Optimal Contracts for Experimentation

Marina Halac[*] Navin Kartik[†] Qingmin Liu[‡]

May 14, 2012

Preliminary and Incomplete[§]

Abstract

This paper studies dynamic contracts for experimentation in a principal-agent setting with adverse selection (pre-contractual hidden information), dynamic moral hazard, and learning. Under full commitment, we show that an optimal (i.e. profit-maximizing) menu of contracts generally induces a low-ability agent to terminate experimentation inefficiently early, whereas there is no distortion in the stopping time for a high-ability agent. The structure of optimal contracts is influenced by a variety of considerations such as dynamic agency costs and post-contractual hidden information. We derive two explicit menus of contracts that can be used to implement the optimal solution: "bonus contracts" and "clawback contracts". Both feature history-contingent dynamic streams of transfers.

---

[*]Graduate School of Business, Columbia University. Email: `mhalac@columbia.edu`.

[†]Department of Economics, Columbia University. Email: `nkartik@columbia.edu`.

[‡]Department of Economics, Columbia University. Email: `ql2177@columbia.edu`.

[§]Prepared for NBER Organizational Economics Meeting. Comments are especially welcome, but please do not circulate more broadly without permission. Kartik gratefully acknowledges the hospitality of and funding from the University of Chicago Booth School of Business.

1 Introduction

This paper is concerned with the following contracting problem: a principal owns a project whose quality — viability, profitability, or difficulty — is unknown. Lacking the skills to directly investigate the project's quality, the principal must rely on an agent to *experiment*, i.e. to exert (costly) effort over time to learn more about the project. Through time, beliefs about the project's quality evolve as a function of not only the agent's effort but also the agent's (persistent) skill. The expected benefit of effort at any time depends on these beliefs: if, at some point, beliefs about the project's quality are sufficiently pessimistic relative to the agent's skill, it would be optimal to abandon the project altogether. Both the agent's skill and his effort choice in any period are unobservable to the principal, inducing both moral hazard and adverse selection.

These features are relevant in many contractual environments. Perhaps the most obvious application is the design of incentives within or across organizations for research and development (R&D) projects. A related application is the testing of a breakthrough product, e.g. investigating potential side effects of a new drug. But there are other quite distinct applications: for example, a firm or university may hire a recruiting agency to search for an external candidate for its CEO or president position. The agency's quality, the agency's effort, and uncertain market conditions determine when it is optimal to stop the search and just go with the best internal candidate.

Although dynamic moral hazard, adverse selection, and learning are essential features of these agency relationships, there is virtually no existing theoretical work on contracting in such environments. It bears emphasis that "learning" here refers not to the principal updating about the agent's type over time, but rather both parties updating about a persistent state variable — the project's quality — that affects the social value of effort. The socially optimal effort profile is non-stationary across time, and private benefits are determined by a conjunction of hidden information, hidden action, and evolving beliefs about the state. Consequently, these are rich environments to contract in. How well can a principal incentivize the agent? What is the nature of distortions, if any, that arise? What are the qualitative properties of optimal incentive contracts?

Our main contribution is to provide answers to these questions in a simple yet canonical model of experimentation, the so-called "exponential bandit" model, which we now briefly summarize.

Modeling framework. The project at hand may either be good or bad (a persistent state). Time is discrete and of infinite horizon. In each period, the agent can either exert effort (work) or not (shirk), a binary choice that is unobservable to the principal. The agent incurs a constant private cost in each period that he exerts effort. If the agent works in a period and the project is good, the project is successful in that period with some constant probability; if either the agent shirks or the project is bad, success cannot obtain in that period. Project success is publicly observable and obviates the need for any further effort.[1] The probability of success in a period (conditional on the agent working and the project being good) depends on the agent's persistent skill or ability, which, as usual, we refer to as his *type*. This is a binary variable — ability is either high or low — that is the agent's private information at the time of contracting. Project success yields a fixed social surplus that is directly accrued by the principal. Both parties are risk neutral and discount the future at a common rate.

Social optimum. Consider the first-best solution, i.e. when the agent's ability and effort are observable. Beliefs about the project being good decline deterministically so long as effort has been exerted but success not obtained. Since effort is costly, the social optimum is characterized by a stopping time for each agent type: as a function of his ability, the agent keeps working (so long as he has failed in the past, i.e. success has not been obtained) up until some point at which the project is permanently abandoned. It turns out that the optimal stopping time is a non-monotonic function of the agent's ability. The intuition stems from two countervailing forces: on the one hand, for any given belief about the project's quality, a higher-ability agent has a greater marginal benefit of effort (since conditional on the project being good, he succeeds with a higher probability); but at any point in time, a higher-ability agent is also more pessimistic about the project's quality (conditional on having exerted effort at all prior periods) because the informativeness of past failures about project quality is increasing in the agent's ability.

Agency issues. The contracting problem entails hidden information at the time of contracting (the agent's ability), dynamic hidden action (effort is costly for the agent and unobservable to the principal), and learning (both parties update over time about project quality and hence the benefits of effort, as function of their beliefs about past effort and ability). The principal's goal is to maximize profits, and to do so, she can commit ex-ante to a dynamic contract that specifies a sequence of transfers to the agent as a function of time and the publicly observable history, viz. project success/failure.[2] More precisely, since there is hidden information at the time of contracting, the principal may offer the agent a *menu* of such dynamic contracts from which the agent can choose one.

Let us highlight some of the complexities that arise in determining the optimal menu

[1]As we show, our analysis carries through without change if project success is only observed by the agent but can be verifiably disclosed.

[2]Thus, we consider a setting with full commitment. Note that we do not impose limited liability, so transfers can be to or from the agent, but there is an ex-ante individual rationality constraint.

of contracts. First, the agent does not commit to an effort profile but rather chooses effort in a sequentially optimal fashion. This implies that after accepting a particular contract, the agent will generally choose different effort profiles depending on his ability, a key distinguishing feature from a standard static adverse selection problem. Second, since the agent is getting more pessimistic about the likelihood of project success over time (so long he has worked in prior periods), the static incentive constraint for effort becomes more demanding over time. While this suggests that it may be optimal to provide increasing rewards for success over time, there is also a dynamic incentive constraint that the agent should not prefer to postpone current effort to the future in order to benefit from a higher future reward for success. In other words, there can be a *dynamic agency cost*: the presence of a future reward for success makes it less costly for the agent to forego a present reward for success and hence harder to prevent the agent form shirking in the present. Third, in addition to the pre-contractual hidden information about the agent's ability, there is also the possibility of post-contractual hidden information regarding beliefs about project quality.[3] In particular, the principal's and the agent's beliefs about project quality would diverge whenever the agent deviates by either choosing a different contract from the menu than he is intended to and/or by shirking when the principal expects him to work.[4] All these elements come into play when determining how to minimize the "information rent" that a menu of contracts provides the agent.

Results. In an optimal menu of contracts, the principal screens the agent types by offering two distinct contracts that satisfy the relevant self-selection or incentive compatibility constraints. Each type's contract induces him to work for a sequence of consecutive periods (so long as success has not been obtained) until when he abandons the project by permanently shirking. Compared to the social optimum, there is a distortion in the stopping time: while the high-ability type's stopping time is efficient, the low-ability type stops experimentation too early. Although this resembles the familiar "no distortion at the top but distortion below" in static models of adverse selection, such an analogy is rather incomplete because of the varied considerations noted above; we elaborate further later in the paper. Note that this implies that it is never optimal to simply "sell the project to the agent", because doing so would induce the socially optimal stopping times by all agent types.

We show how to implement the optimal solution in two different and economically interesting ways: a menu of *bonus contracts* and a menu of *clawback contracts*.[5] In a bonus contract,

[3]Post-contractual hidden information is sometimes referred to as "hidden knowledge" or "moral hazard with hidden information" (e.g. Hart and Holmstrom, 1987).

[4]Of course, in an optimal menu of contracts, such deviations will not occur; however, these off-path considerations matter when determining what the optimal menu is.

[5]To interpret the description below, note that we normalize the agent's reservation utility to zero (independent

the agent pays the principal an up-front fee and is then rewarded with a time-varying bonus for project success. The sequence of bonuses (of which at most one is ever paid) is increasing over time up until some point when it drops to zero; this induces the desired stopping time. We characterize the exact sequence of bonuses that are optimal for the principal. One way of interpreting bonus contracts is that the principal sells the project to the agent at the outset for a specified price, but commits to buy back the project's output (which obtains with a success) at time-dated future prices, so long as output is obtained by a certain date. Note that in this interpretation, the project's output is only valuable to the principal, not otherwise to the agent.

By contrast, in a clawback contract, the principal pays the agent an up-front amount, which can be viewed as a pre-payment for future success. Then, in each period in which a success does not obtain (and has not already obtained in the past), the agent is required to pay the principal some particular amount, up until some point when these payments stop; this induces the desired stopping time. We characterize the exact sequence of payments from the agent to the principal that must be used in such optimal contracts: the payment sequence increases over time with a jump at the termination date. We call this a "clawback contract" based on the idea that clawbacks in practice involve recouping a payment already made to the agent (sometimes with added penalties) when there is some, perhaps inconclusive, evidence of the agent's negligence. In the present context, the evidence is the lack of project success; it is important to note, however, that, in equilibrium, the agent does not shirk before the induced stopping time.

Related Literature. This paper fits into the literature on dynamic moral hazard, to which there have been a number of contributions in recent years, particularly with the proliferation of continuous-time methods. Most papers in this literature ignore adverse selection, but exceptions include Sannikov (2007), Fong (2009), and Gershkov and Perry (2012). The environments considered by these authors do not involve learning or experimentation, which is our focus.

While we are not aware of any other paper that studies contracting with dynamic moral hazard, adverse selection, and experimentation, it is useful to relate our work to other research on contracting for experimentation.[6] Manso (2011) studies a setting in which a principal must not only incentivize an agent to work rather than shirk, but also to work on experimentation rather than exploitation in the usual sense of "two-arm bandit" models. The latter concern — which we do not have — is central to his main insights. He studies a two-period model and does not have adverse selection (instead, he has limited liability); hence, the focus is quite different.

of his type).

[6]Beyond contracting settings, a number of authors have studied games where players experiment with exponential "bandits", such as Keller et al. (2005), Strulovici (2010), Bonatti and Horner (2011), and Garfagnini (2011).

Somewhat closer to our setting is Hörner and Samuleson (2012), who identify the dynamic agency cost mentioned earlier (see also Bhaskar, 2009). Again, they do not have adverse selection, which is an essential component of our analysis. They examine rather different aspects of agency by assuming that the agent cannot exert effort in any period without a fixed amount of funding in that period from the principal. Thus, they are interested in applications such as venture capital financing (Bergemann and Hege, 1998, 2005), whereas we are interested in applications where the principal initially owns the project or the agent cannot directly accrue the full social benefits from project success. In this respect, we are closer to Besanko et al. (2012), but their framework does not have moral hazard and instead focuses on issues of ambiguity. Finally, another related paper is Gerardi and Maestri (2011). They analyze how an agent can be incentivized to acquire and truthfully report information over time using payments that compare the agent's reports with the ex-post observed state of nature. Their baseline model does not have adverse selection; in an extension they consider adverse selection in terms of the agent's beliefs about the state rather than the agent's ability.

2 The Model

Environment. A principal needs to hire an agent to work on a project. The project's quality — synonymous with the state of nature — may either be good or bad, $s \in \{0, 1\}$, where $s = 1$ represents "good". The common prior on the project being good is $\beta_0 \in (0, 1)$. The agent is privately informed about whether his ability is high or low, $\theta \in \{L, H\}$, where $\theta = H$ represents "high". The common prior on ability being high is $\mu_0 \in (0, 1)$. In each period, $t \in \{1, 2, \ldots\}$, the agent can either exert effort (work) or not (shirk); this choice is never observed to the principal. Exerting effort in any period costs the agent $c > 0$. If effort is exerted and the project is good, the project is successful in that period with probability λ^θ; if either the agent shirks or the project is bad, success cannot obtain in that period. Success is observable and once a project is successful, no further effort can be exerted.[7] We assume $1 > \lambda^H > \lambda^L > 0$. A success yields the principal a payoff normalized to 1; the agent does not intrinsically care about project success. Both parties are risk neutral, share a common discount factor $\delta \in (0, 1)$, and are expected utility maximizers.[8]

[7] What is important is that effort has no social value after a single success; it is inessential but convenient to assume that the agent simply cannot exert any further effort. In particular, this means that whenever we make statements about effort being exerted in any period, it is implicit that the project has not already succeeded. More importantly, it will turn out that our results can be applied without change if success is privately observed by the agent but can be verifiably disclosed; see Section 5.

[8] The substance of our analysis and results also hold when $\delta = 1$, as we will discuss; however, we assume $\delta < 1$ to avoid some expositional inconveniences in formal statements.

Contracts. We consider contracting at period zero with full commitment power from the principal. Since there is hidden information at the time of contracting, without loss of generality the principal's problem is to offer the agent a *menu* of dynamic contracts from which the agent chooses one. A dynamic contract specifies a sequence of transfers in each period as a function of the publicly observable history, which is simply whether or not the project has been successful to date. For some applications, it would be natural to assume that once the agent has accepted a contract, he is free to work or shirk in any period (at least up until some termination date). It turns out, however, to be analytically convenient to consider a larger set of contracts in which the principal can stipulate binding "lockout" periods in which the agent is prohibited from working. Our preferred interpretation is that this is a relaxed problem. We will show that the solution to this relaxed problem can be implemented using contracts in which there are no lockout periods; consequently, the solution also solves the problem where the principal cannot prohibit the agent from working in any period. Throughout, we follow the convention that transfers are from the principal to the agent; negative values represent payments in the other direction.

Accordingly, a *contract* is given by $\mathbf{C} = (W_0, \mathbf{b}, \mathbf{w}, \Gamma)$, where W_0 is an up-front transfer at period zero, $\mathbf{b} = (b_t)_{t\in\Gamma}$ is a transfer conditional on the the project being successful in period t, $\mathbf{w} = (w_t)_{t\in\Gamma}$ is a transfer conditional on the project not being successful in period t (nor in any prior period), and Γ is the set of periods at which the agent is not locked out, i.e. at which he is allowed to choose whether to work or shirk, so long as there has not been a prior success.[9] For ease of exposition, we refer to any b_t as a *bonus* and any w_t as a *penalty*; note, however, that b_t is not constrained to be positive nor must w_t be negative. Note also that $\Gamma \subseteq \mathbb{N}$ can be an arbitrary set and, aside from the initial time zero transfer, transfers are only made during periods in Γ, i.e. when the agent makes a choice of whether to work or shirk.[10] We say that a contract is *connected* if Γ is a connected set, i.e. if $\Gamma = \{1, \ldots, T\}$ for some T; in this case we refer to T as the *length* of the contract. The agent's actions are denoted by $\mathbf{a} = (a_t)_{t\in\Gamma}$, where $a_t = 1$ if the agent works in period $t \in \Gamma$ and $a_t = 0$ if the agent shirks.

Payoffs. Given the agent's type θ, a contract $\mathbf{C} = (W_0, \mathbf{b}, \mathbf{w}, \Gamma)$, and a sequence of actions $\mathbf{a} = (a_t)_{t\in\Gamma}$, the principal expected discounted payoff at time zero is

$$\Pi_0^\theta(\mathbf{C}, \mathbf{a}) = \beta_0 \sum_{t\in\Gamma} \delta^t \left[\prod_{s\in\Gamma, s\le t-1} \left(1 - a_s\lambda^\theta\right)\right] \left[a_t\lambda^\theta(1 - b_t) - \left(1 - a_t\lambda^\theta\right) w_t\right] - (1-\beta_0)\sum_{t\in\Gamma}\delta^t w_t - W_0. \tag{1}$$

[9] Throughout this paper, symbols in bold typeface denote vectors.

[10] It is straightforward to confirm that there is no loss of generality in assuming that there are no transfers in any period after the project has succeeded and in also assuming that there are no transfers in lockout periods.

To interpret the above formula, note that W_0 is the up-front payment that is always made. With probability $1-\beta_0$ the state is bad, in which case the project never succeeds and the principal also pays the agent the entire sequence $(w_t)_{t\in\Gamma}$. Conditional on the state being good (which occurs with probability β_0), the probability of project success depends on both the agent's effort choices and his ability; $\prod_{s\in\Gamma, s\leq t-1}\left(1-a_s\lambda^\theta\right)$ is the probability that a success does not obtain between period 1 and $t-1$ conditional on the good state. If the project were to succeed at time t, then the principal would earn a payoff of 1 in that period and would have paid the agent not only the up-front payment but also the sequence $(w_s)_{s\in\Gamma, s\leq t-1}$ followed by b_t.

Through analogous reasoning, the agent's expected discounted payoff at time zero from the contract $\mathbf{C}$ and action profile $\mathbf{a}$ is

$$U_0^\theta\left(\mathbf{C},\mathbf{a}\right)=\beta_0\sum_{t\in\Gamma}\delta^t\left[\prod_{s\in\Gamma, s\leq t-1}\left(1-a_s\lambda^\theta\right)\right]\left[a_t(\lambda^\theta b_t-c)+\left(1-a_t\lambda^\theta\right)w_t\right]+(1-\beta_0)\sum_{t\in\Gamma}\delta^t\left(w_t-a_tc\right)+W_0. \tag{2}$$

If a contract is not agreed upon, we normalize both the principal's and the agent's payoff to be zero.

3 Benchmarks

This section presents preliminaries concerning efficiency benchmarks and simple classes of contracts.

3.1 The First Best

Consider the first-best solution. Since beliefs about the project quality (i.e. the state being good) decline so long as effort has been exerted but success not obtained, the first-best solution is characterized by a stopping rule such that an agent of ability θ keeps exerting effort so long as success has not obtained up until some period t^θ, whereafter effort is no longer exerted. Let β_t^θ be the belief on the state being good, $s=1$, at the beginning of period t, and $\overline{\beta}_t^\theta$ be this belief when the agent has exerted effort in all periods up to time t. The optimal stopping time t^θ is given by

$$t^\theta=\max_{t\geq 0}\left\{t:\overline{\beta}_t^\theta\lambda_\theta\geq c\right\}, \tag{3}$$

where, for each θ, $\overline{\beta}_0^\theta = \beta_0$, and for $t \geq 1$, Bayes' rule yields

$$\overline{\beta}_t^\theta = \frac{\beta_0 \left(1-\lambda^\theta\right)^{t-1}}{\beta_0 \left(1-\lambda^\theta\right)^{t-1} + (1-\beta_0)}. \tag{4}$$

Note that (3) is only well-defined when $\frac{c}{\lambda^\theta} \leq \beta_0$; if $\frac{c}{\lambda^\theta} > \beta_0$, it would be optimal to never experiment, i.e. stop at $t^\theta = 0$. To focus on the most interesting cases, we assume the following:

Assumption 1. *Experimentation is efficient for both types: for $\theta \in \{L, H\}$, $\beta_0 \lambda^\theta > c$.*

Note that in particular, this implies $c < 1$, where 1 is the social benefit from project success. Purely for expositional convenience, we also assume that parameter values are such that $\overline{\beta}_{t^\theta}^\theta = \frac{c}{\lambda^\theta}$. Then, (3) and (4) imply that the optimal stopping time for type θ is

$$t^\theta = 1 + \frac{\log\left(\frac{c}{\lambda^\theta - c}\frac{1-\beta_0}{\beta_0}\right)}{\log\left(1-\lambda^\theta\right)}. \tag{5}$$

The right-hand side of (5) is non-monotonic in λ^θ; it is initially increasing and eventually decreasing. The reason is that there are two countervailing forces: on the one hand, for any given belief about the state, the expected marginal benefit of effort is higher when the agent's ability is higher; on the other hand, the higher is the agent's ability, the more informative about the state is any previous lack of success (given that effort has been exerted), which pushes down the belief about the state (and hence the expected marginal benefit of effort) at any time $t > 1$. Therefore, both $t^H > t^L$ and $t^L < t^H$ are robust possibilities that arise for different parameters.

In this draft, we focus on the subset of the parameter range where it is efficient for the high type to work longer than the low type, condition on no prior success. In other words, we assume:

Assumption 2. *Parameters are such that $t^H > t^L$.*

The case of $t^H < t^L$ is also relevant and will be included in future versions of this paper. Note that the first-best expected discounted surplus at time zero is

$$\sum_{t=1}^{t^\theta} \delta^t \left[\beta_0 \left(1-\lambda^\theta\right)^{t-1}\left(\lambda^\theta - c\right) - (1-\beta_0)c\right].$$

3.2 No Adverse Selection

In the absence of adverse selection the principal can implement the first-best solution and extract all the surplus. A variety of contracts can be used to achieve this; we discuss two classes of relatively simple contracts that will be useful in the subsequent analysis with adverse selection.

3.2.1 Bonus contracts

First consider contracts where once experimentation begins, the only transfers are payments to the agent made when he obtains a success by some deadline.

Definition 1. A *bonus contract* is $\mathbf{C} = (W_0, \mathbf{b}, \mathbf{w}, \Gamma)$ such that $w_t = 0$ for all $t \in \Gamma$. A bonus contract is a *constant-bonus contract* if, in addition, $\Gamma = \{1, \ldots, T\}$ for some T and there is some constant b such that $b_t = b$ for all $t = 1, \ldots, T$.

While a general bonus contract rewards the agent with a time-dependent bonus for success, a constant-bonus contract is connected and pays the same reward independent of when success is obtained up until the deadline. To ease notation, we will denote a bonus contract as just $\mathbf{C} = (W_0, \mathbf{b}, \Gamma)$, a connected bonus contract as $\mathbf{C} = (W_0, \mathbf{b}, T)$, and a constant-bonus contract as $\mathbf{C} = (W_0, b, T)$.

We now argue that when the agent's ability is observable, constant-bonus contracts are optimal for the principal. Suppose the principal offers the agent of type θ a constant-bonus contract $\mathbf{C}^\theta = (W_0^\theta, 1, t^\theta)$, where W_0^θ is chosen such that, conditional on the agent exerting effort in all periods $t = 1, \ldots, t^\theta$, the type θ agent's participation constraint at time 0 binds:

$$U_0^\theta \left(\mathbf{C}^\theta, (1)_{t=1}^{t^\theta}\right) = \sum_{t=1}^{t^\theta} \delta^t \left[\beta_0 \left(1 - \lambda^\theta\right)^{t-1} \left(\lambda^\theta - c\right) - (1 - \beta_0)c\right] + W_0^\theta = 0.$$ [11]

By offering such a constant-bonus contract with a bonus of 1 and making the participation constraint bind, the principal effectively sells the project to the agent at a price that extracts all the surplus. Plainly, this achieves the first-best level of experimentation and the principal cannot improve on this.

[11] On the left hand side, the notation $(1)_{t=1}^{t^\theta}$ denotes the action profile where the agent works in all periods from 1 to t^θ.

3.2.2 Clawback contracts

Now consider contracts where the agent receives no payments for success, and instead is penalized when there is a failure (i.e. success does not obtain). Such a contract will satisfy the agent's participation constraint only if he is paid some positive amount at time zero, which motivates the terminology of "clawbacks". Formally:

Definition 2. A *clawback contract* is $\mathbf{C} = (W_0, \mathbf{b}, \mathbf{w}, \Gamma)$ such that $b_t = 0$ for all $t \in \Gamma$. A clawback contract is a *onetime-clawback contract* if, in addition, $\Gamma = \{1, \ldots, T\}$ for some T and $w_t = 0$ for all $t = 1, \ldots, T-1$.

We highlight that a clawback contract allows for $|w_t| > W_0$ in any t, i.e. clawbacks should be understood as allowing for penalties larger than what the agent initially received from the principal at time zero. Unlike bonus contracts (where at most one bonus is ever paid), a general clawback contract involves the possibility of transfers in multiple periods after experimentation has begun; in particular, in a connected clawback contract of length T, the agent may be penalized for failure in every period from 1 to T. However, in a onetime-clawback contract, this is not the case because the agent makes only one payment, which is at time T (conditional on no success up to that point). To ease notation, we will denote a clawback contract as just $\mathbf{C} = (W_0, \mathbf{w}, \Gamma)$, a connected clawback contract as $\mathbf{C} = (W_0, \mathbf{w}, T)$, and a onetime-clawback contract as $\mathbf{C} = (W_0, w_T, T)$.

Given observable types, suppose the principal offers a type θ agent a onetime-clawback contract $\mathbf{C}^\theta = (W_0^\theta, w_{t^\theta}^\theta, t^\theta)$ where W_0^θ is chosen such that, conditional on the agent exerting effort in all periods $t = 1, \ldots, t^\theta$, the agent's participation constraint at time 0 binds:

$$U_0^\theta\left(\mathbf{C}^\theta, (1)_{t=1}^{t^\theta}\right) = W_0^\theta - c\left[\sum_{t=1}^{t^\theta} \delta^t \left(\beta_0 \left(1-\lambda^\theta\right)^{t-1} + (1-\beta_0)\right)\right] - w_{t^\theta}^\theta \delta^{t^\theta}\left(\beta_0 (1-\lambda^\theta)^{t^\theta} + (1-\beta_0)\right) = 0.$$

First best requires the agent to work in all periods until t^θ so long as success has not been obtained. From the one-step deviation principle, the incentive compatibility conditions for effort are summarized by the following inequality: for any $t \in \{1, \ldots, t^\theta\}$,

$$\lambda^\theta \overline{\beta}_t^\theta \left[\delta^{t^\theta - t}\left(1-\lambda^\theta\right)^{t^\theta - t}\left(-w_{t^\theta}^\theta\right) + \sum_{s=t+1}^{t^\theta} \delta^{s-t}\left(1-\lambda^\theta\right)^{s-t-1} c\right] \geq c. \tag{6}$$

The right-hand side of (6) is the constant cost of effort. The left-hand side is the benefit of effort at any time t (given that effort has been exerted and success not obtained at all prior periods): with probability $\lambda^\theta \overline{\beta}_t^\theta$ there will be a success in period t and the agent saves both the expected discounted penalty, $\delta^{t^\theta - t}\left(1-\lambda^\theta\right)^{t^\theta - t}\left(-w_{t^\theta}^\theta\right)$, and the expected discounted cost of future effort, $\sum_{s=t+1}^{t^\theta} \delta^{s-t}\left(1-\lambda^\theta\right)^{s-t-1} c$. It follows that the agent will work in all periods if $w_{t^\theta}^\theta$ is chosen low enough, i.e., the clawback penalty is large enough. In this case, the first best is again implemented and the principal extracts all the surplus.

3.3 No Moral Hazard

Suppose there were adverse selection but the agent's effort is observable (and contractible). Then, since one can ignore the incentive compatibility for effort constraints, the principal's problem reduces to optimize over how long each type should work conditional on not obtaining success, subject to time zero participation constraints and the incentive compatibility constraints that each type should not choose the other's contract. This is analogous to a standard monopolist screening problem where the stopping time maps into "quantity" and the time-zero transfer maps into "price". It is routine to check that the standard sorting condition of increasing differences holds.[12] Hence, modulo time (i.e. "quantity") being discrete, an optimal contract can be derived from textbook arguments, with the conclusion that the high type's stopping time would be efficient while the low type's stopping time would generally be inefficiently early.

4 Optimal Contracts

With the preceding preliminaries in hand, we turn to our main results on optimal contracts for experimentation with both adverse selection and moral hazard.

Given any contract $\mathbf{C} = (W_0, \mathbf{b}, \mathbf{w}, \Gamma)$, define $\boldsymbol{\alpha}^\theta(\mathbf{C}) = \left(a_t^\theta(\mathbf{C})\right)_{t\in\Gamma}$ as an optimal action

[12]Denote a type θ agent's utility from being required to work until time T (conditional on no prior success) as $v(T,\theta)$; this is gross any time-zero transfer. From (2), we compute

$$v(T,\theta) = -c\left(\beta_0 \sum_{t=1}^{T} \delta^t (1-\lambda^\theta)^{t-1} + (1-\beta_0)\sum_{t=1}^{T}\delta^t\right).$$

Hence, for any $T' > T$,

$$v(T',H) - v(T,H) - [v(T',L) - v(T,L)] = c\beta_0 \sum_{t=T+1}^{T'} \delta^t \left[\left(1-\lambda^L\right)^{t-1} - \left(1-\lambda^H\right)^{t-1}\right] > 0.$$

plan for type θ.[13] That is, for $\theta \in \{H, L\}$,

$$\boldsymbol{\alpha}^{\theta}(\mathbf{C}) \in \underset{\mathbf{a}=(a_t)_{t\in\Gamma}}{\arg\max}\; U_0^{\theta}(\mathbf{C}, \mathbf{a}), \tag{7}$$

where $U_0^{\theta}(\cdot)$ was defined in (2). An optimal menu of contracts, $(\mathbf{C}^H, \mathbf{C}^L)$ maximizes the principal's ex-ante expected profit subject to incentive compatibility constraints for effort (IC_a^{θ} below), participation constraints (IR^{θ} below), and self-selection constraints for the agent's choice of contract ($\text{IC}^{\theta,\theta'}$ below). Formally, recalling the definition of $\Pi_0^{\theta}(\cdot)$ from (1), the principal's program is:

$$\max_{(\mathbf{C}^H, \mathbf{C}^L, \mathbf{a}^H, \mathbf{a}^L)} \mu_0 \Pi_0^H\left(\mathbf{C}^H, \mathbf{a}^H\right) + (1-\mu_0)\,\Pi_0^L\left(\mathbf{C}^L, \mathbf{a}^L\right)$$

subject to, for all $\theta, \theta' \in \{H, L\}$,

$$\mathbf{a}^{\theta} \in \underset{\mathbf{a}=(a_t)_{t\in\Gamma^{\theta}}}{\arg\max}\; U_0^{\theta}\left(\mathbf{C}^{\theta}, \mathbf{a}\right), \tag{IC_a^{θ}}$$

$$U_0^{\theta}\left(\mathbf{C}^{\theta}, \mathbf{a}^{\theta}\right) \geq 0, \tag{IR^{θ}}$$

$$U_0^{\theta}\left(\mathbf{C}^{\theta}, \mathbf{a}^{\theta}\right) \geq U_0^{\theta}\left(\mathbf{C}^{\theta'}, \boldsymbol{\alpha}^{\theta}\left(\mathbf{C}^{\theta'}\right)\right). \tag{$\text{IC}^{\theta,\theta'}$}$$

It is instructive to contrast the above problem with that in the absence of moral hazard (Subsection 3.3). Dynamic moral hazard is reflected directly in the constraints (IC_a^{θ}) and indirectly in the constraints ($\text{IC}^{\theta,\theta'}$) via the term $\boldsymbol{\alpha}^{\theta}(\mathbf{C}^{\theta'})$. To get a sense of how these matter, begin with the obvious point that it is no longer sufficient for the principal to only use time-zero transfers since the agent will not work along the course of the contract unless incentivized to do so at each period. Note that the agent's incentive to work at any time t is shaped not only by period t transfers (b_t and w_t) but also by the subsequent sequences of transfers. For example, *ceteris paribus*, a penalty for failure at period $t+1$, $w_{t+1} < 0$, creates an incentive to work in period t, whereas a bonus for success in period $t+1$, $b_{t+1} > 0$, creates an incentive to shirk in period t. Incentives at any point in time also depend on both the current and future beliefs about the state, which is different for the two agent types. Incentives further differ across the two types because their marginal benefit of effort conditional on the good state is different. Consequently, the optimal plan of action for a given contract will generally be different for the two types of the agent.

Therefore, in the presence of dynamic moral hazard, the analogy suggested in Subsec-

[13] If there are multiple optimal action plans, any selection may be used.

tion 3.3 becomes rather limited: the principal must use a much richer set of instruments than merely setting a "price" to optimally incentivize the agent who intrinsically dislikes consuming more "quantity". Notwithstanding, our main efficiency result is:

Theorem 1. *In any optimal menu of contracts, each type $\theta \in \{H, L\}$ is induced to work in every period up until some second-best stopping time, $\bar{t}^\theta$. Relative to the first-best stopping times, t^H and t^L, the second-best has $\bar{t}^H = t^H$ and $\bar{t}^L \leq t^L$; moreover, at the limit when time becomes continuous, $\bar{t}^L < t^L$.*[14]

Relative to the first best, there is no distortion in the stopping time of the high ability agent whereas the low ability agent stops experimenting too early. In this sense, despite the complications arising from dynamic moral hazard, we recover a familiar "no distortion (only) at the top" result. For typical parameters, the induced stopping time for the low type is some $\bar{t}^L \in (0, t^L)$; however, it is possible that the low type will be induced to not experiment at all ($\bar{t}^L = 0$) and it is also possible to have no distortion in the low type's stopping time ($\bar{t}^L = t^L$). The former possibility arises for reasons akin to exclusion in the standard model (e.g. the prior, μ_0, on the high type is sufficiently high); the latter possibility is because time is discrete. Indeed, when the interval between any two periods gets sufficiently small, some distortion must obtain.

While the efficiency properties of Theorem 1 are familiar, deriving them in our framework is novel. We establish these properties through a characterization of a class of optimal menus:

Theorem 2. *There is an optimal menu in which the principal separates the two types using connected clawback contracts. In particular, the optimum can be implemented using a onetime-clawback contract for type H, $\mathbf{C}^H = (W_0^H, w_{t^H}^H, t^H)$ with $w_{t^H}^H < 0 < W_0^H$, and a connected clawback contract for type L, $\mathbf{C}^L = (W_0^L, \mathbf{w}^L, \bar{t}^L)$, where for all $t \in \{1, \dots, \bar{t}^L\}$,*

$$w_t^L = \begin{cases} -(1-\delta)\frac{c}{\beta_t^L \lambda^L} & \text{if } t < \bar{t}^L, \\ -\frac{c}{\beta_{\bar{t}^L}^L \lambda^L} & \text{if } t = \bar{t}^L, \end{cases} \tag{8}$$

and $W_0^L > 0$ is such that the participation constraint, (IR^L), binds. Type H gets an information rent: $U_0^H(\mathbf{C}^H, \boldsymbol{\alpha}^H(\mathbf{C}^H)) > 0$.

Proof. See Appendix A. *Q.E.D.*

Within the class of clawback contracts, Theorem 2 characterizes what is essentially the

[14] To add: details of how the continuous-time limit is formalized.

(generically) unique optimal contract for the low type.[15] Notice from (8) that for any δ, it is a clawback contract in which the agent pays the principal an increasing penalty in each period $t < \bar{t}^L$ at which the project does not succeed (since $\overline{\beta}_t^L > \overline{\beta}_{t+1}^L$), followed by a larger penalty that "jumps" in the final period $\bar{t}^L$ conditional on no success then.[16] At the limit of $\delta \to 1$, i.e. when there is no discounting, this contract reduces to a *onetime*-clawback contract where the low type only pays the penalty if he has not succeeded by $\bar{t}^L$.[17] For any δ, the high type's contract is a onetime-clawback contract in which he only pays a penalty to the principal if there is no success by the efficient stopping time t^H. Per Theorem 1, both types exert effort in every period until their respective stopping times.

Remark 1. Since the optimal contracts in Theorem 2 are connected clawback contracts in which the agent exerts effort in each period up until some deadline, it follows immediately that the principal does not need to use the instrument of locking the agent out in any period; rather, by simply providing no transfers after each contract's deadline, it would be incentive compatible for the agent to stop exerting effort at the deadline.

4.1 Sketch of Solution

As the proof of Theorem 2 (and hence Theorem 1) is a central contribution of this paper, we now sketch in some detail the steps involved. Recall that a general menu of contracts is $(\mathbf{C}^H, \mathbf{C}^L)$ where for each $\theta \in \{H, L\}$, $\mathbf{C}^\theta = (W_0^\theta, \mathbf{b}^\theta, \mathbf{w}^\theta, \Gamma^\theta)$.

Step 1: It is without loss to focus on contracts for type L that induce him to work in every non-lockout period, i.e. on contracts in the set $\{\mathbf{C}^L : \boldsymbol{\alpha}^L(\mathbf{C}^L) = (1)_{t\in\Gamma^L}\}$. The idea is as follows: fix any contract, $\mathbf{C}^L$, in which there is some period, $t \in \Gamma^L$, such that it would be suboptimal for type L to work in period t. Since the outcome for type L in period t is then deterministic (the project will not succeed in period t), one can modify $\mathbf{C}^L$ to create a new contract, $\widehat{\mathbf{C}}^L$, in which $t \notin \widehat{\Gamma}^L$, and w_t^L is "shifted up" by one period with an adjustment for discounting. This ensures that the incentives for type L in all other periods remain unchanged, and critically, that no matter what behavior would have been optimal for type H under contract $\mathbf{C}^L$, the new contract is less attractive to type H, i.e. for any $\boldsymbol{\alpha}^H(\mathbf{C}^L)$ and any $\boldsymbol{\alpha}^H(\widehat{\mathbf{C}}^L)$, $U_0^H(\mathbf{C}^L, \boldsymbol{\alpha}^H(\mathbf{C}^L)) \geq U_0^H(\widehat{\mathbf{C}}^L, \boldsymbol{\alpha}^H(\widehat{\mathbf{C}}^L))$.

Step 2: It is without loss to focus on contracts in which there are no bonuses, i.e. to

[15]The proof of Theorem 2 reveals that there is generically a unique optimal value of $\bar{t}^L$, even though there is no closed form solution.

[16]By "jump" here we refer to the fact that for any period $t < \bar{t}^L$, $w_t^L \overline{\beta}_t^L = (1-\delta) w_{\bar{t}^L}^L \overline{\beta}_{\bar{t}^L}^L$.

[17]Note that the value of $\bar{t}^L$ generally depends on δ, but always remains weakly smaller than t^L by Theorem 1.

optimize over menus of (not necessarily connected) clawback contracts. To see the intuition for why clawback contracts suffice, consider for simplicity a connected contract $\mathbf{C}^\theta$ for some $\theta \in \{H, L\}$ in which $\Gamma^\theta = \{1, \ldots, T\}$. We build a new contract $\widehat{\mathbf{C}}^\theta$ in which $\widehat{\Gamma}^\theta = \Gamma^\theta$, but $\widehat{b}^\theta_T = 0$ and $\widehat{w}^\theta_T = w^\theta_T - b^\theta_T$. Clearly, for either type, the incentive to work in period T is identical under these two contracts. On the other hand, because the continuation value from reaching period T has been reduced, effort incentives in the previous periods have been affected. But this can be fully corrected by setting $\widehat{w}^\theta_{T-1} = w^\theta_{T-1} + \delta b^\theta_T$. The procedure can now be iterated backward so that $\widehat{b}^\theta_t = 0$ for all $t \in \{1, \ldots, T\}$, yet $\boldsymbol{\alpha}^{\theta'}(\widehat{\mathbf{C}}^\theta) = \boldsymbol{\alpha}^{\theta'}(\mathbf{C}^\theta)$ for each $\theta' \in \{H, L\}$. Furthermore, this implies that the principal's expected payoff evaluated at time zero is the same from either type under both $\mathbf{C}^\theta$ and $\widehat{\mathbf{C}}^\theta$.

The gain in focussing on clawback contracts is that raising incentives for effort in any period t through the penalty for failure in that period, i.e. reducing w_t, has a *positive feedback* on incentives for effort in earlier periods, because this only reduces the continuation value for the agent from reaching period t. By contrast, incentive provision through a bonus has a *negative feedback* on effort incentives in earlier periods: raising the reward for success, b_t, increases the continuation value for the agent from reaching period t. For this reason, dealing with the dynamic moral hazard problem is analytically more convenient under clawback contacts.[18]

Step 3: Based on the above two steps, the principal can optimize over menus of clawback contracts in which the low type's contract induces him to work in every (non-lockout) period, subject to, for each $\theta \in \{H, L\}$, (IC^θ_a), (IR^θ), and ($\mathrm{IC}^{\theta,\theta'}$). Call this program [P]. Since it is daunting to determine *a priori* which constraints in this program bind, we focus instead on a relaxed program, [RP1], that (i) ignores (IR^H) and (IC^{LH}), and (ii) replaces (IC^{HL}) by a relaxed version, called (Weak-IC^{HL}) that only requires that type H should not want to deviate to the L type's contract assuming that type H would work in all periods after taking type L's contract.[19]

In the relaxed program [RP1], it is straightforward to show that Weak-IC^{HL} and IR^L must bind at an optimum: otherwise, time-zero transfers in one of two contracts can be profitably lowered without violating any of the constraints. Consequently, one can substitute from the binding version of these constraints to rewrite the objective function as the sum of total surplus less an "information rent" for the high type, as in the standard approach.[20] We are left with

[18] Although, we will show subsequently that the optimum can also be implemented via bonus contracts.

[19] Formally, in light of other constraints, (IC^{HL}) requires $U^H_0(\mathbf{C}^H, \boldsymbol{\alpha}^H(\mathbf{C}^H)) \geq U^H_0(\mathbf{C}^L, \boldsymbol{\alpha}^H(\mathbf{C}^L))$ whereas (Weak-IC^{HL}) requires only $U^H_0(\mathbf{C}^H, \boldsymbol{\alpha}^H(\mathbf{C}^H)) \geq U^H_0\left(\mathbf{C}^L, \mathbf{1}\right)$. Note that the latter does not imply the former in an arbitrary contract $\mathbf{C}^L$ with $\boldsymbol{\alpha}^L(\mathbf{C}^L) = \mathbf{1}$, because it need not be the case that $\boldsymbol{\alpha}^H(\mathbf{C}^L) = \mathbf{1}$. To illustrate this point, consider a limit case of parameters $\beta_0 = 1$ and $\lambda_H = 1$. Then, for any $\lambda^L < 1$, $c > 0$, and $\delta \in (0, 1)$, in a onetime-clawback contract with a sufficiently negative penalty in period 2, type L would find it optimal to work in both periods whereas type H would work in period 2 but not in period 1.

[20] It is worth emphasizing, however, that this approach only works in the relaxed program, [RP1]. In the full

a relaxed program, [RP2], whose objective is to maximize social surplus less the high type's information rent, and whose only constraints are the direct moral hazard constraints IC_a^H and IC_a^L, where type L must work in all periods. This program is tractable because it can be solved by separately optimizing over each type's clawback contract. The following steps, 4–7, derive an optimal contract for type L in program [RP2] that has useful properties.

Step 4: We show that there is an optimal clawback contract for type L that is connected. A rough intuition is as follows.[21] Because type L is required to work in all non-lockout periods, the value of the objective function in program [RP2] can be improved by removing any gaps in Γ^L in one of two ways: either by "shifting up" the sequence of effort and penalties or by terminating the contract early (suitably adjusting for discounting in either case). Shifting up the sequence of effort and penalties eliminates inefficient delays in type L's experimentation, but it also increases the rent given to type H, because the penalties — which are more likely to be borne by type L than type H — are now paid earlier. Conversely, terminating the contract early reduces the rent given to type H by lowering the total penalties in the contract, but it also shortens experimentation by type L. It turns out that either of these modifications may be beneficial (i.e. increase the value of the objective function of [RP2]), but critically, at least one of them will be.

Step 5: Given any T^L, there are many penalty sequences that can be used by a connected clawback contract of length T^L to induce the low-ability agent to work in each period $1, \ldots, T^L$. We construct the unique sequence, call it $\overline{\mathbf{w}}(T^L)$, that ensures that the low type's incentive constraint for effort binds in each period $1, \ldots, T^L$. The intuition is straightforward: in the final period, T^L, there is obviously a unique such penalty, as it must solve $\overline{w}^L_{T^L}(T^L) = -c + (1 - \overline{\beta}^L_{T^L}\lambda^L)\overline{w}^L_{T^L}(T^L)$. Iteratively working backward using a one-step deviation principle, this pins down penalties in each earlier period through the (forward-looking) incentive constraint for effort in each period. Naturally, for any T^L and $t \in \{1, \ldots, T^L\}$, $\overline{w}^L_t(T^L) < 0$, i.e. as suggested by the term "penalty", the agent pays the principal each time there is a failure.

Step 6: We next show that there is a connected clawback contract for type L that solves program [RP2] using the Step 5 penalty structure $\overline{\mathbf{w}}^L(\cdot)$. An intuition is that if type H were to take type L's clawback contract, he could work in each period and obtain some rent because he is less likely to have to pay the penalty in any period.[22] Any slack left in type L's contract with

program [P], one cannot directly establish that either (IR^L) or (IC^{HL}) must bind. This contrast with the standard approach is because of dynamic moral hazard.

[21] For the intuition that follows, assume that all penalties being discussed are negative transfers, i.e. transfers from the agent to the principal.

[22] As previously noted, it may not be optimal for type H to work in each period when taking type L's contract, but recall that program [RP2] effectively constrains type H to only such deviations.

regards to type L's incentive constraints for effort only increases the rent that can be obtained by type H.

Step 7: In light of Steps 4–6, an optimal contract for type L in program [RP2] can be found by just choosing the optimal length of connected clawback contracts with the penalty structure $\overline{\mathbf{w}}^L(\cdot)$. We show that the optimal length, $\bar{t}^L$, cannot be larger than the first-best stopping time: $\bar{t}^L \leq t^L$. This is derived from a standard monotone comparative statics argument. The intuition is that since the principal distorts the stopping time for type L only to reduce the rent given to type H, incentivizing over-experimentation cannot be optimal since that would only increase the scope for rent.

Step 8: Let $\overline{\mathbf{C}}^L$ be the contract for type L identified in Steps 4–7.[23] The final step is to show that there is a solution to [RP2] that combines $\overline{\mathbf{C}}^L$ with a suitable onetime-clawback contract for the high type and also solves the original program [P]. First, we show that $\boldsymbol{\alpha}^H(\overline{\mathbf{C}}^L) = \mathbf{1}$, i.e. if type H were to take contract $\overline{\mathbf{C}}^L$, it would be optimal for him to work in all periods $1, \ldots, \bar{t}^L$. The intuition is as follows: under contract $\overline{\mathbf{C}}^L$, type H has a higher expected probability of success from working in any period $t \leq \bar{t}^L$, no matter his prior choices of effort, than does type L in period t given that type L has exerted effort in all prior periods (recall $\boldsymbol{\alpha}^L(\overline{\mathbf{C}}^L) = \mathbf{1}$).[24] Consequently, the fact that $\overline{\mathbf{C}}^L$ makes type L indifferent between working and shirking in each period up to $\bar{t}^L$ (given that he has worked in all prior periods) implies through a recursive induction argument that type H would find it strictly optimal to work in each period up to $\bar{t}^L$ no matter his prior history of effort, and hence $\boldsymbol{\alpha}^H(\overline{\mathbf{C}}^L) = \mathbf{1}$. It then follows that if type H's contract is chosen to satisfy (Weak-IC^{HL}) then it will also satisfy (IC^{HL}) and (IR^H); the latter because $U_0^H(\overline{\mathbf{C}}^L, \mathbf{1}) \geq U_0^L(\overline{\mathbf{C}}^L, \mathbf{1})$.

Lastly, we show that by choosing a onetime-clawback contract for type H that imposes a sufficiently severe penalty in period t^H and compensating type H through the initial transfer W_0^H, the principal maximizes the social surplus from the high type, satisfies (Weak-IC^{HL}), and also satisfies (IC^{LH}). In particular, (IC^{LH}) is satisfied because the principal can exploit the two types' differing probabilities of success by making the onetime-clawback contract for type H "risky enough" to deter type L from taking it while still satisfying (Weak-IC^{HL}) and hence (IR^H).

[23]The initial transfer in $\overline{\mathbf{C}}^L$ is set to make the participation constraint for type L bind.

[24]This relies on Step 7 showing that $\bar{t}^L \leq t^L$, because then $t^H > t^L$ implies that for any $t \in \{1, \ldots, \bar{t}^L\}$, $\beta_t^H \lambda^H > \overline{\beta}_t^L \lambda^L$ for any history of effort by type H in periods $1, \ldots, t-1$.

4.2 Bonus Contracts

While Theorem 2 shows that the second-best can be implemented using a menu of clawback contracts, it is also of interest whether bonus contracts can be used.

Theorem 3. *The second-best can also be implemented using a menu of bonus contracts. In particular, the principal can use the following connected bonus contract for type L: $\mathbf{C}^L = (W_0^L, \mathbf{b}^L, \bar{t}^L)$, where, for any $t \in \{1, \ldots, \bar{t}^L\}$,*

$$b_t^L = (1-\delta)c \sum_{s=t}^{\bar{t}^L - 1} \delta^{s-t} \left(\overline{\beta}_t^L \lambda^L\right)^{-1} + \delta^{\bar{t}^L - t} c \left(\overline{\beta}_{\bar{t}^L}^L \lambda^L\right)^{-1}, \tag{9}$$

and W_0^L is such that the participation constraint, (IR^L), binds. For the high type, the principal can use a constant-bonus contract $\mathbf{C}^H = (W_0^H, b^H, t^H)$ with a suitably chosen W_0^H and $b^H > 0$.

Proof. See Appendix B. *Q.E.D.*

Comparing (8) and (9) reveals that type L's bonus sequence in the contract of Theorem 3 and his penalty sequence in the contract of Theorem 2 have the following relationship: for all $t \in \{1, \ldots, \bar{t}^L\}$,

$$b_t^L = \sum_{s=t}^{\bar{t}^L - 1} \delta^{s-t}(-w_s^L) + \delta^{\bar{t}^L - t}(-w_{\bar{t}^L}^L). \tag{10}$$

In particular, as is intuitive, $b_{\bar{t}^L}^L = -w_{\bar{t}^L}^L$, since type L's effort incentive in the final period depends on the difference between the bonus and the penalty in that period. In earlier periods, comparing the incentive constraint for effort in bonus contracts and clawback contracts is more complex due to the differences in how future bonuses and future penalties affect present considerations, for reasons previously noted. The formula (9) obtains from choosing the bonus sequence for type L to make his incentive constraint for effort bind in each period and thereby minimize the rent provided to type H. It is also readily verified that in the bonus sequence (9),

$$b_t^L = \frac{(1-\delta)c}{\overline{\beta}_t^L \lambda^L} + \delta b_{t+1}^L \quad \text{for any } 1 \leq t < \bar{t}^L, \tag{11}$$

and hence the reward for success increases over time. Notice that in the limit as $\delta \to 1$, type L's bonus contract becomes a constant-bonus contract, analogous to how the clawback contract in Theorem 2 becomes a onetime-clawback contract.

5 Discussion

Suppose that project success is privately observed by the agent but can be verifiably disclosed.[25] We assume in this *private-observability* setting that the principal's payoff from project success obtains only when the agent discloses it and contracts are conditioned not on project success but rather the disclosure of project success. Private observability introduces additional constraints for the principal, since the agent must also now be incentivized to not withhold project success. For example, in a bonus contract where $\delta b_{t+1} > b_t$, an agent who obtains success in period t would strictly prefer to withhold it and continue to period $t+1$, shirk in that period, and then reveal the success at the end of period $t+1$.

Theorem 4. *Even if project success is privately observed, the menus of contracts identified in both Theorem 2 and Theorem 3 remain optimal and implement the same outcome as when project success is publicly observable.*

Proof. It suffices to show that in each of the menus, each of the contracts would induce an agent (of either type) to reveal project success immediately when it is obtained.

Consider first the menu of Theorem 2: for each $\theta \in \{H, L\}$, $\mathbf{C}^\theta$, the contract for type θ is a clawback contract in which $w_t^\theta \leq 0$ for all t. Hence, no matter which contract the agent takes and no matter his type, it is optimal to reveal a success when obtained.

For the implementation in Theorem 3, observe from (11) that type L's bonus contract has the property that $\delta b_{t+1}^L < b_t^L$ for all $1 \leq t < \bar{t}^L$; moreover, this property also holds in type H's bonus contract because it is a constant-bonus contract. Hence, under either contract, it is optimal for the agent of either type to disclose success immediately when obtained. *Q.E.D.*

Therefore, project success being privately observed by the agent does not reduce the principal's payoff compared to the baseline setting where project success is publicly observable (and contractible), so long as the agent can verifiable disclose project success. However, unlike the menus of Theorem 2 and Theorem 3, not every optimal menu under public observability is optimal under private observability.[26] In this sense, these optimal menus have a desirable

[25] If disclosure were not verifiable and instead the agent could only make cheap-talk claims, then it is impossible to incentivize the agent.

[26] In particular, there are optimal menus of bonus contracts under public observability that are suboptimal under private observability because the contract given to type H is such that he would have an incentive to delay disclosure of project success. Formally, the dynamic incentive constraint for effort under public observability requires

$$b_t^H \geq \frac{c(1 - \delta\beta_t^H \lambda^H)}{\beta_t^H \lambda^H} + \delta\lambda^H b_{t+1}^H,$$

robustness property that other optimal menus need not.

which can be satisfied with $b_t^H < \delta b_{t+1}^H$, in which case the contract would be suboptimal under private observability as noted earlier. An analogous point applies to menus of clawback contracts.

A Proof of Theorem 2

The proof below for Theorem 2 also proves Theorem 1. We remind the reader that Subsection 4.1 provides an outline and intuition for the proof.

A.1 Step 1

Given any contract for type L, $\mathbf{C}^L = \left(W_0^L, \mathbf{b}^L, \mathbf{w}^L, \Gamma^L\right)$, we claim that there is a contract $\widehat{\mathbf{C}}^L = \left(\widehat{W}_0^L, \widehat{\mathbf{b}}^L, \widehat{\mathbf{w}}^L, \widehat{\Gamma}^L\right)$ such that:

(i) $\mathbf{1} \in \underset{\mathbf{a}=(a_t)_{t\in\widehat{\Gamma}^L}}{\arg\max}\ U_0^L(\widehat{\mathbf{C}}^L, \mathbf{a})$;

(ii) $U_0^L(\mathbf{C}^L, \boldsymbol{\alpha}^L(\mathbf{C}^L)) = U_0^L(\widehat{\mathbf{C}}^L, \mathbf{1})$;

(iii) $\Pi_0^L(\mathbf{C}^L, \boldsymbol{\alpha}^L(\mathbf{C}^L)) = \Pi_0^L(\widehat{\mathbf{C}}^L, \mathbf{1})$; and

(iv) $U_0^H(\mathbf{C}^L, \boldsymbol{\alpha}^H(\mathbf{C}^L)) \geq U_0^H(\widehat{\mathbf{C}}^L, \boldsymbol{\alpha}^H(\widehat{\mathbf{C}}^L))$.

The claim is trivially true if it is optimal for type L to work in each (non-lockout) period of $\mathbf{C}^L$, so assume that it is suboptimal for type L to work in some period t in $\mathbf{C}^L$. Denote the largest preceding period in Γ^L as

$$p(t) = \begin{cases} \max \Gamma^L \setminus \{t, t+1, \ldots\} & \text{if } \exists s \in \Gamma^L \text{ s.t. } s < t, \\ 0 & \text{otherwise.} \end{cases}$$

Construct $\widehat{\mathbf{C}}^L = \left(\widehat{W}_0^L, \widehat{\mathbf{b}}^L, \widehat{\mathbf{w}}^L, \widehat{\Gamma}^L\right)$ as follows:

$$\begin{aligned} \widehat{\Gamma}^L &= \Gamma^L \setminus \{t\}\,; \\ \widehat{w}_s^L &= \begin{cases} w_s^L & \text{if } s \neq p(t) \text{ and } s \in \widehat{\Gamma}^L, \\ w_s^L + \delta^{t-p(t)} w_t^L & \text{if } s = p(t) > 0; \end{cases} \\ \widehat{b}_s^L &= b_s^L \text{ for all } s \in \widehat{\Gamma}^L; \\ \widehat{W}_0^L &= \begin{cases} W_0^L & \text{if } p(t) > 0, \\ W_0^L + \delta^t w_t^L & \text{if } p(t) = 0. \end{cases} \end{aligned}$$

Notice that under contract $\mathbf{C}^L$, type L would have shirked in period t and thus received w_t^L in this period; the new contract $\widehat{\mathbf{C}}^L$ just locks the agent out in period t and shifts the payment w_t^L up to the preceding non-lockout period, suitably discounted. It follows that the incentives for effort for type L remain unchanged in any other period, i.e. that for any $\mathbf{a}^* \in \underset{\mathbf{a}=(a_t)_{t\in\Gamma^L}}{\arg\max}\ U_0^L(\mathbf{C}^L, \mathbf{a})$, $\mathbf{a}^*_{-t} \in \underset{\mathbf{a}=(a_t)_{t\in\widehat{\Gamma}^L}}{\arg\max}\ U_0^L(\widehat{\mathbf{C}}^L, \mathbf{a})$; moreover, plainly both the principal's payoff from type L under this contract and type L's payoff do not change. Finally, observe that for type H, no matter what his optimal action at time t in $\mathbf{C}^L$ would have been (it may have been to work rather than shirk), his payoff from $\widehat{\mathbf{C}}^L$ must be weakly lower because the lockout in period t is effectively as though he has been forced to shirk in period t and receive w_t^L.

Performing the above procedure repeatedly for all periods $t \in \Gamma^L$ in which it would be suboptimal for type L to work under $\mathbf{C}^L$ yields a final contract $\widehat{\mathbf{C}}^L$ which satisfies all the properties stated in the claim.

A.2 Step 2

We claim that for any type $\theta \in \{H, L\}$ and any contract $\mathbf{C}^\theta = (W_0^\theta, \mathbf{b}^\theta, \mathbf{w}^\theta, \Gamma^\theta)$, there exists another contract, $\widehat{\mathbf{C}}^\theta = (\widehat{W}_0^\theta, (0)_{t\in\Gamma^\theta}, \widehat{\mathbf{w}}, \Gamma^\theta)$, that performs identically, i.e. such that for any $\theta' \in \{H, L\}$ and $\boldsymbol{\alpha}^{\theta'}(\mathbf{C}^\theta)$:

(i) $\boldsymbol{\alpha}^{\theta'}(\mathbf{C}^\theta) \in \underset{\mathbf{a}=(a_t)_{t\in\Gamma^\theta}}{\arg\max}\ U_0^{\theta'}\left(\widehat{\mathbf{C}}^\theta, \mathbf{a}\right)$;

(ii) $U_0^{\theta'}(\mathbf{C}^\theta, \boldsymbol{\alpha}^{\theta'}(\mathbf{C}^\theta)) = U_0^{\theta'}(\widehat{\mathbf{C}}^\theta, \boldsymbol{\alpha}^{\theta'}(\widehat{\mathbf{C}}^\theta))$; and

(iii) $\Pi_0^{\theta'}\left(\mathbf{C}^\theta, \boldsymbol{\alpha}^{\theta'}(\mathbf{C}^\theta)\right) = \Pi_0^{\theta'}\left(\widehat{\mathbf{C}}^\theta, \boldsymbol{\alpha}^{\theta'}(\mathbf{C}^\theta)\right)$.

To prove the claim, fix some θ and $\mathbf{C}^\theta = (W_0^\theta, \mathbf{b}^\theta, \mathbf{w}^\theta, \Gamma^\theta)$. The result is trivial if $\Gamma^\theta = \emptyset$, so assume $\Gamma^\theta \neq \emptyset$. For any period $t \in \Gamma^\theta$, define the largest preceding period in Γ^θ as

$$p(t) = \begin{cases} \max \Gamma^\theta \setminus \{t, t+1, \ldots\} & \text{if } \exists s \in \Gamma^\theta \text{ s.t. } s < t, \\ 0 & \text{otherwise.} \end{cases}$$

We modify the contract $\mathbf{C}^\theta$ into $\widehat{\mathbf{C}}^\theta = \left(\widehat{W}_0^\theta, \widehat{\mathbf{b}}^\theta, \widehat{\mathbf{w}}^\theta, \Gamma^\theta\right)$ by only changing transfers for periods t and $p(t)$ as follows:

$$\begin{aligned} \widehat{w}_t^\theta &= w_t^\theta - b_t^\theta, \\ \widehat{b}_t^\theta &= 0, \end{aligned}$$

and, if $p(t) > 0$ then

$$\begin{aligned} \widehat{w}_{p(t)}^\theta &= w_{p(t)}^\theta + \delta^{t-p(t)} b_t^\theta, \\ \widehat{b}_{p(t)}^\theta &= b_{p(t)}^\theta, \end{aligned}$$

while if $p(t) = 0$ then $\widehat{W}_0^\theta = W_0^\theta + \delta^t b_t^\theta$.

Clearly, the incentives for effort for both types from period $t+1$ on are identical in $\mathbf{C}^\theta$ and $\widehat{\mathbf{C}}^\theta$ since no changes have been made from period $t+1$ on. Moreover, since $\left(\widehat{b}_t^\theta, \widehat{w}_t^\theta\right)$ differs from $\left(b_t^\theta, w_t^\theta\right)$ only by the same constant b_t^θ on both coordinates, effort incentives also remain unchanged in period t. Notice, however, that both types' continuation payoff in period t has been reduced by b_t^θ.

Now consider period $p(t)$. If $p(t) > 0$ the agent of either type faces the following considerations under contract $\widehat{\mathbf{C}}^\theta$: (1) if a success is obtained, he obtains the bonus $\widehat{b}_t^\theta = b_t^\theta$ and the game is over;

(2) if a success is not obtained, there is a transfer $w^{\theta}_{p(t)} + \delta^{t-p(t)} b^{\theta}_t$ and the next active period is period t where the continuation payoff is b^{θ}_t lower than it is under contract $\mathbf{C}^{\theta}$; consequently, the discounted continuation payoff from failure is identical to that of the original contract $\mathbf{C}^{\theta}$. In sum, the incentives in period $p(t)$ remain unchanged for both types and, critically, both types have the same continuation payoff at the beginning of period $p(t)$ under either contract. By analogous reasoning, if $p(t) = 0$, then the time zero payoff to both types is the same under either contract.

This modification procedure can be applied to each period $t \in \Gamma^{\theta}$. We eventually obtain a contract with no bonus in any period $t \in \Gamma^{\theta}$, viz. a contract $\widehat{\mathbf{C}}^{\theta} = \left(\widehat{W}^{\theta}_0, \widehat{\mathbf{b}}^{\theta}, \widehat{\mathbf{w}}^{\theta}, \Gamma^{\theta}\right)$ defined as follows:

(i) In any t such that $t < \sup \Gamma^{\theta}$ and $t \in \Gamma^{\theta}$,

$$\begin{aligned} \widehat{b}^{\theta}_t &= 0, \\ \widehat{w}^{\theta}_t &= w^{\theta}_t - b^{\theta}_t + \delta^{s-t} b^{\theta}_s, \text{ where } s = \min \Gamma^{\theta} \backslash \{1, \ldots, t\}. \end{aligned}$$

(ii) If $T = \sup \Gamma^{\theta}$ is finite, then

$$\begin{aligned} \widehat{b}^{\theta}_T &= 0, \\ \widehat{w}^{\theta}_T &= w^{\theta}_T - b^{\theta}_T, \end{aligned}$$

(iii) $\widehat{W}^{\theta}_0 = W^{\theta}_0 + \delta^s b^{\theta}_s$, where $s = \min \Gamma^{\theta}$.

It follows from the construction that this new contract $\widehat{\mathbf{C}}^{\theta}$ has the properties stated in the claim.

A.3 Step 3

By Steps 1 and 2, we can restrict our attention to clawback contracts $\mathbf{C}^{\theta}_w = \left(W^{\theta}_0, \mathbf{w}^{\theta}, \Gamma^{\theta}\right)$, with the L type's contract inducing the L type to exert effort in all periods in Γ^L. Denoting the set of clawback contracts by $\mathcal{C}_w$, the principal faces the following program [P]:

$$\max_{\left(\mathbf{C}^H_w \in \mathcal{C}_w, \mathbf{C}^L_w \in \mathcal{C}_w, \mathbf{a}^H\right)} \mu_0 \Pi^H_0 \left(\mathbf{C}^H_w, \mathbf{a}^H\right) + (1 - \mu_0) \Pi^L_0 \left(\mathbf{C}^L_w, \mathbf{1}\right) \tag{P}$$

subject to

$$\mathbf{1} \in \underset{\mathbf{a}=(a_t)_{t \in \Gamma^L}}{\arg\max}\ U^L_0 \left(\mathbf{C}^L_w, \mathbf{a}\right) \tag{IC^L_a}$$

$$\mathbf{a}^H \in \underset{\mathbf{a}=(a_t)_{t \in \Gamma^H}}{\arg\max}\ U^H_0 \left(\mathbf{C}^H_w, \mathbf{a}\right) \tag{IC^H_a}$$

$$U^L_0 \left(\mathbf{C}^L_w, \mathbf{1}\right) \geq 0 \tag{IR^L}$$

$$U^H_0 \left(\mathbf{C}^H_w, \mathbf{a}^H\right) \geq 0 \tag{IR^H}$$

$$U^L_0 \left(\mathbf{C}^L_w, \mathbf{1}\right) \geq U^L_0 \left(\mathbf{C}^H_w, \boldsymbol{\alpha}^L \left(\mathbf{C}^H_w\right)\right) \tag{IC^{LH}}$$

$$U^H_0 \left(\mathbf{C}^H_w, \mathbf{a}^H\right) \geq U^H_0 \left(\mathbf{C}^L_w, \boldsymbol{\alpha}^H \left(\mathbf{C}^L_w\right)\right). \tag{IC^{HL}}$$

To solve program [P], we solve a relaxed program and later verify that the solution is feasible in (and hence is a solution to) [P]. Specifically, we relax three constraints in [P]: (i) we ignore IC^{LH} and

IR^H, and (ii) we consider a weak version of IC^{HL} in which the H type is assumed to exert effort in all periods in Γ^L if he chooses $\mathbf{C}_w^L$. The relaxed program, [RP1], is therefore:

$$\max_{\left(\mathbf{C}_w^H \in \mathcal{C}_w, \mathbf{C}_w^L \in \mathcal{C}_w, \mathbf{a}^H\right)} \mu_0 \Pi_0^H \left(\mathbf{C}_w^H, \mathbf{a}^H\right) + (1-\mu_0) \Pi_0^L \left(\mathbf{C}_w^L, \mathbf{1}\right) \tag{RP1}$$

subject to

$$\mathbf{1} \in \underset{\mathbf{a}=(a_t)_{t\in\Gamma^L}}{\arg\max}\ U_0^L\left(\mathbf{C}_w^L, \mathbf{a}\right) \tag{IC$_a^L$}$$

$$\mathbf{a}^H \in \underset{\mathbf{a}=(a_t)_{t\in\Gamma^H}}{\arg\max}\ U_0^H\left(\mathbf{C}_w^H, \mathbf{a}\right) \tag{IC$_a^H$}$$

$$U_0^L\left(\mathbf{C}_w^L, \mathbf{1}\right) \geq 0 \tag{IR^L}$$

$$U_0^H\left(\mathbf{C}_w^H, \mathbf{a}^H\right) \geq U_0^H\left(\mathbf{C}_w^L, \mathbf{1}\right). \tag{Weak-IC^{HL}}$$

It is clear that in any solution to program [RP1], (IR^L) must be binding: otherwise, the initial time-zero transfer from the principal to the agent in the contract $\mathbf{C}_w^L$ can be reduced slightly to strictly improve the second term of the objective function while not violating any of the constraints. Similarly, (Weak-IC^{HL}) must also bind because otherwise the time-zero transfer in the contract $\mathbf{C}_w^H$ can be reduced to improve the first term of the objective function without violating any of the constraints.

Using these two binding constraints and substituting in the formulae from equations (1) and (2), we can rewrite the objective function (RP1) as the sum of expected total surplus and type H's "information rent", obtaining the following explicit version of the relaxed program which we call [RP2]:

$$\max_{\left(\mathbf{C}_w^H \in \mathcal{C}_w, \mathbf{C}_w^L \in \mathcal{C}_w, \mathbf{a}^H\right)} \left\{ \begin{array}{l} \mu_0 \left[\beta_0 \sum_{t\in\Gamma^H} \delta^t \left[\prod_{s\in\Gamma^H, s\leq t-1} \left(1-a_s^H\lambda^H\right)\right] a_t^H\left(\lambda^H - c\right) - (1-\beta_0) \sum_{t\in\Gamma^H} \delta^t a_t^H c\right] \\ + (1-\mu_0)\left[\beta_0 \sum_{t\in\Gamma^L} \delta^t \left[\prod_{s\in\Gamma^L, s\leq t-1}\left(1-\lambda^L\right)\right]\left(\lambda^L - c\right) - (1-\beta_0)\sum_{t\in\Gamma^L}\delta^t c\right] \\ \underbrace{- \mu_0\beta_0 \left\{ \begin{array}{l} \sum_{t\in\Gamma^L} \delta^t w_t^L \left[\prod_{s\in\Gamma^L, s\leq t}\left(1-\lambda^H\right) - \prod_{s\in\Gamma^L, s\leq t}\left(1-\lambda^L\right)\right] \\ - \sum_{t\in\Gamma^L}\delta^t c \left[\prod_{s\in\Gamma^L, s\leq t-1}\left(1-\lambda^H\right) - \prod_{s\in\Gamma^L, s\leq t-1}\left(1-\lambda^L\right)\right] \end{array}\right\}}_{\text{Information rent of type H}} \end{array} \right\} \tag{RP2}$$

subject to

$$\mathbf{1} \in \underset{(a_t)_{t\in\Gamma^L}}{\arg\max} \left\{ \begin{array}{c} \beta_0 \sum_{t\in\Gamma^L} \delta^t \left[\prod_{s\in\Gamma^L, s\leq t-1}\left(1-a_s\lambda^L\right)\right]\left[\left(1-a_t\lambda^L\right) w_t^L - a_t c\right] \\ +(1-\beta_0)\sum_{t\in\Gamma^L}\delta^t\left(w_t^L - a_t c\right) + W_0^L \end{array}\right\}, \tag{IC$_a^L$}$$

$$\mathbf{a}^H \in \underset{(a_t)_{t\in\Gamma^H}}{\arg\max} \left\{ \begin{array}{c} \beta_0 \sum_{t\in\Gamma^H} \delta^t \left[\prod_{s\in\Gamma^H, s\leq t-1}\left(1-a_s\lambda^H\right)\right]\left[\left(1-a_t\lambda^H\right) w_t^H - a_t c\right] \\ +(1-\beta_0)\sum_{t\in\Gamma^H}\delta^t\left(w_t^H - a_t c\right) + W_0^H \end{array}\right\}. \tag{IC$_a^H$}$$

A key observation is that this program [RP2] is separable, i.e. it can be solved by separately maximizing (RP2) with respect to $\mathbf{C}_w^L$ subject to (IC_a^L) and separately maximizing (RP2) with respect to $(\mathbf{C}_w^H, \mathbf{a}^H)$ subject to (IC_a^H).

A.4 Step 4

We now claim that in program [RP2], it is without loss to consider solutions in which the low type's contract is a connected clawback contract, i.e. solutions $\mathbf{C}_w^L$ in which $\Gamma^L = \{1, ..., T^L\}$ for some T^L.

To prove this, observe that the optimal $\mathbf{C}_w^L$ is a solution of

$$\max_{\mathbf{C}_w^L} \left\{ \begin{array}{c} (1-\mu_0)\left[\beta_0 \sum_{t\in\Gamma^L} \delta^t \left[\prod_{s\in\Gamma^L, s\le t-1} (1-\lambda^L)\right](\lambda^L - c) - (1-\beta_0)\sum_{t\in\Gamma^L}\delta^t c\right] \\ -\mu_0\beta_0 \left\{ \begin{array}{c} \sum_{t\in\Gamma^L} \delta^t w_t^L \left[\prod_{s\in\Gamma^L, s\le t}(1-\lambda^H) - \prod_{s\in\Gamma^L, s\le t}(1-\lambda^L)\right] \\ -\sum_{t\in\Gamma^L}\delta^t c \left[\prod_{s\in\Gamma^L, s\le t-1}(1-\lambda^H) - \prod_{s\in\Gamma^L, s\le t-1}(1-\lambda^L)\right] \end{array} \right\} \end{array} \right\} \tag{12}$$

subject to the (IC_a^L),

$$(1)_{t\in\Gamma^L} \in \arg\max_{(a_t)_{t\in\Gamma^L}} \left\{ \begin{array}{c} \beta_0 \sum_{t\in\Gamma^L}\delta^t \left[\prod_{s\in\Gamma^L, s\le t-1}(1-a_s\lambda^L)\right]\left[(1-a_t\lambda^L)\, w_t^L - a_t c\right] \\ +(1-\beta_0)\sum_{t\in\Gamma^L}\delta^t\left(w_t^L - a_t c\right) + W_0^L \end{array} \right\}. \tag{13}$$

To avoid trivialities, consider any optimal $\mathbf{C}_w^L$ with $\Gamma^L \neq \emptyset$. First consider the possibility $1 \notin \Gamma^L$. In this case, construct a new clawback contract $\widehat{\mathbf{C}}_w^L$ that is "shifted up by one period":

$$\begin{aligned} \widehat{\Gamma}^L &= \{s : s+1 \in \Gamma^L\}, \\ \widehat{w}_s^L &= w_{s+1}^L \text{ for all } s \in \widehat{\Gamma}^L, \\ \widehat{W}_0^L &= W_0^L. \end{aligned}$$

Clearly it remains optimal for the agent to work in every period in $\widehat{\Gamma}^L$, and since the value of (12) must have been weakly positive under $\mathbf{C}_w^L$, it is now weakly higher since the modification has just multiplied it by $\delta^{-1} > 1$. This procedure can be repeated for all lockout periods at the beginning of the contract, so that without loss, we hereafter assume that $1 \in \Gamma^L$. We are of course done if Γ^L is now connected, so also assume that Γ^L is not connected.

Let t^* be the earliest lockout period in Γ^L, i.e. $t^* = \min\{t : t \notin \Gamma^L \text{ and } t^* - 1 \in \Gamma^L\}$. (Such a $t^* > 1$ exists given the preceding discussion.) We will argue that one of two possible modifications preserves the agent's incentive to work in all periods in the modified contract and weakly improves the principal's payoff. This suffices because the procedure can then be applied iteratively to produce a connected contract.

<u>Modification 1</u>: Consider first a modified clawback contract $\widehat{\mathbf{C}}_w^L$ that removes the lockout period

t^* and shortens the contract by one period as follows:

$$\begin{aligned}
\widehat{\Gamma}^L &= \{1,\ldots,t^*-1\} \cup \{s : s \geq t^* \text{ and } s+1 \in \Gamma^L\}, \\
\widehat{w}_s^L &= \begin{cases} w_s^L & \text{if } s < t^*-1, \\ w_s^L + \Delta_1 & \text{if } s = t^*-1, \\ w_{s+1}^L & \text{if } s \geq t^* \text{ and } s \in \widehat{\Gamma}^L, \end{cases} \\
\widehat{W}_0^L &= W_0^L.
\end{aligned}$$

Note that in the above construction, Δ_1 is a free parameter. We will find conditions on Δ_1 such that the L type's incentives for effort are unchanged and the principal is weakly better off.

For an arbitrary t, define

$$\begin{aligned}
S(t) &= \left[\prod_{s\in\Gamma^L, s\leq t-1} \left(1-\lambda^L\right)\right]\left(\lambda^L - c\right), \\
R(t) &= \prod_{s\in\Gamma^L, s\leq t} \left(1-\lambda^H\right) - \prod_{s\in\Gamma^L, s\leq t} \left(1-\lambda^L\right).
\end{aligned}$$

The value of (12) under $\mathbf{C}_w^L$ is

$$V(\mathbf{C}_w^L) = (1-\mu_0)\left[\beta_0 \sum_{t\in\Gamma^L} \delta^t S(t) - (1-\beta_0)\sum_{t\in\Gamma^L}\delta^t c\right] - \mu_0\beta_0\left[\sum_{t\in\Gamma^L}\delta^t w_t^L R(t) - \sum_{t\in\Gamma^L}\delta^t c R(t-1)\right].$$

The value of (12) after the modification to $\widehat{\mathbf{C}}_w^L$ is

$$\begin{aligned}
V(\widehat{\mathbf{C}}_w^L) &= (1-\mu_0)\left[\begin{array}{c} \beta_0\left(\sum_{t\in\Gamma^L, t<t^*}\delta^t S(t) + \delta^{-1}\sum_{t\in\Gamma^L, t>t^*}\delta^t S(t)\right) \\ -(1-\beta_0)\left(\sum_{t\in\Gamma^L, t<t^*}\delta^t c + \delta^{-1}\sum_{t\in\Gamma^L, t>t^*}\delta^t c\right)\end{array}\right] \\
&\quad -\mu_0\beta_0\left[\begin{array}{c} \sum_{t\in\Gamma^L, t<t^*-1}\delta^t w_t^L R(t) + \delta^{-1}\sum_{t\in\Gamma^L, t>t^*}\delta^t w_t^L R(t) \\ +\delta^{t^*-1}\left(w_{t^*-1}^L + \Delta_1\right)R(t^*-1) \\ -\sum_{t\in\Gamma^L, t<t^*}\delta^t c R(t-1) - \delta^{-1}\sum_{t\in\Gamma^L, t>t^*}\delta^t c R(t-1)\end{array}\right].
\end{aligned}$$

Therefore, the modification benefits the principal if and only if

$$\begin{aligned}
0 \leq V(\widehat{\mathbf{C}}_w^L) - V(\mathbf{C}_w^L) &= (1-\mu_0)\left[\beta_0\left(\delta^{-1}-1\right)\sum_{t\in\Gamma^L, t>t^*}\delta^t S(t) - (1-\beta_0)\left(\delta^{-1}-1\right)\sum_{t\in\Gamma^L, t>t^*}\delta^t c\right] \\
&\quad -\mu_0\beta_0\left[\begin{array}{c}\left(\delta^{-1}-1\right)\sum_{t\in\Gamma^L, t>t^*}\delta^t w_t^L R(t) + \delta^{t^*-1}\Delta_1 R(t^*-1) \\ -\left(\delta^{-1}-1\right)\sum_{t\in\Gamma^L, t>t^*}\delta^t c R(t-1)\end{array}\right],
\end{aligned}$$

or equivalently after rearranging terms, if and only if

$$
\begin{aligned}
& (1-\mu_0)\left[\beta_0 \sum_{t\in\Gamma^L, t>t^*} \delta^t S(t) - (1-\beta_0) \sum_{t\in\Gamma^L, t>t^*} \delta^t c\right] \\
\geq\ & \mu_0\beta_0 \left[\ \sum_{t\in\Gamma^L, t>t^*} \delta^t w_t^L R(t) - \delta^{t^*-1} \tfrac{\Delta_1}{1-\delta^{-1}} R(t^*-1) - \sum_{t\in\Gamma^L, t>t^*} \delta^t c R(t-1)\ \right].
\end{aligned} \tag{14}
$$

Now turn to the incentives for effort for the agent of type L. Clearly, since $\mathbf{C}_w^L$ induces the agent to work in all periods, it remains optimal for the agent to work under $\widehat{\mathbf{C}}_w^L$ in all periods beginning with t^*. Consider the incentive constraint for effort in period t^*-1 under $\widehat{\mathbf{C}}_w^L$. Using (13), this is given by:

$$
-\overline{\beta}_{t^*-1}^L \lambda^L \left\{ w_{t^*-1}^L + \Delta_1 + \delta^{-1} \sum_{t\in\Gamma^L, t>t^*} \delta^{t-(t^*-1)} \left[\prod_{s\in\Gamma^L, t^*-1<s\leq t-1} (1-\lambda^L) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \right\} \geq c. \tag{15}
$$

Analogously, the incentive constraint in period t^*-1 under the original contract $\mathbf{C}_w^L$ is:

$$
-\overline{\beta}_{t^*-1}^L \lambda^L \left\{ w_{t^*-1}^L + \sum_{t\in\Gamma^L, t>t^*-1} \delta^{t-(t^*-1)} \left[\prod_{s\in\Gamma^L, t^*-1<s\leq t-1} (1-\lambda^L) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \right\} \geq c. \tag{16}
$$

If we choose Δ_1 such that the left-hand side of (15) is equal to the left-hand side of (16), then since it is optimal to work under the original contract in period t^*-1, it will also be optimal to work under the new contract in period t^*-1. Accordingly, we choose Δ_1 such that:

$$
\begin{aligned}
\Delta_1 &= \sum_{t\in\Gamma^L, t>t^*-1} \delta^{t-(t^*-1)} \left[\prod_{s\in\Gamma^L, t^*-1<s\leq t-1} (1-\lambda^L) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \\
&\quad -\delta^{-1} \sum_{t\in\Gamma^L, t>t^*} \delta^{t-(t^*-1)} \left[\prod_{s\in\Gamma^L, t^*-1<s\leq t-1} (1-\lambda^L) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \\
&= (1-\delta^{-1}) \sum_{t\in\Gamma^L, t>t^*-1} \delta^{t-(t^*-1)} \left[\prod_{s\in\Gamma^L, t^*-1<s\leq t-1} (1-\lambda^L) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right],
\end{aligned} \tag{17}
$$

where the second equality is because $\{t : t\in\Gamma^L, t>t^*-1\} = \{t : t\in\Gamma^L, t>t^*\}$, since $t^*\notin\Gamma^L$.

Now consider the incentive constraint for effort in any period $\tau < t^*-1$. We will show that because Δ_1 is such that the left-hand side of (15) is equal to the left-hand side of (16), the fact that it was optimal to work in period τ under contract $\mathbf{C}_w^L$ implies that it is optimal to work in period τ under contract $\widehat{\mathbf{C}}_w^L$. Formally, the incentive constraint for effort in period τ under $\mathbf{C}_w^L$ is

$$
-\overline{\beta}_{\tau}^L \lambda^L \left\{ w_{\tau}^L + \sum_{t\in\Gamma^L, t>\tau} \delta^{t-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t-1} (1-\lambda^L) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \right\} \geq c, \tag{18}
$$

which is satisfied since $\mathbf{C}_w^L$ induces the agent to work in all periods. The incentive constraint for effort in period τ under $\widehat{\mathbf{C}}_w^L$ can be written as:

$$\begin{aligned}
c &\leq -\overline{\beta}_\tau^L \lambda^L \left\{ \widehat{w}_\tau^L + \sum_{t\in\widehat{\Gamma}^L, t>\tau} \delta^{t-\tau} \left[\prod_{s\in\widehat{\Gamma}^L, \tau<s\leq t-1} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) \widehat{w}_t^L - c \right] \right\} \\
&= -\overline{\beta}_\tau^L \lambda^L \left\{ \begin{array}{c} w_\tau^L + \sum_{t\in\Gamma^L, \tau<t<t^*-1} \delta^{t-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t-1} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \\ +\delta^{t^*-1-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t^*-2} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) \left(w_{t^*-1}^L + \Delta_1\right) - c \right] \\ +\delta^{-1} \sum_{t\in\Gamma^L, t>t^*-1} \delta^{t-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t-1} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \end{array} \right\} \\
&= -\overline{\beta}_\tau^L \lambda^L \left\{ \begin{array}{c} w_\tau^L + \sum_{t\in\Gamma^L, \tau<t<t^*-1} \delta^{t-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t-1} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \\ +\delta^{t^*-1-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t^*-2} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) w_{t^*-1}^L - c \right] \\ +\left(1-\delta^{-1}\right) \left[\sum_{t\in\Gamma^L, t>t^*-1} \delta^{t-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t-1} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \right] \\ +\delta^{-1} \sum_{t\in\Gamma^L, t>t^*-1} \delta^{t-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t-1} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \end{array} \right\} \\
&= -\overline{\beta}_\tau^L \lambda^L \left\{ w_\tau^L + \sum_{t\in\Gamma^L, t>\tau} \delta^{t-\tau} \left[\prod_{s\in\Gamma^L, \tau<s\leq t-1} \left(1-\lambda^L\right) \right] \left[\left(1-\lambda^L\right) w_t^L - c \right] \right\}
\end{aligned}$$

where the first equality is from the construction of $\widehat{\mathbf{C}}_w^L$, the second equality uses (17), and the third equality follows from algebraic simplification. Since the above constraint is identical to (18), it is satisfied.

Modification 2: Now we consider a modified contract $\tilde{\mathbf{C}}_w^L$ that eliminates all periods after t^*, defined as follows:

$$\begin{aligned}
\tilde{\Gamma}^L &= \{1, \ldots, t^*-1\}, \\
\tilde{w}_s^L &= \begin{cases} w_s^L & \text{if } s < t^*-1, \\ w_s^L + \Delta_2 & \text{if } s = t^*-1, \end{cases} \\
\tilde{W}_0^L &= W_0^L.
\end{aligned}$$

Again, Δ_2 is a free parameter above. We find conditions on Δ_2 such that the L type's incentives are unchanged and the principal is weakly better off.

The value of (12) under the modification $\tilde{\mathbf{C}}_w^L$ is

$$\begin{aligned}
V(\tilde{\mathbf{C}}_w^L) &= (1-\mu_0)\left[\beta_0 \sum_{t\in\Gamma^L, t<t^*} \delta^t S(t) - (1-\beta_0) \sum_{t\in\Gamma^L, t<t^*} \delta^t c\right] \\
&\quad -\mu_0\beta_0 \left[\sum_{t\in\Gamma^L, t<t^*-1} \delta^t w_t^L R(t) + \delta^{t^*-1}\left(w_{t^*-1}^L + \Delta_2\right) R(t^*-1) - \sum_{t\in\Gamma^L, t<t^*} \delta^t c R(t-1) \right].
\end{aligned}$$

Therefore, using the previous formula for $V(\mathbf{C}_w^L)$, this modification benefits the principal if and only if

$$\begin{aligned}
0 \le V(\tilde{\mathbf{C}}_w^L) - V(\mathbf{C}_w^L) &= -(1-\mu_0)\left[\beta_0 \sum_{t\in\Gamma^L, t>t^*} \delta^t S(t) - (1-\beta_0) \sum_{t\in\Gamma^L, t>t^*} \delta^t c\right] \\
&\quad -\mu_0\beta_0 \left[-\sum_{t\in\Gamma^L, t>t^*} \delta^t w_t^L R(t) + \delta^{t^*-1}\Delta_2 R(t^*-1) + \sum_{t\in\Gamma^L, t>t^*} \delta^t c R(t-1) \right],
\end{aligned}$$

or equivalently after rearranging terms, if and only if

$$\begin{aligned}
& (1-\mu_0)\left[\beta_0 \sum_{t\in\Gamma^L, t>t^*} \delta^t S(t) - (1-\beta_0) \sum_{t\in\Gamma^L, t>t^*} \delta^t c\right] \\
\le\ & \mu_0\beta_0 \left[\sum_{t\in\Gamma^L, t>t^*} \delta^t w_t^L R(t) - \delta^{t^*-1}\Delta_2 R(t^*-1) - \sum_{t\in\Gamma^L, t>t^*} \delta^t c R(t-1) \right]. \qquad (19)
\end{aligned}$$

As with the previous modification, the only incentive constraint for effort that needs to be verified in $\tilde{\mathbf{C}}_w^L$ is that of period t^*-1, which since it is the last period of the contract is simply:

$$-\overline{\beta}_{t^*-1}^L \lambda^L \left[w_{t^*-1}^L + \Delta_2\right] \ge c. \qquad (20)$$

We choose Δ_2 so that the left-hand side of (20) is equal to the left-hand side of (16):

$$\Delta_2 = \sum_{t\in\Gamma^L, t>t^*-1} \delta^{t-(t^*-1)} \left[\prod_{s\in\Gamma^L, t^*-1<s\le t-1} \left(1-\lambda^L\right)\right] \left[\left(1-\lambda^L\right) w_t^L - c\right] = \frac{\Delta_1}{1-\delta^{-1}}, \qquad (21)$$

where the second equality follows from (17). But now, observe that (21) implies that either (14) or (19) is guaranteed to hold, and hence either the modification to $\widehat{\mathbf{C}}_w^L$ or to $\tilde{\mathbf{C}}_w^L$ weakly benefits the principal while preserving the agent's effort incentives.

A.5 Step 5

Take any connected clawback contract, $\mathbf{C}_w^L = (W_0^L, \mathbf{w}^L, T^L)$ that induces effort from the low type in each period $t \in \{1, \ldots, T^L\}$. We claim that the low type's incentive constraint for effort binds at all periods if and only if $\mathbf{w}^L = \overline{\mathbf{w}}^L(T^L)$, where $\overline{\mathbf{w}}^L(T^L)$ is defined as follows:

$$\overline{w}_t^L = \begin{cases} -(1-\delta)\frac{c}{\overline{\beta}_t^L \lambda^L} & \text{if } t < T^L, \\ -\frac{c}{\overline{\beta}_{T^L}^L \lambda^L} & \text{if } t = T^L. \end{cases} \tag{22}$$

The proof of this claim is via three sub-steps; for the remainder of this step, since T^L is given and held fixed, we ease notation by just writing $\overline{\mathbf{w}}^L$ instead of $\overline{\mathbf{w}}^L(T^L)$.

Step 5a: First, we argue that with the above penalty sequence, the low type is indifferent between working and shirking in each period $t \in \{1, \ldots, T^L\}$ given that he has worked in all prior periods and will do in all subsequent periods no matter his action at period t. In other words, we need to show that for all $t \in \{1, \ldots, T^L\}$:

$$-\overline{\beta}_t^L \lambda^L \left\{ \overline{w}_t^L + \sum_{s=t+1}^{T^L} \delta^{s-t} \left(1-\lambda^L\right)^{s-(t+1)} \left[\left(1-\lambda^L\right) \overline{w}_s^L - c\right] \right\} = c.^{27} \tag{23}$$

We prove that (23) is indeed satisfied for all t via an induction argument. First, it is obviously true for $t = T^L$. Next, for any $t < T^L$, assume it holds for $t+1$. This is equivalent to

$$\sum_{s=t+2}^{T^L} \delta^{s-(t+1)} \left(1-\lambda^L\right)^{s-(t+2)} \left[\left(1-\lambda^L\right) \overline{w}_s^L - c\right] = -\frac{c}{\overline{\beta}_{t+1}^L \lambda^L} - \overline{w}_{t+1}^L. \tag{24}$$

To show that (23) holds for t, it suffices to show that

$$-\overline{\beta}_t^L \lambda^L \left\{ \overline{w}_t^L + \delta\left[\left(1-\lambda^L\right)\overline{w}_{t+1}^L - c\right] + \delta\left(1-\lambda^L\right) \sum_{s=t+2}^{T^L} \delta^{s-(t+1)} \left(1-\lambda^L\right)^{s-(t+2)} \left[\left(1-\lambda^L\right) \overline{w}_s^L - c\right] \right\} = c$$

Using (24), the above equality is equivalent to

$$-\overline{\beta}_t^L \lambda^L \left\{ \overline{w}_t^L + \delta\left[\left(1-\lambda^L\right)\overline{w}_{t+1}^L - c\right] + \delta\left(1-\lambda^L\right) \left[-\frac{c}{\overline{\beta}_{t+1}^L \lambda^L} - \overline{w}_{t+1}^L \right] \right\} = c,$$

[27] To derive this equality, observe that under the hypotheses, the payoff for type L from working at t is

$$-c + \left(1-\overline{\beta}_t^L\right) \sum_{s=t}^{T^L} \delta^{s-t} \overline{w}_s^L + \overline{\beta}_t^L \left[(1-\lambda^L)\overline{w}_t^L + \sum_{s=t+1}^{T^L} \delta^{s-t} \left(1-\lambda^L\right)^{s-t} \left[\left(1-\lambda^L\right)\overline{w}_s^L - c\right] \right],$$

while the payoff from shirking at time t is

$$\overline{w}_t^L + \left(1-\overline{\beta}_t^L\right) \sum_{s=t+1}^{T^L} \delta^{s-t} \overline{w}_s^L + \overline{\beta}_t^L \left[\sum_{s=t+1}^{T^L} \delta^{s-t} \left(1-\lambda^L\right)^{s-(t+1)} \left[\left(1-\lambda^L\right)\overline{w}_s^L - c\right] \right].$$

Setting these payoffs from working and shirking equal to each other and manipulating terms yields (23).

which simplifies to

$$\overline{w}_t^L = -\frac{c}{\overline{\beta}_t^L \lambda^L} + \delta c + \delta \left(1-\lambda^L\right) \frac{c}{\overline{\beta}_{t+1}^L \lambda^L}. \tag{25}$$

Since $\overline{\beta}_{t+1}^L = \frac{\overline{\beta}_t^L\left(1-\lambda^L\right)}{1-\overline{\beta}_t^L\lambda^L}$, (25) is in turn equivalent to

$$\overline{w}_t^L \quad = \quad -\left(1-\delta\right) \frac{c}{\overline{\beta}_t^L \lambda^L},$$

which is true by the definition of $\overline{\mathbf{w}}^L$.

Step 5b: Next, we show that given the sequence $\overline{\mathbf{w}}^L$, it would be optimal for the low type to work in any period no matter the prior history of effort. The argument is by induction. Consider first the last period, T^L. Since no matter the history of prior effort, the current belief is some $\beta_{T^L}^L \geq \overline{\beta}_{T^L}^L$, and hence

$$-\beta_{T^L}^L \lambda^L \overline{w}_t^L \geq \overline{\beta}_{T^L}^L \lambda^L \overline{w}_t^L = c,$$

it is optimal to work in the last period (note that the equality above is by definition).

Now assume inductively that the assertion is true for period $t+1 \leq T^L$, and consider period $t < T^L$ after any history of prior effort, with current belief β_t^L. Since we already showed that equation (23) holds, it follows from $\beta_t^L \geq \overline{\beta}_t^L$ that

$$-\beta_t^L \lambda^L \left\{ \overline{w}_t^L + \sum_{s=t+1}^{T^L} \delta^{s-t} \left(1-\lambda^L\right)^{s-(t+1)} \left[\left(1-\lambda^L\right) \overline{w}_s^L - c\right] \right\} \geq c,$$

and hence it is optimal for the agent to work in period t.

Step 5c: Finally, we argue that any profile of penalties, $\mathbf{w}^L$, that makes the low type's incentive constraint for effort bind at every period $t \in \{1, \ldots, T^L\}$ must coincide with $\overline{\mathbf{w}}^L$, given that the clawback contract must induce work from the low type in each period up to T^L. Again, we use induction. Since $\overline{w}_{T^L}^L$ is the unique penalty that makes the agent indifferent between working and shirking at period T^L given that he has worked in all prior periods, it follows that $w_{T^L}^L = \overline{w}_{T^L}^L$. Note from Step 5b that it would remain optimal for the agent to work in period T^L given any profile of effort in prior periods.

For the inductive step, pick some period $t < T^L$ and assume that in every period $x \in \{t, \ldots, T^L\}$, the agent is indifferent between working and shirking given that he has worked in all prior periods, and would also find it optimal to work at x following any other profile of effort prior to x. Under these hypotheses, the indifference at period $t+1$ implies that

$$-\overline{\beta}_{t+1}^L \lambda^L \left\{ w_{t+1}^L + \sum_{s=t+2}^{T^L} \delta^{s-(t+1)} \left(1-\lambda^L\right)^{s-(t+2)} \left[\left(1-\lambda^L\right) w_s^L - c\right] \right\} = c. \tag{26}$$

Given the inductive hypothesis, the incentive constraint for effort at period t is

$$-\overline{\beta}_t^L \lambda^L \left\{ w_t^L + \sum_{s=t+1}^{T^L} \delta^{s-t} \left(1-\lambda^L\right)^{s-(t+1)} \left[\left(1-\lambda^L\right) w_s^L - c\right] \right\} \geq c,$$

which, when set to bind, can be written as

$$-\overline{\beta}_t^L \lambda^L \left\{ w_t^L + \delta \left[\left(1-\lambda^L\right) w_{t+1}^L - c \right] + \delta \left(1-\lambda^L\right) \sum_{s=t+2}^{T^L} \delta^{s-(t+1)} \left(1-\lambda^L\right)^{s-(t+2)} \left[\left(1-\lambda^L\right) w_s^L - c \right] \right\} = c. \tag{27}$$

Substituting (26) into (27) , using the fact that $\overline{\beta}_{t+1}^L = \frac{\overline{\beta}_t^L(1-\lambda)}{\overline{\beta}_t^L(1-\lambda)+1-\overline{\beta}_t^L}$, and performing some algebra shows that $w_t^L = \overline{w}_t^L$. Moreover, by the reasoning in Step 5b, this also ensures that the agent would find it optimal to work in period t for any other history of actions prior to period t.

A.6 Step 6

By Step 4, we can restrict attention in solving program [RP2] to connected clawback contracts for the low type. For any T^L, Step 5 identified a particular sequence of penalties, $\overline{\mathbf{w}}^L(T^L)$. We now show that in solving [RP2], we can further restrict attention to the class of connected clawback contracts for the low type with precisely this penalty structure.

The proof involves two sub-steps; throughout, we hold an arbitrary T^L fixed and, to ease notation, drop the dependence of $\overline{\mathbf{w}}^L(\cdot)$ on T^L.

Step 6a: We begin by showing that given any connected clawback contract for the low type of length T^L that satisfies (IC_a^L), the value of [RP2] is weakly higher under a connected clawback contract with the same length that has the property that $w_t^L \leq \overline{w}_t^L(T^L)$ for all $t \in \{1, \dots, T^L\}$.

To show this, consider any connected clawback contract of length T^L that satisfies (IC_a^L) and specifies a penalty $w_{t'}^L > \overline{w}_{t'}^L$ in some period $t' \leq T^L$. We will prove that we can change the penalty structure by lowering $w_{t'}^L$ and raising some subsequent w_s^L for $s \in \{t'+1, \dots, T^L\}$ in a way that keeps type L's incentives for effort unchanged, and yet increase the value of the objective function (RP2). Define

$$\widehat{t} = \max\left\{t : t \leq T^L \text{ and } w_t^L > \overline{w}_t^L\right\}.$$

Observe that we must have $\widehat{t} < T^L$ because otherwise (IC_a^L) would be violated in period T^L. Furthermore, by definition of $\widehat{t}$, $w_t^L \leq \overline{w}_t^L$ for all $T^L \geq t > \widehat{t}$.

Claim: There exists $\tilde{t} \in \{\widehat{t}+1, \dots, T^L\}$ such that (IC_a^L) at $\tilde{t}$ is slack and $w_{\tilde{t}}^L < \overline{w}_{\tilde{t}}^L$.

Proof: Suppose not, then for each $T^L \geq t > \widehat{t}$, either $w_t^L = \overline{w}_t^L$, or $w_t^L < \overline{w}_t^L$ and (IC_a^L) binds. Then since whenever $w_t^L < \overline{w}_t^L$, (IC_a^L) binds by supposition, it must be that in all $t > \widehat{t}$, (IC_a^L) binds (this follows from Step 5). But then (IC_a^L) at $\widehat{t}$ is violated since $w_{\widehat{t}}^L > \overline{w}_{\widehat{t}}^L$. ‖

Claim: There exists $\overline{t} \in \{\widehat{t}+1, \dots, T^L\}$ such that $w_{\overline{t}}^L < \overline{w}_{\overline{t}}^L$ and for any $t \in \{\widehat{t}+1, ..., \overline{t}\}$, ($\text{IC}_a^L$) at t is slack. In particular, we can take $\overline{t}$ to be the first such period after $\widehat{t}$.

Proof: Fix $\tilde{t}$ in the previous claim. Note that (IC_a^L) at $\widehat{t}+1$ must be slack because otherwise (IC_a^L) at $\widehat{t}$ is violated by $w_{\widehat{t}}^L > \overline{w}_{\widehat{t}}^L$ and Step 5. There are two cases. (1) $w_{\widehat{t}+1}^L < \overline{w}_{\widehat{t}+1}^L$; then $\widehat{t}+1$ is the $\overline{t}$ we want. (2) $w_{\widehat{t}+1}^L = \overline{w}_{\widehat{t}+1}^L$ – in this case, since (IC_a^L) is slack at $\widehat{t}+1$, it must be that (IC_a^L) at $\widehat{t}+2$ is slack (otherwise, the claim in Step 5 is violated); now if $w_{\widehat{t}+2}^L < \overline{w}_{\widehat{t}+2}^L$, we are done because $\widehat{t}+2$ is the $\overline{t}$

we are looking for; if $w^L_{\widehat{t}+2} = \overline{w}^L_{\widehat{t}+2}$, then we continue to $\widehat{t}+3$... until we reach $\tilde{t}$ which we know give us a slack (IC^L_a), $w^L_{\tilde{t}} < w^L_{\tilde{t}}$, and we are sure that ($\text{IC}^L_a$) is slack in all periods of this process before reaching $\tilde{t}$. $\|$

Now we shall show that we can slightly reduce $w^L_{\widehat{t}} > \overline{w}^L_{\widehat{t}}$ and slightly increase $w^L_{\bar{t}} < \overline{w}^L_{\bar{t}}$ and meanwhile keep the incentives for effort of the L type satisfied for all periods. We know that we do not violate (IC^L_a) for $t \in \{\widehat{t}+1, ..., \bar{t}\}$ because (IC^L_a) is slack there. We shall show that the modification weakly reduces the H type's rent and meanwhile does not violate (IC^L_a) at $\widehat{t}$ nor any previous period. Therefore, the modified contract weakly dominates the original contract.

We first want to guarantee that (IC^L_a) at $\widehat{t}$ is unchanged. By the same reasoning as used in Step 4, the incentive constraint for effort in period $\widehat{t}$ (given that the agent will work in all subsequent periods no matter his behavior at period t) can be written as

$$-\beta^L_{\widehat{t}}\lambda^L \left\{ w^L_{\widehat{t}} + \sum_{t>\widehat{t}} \delta^{t-\widehat{t}} \left(1-\lambda^L\right)^{t-(\hat{t}+1)} \left[\left(1-\lambda^L\right) w^L_t - c\right] \right\} \geq c. \tag{28}$$

Observe that if we reduce $w^L_{\widehat{t}}$ by $\Delta > 0$ and increase $w^L_{\bar{t}}$ by $\frac{\Delta}{\delta^{\bar{t}-\widehat{t}}(1-\lambda^L)^{\bar{t}-(\widehat{t}+1)}(1-\lambda^L)} = \frac{\Delta}{\delta^{\bar{t}-\widehat{t}}(1-\lambda^L)^{\bar{t}-\widehat{t}}}$, then the left-hand side of (28) does not change. Moreover, it follows that incentives for effort at $t < \widehat{t}$ are also unchanged (see Step 4), and the incentive condition at $\bar{t}$ will be satisfied if Δ is small enough because the original (IC^L_a) at $\bar{t}$ is slack.

We now show that the modification above leads to a reduction of the rent of the H type in (RP2), i.e. raises the value of the objective. The rent is given by

$$\mu_0\beta_0 \left\{ \sum_{t=1}^{T^L} \delta^t w^L_t \left[\left(1-\lambda^H\right)^t - \left(1-\lambda^L\right)^t\right] - \sum_{t=1}^{T^L} \delta^t c \left[\left(1-\lambda^H\right)^{t-1} - \left(1-\lambda^L\right)^{t-1}\right] \right\}.$$

Hence, the change in the rent from reducing $w^L_{\widehat{t}}$ by Δ and increasing $w^L_{\bar{t}}$ by $\frac{\Delta}{\delta^{\bar{t}-\widehat{t}}(1-\lambda^L)^{\bar{t}-\widehat{t}}}$ is

$$\begin{aligned} & \mu_0\beta_0\delta^{\widehat{t}}\Delta \left\{ -\left[\left(1-\lambda^H\right)^{\widehat{t}} - \left(1-\lambda^L\right)^{\widehat{t}}\right] + \left(\frac{1}{\left(1-\lambda^L\right)^{\bar{t}-\widehat{t}}}\right) \left[\left(1-\lambda^H\right)^{\bar{t}} - \left(1-\lambda^L\right)^{\bar{t}}\right] \right\} \\ = \; & \mu_0\beta_0\delta^{\widehat{t}}\Delta \left[\frac{\left(1-\lambda^H\right)^{\widehat{t}}}{\left(1-\lambda^L\right)^{\bar{t}-\widehat{t}}}\right] \left[\left(1-\lambda^H\right)^{\bar{t}-\widehat{t}} - \left(1-\lambda^L\right)^{\bar{t}-\widehat{t}}\right] < 0, \end{aligned}$$

where the inequality is because $\bar{t} - \widehat{t} > 0$ and $1-\lambda^H < 1-\lambda^L$.

Step 6b: Now we show that unless the penalty sequence for the low type is exactly $\overline{\mathbf{w}}^L$, the value of the objective (RP2) can be improved while satisfying the incentive constraint for effort, (IC^L_a).

To show this, recall that the H type's rent is

$$\beta_0 \sum_{t=1}^{T^L} \delta^t w^L_t \left[\left(1-\lambda^H\right)^t - \left(1-\lambda^L\right)^t\right] - \beta_0 c \sum_{t=1}^{T^L} \delta^t \left[\left(1-\lambda^H\right)^{t-1} - \left(1-\lambda^L\right)^{t-1}\right].$$

By Step 6a, (IC_a^L) is satisfied in all periods $t = 1, \ldots, T^L$ whenever $w_t^L = \overline{w}_t^L$. Now, if $w_t^L < \overline{w}_t^L$ for some periods, we can replace w_t^L by $\overline{w}_t^L$ without affecting the effort incentives of the L type, and by doing this we reduce the rent of the H type, thereby raising the value of (RP2).

A.7 Step 7

By Step 6, an optimal contract for the low type that solves program [RP2] can be found by optimizing over T^L, i.e. the length of connected clawback contracts with the penalty structure $\overline{\mathbf{w}}^L(T^L)$. We now argue that the optimal length is no larger than t^L (recall that t^L is the first-best stopping time).

The portion of the objective (RP2) that involves T^L is

$$(1-\mu_0)\left[\beta_0 \sum_{t=1}^{T^L} \delta^t \left(1-\lambda^L\right)^{t-1} \left(\lambda^L - c\right) - (1-\beta_0) \sum_{t=1}^{T^L} \delta^t c\right]$$
$$-\mu_0\beta_0 \left\{\sum_{t=1}^{T^L} \delta^t \overline{w}_t^L(T^L) \left[\left(1-\lambda^H\right)^t - \left(1-\lambda^L\right)^t\right] - \sum_{t=1}^{T^L} \delta^t c \left[\left(1-\lambda^H\right)^{t-1} - \left(1-\lambda^L\right)^{t-1}\right]\right\},$$

where we have used the desired penalty sequence. Now consider the following definition:

$$\begin{aligned}
\Pi\left(z, T^L\right) \quad &= \quad z\left(-\mu_0\beta_0\right)\left\{\sum_{t=1}^{T^L} \delta^t \overline{w}_t^L(T^L) \left[\left(1-\lambda^H\right)^t - \left(1-\lambda^L\right)^t\right] - \sum_{t=1}^{T^L} \delta^t c \left[\left(1-\lambda^H\right)^{t-1} - \left(1-\lambda^L\right)^{t-1}\right]\right\} \\
&\quad + (1-\mu_0)\left[\beta_0 \sum_{t=1}^{T^L} \delta^t \left(1-\lambda^L\right)^{t-1} \left(\lambda^L - c\right) - (1-\beta_0) \sum_{t=1}^{T^L} \delta^t c\right].
\end{aligned}$$

If $z = 0$, the expression above corresponds to surplus maximization; if $z = 1$, the expression corresponds to the principal's optimization problem. Consider the term multiplied by z modulo a negative constant:

$$K\left(T^L\right) = \sum_{t=1}^{T^L} \delta^t \overline{w}_t^L(T^L) \left[\left(1-\lambda^H\right)^t - \left(1-\lambda^L\right)^t\right] - \sum_{t=1}^{T^L} \delta^t c \left[\left(1-\lambda^H\right)^{t-1} - \left(1-\lambda^L\right)^{t-1}\right].$$

If $K(\cdot)$ is shown to be increasing, then $\Pi(z, T^L)$ has decreasing differences, which implies that the optimal T^L when $z = 0$ is no smaller than the optimal T^L when $z = 1$, as desired. To see that $K(\cdot)$ is indeed increasing, observe that

$$\begin{aligned}
K(T+1)-K(T) &= \sum_{t=1}^{T+1}\delta^t \overline{w}_t^L(T+1)\left[\left(1-\lambda^H\right)^t-\left(1-\lambda^L\right)^t\right]-\sum_{t=1}^{T+1}\delta^t c\left[\left(1-\lambda^H\right)^{t-1}-\left(1-\lambda^L\right)^{t-1}\right] \\
&\quad -\sum_{t=1}^{T}\delta^t \overline{w}_t^L(T)\left[\left(1-\lambda^H\right)^t-\left(1-\lambda^L\right)^t\right]+\sum_{t=1}^{T}\delta^t c\left[\left(1-\lambda^H\right)^{t-1}-\left(1-\lambda^L\right)^{t-1}\right] \\
&= \delta^T \overline{w}_T^L(T+1)\left[\left(1-\lambda^H\right)^T-\left(1-\lambda^L\right)^T\right] \\
&\quad +\delta^{T+1}\overline{w}_{T+1}^L(T+1)\left[\left(1-\lambda^H\right)^{T+1}-\left(1-\lambda^L\right)^{T+1}\right] \\
&\quad -\delta^{T+1}c\left[\left(1-\lambda^H\right)^T-\left(1-\lambda^L\right)^T\right]-\delta^T\overline{w}_T^L(T)\left[\left(1-\lambda^H\right)^T-\left(1-\lambda^L\right)^T\right] \\
&= \delta^T\left[\frac{-(1-\delta)c}{\overline{\beta}_T^L\lambda^L}\right]\left[\left(1-\lambda^H\right)^T-\left(1-\lambda^L\right)^T\right] \\
&\quad +\delta^{T+1}\left[\frac{-c}{\overline{\beta}_{T+1}^L\lambda^L}\right]\left[\left(1-\lambda^H\right)^{T+1}-\left(1-\lambda^L\right)^{T+1}\right] \\
&\quad -\delta^{T+1}c\left[\left(1-\lambda^H\right)^T-\left(1-\lambda^L\right)^T\right]-\delta^T\left[\frac{-c}{\overline{\beta}_T^L\lambda^L}\right]\left[\left(1-\lambda^H\right)^T-\left(1-\lambda^L\right)^T\right] \\
&= \delta^{T+1}c\left(\frac{1}{\overline{\beta}_T^L\lambda^L}-1\right)\left(\frac{\lambda^H-\lambda^L}{1-\lambda^L}\right)\left(1-\lambda^H\right)^T \\
&> 0,
\end{aligned}$$

where the third equality uses the definition of $\overline{\mathbf{w}}^L(\cdot)$ and the final inequality is because $\overline{\beta}_T^L\lambda^L<1$.

Note also that it is clear that there is generically a unique T^L that maximizes $\Pi(1,T^L)$; hereafter we denote this solution $\bar{t}^L$.

A.8 Step 8

We have shown so far that there is a solution to program [RP2] in which the low type's contract is a connected clawback contract of length $\bar{t}^L \le t^L$ and in which the penalty sequence is given by $\overline{\mathbf{w}}^L(\bar{t}^L)$. In terms of optimizing over the high type's contract, note that any solution must induce the high type to work in each period up to t^H and no longer: this follows from the fact that the objective in (RP2) involving the high type's contract is social surplus from the high type, and that there is clearly a sequence of (sufficiently low) penalties $\mathbf{w}^H$ to ensure that (IC_a^H) is satisfied.

Recall that solutions to [RP2] produce solutions to [RP1] by choosing W_0^L to make (IR^L) bind and W_0^H to make (Weak-IC^{HL}) bind, which can always be done. Accordingly, let $\overline{\mathbf{C}}_w^L=(\overline{W}_0^L,\overline{\mathbf{w}}^L(\bar{t}^L),\bar{t}^L)$ be the connected clawback contract where $\overline{W}_0^L$ is set to make (IR^L) bind, and consider the solutions to program [RP1] in which the low type's contract is $\overline{\mathbf{C}}_w^L$. We will argue that some of these solutions to [RP1], namely $\overline{\mathbf{C}}_w^L$ combined with a suitable one-time-clawback contract for the high type, also solve the original program [P]. Recall that [RP1] differs from [P] in three ways:

1. it imposes (Weak-IC^{HL}) rather than (IC^{HL});

2. it ignores (IR^{H});

3. it ignores (IC^{LH}).

We address each of these constraints in order.

Step 8a: First, we argue that given any connected clawback contract of length $T^L \leq t^L$ with penalty sequence $\overline{\mathbf{w}}^L(T^L)$, it would be optimal for type H to work in every period $1\ldots,T^L$, no matter the history of prior effort. Consequently, any solution to [RP1] using $\overline{\mathbf{C}}_w^L$ satisfies (IC^{HL}).

To prove the claim, we fix any $T^L \leq t^L$ and write $\overline{w}^L$ as shorthand for $\overline{w}^L(T^L)$. The argument is by induction. Consider first the last period, T^L. Since

$$-\overline{\beta}_{T^L}^L \lambda^L \overline{w}_{T^L}^L = c,$$

it follows from the fact that $t^H > t^L$ (hence $\overline{\beta}_t^H \lambda^H > \overline{\beta}_t^L \lambda^L$ for all $t < t^H$) that no matter the history of effort,

$$-\beta_{T^L}^H \lambda^H \overline{w}_{T^L}^L \geq c,$$

i.e., regardless of the history, the H type will work in period T^L.

Now assume inductively that it is optimal for type H to work in period $t+1 \leq T^L$ no matter the history of effort, and consider period t with belief β_t^H. The inductive hypothesis implies that

$$-\beta_{t+1}^H \lambda^H \left\{ \overline{w}_{t+1}^L + \sum_{s=t+2}^{T^L} \delta^{s-(t+1)} \left(1-\lambda^H\right)^{s-(t+2)} \left[\left(1-\lambda^H\right) \overline{w}_s^L - c\right] \right\} \geq c,$$

or equivalently,

$$\sum_{s=t+2}^{T^L} \delta^{s-(t+1)} \left(1-\lambda^H\right)^{s-(t+2)} \left[\left(1-\lambda^H\right) \overline{w}_s^L - c\right] \leq -\frac{c}{\beta_{t+1}^H \lambda^H} - \overline{w}_{t+1}^L. \tag{29}$$

Therefore, at period t :

$$
\begin{aligned}
& -\beta_t^H\lambda^H\left\{\overline{w}_t^L+\delta\left[\left(1-\lambda^H\right)\overline{w}_{t+1}^L-c\right]+\delta\left(1-\lambda^H\right)\sum_{s=t+2}^{T^L}\delta^{s-(t+1)}\left(1-\lambda^H\right)^{s-(t+2)}\left[\left(1-\lambda^H\right)\overline{w}_s^L-c\right]\right\} \\
\geq\ & -\beta_t^H\lambda^H\left\{\overline{w}_t^L+\delta\left[\left(1-\lambda^H\right)\overline{w}_{t+1}^L-c\right]+\delta\left(1-\lambda^H\right)\left[-\frac{c}{\beta_{t+1}^H\lambda^H}-\overline{w}_{t+1}^L\right]\right\} \\
=\ & -\beta_t^H\lambda^H\left(\overline{w}_t^L-\delta c\right)+\delta\left(1-\lambda^H\right)\left(\frac{\beta_t^H c}{\beta_{t+1}^H}\right) \\
=\ & -\beta_t^H\lambda^H\overline{w}_t^L+\delta c \\
\geq\ & -\overline{\beta}_t^L\lambda^L\overline{w}_t^L+\delta c \\
=\ & -\overline{\beta}_t^L\lambda^L\left(-\left(1-\delta\right)\frac{c}{\overline{\beta}_t^L\lambda^L}\right)+\delta c \\
=\ & c,
\end{aligned}
$$

where the first inequality uses (29), the second equality uses $\beta_{t+1}^H=\frac{\beta_t^H(1-\lambda^H)}{1-\beta_t^H+\beta_t^H(1-\lambda^H)}$, and the penultimate equality uses the definition of $\overline{w}_t^L$.

Step 8b: Next, we show that any solution to [RP1] using $\overline{\mathbf{C}}_w^L$ also satisfies (IR^H). To show this, observe first that

$$
\begin{aligned}
U_0^H\left(\overline{\mathbf{C}}_w^L,\mathbf{1}\right) &= \beta_0\sum_{t=1}^{\bar{t}^L}\delta^t\left(1-\lambda^H\right)^{t-1}\left[\left(1-\lambda^H\right)\overline{w}_t^L(\bar{t}^L)-c\right]+(1-\beta_0)\sum_{t\in\Gamma}\delta^t\left(\overline{w}_t^L(\bar{t}^L)-c\right)+\overline{W}_0^L \\
&\geq \beta_0\sum_{t=1}^{\bar{t}^L}\delta^t\left(1-\lambda^L\right)^{t-1}\left[\left(1-\lambda^L\right)\overline{w}_t^L(\bar{t}^L)-c\right]+(1-\beta_0)\sum_{t\in\Gamma}\delta^t\left(\overline{w}_t^L(\bar{t}^L)-c\right)+\overline{W}_0^L \\
&= U_0^L\left(\overline{\mathbf{C}}_w^L,\mathbf{1}\right),
\end{aligned} \tag{30}
$$

where the inequality follows from the fact that for all $1\leq t\leq\bar{t}^L$, $\overline{w}_t^L\leq 0$.

Consequently, in any solution to [RP1] using $\overline{\mathbf{C}}_w^L$,

$$
\begin{aligned}
U_0^H\left(\mathbf{C}_w^H,\boldsymbol{\alpha}^H\left(\mathbf{C}_w^H\right)\right) &\geq U_0^H\left(\overline{\mathbf{C}}_w^L,\mathbf{1}\right)\ (\text{by IC}_a^H\text{ and Weak-IC}^{HL}) \\
&\geq U_0^L\left(\overline{\mathbf{C}}_w^L,\mathbf{1}\right)(\text{by inequality }30) \\
&\geq 0\,(\text{by IR}^L),
\end{aligned}
$$

and hence (IR^H) is satisfied.

Step 8c: Finally, we show that there is a solution to [RP1] using $\overline{\mathbf{C}}_w^L$ that also satisfies (IC^{LH}) in [P], which completes the proof. As previously noted, any optimal contract for the high type in [RP1] must induce effort from this type in periods $1,\ldots,t^H$ and make (Weak-IC^{HL}) bind. We will construct such a *onetime*-clawback contract, $\mathbf{C}_w^H = (W_0^H, w_{t^H}^H, t^H)$, where given the penalty $w_{t^H}^H$ (a free parameter at this point) and that the high type works in all periods, W_0^H is chosen to make (Weak-IC^{HL}) bind, i.e. by the equation:

$$\left[\left(1-\lambda^H\right)^{t^H}\beta_0 + (1-\beta_0)\right]\delta^{t^H} w_{t^H}^H - \beta_0 \sum_{t=1}^{t^H} \delta^t \left(1-\lambda^H\right)^{t-1} c - (1-\beta_0)\sum_{t=1}^{t^H}\delta^t c + W_0^H = \rho, \tag{31}$$

where

$$\rho = \beta_0 \sum_{t=1}^{\bar{t}^L} \delta^t \overline{w}_t^L(\bar{t}^L)\left[\left(1-\lambda^H\right)^t - \left(1-\lambda^L\right)^t\right] - \beta_0 c \sum_{t=1}^{\bar{t}^L}\delta^t\left[\left(1-\lambda^H\right)^{t-1} - \left(1-\lambda^L\right)^{t-1}\right]$$

is the rent earned by the H type given type L's contract $\overline{\mathbf{C}}_w^L$.

Plainly, the penalty $w_{t^H}^H$ can be chosen to be severe enough (i.e. sufficiently negative) to ensure that it is optimal for an agent of *either type*, H or L, to work in all periods after accepting such a contract $\mathbf{C}_w^H$, i.e. that for all $\theta \in \{H, L\}$, $\boldsymbol{\alpha}^\theta(\mathbf{C}_w^H) = \mathbf{1}$. All that remains is to show that a sufficiently severe $w_{t^H}^H$ and its corresponding W_0^H (determined by (31)) also satisfies (IC^{LH}) given that $\boldsymbol{\alpha}^L(\mathbf{C}_w^H) = \mathbf{1}$.

Using (31), we compute

$$\begin{aligned}
U_0^L\left(\mathbf{C}_w^H, \mathbf{1}\right) &= W_0^H - c\left[\beta_0\sum_{t=1}^{t^H}\delta^t\left(1-\lambda^L\right)^{t-1} + (1-\beta_0)\sum_{t=1}^{t^H}\delta^t\right] + w_{t^H}^H\delta^{t^H}\left[\beta_0\left(1-\lambda^L\right)^{t^H} + (1-\beta_0)\right] \\
&= -\left\{\left[\left(1-\lambda^H\right)^{t^H}\beta_0 + (1-\beta_0)\right]\delta^{t^H} w_{t^H}^H - \beta_0\sum_{t=1}^{t^H}\delta^t\left(1-\lambda^H\right)^{t-1} c - (1-\beta_0)\sum_{t=1}^{t^H}\delta^t c\right\} \\
&\quad + \rho - c\left[\beta_0\sum_{t=1}^{t^H}\delta^t\left(1-\lambda^L\right)^{t-1} + (1-\beta_0)\sum_{t=1}^{t^H}\delta^t\right] + w_{t^H}^H\delta^{t^H}\left[\beta_0\left(1-\lambda^L\right)^{t^H} + (1-\beta_0)\right] \\
&= \beta_0\delta^{t^H}\left[\left(1-\lambda^L\right)^{t^H} - \left(1-\lambda^H\right)^{t^H}\right] w_{t^H}^H + k,
\end{aligned} \tag{32}$$

where $k = \rho + c\left\{\beta_0 \sum_{t=1}^{t^H}\delta^t\left[\left(1-\lambda^H\right)^{t-1} - \left(1-\lambda^L\right)^{t-1}\right]\right\}$ is independent of $w_{t^H}^H$.

The expression (32) is an affine function of $w_{t^H}^H$, with a strictly positive coefficient on $w_{t^H}^H$, since $\lambda^H > \lambda^L$. Hence, we can choose $w_{T^H}^H$ sufficiently low so that (32) is negative, in which case (IC^{LH}) is satisfied because $U_0^L\left(\overline{\mathbf{C}}_w^L, \mathbf{1}\right) = 0$.

B Proof of Theorem 3

We first show that given a connected bonus contract for type L with bonus sequence $\mathbf{b}^L$ and length $T^L \leq \bar{t}^H$, either type of agent who accepts this contract will work in all periods no matter the prior history of effort if the following conditions are satisfied for each $t \in \{1, \ldots, T^L - 1\}$:

$$b^L_{T^L} \geq \frac{c}{\overline{\beta}^L_{T^L} \lambda^L} \quad \text{and} \quad \overline{\beta}^L_t \lambda^L (b^L_t - \delta b^L_{t+1}) \geq (1-\delta)c. \tag{33}$$

To show this, we use the one-step deviation principle with an induction argument. First, note that in the last period, T^L, type L will work no matter the history because by the first inequality of (33), $b^L_{T^L}$ is such that the type L weakly prefers working to shirking given that he has worked in all previous periods, and any other history only induces more favorable beliefs about the state. *A fortiori*, the H type will also work in the last period no matter the history of effort, since for any history of effort, $\beta^H_{T^L}\lambda^H > \overline{\beta}^L_{T^L}\lambda^L$, given $T^L \leq \bar{t}^H$.

Now assume inductively that for any $t < T^L$, a type $\theta \in \{H, L\}$ will work in all periods $s \in \{t+1, \ldots, T^L\}$ given any history of effort in periods $s \in \{1, \ldots, t\}$.

We want to show that, no matter the history of actions $(a_s)^{t-1}_{s=1}$, type θ weakly prefers the action plan $(a_s)^{T^L}_{s=t} = (1, 1, \ldots, 1)$ to action plan $(a_s)^{T^L}_{s=t} = (0, 1, \ldots, 1)$. If we show that type θ weakly prefers plan $(a_t, a_{t+1}) = (1, 0)$ to $(a_t, a_{t+1}) = (0, 1)$, then it follows that type θ prefers $(a_s)^{T^L}_{s=t} = (1, 0, 1, \ldots, 1)$ to $(a_s)^{T^L}_{s=t} = (0, 1, 1, \ldots, 1)$, because β^θ_{t+2} is the same in either case. The desired conclusion then follows because $(a_s)^{T^L}_{s=t+1} = (1, 1, \ldots, 1)$ is weakly preferred to $(a_s)^{T^L}_{s=t+1} = (0, 1, \ldots, 1)$ by the induction hypothesis. But that $(a_t, a_{t+1}) = (1, 0)$ is weakly preferred to $(a_t, a_{t+1}) = (0, 1)$ follows from the second inequality of (33): (i) for the low type, the inequality implies it given the history $(a_s)^{t-1}_{s=1} = (1, \ldots, 1)$, and hence for any other history (since beliefs would be more favorable); and (ii) for the high type, it follows for any history *a fortiori* because $\beta^H_t \lambda^H > \overline{\beta}^L_t \lambda^L$ for all $t < T^L$.

Therefore, given a bonus contract for the L type with length $T^L \leq \bar{t}^H$ and bonus sequence satisfying (33), both types will work in all periods under this contract, and we can focus on a relaxed problem analogous to [RP1] and [RP2].

To derive an optimal contract that minimizes the rent of the H type, we must set $T^L = \bar{t}^L$ and set the bonus sequence to make type L's incentive compatibility constraint for effort in (33) bind in each period $t \leq \bar{t}^L$. Hence, we set:

$$\bar{b}^L_{\bar{t}^L} = \frac{c}{\overline{\beta}^L_{\bar{t}^L}\lambda^L} = -\overline{w}^L_{\bar{t}^L},$$

and, for $t < \bar{t}^L$,

$$\begin{aligned}
\bar{b}^L_t &= \frac{(1-\delta)c}{\overline{\beta}^L_t \lambda^L} + \delta \bar{b}^L_{t+1} \\
&= \sum_{s=t}^{\bar{t}^L - 1} \delta^{s-t} \frac{(1-\delta)c}{\overline{\beta}^L_s \lambda^L} + \delta^{\bar{t}^L - t} \bar{b}^L_{\bar{t}^L} \\
&= \sum_{s=t}^{\bar{t}^L - 1} \delta^{s-t}(-\overline{w}^L_s) + \delta^{\bar{t}^L - t}(-\overline{w}^L_{\bar{t}^L}).
\end{aligned} \tag{34}$$

(Note that above and in the remainder of this Appendix, we suppress the argument of $\overline{w}^L_t(\cdot)$ as it is held fixed at $\bar{t}^L$.)

The key step in proving that this bonus contract achieves the second best is to confirm that the rent to type H under this contract is the same as the rent found in the optimal menu of clawback contracts. Recall that this rent was:

$$\mu_0\beta_0\left\{\sum_{t=1}^{\bar{t}^L}\delta^t\overline{w}_t^L\left[\left(1-\lambda^H\right)^t-\left(1-\lambda^L\right)^t\right]-\sum_{t=1}^{\bar{t}^L}\delta^t c\left[\left(1-\lambda^H\right)^{t-1}-\left(1-\lambda^L\right)^{t-1}\right]\right\}. \tag{35}$$

Under the low type's bonus contract defined above, it is readily seen that the rent is:

$$\mu_0\beta_0\left\{\sum_{t=1}^{\bar{t}^L}\delta^t\left[\left(1-\lambda^H\right)^{t-1}\lambda^H-\left(1-\lambda^L\right)^{t-1}\lambda^L\right]\bar{b}_t^L-\sum_{t=1}^{\bar{t}^L}\delta^t c\left[\left(1-\lambda^H\right)^{t-1}-\left(1-\lambda^L\right)^{t-1}\right]\right\}. \tag{36}$$

The value of (35) is equal to that of (36) if

$$\begin{aligned}
&\sum_{t=1}^{\bar{t}^L-1}\delta^t\left[\left(1-\lambda^H\right)^t-\left(1-\lambda^L\right)^t\right]\overline{w}_t^L+\delta^{\bar{t}^L}\left[\left(1-\lambda^H\right)^{\bar{t}^L}-\left(1-\lambda^L\right)^{\bar{t}^L}\right]\overline{w}_{\bar{t}^L}^L\\
=\;&\sum_{t=1}^{\bar{t}^L-1}\delta^t\left[\left(1-\lambda^H\right)^{t-1}\lambda^H-\left(1-\lambda^L\right)^{t-1}\lambda^L\right]\bar{b}_t^L\\
&+\delta^{\bar{t}^L}\left[\left(1-\lambda^H\right)^{\bar{t}^L-1}\lambda^H-\left(1-\lambda^L\right)^{\bar{t}^L-1}\lambda^L\right]\bar{b}_{\bar{t}^L}^L.
\end{aligned} \tag{37}$$

To prove (37), observe that by (34), the right-hand side of (37) is equivalent to

$$\begin{aligned}
&\sum_{t=1}^{\bar{t}^L-1}\delta^t\left[\left(1-\lambda^H\right)^{t-1}\lambda^H-\left(1-\lambda^L\right)^{t-1}\lambda^L\right]\left(\sum_{s=t}^{\bar{t}^L-1}\delta^{s-t}(-\overline{w}_s^L)+\delta^{\bar{t}^L-t}(-\overline{w}_{\bar{t}^L}^L)\right)\\
&\qquad+\delta^{\bar{t}^L}\left[\left(1-\lambda^H\right)^{\bar{t}^L-1}\lambda^H-\left(1-\lambda^L\right)^{\bar{t}^L-1}\lambda^L\right](-\overline{w}_{\bar{t}^L}^L),
\end{aligned}$$

and hence (37) is equivalent to

$$\begin{aligned}
&\sum_{t=1}^{\bar{t}^L-1}\delta^t\left[\left(1-\lambda^H\right)^t-\left(1-\lambda^L\right)^t\right]\overline{w}_t^L\\
=\;&\sum_{t=1}^{\bar{t}^L-1}\delta^t\left[\left(1-\lambda^H\right)^{t-1}\lambda^H-\left(1-\lambda^L\right)^{t-1}\lambda^L\right]\left(\sum_{s=t}^{\bar{t}^L-1}\delta^{s-t}(-\overline{w}_s^L)+\delta^{\bar{t}^L-t}(-\overline{w}_{\bar{t}^L}^L)\right)\\
&+\delta^{\bar{t}^L}\left[\left(1-\lambda^H\right)^{\bar{t}^L-1}-\left(1-\lambda^L\right)^{\bar{t}^L-1}\right](-\overline{w}_{\bar{t}^L}^L),
\end{aligned}$$

or equivalently, manipulating the right-hand side of the above equality using the fact that for any

$\theta \in \{H, L\}$, $\sum_{t=1}^{\bar{t}^L-1} \left(1-\lambda^\theta\right)^{t-1} \lambda^\theta = -(1-\lambda^\theta)^{\bar{t}^L-1}$, to

$$\sum_{t=1}^{\bar{t}^L-1} \delta^t \left[\left(1-\lambda^H\right)^t - \left(1-\lambda^L\right)^t\right] \overline{w}_t^L = \sum_{t=1}^{\bar{t}^L-1} \left[\left(1-\lambda^H\right)^{t-1} \lambda^H - \left(1-\lambda^L\right)^{t-1} \lambda^L\right] \left(\sum_{s=t}^{\bar{t}^L-1} \delta^s(-\overline{w}_s^L)\right). \quad (38)$$

But now observe that the right-hand side of (38) can be manipulated as follows:

$$\begin{aligned}
& \sum_{t=1}^{\bar{t}^L-1} \left[\left(1-\lambda^H\right)^{t-1} \lambda^H - \left(1-\lambda^L\right)^{t-1} \lambda^L\right] \left(\sum_{s=t}^{\bar{t}^L-1} \delta^s(-\overline{w}_s^L)\right) \\
= \quad & \delta(-\overline{w}_1^L)(\lambda^H - \lambda^L) + \delta^2(-\overline{w}_2^L) \sum_{t=1}^{2} \left[\left(1-\lambda^H\right)^{t-1} \lambda^H - \left(1-\lambda^L\right)^{t-1} \lambda^L\right] \\
& + \ldots + \delta^{\bar{t}^L-1}(-\overline{w}_{\bar{t}^L-1}^L) \sum_{t=1}^{\bar{t}^L-1} \left[\left(1-\lambda^H\right)^{t-1} \lambda^H - \left(1-\lambda^L\right)^{t-1} \lambda^L\right] \\
= \quad & \sum_{s=1}^{\bar{t}^L-1} \delta^s(-\overline{w}_s^L) \sum_{t=1}^{s} \left[\left(1-\lambda^H\right)^{t-1} \lambda^H - \left(1-\lambda^L\right)^{t-1} \lambda^L\right] \\
= \quad & \sum_{s=1}^{\bar{t}^L-1} \delta^s \overline{w}_s^L \left[(1-\lambda^H)^s - (1-\lambda^L)^s\right],
\end{aligned}$$

and since the last expression is equivalent to the left-hand side of (38), we are done.

Finally, we conclude the proof as follows: since type H's rent from the bonus contract constructed for type L is the same as in the optimal clawback menu, this bonus contract for type L solves the analog of program [RP2]. Finally, by an analogous argument as the one used in Step 8 of the proof of Theorem 2, we can choose a constant-bonus contract for type H of length t^H with a high enough bonus $b^H > 0$ and low enough initial transfer W_0^H to induce the first-best experimentation from type H and solve the analog of program [P].

References

Bergemann, Dirk and Ulrigh Hege, "Venture capital financing, moral hazard, and learning," *Journal of Banking & Finance*, August 1998, *22* (6-8), 703–735.

_ **and** _ , "The Financing of Innovation: Learning and Stopping," *RAND Journal of Economics*, Winter 2005, *36* (4), 719–752.

Besanko, David, Jian Tong, and Jianjun Wu, "A Dynamic Game under Ambiguity: Contracting for Delegated Experimentation," January 2012. unpublished.

Bhaskar, V., "Incentives and the Shadow of the Future: Dynamic Moral Hazard with Learning," June 2009. mimeo, University College London.

Bonatti, Alessandro and Johannes Horner, "Collaborating," *American Economic Review*, April 2011, *101* (2), 632–63.

Fong, Kyna, "Evaluating Skilled Experts: Optimal Scoring Rules for Surgeons," 2009. unpublished.

Garfagnini, Umberto, "Delegated Experimentation," 2011. unpublished.

Gerardi, Dino and Lucas Maestri, "A Principal-Agent Model of Sequential Testing," July 2011. unpublished.

Gershkov, Alex and Motty Perry, "Dynamic Contracts with Moral Hazard and Adverse Selection," *Review of Economic Studies*, 2012, *79* (1), 268–306.

Hart, Oliver D. and Bengt Holmstrom, "The Theory of Contracts," in Truman Bewley, ed., *Advances in Economic Theory, Fifth World Congress*, Cambridge: Cambridge University Press, 1987.

Hörner, Johannes and Larry Samuleson, "Incentives for Experimenting Agents," January 2012. mimeo, Yale University.

Keller, Godfrey, Sven Rady, and Martin Cripps, "Strategic Experimentation with Exponential Bandits," *Econometrica*, 01 2005, *73* (1), 39–68.

Manso, Gustavo, "Motivating Innovation," 2011. forthcoming in *Journal of Finance*.

Sannikov, Yuliy, "Agency Problems, Screening and Increasing Credit Lines," 2007. unpublished.

Strulovici, Bruno, "Learning While Voting: Determinants of Collective Experimentation," *Econometrica*, 2010, *78* (3), 933–971.