DISCUSSION PAPER SERIES

DP12572

COMMITMENT VS. FLEXIBILITY WITH COSTLY VERIFICATION

Marina Halac and Pierre Yared

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

COMMITMENT VS. FLEXIBILITY WITH COSTLY VERIFICATION

Marina Halac and Pierre Yared

Discussion Paper DP12572
Published 05 January 2018
Submitted 05 January 2018

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

COMMITMENT VS. FLEXIBILITY WITH COSTLY VERIFICATION

Abstract

A principal faces an agent who is better informed but biased towards higher actions. She can verify the agent's information and specify his permissible actions. We show that if the verification cost is small enough, a threshold with an escape clause (TEC) is optimal: the agent is allowed to choose any action below a threshold or request verification and the efficient action if sufficiently constrained. For higher costs, however, the principal may require verification only for intermediate actions, dividing the delegation set. TEC is always optimal if the principal cannot commit to inefficient allocations following the verification decision and result.

JEL Classification: D02, D82

Keywords: optimal delegation, costly verification, escape clause

Marina Halac - mhalac@columbia.edu
Columbia University and CEPR

Pierre Yared - pyared@columbia.edu
Columbia University

Acknowledgements
We thank Kyle Bagwell, Heski Bar-Isaac, Roland Benabou, Yeon-Koo Che, Eddie Dekel, Wouter Dessein, Simone Galperti, Johannes Horner, Navin Kartik, Francine Lafontaine, Bart Lipman, Andrey Malenko, Marco Ottaviani, Alessandro Pavan, Debraj Ray, Ken Shotts, Paolo Siconolfi, Joel Sobel, Jeff Zwiebel, and various seminar and conference audiences for helpful comments. We also thank Sebastian Di Tella and Niko Matouschek for valuable discussions of the paper. Weijie Zhong provided excellent research assistance.

Commitment vs. Flexibility with Costly Verification*

Marina Halac† Pierre Yared‡

December 26, 2017

Abstract

A principal faces an agent who is better informed but biased towards higher actions. She can verify the agent's information and specify his permissible actions. We show that if the verification cost is small enough, a threshold with an escape clause (TEC) is optimal: the agent is allowed to choose any action below a threshold or request verification and the efficient action if sufficiently constrained. For higher costs, however, the principal may require verification only for intermediate actions, dividing the delegation set. TEC is always optimal if the principal cannot commit to inefficient allocations following the verification decision and result.

*We thank Kyle Bagwell, Heski Bar-Isaac, Roland Benabou, Yeon-Koo Che, Eddie Dekel, Wouter Dessein, Simone Galperti, Johannes Hörner, Navin Kartik, Francine Lafontaine, Bart Lipman, Andrey Malenko, Marco Ottaviani, Alessandro Pavan, Debraj Ray, Ken Shotts, Paolo Siconolfi, Joel Sobel, Jeff Zwiebel, and various seminar and conference audiences for helpful comments. We also thank Sebastian Di Tella and Niko Matouschek for valuable discussions of the paper. Weijie Zhong provided excellent research assistance.

†Columbia University and CEPR. Email: mhalac@columbia.edu.

‡Columbia University and NBER. Email: pyared@columbia.edu.

1 Introduction

The delegation problem was first formally analyzed by Holmström (1977, 1984) and has since been studied by an extensive literature. This problem concerns a principal who faces a better informed but biased agent. Transfers between the parties are infeasible.[1] The principal simply chooses an allowable set of actions from which the agent can select. Optimal delegation reflects a fundamental tradeoff between commitment and flexibility: commitment is valuable to limit biased decisions by the agent, whereas flexibility is valuable to let the agent utilize his private information about the efficient action.

Delegation is central to a variety of applications. In organizations, headquarters delegate investment decisions to division managers who have superior information about project benefits but also a desire for a larger empire. When designing fiscal policy institutions, society delegates spending decisions to a government which can better assess current social needs but tends to overweigh the value of present spending. Real-world delegation rules, however, not only define restrictions on agents' actions, as in the canonical model, but also typically specify "escape clauses." Headquarters impose limits on how much managers can invest, yet managers can request a revision of their budget by providing suitable project documentation.[2] Fiscal rules put constraints on the government's spending and deficit, yet they also include escape clause provisions for the government to break these constraints given evidence of an exceptional circumstance.[3] Rules in practice make use of escape clauses because agents' private information can often be verified, albeit at a cost. Once an agent's information is ascertained, the appropriate action by the agent can be selected.[4]

In this paper, we introduce costly state verification, as in the seminal work of Townsend (1979), into a general delegation framework. Our main goal is to explore how the principal's ability to verify the agent's information affects optimal delegation. How does the principal structure escape clauses to optimally resolve her commitment-versus-flexibility tradeoff while minimizing verification costs?

We examine a principal-agent model with no transfers in which the agent is biased towards higher spending relative to the principal. The agent's private information, or *type*, concerns the value of spending; a higher type corresponds to a higher marginal value of spending for both the principal and the agent. Following the literature, we build upon a setting in which, absent verification, an optimal delegation rule is a *threshold*, allowing the

[1] In various applications, like those described below, (contingent) transfers may be ruled out because of institutional reasons or ethical considerations.

[2] See, for example, Ross (1986) and Taggart (1987).

[3] See Schaechter et al. (2012).

[4] Other applications of the delegation problem where rules make use of verification include international trade agreements (Beshkar and Bond, 2016) and price delegation to sales people in firms (Lo et al., 2016).

agent to choose any spending up to a maximum level. We expand this delegation model by letting the principal verify the agent's type. Verification entails an additive cost for the principal, which may also be partially born by the agent. The agent's type is perfectly revealed if he is verified.[5]

For most of our analysis, we assume that the principal can fully commit to a delegation rule. The problem can be viewed in three steps: first, the principal chooses a mapping from the agent's verification decision and result to a set of allowable spending; second, the agent decides whether to seek verification; third, the agent chooses a spending level from the allowable set. Formally, a delegation rule is a pair of schedules specifying, for each agent type, whether he is verified or not and his spending level.[6] A delegation rule is optimal if it maximizes the principal's expected welfare subject to the incentive compatibility constraint that each agent type prefer his verification assignment and spending level to those of any other type. More precisely, each type must prefer his allocation to that of any other type who is not prescribed verification. Deviations to types who are verified can be trivially deterred as they get revealed by verification. As a consequence, verification can give rise to situations in which local incentive compatibility constraints are slack while non-local constraints bind; our analysis makes use of perturbation methods to address these complications.

Our first main result shows that if the cost of verification is sufficiently small, an optimal rule is a *threshold with an escape clause (TEC)*: the agent can select any action up to a threshold or, if the threshold is sufficiently binding, he can request verification and the efficient action by triggering the escape clause. Importantly, we also show that verifying all agent types is never optimal; hence, no matter how small the verification cost is, an optimal rule prescribes no verification for some types.

The intuition why TEC is optimal is that verifying an upper region of agent types not only allows the principal to improve their spending allocation, but is also an efficient means of imposing discipline on lower agent types who are not verified: these types select from a set of lower spending levels and cannot mimic higher types who are verified. To prove the result, we show that any rule with decreasing verification — prescribing verification for a set of agent types and no verification for a set of higher types — can be dominated. Decreasing verification is expensive for the principal because it requires incentivizing types in the verification region to seek verification rather than mimic a higher type in a no-verification region above them, and this in turn requires inducing significant overspending in the no-verification region. We show that when the verification cost is small enough, a perturbation that verifies all types in the decreasing verification region increases the principal's welfare.

TEC rules are common in applications such as those described above, and rules of this

[5]See Section 6 for a discussion of imperfect verification.

[6]We restrict attention to deterministic rules in our analysis. See Section 6 for a discussion.

form have been shown to be optimal in other contexts with costly state verification (see the literature review below). Existing work, however, has focused on environments in which the objectives of a principal and an agent are in complete disagreement, an assumption that is at odds with the subject of delegation theory. In a delegation problem like the one we study, flexibility is valuable precisely because the agent's bias is not "extreme." We find that this distinction has important implications. Intuitively, an agent with an extreme bias towards higher spending would only pursue verification to increase his spending, whereas one with a moderate bias may pursue verification to increase or *decrease* his spending. In fact, for this reason, we find that while TEC is always optimal under an extreme bias, its optimality relies on verification costs being sufficiently small under a moderate bias.

Specifically, our second main result shows that if the cost of verification is not small enough, decreasing verification can be (strictly) optimal. For example, a rule that verifies only an intermediate set of types can yield the principal higher welfare relative to not verifying any type as well as relative to using a TEC rule. The main reason why verifying only intermediate types can dominate not verifying any type is that the verification region serves to discipline types in the no-verification region below. The main reason why verifying only intermediate types can dominate TEC is that it allows the principal to save on verification costs. We show that these benefits can outweigh the cost of overspending that is needed to incentivize intermediate types to be verified. Thus, when the verification cost is not small (and not large) enough, a rule that involves decreasing verification can be optimal.

The optimality of decreasing verification does not rely on any sort of asymmetry in the payoff or distribution functions. As noted, our setting is one in which threshold rules are always optimal absent verification, and in fact we prove the result by taking the widely-studied case of quadratic preferences and a uniform type distribution. An interior verification region can be beneficial because it allows the principal to divide the delegation set while keeping verification at a minimum. For example, headquarters in organizations may require division managers to either comply with a low budgetary limit meant for relatively small projects or choose from a higher range of investment levels meant for large projects; otherwise documentation is needed to have intermediate levels of investment approved. Such a verification requirement may suffice to discourage overinvestment by managers with small projects: these managers lack proof to justify a small increase in their budget and would not want to increase their investment as much as for a large project.

Whereas a rule with decreasing verification can be optimal, another implication of our construction concerns the need of strong commitment power from the principal when such a rule is used. Consider the aforementioned delegation rule in which the principal verifies only an intermediate set of types. To implement this rule, the principal commits to an allocation that may be inefficient ex post, following the verification decision and result. In particular, the rule may assign an inefficient spending level after the agent's type is verified, both in the case that the agent's seeking verification is "on path" as well as when this verification is part of a deviation. Moreover, the rule may induce an allocation after the agent decides not to seek verification that is inefficient conditional on no verification, i.e. when ignoring the incentives of verified types. What happens if the principal is unable to commit ex ante to these ex-post inefficient allocations?

Our third main result characterizes the optimal rule when the principal's commitment power is limited. In terms of the three-step timing described previously, limited commitment means that the principal now revises the agent's allowable spending set following the agent's verification decision and result. We show that under limited commitment, TEC is optimal whenever verification is optimal. Indeed, we prove that any incentive compatible rule must have weakly increasing verification everywhere. The reason is that inducing decreasing verification requires incentivizing verified types not to deviate and choose a higher spending level in a no-verification region above them, and under limited commitment it also requires incentivizing non-verified types not to seek a verification that guarantees them efficient spending. When unable to fully commit to a rule ex ante, the principal cannot implement the spending levels that would be required to make these deviations unattractive, and thus decreasing verification is not feasible. We obtain that in an environment with limited commitment, the optimality of TEC is restored.

Related literature. Our paper is related to two literatures. First, we contribute to the literature on optimal delegation and self control, starting with Holmström (1977, 1984). Main references include Melumad and Shibano (1991) and Alonso and Matouschek (2008) on delegation under quadratic preferences, Amador, Werning and Angeletos (2006) on consumption-savings problems with hyperbolic preferences, and Amador and Bagwell (2013), which considers a general framework that we take as our baseline.[7] As in this literature, we study a principal-agent environment with no transfers in which the agent is better informed about the efficient action but biased relative to the principal. In contrast to this literature, we allow the principal to verify the agent's information at a cost. By introducing this additional tool, we are able to explore how escape clauses are optimally used, and how optimal delegation depends on the extent of the principal's commitment power.[8]

[7]See also Athey, Atkeson and Kehoe (2005), Ambrus and Egorov (2013, 2015), Halac and Yared (2014, 2017a,b), and Amador, Bagwell and Frankel (2017). Auster and Pavoni (2017) study a delegation problem with limited awareness that gives rise to a non-interval delegation set.

[8]We study the effects of the principal not being able to commit to not changing the agent's allowable spending set following the verification decision and result. A different question that a literature on auditing has investigated concerns a principal's ability to commit to an audit strategy; see, e.g., Reinganum and

Second, we contribute to the literature on costly verification, starting with Townsend (1979). Both that paper and others that followed it, including Gale and Hellwig (1985), Border and Sobel (1987), and Mookherjee and Png (1989), analyze settings with transfers, which we rule out. More recently, Ben-Porath, Dekel and Lipman (2014) and Erlanson and Kleiner (2015) consider costly verification in one-good and collective allocation problems without transfers, and Glazer and Rubinstein (2004, 2006) and Mylovanov and Zapechelnyuk (2014) study related questions using different verification technologies.[9] Our main departure from this literature (in addition to other differences specific to each paper) is that we study a delegation setting in which we allow for different degrees of bias by the agent relative to the principal. This is also a main distinction with respect to Harris and Raviv (1996) — and the dynamic version in Malenko (2016) — which consider costly verification in a delegation model where the agent always benefits from higher actions.[10] The results in Harris and Raviv (1996) are consistent with our benchmark finding that TEC is always optimal under verification if the agent's bias is extreme. As previously discussed, our interest is in understanding optimal delegation and verification when the agent's bias is not extreme: his most preferred action is higher than the principal's but not necessarily the highest possible action. The principal as a result faces a tradeoff between commitment and flexibility, which introduces new conceptual issues into our mechanism design problem. In fact, the agent's bias not being extreme implies that decreasing verification is sometimes optimal in our setting, a feature that does not emerge in this related work.

2 Model

Our baseline model of delegation is the same general principal-agent environment of Amador and Bagwell (2013), where we focus on the case in which the agent's bias is towards higher actions. We extend this delegation model by allowing for costly state verification, following Townsend (1979).

Wilde (1986), Banks (1989), and Chatterjee, Morton and Mukherji (2002). Work on delegation and self-control has also studied lack of commitment to rules; this includes Bernheim, Ray and Yeltekin (2015), Dovis and Kirpalani (2017), and Halac and Yared (2017b).

[9]More broadly, there is a literature on mechanism design and implementation with evidence, including Green and Laffont (1986), Bull and Watson (2007), Deneckere and Severinov (2008), Ben-Porath and Lipman (2012), and Kartik and Tercieux (2012).

[10]In Harris and Raviv (1996), the agent's marginal utility from a higher action depends on his type but is always positive. Their model also differs from ours in other respects: there are only three agent types, the agent receives a non-contingent transfer from the principal, and the principal can choose to verify the agent with an interior probability. Harris and Raviv (1998) consider an extension in which capital is allocated across multiple projects. Malenko (2016) analyzes a dynamic version in which projects of independent and identically distributed quality arrive stochastically over time.

2.1 Environment

There are a principal and an agent. The state is $\gamma \in \Gamma \equiv [\underline{\gamma}, \overline{\gamma}]$ for $\underline{\gamma} > 0$, with continuous density $f(\gamma) > 0$ for all γ. The corresponding distribution function is $F(\gamma)$. The level of spending is denoted by $\pi \in [\underline{\pi}, \overline{\pi}]$.

The principal's welfare is $U_P(\gamma, \pi)$, twice continuously differentiable with $\frac{\partial^2 U_P(\gamma,\pi)}{\partial \pi^2} < 0$. We assume that the principal's optimum, $\pi_P(\gamma) \equiv \arg\max_\pi U_P(\gamma, \pi)$, is interior, and refer to it as the *efficient* level of spending. We impose the following single-crossing condition:

$$\frac{\partial^2 U_P(\gamma, \pi)}{\partial \gamma \partial \pi} > 0. \tag{1}$$

Thus, the efficient level of spending is increasing in the state: $\pi_P'(\gamma) > 0$.

The agent's welfare is $U_A(\gamma, \pi) = \gamma\pi + b(\pi)$, with $b(\pi)$ twice continuously differentiable and $b''(\pi) < 0$. We assume that the agent's optimum, $\pi_A(\gamma) \equiv \arg\max_\pi U_A(\gamma, \pi)$, is interior, and refer to it as the *flexible* level of spending. Note that the agent's welfare satisfies the single-crossing condition $\frac{\partial^2 U_A(\gamma,\pi)}{\partial\gamma\partial\pi} > 0$.[11] We consider an agent who is biased towards higher spending, specifically whose welfare increases by more than the principal's when spending rises. That is, we add the following assumption to the setting of Amador and Bagwell (2013):

$$\frac{\partial U_A(\gamma, \pi)}{\partial \pi} > \frac{\partial U_P(\gamma, \pi)}{\partial \pi}. \tag{2}$$

Thus, conditional on the state, the flexible level of spending always exceeds the efficient level: $\pi_A(\gamma) > \pi_P(\gamma)$ for all $\gamma \in \Gamma$.

The state γ is private information to the agent, i.e. the agent's *type*. The principal can perfectly verify γ by paying an additive cost $\phi > 0$. The agent's cost of verification is $\alpha\phi$ for $\alpha \in [0, 1]$. This formulation allows us to cover situations in which the agent pays no verification cost ($\alpha = 0$) as well as situations in which he pays a cost no larger than the principal's ($\alpha \in (0, 1]$). One could also allow for the agent to pay a higher verification cost than the principal's (namely let $\alpha > 1$); our main results would continue to hold in that case if the agent's bias is sufficiently large.

By featuring both a bias and private information by the agent, our environment gives rise to a tradeoff between commitment and flexibility. If the agent were not biased relative to the principal, the principal could implement the efficient level of spending by providing full flexibility to the agent (who would in this case choose $\pi_A(\gamma) = \pi_P(\gamma)$). Similarly, if the state γ were not the agent's private information, the principal could implement the efficient level of spending by committing the agent to a fully contingent spending plan. In the

[11] For both the principal and the agent's preferences, we will refer to "single-crossing" as the (stronger) supermodularity condition that we have assumed these preferences satisfy.

presence of both a bias and private information, however, the principal cannot implement efficient spending $\pi_P(\gamma)$ for all γ without verification, and she faces a non-trivial tradeoff between commitment and flexibility.

Special cases. The model of delegation described above encompasses specific cases commonly studied in the literature. One example is the case of quadratic preferences with a constant bias (which we will refer to as simply quadratic preferences), examined by Melumad and Shibano (1991) and Alonso and Matouschek (2008) and used extensively in applied work. Under these preferences, the principal's welfare is $-\frac{(\gamma-\pi)^2}{2}$ and the agent's welfare is $-\frac{(\gamma+\beta-\pi)^2}{2}$, for some $\beta > 0$ representing the agent's bias. This formulation is equivalent to letting $U_P(\gamma, \pi) = \gamma\pi + b(\pi) - \beta\pi$ and $U_A(\gamma, \pi) = \gamma\pi + b(\pi)$ for $b(\pi) = \beta\pi - \frac{\pi^2}{2}$, and is therefore a special case of our model. We will use the quadratic preferences case to illustrate some of our results.

Another example is the model of consumption under hyperbolic preferences, analyzed by Amador, Werning and Angeletos (2006) and Halac and Yared (2014, 2017a,b).[12] The principal's welfare in this case is $\gamma u(c) + w(y-c)$ and the agent's welfare is $\gamma u(c) + \beta w(y-c)$, where u and w are utility functions, c and y represent consumption and exogenous income respectively, and $\beta \in (0, 1)$ captures the degree of present bias by the agent. This formulation is equivalent to letting $U_P(\gamma, \pi) = \gamma\pi + \frac{1}{\beta} b(\pi)$ and $U_A(\gamma, \pi) = \gamma\pi + b(\pi)$ with $\pi = u(c)$ and $b(\pi) = \beta w(y - u^{-1}(\pi))$, and is thus also encompassed by our model.

2.2 Timing

The order of events is as follows:

1. The principal sets a rule, which maps a verification decision and result into an allowable spending set Π.

2. The agent chooses whether or not to seek verification, $a \in \{0, 1\}$, and the principal perfectly verifies his type γ if $a = 1$.

3. The agent chooses a spending level π from the allowable set Π.

The above timing assumes that the agent learns his type γ before the principal sets a rule in Step 1. Our analysis is unchanged if instead the agent learns his type after the rule has been set, i.e. at the beginning of Step 2.

[12] Halac and Yared use this model to study fiscal rules, where a government's deficit bias may emerge from the aggregation of heterogeneous, time-consistent citizens' preferences (Jackson and Yariv, 2015, 2014) or from turnover in a political economy setting (Aguiar and Amador, 2011; Alesina and Passalacqua, 2016).

2.3 Delegation Rules

Given the game form described above, we can analyze the principal's problem as that of choosing a delegation rule M which consists of a pair of schedules $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$, specifying a verification decision and spending level for each type γ. The principal chooses a rule M to maximize her expected welfare:

$$\max_{\{a(\gamma),\pi(\gamma)\}_{\gamma\in\Gamma}} \int_{\underline{\gamma}}^{\overline{\gamma}} \big(U_P(\gamma, \pi(\gamma)) - a(\gamma)\phi\big) f(\gamma)\, d\gamma \tag{3}$$

subject to

$$U_A(\gamma, \pi(\gamma)) - a(\gamma)\alpha\phi \geq U_A(\gamma, \pi(\widehat{\gamma})) \text{ for all } \gamma, \widehat{\gamma} \text{ for which } a(\widehat{\gamma}) = 0. \tag{4}$$

The objective (3) is the principal's expected welfare under a given rule, taking into account the additive verification costs. The constraint (4) is an incentive compatibility (or truthtelling) constraint: it guarantees that an agent of type γ prefers his assigned verification decision and spending level, $a(\gamma)$ and $\pi(\gamma)$, to a different allocation $a(\widehat{\gamma})$ and $\pi(\widehat{\gamma})$ for some type $\widehat{\gamma}$ who is not verified (that is, with $a(\widehat{\gamma}) = 0$). Note that it is sufficient to consider deviations to non-verified types: since a deviation in which an agent of type γ mimics a verified type $\widehat{\gamma}$ would be detected by the principal (as verification reveals the true type) and the principal can arbitrarily punish the agent when she learns that he has deviated (off path), we do not need to consider such a deviation.[13]

We also note that the formulation above does not rule out mixed strategies by the agent. If the agent were willing to mix over verification and no verification or over different spending levels, he would be indifferent over these allocations, and thus the principal can select one of these that maximizes her expected welfare.[14] In fact, building on this observation, we can show that our results are not limited to the game form in Section 2.2 but continue to hold when allowing for any indirect mechanism specifying a message space for the agent and a deterministic allocation function to which the principal commits. Such a mechanism induces a game in which the agent sends a message, is either verified or not as a function of the message, and is assigned a spending level as a function of the message and verification result. We show in the Online Appendix that a version of the Revelation Principle in terms of payoffs holds in our setting, implying that to study the optimal deterministic mechanism for the principal, it is without loss to restrict attention to deterministic direct mechanisms

[13] The principal can punish a deviation of a type γ in which he mimics a type $\widehat{\gamma} \neq \gamma$ with $a(\widehat{\gamma}) = 1$ by assigning following verification some spending level $\pi(\widehat{\gamma}, \gamma)$ such that $U_A(\gamma, \pi(\widehat{\gamma}, \gamma)) \leq U_A(\gamma, \pi(\gamma))$. It is clear that such a spending level exists; in fact, setting $\pi(\widehat{\gamma}, \gamma) = \pi(\gamma)$ would be a sufficient punishment.

[14] While this selection relaxes the principal's problem, it is not used under the optimal rule described in our main result in Proposition 3, which induces a unique best response by the agent. Hence, the result does not rely on selection of equilibria of the game in Section 2.2.

(i.e. where the message space coincides with the agent's type space) that induce truthful reporting by the agent, as considered in program (3)-(4) above.

Because there is a continuum of types, it is possible that the problem in (3)-(4) admit multiple solutions that are identical everywhere except for a measure-zero set of types. As a means of selecting the optimum in such a situation, we say that a rule M is *optimal* if it solves (3)-(4) *and* there is no other solution $\widetilde{M}$, with associated verification and spending schedules $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma\in\Gamma}$, such that

$$U_P\left(\gamma, \widetilde{\pi}\left(\gamma\right)\right) - \widetilde{a}\left(\gamma\right)\phi \geq U_P\left(\gamma, \pi\left(\gamma\right)\right) - a\left(\gamma\right)\phi \tag{5}$$

for all γ and strictly for some $\gamma \in \Gamma$.[15]

3 No Verification Benchmark

Before analyzing the optimal delegation rule with verification, we review the results of the literature by considering the optimal rule in the absence of verification. Suppose the principal faces the constraint that $a\left(\gamma\right) = 0$ for all γ.[16] The problem in (3)-(4) subject to this additional constraint is studied by Amador and Bagwell (2013). To solve this problem, they make the following Assumption 1 on the distribution of γ; we extend this assumption to any truncation from above, with support $[\underline{\gamma}, \overline{\gamma}']$ for $\overline{\gamma}' \leq \overline{\gamma}$, density $f\left(\gamma\right)/F\left(\overline{\gamma}'\right)$, and distribution function $F\left(\gamma\right)/F\left(\overline{\gamma}'\right)$:

Assumption 1. *Take the distribution of γ truncated from above by $\overline{\gamma}' \leq \overline{\gamma}$. For each such truncated distribution, there exists γ^* such that for $\kappa \equiv \inf_{\{\gamma,\pi\}}\left(\frac{\partial^2 U_P(\gamma,\pi)/\partial^2\pi}{b''(\pi)}\right)$,*

(i) $\kappa F\left(\gamma\right) - \frac{\partial U_P(\gamma,\pi_A(\gamma))}{\partial\pi} f\left(\gamma\right)$ is nondecreasing for all $\gamma \in \left[\underline{\gamma}, \gamma^*\right]$, and

(ii) $\left(\gamma - \gamma^*\right)\kappa \geq \int_{\gamma}^{\overline{\gamma}'} \frac{\partial U_P(\widetilde{\gamma},\pi_A(\gamma^*))}{\partial\pi} \frac{f(\widetilde{\gamma})}{1-F(\gamma)} d\widetilde{\gamma}$ for all $\gamma \in [\gamma^*, \overline{\gamma}']$, with equality at γ^*.

One can verify that for the special cases typically studied in the literature, such as those with quadratic or hyperbolic preferences, Assumption 1 is satisfied under commonly used distribution functions, including exponential, log-normal, and any nondecreasing density.[17] Given Assumption 1, the results in Amador and Bagwell (2013) yield:

[15] Although multiple solutions can in principle continue to exist under this condition, this criterion turns out to be sufficient for our characterization.

[16] In this case, the constraint (4) becomes $U_A\left(\gamma, \pi\left(\gamma\right)\right) \geq U_A\left(\gamma, \pi\left(\widehat{\gamma}\right)\right)$ for all $\gamma, \widehat{\gamma}$.

[17] We note also that Assumption 1 on the original distribution implies that the assumption is satisfied for all truncations from above if the conditions in Proposition 2 of Amador and Bagwell (2013) hold.

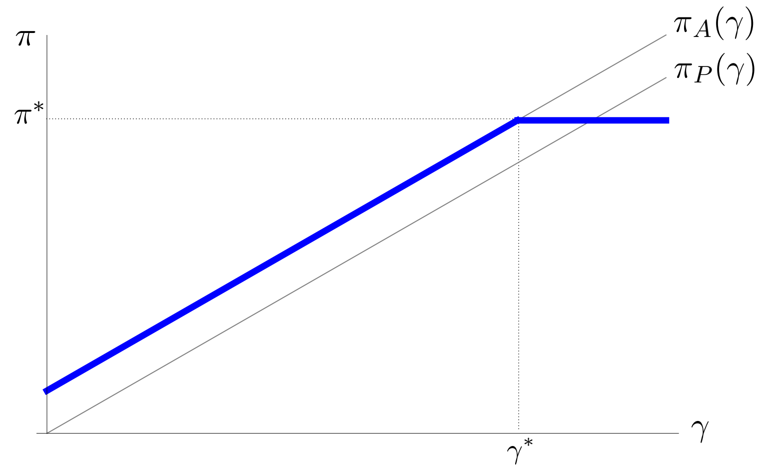

Figure 1: An optimal rule under no verification. The figure is drawn for the quadratic preferences case (see Section 2.1), where we let $\underline{\gamma} = 0.5, \overline{\gamma} = 1.5$, $\beta = 0.12$, and $F(\gamma)$ uniform.

Proposition 1 (no verification)**.** *Take the distribution of γ truncated from above by $\overline{\gamma}' \leq \overline{\gamma}$. If the principal is constrained to $a(\gamma) = 0$ for all $\gamma \in [\underline{\gamma}, \overline{\gamma}']$, an optimal rule is a threshold $\gamma^* < \overline{\gamma}'$ such that*

$$\pi(\gamma) = \min\{\pi_A(\gamma), \pi_A(\gamma^*)\} \quad \textit{for } \gamma \in [\underline{\gamma}, \overline{\gamma}'].$$

Under no verification, an optimal rule is a threshold γ^* such that all types $\gamma \leq \gamma^*$ spend at their flexible level and all types $\gamma > \gamma^*$ are bunched at the flexible spending level of γ^*. The principal can implement this rule by setting a spending limit $\pi^* = \pi_A(\gamma^*)$ and allowing the agent to choose any spending level up to this limit.

Figure 1 illustrates an optimal rule with no verification for the case of quadratic preferences. The level of spending is on the vertical axis and the agent's type on the horizontal axis. In this simple example, both efficient and flexible spending are increasing linear functions of the state γ, and flexible spending exceeds efficient spending by a constant amount representing the agent's bias. The rule characterized in Proposition 1 specifies a spending level that coincides with the agent's flexible level for $\gamma \leq \gamma^*$ and equals $\pi_A(\gamma^*)$ for $\gamma > \gamma^*$.

A key insight behind the result in Proposition 1 is that "holes" are suboptimal. More precisely, the principal can always improve upon a rule as that depicted in Figure 2, which does not allow the agent to choose a spending level $\pi \in [\pi_L, \pi_H]$, for some $\underline{\pi} < \pi_L < \pi_H < \overline{\pi}'$, but allows the agent to choose spending immediately below π_L and immediately above π_H. The hole $[\pi_L, \pi_H]$ implies that an agent of type γ for whom $\pi_A(\gamma) \in (\pi_L, \pi_H)$ is not allowed to spend at his flexible level. Such an agent spends at the lower limit of the hole $\pi_L < \pi_A(\gamma)$ if his type is relatively low, but he spends at the upper limit of the hole $\pi_H > \pi_A(\gamma)$ if his type is higher. The role of Assumption 1 is to guarantee that if the principal removes the

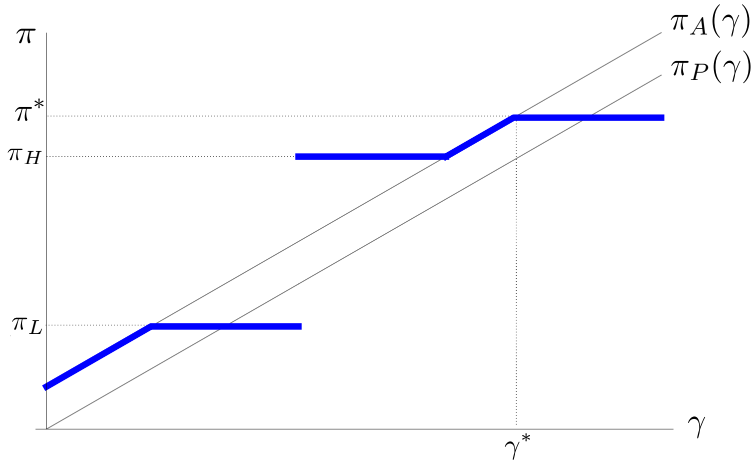

Figure 2: A rule without verification with a hole $[\pi_L, \pi_H]$. Parameters are the same as in Figure 1.

hole and allows full flexibility over $[\pi_L, \pi_H]$, the benefit of reducing overspending for the types that bunch at π_H would outweigh the (potential) cost of increasing spending for the types that bunch at π_L. Therefore, the principal is better off by closing the hole.

4 Optimal Rule

We now turn to the study of optimal delegation when costly verification is possible. The following class of rules will play an important role in our analysis:

Definition 1. *A rule is a threshold with an escape clause (TEC) if it consists of* $\{\gamma^*, \gamma^{**}\}$ *with* $\gamma^* < \gamma^{**}$ *and* $\underline{\gamma} < \gamma^{**} < \overline{\gamma}$ *such that*

(i) (threshold) If $\gamma \leq \gamma^{**}$, $a(\gamma) = 0$ *and* $\pi(\gamma) = \min\{\pi_A(\gamma), \pi_A(\gamma^*)\}$, *and*

(ii) (escape clause) if $\gamma > \gamma^{**}$, $a(\gamma) = 1$ *and* $\pi(\gamma) = \pi_P(\gamma)$.

Figure 3 illustrates a TEC rule using the quadratic preferences example. Under TEC, types $\gamma \leq \gamma^*$ are not verified and spend at their flexible level, types $\gamma \in (\gamma^*, \gamma^{**}]$ are not verified and are bunched at the flexible spending level of γ^*, and types $\gamma > \gamma^{**}$ are verified and are assigned their efficient spending level. The principal can implement this rule by allowing the agent to choose any spending level up to a limit $\pi^* = \pi_A(\gamma^*)$ or request verification by triggering an escape clause. When the agent is verified, he is assigned his efficient spending level provided that it is above a specified level $\pi^{**} = \pi_P(\gamma^{**})$ (and is otherwise punished).

An important feature of TEC is that the verification function $a(\gamma)$ is weakly increasing, that is, there is no decreasing verification:

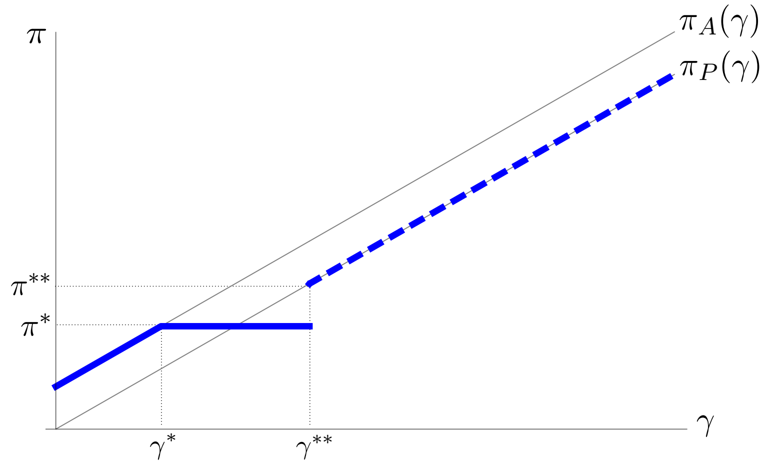

Figure 3: A TEC rule. Parameters are the same as in Figure 1, with $\phi = 0.008$ and $\alpha = 0$. The solid line depicts the allocation of non-verified types; the dashed line corresponds to verified types.

Definition 2. *A rule features decreasing verification at* γ' *if either (i)* $a(\gamma') < \limsup_{\gamma \uparrow \gamma'} a(\gamma)$ *or (ii)* $a(\gamma') > \liminf_{\gamma \downarrow \gamma'} a(\gamma)$*. A rule features weakly increasing verification at* γ' *if neither* (i) *nor* (ii) *holds.*

Note that we will refer to decreasing/increasing verification in the strict sense, and we will clarify whenever we use decreasing/increasing verification in the weak sense. Figure 4 depicts an example of a rule with decreasing verification. This rule specifies verification only for types between two interior cutoffs, γ_L and $\gamma_H > \gamma_L$. Types above and below this region are not verified, and hence the rule features decreasing verification at γ_H. We will return to this example in Section 4.3.

Another feature of TEC is that it specifies verification for some agent types but not for all. We begin by showing in Section 4.1 that inducing no verification for some types is in fact a property of any optimal rule. Furthermore, building on this result, we show that TEC is optimal whenever optimal verification is everywhere weakly increasing. We study a simple extreme bias case in Section 4.2 and provide a characterization for our general setting in Section 4.3.

4.1 Preliminaries

The next lemma shows that verifying all agent types is never optimal for the principal:

Lemma 1. *A rule with* $a(\gamma) = 1$ *for all* $\gamma \in \Gamma$ *is not optimal.*

The logic is simple. Suppose that a rule that verifies all types is optimal. Such a rule must trivially assign efficient spending to all types. Now consider a perturbation in which

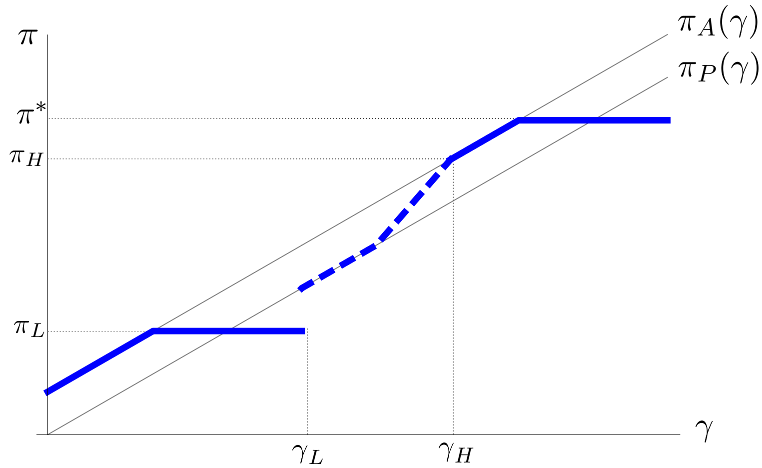

Figure 4: A rule with decreasing verification. Parameters are the same as in Figure 3. The solid line depicts the allocation of non-verified types; the dashed line corresponds to verified types.

the principal allows the agent to choose $\pi_P(\underline{\gamma})$ without verification. Under the perturbed rule, a set of types $[\underline{\gamma}, \gamma']$, for $\gamma' \geq \underline{\gamma}$, will prefer $\pi_P(\underline{\gamma})$ over being verified and assigned efficient spending. Moreover, since the agent is biased towards higher spending and pays a verification cost no larger than the principal's, it must be that the principal is strictly better off by not verifying these types. Hence, we find that incentivizing low types to not overspend is cheaper than verifying them, and thus verifying all types cannot be optimal.

Given Lemma 1, we establish:

Lemma 2. *If an optimal rule features verification that is weakly increasing everywhere, then TEC is optimal.*

Since verifying all agent types is suboptimal, an optimal rule with verification that is weakly increasing everywhere must feature a no-verification region followed by a verification region, i.e. there must be a type γ^{**} such that $a(\gamma) = 0$ for $\gamma < \gamma^{**}$ and $a(\gamma) = 1$ for $\gamma > \gamma^{**}$. Conditional on the agent's type being in the no-verification region, an optimal rule is a threshold $\gamma^* < \gamma^{**}$ (by Proposition 1), and conditional on the agent's type being in the verification region, an optimal rule assigns efficient spending to all types. To prove Lemma 2, we show that the rule that results from optimizing over each region separately is incentive compatible, and thus optimal, over the whole set of types. Specifically, we show that an optimal rule conditional on no-verification sets a maximum allowable spending level $\pi_A(\gamma^*) \leq \pi_P(\gamma^{**})$, and by optimality of γ^{**} the principal prefers to pay the cost of verifying type $\gamma > \gamma^{**}$ to assign him $\pi_P(\gamma)$ rather than bunch him at $\pi_A(\gamma^*)$. Since the agent is biased towards higher spending and pays a verification cost no larger than the principal's, it follows that types $\gamma > \gamma^{**}$ also prefer to be verified rather than deviating to $\pi_A(\gamma^*)$.

Therefore, the resulting rule is incentive compatible and thus optimal, and it is TEC.

4.2 Extreme Bias

Before turning to our main results, we consider a setting in which the agent's bias is extreme. Suppose $b(\pi) = 0$ for all $\pi \in [\underline{\pi}, \overline{\pi}]$, so that the agent's welfare is simply $U_A(\gamma, \pi) = \gamma\pi$. The agent in this case always prefers higher levels of spending: his flexible spending level is $\pi_A(\gamma) = \overline{\pi}$ for all $\gamma \in \Gamma$.[18] As mentioned in the Introduction, such an extreme bias corresponds to what is assumed in other models of costly verification, including the seminal work of Townsend (1979), the delegation model of Harris and Raviv (1996, 1998), and more recent contributions such as Ben-Porath, Dekel and Lipman (2014).

An extreme bias implies that if the agent is not verified, he will choose the highest allowable level of spending, regardless of his type. Moreover, the agent will seek verification only if that allows him to spend *more* than under no verification. The analysis therefore is significantly simplified. The only incentive compatible rule for an agent with an extreme bias involves bunching all non-verified types at one spending level; that is, flexibility has no value in this setting. Furthermore, any type that is verified must be assigned a higher spending level than that at which non-verified types are bunched. As a result:[19]

Proposition 2 (extreme bias). *Suppose $b(\pi) = 0$ for all $\pi \in [\underline{\pi}, \overline{\pi}]$. Then if verification is optimal, TEC is optimal.*

When the agent's bias is extreme and verifying some types is optimal, an optimal rule is TEC, with non-verified types $\gamma \leq \gamma^{**}$ bunched and awarded no flexibility and verified types $\gamma > \gamma^{**}$ spending at their efficient level. The optimality of TEC follows from the optimality of weakly increasing verification. Suppose by contradiction that an optimal rule featured decreasing verification. Take γ' to be a marginal non-verified type splitting a verification region and a higher no-verification region, so $a(\gamma') = 0$ and $a(\gamma' - \varepsilon) = 1$ for some $\varepsilon > 0$ arbitrarily small. Let $\pi_A(\gamma^*)$ be the level of spending at which non-verified types are bunched. The optimality of verifying $\gamma' - \varepsilon$ implies

$$U_P(\gamma' - \varepsilon, \pi(\gamma' - \varepsilon)) - U_P(\gamma' - \varepsilon, \pi_A(\gamma^*)) \geq \phi, \tag{6}$$

where, as noted, incentive compatibility requires $\pi(\gamma' - \varepsilon) \geq \pi_A(\gamma^*)$, and since $\phi > 0$, (6)

[18] As assumed in Section 2.1, we are primarily interested in the case in which $\pi_A(\gamma)$ is interior rather than a corner; however, we find it is instructive to study this corner case first.

[19] Proposition 2, as well as Proposition 3 and Proposition 5, describes an optimal rule when verification is optimal. Clearly, verification is optimal if and only if the verification cost ϕ is not too high.

yields $\pi(\gamma' - \varepsilon) > \pi_A(\gamma^*)$. The optimality of not verifying γ' then implies

$$U_P(\gamma', \pi(\gamma' - \varepsilon)) - U_P(\gamma', \pi_A(\gamma^*)) \leq \phi. \tag{7}$$

However, (6) and (7) together with $\pi(\gamma' - \varepsilon) > \pi_A(\gamma^*)$ violate the single-crossing condition (1), yielding a contradiction. Intuitively, the principal can improve upon a rule with decreasing verification by verifying a higher agent type instead of a lower type, as the marginal benefit of letting the higher type spend more is higher. Note that such a perturbation is always incentive compatible for the agent because all non-verified types are bunched at the same spending level, which (by incentive compatibility) is lower than the spending level assigned to any verified type. This feature is of course due to the agent's bias being extreme.

4.3 Optimal Rule with Verification

We next study optimal delegation with verification in our general setting in which the agent's bias is not extreme. To this end, it is useful to consider a relaxed version of the problem in (3)-(4), in which we assume that the agent pays no verification cost ($\alpha = 0$):

$$\max_{\{a(\gamma),\pi(\gamma)\}_{\gamma\in\Gamma}} \int_{\underline{\gamma}}^{\overline{\gamma}} \big(U_P(\gamma, \pi(\gamma)) - a(\gamma)\,\phi\big) f(\gamma)\, d\gamma \tag{8}$$

subject to

$$U_A(\gamma, \pi(\gamma)) \geq U_A(\gamma, \pi(\widehat{\gamma})) \text{ for all } \gamma, \widehat{\gamma} \text{ for which } a(\widehat{\gamma}) = 0. \tag{9}$$

Since the original incentive compatibility constraint (4) is tighter than the relaxed constraint (9), if a solution to (8)-(9) satisfies (4), then it is also a solution to the problem in (3)-(4). Furthermore, we can show that if a solution to (8)-(9) is TEC, then it will indeed satisfy (4), implying:

Lemma 3. *If a TEC rule is a solution to* (8)-(9)*, it is also a solution to* (3)-(4)*.*

The logic is similar to that behind Lemma 2. To verify that a TEC rule $\{\gamma^*, \gamma^{**}\}$ that solves (8)-(9) satisfies the original constraint (4), we must check that an agent of type $\gamma > \gamma^{**}$ would prefer to pay the verification cost $\alpha\phi$ and spend at his efficient level $\pi_P(\gamma)$ rather than pay no verification cost and choose the threshold flexible spending level $\pi_A(\gamma^*)$. Now TEC being a solution to (8)-(9) implies that the principal prefers verifying such a type γ to assign him $\pi_P(\gamma)$ rather than bunching this type at $\pi_A(\gamma^*)$, where $\pi_A(\gamma^*) \leq \pi_P(\gamma)$ for all $\gamma > \gamma^{**}$. Since the agent is biased towards higher spending and pays a verification cost no larger than the principal's, the optimality of verifying γ for the principal therefore yields that verifying γ is incentive compatible for the agent. This establishes Lemma 3, and it implies that to study whether TEC is optimal, it is without loss to focus on the relaxed problem in (8)-(9).[20] We analyze this problem for the remainder of this section.

The following two lemmas establish useful properties of any solution:

Lemma 4. *If a solution to* (8)-(9) *prescribes verification for type* γ*, it has* $\pi_P(\gamma) \leq \pi(\gamma) \leq \pi_A(\gamma)$*. If* (9) *does not bind for* γ*, then* $\pi(\gamma) = \pi_P(\gamma)$*.*

Lemma 5. *In any solution to* (8)-(9)*,* $\pi(\gamma)$ *is weakly increasing.*

Lemma 4 states that if a type γ is verified, his assigned spending level is (weakly) between his efficient level and his flexible level. The argument is straightforward. If assigned spending for type γ is either below efficient or above flexible, then either increasing or decreasing this spending, respectively, makes the principal better off and is incentive compatible for the agent. Since the principal maximizes her expected welfare subject to incentive compatibility, if a verified type's incentive compatibility constraint is slack, the principal assigns this type efficient spending.

Lemma 5 shows that the principal assigns a spending level that is weakly increasing in the agent's type γ. When comparing two agent types that are not verified, the result naturally follows from incentive compatibility: a type γ cannot be assigned higher spending than a higher type $\gamma' > \gamma$, as at least one of them would have an incentive to deviate given that preferences satisfy single-crossing. When comparing two agent types such that (at least) one of them is verified, the result follows from optimality: if a type γ is assigned higher spending than a higher type $\gamma' > \gamma$, the principal can improve welfare by swapping these types' spending levels and verification assignments, and if incentive compatibility was initially satisfied, it will continue to be satisfied after the swap, given single-crossing.

By definition of TEC and Lemma 2, whether a TEC rule is optimal depends on whether the principal can instead benefit from inducing decreasing verification, namely a situation in which a set of types is verified and a set of higher types is not verified. Using Lemma 4 and Lemma 5, we next show that any rule featuring decreasing verification must induce significant overspending, limiting the welfare that such a rule can provide to the principal:

Lemma 6. *Suppose a solution to* (8)-(9) *features decreasing verification at* $\gamma' < \overline{\gamma}$*. Then the solution satisfies*

$$\frac{\int_{\gamma'}^{\overline{\gamma}} \left(U_P\left(\gamma, \pi_P\left(\gamma\right)\right) - U_P\left(\gamma, \pi\left(\gamma\right)\right)\right) f\left(\gamma\right) d\gamma}{1 - F\left(\gamma'\right)} \geq \eta\left(\gamma'\right) \tag{10}$$

[20] We maintain our optimality condition (5), so that to prove the optimality of TEC, it is sufficient to show that TEC solves (8)-(9) and no other solution provides the principal weakly larger welfare from each type γ and strictly larger from some type γ.

for

$$\eta(\gamma') = \frac{\int_{\gamma'}^{\min\left\{\pi_P^{-1}(\pi_A(\gamma')),\overline{\gamma}\right\}} \big(U_P(\gamma,\pi_P(\gamma)) - U_P(\gamma,\pi_A(\gamma'))\big) f(\gamma)\, d\gamma}{1-F(\gamma')} > 0. \tag{11}$$

If a rule with decreasing verification at $\gamma' < \overline{\gamma}$ is optimal, the principal's expected welfare in the region above γ' is strictly bounded away from that achieved under efficient spending. The reason is that such a rule induces strict overspending by a positive mass of types $\gamma \geq \gamma'$. To see the intuition, let $a(\gamma') = 1$ and thus $a(\gamma' + \varepsilon) = 0$ for $\varepsilon > 0$ arbitrarily small. The principal must incentivize types in the verification region below γ' to seek verification rather than deviate and mimic a type in the no-verification region above γ'. By Lemma 5, all types above γ' spend more than those below, and by Lemma 4, verified types γ spend no more than their flexible amount $\pi_A(\gamma)$. Thus, for types in the verification region below γ' not to deviate to the no-verification region above γ', we must have $\pi(\gamma' + \varepsilon) \geq \pi_A(\gamma')$; in fact, by optimality, this inequality must be strict.[21] Given that by Lemma 5 all types $\gamma \geq \gamma' + \varepsilon$ spend weakly above $\pi(\gamma'+\varepsilon)$, it follows that all types $\gamma \in (\gamma', \min\left\{\pi_P^{-1}(\pi_A(\gamma')), \overline{\gamma}\right\})$ spend strictly above $\pi_A(\gamma') > \pi_P(\gamma)$. This yields the bound in (11).

The properties shown in Lemma 4-Lemma 6 are satisfied in the examples of Figure 3 and Figure 4. Importantly, Lemma 6 shows that the principal's expected welfare under a rule featuring decreasing verification is bounded away from efficient welfare by a bound that is independent of the verification cost ϕ. This allows us to establish our first main result. In what follows, let $\eta(\overline{\gamma}) = \lim_{\gamma \uparrow \overline{\gamma}} \eta(\gamma)$.

Proposition 3 (small verification cost)**.** *Let* $\overline{\phi} \equiv \min_{\gamma \in \Gamma} \eta(\gamma) > 0$. *If* $\phi < \overline{\phi}$ *and verification is optimal, TEC is optimal.*

Given Lemma 2, the proof of Proposition 3 rests on showing that, for any verification cost $\phi < \overline{\phi}$, an optimal rule induces weakly increasing verification everywhere, i.e. decreasing verification is suboptimal. To see why this must be true, suppose by contradiction that an optimal rule induces decreasing verification at some point, and let γ^{**} be the lowest verified type under this rule. We show that the principal can improve upon such a rule by performing a global perturbation: in the perturbed rule, the principal verifies all types $\gamma \geq \gamma^{**}$ and assigns them efficient spending, while solving for an optimal rule without verification for types $\gamma < \gamma^{**}$. By Proposition 1, an optimal rule for the no-verification region is a threshold $\gamma^* < \gamma^{**}$, and since $\pi_A(\gamma^*) \leq \pi_P(\gamma^{**})$ (by optimality of γ^*) and $\alpha = 0$, it is easy to verify that the perturbed rule is incentive compatible.

[21] If $\pi(\gamma' + \varepsilon) = \pi_A(\gamma')$, incentive compatibility requires $\pi(\gamma') = \pi_A(\gamma')$, but then the principal can improve upon the rule by setting $a(\gamma') = 0$ while keeping everything else unchanged.

To show that the perturbation strictly raises the principal's welfare, note first that expected welfare conditional on $\gamma < \gamma^{**}$ weakly increases because it is now maximized subject to fewer constraints: under the perturbed rule, types $\gamma < \gamma^{**}$ cannot mimic a type $\widehat{\gamma} \geq \gamma^{**}$. Thus, all we need to show is that expected welfare conditional on $\gamma \geq \gamma^{**}$ increases strictly, namely that the (allocative) benefit of verifying these types is strictly greater than the additional verification cost the principal incurs. Because verified types are assigned efficient spending, the benefit of verifying $\gamma \geq \gamma^{**}$ is weakly positive for all such types. Moreover, note that by the contradiction assumption, there exists a type above γ^{**} at which the original rule features decreasing verification. Thus, if $\gamma' < \overline{\gamma}$ is the lowest such type, Lemma 6 implies that the benefit of verifying types $\gamma \geq \gamma^{**}$ is bounded from below by $(1 - F(\gamma'))\eta(\gamma')$, where $\eta(\cdot)$ is defined in (11). The claim then follows in this case from the fact that, given $\phi < \overline{\phi}$, the additional cost of verifying types $\gamma \geq \gamma^{**}$ is strictly smaller than $(1 - F(\gamma'))\overline{\phi} = (1 - F(\gamma')) \min_{\gamma \in \Gamma} \eta(\gamma)$, and hence strictly smaller than the benefit of verifying these types. If the lowest type above γ^{**} at which the original rule features decreasing verification is $\gamma' = \overline{\gamma}$, an analogous argument applies, as in this case the original rule induces strict overspending by $\overline{\gamma}$ and the benefit of verifying this type is no smaller than $\overline{\phi}$.

Proposition 3 shows that if the cost of verification ϕ is small enough, an optimal rule takes the form of TEC, which resembles the delegation rules often used in applications.[22] But what happens if ϕ is higher? Our next result shows that there exist environments and verification costs for which the principal induces verification but not in the form of TEC:

Proposition 4 (intermediate verification cost)**.** *There exist $\{U_P, b, f, \phi, \alpha\}$ for which any optimal rule features decreasing verification.*

To prove this result, we construct examples in which verifying only an intermediate range of types $[\gamma_L, \gamma_H]$ dominates both not verifying any type as well as using TEC.[23] The main reason why verifying only intermediate types can dominate not verifying any type is that an intermediate verification region imposes discipline on the no-verification region below. That is, even when the verification cost is large enough that the principal would not benefit from verifying types in $[\gamma_L, \gamma_H]$ only to improve their allocation relative to flexible spending, she may benefit from verifying these types to discipline lower types: with the intermediate verification region, types $\gamma < \gamma_L$ can no longer mimic types in $[\gamma_L, \gamma_H]$. The main reason why verifying only intermediate types can dominate using a TEC rule is that it allows

[22] Comparative statics are as one would expect. In particular, the lower is $\phi < \overline{\phi}$, other things equal, the larger is the verification region (i.e. the smaller is γ^{**}) in the optimal TEC rule.

[23] By Lemma 2, any other rule with verification that is weakly increasing everywhere is thus also dominated. Hence, the claim in Proposition 4 follows.

the principal to save on verification costs. Specifically, with intermediate verification, the principal may be able to impose discipline on types $\gamma < \gamma_L$ without prescribing verification for types $\gamma > \gamma_H$ as she would under a TEC rule; this will be the case if γ_L has no incentive to deviate to mimic a type as high as γ_H. In such a situation, intermediate verification allows the principal to save on the cost of verifying types above γ_H.

These arguments yield that a rule with decreasing verification as that depicted in Figure 4 can dominate any no-verification rule (as that in Figure 1) and any TEC rule (as that in Figure 3), provided that the cost of verification ϕ is not small (nor large) enough. We emphasize that Proposition 4 does not rely on non-uniformity of the principal's objective across types or any other sort of asymmetry; we prove the result by constructing examples as those depicted in our figures, with quadratic preferences and a uniform distribution of types. We also note that while these examples imply the optimality of decreasing verification under some parameters with $\phi > \overline{\phi}$, the optimal rule in this case may not take the simple intermediate-verification structure that we consider to prove the result. In fact, we can show that even when restricting attention to quadratic preferences and a uniform distribution, there exist parameters for which TEC, no verification, and intermediate verification are all dominated by a rule featuring multiple interior verification regions.[24] Intuitively, intercalating verification regions to further divide the delegation set can allow the principal to improve discipline while keeping verification costs at a minimum.[25]

Whereas a rule with decreasing verification can be optimal, our construction also highlights the need of strong commitment power from the principal to implement such a rule. Take for example the rule depicted in Figure 4. The principal assigns spending strictly above the efficient level to some agent types $\gamma \in [\gamma_L, \gamma_H]$ who are verified. By doing this, the principal incentivizes those types to be verified: if they were instead assigned efficient spending following verification, they would not seek verification in the first place. The principal must be committed to allowing this inefficient spending despite her learning the true type of the agent following verification. Strong commitment power from the principal is also required to incentivize types $\gamma < \gamma_L$ sufficiently close to γ_L to not seek verification. In the rule of Figure 4, these types are punished if they seek verification, even though ex post, once verification took place, both the principal and the agent would strictly prefer efficient spending to punishment. Without the threat of punishment, the principal may not be able to prevent an agent of type $\gamma < \gamma_L$ sufficiently close to γ_L from seeking verification, as an efficient allocation following verification would allow this agent to increase his spending.

[24]In particular, the rule constructed in Lemma 9 in the proof of Proposition 4 in the Online Appendix is not optimal for some parameter values satisfying the assumptions of the lemma. Details are available from the authors upon request.

[25]For this reason, when decreasing verification is optimal, the optimal rule is very sensitive to parameters, such as for instance the value of $\overline{\gamma} - \underline{\gamma}$.

In practice, principals may not have sufficient commitment power to implement allocations that are inefficient ex post, following verification. We explore the implications of limited commitment power in the next section.

5 Limited Commitment

We study a setting in which the principal has limited commitment power. The order of events is as follows:

1. The principal sets a rule, which maps a verification decision and result into an allowable spending set Π.
2. The agent chooses whether or not to seek verification, $a \in \{0, 1\}$, and the principal perfectly verifies his type γ if $a = 1$.
3. The principal revises the allowable spending set Π to Π'.
4. The agent chooses a spending level π from the allowable set Π'.

The first two steps are the same as those in our environment of Section 2 with full commitment power. What is new is Step 3: after observing the agent's verification decision and the result if verification is chosen, the principal now revises the allowable spending set for the agent. This is a mild form of limited commitment. In particular, in Step 2 we maintain the assumption that the principal is able to commit to a verification plan, so the agent's type is verified if and only if the agent requests verification, and in Step 4 we maintain the assumption that the principal is able to commit to allowing the agent to choose any spending level from the allowable spending set.[26] Our problem is therefore still one of delegation rather than cheap talk. The only assumption that we relax is about the principal's commitment to not changing the allowable spending set following the verification decision and result.[27] We believe lack of commitment in this respect often shapes delegation rules in the real world. For example, managers in organizations may request a revision of

[26]As noted in fn. 8, there is a literature that studies auditing when the principal cannot commit to an audit strategy. In many of the applications of our problem, however, we find that there are often institutions ensuring that principals cannot deny verification once it has been requested. In this sense, the agent can always choose to *trigger* verification. Lack of commitment in this respect would change the nature of our problem; we leave its analysis for future work.

[27]It is worth noting that our results in this section are not limited to the exact game described above; analogous to our claims in Section 2.3, our findings can be extended to variations of this game that allow messages between principal and agent (while maintaining our assumptions on the principal's limited commitment). We also emphasize that throughout this section, we maintain our optimality condition (5), so that a rule is optimal if it maximizes the principal's expected welfare and no other rule provides weakly larger welfare from each type γ and strictly larger welfare from some type γ.

their budgets for the next period; can headquarters commit to not changing their allocation ex post when no request is submitted? And in the case of a request, can headquarters commit to an inefficient budget after verifying the benefits of the manager's projects?

Limited commitment on the side of the principal matters for two reasons. First, conditional on no verification, the principal must choose an allocation that is optimal for the non-verified types. More precisely, when the agent chooses not to seek verification, the principal assigns spending taking into account the distribution of non-verified types and ignoring the incentives of verified types. A second implication of limited commitment is that conditional on verification, the principal learns the agent's true type γ and must assign the agent the efficient spending level $\pi_P(\gamma)$. This is true both when the agent's seeking verification is on path as well as when this verification decision is part of a deviation. As such, the agent can always choose to be verified to guarantee himself the efficient level of spending. Importantly, this means that all agent types who are not verified must weakly prefer their allocation under no verification to being verified and spending efficiently.

Limited commitment as a result implies certain conditions that any incentive compatible rule must satisfy. In what follows, we restrict attention to strategies that specify piecewise continuous mappings $\{a(\gamma), \pi(\gamma)\}$.

Lemma 7. *Under limited commitment, any incentive compatible rule satisfies:*

(i) If there is decreasing verification at γ_H, then

$$U_A(\gamma_H, \pi_P(\gamma_H)) - \alpha\phi = U_A(\gamma_H, \pi(\gamma_H)), \tag{12}$$

where $\pi(\gamma_H) \equiv \lim_{\varepsilon \downarrow 0} \pi(\gamma_H + \varepsilon)$ if $a(\gamma_H) = 1$. Moreover,

$$\pi(\gamma_H) > \pi_A(\gamma_H). \tag{13}$$

(ii) If there is increasing verification at γ_L, then

$$U_A(\gamma_L, \pi_P(\gamma_L)) - \alpha\phi = U_A(\gamma_L, \pi(\gamma_L)), \tag{14}$$

where $\pi(\gamma_L) \equiv \lim_{\varepsilon \downarrow 0} \pi(\gamma_L - \varepsilon)$ if $a(\gamma_L) = 1$.

Part (i) shows that if γ_H splits a verification region from a higher no-verification region, then γ_H must be indifferent between being verified and spending at the efficient level versus not being verified and spending at $\pi(\gamma_H)$ as allowed in the no-verification region above this type. Likewise, part (ii) shows that if γ_L splits a no-verification region from a higher verification region, then γ_L must be indifferent between being verified and spending at the efficient level versus not being verified and spending at $\pi(\gamma_L)$ as allowed in the no-verification region below this type. This result follows from the fact that a principal with limited commitment power assigns efficient spending whenever the agent seeks verification, both on and off path. Therefore, if there is a point at which a verification region either ends or starts, the marginal verified type at such point must weakly prefer verification with efficient spending to no verification, and the marginal non-verified type must weakly prefer no verification to verification with efficient spending. The marginal type must thus be indifferent.

Lemma 7 also shows that for type γ_H as defined in the lemma, an incentive compatible rule must set $\pi(\gamma_H) > \pi_A(\gamma_H)$. This is required to make γ_H indifferent between verification and no verification: if this inequality is not satisfied, the marginal verified type would instead prefer to deviate and not seek verification.

For the remainder of our analysis, we require:

Assumption 2. *If*

$$R(\gamma, \pi_H) \equiv U_A(\gamma, \pi_P(\gamma)) - \alpha\phi - U_A(\gamma, \pi_H) \geq 0$$

for $\pi_H > \pi_P(\gamma)$, *then*

$$R(\gamma', \pi_H) > 0 \text{ for all } \gamma' < \gamma.$$

This is a single-crossing property: we assume that if a type γ weakly prefers verification with efficient spending $\pi_P(\gamma)$ to no verification with a higher spending level $\pi_H > \pi_P(\gamma)$, then any lower type $\gamma' < \gamma$ strictly prefers verification with efficient spending $\pi_P(\gamma')$ to no verification with the higher spending level π_H.[28] A sufficient condition for this assumption is that if $R(\gamma, \pi_H) \geq 0$ for some $\gamma \in \Gamma$ and $\pi_H > \pi_P(\gamma)$, then $U_A(\gamma, \pi_P(\gamma))$ be convex in γ for all $\gamma \in \Gamma$; it can be established that in this case $U_A(\gamma, \pi_P(\gamma))$ is convex somewhere, and a sufficient condition is that it be convex everywhere. This convexity assumption is in fact satisfied in the cases commonly studied in the literature, such as those with quadratic preferences or with hyperbolic preferences under common parameterizations.[29]

Given Assumption 2, we obtain:

Proposition 5 (limited commitment)**.** *Under limited commitment, any incentive compatible rule features weakly increasing verification everywhere. Moreover, if verification is optimal, TEC is optimal.*

[28] Our single-crossing conditions on preferences imply that if a type γ weakly prefers verification with efficient spending $\pi_P(\gamma)$ to no verification with a lower spending level $\pi_L < \pi_P(\gamma)$, then any higher type $\gamma' > \gamma$ strictly prefers verification with efficient spending $\pi_P(\gamma')$ to no verification with the lower spending level π_L. Assumption 2 requires that this property be maintained in the opposite direction as well.

[29] For example, in the hyperbolic preferences case described in Section 2.1, $U_A(\gamma, \pi_P(\gamma))$ will be convex in γ if the utility functions for present and future consumption are the same and either exponential or CRRA with a coefficient weakly greater than 1.

Under limited commitment, decreasing verification is not incentive compatible for the principal. As we discussed in Section 4.3, decreasing verification requires that the principal commit to allowing the agent to spend at a level that is inefficient ex post, following the agent's verification decision and result. Without this commitment, the principal cannot induce decreasing verification, and hence any incentive compatible rule must feature weakly increasing verification at all types $\gamma \in \Gamma$. Analogous arguments to those behind Lemma 1 and Lemma 2 in our full-commitment environment then imply that if verifying some agent types is optimal, a TEC rule is optimal.

A sketch of the proof of Proposition 5 is as follows. Suppose by contradiction that there is an incentive compatible rule that induces decreasing verification, with γ_H being a type splitting a verification region from a higher no-verification region. Given limited commitment, verified types immediately below γ_H are assigned efficient spending, and types γ immediately above γ_H spend at a level $\pi_H > \pi_A(\gamma)$ that makes γ_H indifferent between verification and no verification (cf. Lemma 7). This means that types immediately above γ_H must be strictly overspending, in fact spending above their flexible level. The heart of the proof is showing that the principal cannot commit to allowing such overspending.

It is clear that conditional on the agent not seeking verification, the principal would like to reduce the overspending by types immediately above γ_H. Reducing this overspending is ex post incentive compatible for these types: having chosen no verification, types $\gamma > \gamma_H$ would prefer $\pi_A(\gamma)$ to $\pi_H > \pi_A(\gamma)$. Hence, the only reason the principal would not reduce the overspending immediately above γ_H once the agent chooses no verification is if doing so would violate incentive compatibility for some other non-verified type. Such a non-verified type must be below γ_H; specifically, there must exist a type $\gamma_L < \gamma_H$ who is not verified and is exactly indifferent between his assigned spending level, call it π_L, and the spending level $\pi_H > \pi_L$. In fact, because of single-crossing, this type must be the marginal type right below the verification region that ends at γ_H, i.e. the rule must induce verification for types $\gamma \in [\gamma_L, \gamma_H]$ and no verification for types immediately below and above this set. An example is the rule depicted in Figure 4.

Now if the principal induces such an interior verification region $[\gamma_L, \gamma_H]$, then by Lemma 7 type γ_L must be indifferent between no verification with spending π_L and verification with efficient spending. Since we have defined γ_L as being indifferent between spending at π_L and spending at π_H under no verification, by transitivity, we obtain that γ_L must be indifferent between no verification with spending π_H and verification with efficient spending. However, recall that type γ_H is also indifferent between no verification with spending π_H and verification with efficient spending. Hence, by Assumption 2, $\gamma_L < \gamma_H$ cannot hold,[30]

[30] If $\gamma_L < \gamma_H$, the indifference of type γ_H between verification with efficient spending and no verification

and we must have $\gamma_L = \gamma_H$. This means that the principal verifies a single type at this point who is indifferent between verification with efficient spending, no verification with higher spending at π_H, and no verification with lower spending at π_L. Conditional on no verification, this is thus an allocation in which the agent faces a hole $[\pi_L, \pi_H]$, namely he is not allowed to choose spending in this set but can choose spending immediately below and above this set. But our analysis in Section 3 shows that such a hole is suboptimal conditional on no verification; hence, following no verification, the principal would have a strict incentive to close the hole. This shows that a rule with decreasing verification cannot be incentive compatible when the principal has limited commitment power, allowing us to establish that TEC is optimal in this case.

It is worth pointing out that while TEC is optimal both when the principal has full commitment power and a small verification cost (as shown in Proposition 3) as well as when she has limited commitment power (as shown in Proposition 5), the specific details of an optimal TEC rule vary with each case. Under full commitment, an optimal TEC rule $\{\gamma^*, \gamma^{**}\}$ is such that the principal prefers to verify types $\gamma > \gamma^{**}$ to assign them efficient spending rather than bunch them at $\pi_A(\gamma^*)$ without verification, whereas the opposite is true for types $\gamma \in [\gamma^*, \gamma^{**}]$. Hence, the principal is indifferent between verifying and not verifying the threshold type γ^{**}; that is, the increase in assigned spending at γ^{**} exactly compensates the principal for the cost ϕ of verifying this type. In contrast, under limited commitment, it is the agent who is indifferent at γ^{**}: as implied by Lemma 7, type γ^{**} must be indifferent between being verified and assigned efficient spending versus not being verified and assigned $\pi_A(\gamma^*)$, and thus any increase in assigned spending at γ^{**} must exactly compensate this type for his verification cost $\alpha\phi$.

6 Conclusion

This paper has studied the tradeoff between commitment and flexibility in the presence of costly state verification. We have examined a general delegation problem in which a principal delegates decision making to an agent who has superior information about the efficient action but is biased towards higher actions. A novel element of our framework is that the principal can verify the agent's private information. Because verification is costly, the principal wishes to use this technology selectively, and in a way that supplements delegation and improves her commitment-versus-flexibility tradeoff.

Our results provide insight into how the principal achieves this by designing an optimal

with spending π_H would imply that γ_L strictly prefers verification with efficient spending to no verification with spending π_H, a contradiction.

delegation rule. We have shown that if the cost of verification is small enough, an optimal rule is a threshold with an escape clause (TEC), allowing the agent to freely select any action up to a threshold or to request verification and the efficient action if the threshold is sufficiently binding. If the cost of verification is larger, the principal may instead prefer to require verification only for intermediate actions, still imposing some discipline on the agent but saving on verification costs. The optimality of TEC is restored under mild limitations to the principal's commitment power: if the principal is unable to commit to not changing the agent's permissible action set following the verification decision and result, TEC is optimal for any verification cost for which verification is optimal.

As discussed in the Introduction, there is a variety of applications where delegation is central and rules make use of verification by specifying escape clauses. Our analysis sheds light on the optimal structure of escape clauses and provides a theoretical foundation for the common use of TEC rules. More broadly, our framework may help inform the empirical analysis of real-world rules. Data on delegation policies and the way verification is used is increasingly available and offers an opportunity to explore the design of these rules in more detail. For instance, data on fiscal rules around the world may be used to study how delegation varies with the institutional and macroeconomic context, which may affect both the cost of verifying a government's information and the importance of flexibility in responding to shocks.

Lastly, by uncovering a new set of issues that arise when verification is introduced to a setting in which both commitment and flexibility are valuable, our paper opens the door for further work that can help understand the optimal joint design of delegation and verification. We have focused on a simple model that emphasizes the main forces at play but abstracts from other potentially relevant aspects, for instance associated with more complex verification technologies. We close by discussing some possible extensions and variations of our work.

Random verification. As in the seminal work of Townsend (1979), we have considered deterministic verification, namely, we assumed that the principal's rule assigns $a(\gamma) \in \{0, 1\}$ to each agent type γ. More generally, one could allow for mechanisms in which the principal randomizes over the verification assignment, choosing a probability of verification for each type. The literature on financial contracting and tax collection finds that random verification can yield different results compared to deterministic verification; see Border and Sobel (1987) and Mookherjee and Png (1989).

Our focus on deterministic verification is motivated by the applications we study. Take capital budgeting or fiscal policy. As captured by the game form that we have proposed, here the agent (manager, policymaker) decides whether to request verification to obtain approval to choose actions that are not allowed by the principal (headquarters, society) without verification. The principal commits to following the agent's request, and so it is the agent's choice whether to trigger the verification process or not. Unlike in other applications where verification/audit is used to determine fines for misbehavior (e.g. tax collection), random verification is not natural in these contexts. Using the timing of Section 2.2, random verification would mean that the agent chooses in Step 2 not between verification and no-verification but rather between different lotteries over verification — something that is rarely observed in practice, possibly because committing to a non-degenerate lottery can be difficult for a principal.[31]

That said, the study of random verification in delegation could be an interesting extension of our work. As noted in the aforementioned literature, one issue is that an optimal randomized mechanism would depend on the extent to which the agent can be punished following verification, which in turn would depend on preference assumptions in our setting given that punishments are imposed through the spending allocation. Importantly, these punishments must be bounded: otherwise the efficient allocation can be approached with a rule that verifies all agent types with very low probability and arbitrarily punishes the agent when verification reveals that he has deviated.[32] Such a possibility not only yields rather implausible predictions, but also implies that an optimal rule in general will fail to exist unless a bound on punishments is imposed.

Imperfect verification. Also following Townsend (1979), our analysis assumed that verification reveals the agent's type perfectly. An alternative would be to consider imperfect verification, namely verification that provides only imperfect information about the agent's type. For example, in the context of capital budgeting in organizations, headquarters may review information about the benefits of a project that a manager advocates, but the available documentation may be incomplete and fail to reveal the full merits of the project.

A simple specification that may be possible to accommodate within our framework is when imperfect verification either reveals the agent's type perfectly or provides no information (i.e., when there are no "false" results). Provided that the principal can severely punish the agent (through the spending allocation), she would be able to prevent, at no cost, any deviation in which an agent type mimics another type who is verified, as is true in our problem with perfect verification. Yet, a difference introduced by imperfect verification

[31]When the decision is simply over verification or no-verification, commitment to the verification policy would in principle be facilitated by the fact that the principal's execution of the agent's request can be easily monitored. However, checking that the principal implements a specific lottery is harder, as it requires monitoring of the randomization itself rather than its outcome.

[32]In our game form, a rule that approaches the efficient allocation would be implemented by inducing each agent type to choose a different lottery over verification.

is that the principal may not observe the agent's type and thus may not be able to assign a type-dependent spending level following verification; the principal's rule must specify a spending allocation for the case of verification and no information. Allowing for imperfect verification that may produce false results would naturally introduce further issues, as now punishing an agent type for mimicking another type who is verified would require imposing punishments on path.

How imperfect is imperfect verification? At one extreme, if verification is sufficiently accurate, we conjecture that our qualitative results would remain valid. At the other extreme, if verification is sufficiently inaccurate, it would become equivalent to money burning, and the results of the literature on when money burning is used in an optimal delegation rule would then apply (see Amador, Werning and Angeletos, 2006; Amador and Bagwell, 2013; Ambrus and Egorov, 2015). More generally, it would be of interest to explore the role of verification in delegation away from these two extremes.

Verification costs. We have considered verification costs that are both type-independent and exogenous. An extension of our problem could explore the effects of type-dependent verification costs: the principal's cost of verifying the agent's private information may be increasing in his type, for example because more evidence is needed to verify larger project benefits, or one may take the view that verification costs are actually lower for extreme types, as these states are more "visible." One possible difficulty is that monotonicity of the spending allocation (as shown in Lemma 5) may fail to hold if verification costs increase very rapidly with the agent's type. But if the verification cost function is such that the principal would still prefer to swap the verification and spending allocations of two types γ and $\gamma' > \gamma$ whenever type γ has higher spending than γ', monotonicity will be satisfied and our analysis could be extended to allow for type-dependent verification costs.

Another variation would be to endogenize α, so that the principal can affect the agent's cost of verification. In our problem with full commitment power, the principal would optimally set $\alpha = 0$, as a zero cost of verification for the agent maximally relaxes his incentive compatibility constraint (4). Things are less straightforward in the setting of Section 5 where the principal has limited commitment power: here the principal may want to set a strictly positive verification cost for the agent in order to limit the set of agent types that may want to demand verification and efficient spending.

Transfers. Our focus has been on a canonical delegation problem in which transfers between the principal and the agent are not feasible. There are various ways in which transfers could be introduced in our framework and used to alter the feasibility and cost of inducing different allocations. Transfers could be contingent on the agent's verification decision and/or the verification result; moreover, the principal could offer different allowable spending sets for the agent to choose from and specify transfers associated with each set. These questions are beyond the scope of our paper and so we leave them for future research.

References

Aguiar, Mark and Manuel Amador, "Growth under the Shadow of Expropriation," *Quarterly Journal of Economics*, 2011, *126*, 651–97.

Alesina, Alberto and Andrea Passalacqua, "The Political Economy of Government Debt," in "Handbook of Macroeconomics," Vol. 2, North Holland, 2016, pp. 2599–2651.

Alonso, Ricardo and Niko Matouschek, "Optimal Delegation," *Review of Economic Studies*, 2008, *75*, 259–293.

Amador, Manuel and Kyle Bagwell, "The Theory of Delegation with an Application to Tariff Caps," *Econometrica*, 2013, *81* (4), 1541–600.

_ , Iván Werning, and George-Marios Angeletos, "Commitment Vs. Flexibility," *Econometrica*, 2006, *74*, 365–396.

_ , Kyle Bagwell, and Alex Frankel, "A Note on Interval Delegation," *Economic Theory Bulletin*, 2017, *Forthcoming.*

Ambrus, Attila and Georgy Egorov, "Comment on 'Commitment vs. Flexibility'," *Econometrica*, 2013, *81* (5), 2113–2124.

_ and _ , "Delegation and Nonmonetary Incentives," 2015. Working Paper.

Athey, Susan, Andrew Atkeson, and Patrick J. Kehoe, "The Optimal Degree of Discretion in Monetary Policy," *Econometrica*, 2005, *73*, 1431–1475.

Auster, Sarah and Nicola Pavoni, "Optimal Delegation and Limited Awareness, with an Application to Financial Intermediation," 2017. Working Paper.

Banks, Jeffrey, "Agency Budgets, Cost Information, and Auditing," *American Journal of Political Science*, 1989, *33* (3), 670–699.

Ben-Porath, Elchanan and Barton L. Lipman, "Implementation with Partial Provability," *Journal of Economic Theory*, 2012, *147* (5), 1689–1724.

— , **Eddie Dekel, and Barton Lipman**, "Optimal Allocation with Costly Verification," *American Economic Review*, 2014, *104* (12), 3779–3813.

Bernheim, B. Douglas, Debraj Ray, and Sevin Yeltekin, "Poverty and Self-Control," *Econometrica*, 2015, *83* (5), 1877–1911.

Beshkar, Mostafa and Eric Bond, "Cap and Escape in Trade Agreements," 2016. Working Paper.

Border, Kim C. and Joel Sobel, "Samurai Accountant: A Theory of Auditing and Plunder," *Review of Economic Studies*, 1987, *54* (4), 525–540.

Bull, Jesse and Joel Watson, "Hard Evidence and Mechanism Design," *Games and Economic Behavior*, 2007, *58*, 75–93.

Chatterjee, Kalyan, Sanford Morton, and Arijit Mukherji, "Strategic Audit Policies Without Commitment," 2002. Working Paper.

Deneckere, Raymond and Sergei Severinov, "Mechanism Design with Partial State Verifiability," *Games and Economic Behavior*, 2008, *64*, 487–513.

Dovis, Alessandro and Rishabh Kirpalani, "Fiscal Rules, Bailouts, and Reputation in Federal Governments," 2017.

Erlanson, Albin and Andreas Kleiner, "Costly Verification in Collective Decisions," November 2015. Working Paper.

Gale, Douglas and Martin Hellwig, "Incentive-Compatible Debt Contracts: The One-Period Problem," *Review of Economic Studies*, 1985, *52* (4), 647–663.

Glazer, Jacob and Ariel Rubinstein, "On Optimal Rules of Persuasion," *Econometrica*, 2004, *72* (6), 1715–1736.

— **and** — , "A Study in the Pragmatics of Persuasion: A Game Theoretical Approach," *Theoretical Economics*, 2006, *1*, 395–410.

Green, Jerry R. and Jean-Jacques Laffont, "Partially Verifiable Information and Mechanism Design," *Review of Economic Studies*, 1986, *53* (3), 447–456.

Halac, Marina and Pierre Yared, "Fiscal Rules and Discretion under Persistent Shocks," *Econometrica*, 2014, *82* (5), 1557–614.

— **and** — , "Fiscal Rules and Discretion in a World Economy," 2017. Working Paper.

— **and** — , "Fiscal Rules and Discretion under Self-Enforcement," 2017. Working Paper.

Harris, Milton and Artur Raviv, "The Capital Budgeting Process: Incentives and Information," *Journal of Finance*, 1996, *51* (4), 1139–1174.

— **and** — , "Capital budgeting and delegation," *Journal of Financial Economics*, 1998, *50*, 259–289.

Holmström, Bengt, "On Incentives and Control in Organizations," 1977. Ph.D. Thesis, Stanford University.

— , "On the Theory of Delegation," in M. Boyer and R. Kihlstrom, eds., *Bayesian Models in Economic Theory*, North-Holland; New York, 1984.

Jackson, Matthew O. and Leeat Yariv, "Present Bias and Collective Dynamic Choice in the Lab," *American Economic Review*, 2014, *104* (12), 4184–204.

— **and** — , "Collective Dynamic Choice: The Necessity of Time Inconsistency," *American Economic Journal: Microeconomics*, 2015, *7* (4), 150–178.

Kartik, Navin and Olivier Tercieux, "Implementation with Evidence," *Theoretical Economics*, 2012, *7*, 323–355.

Lo, Desmond (Ho-Fu), Wouter Dessein, Mrinal Ghosh, and Francine Lafontaine, "Price Delegation and Performance Pay: Evidence from Industrial Sales Forces," *Journal of Law, Economics, and Organization*, 2016, *32* (3), 508–544.

Malenko, Andrey, "Optimal Dynamic Capital Budgeting," 2016. Working Paper.

Melumad, Nahum D. and Toshiyuki Shibano, "Communication in Settings with No Transfers," *RAND Journal of Economics*, 1991, *22* (2), 173–198.

Mookherjee, Dilip and Ivan Png, "Optimal Auditing, Insurance, and Redistribution," *Quarterly Journal of Economics*, 1989, *104* (2), 399–415.

Mylovanov, Tymofiy and Andriy Zapechelnyuk, "Mechanism Design with Ex-Post Verification and Limited Punishments," 2014. Working Paper.

Reinganum, Jennifer and Louis Wilde, "Equilibrium Verification and Reporting Policies in a Model of Tax Compliance," *International Economic Review*, 1986, *27*, 739–760.

Ross, Marc, "Capital Budgeting Practices of Twelve Large Manufacturers," *Financial Management*, 1986, *15* (4), 15–22.

Schaechter, Andrea, Tidiane Kinda, Nina Budina, and Anke Weber, "Fiscal Rules in Response to the Crisis – Toward the 'Next-Generation' Rules. A New Dataset," 2012. IMF Working Paper WP/12/187.

Strausz, Roland, "Deterministic Mechanisms and the Revelation Principle," *Economics Letters*, 2003, *79*, 333–337.

Taggart, Robert, "Allocating Capital among a Firm's Divisions: Hurdle Rates vs. Budgets," *Journal of Financial Research*, 1987, *10*, 177–190.

Townsend, Robert M., "Optimal Contracts and Competitive Markets with Costly State Verification," *Journal of Economic Theory*, 1979, *21*, 265–293.

_ , "Information Constrained Insurance: The Revelation Principle Extended," *Journal of Monetary Economics*, 1988, *21*, 411–450.

A Appendix

A.1 Proof of Proposition 1

The claim follows from Proposition 1(a) in Amador and Bagwell (2013, p. 1551).

A.2 Proof of Lemma 1

Suppose by contradiction that a rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ with $a(\gamma) = 1$ for all $\gamma \in \Gamma$ is optimal. Since the incentive compatibility constraint (4) is trivially satisfied under this rule, it must be that $\pi(\gamma) = \pi_P(\gamma)$ for all $\gamma \in \Gamma$. Define $\gamma' \in \Gamma$ as the solution to

$$U_A(\gamma', \pi_P(\gamma')) - \alpha\phi = U_A\left(\gamma', \pi_P\left(\underline{\gamma}\right)\right) \tag{15}$$

if such a solution exists and $\gamma' = \overline{\gamma}$ otherwise. Consider now a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma\in\Gamma}$ with $\widetilde{a}(\gamma) = 0$, $\widetilde{\pi}(\gamma) = \pi_P\left(\underline{\gamma}\right)$ for $\gamma \leq \gamma'$, and $\widetilde{a}(\gamma) = a(\gamma)$, $\widetilde{\pi}(\gamma) = \pi(\gamma)$ for $\gamma > \gamma'$. By single-crossing and the definition of γ' in (15), the perturbed rule satisfies the incentive compatibility constraint (4). Conditional on $\gamma > \gamma'$, this rule yields the same expected welfare to the principal and the agent as the original rule. However, conditional on $\gamma \leq \gamma'$, the perturbed rule yields the agent a higher welfare than the original one, since, by (15),

$$U_A(\gamma, \pi_P(\gamma)) - \alpha\phi \leq U_A\left(\gamma, \pi_P\left(\underline{\gamma}\right)\right) \tag{16}$$

for all $\gamma \leq \gamma'$. Moreover, note that (2) implies

$$U_A\left(\gamma, \pi_P\left(\gamma\right)\right) - U_A\left(\gamma, \pi_P\left(\underline{\gamma}\right)\right) > U_P\left(\gamma, \pi_P\left(\gamma\right)\right) - U_P\left(\gamma, \pi_P\left(\underline{\gamma}\right)\right)$$

for all $\gamma > \underline{\gamma}$. Hence, using (16) and the fact that $\alpha \in [0, 1]$ and $\phi > 0$,

$$U_P\left(\gamma, \pi_P\left(\gamma\right)\right) - \phi < U_P\left(\gamma, \pi_P\left(\underline{\gamma}\right)\right)$$

for all $\gamma \leq \gamma'$. Conditional on $\gamma \leq \gamma'$, the principal is therefore strictly better off under the perturbed rule than under the original rule. It follows that the perturbed rule with no verification below γ' strictly dominates the original rule, contradicting the optimality of a rule that verifies all types.[33]

A.3 Proof of Proposition 2

Assume $b\left(\pi\right) = 0$ for all $\pi \in [\underline{\pi}, \overline{\pi}]$. Suppose by contradiction that an optimal rule specifies $a\left(\gamma\right) = 1$ for some $\gamma \in \Gamma$ but TEC is not optimal. By Lemma 2, this rule must feature decreasing verification. We proceed by showing that an optimal rule cannot feature decreasing verification at any $\gamma' \in \Gamma$.

Consider first decreasing verification at some $\gamma' \in \Gamma$ with $a\left(\gamma'\right) = 0$, so $a\left(\gamma' - \varepsilon\right) = 1$ for some $\varepsilon > 0$ arbitrarily small. As shown in the text, the optimality of verifying type $\gamma' - \varepsilon$ implies (6) and $\pi(\gamma' - \varepsilon) > \pi_A\left(\gamma^*\right)$, whereas the optimality of not verifying γ' implies (7). However, the two equations together with $\pi(\gamma' - \varepsilon) > \pi_A\left(\gamma^*\right)$ violate the single-crossing condition (1). Contradiction.

Consider next decreasing verification at some $\gamma' \in \Gamma$ with $a\left(\gamma'\right) = 1$, so $a\left(\gamma' + \varepsilon\right) = 0$ for some $\varepsilon > 0$ arbitrarily small. Analogous arguments to those above apply to this case and yield a contradiction.

A.4 Proof of Lemma 3

Suppose a TEC rule with cutoffs γ^* and γ^{**} is a solution to (8)-(9). Note that any rule satisfying constraint (4) will satisfy constraint (9). Hence, (8)-(9) is a relaxed version of (3)-(4), implying that any solution to (8)-(9) that satisfies (4) will also be a solution to (3)-(4). It follows that to prove the claim, all we need to show is that the TEC rule that

[33] Note that in the case of $\alpha = 0$, $\gamma' = \underline{\gamma}$, and the perturbed rule increases the principal's welfare from type $\underline{\gamma}$. Moreover, in this case, we can also consider a perturbed rule that prescribes $\widetilde{a}(\gamma) = 0$ for all $\gamma \in [\underline{\gamma}, \underline{\gamma} + \varepsilon]$, for $\varepsilon > 0$ arbitrarily small, bunching all such γ at their average efficient spending level. This rule is incentive compatible and increases the principal's welfare relative to verifying all types: given ε small enough, the welfare loss from not assigning efficient spending to $\gamma \in [\underline{\gamma}, \underline{\gamma} + \varepsilon]$ is second order, while the gain from saving on verification costs is first order.

solves (8)-(9) will satisfy constraint (4). It is immediate that for any γ with $a(\gamma) = 0$, (9) being satisfied implies that (4) will be satisfied. Now consider γ with $a(\gamma) = 1$. Optimality of verifying type γ under a TEC rule that solves (8)-(9) implies

$$U_P(\gamma, \pi_P(\gamma)) - \phi \geq U_P(\gamma, \pi_A(\gamma^*)), \tag{17}$$

since a perturbation that assigns no verification and spending level $\pi_A(\gamma^*)$ to a type $\gamma > \gamma^{**}$ is incentive compatible. Note that by the arguments in the proof of Lemma 2, a TEC rule that solves (8)-(9) satisfies $\pi_P(\gamma) \geq \pi_A(\gamma^*)$ for all $\gamma > \gamma^{**}$. Hence, combining (17) with (2) and the fact that $\alpha \in [0,1]$ implies

$$U_A(\gamma, \pi_P(\gamma)) - \alpha\phi \geq U_A(\gamma, \pi_A(\gamma^*)).$$

It follows that (4) is satisfied for type γ with $a(\gamma) = 1$.

A.5 Proof of Lemma 6

Suppose a rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ solves (8)-(9) and features decreasing verification at some $\gamma' < \overline{\gamma}$ with $a(\gamma') = 1$. Then $a(\gamma' + \varepsilon) = 0$ for some $\varepsilon > 0$ arbitrarily small. Suppose it were the case that $\pi(\gamma' + \varepsilon) = \pi(\gamma')$. Then optimality of this rule would be violated, as a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma\in\Gamma}$ which sets $\widetilde{a}(\gamma') = 0$ and $\widetilde{\pi}(\gamma') = \pi(\gamma')$ while keeping the allocation unchanged for $\gamma \neq \gamma'$ would be incentive compatible and strictly increase the principal's welfare (recall $\phi > 0$). It follows that $\pi(\gamma' + \varepsilon) \neq \pi(\gamma')$, and hence by Lemma 5, $\pi(\gamma' + \varepsilon) > \pi(\gamma')$. Moreover, by Lemma 4, $\pi(\gamma') \leq \pi_A(\gamma')$, and thus incentive compatibility for γ' would be violated if it were the case that $\pi_A(\gamma') \geq \pi(\gamma' + \varepsilon) > \pi(\gamma')$. It therefore follows that

$$\pi(\gamma' + \varepsilon) > \pi_A(\gamma') \tag{18}$$

for $\varepsilon > 0$ arbitrarily small. Lemma 5 then implies $\pi(\gamma) > \pi_A(\gamma')$ for all $\gamma \in (\gamma', \gamma'')$, $\gamma'' \equiv \min\left\{\pi_P^{-1}(\pi_A(\gamma')), \overline{\gamma}\right\}$, which implies

$$\int_{\gamma'}^{\gamma''} U_P(\gamma, \pi(\gamma)) f(\gamma)\, d\gamma < \int_{\gamma'}^{\gamma''} U_P(\gamma, \pi_A(\gamma')) f(\gamma)\, d\gamma. \tag{19}$$

Moreover, by definition,

$$\int_{\gamma''}^{\overline{\gamma}} U_P(\gamma, \pi(\gamma)) f(\gamma)\, d\gamma \leq \int_{\gamma''}^{\overline{\gamma}} U_P(\gamma, \pi_P(\gamma)) f(\gamma)\, d\gamma. \tag{20}$$

Combining (19) and (20), and taking into account that $1 - F(\gamma') > 0$, yields (10).

Suppose next that a rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ solves (8)-(9) and features decreasing verification at some $\gamma' < \overline{\gamma}$ with $a(\gamma') = 0$. Then $a(\gamma' - \varepsilon) = 1$ for some $\varepsilon > 0$ arbitrarily small and arguments analogous to those above yield (10).

A.6 Proof of Proposition 3

The arguments in the proofs of Lemma 1 and Lemma 2 apply to the relaxed problem, implying that if a solution to (8)-(9) involves verifying some type $\gamma \in \Gamma$, this solution is either a TEC rule or a rule that features decreasing verification at some $\gamma' \in \Gamma$. To prove the optimality of TEC for $\phi < \overline{\phi}$, we thus proceed by showing that for any such verification cost a rule featuring decreasing verification cannot be a solution to (8)-(9).

Suppose a rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ solves (8)-(9) and features decreasing verification. Denote by γ^{**} the infimum of the lowest verification region under this rule. Now consider a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma\in\Gamma}$ which sets $\widetilde{a}(\gamma) = 0$ for $\gamma < \gamma^{**}$, $\widetilde{a}(\gamma^{**}) = a(\gamma^{**})$, and $\widetilde{a}(\gamma) = 1$ for $\gamma > \gamma^{**}$. If $\widetilde{a}(\gamma) = 0$, let $\widetilde{\pi}(\gamma) = \min\{\pi_A(\gamma), \pi_A(\gamma^*)\}$ for γ^* as defined in Proposition 1 under $\overline{\gamma}' = \gamma^{**}$. If $\widetilde{a}(\gamma) = 1$, let $\widetilde{\pi}(\gamma) = \pi_P(\gamma)$. By the arguments in the proof of Lemma 2, this rule is incentive compatible for types prescribed no verification and sets $\pi_A(\gamma^*) \leq \pi_P(\gamma^{**})$. Moreover, given this inequality and the fact that $\alpha = 0$, it follows that the rule is also incentive compatible for types prescribed verification. We now show that this rule strictly increases the principal's expected welfare for $\phi < \overline{\phi}$, contradicting the optimality of the original rule. Denote by γ' the lowest type above γ^{**} featuring decreasing verification in the original rule. Then the change in the principal's expected welfare from using the perturbed rule instead of the original rule is

$$\int_{\underline{\gamma}}^{\gamma^{**}} \big(U_P(\gamma, \min\{\pi_A(\gamma), \pi_A(\gamma^*)\}) - U_P(\gamma, \pi(\gamma))\big) f(\gamma)\, d\gamma \tag{21}$$
$$+ \int_{\gamma^{**}}^{\overline{\gamma}} \big(U_P(\gamma, \pi_P(\gamma)) - U_P(\gamma, \pi(\gamma))\big) f(\gamma)\, d\gamma$$
$$- \int_{\gamma'}^{\overline{\gamma}} [\phi(1 - a(\gamma))] f(\gamma)\, d\gamma.$$

Note that since all types above γ^{**} are verified, the principal's welfare conditional on the agent's type being in the no-verification region of the perturbed rule is optimized subject to fewer incentive compatibility constraints in this rule compared to the original rule. Hence, the first term in (21) is weakly positive.

To evaluate the second and third terms in (21), suppose first that $\gamma' < \overline{\gamma}$. Then by Lemma 6, the second term in (21) satisfies

$$\begin{aligned}\int_{\gamma^{**}}^{\overline{\gamma}}\big(U_P(\gamma,\pi_P(\gamma))-U_P(\gamma,\pi(\gamma))\big)f(\gamma)\,d\gamma &\geq \int_{\gamma'}^{\overline{\gamma}}\big(U_P(\gamma,\pi_P(\gamma))-U_P(\gamma,\pi(\gamma))\big)f(\gamma)\,d\gamma \\ &\geq (1-F(\gamma'))\,\eta(\gamma')\,. \end{aligned}\tag{22}$$

Moreover, the third term in (21) satisfies

$$\begin{aligned}-\int_{\gamma'}^{\overline{\gamma}}[\phi(1-a(\gamma))]\,f(\gamma)\,d\gamma &> -(1-F(\gamma'))\,\overline{\phi} \\ &= -(1-F(\gamma'))\min_{\gamma\in\Gamma}\eta(\gamma)\,. \end{aligned}\tag{23}$$

Together, (22) and (23) imply that the perturbation strictly increases welfare.

Suppose next that $\gamma' = \overline{\gamma}$. Analogous arguments to those above imply that the perturbation makes the principal weakly better off conditional on $\gamma < \overline{\gamma}$. To evaluate the change in welfare conditional on $\gamma = \overline{\gamma}$, note that in this case we must have $a(\overline{\gamma}) = 0$ and $a(\overline{\gamma}-\varepsilon) = 1$ for $\varepsilon > 0$ arbitrarily small. Analogous arguments to those in the proof of Lemma 6 then imply $\pi(\overline{\gamma}) \geq \pi_A(\overline{\gamma})$. Moreover, by (11),

$$\eta(\overline{\gamma}) = \lim_{\gamma\uparrow\overline{\gamma}}\eta(\gamma) = U_P(\overline{\gamma},\pi_P(\overline{\gamma})) - U_P(\overline{\gamma},\pi_A(\overline{\gamma})) \geq \overline{\phi} > \phi, \tag{24}$$

where we have appealed to the definition of $\overline{\phi}$. It thus follows from (24) that the perturbation strictly increases the principal's welfare conditional on $\gamma = \overline{\gamma}$.

Online Appendix

B Omitted Proofs

B.1 Proof of Lemma 2

Suppose an optimal rule features verification which is weakly increasing everywhere. By Lemma 1, $a(\gamma) = 0$ for some $\gamma \in \Gamma$, and hence this rule must feature a no-verification region followed by a verification region. That is, the principal solves (3)-(4) by choosing a threshold γ^{**} such that $a(\gamma) = 0$ for $\gamma < \gamma^{**}$ and $a(\gamma) = 1$ for $\gamma > \gamma^{**}$, and a spending allocation $\pi(\gamma)$ for each $\gamma \in \Gamma$.

Now consider a relaxed version of this problem in which the principal chooses an optimal allocation in the no-verification and verification regions separately, ignoring the incentives of types in one region to deviate to the other region. Taking the no-verification region to be $[\underline{\gamma}, \gamma^{**}]$, it follows from Proposition 1 that an optimal allocation is a threshold $\gamma^* < \gamma^{**}$ such that $\pi(\gamma) = \min\{\pi_A(\gamma), \pi_A(\gamma^*)\}$ for each $\gamma \in [\underline{\gamma}, \gamma^{**}]$. For the verification region $(\gamma^{**}, \overline{\gamma}]$, since incentive compatibility is trivially satisfied, an optimal allocation assigns $\pi_P(\gamma)$ to each $\gamma \in (\gamma^{**}, \overline{\gamma}]$. Note that the resulting rule for the whole set Γ is TEC. Moreover, because this rule solves a relaxed problem, it is sufficient to show that it is incentive compatible over the whole set Γ to prove its optimality in the original problem.

To show incentive compatibility, note first that incentive compatibility within each region is guaranteed by construction. Furthermore, since, as explained in Section 2.3, no type would have incentives to deviate to mimic a different type which is verified, incentive compatibility is satisfied for all $\gamma \in [\underline{\gamma}, \gamma^{**}]$. All is left to be shown is that no type $\gamma \in (\gamma^{**}, \overline{\gamma}]$ has incentives to deviate to mimic a type $\widehat{\gamma} \in [\underline{\gamma}, \gamma^{**}]$:

$$U_A(\gamma, \pi_P(\gamma)) - \alpha\phi \geq U_A(\gamma, \pi(\widehat{\gamma})) \text{ for all } \gamma > \gamma^{**}, \widehat{\gamma} \leq \gamma^{**}.$$

The single-crossing condition in U_A implies that a sufficient condition for the above inequality to hold is

$$U_A(\gamma, \pi_P(\gamma)) - \alpha\phi \geq U_A(\gamma, \pi_A(\gamma^*)) \text{ for all } \gamma > \gamma^{**}. \tag{25}$$

Now note that optimality of γ^{**} for the principal implies

$$U_P(\gamma, \pi_P(\gamma)) - \phi \geq U_P(\gamma, \pi_A(\gamma^*)) \text{ for all } \gamma > \gamma^{**}. \tag{26}$$

Given the agent's bias (2) and $\alpha \in [0, 1]$, (26) implies (25) if $\pi_P(\gamma) \geq \pi_A(\gamma^*)$ for all $\gamma > \gamma^{**}$, or equivalently since $\pi'_P(\gamma) > 0$, if

$$\pi_P(\gamma^{**}) \geq \pi_A(\gamma^*). \tag{27}$$

We prove that the TEC rule that we constructed satisfies (27). The optimal threshold γ^* in the no-verification region solves

$$\max_{\gamma^* \in [0,\gamma^{**})} \left\{ \int_{\underline{\gamma}}^{\gamma^*} U_P\left(\gamma, \pi_A(\gamma)\right) f\left(\gamma\right) d\gamma + \int_{\gamma^*}^{\gamma^{**}} U_P\left(\gamma, \pi_A(\gamma^*)\right) f\left(\gamma\right) d\gamma \right\}.$$

The first-order condition yields

$$\int_{\gamma^*}^{\gamma^{**}} \frac{\partial U_P\left(\gamma, \pi_A(\gamma^*)\right)}{\partial \pi_A(\gamma^*)} \pi'_A(\gamma^*) f\left(\gamma\right) d\gamma = 0.$$

Note that $\pi'_A(\gamma^*) > 0$, $\frac{\partial U_P(\gamma, \pi_A(\gamma^*))}{\partial \pi_A(\gamma^*)} < 0$ if $\pi_P(\gamma) < \pi_A(\gamma^*)$, and $\frac{\partial U_P(\gamma, \pi_A(\gamma^*))}{\partial \pi_A(\gamma^*)} > 0$ if $\pi_P(\gamma) > \pi_A(\gamma^*)$. Hence, the first-order condition requires $\pi_P(\gamma) > \pi_A(\gamma^*)$ for some $\gamma \in [\gamma^*, \gamma^{**}]$, implying that (27) must hold.

B.2 Proof of Lemma 4

Suppose a rule $\{a(\gamma), \pi(\gamma)\}_{\gamma \in \Gamma}$ solving (8)-(9) specifies $a\left(\gamma'\right) = 1$ for some $\gamma' \in \Gamma$.

To prove that the rule specifies $\pi\left(\gamma'\right) \leq \pi_A\left(\gamma'\right)$, suppose by contradiction that $\pi\left(\gamma'\right) > \pi_A\left(\gamma'\right)$. Consider a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma \in \Gamma}$ which sets $\widetilde{a}\left(\gamma'\right) = 1$ and $\widetilde{\pi}\left(\gamma'\right) = \pi_A\left(\gamma'\right)$ while keeping the allocation unchanged for all $\gamma \neq \gamma'$. This perturbation strictly increases the principal's welfare conditional on γ', leaves the principal's welfare conditional on $\gamma \neq \gamma'$ unchanged, and is incentive compatible for the agent.

Similarly, to prove that the rule specifies $\pi\left(\gamma'\right) \geq \pi_P\left(\gamma'\right)$, suppose by contradiction that $\pi\left(\gamma'\right) < \pi_P\left(\gamma'\right)$. Consider a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma \in \Gamma}$ which sets $\widetilde{a}\left(\gamma'\right) = 1$ and $\widetilde{\pi}\left(\gamma'\right) = \pi_P\left(\gamma\right)$ while keeping the allocation unchanged for all $\gamma \neq \gamma'$. This perturbation strictly increases the principal's welfare conditional on γ', leaves the principal's welfare conditional on $\gamma \neq \gamma'$ unchanged, and is incentive compatible for the agent.

Finally, we prove that the rule must specify $\pi\left(\gamma'\right) = \pi_P\left(\gamma'\right)$ if (9) does not bind for γ'. Suppose by contradiction that (9) does not bind for γ' and $\pi\left(\gamma'\right) \neq \pi_P\left(\gamma'\right)$. By the claim above, $\pi\left(\gamma'\right) \geq \pi_P\left(\gamma'\right)$, and thus the rule must set $\pi\left(\gamma'\right) > \pi_P\left(\gamma'\right)$. But then a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma \in \Gamma}$ which sets $\widetilde{a}\left(\gamma'\right) = 1$ and $\widetilde{\pi}\left(\gamma'\right) = \pi\left(\gamma'\right) - \varepsilon$ for $\varepsilon > 0$ arbitrarily small, while keeping the allocation unchanged for all $\gamma \neq \gamma'$, strictly increases the principal's welfare conditional on γ', leaves the principal's welfare conditional on $\gamma \neq \gamma'$ unchanged, and is incentive compatible for the agent.

B.3 Proof of Lemma 5

Suppose by contradiction that a rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ that solves (8)-(9) specifies $\pi(\gamma') > \pi(\gamma'')$ for some $\gamma' < \gamma''$. We consider four cases separately.

Case 1. Suppose $a(\gamma') = a(\gamma'') = 0$. Then (9) for γ' and γ'' requires

$$\begin{aligned} U_A(\gamma', \pi(\gamma')) &\geq U_A(\gamma', \pi(\gamma'')), \\ U_A(\gamma'', \pi(\gamma'')) &\geq U_A(\gamma'', \pi(\gamma')), \end{aligned}$$

which together imply

$$U_A(\gamma', \pi(\gamma')) - U_A(\gamma', \pi(\gamma'')) \geq U_A(\gamma'', \pi(\gamma')) - U_A(\gamma'', \pi(\gamma'')). \tag{28}$$

However, given $\gamma' < \gamma''$ and $\pi(\gamma') > \pi(\gamma'')$, (28) violates the single-crossing condition in U_A. Contradiction.

Case 2. Suppose $a(\gamma') = a(\gamma'') = 1$. By Lemma 4, $\pi(\gamma'') \geq \pi_P(\gamma'')$, and thus $\pi(\gamma') > \pi(\gamma'')$ implies $\pi(\gamma') > \pi_P(\gamma'') > \pi_P(\gamma')$. Using Lemma 4 again, it then follows that (9) binds for γ', that is, there exists $\widehat{\gamma} \in \Gamma$ with $a(\widehat{\gamma}) = 0$ such that

$$U_A(\gamma', \pi(\gamma')) = U_A(\gamma', \pi(\widehat{\gamma})). \tag{29}$$

Furthermore, note that we must have $\pi(\widehat{\gamma}) \geq \pi(\gamma')$, since $\pi(\gamma') \leq \pi_A(\gamma')$ and U_A is strictly concave. Incentive compatibility for γ'' requires

$$U_A(\gamma'', \pi(\gamma'')) \geq U_A(\gamma'', \pi(\widehat{\gamma})),$$

which, combined with the observation that

$$\pi(\gamma'') < \pi(\gamma') \leq \pi_A(\gamma') < \pi_A(\gamma''), \tag{30}$$

implies

$$U_A(\gamma'', \pi(\gamma')) > U_A(\gamma'', \pi(\widehat{\gamma})). \tag{31}$$

Combining (29) and (31) yields

$$U_A(\gamma', \pi(\widehat{\gamma})) - U_A(\gamma', \pi(\gamma')) > U_A(\gamma'', \pi(\widehat{\gamma})) - U_A(\gamma'', \pi(\gamma')). \tag{32}$$

However, given $\gamma' < \gamma''$ and $\pi(\widehat{\gamma}) \geq \pi(\gamma')$, (32) violates the single-crossing condition in U_A. Contradiction.

Case 3. Suppose $a(\gamma') = 1$ and $a(\gamma'') = 0$. Note that (30) must hold. Then consider a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma\in\Gamma}$ which sets $\widetilde{a}(\gamma'') = 1$ and $\widetilde{\pi}(\gamma'') = \pi(\gamma')$ while leaving the allocation for types $\gamma \neq \gamma''$ unchanged. Since incentive compatibility was initially satisfied and $\gamma' < \gamma''$ while (30) holds, this perturbation is incentive compatible. Optimality of the original rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ therefore requires that this perturbation do not strictly increase the principal's welfare, which requires

$$U_P(\gamma'', \pi(\gamma'')) \geq U_P(\gamma'', \pi(\gamma')) - \phi.$$

The single-crossing condition in U_P then implies

$$U_P(\gamma', \pi(\gamma'')) > U_P(\gamma', \pi(\gamma')) - \phi. \tag{33}$$

Now consider a different perturbed rule $\{\widehat{a}(\gamma), \widehat{\pi}(\gamma)\}_{\gamma\in\Gamma}$ which sets $\widehat{a}(\gamma') = 0$ and $\widehat{\pi}(\gamma') = \pi(\gamma'')$ while leaving the allocation for types $\gamma \neq \gamma'$ unchanged. Equation (33) implies that this perturbation would strictly increase the principal's welfare. Hence, optimality of the original rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ requires that this perturbation violate incentive compatibility, that is, there must exist $\widehat{\gamma} \in \Gamma$ with $a(\widehat{\gamma}) = 0$ such that

$$U_A(\gamma', \pi(\widehat{\gamma})) > U_A(\gamma', \pi(\gamma'')). \tag{34}$$

Note that since $\pi(\gamma'') < \pi_A(\gamma')$, we must have $\pi(\widehat{\gamma}) > \pi(\gamma'')$. Moreover, by incentive compatibility being satisfied under the original rule, we have

$$U_A(\gamma'', \pi(\gamma'')) \geq U_A(\gamma'', \pi(\widehat{\gamma})).$$

Combining this equation with (34) yields

$$U_A(\gamma', \pi(\widehat{\gamma})) - U_A(\gamma', \pi(\gamma'')) > U_A(\gamma'', \pi(\widehat{\gamma})) - U_A(\gamma'', \pi(\gamma'')). \tag{35}$$

However, given $\gamma' < \gamma''$ and $\pi(\widehat{\gamma}) > \pi(\gamma'')$, (35) violates the single-crossing condition in U_A. Contradiction.

Case 4. Suppose $a(\gamma') = 0$ and $a(\gamma'') = 1$. By Lemma 4, $\pi(\gamma'') \leq \pi_A(\gamma'')$, and hence given $\pi(\gamma') > \pi(\gamma'')$, incentive compatibility for type γ'' requires $\pi(\gamma') > \pi_A(\gamma'')$. Consider a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma\in\Gamma}$ which sets $\widetilde{a}(\gamma') = 1$ and $\widetilde{\pi}(\gamma') = \pi(\gamma'')$ while leaving the allocation for types $\gamma \neq \gamma'$ unchanged. Since the original rule satisfies incentive compatibility for γ'', single-crossing implies that this perturbation is incentive compatible for γ'. Optimality of the original rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ then requires that this perturbation do not strictly increase the principal's welfare, which requires

$$U_P(\gamma', \pi(\gamma')) \geq U_P(\gamma', \pi(\gamma'')) - \phi.$$

The single-crossing condition in U_P then implies

$$U_P(\gamma'', \pi(\gamma')) > U_P(\gamma'', \pi(\gamma'')) - \phi. \tag{36}$$

Now consider a different perturbed rule $\{\widehat{a}(\gamma), \widehat{\pi}(\gamma)\}_{\gamma\in\Gamma}$ which sets $\widehat{a}(\gamma'') = 0$ and $\widehat{\pi}(\gamma'') = \pi(\gamma')$ while leaving the allocation for types $\gamma \neq \gamma''$ unchanged. Equation (36) implies that such a perturbation would strictly increase the principal's welfare. Hence, optimality of the original rule $\{a(\gamma), \pi(\gamma)\}_{\gamma\in\Gamma}$ requires that this perturbation violate incentive compatibility, that is, there must exist $\widehat{\gamma} \in \Gamma$ with $a(\widehat{\gamma}) = 0$ such that

$$U_A(\gamma'', \pi(\widehat{\gamma})) > U_A(\gamma'', \pi(\gamma')). \tag{37}$$

Note that since $\pi(\gamma') > \pi_A(\gamma'')$, we must have $\pi(\widehat{\gamma}) < \pi(\gamma')$. Moreover, by incentive compatibility being satisfied under the original rule, we have

$$U_A(\gamma', \pi(\gamma')) \geq U_A(\gamma', \pi(\widehat{\gamma})).$$

Combining this equation with (37) yields

$$U_A(\gamma', \pi(\gamma')) - U_A(\gamma', \pi(\widehat{\gamma})) > U_A(\gamma'', \pi(\gamma')) - U_A(\gamma'', \pi(\widehat{\gamma})). \tag{38}$$

However, given $\gamma' < \gamma''$ and $\pi(\widehat{\gamma}) < \pi(\gamma')$, (38) violates the single-crossing condition in U_A. Contradiction.

B.4 Proof of Proposition 4

Consider the following *quadratic-uniform setting*: preferences satisfy $U_P(\gamma, \pi) = \gamma\pi - \pi^2/2$ and $U_A(\gamma, \pi) = (\gamma + \beta)\pi - \pi^2/2$ for $\beta > 0$, and $f(\gamma) = 1$ for all $\gamma \in \Gamma$. In this setting, the efficient and flexible spending levels are given by $\pi_P(\gamma) = \gamma$ and $\pi_A(\gamma) = \gamma + \beta$ respectively. Let $\alpha = 0$, so that the agent pays no verification cost.

We first establish that in this setting, if the verification cost satisfies $\phi > \beta^2/2$, TEC is suboptimal, as it is dominated by a rule without verification.

Lemma 8. *Consider the quadratic-uniform setting with* $\alpha = 0$*. If* $\phi > \beta^2/2$*, then TEC is not optimal.*

Proof. Take the quadratic-uniform setting with $\alpha = 0$ and $\phi > \beta^2/2$. Consider the following problem:

$$\max_{\{\gamma^*,\gamma^{**}\}} \left\{ \begin{array}{c} \int_{\underline{\gamma}}^{\gamma^*} U_P(\gamma, \pi_A(\gamma)) f(\gamma) d\gamma + \int_{\gamma^*}^{\gamma^{**}} U_P(\gamma, \pi_A(\gamma^*)) f(\gamma) d\gamma \\ + \int_{\gamma^{**}}^{\overline{\gamma}} (U_P(\gamma, \pi_P(\gamma)) - \phi) f(\gamma) d\gamma \end{array} \right\}. \tag{39}$$

Note that the solution to this program coincides with a rule without verification if it sets $\gamma^{**} = \overline{\gamma}$, and it coincides with a rule that verifies all types if it sets $\gamma^{**} = \underline{\gamma}$. By the definition of TEC, a necessary condition for a TEC rule to be optimal is that the solution to program (39) specify $\underline{\gamma} < \gamma^{**} < \overline{\gamma}$. We show that this cannot be satisfied when $\phi > \beta^2/2$.

The first-order condition for γ^*, given our assumptions on preferences and the distribution of γ, implies

$$\gamma^* = \max\left\{\frac{\underline{\gamma} + \gamma^{**}}{2} - \beta, \gamma^{**} - 2\beta\right\}, \tag{40}$$

where we have taken into account the fact that γ^* may be lower than $\underline{\gamma}$. If the solution to (39) sets γ^{**} strictly interior, then the first-order condition for γ^{**} implies

$$-\gamma^{**}(\gamma^* + \beta) + \frac{(\gamma^* + \beta)^2}{2} + \frac{\gamma^{**2}}{2} = \phi.$$

Substituting with (40) and rearranging terms yields

$$\frac{\left(\gamma^{**} - \max\left\{\frac{\underline{\gamma}+\gamma^{**}}{2} - \beta, \gamma^{**} - 2\beta\right\} - \beta\right)^2}{2} = \phi. \tag{41}$$

Note that if $\gamma^* \geq \underline{\gamma}$, (41) implies $\phi = \beta^2/2$, contradicting the assumption that $\phi > \beta^2/2$. Therefore,

$$\gamma^* < \underline{\gamma}, \tag{42}$$

and thus (41) implies

$$\gamma^{**} = \underline{\gamma} + 2\sqrt{2\phi}.$$

Substituting back into (40), we obtain

$$\gamma^* = \underline{\gamma} + \sqrt{2\phi} - \beta. \tag{43}$$

However, combined with (42), equation (43) implies $\phi < \beta^2/2$, contradicting the assumption that $\phi > \beta^2/2$. Therefore, the solution to (39) cannot set γ^{**} strictly interior when $\phi > \beta^2/2$. $\square$

We next show that there exists $\phi > \beta^2/2$ under which a rule with verification is optimal.

Lemma 9. *Consider the quadratic-uniform setting with $\alpha = 0$. If $\beta^2/2 < \phi < 2\beta^2/3$ and $6\beta < \overline{\gamma} - \underline{\gamma}$, then a rule with verification is optimal.*

Proof. Take the quadratic-uniform setting with $\alpha = 0$, $\beta^2/2 < \phi < 2\beta^2/3$, and $6\beta < \overline{\gamma} - \underline{\gamma}$. An optimal rule without verification sets $\pi(\gamma) = \min\{\pi_A(\gamma), \pi_A(\gamma^*)\}$, where using (40) (with $\gamma^{**} = \overline{\gamma}$) and the fact that $\overline{\gamma} - 2\beta > \underline{\gamma} + 4\beta > \underline{\gamma}$, we have

$$\gamma^* = \overline{\gamma} - 2\beta.$$

We construct a perturbed rule $\{\widetilde{a}(\gamma), \widetilde{\pi}(\gamma)\}_{\gamma\in\Gamma}$ that features verification and yields the principal strictly higher expected welfare than this optimal rule without verification. For any given $\gamma_H < \gamma^*$, define γ_L as the solution to

$$U_A(\gamma_L, \gamma_L - \beta) = U_A(\gamma_L, \pi_A(\gamma_H)),$$

which after some algebra yields

$$\gamma_L = \gamma_H - 2\beta. \tag{44}$$

Take $\gamma_H < \gamma^*$ sufficiently close to γ^* so that γ_L satisfies $\gamma_L - 2\beta > \underline{\gamma}$ (note that the assumption that $6\beta < \overline{\gamma} - \underline{\gamma}$ ensures that such a γ_H exists). Type γ_L is defined so that he is indifferent between the flexible spending level of γ_H and the optimal spending limit under no verification for a distribution truncated at γ_L (which is given by $\pi_A(\gamma_L - 2\beta) = \gamma_L - \beta$). Now construct the perturbed rule as follows: if $\gamma < \gamma_L - 2\beta$ or $\gamma > \gamma_H$, then $\widetilde{a}(\gamma) = 0$ and $\widetilde{\pi}(\gamma) = \pi(\gamma)$; if $\gamma \in [\gamma_L - 2\beta, \gamma_L)$, then $\widetilde{a}(\gamma) = 0$ and $\widetilde{\pi}(\gamma) = \gamma_L - \beta$; and if $\gamma \in [\gamma_L, \gamma_H]$, then $\widetilde{a}(\gamma) = 1$ and $\widetilde{\pi}(\gamma)$ satisfies

$$U_A(\gamma, \widetilde{\pi}(\gamma)) = U_A(\gamma, \pi_A(\gamma_H)),$$

which after some algebra yields

$$\widetilde{\pi}(\gamma) = 2\gamma - \gamma_H + \beta.$$

Note that given the definition of γ_L, this rule is incentive compatible. The perturbation only changes the principal's welfare for types $\gamma \in [\gamma_L - 2\beta, \gamma_H]$. The change in welfare is equal to

$$\int_{\gamma_{L-2\beta}}^{\gamma_L} \big(U_P(\gamma,\gamma_L-\beta)-U_P(\gamma,\gamma+\beta)\big) f(\gamma)\,d\gamma$$
$$+\int_{\gamma_L}^{\gamma_H} \big(U_P(\gamma,2\gamma-\gamma_H+\beta)-\phi-U_P(\gamma,\gamma+\beta)\big) f(\gamma)\,d\gamma.$$

After some algebra and substitution of (44), using our assumptions on preferences and the distribution of γ, this simplifies to

$$-\int_{\gamma_{H-4\beta}}^{\gamma_{H-2\beta}} \frac{(\gamma-\gamma_H+3\beta)^2}{2}d\gamma-\int_{\gamma_{H-2\beta}}^{\gamma_H} \frac{(\gamma_H-\gamma-\beta)^2}{2}d\gamma-\int_{\gamma_{H-2\beta}}^{\gamma_H} \phi d\gamma+\int_{\gamma_H-4\beta}^{\gamma_H} \frac{\beta^2}{2}d\gamma.$$

Simplifying further yields that the change in welfare is equal to

$$\frac{4}{3}\beta^3-2\beta\phi>0,$$

where the inequality follows from the assumption that $\phi<2\beta^2/3$. Therefore, the perturbed rule with verification strictly increases the principal's expected welfare relative to no verification. □

It follows from Lemma 8 and Lemma 9 that in a quadratic-uniform setting with $\alpha=0$, $\beta^2/2<\phi<2\beta^2/3$, and $6\beta<\overline{\gamma}-\underline{\gamma}$, verification is optimal but TEC is not. By Lemma 2, any optimal rule must therefore feature decreasing verification.

B.5 Proof of Lemma 7

Part (i). Suppose an incentive compatible rule induces decreasing verification at γ_H. Consider first the case in which $a(\gamma_H)=0$ and thus $a(\gamma_H-\varepsilon)=1$ for $\varepsilon>0$ arbitrarily small. Incentive compatibility for type γ_H requires

$$U_A(\gamma_H,\pi(\gamma_H))\geq U_A(\gamma_H,\pi_P(\gamma_H))-\alpha\phi, \tag{45}$$

since γ_H can choose to be verified and guarantee himself the efficient level of spending. Incentive compatibility for type $\gamma_H-\varepsilon$ requires

$$U_A(\gamma_H-\varepsilon,\pi_P(\gamma_H-\varepsilon))-\alpha\phi\geq U_A(\gamma_H-\varepsilon,\pi(\gamma_H)), \tag{46}$$

since $\gamma_H-\varepsilon$ can choose not to be verified and spend at $\pi(\gamma_H)$. Given the continuity of U_A and π_P in their respective arguments, we can take the limit of both sides of (46) as ε

approaches 0 to obtain

$$U_A(\gamma_H, \pi_P(\gamma_H)) - \alpha\phi \geq U_A(\gamma_H, \pi(\gamma_H)). \tag{47}$$

Combining (45) and (47) yields (12).

Consider next the case in which $a(\gamma_H) = 1$ and thus $a(\gamma_H + \varepsilon) = 0$ for $\varepsilon > 0$ arbitrarily small. Analogous arguments to those above imply the following incentive compatibility constraints for γ_H and $\gamma_H + \varepsilon$, respectively:

$$\begin{aligned} U_A(\gamma_H, \pi_P(\gamma_H)) - \alpha\phi &\geq U_A(\gamma_H, \pi(\gamma_H + \varepsilon)), & (48)\\ U_A(\gamma_H + \varepsilon, \pi(\gamma_H + \varepsilon)) &\geq U_A(\gamma_H + \varepsilon, \pi_P(\gamma_H + \varepsilon)) - \alpha\phi. & (49)\end{aligned}$$

Since the rule is piecewise continuous, $\lim_{\varepsilon\downarrow 0} \pi(\gamma_H + \varepsilon)$ exists and can be defined as $\pi(\gamma_H)$. Taking the limit of both sides of (48) and (49) as ε goes to 0 yields (47) and (45), and combining these two inequalities yields (12).

To complete the proof of part (i), we show that $\pi(\gamma_H) > \pi_A(\gamma_H)$ must hold. Note that by (12), either $\pi(\gamma_H) > \pi_A(\gamma_H)$ or $\pi(\gamma_H) \leq \pi_P(\gamma_H)$. For the purpose of contradiction, suppose it were the case that $\pi(\gamma_H) \leq \pi_P(\gamma_H)$. Consider the incentive compatibility constraint of type $\gamma_H - \varepsilon$ for $\varepsilon > 0$ arbitrarily small. Take first the case in which $a(\gamma_H - \varepsilon) = 1$. Then $\gamma_H - \varepsilon$ must weakly prefer verification to no verification, which requires

$$U_A(\gamma_H - \varepsilon, \pi_P(\gamma_H - \varepsilon)) - \alpha\phi \geq U_A(\gamma_H - \varepsilon, \pi(\gamma_H)). \tag{50}$$

Since $\pi_P(\gamma_H - \varepsilon) < \pi_P(\gamma_H) < \pi_A(\gamma_H - \varepsilon)$, (50) implies

$$U_A(\gamma_H - \varepsilon, \pi_P(\gamma_H)) - \alpha\phi > U_A(\gamma_H - \varepsilon, \pi(\gamma_H)). \tag{51}$$

Combining (12) and (51) yields

$$U_A(\gamma_H - \varepsilon, \pi_P(\gamma_H)) - U_A(\gamma_H - \varepsilon, \pi(\gamma_H)) > U_A(\gamma_H, \pi_P(\gamma_H)) - U_A(\gamma_H, \pi(\gamma_H)).$$

Given $\pi(\gamma_H) \leq \pi_P(\gamma_H)$, this inequality violates the single-crossing condition in U_A, thus yielding a contradiction.

Consider next the case in which $a(\gamma_H - \varepsilon) = 0$. Given decreasing verification at γ_H, in this case we must have $a(\gamma_H) = 1$ and $a(\gamma_H + \varepsilon) = 0$ for $\varepsilon > 0$ arbitrarily small. Moreover, given our definition of $\pi(\gamma_H)$, $\pi(\gamma_H) \leq \pi_P(\gamma_H)$ implies $\lim_{\varepsilon\downarrow 0} \pi(\gamma_H + \varepsilon) \leq \pi_P(\gamma_H)$. By incentive compatibility, type γ_H must weakly prefer verification to no verification, which requires

$$U_A(\gamma_H, \pi_P(\gamma_H)) - \alpha\phi \geq U_A(\gamma_H, \pi(\gamma_H + \varepsilon)), \tag{52}$$

whereas type $\gamma_H + \varepsilon$ must weakly prefer no verification to verification, which requires

$$U_A(\gamma_H + \varepsilon, \pi(\gamma_H + \varepsilon)) \geq U_A(\gamma_H + \varepsilon, \pi_P(\gamma_H + \varepsilon)) - \alpha\phi. \tag{53}$$

Combining (52) and (53) and using the fact that $\pi_A(\gamma_H) > \pi_P(\gamma_H + \varepsilon) > \pi_P(\gamma_H)$ yields

$$U_A(\gamma_H, \pi_P(\gamma_H + \varepsilon)) - U_A(\gamma_H, \pi(\gamma_H + \varepsilon)) > U_A(\gamma_H + \varepsilon, \pi_P(\gamma_H + \varepsilon)) - U_A(\gamma_H + \varepsilon, \pi(\gamma_H + \varepsilon)).$$

Since $\pi(\gamma_H + \varepsilon) \leq \pi_P(\gamma_H) \leq \pi_P(\gamma_H + \varepsilon)$ for ε approaching 0, this inequality violates the single-crossing condition in U_A, thus yielding again a contradiction.

Therefore, we obtain that $\pi(\gamma_H) \leq \pi_P(\gamma_H)$ cannot hold and we must thus have $\pi(\gamma_H) > \pi_A(\gamma_H)$.

Part (ii). Suppose an incentive compatible rule induces increasing verification at γ_L. Then analogous arguments to those used to prove part (i) can be applied to show that (14) must hold at γ_L. Since the steps are analogous, we omit the details.

B.6 Proof of Proposition 5

To prove this result, we first establish the following lemmas.

Lemma 10. *Under limited commitment, if an incentive compatible rule features increasing verification at γ_L, then*

$$\pi(\gamma_L) \leq \pi_P(\gamma_L), \tag{54}$$

where $\pi(\gamma_L) \equiv \lim_{\varepsilon \downarrow 0} \pi(\gamma_L - \varepsilon)$ if $a(\gamma_L) = 1$.

Proof. Suppose an incentive compatible rule features increasing verification at γ_L. By equation (14) in Lemma 7, either $\pi(\gamma_L) > \pi_A(\gamma_L)$ or $\pi(\gamma_L) \leq \pi_P(\gamma_L)$. For the purpose of contradiction, suppose $\pi(\gamma_L) > \pi_A(\gamma_L)$ holds. Take first the case in which $a(\gamma_L) = 1$, so that $a(\gamma_L - \varepsilon) = 0$ for $\varepsilon > 0$ arbitrarily small and, given our definition of $\pi(\gamma_L)$, $\lim_{\varepsilon \downarrow 0} \pi(\gamma_L - \varepsilon) > \pi_A(\gamma_L)$. By incentive compatibility, type $\gamma_L - \varepsilon$ must weakly prefer no verification to verification, which requires

$$U_A(\gamma_L - \varepsilon, \pi(\gamma_L - \varepsilon)) \geq U_A(\gamma_L - \varepsilon, \pi_P(\gamma_L - \varepsilon)) - \alpha\phi. \tag{55}$$

However, (14) and (55) together with the fact that $\pi(\gamma_L) > \pi_A(\gamma_L)$ imply that Assumption 2 is violated, yielding a contradiction.

Consider next the case in which $a(\gamma_L) = 0$, so that $a(\gamma_L + \varepsilon) = 1$ for $\varepsilon > 0$ arbitrarily small. By incentive compatibility, type $\gamma_L + \varepsilon$ must weakly prefer verification to no verification, which requires

$$U_A(\gamma_L + \varepsilon, \pi_P(\gamma_L + \varepsilon)) - \alpha\phi \geq U_A(\gamma_L + \varepsilon, \pi(\gamma_L)). \tag{56}$$

Note that in this case, $\pi(\gamma_L) > \pi_A(\gamma_L) > \pi_P(\gamma_L + \varepsilon)$ requires $\pi(\gamma_L) > \pi_A(\gamma_L + \varepsilon)$. However, (14) and (56) together with $\pi(\gamma_L) > \pi_A(\gamma_L + \varepsilon)$ imply that Assumption 2 is violated, yielding again a contradiction.

Therefore, we obtain that $\pi(\gamma_L) > \pi_A(\gamma_L)$ cannot hold and we must thus have $\pi(\gamma_L) \leq \pi_P(\gamma_L)$. $\square$

Lemma 11. *Under limited commitment, if an incentive compatible rule features decreasing verification at γ_H, then there exists $\gamma' \leq \gamma_H$ satisfying*

$$U_A(\gamma', \pi(\gamma')) = U_A(\gamma', \pi(\gamma_H)) \tag{57}$$

for $\pi(\gamma') < \pi_A(\gamma')$, $\pi(\gamma_H) \equiv \lim_{\varepsilon \downarrow 0} \pi(\gamma_H + \varepsilon)$ if $a(\gamma_H) = 1$, and either $a(\gamma') = 0$, or $a(\gamma') = 1$, $\lim_{\varepsilon \downarrow 0} a(\gamma' - \varepsilon) = 0$, and $\pi(\gamma') \equiv \lim_{\varepsilon \downarrow 0} \pi(\gamma' - \varepsilon)$.

Proof. Suppose an incentive compatible rule features decreasing verification at γ_H. By condition (13) in Lemma 7, $\pi(\gamma_H) > \pi_A(\gamma_H)$. Consider the problem of the principal *after* the verification decision $a(\gamma)$ has been made and the verification result (in case of verification) has been obtained:

$$\max_{\{\pi(\gamma)\}_{\gamma\in\Gamma}} \int_{\underline{\gamma}}^{\overline{\gamma}} U_P(\gamma, \pi(\gamma)) f(\gamma) d\gamma \tag{58}$$

subject to

$$\pi(\gamma) = \pi_P(\gamma) \text{ if } a(\gamma) = 1, \tag{59}$$

$$U_A(\gamma, \pi(\gamma)) \geq U_A(\gamma, \pi(\widehat{\gamma})) \text{ for all } \gamma, \widehat{\gamma} \text{ for which } a(\gamma) = a(\widehat{\gamma}) = 0. \tag{60}$$

This program takes into account that the principal will assign the efficient spending level to any agent type who chooses to be verified, and she will ignore the incentives of verified types when deciding the spending allocation of types who choose not to be verified. We now consider the optimal level of $\pi(\gamma_H)$ given decreasing verification at γ_H and the conditions that are necessary for the principal to choose $\pi(\gamma_H) > \pi_A(\gamma_H)$.

Step 1. Consider the spending allocation conditional on no verification. Note that analogous arguments to those used in the proof of Lemma 5 imply that $\pi(\gamma)$ must be weakly increasing for non-verified types γ. For each non-verified type γ, denote by $\underline{\pi}(\gamma)$ the spending level closest to $\pi_A(\gamma)$ from below in the allowable spending set for non-verified types (i.e., among all spending levels assigned to types who choose no verification). Analogously, denote by $\overline{\pi}(\gamma)$ the closest spending level to $\pi_A(\gamma)$ from above in the allowable spending set for non-verified types. Clearly, if $\pi_A(\gamma)$ is in this allowable spending set, then $\pi_A(\gamma) = \underline{\pi}(\gamma) = \overline{\pi}(\gamma)$. The incentive compatibility constraint (60) together with the concavity of U_A require that if $a(\gamma) = 0$, then

$$\pi(\gamma) = \arg \max_{\pi \in \{\underline{\pi}(\gamma), \overline{\pi}(\gamma)\}} U_A(\gamma, \pi). \tag{61}$$

Step 2. As noted, given decreasing verification at γ_H, the rule must set $\pi(\gamma_H) > \pi_A(\gamma_H)$. We show that as a result, the rule must induce $a(\gamma) = 0$ and $\pi(\gamma) = \pi(\gamma_H)$ for all types $\gamma \in \left(\gamma_H, \pi_A^{-1}(\pi(\gamma_H))\right)$. To see why, note first that by (61) and the single-crossing condition in U_A, any type $\gamma \in \left(\gamma_H, \pi_A^{-1}(\pi(\gamma_H))\right)$ who is not verified necessarily chooses spending $\pi(\gamma) = \pi(\gamma_H)$. Therefore, it is sufficient to show that any type $\gamma \in \left(\gamma_H, \pi_A^{-1}(\pi(\gamma_H))\right)$ must have $a(\gamma) = 0$. Suppose by contradiction that this were not the case. Then incentive compatibility for a type $\gamma \in \left(\gamma_H, \pi_A^{-1}(\pi(\gamma_H))\right)$ with $a(\gamma) = 1$ requires that this type weakly prefer verification to no verification, which requires

$$U_A(\gamma, \pi_P(\gamma)) - \alpha\phi \geq U_A(\gamma, \pi(\gamma_H)). \tag{62}$$

However, (12) and (62) together with the fact that $\gamma > \gamma_H$ and $\pi(\gamma_H) > \pi_A(\gamma)$ violate Assumption 2. The claim therefore follows.

Step 3. We show that in an incentive compatible rule, constraint (60) cannot be uniformly slack for all $\gamma \leq \gamma_H$ and $\widehat{\gamma} = \gamma_H$, where recall $\pi(\gamma_H) > \pi_A(\gamma_H)$ by decreasing verification at γ_H. Suppose by contradiction that this is true. Note that from Step 2, (60) is then uniformly slack for all $\gamma \leq \gamma_H$ and $\widehat{\gamma} \in \left(\gamma_H, \pi_A^{-1}(\pi(\gamma_H))\right)$, where $a(\widehat{\gamma}) = 0$ for all such $\widehat{\gamma}$. Now consider the following perturbation $\{\widetilde{\pi}(\gamma)\}_{\gamma\in\Gamma}$: for $\varepsilon > 0$ arbitrarily small and all $\gamma \in \left(\gamma_H, \pi_A^{-1}(\pi(\gamma_H) - \varepsilon)\right]$, set $\widetilde{\pi}(\gamma) = \pi(\gamma_H) - \varepsilon$; for all $\gamma \in \left(\pi_A^{-1}(\pi(\gamma_H) - \varepsilon), \pi_A^{-1}(\pi(\gamma_H))\right)$, set $\widetilde{\pi}(\gamma) = \pi_A(\gamma)$; and for all other types leave the spending allocation unchanged. This perturbation strictly increases the principal's welfare as it reduces overspending by types $\gamma \in \left(\gamma_H, \pi_A^{-1}(\pi(\gamma_H))\right)$. Moreover, since (by the contradiction assumption) (60) was uniformly slack before the perturbation for all $\gamma \leq \gamma_H$, it is still satisfied after the perturbation, and incentive compatibility for all types $\gamma \geq \gamma_H$ is guaranteed as the perturbation satisfies (61).

Therefore, we obtain that if (60) is uniformly slack for all $\gamma \leq \gamma_H$ and $\widehat{\gamma} = \gamma_H$, the principal can strictly improve upon the original rule by reducing $\pi(\gamma_H)$ after the verification decision has been made, and hence the original rule violates incentive compatibility for the principal. The claim follows.

Step 4. By Step 3, in any incentive compatible rule with decreasing verification at γ_H, there exists $\gamma' \leq \gamma_H$ satisfying (57). Moreover, since decreasing verification at γ_H implies $\pi(\gamma_H) > \pi_A(\gamma_H) \geq \pi_A(\gamma')$, this requires $\pi(\gamma') < \pi_A(\gamma')$. This proves the lemma. □

Lemma 12. *Under limited commitment, if an incentive compatible rule features decreasing verification at γ_H, then there exists $\gamma_L \leq \gamma_H$ at which the rule features increasing verification. Moreover, $a(\gamma) = 1$ for all $\gamma \in (\gamma_L, \gamma_H)$ and*

$$U_A(\gamma_L, \pi(\gamma_L)) = U_A(\gamma_L, \pi(\gamma_H)) \tag{63}$$

for $\pi(\gamma_L) \equiv \lim_{\varepsilon \downarrow 0} \pi(\gamma_L - \varepsilon)$ if $a(\gamma_L) = 1$ and $\pi(\gamma_H) \equiv \lim_{\varepsilon \downarrow 0} \pi(\gamma_H + \varepsilon)$ if $a(\gamma_H) = 1$.

Proof. Suppose an incentive compatible rule features decreasing verification at γ_H. By Lemma 11, there exists a type $\gamma' \leq \gamma_H$ satisfying (57) either with $a(\gamma') = 0$ or at which there is increasing verification. We can establish that such a type is unique. Suppose by contradiction that there are two types, $\gamma'' \leq \gamma_H$ and $\gamma' < \gamma''$, satisfying the condition in Lemma 11. Then

$$U_A(\gamma'', \pi(\gamma'')) = U_A(\gamma'', \pi(\gamma_H)), \tag{64}$$

$$U_A(\gamma', \pi(\gamma')) = U_A(\gamma', \pi(\gamma_H)). \tag{65}$$

Incentive compatibility requires

$$U_A(\gamma'', \pi(\gamma'')) \geq U_A(\gamma'', \pi(\gamma')), \tag{66}$$

$$U_A(\gamma', \pi(\gamma')) \geq U_A(\gamma', \pi(\gamma'')). \tag{67}$$

Combining (64)-(67) yields

$$U_A(\gamma', \pi(\gamma_H)) - U_A(\gamma', \pi(\gamma'')) \geq U_A(\gamma'', \pi(\gamma_H)) - U_A(\gamma'', \pi(\gamma'')).$$

Since $\gamma' < \gamma''$ and $\pi(\gamma_H) > \pi_A(\gamma_H) \geq \pi_A(\gamma'') > \pi(\gamma'')$ by decreasing verification at γ_H and Lemma 11, this inequality violates the single-crossing condition in U_A, yielding a contradiction. Therefore, there exists a unique type below γ_H for which (57) holds, and denoting this type by γ_L yields (63).

Next, we show that $a(\gamma) = 1$ for all $\gamma \in (\gamma_L, \gamma_H)$. Note first that a spending level $\pi \in (\pi(\gamma_L), \pi(\gamma_H))$ cannot be allowed by the rule under no verification, since otherwise type γ_L would have a strict incentive to deviate to such a spending level. Consider the relevant case in which $\gamma_L < \gamma_H$ and suppose by contradiction that $a(\gamma) = 0$ for some type $\gamma \in (\gamma_L, \gamma_H)$. Let γ' denote the highest such type γ. Since, as noted, spending levels strictly between $\pi(\gamma_L) < \pi_A(\gamma')$ and $\pi(\gamma_H) > \pi_A(\gamma')$ are not allowed, it follows from (63) and $\gamma' > \gamma_L$ that the rule must set $\pi(\gamma') = \pi(\gamma_H)$. Moreover, since by construction the rule features increasing verification at γ', condition (14) in Lemma 7 implies

$$U_A(\gamma', \pi_P(\gamma')) - \alpha\phi = U_A(\gamma', \pi(\gamma')) = U_A(\gamma', \pi(\gamma_H)). \tag{68}$$

However, given (12) and (13), equation (68) violates Assumption 2. It follows that $a(\gamma) = 1$ for all $\gamma \in (\gamma_L, \gamma_H)$. □

We can now prove the proposition. We begin by ruling out decreasing verification. Suppose by contradiction that an incentive compatible rule features decreasing verification at some $\gamma_H \in \Gamma$. By Lemma 12, there must exist a type $\gamma_L \leq \gamma_H$ satisfying the conditions in the lemma. We proceed in two steps.

Step 1. Suppose $\gamma_L < \gamma_H$. Then it follows from (14) and (63) that

$$U_A(\gamma_L, \pi_P(\gamma_L)) - \alpha\phi = U_A(\gamma_L, \pi(\gamma_H)). \tag{69}$$

However, (12) and (69) together with the fact that $\gamma_L < \gamma_H$ and $\pi(\gamma_H) > \pi_A(\gamma_H)$ (by (13)) imply that Assumption 2 is violated. Contradiction.

Step 2. By Step 1, any incentive compatible rule with decreasing verification must have $\gamma_L = \gamma_H$ at each point γ_H at which there is decreasing verification. Now consider the principal's problem (58)-(60). Let $\overline{\gamma}' \leq \overline{\gamma}$ be the highest non-verified type. Since the types with decreasing verification are atomistic and the rule is piecewise continuous, following a decision of no verification the principal solves

$$\max_{\{\pi(\gamma)\}_{\gamma\in\Gamma}} \int_{\underline{\gamma}}^{\overline{\gamma}'} U_P\left(\gamma, \pi\left(\gamma\right)\right) f\left(\gamma\right) d\gamma$$
$$\text{subject to}$$
$$U_A\left(\gamma, \pi\left(\gamma\right)\right) \geq U_A\left(\gamma, \pi\left(\widehat{\gamma}\right)\right) \text{ for all } \gamma, \widehat{\gamma} \text{ for which } a\left(\gamma\right) = a\left(\widehat{\gamma}\right) = 0.$$

By Proposition 1, the solution assigns $\pi\left(\gamma\right) = \min\left\{\pi_A\left(\gamma\right), \pi_A\left(\gamma^*\right)\right\}$ for $\gamma \in [\underline{\gamma}, \overline{\gamma}']$ and some $\gamma^* < \overline{\gamma}'$. However, in this case, conditions (13) and (54) (which require $\pi\left(\gamma_H\right) > \pi_A\left(\gamma_H\right)$ and $\pi\left(\gamma_L\right) \leq \pi_P\left(\gamma_L\right)$ respectively) cannot be satisfied at a point $\gamma_H \in [\underline{\gamma}, \overline{\gamma}']$ at which there is decreasing verification and thus $\gamma_L = \gamma_H$. Contradiction.

The claims above show that under limited commitment, any incentive compatible rule features weakly increasing verification everywhere. Analogous arguments to those in the proofs of Lemma 1 and Lemma 2 can then be applied to show that a TEC rule is optimal if a rule with verification that is weakly increasing everywhere is optimal. Therefore, under limited commitment, if verification is optimal, TEC is optimal.

C Other Mechanisms

In this Appendix, we show that our results under full commitment are not limited to the game form in Section 2.2 but apply more generally when allowing for any indirect mechanism specifying a message space for the agent and a deterministic allocation function to which the principal commits. Specifically, we prove that a Revelation Principle in terms of payoffs holds in our setting, implying that to study the optimal deterministic mechanism for the principal, it is without loss to restrict attention to deterministic direct mechanisms in which the agent reports his type truthfully, as in program (3)-(4).

The usual version of the Revelation Principle cannot be directly applied to our problem because we consider a game with verification and limit attention to deterministic allocations. We note that Townsend (1988) provides an extension of the Revelation Principle to a class of models with verification and Strausz (2003) provides an extension to a setting with deterministic mechanisms and one agent. Our results below build on this work, particularly the latter.

Consider a general problem in which the principal must select a deterministic allocation specifying whether the agent is verified or not, $a \in A \equiv \{0, 1\}$, and the spending level that the agent is assigned, $\pi \in [\underline{\pi}, \overline{\pi}]$. A mechanism (S, a, π) consists of a message space for the agent S, a verification function $a : S \rightarrow A$ that commits the principal to implement the verification assignment $a(s)$ when the agent sends message s, and a spending function $\pi : S \times R \to \pi \in [\underline{\pi}, \overline{\pi}]$ that commits the principal to implement the spending level $\pi(s,r)$ when the agent sends message s and the verification result is r. Without loss given our assumption that verification reveals the agent's type perfectly, we let $r = a(s)\gamma$, that is, the verification result is equal to the agent's type γ if the agent is verified and it is equal to 0 if the agent is not verified (recall $\underline{\gamma} > 0$).

Given a mechanism (S, a, π), the agent chooses a message. The agent's reporting strategy $\mu(\gamma) : \Gamma \to S$ selects the message s with probability $\mu(s|\gamma)$. A mechanism is a direct mechanism if $S = \Gamma$.

Proposition 6. *Consider an equilibrium of a game induced by a deterministic indirect mechanism (S, a, π). There exists a deterministic direct mechanism, (Γ, a, π), that induces an equilibrium with truthful revelation yielding the principal a weakly larger expected welfare than that in the equilibrium under the indirect mechanism. Hence, an optimal deterministic mechanism for the principal solves program (3)-(4).*

Proof. Consider a deterministic mechanism (S, a, π) and equilibrium reporting strategies $\mu(\gamma)$ for each type $\gamma \in \Gamma$. For each γ, let S_γ be the set of messages that the agent sends with positive probability, i.e. $S_\gamma = \{s|\mu(s|\gamma) > 0\}$. Since $\mu(\gamma)$ is an equilibrium strategy, it satisfies

$$U_A(\gamma, \pi(s, a(s)\gamma)) - a(s)\alpha\phi \geq U_A(\gamma, \pi(\widehat{s}, a(\widehat{s})\gamma)) - a(\widehat{s})\alpha\phi \text{ for all } s \in S_\gamma, \widehat{s} \in S, \quad (70)$$

$$U_A(\gamma, \pi(s, a(s)\gamma)) - a(s)\alpha\phi = U_A(\gamma, \pi(\widehat{s}, a(\widehat{s})\gamma)) - a(\widehat{s})\alpha\phi \text{ for all } s, \widehat{s} \in S_\gamma. \quad (71)$$

Define the set $S_{\gamma,P}$ as the set of messages that, given the allocation specified by the mechanism, yield the principal the highest welfare from type γ among the messages that are sent with positive probability under this type's reporting strategy. That is, $S_{\gamma,P} = \{s \in S_\gamma | U_P(\gamma, \pi(s, a(s)\gamma)) - a(s)\phi = \max_{z \in S_\gamma}\{U_P(\gamma, \pi(z, a(z)\gamma)) - a(z)\phi\}\}$. Then construct a direct mechanism (Γ, a, π) specifying: for a given $s \in S_{\gamma,P}$ (arbitrarily chosen if $|S_{\gamma,P}| > 1$), $a(\gamma) = a(s)$, $\pi(\gamma, r) = \pi(s, a(s)\gamma)$ if $r \in \{0, \gamma\}$, and $\pi(\gamma, r) = \pi(s, \widehat{\gamma})$ if $r = \widehat{\gamma} \in \Gamma$, $\widehat{\gamma} \neq \gamma$. By construction, given this direct mechanism, it is an optimal strategy for each type γ to report his type truthfully. This equilibrium yields each type γ the same welfare as the equilibrium under the original indirect mechanism and it yields the principal weakly larger expected welfare than that equilibrium. The latter follows from the fact that the principal receives weakly larger welfare conditional on any type γ in the equilibrium of the direct mechanism with truthful revelation. The principal's expected welfare is the same in the two equilibria if $S_\gamma = S_{\gamma,P}$ for all γ and is strictly larger in the equilibrium of the direct mechanism with truthful revelation if $S_\gamma \neq S_{\gamma,P}$ for some γ. This proves the first part of the proposition.

We now prove the second part, namely that an optimal deterministic mechanism solves program (3)-(4). By the result just established, we can restrict attention to direct mechanisms in which the agent reports his type truthfully. With some abuse of notation, let $\pi(\gamma) \equiv \pi(\gamma, a(\gamma)\gamma)$. The principal's problem is

$$\max_{\{a(\gamma),\pi(\gamma),\pi(\widehat{\gamma},\gamma)\}_{\gamma,\widehat{\gamma}\in\Gamma}} \int_{\underline{\gamma}}^{\overline{\gamma}} \big(U_P\left(\gamma, \pi\left(\gamma\right)\right) - a\left(\gamma\right)\phi\big) f\left(\gamma\right) d\gamma \tag{72}$$

subject to

$$U_A\left(\gamma, \pi\left(\gamma\right)\right) - a\left(\gamma\right)\alpha\phi \geq U_A\left(\gamma, \pi\left(\widehat{\gamma}\right)\right) \text{ for all } \gamma, \widehat{\gamma} \text{ for which } a\left(\widehat{\gamma}\right) = 0, \tag{73}$$

$$U_A\left(\gamma, \pi\left(\gamma\right)\right) - a\left(\gamma\right)\alpha\phi \geq U_A\left(\gamma, \pi\left(\widehat{\gamma}, \gamma\right)\right) - \alpha\phi \text{ for all } \gamma, \widehat{\gamma} \text{ for which } a\left(\widehat{\gamma}\right) = 1. \tag{74}$$

The only difference between this program and program (3)-(4) is that the incentive compatibility constraint (74) is absent in (3)-(4). However, the principal can trivially prevent a deviation of a type γ in which he mimics a type $\widehat{\gamma}$ with $a(\widehat{\gamma}) = 1$: as this type is verified following the deviation, the principal can verify that he has deviated and punish him by assigning a spending level $\pi\left(\widehat{\gamma}, \gamma\right)$ such that $U_A(\gamma, \pi\left(\widehat{\gamma}, \gamma\right)) \leq U_A(\gamma, \pi(\gamma))$ for $\gamma \neq \widehat{\gamma}$ (where it is clear that such a spending level $\pi\left(\widehat{\gamma}, \gamma\right)$ exists, and in fact $\pi\left(\widehat{\gamma}, \gamma\right) = \pi(\gamma)$ would be a sufficient punishment). Therefore, the principal can satisfy (74) at no cost, and hence the solution to this program coincides with the solution to (72)-(73), which is equivalent to program (3)-(4). □