Delegated Costly Screening

Suraj Malladi*

September 21, 2021

Abstract

A policymaker relies on regulators or bureaucrats to screen agents through a costly instrument on her behalf. How can she maintain some control over the design of the screening process? She solves a two-layer mechanism design problem: she restricts the set of allowable allocations, after which a screener picks a menu that maps an agent's costly evidence to this restricted set. In general, the policymaker can set a floor in a way that dominates full delegation no matter how the screener's objectives are misaligned. When this misalignment is only over the relative importance of reducing allocation errors or agent's screening costs, the effectiveness of this restriction hinges sharply on the direction of the screener's bias. In the min-max optimal mechanism, if the screener is more concerned with reducing errors, setting this floor is in fact robustly optimal for the policymaker. But if the screener is more concerned with keeping costs down, not only does this particular floor have no effect: any restriction that strictly improves over full delegation is complex and sensitive to the details of the screener's preferences. I consider the implications for regulatory governance.

*Stanford Graduate School of Business. Email: surajm@stanford.edu. I thank my advisor, Andrzej Skrzypacz, and committee members Matthew Jackson, Michael Ostrovsky and Mohammad Akbarpour for their close guidance. I am grateful for valuable suggestions from Kyle Bagwell and Yonatan Gur. I also thank Jeremy Bulow, Daniel Chen, Zi Yang Kang, Edward Lazear, Yucheng Liang, Andres Perlroth, Marcos Salgado, David Yang, participants at the Political Economy Theory Seminar at Stanford GSB, SIOE 2020, EC'20, and especially Anirudha Balasubramanian for many helpful discussions.

1 Introduction

Policymakers rely on intermediaries to screen agents on their behalf: Lawmakers rely on local government officials to screen workfare applicants by their willingness to engage in menial labor. Managers of a public healthcare program rely on hospitals to screen patients by their willingness to wait for procedures. Governments let consumer protection regulators and patent offices screen innovators by their willingness to undergo costly applications.

These intermediaries might not be better informed than policymakers ex-ante about which applicant is relatively well off or poor or which innovations are good or bad. But they may be better placed to observe the costly evidence or effort that applicants or innovators produce, and to therefore screen these agents on that basis. But this advantage may also allow intermediaries to seize control over mechanism design. How can a policymaker influence an intermediary's design? When is this easy to do? How much does she need to know about an intermediary's biases to do it?

To answer these questions, I study a model of delegated costly screening. There are three players: a policymaker (she), a screener (he) and an agent. First, the policymaker limits the screener to some set of allowable allocations. Next, the screener designs a menu that offers the agent a larger allowable allocation for more costly evidence. Finally, the agent takes stock of its private information and decides how much costly evidence to generate. For example, a policymaker can require that workfare applicants receive either nothing, between \$50 and \$100 per week, or between \$200 and \$400 per week. A local government official or bureaucrat can respond by offering any applicant \$70 if they meet basic eligibility requirements, and \$300 if they also complete certain low-productivity tasks. An eligible agent considers his costs of labor and benefits of receiving payments to decide whether to take the welfare check or work for the larger payment.

A good screening program has high targeting efficiency so that relatively worse (well) off applicants are not too under (over) allocated. At the same time, the program should not be too onerous for applicants. The policymaker and screener minimize a weighted sum of their (agent type-dependent) disutility from allocation errors and agent screening costs but disagree about the relative importance of these objectives. For tractability, the benchmark setting considers the case when the policymaker and screener evaluate menu choices by the worst-case realization over the agent's type[1]

[1] One interpretation is that policymakers and regulators may be min-max ambiguity averse in a setting like innovation approval and seek mechanisms with good payoff guarantees. But they could also worry about worst-case outcomes even when they have priors over agent types. For example, managers of a workfare or publicly funded health care program may try to defend against political opponents who would cherry-pick cases to attack those programs, e.g., by highlighting patients who undergo significant hurdles only to receive poor healthcare coverage or beneficiaries who have little need and are still awarded substantial welfare.

Screener Bias:	Reducing Allocation Errors	Reducing Screening Costs
Payoffs and Form:	optimal delegation is *simple* and *perfect*	delegation is either *complex or futile*
Information:	optimal delegation is *robustly optimal*	no delegation set is *robustly improving*

Figure 1: This table summarizes the main results for the benchmark setting.

To paraphrase the main result, the policymaker can easily align the screener's behavior if the latter cares more about keeping down allocation errors than she does. But aligning the screener's behavior is hard if he worries more about keeping down applicant's screening costs (e.g., would prefer moving closer to a welfare rather than workfare scheme).

More precisely, Theorem 1 (summarized in Figure 1) shows that when the screener places more weight on reducing allocation errors than the policymaker, optimal delegation rules are *simple*: they take the form of intervals. One optimal delegation rule for the policymaker is to set a floor for screener at the lowest allocation level she would award, were she the one designing the menu. This delegation set is *perfect*: the constrained screener can do no better than to choose the policymaker's favorite menu. Finally, this delegation set is *robust*: the policymaker can identify her optimal rule knowing only the direction of the screener's bias.

Theorem 2 shows that these results all fail when the screener places more weight on reducing screening costs. Unless the players' preferred menus coincide, delegation is either *futile*, i.e, no restriction to the agent's action space strictly improves over full delegation, or delegation is *complex*, i.e., strict improvements exist but none of them are simple. Moreover, a strict improvement over full delegation, even if it exists, can never be *robust*: there always exists a screener, biased in the same direction but to a different extent, for whom this restriction produces a worse outcome for the policymaker than full delegation would.

The difficulty of delegation depends on whether the screener wants to offer a steeper or flatter menu than the policymaker. In the former case, the policymaker can compress his menu from the bottom. The min-max screener complies with the policymaker's preferences on the interior. In the latter case, the policymaker has to stretch the screener's menu by perforating the allocation space. This can backfire if the allowable allocations are not precisely chosen.

The results imply that floors are min-max optimal and undominated, when even the direction of the screener's bias is unknown. A practical implication is that policymakers should always push to make the screener worry more about allocation errors and set a floor, *even if they care a lot about screening costs.* If a screener still cares more about costs afterward, the floor is not binding and this move at least closes the gap in preferences. And if the move exacerbates the divergence in preference, the floor ensures that the screener uses the policymaker's preferred menu all the same.

The next set of results generalize beyond the benchmark setting. First, policymakers and screeners may disagree on many dimensions, such as the relative importance of type I or type II errors or what even constitutes an allocation error. Proposition 1 shows that even when both have their own distinct loss functions, the policymaker can weakly improve over full delegation by setting a floor at the allocation level she would offer for free. Next, Proposition 2 shows that if allocation error losses are strictly convex and screening cost losses are linear, the same result holds in a model where the players care about expected losses. Therefore, setting floors is a sound policy under fully general preference divergence, in max-min and expected-loss settings.

1.1 Related Literature

Describing the role of dissipative costs for screening users of publicly funded healthcare services, Zeckhauser (2019) writes:

> Ordeals currently play a prominent and critical role in directing resources to high-value users... Unlike pricing, the primary instrument of resource allocation in developed societies, ordeals are scarcely studied, little understood, and often accepted without thought... conscious attention to their design and operation could greatly enhance that value.

One way of framing the exercise here is to note that a policymaker's inability to verify agent costs would distinguish ordeals from pricing. So to give "conscious attention" to the design of ordeals, policymakers need to consider how their ideal mechanism differ from those they task with implementing it.

This paper can speak to the various applications where policymakers might try to introduce ordeals and costly screens. One such application is innovation approval. Masur (2010) and De Rassenfosse and Jaffe (2018) consider how applications costs and high attorney fees can produce a positive selection effect in patents. Lemley and Shapiro (2005) goes further by arguing for a reformed system that "thinks of the process of issuing patents in terms of designing a mechanism", for example by "[letting] patent applicants select either the normal, brief examination process, which would lead to a Standard Patent if the application were approved, or a more rigorous application process, which would lead to Super Patent if the application were approved". Other applications include targeting welfare disbursement, where recent empirical work measures the efficacy of different ordeals (see Alatas et al. (2016) and Deshpande and Li (2019)).

This paper belongs to the literature on delegation, which, starting with Holmstrom (1977, 1980), focuses on aligning the behavior of experts who are better informed (e.g., Frankel (2014)) or can more cheaply acquire information (e.g., Szalay (2005); Chade and Kovrijnykh (2016)). Here I instead focus on how to align the behavior of intermediaries whose advantage lies in being able to design richer mechanisms to elicit applicants' private information than the policymaker can design alone.[2]

The closest delegation papers in this sense are Amador and Bagwell (2016), who study the Baron and Myerson (1982) monopolist regulation problem in a setting without transfers; and Guo and Shmaya (2019), who study this problem when the regulator minimizes worst-case regret. There, the monopolist's choice is assumed to be a single price-quantity pair; for example, Amador and Bagwell (2016) restrict the monopolist from using schemes like two-part tariffs. In contrast, the intermediary here is free to choose any implementable menu, which gives rise to a much richer delegation and contracting space.

This paper is also closely related to the literature on mechanism design and collusion, where one agent (like the screener in the present model) has all the bargaining power to offer side contracts (Laffont and Martimort, 1998; Mookherjee and Tsumagari, 2004; Faure-Grimaud et al., 2003; Celik, 2009). By considering environments with transfers (where the intermediary may be made a residual claimant) and binary types or perfectly informed intermediaries, these earlier models shift the focus away from optimal delegation. They instead address the more primitive question of whether or not delegation to a subcontractor can be as effective as directly contracting with all parties. In the present paper, a form of the taxation principle immediately shows that some delegated screening mechanism is optimal. The focus is instead on characterizing optimal rules and study when delegation leads to only a nominal or a real loss of control.

To make progress on this problem of delegated mechanism design, this paper takes a min-max approach, as in Hurwicz and Shapiro (1978), Frankel (2014), and Carroll (2015, 2017)[3]. It is especially related to those with min-max regret criteria like Bergemann and Schlag (2008) and Guo and Shmaya (2019), who respectively study robust monopoly pricing and monopoly regulation. Minimizing worst-case regret can be seen as a special case of the model where allocation error losses are linear.

[2]The general results here on the robustness of floors have no analog in the delegated information acquisition literature. There, whether floors or any other policy would help or backfire depends on the specific assumptions made about the intermediary's information acquisition technology (see Section 5.4).

[3]See Carroll (2019) for an overview of the robust mechanism design literature.

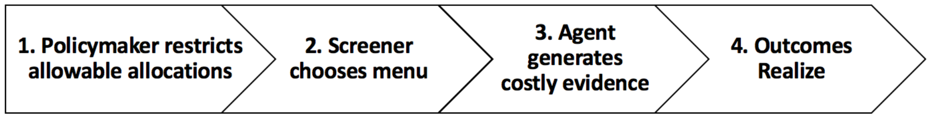

Figure 2: The game proceeds in four periods, with the policymaker committing to a delegation set first and the screener committing to a menu thereafter.

2 The Model

This section develops the benchmark model of delegated screening. Section 2.1 does a preliminary analysis to restate player objectives in a more convenient way. Section 2.2 discusses interpretations of the model.

Players, Actions and Timing There is a policymaker (she), a screener (he), and an agent (it).

In the first period, the policymaker chooses an allowable allocation space, which is *any* closed subset Y of the $[0, 1]$ interval. That is, the policymaker's action space is any closed set in $\mathcal{P}([0, 1])$.

In the second period, the screener chooses a menu that maps costly evidence generated by the agent to an allocation in Y. That is, the screener is free to choose any right-continuous (r.c.) menu $y : \mathbb{R}_+ \to Y$.

In the third period, the agent generates costly evidence $n \in \mathbb{R}$, given its private information and the menu selected by the screener.

In the last period, payoffs of all agents are realized. The timing is summarized in Figure 2.

Other Primitives and Information The agent has a multidimensional type $\theta \in \Theta$. The agent's type determines its marginal cost $c(\theta) > 0$ of generating costly evidence, and its marginal benefit $b(\theta) \in \mathbb{R}_+$ of receiving an allocation. The agent's payoff is normalized to 0 if it is allocated nothing. The agent's type also determines the type-specific losses that the policymaker and screener face from misallocating to the agent, $f_\theta : [0, 1] \to \mathbb{R}_+$. Finally, the agent's type determines the type-specific losses that the policymaker and screener face from the agent having to generate costly effort, $h_\theta : \mathbb{R}_+ \to \mathbb{R}_+$.

$\Theta, c, b, f_\theta, h_\theta$ are common knowledge, but θ is the agent's private information. The policymaker and screener evaluate their choices by considering their payoffs under their worst-case realization of the agent's type, so it is immaterial for the results whether one assumes that they have ambiguity over Θ or full support priors (however, see Section 5.2.1 for two interpretations).

I assume that the allocative loss functions $\{f_\theta\}_{\theta\in\Theta}$ are continuous and strictly quasiconvex, so deviations from the unique optimal allocation level for each type lead to losses. Next, I assume that the screening cost loss functions $\{h_\theta\}_{\theta\in\Theta}$ are continuous and strictly increasing and that $h_\theta(0) = 0$ for all $\theta \in \Theta$. Finally, I assume that Θ is finite to sidestep issues like the existence of optimal menus.[4]

Agent's Objective The agent produces an amount of costly evidence n to solve:

$$\max_{n\in\mathbb{R}_+} b(\theta)y(n) - c(\theta)n.$$

Note that b and c need not be co-monotone. Let $N_y(\theta)$ denote the set of optimal evidence levels for a type θ agent facing menu y.

Screener's Objective The screener minimizes the worst-case weighted sum of losses due to allocation errors and screening costs, where the worst case is taken over the realization of the agent's type. The screener chooses a menu y to minimize:

$$\max_{\theta\in\Theta}\ \min_{n\in N_y(\theta)} \alpha f_\theta(y(n)) + (1-\alpha)h_\theta(n),$$

where $\alpha \in (0, 1)$ is the screener's *weight on allocation errors*. The solution concept assumes that the agent breaks ties in favor of the screener when indifferent among multiple menu choices. Let $\mathcal{X}(Y)$ denote the set of screener optimal menus when she is restricted to menus with co-domain Y.[5] Among evidence levels in $N_y(\theta)$, let $n_y(\theta)$ denote the screener's favored choice for a type θ agent.

Policymaker's Objective The policymaker is also motivated by a desire to minimize worst-case weighted sum of losses. She differs from the screener only by the relative weights she places on the losses from allocation errors versus screening costs. She chooses a delegation set $Y \subset [0, 1]$ to solve:

$$\inf_{Y}\ \min_{y\in\mathcal{X}(Y)}\ \max_{\theta\in\Theta} \alpha_P f_\theta(y(n_y(\theta))) + (1-\alpha_P)h_\theta(n_y(\theta))$$

The policymaker's choice of delegation set, Y, changes the set of screener optimal menus, $\mathcal{X}(Y)$. Of these menus, the screener chooses the one that minimizes the policymaker's worst-case loss. Note that when $\alpha_P \neq \alpha$, the worst-case type realization for the policymaker need not coincide with the worst-case type realization for the screener.

[4] Appendix A discusses how the results extend to the case with infinitely many types.

[5] The screener's preference for reducing evidence production costs for the agent endogenously puts a cap on the largest evidence level that any agent will be required to show in equilibrium. This coupled with the facts that Y is closed and the loss functions are continuous ensures that $\mathcal{X}(Y)$ is nonempty.

2.1 Implemenetability and a Reformulations of the Screener's and Policymaker's Problem

This section restates the policymaker's and screener's objective in a more convenient way by characterizing implementable menus. One complication that needs to be accounted for is that the agent has a multi-dimensional type but is screened along a single dimension.

Effective Types Since agent types vary in both costs of producing evidence and benefits of receiving allocations, different types can have the same preferences over the screener's menu. Consider two menu items $(n, y(n))$ and $(n', y(n'))$, and suppose there are two types of agents θ and θ' such that $\frac{b(\theta)}{c(\theta)} = \frac{b(\theta')}{c(\theta')}$. Note that,

$$\begin{aligned}
& b(\theta)y(n) - c(\theta)n \geq b(\theta)y(n') - c(\theta)n' \\
\iff & \frac{b(\theta)}{c(\theta)}y(n) - n \geq \frac{b(\theta)}{c(\theta)}y(n') - n' \\
\iff & \frac{b(\theta')}{c(\theta')}y(n) - n \geq \frac{b(\theta')}{c(\theta')}y(n') - n' \\
\iff & b(\theta')y(n) - c(\theta')n \geq b(\theta')y(n') - c(\theta')n'.
\end{aligned}$$

Define $\tau(\theta) \equiv \frac{b(\theta)}{c(\theta)}$ to be *effective type* of a type θ agent. Let $\Theta_\tau \subset \Theta$ be the set of all types θ with effective type τ, and let $\mathcal{T}$ denote the set of all effective types. Let $\underline{\tau}$ and $\overline{\tau}$ denote the lowest and highest effective types, respectively. The preceding shows that all types in Θ_τ for a given $\tau \in \mathcal{T}$ have identical preferences over the screener's menu (see the left panel in Figure 3).

It is convenient to assume the following tie-breaking rule: for every τ, all types in Θ_τ break indifferences among choices in the screener's menu in the same way. Screener optimal menus may sometimes involve ties being broken in a different way, but the results hold with or without this assumption.

Revelation Principle The convenience of making this assumption is that we can now apply the *revelation principle* to the one-dimensional space of effective types, $\mathcal{T}$. That is, we can equivalently think of the screener as choosing direct mechanisms that map an agent's effective type to an allocation and the evidence the agent needs to show to obtain it. Abusing notation, the screener chooses an *allocation rule* and *standard of proof*, $(y, n) : \mathcal{T} \to [0, 1] \times \mathbb{R}_+$, that satisfy incentive compatibility (IC): for all $\tau, \tau' \in \mathcal{T}$,

$$\tau y(\tau) - n(\tau) \geq \tau y(\tau') - n(\tau') \tag{IC}$$

An allocation rule y is *implementable* if there exists some standard of proof n such that (y, n) is incentive compatible. We now have the familiar characterization of implementability.

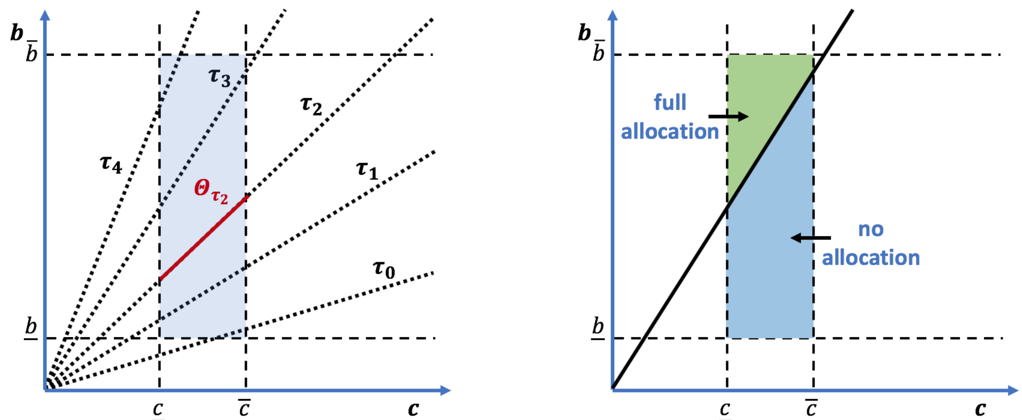

Figure 3: In this example, the agent's costs lie between $\underline{c}$ and $\bar{c}$ and benefits between $\underline{b}$ and $\bar{b}$. The left figure shows iso effective type curves. The highlighted line segment in red consists of all types with effective type τ_2. The right figure depicts participating (green) and nonparticipating (blue) types under a cutoff allocation rules which awards full allocation if the agent shows a certain amount of evidence and no allocation otherwise. Types with sufficiently low benefits or sufficiently high costs choose not to show evidence.

Lemma 1 (Myerson's Lemma). *An allocation rule is implementable if and only if it is non-decreasing.*

Let $\mathcal{I}$ denote the set of implementable allocation rules, and let $\mathcal{I}_Y$ denote the implementable allocation rules with co-domain $Y \subset [0, 1]$. Finally, abusing notation, let $\mathcal{X}(Y)$ now refer to the set of screener optimal *allocation rules* in $\mathcal{I}_Y$.

Optimal Standard of Proof For an allocation rule $y \in \mathcal{I}_Y$, let $y_0 < \ldots < y_k$ denote the set of allocation levels in its range. For $i < k$, let τ_i denote the highest effective type $\tau \in \mathcal{T}$ for which $y(\tau) = y_i$. Define a standard of proof n inductively: $n(\tau_0) = 0$, and $n(\tau_{i+1})$ is set so that all agents in Θ_{τ_i} are indifferent between (y_i, n_i) and (y_{i+1}, n_{i+1}). It is straightforward to check that n is pointwise (weakly) smaller than any other standard of proof that implements y.

Since the agent's costly evidence generation contributes to screener's loss, I assume without loss of generality that the screener chooses a standard of proof with upward IC constraints binding in this way. Therefore, the screener's choice of allocation rule y pins down the standard of proof n.

Screener's Problem and Policymaker's Problem We can now re-express the screener's and policymaker's objective in a more manageable form.

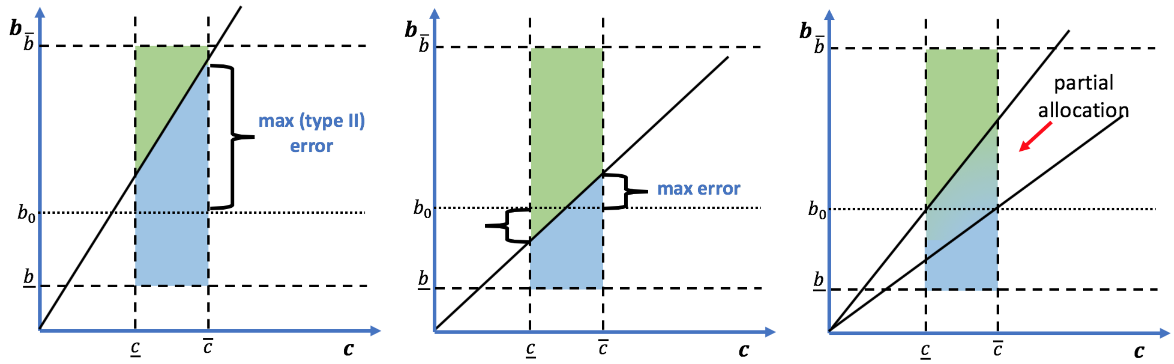

Figure 4: The first two panels depict cutoff allocation rules. By lowering the requisite evidence (moving from the left to the center panel), the screener reduces the size of the worst-case type II (i.e., failing to approve certain types with high benefits of allocation), but introduces the possibility of making type I errors (i.e., allocating to types with low benefits of allocation). The screener can reduce worst-case errors further by awarding partial allocations to intermediate effective types, as depicted in the third panel.

The screener chooses $y \in \mathcal{I}_Y$ to minimize:

$$R(y) \equiv \max_{\tau \in \mathcal{T}} R_\tau(y) \equiv \max_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} \alpha f_\theta(y(\tau(\theta))) + (1-\alpha) h_\theta(n(\tau(\theta))) \tag{1}$$

The policymaker chooses a closed delegation set $Y \subset [0,1]$ to minimize:

$$\min_{y \in \mathcal{X}(Y)} R^P(y) \equiv \min_{y \in \mathcal{X}(Y)} \max_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} \alpha_P f_\theta(y(\tau(\theta))) + (1-\alpha_P) h_\theta(n(\tau(\theta))) \tag{2}$$

When we say an allocation rule y is screener-optimal (policymaker-optimal) without reference to Y, we mean it minimizes $R(y)$ ($R^P(y)$) among all allocation rules in $\mathcal{I}$.

Example Figure 4 shows a simple example where the screener wishes to give a full allocation to types who have a benefit above b_0 an no allocation to those below. Losses due to misallocations are greater when the agent's benefits are further from b_0. In particular, the screener chooses an allocation rule y to solve

$$\min_{y \in \mathcal{I}} \max_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} \alpha \max\{(1-y(\tau))(b(\theta)-b_0), y(\tau)(b_0 - b(\theta))\} + (1-\alpha) h_\theta(n(\tau)).$$

Here, $f_\theta(y) \equiv \max\{(1-y)(b(\theta)-b_0), y(b_0-b(\theta))\}$. The screener stands to gain by using partial allocations since it minimizes the size of the largest error he can make.

2.2 Discussion

Optimality of Delegation The model studies a game where the design of a screen is disaggregated between two players. Every choice by the policymaker and screener gives one particular indirect mechanism in an alternative model where (1) the policymaker can write any contract with the screener and agent directly, (2) the screener can offer the agent a side contract before either responds to the policymaker, and (3) the agent has a hidden type, while its action (evidence generation) is hidden only to the policymaker. Can the focus on delegated screening mechanisms be justified in the context of this more general model? Appendix B answers affirmatively: restricting attention to delegated screening mechanisms is without loss of generality in this alternative model in the search for policymaker optimal (deterministic) mechanisms. The policymaker loses nothing from her inability to contract with the agent directly.

Soft Evidence In this model, evidence is a one dimensional choice variable that the agent produces at a cost. This captures the notion of an ordeal that is verifiable by the screener but not the policymaker. For example, in the case of welfare disbursement, this can be the number of times or days the agent spends appealing an initial rejection decision.

Soft evidence reveals the agent's private information and is particularly useful when readily observable information about applicants is inconclusive. For example, the Social Security Disability Insurance (SSDI) program determines eligible benefits for each impairment, income level and other observable characteristics. But an applicant may claim a disability that is not automatically ruled in or out for benefits. In these cases, the costs of filling out applications and getting additional tests may screen out applicants who are less likely to pass such tests or more easily obtain employment.[6]

This type of screening can also be helpful when agents can manipulate hard evidence. For example, a pharmaceutical company may find some limited opportunities to inflate its apparent efficacy and safety metrics in trials.[7] Of the firms that show a middling average treatment effect, those firms that are willing to spend more and run larger trials may be the ones that manipulate less and so will sell better post-approval. Alternatively, an automated vehicle company that self-reports data to a regulator may hide some incidents where human intervention was required to override the computer's algorithm. Even if they can bury negative evidence, higher quality firms have an easier time producing trips without incident. In these examples, trials or successful trips are soft evidence that firms use to signal their type.

Agent's Private Information The agent is assumed to know its marginal cost of showing evidence $c(\theta)$ and benefit of receiving an allocation $b(\theta)$. Welfare applicants may know both

[6] Keiser (1999) documents substantial variation in how bureaucrats further screen such applicants, showing that the SSDI program leaves them with substantial leeway in deciding many cases.

[7] One way is to halt trials prematurely when positive results are noticed without checking if these signals persist.

quantities. But an innovating firm may face substantial uncertainty over its benefit of getting approval from a regulator. The model can be easily reinterpreted as one in which the firm knows much less. For example, it may be min-max ambiguity averse and know only some range of its possible benefits, with $b(\theta)$ being the lowest benefit in this range.

Policymaker and Screener's Objectives[8] The allocation error loss functions f_θ are strictly quasiconvex. So for each type, there is a unique optimal allocation and larger deviations from it lead to larger losses. For example, policymakers and bureaucrats may worry about disbursing too little to a deserving welfare applicant or too much to an undeserving one, approving a bad drug or delaying the approval of a good one, etc.

The screening cost loss functions h_θ are strictly increasing, since agents signal through dissipative costs rather than transfers. The policymaker and screener may internalize the strain that a long appeals process places on a family waiting for welfare; or they may face criticism for the trial costs they impose on innovating firms. These type specific loss functions may also represent the fact that the players place more weight on the losses of some type of agents than they do others.

Preference Misalignment The disagreement between the policymaker and screener is over the weights on allocation errors versus screening costs, i.e., α versus α_P. Consumer protection regulators, for example, may receive more criticism on the basis of approval errors while the lawmakers who appointed them may face more criticism for the costs that firms incur during approval. For example, the largest component of drug development costs are FDA trials, and part of these costs are borne by drug users in the form of higher prices. In other cases, regulators may care more about imposing intrusive costs on the firms they work closely with. Alternatively, bureaucrats in a welfare program or the patent office have to spend more time and do more paperwork every time to review the additional evidence that agents submit. They may mechanically care more about screening costs than policymakers who would worry mainly about the correctness of these decisions.

3 Optimal Menus for the Unconstrained Screener

This section characterizes optimal allocation rules for an unconstrained screener. Section 4, which presents the main results on delegation, shows that the characterization for this case alone is the key to solving the policymaker's problem.

Example Figure 5 continues with an example where the allocative loss functions are linear, with slope determined by the distance between the agent's marginal benefit of allocation and

[8] A discussion of max-min assumption is deferred to Section 5.2 where a Bayesian model is introduced.

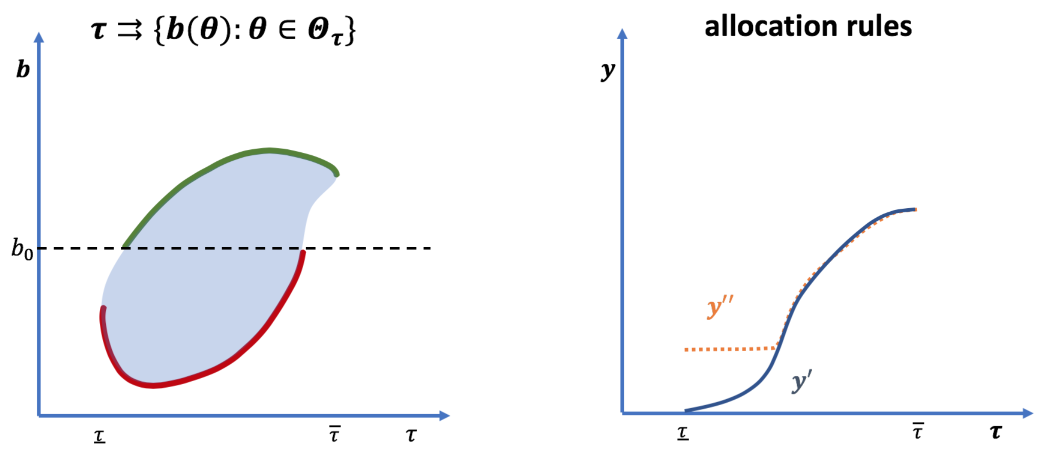

Figure 5: The first panel shows the correspondence from effective types to the range of benefits of allocation that types with those effective types may have. The red (green) line highlights the worst-case (best-case) benefit level below (above) b_0 for a given effective type. The distance between b_0 and the red (green) line is the largest allocative loss to full (no) allocation. The second panel shows two implementable allocation rules for this type-space.

a benchmark b_0: $f_\theta(y) = \max\{(1-y)(b(\theta)-b_0), y(b_0 - b(\theta))\}$.

Can the allocation rule y' shown on the right of Figure 5 be an optimal allocation rule for the type space depicted to the left? The answer is no.

By increasing the allocation awarded to a set of low effective types (call it $\mathcal{T}_0$) as shown in y'', the screener introduces larger type I errors on these low types. But suppose this increase is small so that the worst-case loss due to errors on these low effective types, $\max_{\tau\in\mathcal{T}_0} \max_{\theta\in\Theta_\tau} y''(\underline{\tau})(b_0 - b(\theta))$, is less than $R(y')$. Therefore the move from y' to y'' cannot increase total loss.

But since types in $\mathcal{T}_0$ get a higher allocation for free, types with effective types just above that need only show a little more evidence to receive their allocation. This logic extends all the way up, so everyone in $\mathcal{T} \setminus \mathcal{T}_0$ shows less evidence under y'' to obtain the same allocation as they did under y'. Therefore, the worst-case loss under y'' is less than under y', so the latter is not optimal

This logic extends to the general setting, and some useful comparative statics arise from it.

Terminology Let $y \in \mathcal{I}$, $y_0 \in \text{Range}(y)$, and let $\mathcal{T}_0 = \{\tau \in \mathcal{T} : y(\tau) = y_0\}$. The allocation rule y *attains its worst-case payoff on* y_0 if $\sup_{\tau\in\mathcal{T}_0} R_\tau(y) = R(y)$.

The lowest allocation level, in y (i.e., $y(\underline{\tau})$) is the *free option* (as $n(\underline{\tau}) = 0$ in the optimal standard of proof). Note that there may of course be $\tau > \underline{\tau}$ which are also allocated the free option.

Next, we say *screening costs matter* if

$$\min_{y\in\mathcal{I}}\max_{\tau\in\mathcal{T}}\max_{\theta\in\Theta_\tau}\alpha f_\theta(y(\tau(\theta)))+(1-\alpha)h_\theta(n(\tau(\theta)))>\min_{y\in\mathcal{I}}\max_{\tau\in\mathcal{T}}\max_{\theta\in\Theta_\tau}\alpha f_\theta(y(\tau(\theta))).$$

That is, dropping the losses from screening costs strictly reduces the screener's loss. Screening costs would not matter, for example, if the optimal allocation rule is constant, independent of α.

LEMMA **2.** *Let y be a screener-optimal allocation rule, and let y_0 denote the free option in y.*

1. *The screener attains his worst-case payoff on y_0.*
2. *If screening costs matter, then $y_0 > 0$.*
3. *The size of the free option is equal across all screener-optimal allocation rules,*
4. *The size of the free option is weakly decreasing in the screener's weight on allocation errors.*

Proof. Let $\mathcal{T}_0^y$ be the set of effective types which take the free option. Let y_1 be the next lowest allocation level after the free option, if y is not constant, and let $y_1 = y_0$ otherwise. We show the statements in turn.

Part 1: Suppose for contradiction that y does not attain its worst-case payoff at y_0, i.e., $R(y) - \max_{\tau\in\mathcal{T}_0^y} R_\tau(y) > \epsilon > 0$. By this assumption, $y_1 > y_0$.

There exists a $0 < \delta < y_1 - y_0$ such that $|f_\theta(y_0) - f_\theta(y_0+\delta)| < \frac{\epsilon}{2}$ for all θ with $\tau(\theta) \in \mathcal{T}_0^y$. Consider the new allocation rule y', where $y'(\tau) \equiv \max\{y(\tau), y_0+\delta\}$ for all $\tau \in \mathcal{T}$. This reduces the evidence every effective type in $\mathcal{T} \setminus \mathcal{T}_0^y$ has to show by $\epsilon\underline{\mathcal{T}}$, under the optimal standard of proof that implements y'. Since h_θ is strictly increasing for all θ, $R_\tau(y') < R_\tau(y)$ for $\tau \in \mathcal{T} \setminus \mathcal{T}_0^y$. Therefore, $R(y') < R(y)$, a contradiction.

Part 2: Next, suppose that screening costs matter and for contradiction that $y_0 = 0$. Let $L \equiv \min_{y'\in\mathcal{I}}\max_{\tau\in\mathcal{T}}\max_{\theta\in\Theta_\tau}\alpha f_\theta(y'(\tau(\theta)))$, and let $\mathcal{Y} \subset \mathcal{I}$ be the corresponding set of optimizers. Let $\mathcal{T}_0 = \{\tau \in \mathcal{T} : y'(\tau) = 0 \forall y' \in \mathcal{Y}\}$.

First, we show that $\mathcal{T}_0$ must be empty.

Suppose first $\mathcal{T}_0$ is non-empty and $\mathcal{T}_0^y \subset \mathcal{T}_0$. Then the worst-case loss that the screener faces among agents that take the free option is no more than L, which by the assumption that costs matter is smaller than $R(y)$. This contradicts the first part of this theorem.

Suppose next that $\mathcal{T}_0$ is non-empty and $\mathcal{T}_0 \subset \mathcal{T}_0^y$. By the first part of the theorem, combined with the assumption that costs matter, the screener faces worst-case loss on $\mathcal{T}_0^y \setminus \mathcal{T}_0$

and a loss of no more than L on $\mathcal{T}_0$. Note that $f \equiv \max_{\tau \in \mathcal{T}_0^y \setminus \mathcal{T}_0} \max_{\theta \in \Theta_\tau} f_\theta$ is strictly quasiconvex, being the max of strictly quasiconvex functions. Combined with the fact that $y'(\tau) > 0$ for all $y' \in \mathcal{Y}$, this implies that there exists a region $[0, \epsilon]$ for some $0 < \epsilon < y_1$ where f is strictly decreasing. Finally, let $0 < \delta < \epsilon$ be such that $|f_\theta(y_0) - f_\theta(y_0 + \delta)| < R(y) - L$. Then the allocation rule y' where $y'(\tau) \equiv \{y(\tau), y_0 + \delta\}$ for all τ has a lower worst-case loss loss than y by the same arguments as before, a contradiction to the optimality of y.

This proves that $\mathcal{T}_0$ is empty, which means there exists a $y' \in \mathcal{Y}$ such that $y'(\underline{\tau}) > 0$. Let $\underline{y} \equiv y'(\underline{\tau})$. It can be shown, by similar arguments to those above, that y', where $y'(\tau) \equiv \max\{y(\tau), \underline{y}\}$ has strictly lower worst-case loss for the screener. This implies that $y_0 > 0$.

Part 3: Next, the fact that the size of the free option is equal across all screener optimal allocation rules follows as a simple consequence of part 4 of the theorem: if y' is an optimal allocation rule for a screener with weight α' on allocation errors and $\alpha \geq \alpha'$, then the free option is weakly larger under y' than under y. Applying this statement to $\alpha' = \alpha$ proves part 3. Therefore, we only need to prove part 4.

Part 4: Suppose for contradiction that the free option is larger under y than under y'. Let $0 < \epsilon < y(\underline{\tau}) - y'(\underline{\tau})$. Let $\tilde{y} = \max\{y(\underline{\tau}) - \epsilon, y'\}$ and let $\tilde{n}$ be the corresponding optimal standard of proof. Let $\tilde{\mathcal{T}}_0 \subset \mathcal{T}$ be the set of effective types who are allocated the free option under $\tilde{y}$. Let $R^s \equiv R(y)\frac{\alpha'}{\alpha}$. Let $R' \equiv \max_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} \alpha' f_\theta(y'(\tau(\theta))) + (1 - \alpha')h_\theta(n'(\tau(\theta)))$ be the type-α' screener's worst-case loss under y' (n' is the optimal standard of proof corresponding to y').

The optimality of y over y' for the type-α screener and the fact that $\frac{\alpha'}{\alpha}(1 - \alpha) < 1 - \alpha < 1 - \alpha'$ implies that $R^s \leq R'$.

Suppose θ is such that $\tau(\theta) \in \tilde{\mathcal{T}}_0$. Then,

$$
\begin{aligned}
\alpha' f_\theta(\tilde{y}(\tau(\theta))) &< \alpha' \max\{f_\theta(y(\tau(\theta))), f_\theta(y'(\tau(\theta)))\} \\
&\leq \max\{R^s, R'\} \\
&= R',
\end{aligned}
$$

where the first inequality follows from the strict quasi-convexity of f_θ. This implies $R' > \max_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} \alpha' f_\theta(\tilde{y}(\tau(\theta))) + (1 - \alpha')h_\theta(\tilde{n}(\tau(\theta)))$, contradicting the optimality of y' for the type-α' screener. □

The property that the size of the free option in an optimal allocation rule is decreasing in α has an intuitive explanation. Due to the restriction imposed by implementability, higher effective types receive weakly more allocation than low effective types, even if the screener would use a non-monotonic rule under complete information. Given the incentive to reduce screening costs for types that receive larger allocations, the only type of distortion that would happen on low effective types is an upward distortion. When the screener cares less about minimizing costs and more about minimizing allocation errors, she gains less from making these upward distortions and thereby chooses a smaller free option.

Interpretation Lemma 2, part 2 has the clearest implication for the design of costly screens when the agent's costs matter to the screener: even if introducing certain costs induce advantageous selection, their value is always enhanced by giving away some part of the allocation for free. The positive effect of reducing the evidence procurement burden on types who receive larger allocations initially outweighs the allocation errors that may arise, if any.

This is to be interpreted as a result pertaining to agents whose observable characteristics do not pin down their optimal allocation. Indeed, it is only for this pool of borderline candidates that costly screening is valuable. By conditioning on these characteristics, the screener stymies free entry of undeserving agents who may otherwise try to take advantage of the free option.

An example would be a welfare program applicant whose readily observable traits neither outright qualify or disqualify him from receiving benefits. Suppose the applicant's willingness to participate in a costly appeals process is typically an indication of greater need. One design the screener may adopt is to ask all borderline candidates to either participate in the appeals process or give up on their application. Lemma 2 suggests that the program should instead have options for such an applicant to either (1) walk away with a small amount of benefits (or alternatively, lower quality in-kind transfers), or (2) participate in a less stringent appeals process to receive more benefits.[9]

4 Main Results: Simplicity, Effectiveness and Robustness of Delegation

This section uses the characterization of the unconstrained screener's optimal allocation rules from Section 3 to solve the policymaker's problem. Section 4.1 considers the case where the screener is biased towards minimizing allocation errors, while Section 4.2 considers the case where the screener is biased towards keeping down screening costs. The results aim to answer the following questions: When is the optimal delegation set just an interval, and when is it a more complex set? Can these restrictions be designed in such a way that the screener

[9] Alternatively, in the case of innovation approval processes, Lemma 2 suggests including a regulatory sandbox, where ex-ante qualified firms can test or market their products in a restricted setting and show costly evidence to obtain more approval from regulators.

decides to offer the policymaker's favorite menu? How much does the policymaker need to know about the screener's preferences to design these restrictions?

Terminology Recall that α and α_P are the weights that the screener and policymaker place on allocation errors, respectively.

The screener is *biased towards reducing allocation errors* if $\alpha > \alpha_P$. And the screener is *biased towards reducing screening costs* if $1-\alpha > 1-\alpha_p$.

If the set of screener optimal allocation rules in $\mathcal{I}_Y$ and policymaker optimal allocation rules in $\mathcal{I}$ intersect for some Y, then delegation set Y is *perfect*.

Delegation is futile if full delegation is optimal for the policymaker but not perfect. In other words, there are no restrictions that the policymaker can set in place to improve the screener's behavior.

A delegation set Y is *simple* if it is connected. A particular simple delegation set is *full delegation*: $Y = [0,1]$. The policymaker chooses *no delegation* if Y is a singleton.

Delegation is complex if full delegation is not optimal for the policymaker, but no simple delegation set strictly improves (i.e., reduces the policymaker's loss as given by eq. (2)) over full delegation.

For a fixed α_P, a delegation set Y is *robustly optimal* if it solves the policymaker's problem either for any $\alpha > \alpha_P$ or all $\alpha < \alpha_P$. This captures the notion that all the policymaker needs to know about the screener's preferences is contained in just the direction of his bias.

For a fixed α_P, a delegation set Y is *robustly improving* if it is a strict improvement over full delegation for some $\alpha > \alpha_P$ ($\alpha < \alpha_P$) and a weak improvement over full delegation for all $\alpha > \alpha_P$ ($\alpha < \alpha_P$). Full delegation is dominated for a policymaker who knows the direction of the screener's bias and has robustly improving delegation sets at his disposal.

4.1 Screener is Biased Towards Reducing Allocation Errors

When $\alpha > \alpha_P$, Theorem 1 shows that the optimal delegation set has a very simple form. The policymaker simply restricts the screener from awarding smaller allocations than she would ever award, were she directly in charge of screening. That is, if y_P is optimal for the policymaker, she sets a floor at $y_P(\underline{\tau})$. With this restriction, y_P is also optimal for the screener in the set of allowable allocation rules, and this result holds regardless of how much larger α might be than α_P.

THEOREM **1.** *Suppose the screener is biased towards reducing allocation errors. Let y_P be an optimal allocation rule for the policymaker, and let $Y \equiv [y_P(\underline{\tau}), 1]$. Y is a simple, perfect and robustly optimal delegation set for the policymaker.*

Proof. Suppose the policymaker uses delegation set Y. Let $y_0 \equiv \min Y$. If y_P is constant, then Y is a singleton, and the theorem follows trivially. Suppose y_P is not constant, and let y_1 be the second lowest allocation under y_P.

The screener is worse off by choosing an allocation rule y with a larger free option than y_0 than she would be by choosing the best allocation rule with a free option of this size.[10]

To complete the proof, we have to show that y_P is optimal for the screener among all allocation rules y with $y(\underline{\tau}) = y_0$.

Let L denote the policymaker's worst-case loss when the allocation rule y_P is used. Let $\tau \in \mathcal{T}$.

$$\begin{aligned}\frac{\alpha_P}{\alpha} R_\tau(y_P) &= \max_{\theta \in \Theta_\tau} \alpha_P f_\theta(y_P(\tau(\theta))) + \frac{\alpha_P}{\alpha}(1-\alpha) h_\theta(n_P(\tau(\theta))) \\ &\leq L.\end{aligned}$$

The inequality is strict for any τ for which $y_P(\tau) > y_0$, since $\frac{\alpha_P}{\alpha}(1-\alpha) < 1 - \alpha_P$ and $n_P(\tau) > 0$. Next, the inequality is an equality for any τ such that $y_P(\tau) = y_0$, since $n_P(\tau) = 0$ and $h_\theta(0) = 0$. Therefore $R(y_P) = \frac{\alpha}{\alpha_P} L$, but the worst-case loss for the screener only occurs at the free option.

Letting $\mathcal{T}_0$ be the set of effective types for which $y_P = 0$, the above shows that for any allocation rule y with $y(\tau) = 0$ on $\mathcal{T}_0$, $R(y) \leq R(y_P)$.[11]

The only remaining case to consider address is whether or not an allocation rule y with $y(\underline{\tau}) = y_P(\underline{\tau})$ but $y(\tau) > y_P(\tau)$ for some $\tau \in \mathcal{T}_0$ can have strictly lower worst-case loss than y_P for the screener.

Suppose for contradiction that this is indeed the case. Then let $\tau_1 \in \mathcal{T}_0$ denote the smallest effective type for which $y(\tau) > y(\underline{\tau})$.

By assumption, the worst-case loss on $\tau < \tau_1$ is strictly less than $\frac{\alpha}{\alpha_P} L$ for the screener, under y. And since y and y_P coincide on this region, the worst-case loss for the policymaker is also strictly under L on this region.

Then let $0 < \epsilon < \min y(\tau_1), y_1 - y_0$ be such that $\alpha_P f_\theta(y_0 + \epsilon) < L$ for all θ with $\tau(\theta) < \tau_1$.

[10] This follows from entirely symmetric arguments to those in the proof of Lemma 2, part 4: were this not true, the policymaker would be strictly better off using some other allocation rule with a larger free option as well, contradicting the optimality of y_P.

[11] Therefore, if y_P was fully separating on $\mathcal{T}$, for example, this is enough to conclude that it is an optimal choice for the screener in $\mathcal{I}_Y$.

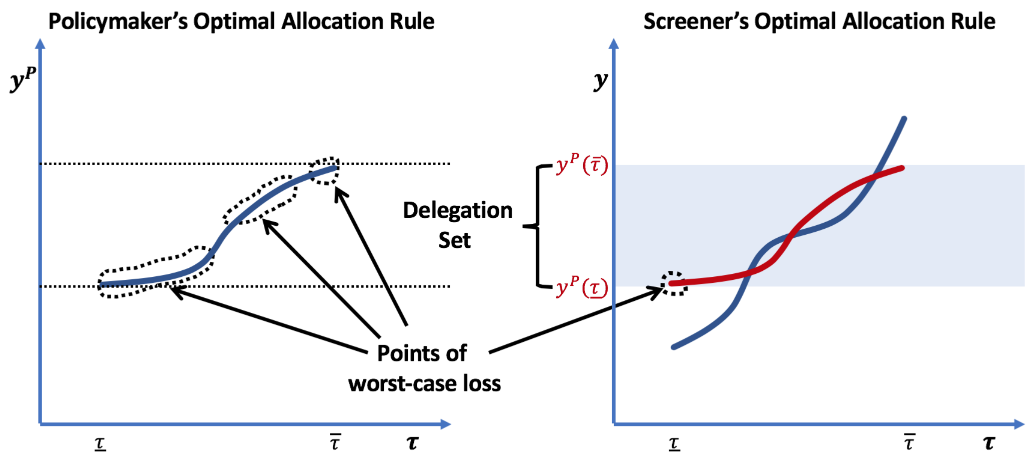

Figure 6: The blue curve on the left panel is a policymaker optimal allocation rule, y_P. The blue curve on the right panel is an optimal allocation rule for the unconstrained screener. Restricted to allocation rules with co-domain $[y_P(\underline{\tau}), y_P(\overline{\tau})]$, the screener finds y_P to be optimal (red curve).

For θ with $\tau(\theta) \in [\tau_1, \overline{\tau}] \cap \mathcal{T}_0$,

$$
\begin{aligned}
\alpha_P f_\theta(y_0 + \epsilon) &< \max\{\alpha_P f_\theta(y_0), \alpha_P f_\theta(y(\tau_1))\} \\
&\le \max\{\alpha_P f_\theta(y_0), \alpha_P f_\theta(y(\tau(\theta)))\} \\
&\le \max\{L, \frac{\alpha_P}{\alpha} \alpha f_\theta(y(\tau(\theta)))\} \\
&\le \max\{L, \frac{\alpha_P}{\alpha} R(y)\} \\
&\le \max\{L, \frac{\alpha_P}{\alpha} R(y_P)\} \\
&= L
\end{aligned}
$$

This raises the contradiction that $y'_P = \max\{y_0 + \epsilon, y_P\}$ has lower worst-case loss for the policymaker than y_P.

Therefore y_P is optimal for the screener. $\square$

The intuition for this result is easiest to see in the case where the policymaker's optimal allocation rule is fully separating in effective types, as shown in Figure 6. The screener's unconstrained optimal allocation rules typically have a smaller free option (less upward distortion on low effective types) than the policymaker's optimal allocation rule (Lemma 2, part 4). Once the screener is forced to use an allocation rule with a larger free option, the screener has no incentive to increase this upward distortion even further.

Next, when adopting the policymaker's favorite screening rule, the screener only faces worst-case losses at the lowest effective type. Elsewhere, up to a re-scaling of his objective function, the screener faces smaller costs than the policymaker. So any other approval rule with a free option of the same size can only have a weakly larger worst-case loss for the screener, as he is already weighed down by the loss faced on the lowest effective type.

This last argument does not work when the screener optimal rule pools multiple effective types at the lowest allocation level. The proof extends the argument to cover this case as well.

4.2 Screener is Biased Towards Reducing Screening Costs

The next theorem shows that the results on the simplicity, effectiveness and robustness of delegation are quite dramatically reversed when $\alpha < \alpha_P$. To begin with, the optimal delegation set is typically a non-interval set and fails to be robustly optimal.

But the problem runs deeper. In fact, *any* restriction that improves over full delegation is a non-interval set. And moreover, *any* improvement over full delegation for some $\alpha < \alpha_P$ is guaranteed to worsen the policymaker's outcome if the screener's weight on allocation errors was instead some other $\alpha' < \alpha_P$. There are therefore no easy restrictions that the policymaker can impose to improve her outcome over the status-quo option of fully delegating to the screener.

Theorem **2.** *Suppose the screener is biased towards reducing screening costs.*

1. *If the unconstrained screener's optimal allocation rules do not coincide with the Policymaker's, delegation is either complex or futile.*

2. *No delegation set is robustly improving.*

We prove the first part of this theorem sketch the second, leaving the formal proof to the appendix.

Proof of Theorem 2, part 1. Let Y be any delegation set. Let y be the screener's unconstrained optimal allocation rule.

If Y were an interval, then either $Y \subset [0, y(\underline{\tau}))$, $Y \subset (y(\underline{\tau}), 1]$ or $y(\underline{\tau}) \in Y$.

Claim: in each of these cases, the policymaker is weakly better off under full delegation rather than using delegation set Y.

First, suppose that $Y \subset [0, y(\underline{\tau}))$. Let y' be the constant allocation rule with $y'(\tau) = \max Y$ for all $\tau \in \mathcal{T}$. Let $y'' \in \mathcal{I}_Y$ be any other rule; note y'' must have a smaller free option, since it is weakly increasing. At any point where $y''(\tau) = y'(\tau)$, $R_\tau(y') < R_\tau(y'')$, since the screening costs under y' are 0. If $y''(\tau(\theta)) < y'(\tau(\theta))$, then:

$$\begin{aligned}\alpha f_\theta(y'(\tau(\theta))) &< \max\{\alpha f_\theta(y(\tau(\theta))), \alpha f_\theta(y''(\tau(\theta)))\} \\ &\leq \max\{R(y), R(y'')\} \\ &= R(y'')\end{aligned}$$

This shows that $R(y') < R(y'')$. Therefore, the screener's unique optimal allocation rule in $\mathcal{I}_Y$ is y', a constant delegation rule. But full delegation is better than no delegation for the policymaker:

$$\max_\tau \max_{\theta\in\Theta_\tau} \alpha_P f_\theta(y(\tau)) + (1-\alpha_P)h_\theta(n(\tau)) < \frac{\alpha_p}{\alpha}R(y) < \frac{\alpha_p}{\alpha}R(y') = \max_{\theta\in\Theta_\tau} \alpha_P f_\theta(y'(\tau)).$$

Next suppose $Y \subset (y(\underline{\tau}), 1]$. Take any allocation rule y' in $\mathcal{I}_Y$. The policymaker prefers the screener's chosen rule, y, to y'.[12] Therefore, full delegation is better than Y for the policymaker.

Finally, suppose that Y is such that $y(\underline{\tau}) \in Y$. Then, for any y' with $y'(\underline{\tau}) < y(\underline{\tau})$, the allocation rule y'' with $y''(\tau) \equiv \max\{y(\underline{\tau}), y'(\tau)\}$ is a strict improvement for the screener. This means the screener will only choose those allocation rules with free option weakly larger than $y(\underline{\tau})$, leaving the policymaker weakly worse off.

This proves the claim that full delegation is weakly better than any interval restrictions. Now if the unconstrained screener's optimal allocation rules do not coincide with the policymaker's, the screener implements a sub-optimal allocation rule for the policymaker under full delegation. Therefore, delegation is either futile (e.g., if full-delegation is optimal) or complex (if full delegation is not optimal). □

Next, we argue that any delegation set which might improve over full delegation for some $\alpha > \alpha_P$ can make matters worse for some other $\alpha' > \alpha_P$.

To rule out irrelevant delegation sets, we identify a property that all strict improvements over full delegation share. The necessary characterization comes from the proof of Theorem 2, part 1: the only way that the policymaker might improve over full delegation is by using a delegation set which excludes some open interval (c, d) about $y(\underline{\tau})$, as shown in Figure 7.

There are two ways that removing an interval (c, d) can backfire. Consider the optimal allocation rule for a screener who places weight $\alpha' \approx 0$ on allocation errors. Such a screener is primarily concerned about allocation costs, and therefore uses a (nearly) flat allocation rule at some level y_0. The two cases to consider are $y_0 \in (c, d)$ and $y_0 \geq d$, as pictured in Figure 8.[13]

[12]This follows by the same arguments as in the proof of Theorem 1, but with the role of the policymaker and the screener reversed: if the policymaker strictly preferred some allocation rule with a larger free option, then the screener prefers some other rule with a larger free option as well, contradicting the optimality of y

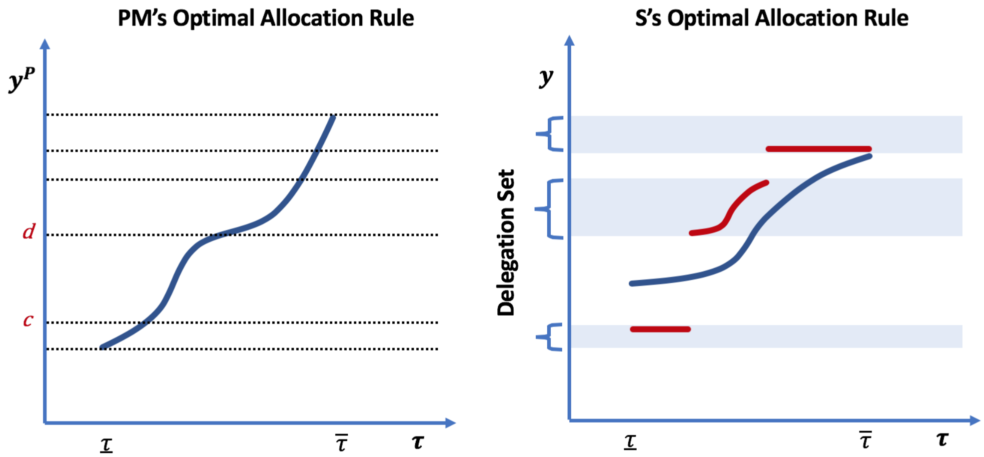

Figure 7: The blue curve on the left panel is a policymaker optimal allocation rule, y_P. The blue curve on the right panel is an optimal allocation rule for the unconstrained screener, y. When restricted to allocation rules with co-domain in the delegation set shown, the screener finds the red curve to be optimal.

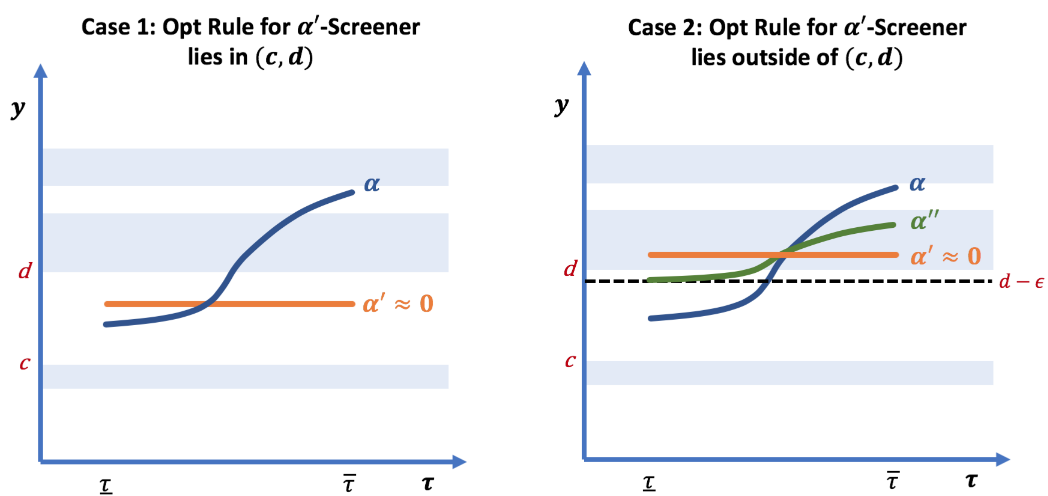

Figure 8: On the left is the case where an $\alpha' \approx 0$ type screener uses a flat allocation rule in (c, d). On the right is the case where an α' type screener uses a rule in $[d, 1]$. In the second case, there exists some other type, a screener with weight α'' on allocation errors, whose unconstrained optimal allocation rule has a free option in (c, d) that is very close to d.

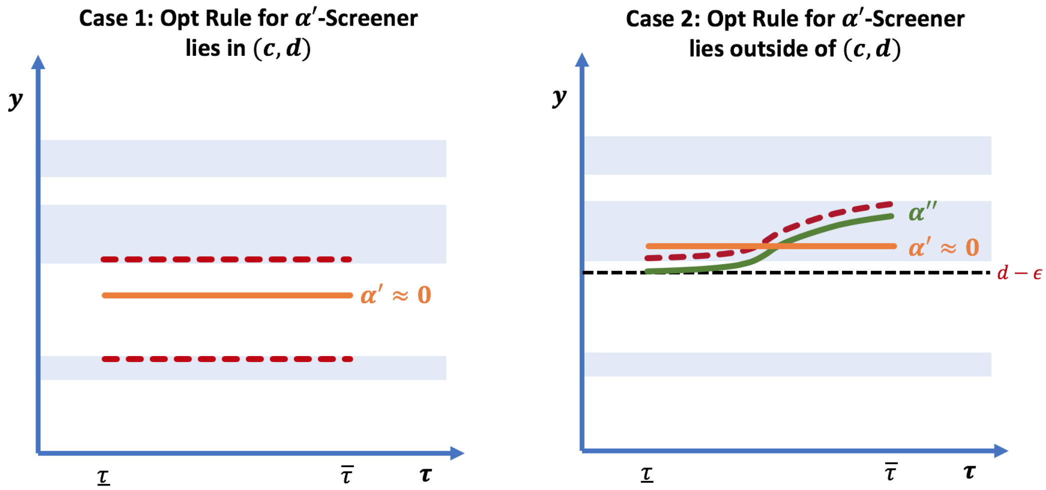

Figure 9: In the first case, an α' type screener moves chooses one of the two dashed red lines after the policymaker's restriction is imposed. In the second case, an α'' type screener moves to the dashed red line after the restriction. The free option is larger in this new allocation rule, making the policymaker worse off.

If $y_0 \in (c, d)$, then removing this interval forces the type α' screener to pick the flat allocation rule at c or d instead (see Section 4.2). But among the constant allocation rule, the type α' screener chooses the one that minimizes allocation errors. Therefore, removing the interval (c, d) makes the policymaker worse off when facing a type α' screener.

If $y_0 \geq d$, it may or may not belong to the policymaker's delegation set. If it does belong, then the type α' screener does not serve as the necessary counterexample.

However, the fact that there is a screener type who would choose a free option in (c, d) and another who chooses a free option in $[d, 1]$ implies that there is a screener type who would choose a free option for every level in between. [14] So an unconstrained screener who would choose a free option of size $d - \epsilon$ would react to the policymaker's delegation rule by choosing an allocation rule with a free option of size d or greater rather than drop down to c, as shown in Section 4.2. But increasing the size of the free option makes the policymaker worse off.

Theorem 2, part 1 says that even if the policymaker is certain about α, simple interval restrictions would not improve outcomes from her perspective. Theorem 2, part 2 implies

(by Lemma 2, part 3).

[13] Lemma 2, part 4 rules out the case where $y_0 < c$: since the optimal allocation rule of a type α screener has a free option in (c, d), the free option for the type α' screener must be weakly larger.

[14] The mapping from screener types to the size of the free option under optimal allocation rules is continuous. This follows from the continuity of the screener's value function in α.

that knowing only the direction of the screener's bias is grossly insufficient for the policymaker to be able to effect positive change.

4.3 Screener Private Information

The results so far discuss properties of min-max optimal delegation sets when the policymaker knows the relative weight, α, that the screener places on allocation errors. In some cases, the policymaker might not know α or even the direction of the screener's bias. Here we solve for the policymaker's min-max optimal delegation set, taking worst case over both the screener's type and the agent's type.

Let $\mathcal{X}(Y, \alpha)$ denote the set of screener optimal menus when her private type is α and she is restricted to menus with co-domain Y. The policy maker chooses a delegation set Y to solve:

$$\inf_{Y} \min_{y \in \mathcal{X}(Y,\alpha)} \max_{\alpha \in [0,1]} \max_{\theta \in \Theta} \alpha_P f_\theta(y(n_y(\theta))) + (1-\alpha_P) h_\theta(n_y(\theta))$$

A delegation set which solves this objective is *min-max optimal with respect to the screener's type.*

Say a a delegation set Y is *undominated* if there is no delegation set Y' that gives the policymaker weakly lower loss for some α and θ and strictly lower loss for some α' and θ'.

COROLLARY **1.** *Let y_P be an optimal allocation rule for the policymaker, and let $Y = [y_P(\underline{\tau}, 1)]$. Y is undominated and min-max optimal with respect to the screener's type for the policymaker.*

Proof. Theorem 1 and Theorem 2 immediately imply that Y is undominated: Y is optimal for any $\alpha > \alpha_P$, and any Y' that gives the policymaker a lower loss than Y does for some $\alpha < \alpha_P$ must be strictly worse than Y for some other $\alpha' < \alpha_P$.

To see that this strategy is min-max optimal with respect to the screener's type, note first that for any delegation set Y', the optimal choice of a screener for whom $\alpha = 0$ is some constant allocation rule y. Consider another screener $\alpha' \in (0, \alpha_P)$ for whom the optimal allocation rule is y' when the delegation set is Y'. The optimality of y' for such a screener implies that there is some $\theta' \in \Theta$ such that for all $\theta \in \Theta$,

$$\alpha' f_{\theta'}(y(n_y(\theta'))) + (1-\alpha')h_{\theta'}(n_y(\theta')) \geq \alpha' f_{\theta'}(y'(n_{y'}(\theta))) + (1-\alpha')h_\theta(n_{y'}(\theta))$$

Since y is flat, $n_y = 0$, so this inequality becomes,

$$\alpha'(f_{\theta'}(y(n_y(\theta'))) - f_{\theta'}(y'(n_{y'}(\theta)))) \geq (1-\alpha')h_\theta(n_{y'}(\theta))$$

which implies the same inequality holds when replacing α' with $\alpha_P > \alpha'$. In other words, the policymaker prefers y' to y. Therefore the worst case loss for the policymaker for any delegation set occurs when the choice of allocation rule is being made by a screener with weight zero on allocation errors. Then the policymaker's min-max optimal delegation set is one which allows such a screener to choose the best constant allocation rule. Since a floor at $y_P(\underline{\tau})$ does not interfere with such a screener's choice, it is min-max optimal with respect to the screener's type for the policymaker. □

4.4 Implications for Regulatory Governance

Theorem 1 and Theorem 2 present a stark dichotomy in the difficulty of delegation between the cases where the screener is more biased towards minimizing approval errors versus screening costs. Depending on the context in which they find themselves, policymakers may or may not be able to effectively delegate the task of costly screening.

However, policymakers often have many coarse tools for altering the incentives of regulators, even if they abstain from offering high powered contracts out of concern for adverse reactions.

For example, Congress delegates authority to regulatory agencies and partially shapes its objectives by determining the agency's mandate. They can make the agency's stated objective to ensure the safety of new innovations or they may declare that the agency should consider both consumer protection and costs of regulated business.[15]

We can think of a coarse tool as the policymaker having the ability to increase or decrease α by some $x\%$. How should the policymaker use this additional lever in conjunction with her capacity to restrict the space of allocation rules?

Result. *A robust policy for the policymaker is to increase α by $x\%$ and set a floor at the minimum of her optimal allocation rule, $y_P(\underline{\tau})$.*

If $(1 - \frac{x}{100})\alpha$ is greater than α_P, then this policy allows the screener to achieve her first best allocation rule.

But even if $(1 - \frac{x}{100})\alpha < \alpha_P$, the screener is better off. The floor is nonbinding by Lemma 2 part 4, so effectively, the policymaker maintained full delegation while decreasing the extent of preference divergence between herself and the screener. Again, she ends up better off. This result is driven by the fact that preference divergence is inconsequential in one direction but not the other.

An implication for innovation approval processes is that regulators should not be incentivized to care about screening costs for regulated firms, even if lawmakers themselves care. Emphasizing consumer protection in a regulatory agency's mandate while limiting its capacity to avoid approval errors is a better form of governance than holding it accountable for

[15] To give another example, policymakers can automate certain processes to alter the incentives of bureaucrats to patiently review multiple appeals decisions by the same applicant.

approval costs directly. In the latter situation, lawmakers are forced to either suffer the costs of full delegation or make restrictions that have the potential of backfiring if the agency's preferences are misestimated.

5 Extensions: The Case for Setting Floors

This section extends the model in different directions. Section 5.1 considers fully general differences in objectives between the policymaker and the screener. Section 5.2 considers the model of delegated screening where the players are Bayesian and minimize expected rather than worst-case loss. Section 5.3 studies how results change in a model where additional conditions are imposed on how the screener breaks inferences. A consistent conclusion arises in all three extensions: setting floors is a simple and robust policy that always dominates full delegation. Section 5.4 shows that such floors can backfire in a model of delegated information acquisition, distinguishing it from delegated screening.

5.1 General Divergence in Objectives

Section 2 and Section 4 model the disagreement between the policymaker and the screener as arising solely from placing different weights on allocation errors and screening costs. This may be a salient friction in many contexts where costly screening is delegated, as discussed in the introduction.

But other disagreements are also possible. Consumer protection regulators may be more worried about type I errors than type II errors, where policymakers might have a more balanced weight on both. Alternatively, street-level bureaucrats may be intrinsically biased toward extending welfare benefits to applicants, while policymakers might prefer them to be more discerning.[16]

We can extend the model to capture fully general preference divergence, imposing only that both the policymaker and screener are still averse to screening costs.

Let $\{f_\theta\}$ and $\{h_\theta\}$ denote the allocation loss functions for the screener as before. But now let $\{f_\theta^P\}$ and $\{h_\theta^P\}$ denote the corresponding loss functions for the policymaker, where f_θ is strictly quasiconvex and h_θ is strictly increasing for all $\theta \in \Theta$.

As before, the screener chooses $y \in \mathcal{I}_Y$ to minimize:

$$R(y) \equiv \max_{\tau \in \mathcal{T}} R_\tau(y) \equiv \max_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} \alpha f_\theta(y(\tau(\theta))) + (1-\alpha) h_\theta(n(\tau(\theta))) \tag{3}$$

[16] See Prendergast (2007) for a discussion on various forms of biases that bureaucrats may have relative to policymakers.

The policymaker chooses a closed delegation set $Y \subset [0,1]$ to minimize:

$$\min_{y \in \mathcal{X}(Y)} R^P(y) \equiv \min_{y \in \mathcal{X}(Y)} \max_{\tau \in \mathcal{T}} \max_{\theta \in \Theta_\tau} \alpha_P f^P_\theta(y(\tau(\theta))) + (1-\alpha_P) h^P_\theta(n(\tau(\theta))) \tag{4}$$

We refer to this as *the general preference divergence formulation of the model.*

Even at this generality, the policymaker has a policy at her disposal that can only make her weakly better off and does not require knowing anything about the screener's preferences. She simply considers her optimal allocation rule and sets a floor at the minimum of the range for that rule. If the floor is not binding, there was no harm in applying it. If it is binding, she is better off for it.

PROPOSITION **1.** *Consider the general preference divergence formulation of the model. Let y_P be an optimal allocation rule for the policymaker. Setting a floor at $y_P(\underline{\tau})$ is a weak improvement over full delegation for the policymaker.*

5.2 Delegated Costly Screening in the Expected Loss Case

Consider the expected loss analog of the problem, where both the screener and the policymaker with a common prior, letting p_θ be the probability that the agent is of type θ. Let $u, u^P : \mathbb{R}_+ \to \mathbb{R}_+$ be increasing, concave functions.

The screener chooses $y \in \mathcal{I}_Y$ to minimize:

$$B(y) \equiv \sum_{\tau \in \mathcal{T}} \sum_{\theta \in \Theta_\tau} p_\theta u(\alpha f_\theta(y(\tau(\theta))) + (1-\alpha) h_\theta(n(\tau(\theta)))) \tag{5}$$

The policymaker chooses a closed delegation set $Y \subset [0,1]$ to minimize:

$$\min_{y \in \mathcal{X}(Y)} \sum_{\tau \in \mathcal{T}} \sum_{\theta \in \Theta_\tau} p_\theta u^P(\alpha_P f^P_\theta(y(\tau(\theta))) + (1-\alpha_P) h^P_\theta(n(\tau(\theta)))) \tag{6}$$

Notice here that we are allowing the preference divergence between the policymaker and the screener to be fully general, as in Section 5.1. We call this *the expected loss formulation of the model.*

5.2.1 Relationship to the Max-Min Problem

There are two ways of going from the expected-loss to the max-min formulation of the problem, and each has different interpretations.

One way is to take the curvature of u to infinity so that the players are infinitely risk-averse in the max-min case. While they have a prior over the state space, the worst-case loss is particularly salient in determining their choices. For example, regulators and politicians may worry about the poor optics of very deserving candidates being given insufficient care under a public healthcare program or undeserving candidates getting too many benefits. Media sources and political opponents to the program may have incentives to cherry-pick evidence about the program's efficacy and report the most egregious errors.

Another way of reaching the model of in Section 2 is to simply replace the double summation and the priors with a max over the realization of agent types. This can be interpreted as a model of ambiguity and may be an appropriate way to think about an agency making decisions in the face of large uncertainties, as is the case when regulators are charged with approving novel innovations.

5.2.2 Floors are (Weakly) Robustly Improving in the Linear Costs Case

Optimal delegation need not be simple in the expected loss case, even when $\alpha > \alpha_P$. The easiest way to see this is to consider a three type case. Even if the cap and floor are set optimally, the screener's optimal allocation to the middle type generally will not coincide with the policymaker's choice.

Indeed, proving optimality results with any generality is difficult in the expected loss case due to the richness of the policymaker's action space. However, when certain additional conditions on the loss functions hold, setting floors in the same way as Proposition 1 will at least be robustly improving.

Proposition 2. *Consider the expected loss formulation of the model. Suppose the allocation loss functions $\{f_\theta\}_{\theta\in\Theta}$ and $\{f^P_\theta\}_{\theta\in\Theta}$ are strictly convex*[17]*, and $\{h_\theta\}_{\theta\in\Theta}$ and $\{h^P_\theta\}_{\theta\in\Theta}$ are linear*[18]*, and both the policymaker and screener are risk-neutral (i.e., u and u^P are linear). Let y_P be an optimal allocation rule for the policymaker. Setting a floor at $y_P(\underline{\tau})$ is a weak improvement over full delegation for the policymaker.*

While the model focuses on examples where max-min analysis is applicable, many situations of delegated screening involve neither ambiguity nor worst-case payoffs. Proposition 2 shows that the same simple and robust policies can be used to improve outcomes in this case as well.

[17]The stronger assumption of strict convexity rather than strict quasiconvexity is required to retain the notion of a 'right' allocation level for any given effective type. Quasiconvexity suffices in the max-min model since the max of quasiconvex loss functions of type in some Θ_τ is quasiconvex. In the expected loss case, the needed condition is that convex functions are closed under addition, so that the probability weighted sum of errors is convex.

[18]The linearity in screening costs ensures that changes to the allocation for some effective type have constant marginal benefits in terms of lower costs for higher type. This implies that the level of allocation awarded to lower types does not affect the choice of allocation to award to higher types.

5.3 Additional Constraints on Implementation

One feature of the max-min analysis is that the screener may have many optimal allocation rules. In our solution concept, we assumed that these ties are broken in favor of the policymaker. While retaining this assumption, I consider two additional considerations in the screener's choice, aimed at thinning out the size of the indifference sets.

One 'natural' way for the screener to choose a constrained optimal allocation rule is to pick the one closest to the allocation rule he was using under full delegation. That is, if the screener uses allocation rule y under full delegation, he picks a rule in $\arg\min_{y' \in \mathcal{X}(Y)} d_\infty(y, y')$ when delegation set Y is imposed. We refer to this as the *least-change criterion*.

Another consideration for the screener may be to select among the optimal allocation rules only those which are not pointwise dominated in terms of worst-case loss. That is, the screener will not pick $y \in \mathcal{X}(Y)$ if there exists a $y' \in \mathcal{X}(Y)$ such that $R_\tau(y') \leq R_\tau(y)$ for all $\tau \in \mathcal{T}$ and $R_{\tau'}(y') \leq R_{\tau'}(y)$ for some $\tau' \in \mathcal{T}$. This is the *undominated criterion*.

We can modify Proposition 1 to allow for these stronger solution concepts.

Result. *Consider the general preference divergence formulation of the model. Let y_P be an optimal allocation rule for the policymaker. Setting a floor at $y_P(\underline{\tau})$ is a weak improvement over full delegation for the policymaker even when the screener applies the least-change or undominated criteria.*

The proof of this result follows directly from the proof of Proposition 1, which shows that if y is optimal for the screener under full delegation, $y' \equiv \{y_P(\underline{\tau}), y\}$ is optimal for the screener in $\mathcal{I}_{[y_P(\underline{\tau}),1]}$. Notice y' is the closest allocation rule to y, so the least-change criterion does not affect the result. Next, it is straightforward to check that if y is undominated in the sense of pointwise worst-case loss, y' is undominated as well. This implies that the policymaker's preferred undominated allocation rule in $\mathcal{X}([y_P(\underline{\tau}), 1])$, is a (weak) improvement over y as well.

Consider again the model of Section 2 where the policymaker and screener diverge only on the weights they place on allocation errors and screening costs. When the screener is biased towards minimizing screening costs, introducing new conditions on the screener's choice does not change the results on the complexity or cost of delegation or the non-existence of robustly improving policies. But when the screener is biased towards minimizing allocation errors, the policymaker may no longer be able to induce the screener to utilize her favorite allocation rule.

However, Lemma 2, part 4 combined with Section 5.3 guarantee that simple and robust improvements continue exist in the case where $\alpha > \alpha_P$. Unless the policymaker's and unconstrained screener's optimal allocation rules already overlap, these improvements are strict.

5.4 Paying for Versus Eliciting Information

Setting floors is a sound prescription for policymakers: it is a robust policy with fully general preference divergence, in max-min and Bayesian settings, and without regularity assumptions on the agent's type space. These results are driven by the assumption that the intermediary has to elicit information from agents in an incentive compatible way.

This sets the theory of delegated screening apart from delegated information acquisition, where the intermediary can acquire signals about the agent at a cost to her and the policymaker (but the costs are not contractible). There, the assumptions on the signal structure drive the form of delegation, precluding more unified results.

Appendix C shows a simple counterexample that illustrates why a result analogous to Proposition 2 does not hold. That is, suppose the intermediary chooses costly signals about the agent's type. Setting floors at the minimum allocation level that the policymaker would award can strictly worsen her outcome from full delegation.

6 Conclusion

When policymakers create a new welfare program or an innovation approval process, various intermediaries invariably take control of certain aspects of the design. Costly screening of agents downstream is typically one of those aspects. This paper introduces a model of delegated screening to study when and how a policymaker can influence a screener's choices.

Two themes emerge from the analysis. First, aligning the behavior of a screener who cares more about reducing errors is easy (i.e., simple restrictions are optimal and robust) while dealing with screeners who care more about costs is not. Second, the policymaker can always set floors in a way that does not require knowing anything about the screener's preferences and weakly dominates setting no restrictions. Together, these results suggest that a policymaker should take any measures available to reduce a screener's concern for an agent's costs and increase his concern for errors. Even if such a measure runs the risk of pushing the screener's objective further away from the policymaker's, divergence in this direction is easier to fix through simple restrictions than divergence in the other.

While this paper studies the question of delegated screening, there are many other settings where policymakers rely on intermediaries who have a richer contract space than they do. For example, lawmakers delegate sentencing to judges who commit to their own preferred mappings from evidence to sentences. How should lawmakers set allowable ranges of punishments if they disagree with judges on the importance of fairness versus deterrence? Other examples include firms delegating the design of incentive schemes and promotion schedules to managers who can observe worker output; or a state government delegating the design of a procurement process to a local government that can verify product quality.

The intractability of an expected utility framework may have been a limiting factor in studying questions about delegated mechanism design. The min-max approach taken in this paper may also prove useful in these other settings.

References

Alatas, Vivi, Ririn Purnamasari, Matthew Wai-Poi, Abhijit Banerjee, Benjamin A Olken, and Rema Hanna, "Self-targeting: Evidence from a field experiment in Indonesia," *Journal of Political Economy*, 2016, *124* (2), 371–427.

Amador, Manuel and Kyle Bagwell, "Regulating a monopolist with uncertain costs without transfers," Technical Report, Working paper 2016.

Baron, David P and Roger B Myerson, "Regulating a monopolist with unknown costs," *Econometrica: Journal of the Econometric Society*, 1982, pp. 911–930.

Bergemann, Dirk and Karl H Schlag, "Pricing without priors," *Journal of the European Economic Association*, 2008, *6* (2-3), 560–569.

Carroll, Gabriel, "Robustness and linear contracts," *American Economic Review*, 2015, *105* (2), 536–63.

_ , "Robustness and separation in multidimensional screening," *Econometrica*, 2017, *85* (2), 453–488.

_ , "Robustness in mechanism design and contracting," *Annual Review of Economics*, 2019, *11*, 139–166.

Celik, Gorkem, "Mechanism design with collusive supervision," *Journal of Economic Theory*, 2009, *144* (1), 69–95.

Chade, Hector and Natalia Kovrijnykh, "Delegated information acquisition with moral hazard," *Journal of Economic Theory*, 2016, *162*, 55–92.

Deshpande, Manasi and Yue Li, "Who is screened out? application costs and the targeting of disability programs," *American Economic Journal: Economic Policy*, 2019, *11* (4), 213–48.

Faure-Grimaud, Antoine, Jean-Jacques Laffont, and David Martimort, "Collusion, delegation and supervision with soft information," *The Review of Economic Studies*, 2003, *70* (2), 253–279.

Frankel, Alexander, "Aligned delegation," *American Economic Review*, 2014, *104* (1), 66–83.

Guo, Yingni and Eran Shmaya, "Robust Monopoly Regulation," *arXiv preprint arXiv:1910.04260*, 2019.

Holmstrom, Bengt, "On the theory of delegation," Technical Report, Discussion Paper 1980.

Holmstrom, Bengt R, "On Incentives and Control in Organizations." PhD dissertation 1977.

Hurwicz, Leonid and Leonard Shapiro, "Incentive structures maximizing residual gain under incomplete information," *The Bell Journal of Economics*, 1978, pp. 180–191.

Keiser, Lael R, "State bureaucratic discretion and the administration of social welfare programs: The case of social security disability," *Journal of Public Administration Research and Theory*, 1999, *9* (1), 87–106.

Laffont, Jean-Jacques and David Martimort, "Collusion and delegation," *The Rand Journal of Economics*, 1998, pp. 280–305.

Lemley, Mark A and Carl Shapiro, "Probabilistic patents," *Journal of Economic Perspectives*, 2005, *19* (2), 75–98.

Masur, Jonathan S, "Costly Screens and Patent Examination," *Journal of Legal Analysis*, 2010, *2* (2), 687–734.

Mookherjee, Dilip and Masatoshi Tsumagari, "The organization of supplier networks: effects of delegation and intermediation," *Econometrica*, 2004, *72* (4), 1179–1219.

Prendergast, Canice, "The motivation and bias of bureaucrats," *American Economic Review*, 2007, *97* (1), 180–196.

Rassenfosse, Gaétan De and Adam B Jaffe, "Are patent fees effective at weeding out low-quality patents?," *Journal of Economics & Management Strategy*, 2018, *27* (1), 134–148.

Szalay, Dezsö, "The economics of clear advice and extreme options," *The Review of Economic Studies*, 2005, *72* (4), 1173–1198.

Zeckhauser, Richard J, "Strategic sorting: The role of ordeals in health care," Technical Report, National Bureau of Economic Research 2019.

A Infinite Type Spaces

This sections gives the auxiliary assumptions needed to extend the results to the case where Θ is infinite. I give a proof of the analogs of Lemma 2 parts 1 and 2, which is the crucial step in translating the other theorems.

In addition to the assumptions in Section 2, assume:

1. $\mathcal{T}$ is compact.
2. The family of allocative loss functions $\{f_\theta\}_{\theta\in\Theta}$ is uniformly bounded, equicontinuous.
3. The correspondence $\tau \rightrightarrows \{f_\theta : \theta \in \Theta_\tau\}$ is upper-hemicontinuous, under the topology induced by the sup-norm metric.
4. The family of loss functions for costly evience provision, $\{h_\theta\}_{\theta\in\Theta}$, is uniformly bounded.

Theorem **3.** *Let y be a screener-optimal allocation rule, and let y_0 denote the free option in y.*

1. *The screener attains his worst-case payoff on y_0.*
2. *If screening costs matter, then $y_0 > 0$.*

Proof. Suppose y is an optimal allocation rule and let $\mathcal{T}_0$ be the set of effective types which take the free option (i.e., the allocation requiring no costly evidence), $y(\underline{\tau})$. Suppose for contradiction that y does not attain its worst-case payoff for $\tau \in \mathcal{T}_0$, i.e., $R(y) - sup_{\tau\in\underline{T}} R_\tau(y) > \epsilon > 0$. There are two cases to consider.

Case 1: Suppose $y(\underline{\tau})$ is an isolated point in Range(y), so that a $\delta' > 0$ radius ball around $y(\underline{\tau})$ does not intersect with Range(y).

By equicontinuity, there exists a $\delta < \delta'$ sufficiently small such that $|f_\theta(y(\underline{\tau})) - f_\theta(y(\underline{\tau}) + \delta)| < \frac{\epsilon}{2}$ for all θ with $\tau(\theta) \in \mathcal{T}_0$. Consider the new approval rule $y_1(\tau) = \max\{y(\tau), y(\underline{\tau}) + \delta\}$ for all $\tau \in \mathcal{T}$. This reduces the evidence every effective type in $\mathcal{T} - \mathcal{T}_0$ has to show by $\epsilon\underline{T}$. Therefore $R(y_1) < R(y)$, a contradiction.

Case 2: Suppose $y(\underline{\tau})$ is not an isolated point in Range(y). Then $\mathcal{T}_0$ is a closed set; let $\tau' = \max \mathcal{T}_0$.

Claim: there is an open interval $O \subset \mathbb{R}$ such that $O \cap \underline{T}$ contains τ', and such that $R_\tau(y(\underline{\tau})) < R(y) - \epsilon$. Suppose not and consider the correspondence $F : \tau \rightrightarrows \{f_\theta : \theta \in \Theta_\tau\}$. Then there exists a sequence of effective types $\tau_n \in \tau$, such that $\lim_{n\to\infty} \tau_n = \tau'$, and $f_{\theta_n}(\underline{\tau}) \geq R(y) - \epsilon$ for some θ_n such that $\tau(\theta_n) = \tau_n$. Then by equicontinuity and uniform boundedness, there exists a subsequence of $\{f_{\theta_n}\}_{n\in\mathbb{N}}$ that converges uniformly to some limit

f. By upper-hemicontinuity of F, $f \in F(\tau')$, so $f(y(\underline{\tau})) \geq R(y) - \epsilon$, i.e., there is some θ with $\tau(\theta) = \tau'$ such that $f_\theta(y(\underline{\tau})) \geq R(y) - \epsilon$, a contradiction.

Therefore, there exists a closed ball centered at τ' with radius $\delta'' > 0$ such that the claim above holds. Let $\epsilon' \equiv y(\underline{\tau} + \delta'') - y(\underline{\tau})$ ($\epsilon' > 0$ by assumption of Case 2). Appealing to equicontinuity again, let $0 < \delta < \epsilon'$ be such that $|f_\theta(y(\underline{\tau})) - f_\theta(y(\underline{\tau}) + \delta)| < \frac{\epsilon}{2}$ for all $\theta \in \Theta$. The rule $y_2(\tau) = \max\{y(\tau), y(\underline{\tau}) + \delta\}$ for all $\tau \in \mathcal{T}$ has $R(y_2) < R(y)$, contradicting the optimality of y. $\square$

B Optimality of Delegated Screening

Section 2 describes a game where the policymaker chooses sets of allocations and the screener chooses menus satisfying this co-domain restriction. This section shows that if we instead suppose that both the policymaker and screener can write complete contracts, the policymaker's optimal deterministic grand contract can be implemented using an indirect mechanism of the form in Section 2.

To establish this result, we appeal to a form of the revelation principle that applies when one of the participants in the mechanism (i.e., the screener) has commitment power, and there are both hidden actions and hidden types (i.e., the agent's evidence choice and type).

The timing of contracting is as follows.

1. The policymaker chooses a communication protocol: M_s and M_a are the sets of all message strategies for the screener and agent, and $q : M_s \times M_a \to [0, 1]$ maps message strategies to an allocation for the agent.[19]

2. Before either player participates in the communication protocol, the screener can offer a side contract to the agent. The screener chooses a direct mechanism $m_s : \Theta \times \mathbb{R}_+ \to M_s$ that maps the agent's reported type and costly evidence to a messaging strategy that the screener will subsequently follow in the policymaker's protocol.

3. The agent chooses what to report and how much costly evidence to generate for the screener.

4. The agent then chooses a message strategy $m_a \in M_a$.

Note the following features of this contracting environment.

[19] As usual, communication may be dynamic and among the players or between the players and the policymaker at various stages. Here, we treat any messages sent by the policymaker as part of the communication protocol and assume q is measurable with respect to what she knows. We restrict attention to deterministic mechanisms.

We allow the screener to offer a contract after the policymaker offers hers. This timing seems natural for many instances and only serves to increase the agency problems faced by the policymaker.[20]

Next, the outcomes of the policymaker's communication protocol depend only on the messages sent by the screener and agent, since the agent's costly action is hidden to her. On the other hand, the screener observes and can contract on the evidence generated by the agent. There is no loss of generality in restricting attention to direct mechanisms for the screener, as the usual revelation principle applies at this stage.

To analyze this game, note that the screener's messaging strategy $m_s \in M_s$ is already determined before the players join the communication protocol. The agent chooses the messaging strategy $\arg\max_{m_a \in M_a} q(m_s, m_a)$, as its payoff is strictly increasing in allocation. Therefore, $Y = \{\max_{m_a \in M_a} q(m_s, m_a) | m_s \in M_s\}$ is the set of possible allocations any agent may receive in equilibrium through the policymaker's communication protocol.

The policymaker's communication protocol composed with the screener's direct mechanism induces a mapping from the agent's type and costly evidence generation to Y. Indeed, the screener may choose any incentive compatible and individually rational mechanism with co-domain Y. Therefore, restricting attention to the game of delegated screening in 2 is without loss of generality.

C Comparison to Costly Information Acquisition

When the intermediary pays to learn exogenous information about the agent's private type, setting floor as in Proposition 1 or Proposition 2 can leave the policymaker strictly worse off. This highlights an important difference between this and the model of delegated screening.

This section gives a simple Bayesian example to illustrate this difference (compared to Proposition 2). In place of eliciting information about the agent's type, an intermediary (rather than screener) can choose costly signals, or tests, to learn this information.

There are four agent types, $\theta_1, \theta_2, \theta_3, \theta_4$. If an agent of type θ_i is allocated y units, the policymaker and screener face an allocation error loss of $f_{\theta_i}(y) = |y - \frac{i}{5}|$. Types θ_1 and θ_2 occur with equal probability and with probability 0.99.

There are two tests, t_1 and t_2. Test t_1 reveals that the agent's type is a member of $\{\theta_1, \theta_2, \{\theta_3, \theta_4\}\}$. Test t_2 reveals that the agents type is a member of $\{\{\theta_1, \theta_2\}, \theta_3, \theta_4\}$.

The loss each player when the intermediary awards allocation y to a type θ_i agent is $f_{\theta_i}(y)$ plus that player's costs for the tests the intermediary used. Test t_1 costs 1 to the policymaker, whereas test t_2 costs 2. Both tests cost $c > 0$ to the intermediary: assume that

[20] Mookherjee and Tsumagari (2004) model collusion similarly in a setting where a principal contracts with two agents, one of which can offer a side contract with the other.

c is small enough that the unconstrained intermediary would strictly prefer to use either test over making a decision with no test at all.

The policymaker would not use either test if she was in charge of picking allocations. She would instead pick an allocation just slightly to the right of $\frac{1}{2}(\frac{1}{5}+\frac{2}{5})=\frac{3}{10}$, say $\frac{3}{10}+\delta$.

An optimal policy for the intermediary, however, would be to first use test t_1, and stop there 99% of the time (it is immaterial to the example whether or not c is small enough such that he would continue on to use t_2 in the 1% this case arises).

Suppose that the policymaker places a floor on the allocations that the intermediary can award at the lowest level she would ever award, $\frac{3}{10}+\delta$. Since this ties the intermediary's decision in the case where the state is θ_1 or θ_2, t_2 is now a strictly more valuable test than t_1. Moreover, its value is not much diminished from before (assuming without loss of generality that δ is sufficiently close to 0), so he would still prefer using t_2 over no test at all.

Note that after imposing the floor, the policymaker's loss from tests increased from at most $0.99+0.01\times 2$ to 2. The change in the accuracy of the intermediary's decision is not enough to compensate for this. Therefore, by setting a floor in this way, the policymaker is worse off than under full delegation.

D Omitted Proofs

Proof of Theorem 2, part 2. Suppose that for some $\alpha_p>\alpha$, there exists a delegation set which is a strict improvement over full delegation. Let y be the screener's unconstrained optimal allocation rule. Following the proof of Theorem 2, part 1, this rule must exclude an interval (c,d) containing $y(\underline{\tau})$. There are two cases to consider. Let y_0 denote the optimal constant allocation rule (by strict quasiconvexity of every f_θ, $f(a)\equiv\sup_{\theta\in\Theta}f_\theta(a)$ is also strictly quasiconvex and therefore has a unique minimum).

Case 1: $y_0\in(c,d)$. This implies that a screener with a weight α' close enough to 0 would have chosen an allocation rule in (c,d) if unconstrained, but is forced to choose between the allocaiton rule that is constant at c or an allocation rule with free option of size d instead. Full delegation is superior to both outcomes for such a screener.

Case 2: $y_0\geq d$. By Lemma 2 part 3, the correspondence from the screener's weight on allocation errors to the sizes of free options in optimal menus is singleton valued.

Claim: this mapping is continuous. To see this, let $\alpha'<\alpha$ be two weights the screener may place on allocative errors, and let y and y' be corresponding optimal allocation rules. By Lemma 2 part 4, $y(\underline{\tau})\leq y'(\underline{\tau})$. Let $\epsilon>0$. By Lemma 2, there is an $\epsilon'>0$ such that if the free option under an allocation rule y'' is greater than or equal to $y(\underline{\tau})+\epsilon$, then the worst-case loss to a type-α screener increases by at least ϵ'. By continuity of the value function in α, there is a $\delta>0$ such that if $\alpha-\alpha'<\delta$, then $\frac{\alpha}{\alpha'}$ times the worst-case loss of the α'-type screener under y' minus the worst-case loss of the α-type screener under y is less than ϵ'. This difference is an upper-bound on difference in worst-case loss for the type α screener under y versus y'. This means if $\alpha - \alpha' < \delta$, $y'(\underline{\tau}) - y(\underline{\tau}) < \epsilon$, proving the claim

Since the mapping is continuous, by the intermediate value theorem, then there exist a sequence $\alpha'_n > \alpha$ for n large enough, who would choose a free option of size $d - \frac{1}{n}$ if unconstrained. For large enough n, the type-α'_n screener would choose a free option of size d or larger under the policymaker's restriction. Since the policymaker increases the size of the free option for such a screener, she increases her worst-case error on this type.

By Lemma 2, part 4, the case where $y_0 \leq c$ is an impossibility.

Therefore case 1 and case 2 collectively show that robust improvement are not possible.

□

Lemma **3.** *Let y be an optimal allocation rule for the screener, and suppose that a floor at $y_0 \in [0,1]$ is set. Then the allocation rule y' where $y'(\tau) \equiv \max\{y_0, y(\tau)\}$ is optimal for the screener.*

Proof. The statement is clearly true if the floor in non-binding, so suppose it is binding.

Note first that $R(y) < R(y')$ by Lemma 2, part 3. Therefore, the worst-case loss on the set of effective types $\mathcal{T}_0$ for which $y' = y$ is larger under y', since the worst-case loss on $\mathcal{T} \setminus \mathcal{T}_0$ is smaller under y'. To show an allocation rule $y'' \geq y_0$ cannot have a lower worst-case loss than y', it suffices to show it cannot have a lower worst-case loss on $\mathcal{T}_0$.

If the worst-case loss is greater under y' than under y at some $\tau \in \mathcal{T}_0$, by strict quasi-convexity of $\max_{\theta \in \Theta_\tau} f_\theta$, increasing allocation to this effective type further only increases worst-case loss. If worst-case loss at some $\tau \in \mathcal{T}_0$ is smaller under y' than under y, then $R_\tau(y') < R_\tau(y) < R(y)$.

This implies that y'' cannot have lower worst-case loss than y': if $\tau \in \mathcal{T}_0$ is such that $R_\tau(y'') < R_\tau(y')$, then $R_\tau(y'') < R(y) < R(y'')$, so this does not affect the worst-case loss at y''. The worst-case loss for y'' happens at an effective type τ where $R_\tau(y') \leq R_\tau(y'')$. □

Proof. By Lemma 3 (and using the same notation as in the proof there), y' (with the floor set at $y_P(\underline{\tau})$) is an optimal allocation rule for the screener. In the case where the policymaker's floor is binding, the policymaker faces lower worst-case loss on every effective type for whom the allocation stayed the same between y' and y. Among those effective types for whom the floor was binding, the policymaker's loss due to screening costs is reduced to 0 under y'. Finally, consider those effective types τ for whom worst-case losses are larger under y' than under y. The allocation error losses must be weakly larger still under y_P, since $y(\tau) < y'(\tau) \leq y_P(\tau)$, and the screening costs are weakly larger at τ under y_P as well. But the worst-case loss here is less than or equal to the worst-case loss of the policymaker under y_P. By the optimality of y_P for the policymaker, the policymaker does not attain her worst-case loss under y' or y at τ.

Since y' is therefore a weak improvement over y for the policymaker, the screener's choice of allocation rule (which is the one the policymaker prefers among those optimal for the screener) after imposing the floor is a weak improvement over y. $\square$

Lemma **4.** *Consider the Bayesian formulation of the delegated screening problem and suppose the assumptions in Proposition 2 hold. Let y be an optimal allocation rule for the screener, and suppose that a floor at $y_0 \in [0,1]$ is set. Then the allocation rule y' where $y'(\tau) \equiv \max\{y_0, y(\tau)\}$ is optimal for the screener.*

Proof of Lemma 4. We first make some definitions to save on notation. Let $k(\theta) = p_\theta \alpha$. Since $h_\theta(n(\theta))$ is linear, let $l_\theta n(\theta) \equiv p_\theta(1-\alpha)h_\theta(n(\theta))$.

If the floor is not binding, the statement holds.

Suppose the floor is binding. Let $\mathcal{T}_0$ be the set of effective types for whom $y(\tau) < y'(\tau)$. Let $y'' \in \mathcal{I}_{[y_0,1]}$ be any implementable allocation rule such that $y'(\tau) \geq y_0$ for all τ. There are two cases to consider.

Case 1: $y''(\tau) = y_0$ for all $\tau \in \mathcal{T}_0$. Define the allocation rule y''' to be equal to y on $\mathcal{T}_0$ and y'' on $\mathcal{T} \setminus \mathcal{T}_0$.

In this case,

$$\begin{aligned}
B(y') - B(y'') &= \sum_{\tau \in \mathcal{T} \setminus \mathcal{T}_0} \sum_{\theta \in \Theta_\tau} k_\theta(f_\theta(y'(\tau(\theta))) - f_\theta(y''(\tau(\theta)))) + l_\theta(n'(\tau(\theta)) - n''(\tau(\theta))) \\
&= \sum_{\tau \in \mathcal{T} \setminus \mathcal{T}_0} \sum_{\theta \in \Theta_\tau} k_\theta(f_\theta(y(\tau(\theta))) - f_\theta(y'''(\tau(\theta)))) + l_\theta(n(\tau(\theta)) - n'''(\tau(\theta))) \\
&= \sum_{\tau \in \mathcal{T}} \sum_{\theta \in \Theta_\tau} k_\theta(f_\theta(y(\tau(\theta))) - f_\theta(y'''(\tau(\theta)))) + l_\theta(n(\tau(\theta)) - n'''(\tau(\theta))) \\
&= B(y) - B(y''') \\
&\geq 0
\end{aligned}$$

The second equality follows from the fact that n differs from n' by a constant on $\mathcal{T} \setminus \mathcal{T}_0$, and n'' differs from n''' by the same constant. The third equality follows from the facts that y''' and y agree on $\mathcal{T}_0$. The last inequality follows from the optimality of y for the screener.

The screener therefore weakly prefers y' to y''.

Case 2: $y''(\tau) > y_0$ for some $\tau \in \mathcal{T}_0$. Let τ_0 denote the smallest effective type for which this holds. Let τ_1 be the largest effective type for which $y(\tau_0) = y(\tau_1)$. Note that by definition, τ_0 and τ_1 are in $\mathcal{T}_0$. Finally, let τ_2 be the smallest effective type for which $y''(\tau_2) = y''(\tau_1)$.

To recap: $\tau_0 \leq \tau_2 \leq \tau_1$, y is constant on $[\tau_0, \tau_1]$, and y'' is constant on $[\tau_2, \tau_1]$. Therefore, both y and y'' are constant on $[\tau_2, \tau_1]$. Next, $y''(\tau) > y_0$ on this region, while $y(\tau) < y_0$ for $\tau \in [\tau_2, \tau_1]$. Denote the constant values these allocation rules take on this interval by y''_c and y_c . Finally, compared to the values they take on $[\tau_2, \tau_1]$, $y''(\tau)$ is strictly smaller for $\tau \in [0, \tau_2)$ while y is strictly larger for $\tau \in (\tau_1, \overline{\tau}]$.

Let $\Theta_{[\tau_2,\tau_1]}$ denote the set of all types θ with $\tau(\theta) \in [\tau_2, \tau_1]$. Define $\mu(x) \equiv \sum_{\theta \in \Theta_{[\tau_2,\tau_1]}} k_\theta f_\theta(x)$ for all $x \in [y_c, 1]$. Note μ is strictly convex, since it is a sum of strictly convex functions. Define $\eta(x) \equiv \sum_{\theta \in \Theta_{[\tau_2,\tau_1]}} l_\theta \tau_3 x$ for all $x \in [y_c, 1]$ Next, let $\Theta_{(\tau_1,\overline{\tau}]}$ denote the set of all types θ with $\tau(\theta) > \tau_1$. Let τ_3 denote the next lowest effective type to τ_3. Define $\phi(x) \equiv \sum_{\theta \in \Theta_{(\tau_1,\overline{\tau}]}} l_\theta(\tau_3 - \tau_2)x$ for $x \in [y_c, 1]$. Notice that both η and ϕ are linear in x.

Consider the allocation rule y_ϵ, which is equal to y on $[\underline{\tau}, \overline{\tau}] \setminus [\tau_2, \tau_1]$ and equal to $y_c + \epsilon$ on $[\tau_2, \tau_1]$, where ϵ is sufficiently small so that y_ϵ is still non-decreasing. Then it follows from the definitions that

$$B(y_\epsilon) - B(y) = (\eta(y_c + \epsilon) - \eta(y_c)) + (\phi(y_c + \epsilon) - \phi(y_c)) + (\mu(y_c + \epsilon) - \mu(y_c)) \geq 0$$

. by the optimality of y.

But now consider the the allocation rule y''_ϵ, which is equal to y'' on $[\underline{\tau}, \overline{\tau}] \setminus [\tau_2, \tau_1]$ and equal to $y_c - \epsilon$ on $[\tau_2, \tau_1]$, where ϵ is the same as before and we assume without loss of generality that it was chosen to be sufficiently small so that y''_ϵ is still non-decreasing. Then

$$\begin{aligned} B(y'') - B(y''_\epsilon) &= (\eta(y''_c) - \eta(y''_c - \epsilon)) + (\phi(y''_c) - \phi(y''_c - \epsilon)) + (\mu(y''_c) - \mu(y''_c - \epsilon)) \\ &= (\eta(y_c + \epsilon) - \eta(y_c)) + (\phi(y_c + \epsilon) - \phi(y_c)) + (\mu(y''_c) - \mu(y''_c - \epsilon)), \end{aligned}$$

where the equality follows from the linearity of ϕ and η. Note, however, that by strict convexity, $(\mu(y''_c) - \mu(y''_c - \epsilon)) > (\mu(y_c + \epsilon) - \mu(y_c))$. This implies $B(y'') - B(y''_\epsilon) > 0$, so y'' is not an optimal allocation rule among those that have a floor at y_0.

Case 1 and case 2 together imply (since optimal allocation rules exist) that y' is an optimal allocation rule for the screener among those that have a floor at y_0. □

Proof of Proposition 2. Let y_P be a policymaker optimal allocation rule, and let $y_0 \equiv y_P(\underline{\tau})$. We also carrying over all the notation from Lemma 4.

y' is optimal for the screener by the same Lemma 4. We want to show that the policymaker is weakly prefers y' to y.

We prove this statement under the assumption that y_P and y are fully separating in effective types, as the case where there is some pooling is handled precisely as in case 2 of the proof of Lemma 4.

Let $\tau \in \mathcal{T}_0$. Let y_ϵ be a perturbation which is equal to y everywhere but τ, where it equals $y(\tau) + \epsilon$. Let $y_{P,\epsilon}$ be a perturbation which is equal to y_P everywhere but τ, where it equals

$y_P - \epsilon$. Assume that $\epsilon > 0$ is sufficiently small so that y_ϵ and $y_{P,\epsilon}$ are non-decreasing (such an ϵ exists by the assumption of full separation). The reduction in loss due to corresponding evidence costs on types with effective types in $(\tau, \overline{\tau}]$ is the same when moving from y to y_ϵ, as when moving from $y_{P,\epsilon}$ to y_P. The same is true for the increase in loss due to evidence costs on types with effective type τ. However, the increase in loss due to total allocation errors is larger moving from $y_{P,\epsilon}$ to y_P than moving from y to y_ϵ, by strict convexity of the allocation loss functions. The optimality of y_P implies that the the policymaker prefers y_P to $y_{P,\epsilon}$. Then the policymaker should prefer y_ϵ to y.

By applying this argument to each τ in $\mathcal{T}_0$, starting by perturbing the allocation function upward to y_0 on the largest effective type in this set, we see that each of these moves is beneficial to the policymaker. Therefore, the policymaker prefers y' to y. $\square$