u_C := u(x', \theta) \geq 0$. The difference between the payoff from compliance and that from breach satisfies

$$\begin{aligned} (1 - \beta)(U_C - U_B) &= (1 - \beta)(u_C - u_B) + (1 - \beta)(\beta - \alpha)(\bar{P} - T(y)) \\ &\geq (1 - \beta)u_C - (1 - \alpha)u_B + (\beta - \alpha)\bar{P} \\ &\geq (\beta - \alpha)\bar{P} - (1 - \alpha)\bar{u} > 0. \end{aligned}$$

where the second line plugs in $(1 - \beta)T(y) \leq u_B - \beta\bar{P}$ (from participation) and simplifies, the third line uses $u_B \leq \bar{u}$ and $u_C \geq 0$, and the final inequality asserts the assumed condition. The agent therefore strictly prefers compliance, a contradiction. If $\beta = 1$, the stated condition reduces to $\bar{P} > \bar{u}$, in which case every detectable breach gives the agent a strictly negative payoff. Hence, no agent ever participates and then commits a detectable breach. Type II errors consequently do not occur on path and are immaterial to the remainder of the analysis.

We now study how type I errors affect the principal's problem. As discussed before, these transform the principal and agent gross payoffs into the following form $\tilde{\pi}(x, \theta) = \pi(x, \theta) + \alpha\bar{P}$ and $\tilde{u}(x, \theta) = u(x, \theta) - \alpha\bar{P}$ conditional on the agent participating. This violates our maintained assumption that $\tilde{\pi}(0, \theta) = \tilde{u}(0, \theta) = 0$ for all $\theta \in \Theta$ but does not affect any of our other maintained assumptions on the gross payoffs. We must therefore adapt Lemma 3 to handle this change. A mechanism now comprises (ϕ, ξ, r, T) , where $r : \Theta \rightarrow \{0, 1\}$ is a mapping from types to inclusion, with $r(\theta) = 1$ denoting inclusion and $r(\theta) = 0$ denoting exclusion. The agent's IR constraint now changes to $u(\phi(\theta), \theta) - (1 - \alpha)T(\xi(\theta)) \geq \alpha\bar{P}$ if $r(\theta) = 1$. The IC constraint is the same, except the agent's payoff is replaced by $u(x, \theta) - (1 - \alpha)T(y)$ and it only need hold when $r(\theta) = 1$. The Obedience constraint is unchanged, except that it need only hold when $r(\theta) = 1$. Finally, we require that the agent's decision not to participate is incentive compatible:

$$r(\theta) = 0 \implies \mathcal{L}(\theta) \equiv \max_{y \in X, x \in C(y)} u(x, \theta) - (1 - \alpha)T(y) - \alpha\bar{P} \leq 0. \quad (130)$$

$\mathcal{L}(\theta)$ is a weakly increasing function of θ . Hence, $r(\theta)$ may be taken without loss of optimality

as monotone increasing, and so there exists a $\bar{\theta} \in [0, 1]$ such that $r(\theta) = \mathbb{I}[\theta \geq \bar{\theta}]$.

Lemmas 8 and 9 must be updated to replace the IR constraints with the modified IR constraint. Lemma 9 holds as stated except that the tariff must instead be given, for all $x \in X$ such that there exists $\theta \in \Theta$ with $r(\theta) = 1$ and $\phi(\theta) = x$, by:

$$(1 - \alpha)T(x) = Z + u(x, \phi^{-1}(x)) - \int_{\bar{\theta}}^{\phi^{-1}(x)} u_{\theta}(\phi(s), s) ds \quad (131)$$

with $Z \leq -\alpha\bar{P}$. In Lemma 10, the principal's problem then becomes:

$$\begin{aligned} \max_{\phi, \bar{\theta}, Z \leq -\alpha\bar{P}} \int_{\bar{\theta}}^1 & \left(\pi(\phi(\theta), \theta) + \alpha\bar{P} + Z + u(\phi(\theta), \theta) - \int_{\bar{\theta}}^{\theta} u_{\theta}(\phi(s), s) ds - \int_{\phi(\theta)}^{\bar{x}} g_b(\phi(\theta), z) dz \right) dF(\theta) \\ \text{s.t. } & \phi(\theta') \geq \phi(\theta), \phi(\theta) \in \bar{D} \quad \forall \theta, \theta' \in \Theta : \theta' \geq \theta \\ & u(\phi(\theta), \theta) - \left(\alpha\bar{P} + Z + u(\phi(\theta), \theta) - \int_{\bar{\theta}}^{\theta} u_{\theta}(\phi(s), s) ds \right) \geq 0 \quad \forall \theta \geq \bar{\theta}. \end{aligned} \quad (132)$$

It follows that it is optimal to set $Z = -\alpha\bar{P}$. Adapting the standard integration by parts argument, we then obtain that:

$$\mathcal{R}(\bar{D}) = \max_{\bar{\theta}} \int_{\bar{\theta}}^1 \left(\max_{x \in \bar{D}} H(x, \theta) \right) dF(\theta). \quad (133)$$

The fact that H is strictly supermodular is what allows us to drop the monotonicity constraint. Let $\bar{\Theta}$ denote the set of types such that $\check{H}(\theta) \equiv \max_{x \in \bar{D}} H(x, \theta) = 0$. As H is supermodular, $\check{H}$ is increasing. Moreover, for all $\theta \geq \max \bar{\Theta}$, $\check{H}$ is strictly increasing. This follows because for any $\theta, \theta' \geq \max \bar{\Theta}$ such that $\theta' > \theta$, we have that $\check{H}(\theta') = H(x^*(\theta'), \theta') \geq H(x^*(\theta), \theta') > H(x^*(\theta), \theta) = \check{H}(\theta)$, where the strict inequality follows from the fact that H is strictly increasing in θ for $x > 0$. With this, we may take $\bar{\theta} = \max \bar{\Theta}$ if $\bar{\Theta}$ is non-empty. Otherwise, if $\check{H}(\theta) < 0$ for all $\theta \in \Theta$ we take $\bar{\theta} = 1$ and if $\check{H}(\theta) > 0$ for all $\theta \in \Theta$ we take $\bar{\theta} = 0$. By non-triviality (with the stated strengthening in Footnote 24), the same arguments as those of Theorem 2 imply that any $\bar{D}$ that maximizes $\mathcal{R}$ must be infinite.

We now establish that the conclusion of Lemma 7 holds as stated. We define $\bar{\theta}(\bar{D})$ as the solution $\bar{\theta}$ that we just described. We also define $\hat{x}(\bar{D}) \equiv \max(\arg \max_{x \in \bar{D}} H(x, \bar{\theta}(\bar{D})))$. It is immediate that if $\bar{D} \cap (0, \hat{x}(\bar{D})) \neq \emptyset$, then $\bar{D} \setminus (0, \hat{x}(\bar{D}))$ yields a strictly lower front-end cost and the same value of $\mathcal{R}$. Hence, any optimal $\bar{D}$ contains no elements between 0 and $\hat{x}(\bar{D})$. Moreover, for any (a, b) such that $\hat{x}(\bar{D}) < a < b$, we have that $\bar{\theta}(\bar{D} \setminus (a, b)) = \bar{\theta}(\bar{D})$. Thus, the arguments of Lemma 7 hold as stated. Theorem 1 and Corollary 1 follow. $\square$

Remark 3. This proof addresses the variant of the model that we flagged in Footnote 9. In particular, if we allow the principal not to pay back-end costs for the agents that it excludes, then it may wish to exclude agents whenever $\max_{x \in X} H(x, \theta) < 0$. The analysis above shows that this does not affect how front and back-end costs affect optimal contracts. $\triangle$

F Extending the Results to More General Costs

In this appendix, we describe a richer family of front-end and back-end cost functions and provide general conditions under which the conclusions of Theorems 1 and 2 hold.

F.1 Generalizing Front-End Costs

We first define an order over regular evidentiary correspondences, the set of which we denote by $\mathcal{E}$. We say that $\mathcal{E}'$ generates more refined evidence than $\mathcal{E}$, which we denote by $\mathcal{E}' \succsim \mathcal{E}$, if for all $x, y \in X$ such that $\mathcal{E}'(x) \subseteq \mathcal{E}'(y)$ we also have that $\mathcal{E}(x) \subseteq \mathcal{E}(y)$. In words, this means that $\mathcal{E}'$ generates more refined evidence than $\mathcal{E}$ if every time that x cannot be proven by the principal to be inconsistent with y using evidence generated by $\mathcal{E}'$, the same is true if evidence were generated by $\mathcal{E}$. Observe that $\mathcal{E}' \succsim \mathcal{E}$ if and only if $C_{\mathcal{E}'} \subseteq C_{\mathcal{E}}$. Thus, this order over evidence is equivalent to inducing a greater set of self-enforcing recommendations.

We say that a (front-end) cost function $\Gamma : \mathcal{E} \rightarrow \mathbb{R}_+$ is monotone if whenever $\mathcal{E}' \succsim \mathcal{E}$, then we have that $\Gamma(\mathcal{E}') \geq \Gamma(\mathcal{E})$. We argue that this is a natural property for a cost to possess: if an evidentiary correspondence generates more refined evidence, then it incurs a higher cost. We now show that monotonicity of Γ justifies writing costs directly over contractibility correspondences. Recall that we let $\mathcal{C}$ denote the set of regular contractibility correspondences. Formally, we define what it means for a cost function to be measurable in the induced contractibility correspondence as:

Definition 5 (C-measurability). Γ is C -measurable if there exists a $\check{\Gamma} : \mathcal{C} \rightarrow [0, \infty]$ such that for all $\mathcal{E} \in \mathcal{E}$, we have that $\Gamma(\mathcal{E}) = \check{\Gamma}(C_{\mathcal{E}})$.

If Γ is C -measurable, we call the corresponding $\check{\Gamma}$ the induced cost. We have that:

Lemma 15. If Γ is monotone, then it is C -measurable.

Proof. Fix an arbitrary pair of evidentiary correspondences $\mathcal{E}, \mathcal{E}' \in \mathcal{E}$ such that $C_{\mathcal{E}} = C_{\mathcal{E}'}$. We have that $C_{\mathcal{E}} \subseteq C_{\mathcal{E}'}$, which implies that $\mathcal{E} \succsim \mathcal{E}'$. By monotonicity of Γ we have that $\Gamma(\mathcal{E}) \geq \Gamma(\mathcal{E}')$. By the reverse argument, we have that $\Gamma(\mathcal{E}') \geq \Gamma(\mathcal{E})$. Hence, we have that $\Gamma(\mathcal{E}') = \Gamma(\mathcal{E})$ for all $\mathcal{E}, \mathcal{E}' \in \mathcal{E}$ that induce the same contractibility correspondence. Take $\check{\Gamma}(C_{\mathcal{E}}) = \Gamma(\mathcal{E})$. Thus, Γ is C -measurable with induced $\check{\Gamma}$. $\square$

This result implies that whenever costs are monotone in the natural sense over underlying evidentiary correspondences, it is without loss of optimality to write the cost in the space of contractibility correspondences directly. It is immediate to see that all the front-end costs of distinguishing that we considered in the main text are monotone.

As in the main analysis, we call a front-end evidentiary cost of distinguishing linear if $g_f(x, y) = \kappa > 0$ and we write this cost function as Γ^κ . We can now define strong monotonicity:

Definition 6 (Strong Monotonicity). *A cost function Γ is strongly monotone if there exists $\varepsilon > 0$ such that, for all $\mathcal{E}, \mathcal{E}' \in \mathcal{E}$ such that $\mathcal{E}' \succsim \mathcal{E}$:*

$$\Gamma(\mathcal{E}') - \Gamma(\mathcal{E}) \geq \varepsilon (\Gamma^1(\mathcal{E}') - \Gamma^1(\mathcal{E})). \quad (134)$$

With this, we show that strong monotonicity implies a certain property for the induced cost defined over contractibility correspondences:

Proposition 6. *If Γ is strongly monotone with constant $\varepsilon > 0$, then the induced $\check{\Gamma}$ satisfies:*

$$\check{\Gamma}(C') - \check{\Gamma}(C) \geq \varepsilon (\check{\Gamma}^1(C') - \check{\Gamma}^1(C)), \quad \forall C' \subseteq C \quad (135)$$

Proof. We first observe that strong monotonicity of Γ implies monotonicity of Γ . Thus, by Lemma 15, we have that Γ is C -measurable and therefore has an induced $\check{\Gamma}$. Now fix an arbitrary pair $\mathcal{E}, \mathcal{E}' \in \mathcal{E}$ such that $\mathcal{E}' \succsim \mathcal{E}$, we have that:

$$\check{\Gamma}(C_{\mathcal{E}'}) - \check{\Gamma}(C_{\mathcal{E}}) = \Gamma(\mathcal{E}') - \Gamma(\mathcal{E}) \geq \varepsilon (\Gamma^1(\mathcal{E}') - \Gamma^1(\mathcal{E})) = \varepsilon (\check{\Gamma}^1(C_{\mathcal{E}'}) - \check{\Gamma}^1(C_{\mathcal{E}})), \quad (136)$$

where the first equality is by C -measurability, the first inequality is by strong monotonicity of Γ and the final equality is by definition of the linear front-end cost of distinguishing. Thus, as $\mathcal{E}' \succsim \mathcal{E} \iff C_{\mathcal{E}'} \subseteq C_{\mathcal{E}}$, we have established the required condition on $\check{\Gamma}$. $\square$

By Lemma 12, we have that $C \in \mathcal{C}$ if and only if $C(y) = [\underline{\delta}_C(y), \bar{\delta}_C(y)]$ for two semicontinuous, monotone functions with the escalator-step property. We say that $\check{\Gamma} : \mathcal{C} \rightarrow \mathbb{R}_+$ is continuous if it is continuous in the product L_1 -topology with respect to $(\underline{\delta}_C, \bar{\delta}_C)$.

Given a monotone front-end cost function Γ , we say that Γ is continuous if its induced $\check{\Gamma}$ is continuous. In this case, Proposition 3 implies that the induced cost function over the set of self-enforcing recommendations is continuous in the Hausdorff topology. We now show that strong monotonicity and continuity of Γ are sufficient for the conclusion that optimal contracts are necessarily coarse.

Proposition 7. *Suppose that Γ is continuous and strongly monotone with constant $\varepsilon > 0$ and the principal faces no back-end costs. An optimal evidence structure exists and every optimal evidence structure is coarse. Moreover, the bound on the richness of any optimal evidence structure and menu from Theorem 1 holds, with ε replacing $\underline{g}_f$.*

Proof. By continuity, Lemma 4 applies and so an optimal evidence structure exists. By Proposition 6, Lemma 6 holds with ε replacing $\underline{g}_f$. The remaining arguments follow the proof of Theorem 1. $\square$

F.2 Generalizing Back-End Costs

We consider back-end cost functions $\Gamma : \mathcal{E} \times \Delta(X) \rightarrow \mathbb{R}_+$ given by:

$$\Gamma(\mathcal{E}, Q) = \int_X \gamma(\mathcal{E}, x) dQ(x), \quad (137)$$

for some $\gamma : \mathcal{E} \times X \rightarrow \mathbb{R}_+$. The interpretation here is that $\gamma(\mathcal{E}, x)$ is the cost of generating the evidence according to the action x given the evidence structure $\mathcal{E}$. Given this, Γ is then the expected cost that will be incurred if the distribution over actions that is taken is given by Q . We say that γ is monotone if whenever $\mathcal{E}' \succeq \mathcal{E}$, then we have that $\gamma(\mathcal{E}', x) \geq \gamma(\mathcal{E}, x)$ for all $x \in X$. By the exact same arguments as those in Lemma 15, if γ is monotone, then it is C -measurable. We denote by $\tilde{\gamma}$ the induced cost over contractibility correspondences, *i.e.*, $\tilde{\gamma}(C_{\mathcal{E}}, x) \equiv \gamma(\mathcal{E}, x)$. By monotonicity of γ in $\mathcal{E}$, we also have that $\tilde{\gamma}$ is monotone with respect to set inclusion in the space of regular contractibility correspondences, that is, $C' \subseteq C$ implies that $\tilde{\gamma}(C', x) \geq \tilde{\gamma}(C, x)$ for all $x \in X$. With this, we have that the induced back-end cost is also C -measurable and we denote the induced cost over contractibility correspondences by $\tilde{\Gamma}$, *i.e.*, $\tilde{\Gamma}(C_{\mathcal{E}}, Q) = \int_X \tilde{\gamma}(C_{\mathcal{E}}, x) dQ(x) = \Gamma(\mathcal{E}, Q)$. We say that $\tilde{\gamma}$ is *separable* if $\tilde{\gamma}(C, x) = \hat{\gamma}(C(x), x)$ for some $\hat{\gamma} : \Xi \times X \rightarrow \mathbb{R}_+$, where $\Xi \subseteq 2^X$ is the collection of closed and convex subsets of X , and we endow it with the Hausdorff topology. We note that as $\tilde{\gamma}$ is monotone with respect to set inclusion, so too is $\hat{\gamma}$. We say that a separable $\tilde{\gamma}$ is continuous if the induced function $\hat{\gamma}$ is continuous with respect to the product topology. Finally, we define $\zeta(x) = \hat{\gamma}([0, x], x)$ and $H(x, \theta) = J(x, \theta) - \zeta(x)$. With these definitions, we can state the following result.

Proposition 8. *Suppose that the principal faces no front-end costs and γ is monotone with an induced $\tilde{\gamma}$ that is separable and continuous. An optimal evidence structure $\mathcal{E}^*$ with cardinality $|\mathcal{E}^*(X)| = \infty$ exists. Moreover, if H is non-trivial and twice continuously differentiable, then every optimal evidence structure $\mathcal{E}^*$ and every optimal menu $\mathcal{M}^*$ have cardinality $|\mathcal{E}^*(X)| = |\mathcal{M}^*| = \infty$.*

Proof. If γ is monotone with a separable $\tilde{\gamma}$, then we may write the back-end cost as:

$$\tilde{\Gamma}(C, Q) = \int_X \hat{\gamma}(C(x), x) dQ(x) = \int_X \hat{\gamma}([\underline{\delta}(x), \bar{\delta}(x)], x) dQ(x), \quad (138)$$

where $\underline{\delta}$ and $\bar{\delta}$ are the lower and upper envelopes of C from the representation of Lemma 12. The arguments of Lemma 2 then apply here and we have that $\underline{\delta}(x) \equiv 0$ is optimal. Thus:

$$\begin{aligned} \tilde{\Gamma}(C, Q_\phi) &= \int_X \hat{\gamma}([0, \bar{\delta}(x)], x) dQ_\phi(x) = \int_\Theta \hat{\gamma}([0, \bar{\delta}(\phi(\theta))], \phi(\theta)) dF(\theta) \\ &= \int_\Theta \hat{\gamma}([0, \phi(\theta)], \phi(\theta)) dF(\theta) = \int_\Theta \zeta(\phi(\theta)) dF(\theta), \end{aligned} \quad (139)$$

where the penultimate line follows (as in Lemma 3) from Obedience and the last uses the definition of $\zeta(x) = \hat{\gamma}([0, x], x)$. We can now apply the remaining arguments of Lemma 3 where the statement now holds with $H(x, \theta) = J(x, \theta) - \zeta(x)$. By the hypothesis of continuity of $\tilde{\gamma}$, we have that ζ (and therefore H) is continuous. From this, the arguments from Appendix A.2 apply and the conclusion of Theorem 2 then follows. For the second conclusion, non-triviality and twice continuous differentiability yield the conclusion by the arguments of Theorem 2. $\square$

References

- ALIPRANTIS, C. D., AND K. BORDER (2006): *Infinite dimensional analysis: A Hitchhiker's Guide*. Springer.
- KULKARNI, D., D. SCHMIDT, AND S.-K. TSUI (1999): "Eigenvalues of tridiagonal pseudo-Toeplitz matrices," *Linear Algebra and its Applications*, 297, 63–80.