Implications of unequal discounting in dynamic contracting*

Ilia Krasikov Rohit Lamba Thomas Mettral

November 2020

Abstract

A principal and agent sign on a dynamic contract where (i) the agent has Markovian private information that affects total surplus, (ii) the principal can commit to the contract and the agent has limited commitment, and (iii) the principal is more patient than the agent. The interaction of these three forces, which captures many applications in financial contracting, produces permanent distortions that go through cycles. The standard rent-versus-efficiency tradeoff that determines the optimal distortion is now enriched by two competing dynamic considerations: The principal backloads agent's information rent as much as possible to relax incentive constraints, but unequal discounting introduces inter-temporal costs of incentive provision which front-load agent's payoffs. The optimal contract pins down this tradeoff. Persistence of private information creates technical challenges in determining the set of binding incentive constraints– to deal with it, a notion of simplicity and approximate optimality is introduced.

1 Introduction

In their treatise on the theory of incentives, Laffont and Martimort [2002] define the quintessential rent-efficiency tradeoff in contract theory thus:

> [T]he information gap between the principal and the agent has some fundamental implications for the design of the bilateral contract they sign... At the optimal second-best contract, the principal trades-off his desire to reach allocative efficiency against the costly information rent given up to the agent to induce information revelation.

The objective of this paper is to understand how the aforementioned tradeoff evolves when the principal and agent contract over time and the principal faces more favorable interest rates than the agent. It studies the interaction of three *forces*– (i) the agent has payoff relevant private information

*Krasikov: Higher School of Economics Moscow, krasikovis.main@gmail.com; Lamba: Penn State University, rlamba@psu.edu; Mettral: Allianz Life Insurance, t.mettral@gmx.de. We are indebted to Leeat Yariv for a conversation that instigated this paper, Vitor Farinha Luz for fruitful initial discussions, Daniel Barron for detailed feedback, and Shomak Chakrabarti for excellent research assistance. We are also grateful to four anonymous referees who helped improve the paper, and to Nageeb Ali, Marco Battaglini, Vijay Krishna, Roland Strausz, Harald Uhlig, and participants at the third annual conference on relational contracting at Northwestern University and micro theory workshop at the University of Pennsylvania for their comments. A previous draft of this paper was circulated under the title *Of restarts and shutdowns: dynamic contracts with unequal discounting*.

that follows a Markov process, (ii) the principal has commitment power while the agent has limited commitment, and (iii) the principal is more patient than the agent. Taking away the third force, novel to this paper, would put us in the rubric of standard dynamic mechanism design models (see Bergemann and Välimäki [2019]), and in addition, assuming that the principal lacks commitment would make our setup akin to a stochastic game (see Hörner et al. [2011]).

The formal model entails a “small" firm (agent) with a production technology whose productivity changes periodically according to a two state Markov process, and a “large" supplier (principal) of capital that is critical for production. The principal is more patient than the agent, and the realization of the productivity shock is privately observed by the agent. The principal commits to a dynamic menu of capital allocations in return for periodic payments. We solve for the profit maximizing contract of the principal subject to incentive compatibility and individual rationality constraints for the agent, where the former captures agency frictions and the latter limited commitment on part of the agent. At a conceptual level, we study a dynamic screening model analogous to Battaglini [2005] but with unequal discounting.

Four main results are presented. First, as is standard in contract theory, we solve the relaxed problem, and show that the solution delivers what we call a *restart contract*.[1] The high productivity type is always provided the efficient (or surplus maximizing) allocation, and the low productivity type is meted out a distortion.[2] These distortions feature cycles: they are a function of the number of low shocks since the last high shock. The contract starts with some initial distortion for the low type that monotonically converges to a positive level for successive low shocks, i.e. it does not disappear. Once a high shock arrives, it erases the memory of past distortions, and then every successive low type repeats the cycles of previous distortions, and so on. This result is in striking contrast to dynamic mechanism design models with equal discounting that predict vanishing distortions in the long-run.[3]

The cyclicality of optimal distortions is driven by a tussle between a push and a pull force. Typically, dynamic contracting allows the principal to gradually resolve the rent-versus-efficiency tradeoff in favor of the latter by *backloading* the agent's payoffs– this reduces the shadow price of incentives in the long-run.[4] But, unequal discounting introduces a novel *intertermporal cost of incentive provision*. Intuitively, if the principal promises to pay the agent an expected information rent of x tomorrow, limited commitment implies she can extract a maximum of $\delta_A x$ today. This generates an account of $\delta_A x - \delta_P x = -(\delta_P - \delta_A)x$, which is negative when the principal is more patient than the agent. So, while backloading wants x to be pushed as much in the future as possible, unequal discounting pulls against it since recouping this today now features the intertermporal cost

[1]Relaxed problem refers to the maximization of the objective (principal's profit) subject to a subset of constraints which are the only ones that bind at the optimum in the static version of the problem.

[2]Throughout, the word "distortion" means the wedge between the optimal allocation and the efficient allocation.

[3]For example, in Battaglini [2005], the contract becomes efficient the moment a high shock arrives and converges to the efficient allocation along the constant low shock history. Further, discussing Garrett, Pavan, and Toikka [2018], Bergemann and Välimäki [2019] write: "They show that regardless of whether the first-order approach is applicable or not, the optimal contract must have vanishing distortions as long as the underlying process on types is sufficiently mixing, in the sense that the impact of initial information on future types vanishes. Hence this paper confirms, for a larger class of models, one of the key findings in Battaglini [2005] derived for models with binary types."

[4]We use the standard economics terminology in referring to the Lagrange multiplier of a constraint in an optimization problem as the associated shadow price (Dixit [1990]).

of incentive provision. The tussle settles onto a compromise culminating in the restart contract.

The second result pins down down the validity of the relaxed problem (or first-order approach) approach in terms of the primitives of the model. The relaxed problem does achieve the full optimum for a large constellation of parameters. Unlike the standard model with equal discounting, however, it can fail even when the agent's type follows a two-state Markov process. The rough intuition for this is as follows: Subsequent to the erasure of past distortions on the arrival of a high shock, the first low shock creates a seed distortion which can be quite large; so the distortion for a history of shocks LHL can be much larger than the distortion for LLL. This introduces a non-monotonicity in the capital allocation rule for LHL is a better history in terms of productivity shocks but it may produce lower capital allocations. When this force is strong, incentive constraints not included in the relaxed problem starts binding, resulting in a distorted allocation for the high type as well.

In the third contribution of the paper, we ask: What can the principal do if she faces parameters for which the relaxed problem approach may not be valid? The first answer is of course to brute-force her way through (like the modeler here) and solve for the optimal contract. We provide this solution in the recursive format. While the solution is completely specified, it is quite complicated as the support of the contract set grows exponentially with time.

Complimentarily, we propose a "simpler" contract that finds the optimum in the restricted class of restart contracts, it is termed the optimal restart contract. We then construct a theoretical bound that satisfies the following two properties: (i) it shows that there is no gap between the optimal restart contract and the global optimum when the relaxed problem approach is valid, and (ii) the general loss from focussing on restart contracts is small even when the relaxed problem approach fails. This exercise is related in spirit to Chassang [2013] in that it looks for an appropriately defined approximately optimal dynamic mechanism, but it is also different from it in that it still operates within a Bayesian paradigm and demands incentive compatibility.

The fourth result formalizes this notion of simplicity that we seek in writing down the contract. Here we take the cue from Abreu and Rubinstein [1988] and frame simplicity in the language of automaton. However, unlike their approach, we do not restrict the contract to be a finite automaton, rather we allow the space of contracts to grow linearly with time. This is done both for tractability and to allow the contract to at least depend on time (as opposed to the entire history). Then, through the recursive approach to contract design, we show that restart contracts are simple and the optimal contract is simple if and only if it is restart. To the best of our knowledge, this is the first formalization of simplicity in dynamic contracts or dynamic mechanism design.

Finally, we also explore some comparative statics with respect to differential discounting and persistence in agency frictions. What happens when the principal becomes arbitrarily patient? Does the optimal contract converge to efficiency? We show this happens if and only if the agent too becomes arbitrarily patient. If the agent continues to be impatient, then the novel intertemporal costs associated with unequal discounting have a permanent and cyclical effect which does not disappear by simply making the principal arbitrarily patient.

Moreover, we also ask the question: in terms of optimal profits, does the principal prefer a myopic or patient agent? The answer here depends on the extent of agency friction as measured by the persistence in the agent's private information. If the persistence is very high and hence the information rent the principal has to pay quite large, the principal wants to bring down this cost by preferring a myopic agent. On the other hand if the persistence is low, and hence information rent quite small, the principal wants a patient agent so that she can extract future surplus upfront.

Motivation, related literature and comparative contribution. There are at least two reasons to explore principal agent models with unequal discounting. First, a burgeoning literature on dynamic mechanism design seeks to explore the implications of private and evolving information on the design of contracts, with applications such as dynamic pricing, managerial compensation and optimal taxation in mind, along with more abstract considerations of how to mitigate problem of agency in the design of institutions when the principal has some commitment power. Interestingly, and to the best of our knowledge, none of the papers thus far consider the question of how the qualitative predictions therein would change if the principal is more patient than the agent(s).

Second, unequal discounting itself captures at a high level the fact that the contractual relationship being modeled is between two parties with different capacities to last in the long-run. In more concrete terms, we have in mind the interpretation that the principal has deeper pockets or better access to capital markets and hence more favorable interest rates. In fact, Krueger and Uhlig [2006] write that different discount factors in principal-agent models can be interpreted as "the gross real interest rate or the return to some storage technology the principal has access to."

In addition, our interest in restart contracts is threefold. First, it arises naturally as the solution to the relaxed problem, which is quite the standard in contract theory. Even though we provide conditions on the primitives of the model as to when this delivers the global optimum, we choose not to ignore other parts of the parametric space, where it doesn't. Second, for these cases, we appeal to optimal restart contracts because there is an inherent normative appeal to the idea of restartness in the form of "let bygones be bygones". The contract is history dependent, but allows for the erasure of history upon realization of good outcomes, only for distortions to be reinstated on the arrival of new bad outcomes. Of course, in our model the erasure happens rather starkly, upon the realization of one high shock. Third, the implications of unequal discounting and the idea of restart contracts connect to a sizeable literature in economics, to which we now turn.

It is well known that in repeated games with differential rate of time preference the set of equilibrium payoffs expands favoring the patient player (see the classic Lehrer and Pauzner [1999]). More specifically, Opp and Zhu [2015] analyze the general relational contracting model of Ray [2002] with unequal discounting and show that all Pareto efficient contracts follow a cyclical pattern similar to our paper. The main frictions there are however different– there is no private information or moral hazard, rather two-sided limited commitment.

Further, the literatures on political economy and public finance have used unequal discounting to generate allocations that reminisce restart contracts either in their cyclicality or persistence of long-run distortions. Acemoglu, Golosov, and Tsyvinski [2008] show that when politicians are less patient than the citizens, positive aggregate labor and capital taxes are charged forever to correct for political economy distortions. Farhi and Werning [2007] find that in an Atkeson and Lucas [1992] style risk sharing model with taste shocks, when the social discount factor is higher than the private one, consumption exhibits mean reversion with no immiseration.[5] Krueger and Uhlig [2006] study a risk sharing model with a risk averse agent and competing risk neutral principals. For equal discounting the model generates full risk-sharing in the long-run, for moderate differences in discounting it generates permanent partial insurance, and when difference in discount factors is very large, the result is autarky. This again compares to our results that long-run efficiency is unattainable with even a small difference in discounting.

As argued above, for dynamic models of agency, unequal discounting can be interpreted as a sort of financial constraint. Start with the vanilla dynamic model of agency where the underlying friction is either cash flow diversion (as in literature on dynamic financial contracting) or private information (as in the literature on dynamic mechanism design). With iid shocks, pure backloading of payoffs helps mitigate the agency friction and generate long-term efficiency results.[6]

	Agency	Persistence	Discounting	Limited liability	Liquidation
Battaglini [2005]	private info	✓	equal	✗	✗
Krasikov and Lamba [2018]	private info	✓	equal	✓	✓
Biais et al. [2007]	cash flow divert	✗	unequal	✓	✓
This paper	private info	✓	unequal	✗	✗

Table 1: Dynamic models of agency with financial constraints

The vanilla model though predicts a quick dissolution of agency concerns, which does not match empirical realities. So, the cash flow diversion types of models introduced limited liability (eg. Clementi and Hopenhayn [2006]) and screening models introduced persistence in private information (eg. Battaglini [2005]) to bring some positivity to the analysis. In either case, long-term efficiency is eventually achieved, but slowed by these respective new features.

Biais et al. [2007] (see also DeMarzo and Sannikov [2006]) perhaps went the furthest in incorporating various aspects of financial constraints by adding to the iid moral hazard model with limited liability both unequal discounting and possibility of liquidation– the first feature introduces a reflective boundary below the efficient level of output which lends a cyclical element to the optimal contract, and the the second feature creates a lower boundary above zero output, below which the contract is liquidated. So, in their framework, as in our paper, efficiency is not possible, however, liquidation becomes a certainty.[7]

We contribute to this line of work on financial contracting by pushing in the realm of dynamic mechanism design where persistence of agency friction is a central facet. In our earlier work, Krasikov and Lamba [2018], we introduced both limited liability and liquidation to this model. The agent simply cannot borrow and is cash-strapped– the principal binds the limited liability constraints for as long as information rent to be paid out to the agent is recouped, and then eventually implements the efficient contract. Persistence prolongs the path to efficiency– financial constraints bind for much longer. If in addition the setup features the possibility of liquidation,

[5] A similar mechanism is generated through the interaction of aggregate shocks and unequal discounting in Aguiar, Amador, and Gopinath [2009] with an application to foreign direct investment and sovereign debt.

[6] Baron and Besanko [1984] and Laffont and Tirole [1996] are two early papers that exposit the advantages of dynamic contracts in overcoming frictions of private information in the context of regulation and pollution permits respectively.

[7] Note that Biais et al. [2007] also invoke unequal discounting for a technical reason- the continuous time limit of their discreet time model is not well defined for equal discounting. No such problem exists in our framework.

then two absorbing state emerge: efficiency or liquidation.[8]

However, we show here that a "softer" financial constraint modeled as the difference in "access to capital" in the form of unequal interest rates creates a permanent cost in generating the requisite room to relax future incentive constraints, which culminates in cyclical and non-vanishing distortions. Moreover, the interaction of this constraint of unequal access to capital with persistent private information makes the problem quite challenging, and hence we explore new technical results in the form of simple contracts. Table 1 summarizes the comparison of key modeling ingredients with the most closely related papers under the rubric of dynamic models of agency.

One way in which to view our results is in the context of the famed Modigilliani-Miller Theorem (Modigliani and Miller [1958]). Our analysis suggests that the wide-spread prevalence of inefficiencies in financial contracts could, at least at a high level, be explained by the stubborn inequality in access to capital markets. In such a situation, it is never optimal for the principal to take over the technology from the agent, or conversely "sell the firm" to the agent. Thus, in the corporate finance view of the model, the Modigilliani-Miller Theorem never holds, even in the long-run, which means capital structure always matters.

2 Model

2.1 Primitives

A firm (agent) with access to a production technology approaches a supplier (principal) of a key input (say capital); the former is a small player while the latter is a big player in the market.[9] The productivity of the firm is its private information. They agree to sign a (dynamic) contract whereby endogenous levels of input are supplied by the principal every period, in return for monetary payments by the agent. Formally, the agent's stage (or per-period) preferences are given by $\theta R(k) - p$ where k is the input supplied by the principal, p is the payment made by the agent, θ is the productivity shock, and R is a concave production function that satisfies Inada conditions.[10] The principal's stage utility is simply $p - k$.[11]

Productivity shocks can take values in $\Theta := \{\theta_H, \theta_L\}$, where $\theta_L > 0$ and $\theta_H - \theta_L = \Delta\theta > 0$. This will be referred as the agent's type. The types follow a Markov chain $\mathbb{P}(\theta_H|\theta_i) = \alpha_i$, which satisfies first-order stochastic dominance and full support: $1 > \alpha_H \geqslant \alpha_L > 0$. To simplify calculations, we assume that the prior distribution coincides with the invariant distribution of Markov process, that is $\mathbb{P}(\theta_H) = \frac{\alpha_L}{1-\alpha_H+\alpha_L}$ and $\mathbb{P}(\theta_L) = \frac{1-\alpha_H}{1-\alpha_H+\alpha_L}$. All of this information about preferences and stochastic evolution of types is common knowledge, however, the exact type realization is privately observed by the agent, and therein lies the asymmetric information or agency friction.

[8] See also Krishna, Lopomo, and Taylor [2013] who consider an iid model of private information and Krähmer and Strausz [2015]) who look at a two period sequential screening model, both with financial constraints.

[9] Throughout the agent will be referred to as a he and the principal as a she.

[10] Technically: (i) $R'(k) > 0$, $R''(k) < 0$ for all $k \geqslant 0$, (ii) $R(0) = 0$ and (iii) $\lim_{k\to 0} R'(k) = \infty$, $\lim_{k\to\infty} R'(k) = 0$.

[11] Note that other dynamic screening models can mapped into our framework and all the results in the paper can be analogously stated. For example, we can also consider the regulation model à la Laffont and Tirole [1993] where the principal and agent have preferences $V(k) - p$ and $p - \theta k$ respectively, or the monopolistic screening model à la Mussa and Rosen [1978] where the principal and agent have preferences $p - k^2/2$ and $\theta k - p$, respectively.

We consider an infinite horizon setting where the principal and agent discount future utility. However, critically, we *do not restrict them to have the same discount factor*; these are denoted by δ_P and δ_A, respectively, where $\delta_P \geqslant \delta_A$. The concept of discounting is closely connected to the idea of interest rates. For example, we can write $\delta_P = e^{-r}$ and $\delta_A = e^{-s}$ where r and s are respectively the interest rates faced by the principal and agent in the market with $s \geq r$, and the exponential representation approximates a continuously compounded rate.[12]

The principal can commit to a long-term contract. Then, invoking the revelation principle, it is without loss of generality, to focus on direct mechanisms. A direct mechanism is denoted by $\langle \mathbf{k}, \mathbf{p} \rangle := \{(k_t, p_t)\}_{t=1}^{\infty}$ where (k_t, p_t) is a function of reports up to time t: $\hat{\theta}^{\mathbf{t}} := (\hat{\theta}_1, \ldots, \hat{\theta}_t)$. Denote a history with t consecutive reports of type θ_j by $\theta^{\mathbf{t}}_{\mathbf{j}}$.[13] The principal's objective is to maximize her profit subject to incentive compatibility and participation constraints for the agent. For a fixed mechanism, the agent faces a dynamic decision problem in which her strategy is simply a function that maps his private history into an announcement every period.[14]

2.2 Constraints

Define the stage and expected utility of the agent (under truthful reporting) at any history of the contract tree to be

$$u_t(\theta^{\mathbf{t}}) := \theta_t R\left(k_t(\theta^{\mathbf{t}})\right) - p_t(\theta^{\mathbf{t}}), \quad U_t(\theta^{\mathbf{t}}) := u_t(\theta^{\mathbf{t}}) + \delta_A \mathbb{E}\left[U_{t+1}(\theta^{\mathbf{t+1}}) | \theta^{\mathbf{t}}\right].$$

It is straightforward to note that a contract can then be expressed as $\langle \mathbf{k}, \mathbf{u} \rangle$ or $\langle \mathbf{k}, \mathbf{U} \rangle$. We shall use the three formulations interchangeably.

A contract is said to be *incentive compatible* if truthful reporting by the agent is always profitable for him. Using the one shot deviation principle, incentive compatibility can be formally expressed as:[15]

$$U_t(\theta^{\mathbf{t}}) \geqslant \theta_t R\left(k_t(\theta^{\mathbf{t-1}}, \hat{\theta}_t)\right) - p_t(\theta^{\mathbf{t-1}}, \hat{\theta}_t) + \delta_A \mathbb{E}\left[U_{t+1}(\theta^{\mathbf{t-1}}, \hat{\theta}_t, \theta_{t+1}) | \theta^{\mathbf{t}}\right].$$

Equivalently, incentive compatibility can be expressed directly in terms of $\langle \mathbf{k}, \mathbf{U} \rangle$:

$$\begin{aligned} U_t(\theta^{\mathbf{t-1}}, \theta_t) - U_t(\theta^{\mathbf{t-1}}, \hat{\theta}_t) \geqslant (\theta_t - \hat{\theta}_t) R\left(k_t(\theta^{\mathbf{t-1}}, \hat{\theta}_t)\right) + \\ + \delta_A \left(\mathbb{P}(\theta_H | \theta_t) - \mathbb{P}(\theta_H | \hat{\theta}_t)\right)\left(U_{t+1}(\theta^{\mathbf{t-1}}, \hat{\theta}_t, \theta_H) - U_{t+1}(\theta^{\mathbf{t-1}}, \hat{\theta}_t, \theta_L)\right). \end{aligned}$$

where $\theta_t - \hat{\theta}_t$ is the measure of static information rents and $\mathbb{P}(\theta_H|\theta_t) - \mathbb{P}(\theta_H|\hat{\theta}_t)$ is its dynamic

[12] Abreu [1988], in his classic work, motivates the study of dynamic games under discounting as follows: "Indeed, in most economic applications, the assumption of a zero interest rate is inappropriate; we are typically concerned with situations in which the future is less important than the present." So under his interpretation, different discount factors automatically refer to access to different interest rates.

[13] At the cost of minimal confusion, the subscript will be used interchangeably for time and type. Also, as is standard, a contract is restricted to lie in l^{∞}.

[14] The private history of the agent includes the previous reported types $\hat{\theta}^{\mathbf{t-1}}$ as well as actual types $\theta^{\mathbf{t}} := (\theta_1, \ldots, \theta_t)$.

[15] The Markovian (full support) assumption on stochastic evolution of types ensures that the agent wants to report truthfully even if he has lied in the past; incentives are preserved both on and off-path.

counterpart; the latter records the fact with Markovian shocks, knowing his type today also gives some information to the agent about his types in the future. It is useful to partition the set of incentive compatibility constraints into *"downward"* (IC_H) corresponding to $\theta_t = \theta_H$ and $\hat{\theta}_t = \theta_L$, and *"upward"* ($IC_L$) corresponding to $\theta_t = \theta_L$ and $\hat{\theta}_t = \theta_H$.

A contract is said to be *individually rational* if it offers each type of the agent a non-negative expected utility after every history, that is $U_t(\theta^{\mathbf{t}}) \geqslant 0$. Individual rationality ensures that the agent is provided with a minimum expected utility at each stage, its normalization to zero is done for simplicity. This corresponds to a limited commitment assumption for the agent– he cannot be forced to be in the contractual relationship. The set of participation constraints are analogously partitioned into IR_H for $\theta_t = \theta_H$ and IR_L for $\theta_t = \theta_L$.

2.3 Optimization problem

The principal's objective is to maximize her profits subject to incentive and individual rationality constraints for the agent. This problem is now formally stated.

The static surplus (under truthful revelation) is denoted by $s(\theta, k) := \theta R(k) - k$. Thus, the (ex ante) expected surplus generated by a given contract is $\bar{S} := \sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{E}\left[s\left(\theta_t, k_t(\theta^{\mathbf{t}})\right)\right]$. Moreover, define

$$\bar{U}_P := \sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{E}\left[u_t(\theta^{\mathbf{t}})\right], \quad \bar{U}_A = \sum_{t=1}^{\infty} \delta_A^{t-1} \mathbb{E}\left[u_t(\theta^{\mathbf{t}})\right].$$

to be the expected net present value of the agent's utility using the principal and agent's discount factors, respectively. For $\delta_P = \delta_A$, we have $\bar{U}_P = \bar{U}_A$. However, in our framework, the principal and agent evaluate the agent's utility stream differently.

To express $\bar{U}_P$ only in terms of $\mathbf{U}$, parse it out into two components: $\bar{U}_P = \bar{U}_A + I$, where

$$\bar{U}_A = \mathbb{E}\left[U_1(\theta_1)\right], \quad I := \sum_{t=1}^{\infty} (\delta_P^{t-1} - \delta_A^{t-1}) \mathbb{E}\left[u_t(\theta^{\mathbf{t}})\right] = (\delta_P - \delta_A) \sum_{t=2}^{\infty} \delta_P^{t-2} \mathbb{E}\left[U_t(\theta^{\mathbf{t}})\right].$$

$\bar{U}_A$ is the *standard information rent* and I is the *intertemporal cost of incentive provision*. Then, the principal's problem, say ($\star$), can be stated as:

$$(\star) \qquad \Pi^\star = \max_{\langle \mathbf{k}, \mathbf{U} \rangle} \quad \bar{S} - \bar{U}_A - I \quad \text{subject to} \quad \mathbf{k} \geq 0 \text{ and } IC_H, IR_H, IC_L, IR_L.$$

We will refer to the solution to problem ($\star$) by $\langle \mathbf{k}^\star, \mathbf{U}^\star \rangle$ (or alternatively $\langle \mathbf{k}^\star, \mathbf{u}^\star \rangle$ or $\langle \mathbf{k}^\star, \mathbf{p}^\star \rangle$).

3 Building blocks

First, we introduce the idea of virtual value. Then, to fix ideas, we look at the solutions to the one and two period versions of the problem.

3.1 Virtual value

The main building block of the solution to the problem we study is the notion of Myersonian virtual value (Myerson [1981]). In our *quasi-linear* environment, define

$$\phi_H(\rho) := \theta_H + \rho\Delta\theta, \quad \phi_L(\rho) := \theta_L - \rho\Delta\theta,$$

to be the virtual values of the high and low types respectively. Here $\rho \geq 0$ measures the level of distortion arising out of information asymmetry, and is pinned down by the set of binding constraints at the optimum. Then, the optimal allocations are recorded as

$$\mathcal{K}_H(\rho) = arg \max_{k \geqslant 0} \quad \phi_H(\rho)R(k) - k, \quad \mathcal{K}_L(\rho) = arg \max_{k \geqslant 0} \quad \phi_L(\rho)R(k) - k.$$

Concavity of R implies that $\mathcal{K}_H$ is an increasing and $\mathcal{K}_L$ a decreasing function of ρ. The efficient allocations are, of course, given by $k^e_H = \mathcal{K}_H(0)$ and $k^e_L = \mathcal{K}_L(0)$.

3.2 Static problem

To fix ideas, and describe the basic rent-versus-efficiency tradeoff, we start with the static problem. Here discounting is of course irrelevant. The principal solves:

$$\max_{\langle \mathbf{k}, \mathbf{u} \rangle} \quad \sum_{i=H,L} \mathbb{P}(\theta_i) \left[s(\theta_i, k_i) - u_i \right] \quad \text{subject to} \quad \mathbf{k} \geqslant 0 \text{ and } IC_H, IR_H, IC_L, IR_L.$$

where we have simplified k_i for $k_1(\theta_i)$, etc. It is well known that we can look at a relaxed problem where we maximize the objective subject only to $IC_H : u_H \geqslant \Delta\theta R(k_L) + u_L$ and $IR_L : u_L \geqslant 0$. Both these constraints hold as equalities, thus the objective can be re-written as

$$\max_{\mathbf{k}} \quad \mathbb{P}(\theta_H) \left(s(\theta_H, k_H) - \Delta\theta R(k_L) \right) + \mathbb{P}(\theta_L) \left(s(\theta_L, k_L) \right) \quad \text{subject to} \quad \mathbf{k} \geqslant 0.$$

Using the notation of virtual valuation above, the optimal allocation rule is then given by: $k^\star_H = k^e_H = \mathcal{K}_H(0)$, and $k^\star_L = \mathcal{K}_L(\rho)$ for $\rho = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)} = \frac{\alpha_L}{1-\alpha_H}$. The rent-versus-efficiency essentially boils down to offering an optimal distortion to the low type that binds the IC_H constraint, and payments are further pinned down by the binding IR_L constraint. The reader can verify that this contract satisfies the remaining constraints, namely IR_H and IC_L.

In addition to static considerations, in the dynamic problem, the distortion offered to the low type (and potentially the high type), captured by ρ, evolves over time as a function of information asymmetry driven by the Markov process, and extent of differential discounting $\delta_P - \delta_A$.

3.3 Two-period problem

To understand the basics of dynamics, we start first with the two period problem. As in the static model above, we invoke the relaxed problem approach (sometimes referred to as the first-order approach), wherein we maximize the objective subject to the downward incentive constraints and the individual rationality constraint of the low type:

$$\max_{\langle \mathbf{k},\mathbf{u}\rangle} \quad \bar{S} - \bar{U}_P \quad \text{subject to} \quad \mathbf{k} \geqslant 0 \text{ and } IC_H, IR_L.$$

Unlike the equal discounting model, stage utility streams u_1 and u_2 are aggregated into different evaluations for the principal and agent to $\bar{U}_P$ and $\bar{U}_A$ respectively. Thus, in order to employ the Myersonian pointwise maximization of virtual surplus (that is, surplus minus information rents), simply calculating $\bar{U}_A$ will not do. Instead, we need to calculate the vector of stage payoffs $\mathbf{u}$ and then aggregate them to $\bar{U}_P$ using the principal's discount factor.

We show in the appendix that IC_H and IR_L both bind in the solution to the relaxed problem. Akin to the static model, the second period constraints give:

$$u_2(\theta_i,\theta_H) = u_2(\theta_i,\theta_L) + \Delta\theta R(k_2(\theta_i,\theta_L)), \quad u_2(\theta_i,\theta_L) = 0 \quad \text{for } i = H, L.$$

The binding first period incentive constraint is a bit more nuanced:

$$\begin{aligned} U_1(\theta_H) &= \theta_H R(k_1(\theta_L)) - p_1(\theta_L) + \delta\left(\alpha_H u_2(\theta_L,\theta_H) + (1-\alpha_H) u_2(\theta_{\mathrm{L}}^2)\right) \\ &= U_1(\theta_L) + \Delta\theta R(k_1(\theta_L)) + \delta_A(\alpha_H - \alpha_L)\left(u_2(\theta_L,\theta_H) - u_2(\theta_{\mathrm{L}}^2)\right) \\ &= U_1(\theta_L) + \Delta\theta R(k_1(\theta_L)) + \delta_A(\alpha_H - \alpha_L)\Delta\theta R(k_2(\theta_{\mathrm{L}}^2)) \end{aligned}$$

where in the second equality, we simply add and subtract $U(\theta_L)$ and in the third equality we substitute from the binding second period incentive constraint. Finally, the first period individual rationality constraint gives $U(\theta_L) = 0$.

The next goal is to write down the total information rent from the perspective of the principal in terms of allocations:

$$\begin{aligned} \bar{U}_P &= \sum_{i=H,L} \left(u_1(\theta_i) + \delta_P\left(\alpha_i u_2(\theta_i,\theta_H) + (1-\alpha_i) u_2(\theta_i,\theta_L)\right)\right) \\ &= \sum_{i=H,L} \mathbb{P}(\theta_i)\left(U_1(\theta_i) + (\delta_P - \delta_A)\left(\alpha_i u_2(\theta_i,\theta_H) + (1-\alpha_i) u_2(\theta_i,\theta_L)\right)\right) \\ &= \underbrace{\mathbb{P}(\theta_H)U_1(\theta_H) + \mathbb{P}(\theta_L)U_1(\theta_L)}_{\bar{U}_A} + \underbrace{(\delta_P - \delta_A)\sum_{i=H,L} \mathbb{P}(\theta_i)\left(\alpha_i u_2(\theta_i,\theta_H) + (1-\alpha_i) u_2(\theta_i,\theta_L)\right)}_{I}. \end{aligned}$$

Substitute U_1 and u_2 into the information rent and the inter-temporal costs of incentive provision:

$$\begin{aligned} \bar{U}_A &= \mathbb{P}(\theta_H)\left(\Delta\theta R(k_1(\theta_L)) + \delta_A(\alpha_H - \alpha_L)\Delta\theta R(k_2(\theta_{\mathrm{L}}^2))\right) + \mathbb{P}(\theta_L)\cdot 0 = \\ &= \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)}\Delta\theta R(k_1(\theta_L))\mathbb{P}(\theta_L) + \delta_P \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)}\Bigg(\underbrace{\frac{\delta_A}{\delta_P}\frac{\alpha_H - \alpha_L}{1-\alpha_L}}_{=:b}\Bigg)\Delta\theta R(k_2(\theta_{\mathrm{L}}^2))\mathbb{P}(\theta_{\mathrm{L}}^2), \end{aligned} \tag{1}$$

and

$$I = (\delta_P - \delta_A) \sum_{i=H,L} \mathbb{P}(\theta_i)\big(\alpha_i u_2(\theta_i, \theta_H) + (1-\alpha_i) u_2(\theta_i, \theta_L)\big)$$
$$= \delta_P \left(\Big(\underbrace{\frac{\delta_P - \delta_A}{\delta_P} \frac{\alpha_H}{1-\alpha_H}}_{=:a_H} \Big) \Delta\theta R(k_2(\theta_H, \theta_L)) \mathbb{P}(\theta_H, \theta_L) + \Big(\underbrace{\frac{\delta_P - \delta_A}{\delta_P} \frac{\alpha_L}{1-\alpha_L}}_{=:a_L} \Big) \Delta\theta R(k_2(\theta_{\mathrm{L}}^2)) \mathbb{P}(\theta_{\mathrm{L}}^2) \right). \tag{2}$$

Finally, plug in Equations (1) and (2) back into the objective, and optimize to obtain the allocations that solves the relaxed problem. The solution is presented in next result. Recollect that we define $\mathcal{K}_L(\rho) = (R')^{-1}\left(\frac{1}{\theta_i - \rho\Delta\theta}\right)$ for $\rho\Delta\theta < \theta_L$, zero otherwise.

Proposition 1. *The following supply contract* $\mathbf{k}^{\#}$ *characterizes the solution to the relaxed problem:*

$$\begin{cases} k_t^{\#}(\theta^{t-1}, \theta_H) = k_H^e = \mathcal{K}_H(0) \\ k_t^{\#}(\theta^{t-1}, \theta_L) = \mathcal{K}_L(\rho_t(\theta^{t-1}, \theta_L)) \end{cases} \qquad \text{for } \theta^{t-1} = \emptyset, \theta_H, \theta_L,$$

where $\rho_1(\theta_L) = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)} = \frac{\alpha_L}{1-\alpha_H}$, $\rho_2(\theta_{\mathrm{L}}^2) = \rho_1(\theta_L) b + a_L$ *and* $\rho_2(\theta_H, \theta_L) = a_H$.

This result pins down dynamic distortions in the two period screening contract with unequal discounting. The high type is always supplied the efficient allocation, and the supply to the low type is distorted downwards. Distortions are pervasive in that $k_t(\theta^{t-1}, \theta_L) < k_L^e$ for all θ^{t-1}. We present the intuition for how these distortions, $\rho_1(\theta_L)$, $\rho_2(\theta_{\mathrm{L}}^2)$ and $\rho_2(\theta_H.\theta_L)$, are generated.

Start with the case where $\delta_P = \delta_A$, then in Equation (2), $a_H = a_L = 0$, thus $I = 0$. The total rent for the low type, $U_1(\theta_L)$, is given by the binding IR_L constraint, and that for the high type, $U_1(\theta_H)$, is given by the binding IC_H constraint. The former gives us:

$$U_1(\theta_L) = u_1(\theta_L) + \delta_A \left[\alpha_L u_2(\theta_L, \theta_H) + (1-\alpha_L) u_2(\theta_{\mathrm{L}}^2) \right] = 0.$$

Since second period utility is non-negative, it must be $u_1(\theta_L) < 0$. Thus, the principal is *backloading* the agent's payoffs and extracting upfront some fraction of the information rent she needs to pay in the second period, this relaxes the incentive and participation constraints. In fact, once the total information rent for each type, viz. $U_1(\theta_H)$ and $U_1(\theta_L)$, has been determined, the exact timing of payments in not uniquely pinned down.[16]

The binding incentive constraints, IC_H, introduce distortions in the allocations to the low type. In the first period, this is captured by the coefficient $\rho_1(\theta_L)$. Due to persistence in private information, this distortion produced by interaction of IC_H and IR_L in the first period propagates to the second period along the history of consecutive low shocks. The propagation is captured by the coefficient b, which defines the connection between $\rho_1(\theta_L)$ and $\rho_2(\theta_{\mathrm{L}}^2)$. Note that if $\alpha_H = \alpha_L$, so the model is iid, $b = 0$ and there are no distortions in the second period. Hence, persistence is

[16] Taking $u_2(\theta_i, \theta_L) = 0$ is one possible implementation of the optimal allocation. In the formula for $\bar{U}_A$ in terms of capital allocations, we substituted for $u_2(\theta_L, \theta_H) - u_2(\theta_{\mathrm{L}}^2)$, the exact value of $u_2(\theta_{\mathrm{L}}^2)$ was not relevant.

critical for the propagation of distortions. Moreover, if the high shock is realized in the first-period, there are no distortions for the low type in the second period, $\rho_2(\theta_H, \theta_L) = 0$. The principal had already managed to extract upfront the information rent to be paid at this history, and thus the shadow price of providing incentives here is zero.

Now, let $\delta_P > \delta_A$. For a rent of x to be paid in period 2, participation constraints demand that the principal can only extract $\delta_A x$ upfront, so in comparison to the equal discounting, there is always a deficit in the principal payoff given by $\delta_A x - \delta_P x = -(\delta_P - \delta_A)x < 0$. Thus, the power of the backloading force is reduced to some extent by the *inter-temporal cost of incentive provision*, captured by the term I in Equation (2). This results in IC_H and IR_L always binding: the shadow price of incentives is positive at each history and the time structure of payments in uniquely determined. This interaction of IC_H and IR_L in the second period generates new distortions. In Proposition 1, a_H represents the coefficient of the distortion for allocation $k_2(\theta_H, \theta_L)$, and a_L represents the added distortion for allocation $k_2(\theta_{\mathrm{L}}^2)$.

To summarize, when the low type is realized in the first period, it is delivered an allocation with a distortion $\rho(\theta_L)$. Then, if another low shock is realized in the second period, a new distortion a_L is added, and the previous distortion is multiplied by b, resulting in $\rho_2(\theta_{\mathrm{L}}^2) = b\rho_1(\theta_L) + a_L$. If on the other hand a low shock is realized in the second period after a high type, then there is no propagation from the previous period, just a new seed distortion, given by $\rho_2(\theta_H, \theta_L) = a_H$.

In the next section we will see how these four characters– the *starter* $\rho(\theta_L)$, the *propagator* b, the *adder* a_L and the *seed* a_H– define all the distortions for the infinite horizon model. But before that, two final thoughts for the two-period problem.

First, once $\mathbf{k}^\#$ is determined by Proposition 1, the payments, $\mathbf{u}^\#$, are uniquely pinned down by the six binding constraints. This is in contrast to the model with equal discounting where only the first period expected utilities $U_1(\theta_H)$ and $U_1(\theta_L)$ are uniquely pinned down. Second, it is possible that a_H is large enough so that at the optimum, $k_2^\#(\theta_H, \theta_L) \ll k_2^\#(\theta_{\mathrm{L}}^2)$. Then, even though HL is better history than LL in terms of productivity shocks, the capital allocations are switched in a strong way. In this case the upward constraint IC_L can start binding in the first period. violating the validity of the relaxed problem approach. For the two period model a necessary and sufficient condition for the validity for the relaxed problem approach can be immediately generated by plugging the allocations, $\mathbf{k}^\#$ into IC_L, which delivers a condition on the primitives.

4 Restart contracts

In this section, we provide a solution to problem ($\star$) by focusing on a specific class of contracts. We call these *restart contracts*, because they reset their terms and start at the same allocation after a high type is reported. The main focus of this section is on a pair of restart contracts that provide tight upper and lower bounds on the optimal profit. For a large measure of parameters the upper and lower bounds coincide, and thus our characterization is the exact optimum. More generally, the loss from using the *optimal restart contract* is shown to be relatively small.

Definition 1. *A contract* $\langle \mathbf{k}, \mathbf{U} \rangle$ *is called* **restart** *if there exists a number* k_H *and a sequence* $\{k_t\}$ *such*

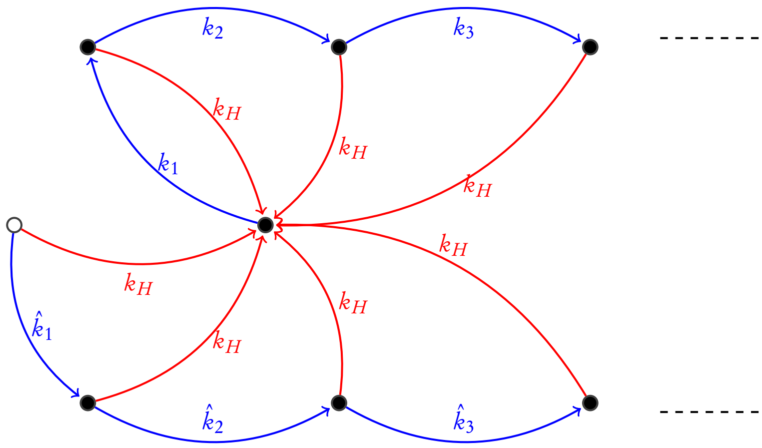

Figure 1: The evolution of allocation in a restart contract. A red/blue arrow indicates a transition, because of a high/low report.

that for all θ^{t-1}, we have

$$k_t(\theta^{t-1}, \theta_H) = k_H, \quad k_{t+s}(\theta^{t-1}, \theta_H, \theta_L^s) = k_s \quad \forall s.$$

The restart property is modeled as a measurability restriction on the allocation rule: all relevant history dependence is encoded in the number of consecutive low shocks since the last high realization. The allocation is completely characterized by the number k_H and two sequences $\{k_t\}$ and $\{\hat{k}_t\}$. The first sequence, $\{k_t\}$, defines the allocation for consecutive low shocks after a high shock has been realized, and the second sequence, $\{\hat{k}_t\}$, defines the allocation to the low type along the "lowest" history, where the high type has never been realized in the past.[17]

Figure 1 exposits the evolution of restart contracts. The contract starts in the white circle on the left, and then evolves dynamically. If the agent reports θ_H, then k_H is supplied irrespective of the previous history. If θ_L is reported in the first period then the allocation is $\hat{k}_1$, followed by $\hat{k}_t$ for every subsequent announcement of θ_L. If θ_L is reported immediately after θ_H, then k_1 is allocated, followed by k_t for every subsequent announcement of θ_L. The restart feature is captured by the fact the allocation always returns to k_H on the realization of a high shock, and remains there until a low shock is realized, which triggers the sequence $\{k_t\}$.

Now, we solve two problems that are easier to characterize than the original problem ($\star$), and that provide upper and lower bounds on the optimal profit.

[17] The second sequence is left out in Definition 1 for simplicity, because it is implicit that since the lowest history is the only one which cannot be written in the form $(\theta^{t-1}, \theta_H, \theta_L^s)$, it will have its own sequence of allocations.

4.1 Relaxed problem

We start with the standard relaxed problem approach from contract theory, wherein the incentive constraint for the low type and the individual rationality constraint for the high type are *ignored*:

$$(\#) \qquad \Pi^{\#} = \max_{\langle \mathbf{k},\mathbf{U}\rangle} \quad \bar{S} - \bar{U}_A - I \quad \text{subject to} \quad \mathbf{k} \geq 0 \text{ and } IC_H, IR_L.$$

We will denote the solution to this problem by $\langle \mathbf{k}^{\#}, \mathbf{U}^{\#}\rangle$ and its profit by $\Pi^{\#}$. This is often referred to in the literature as the *first-order optimum*, because it only takes the "first-order constraints" into account. Interestingly, $k^{\#}$ satisfies the restart property, and by construction it provides an upper bound for the optimal profit, that is $\Pi^{\star} \leqslant \Pi^{\#}$.[18] In what follows we illustrate how to obtain the first-order optimum and provide a closed-form solution.

Start by rewriting IC_H as it follows:

$$U_t(\theta^{t-1},\theta_H) - U_t(\theta^{t-1},\theta_L) \geqslant \Delta\theta R(k_t(\theta^{t-1},\theta_L)) + \delta_A(\alpha_H - \alpha_L)\left(U_{t+1}(\theta^{t-1},\theta_L,\theta_H) - U_{t+1}(\theta^{t-1},\theta_L^2)\right).$$

In the appendix, we show that IC_H and IR_L always bind at the optimum. Then, the following identity is generated by the inductive application of binding constraints:

$$U_t(\theta^{t-1},\theta_H) = \sum_{s=1}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-1}\, \Delta\theta R(k_{t-1+s}(\theta^{t-1},\theta_L^{s})). \tag{3}$$

Equation (3) gives the expression for the principal's expected profit in terms of the allocation:

$$\begin{aligned}\bar{U}_P = \bar{U}_A + I &= \mathbb{E}[U_1(\theta_1)] + (\delta_P - \delta_A)\sum_{t=2}^{\infty}\delta_P^{t-2}\mathbb{E}\left[U_t(\theta^{t})\right] = \\ &= \sum_{t=1}^{\infty}\delta_P^{t-1}\cdot\hat{\rho}_t\cdot\Delta\theta R\left(k_t(\theta_L^{t})\right)\mathbb{P}(\theta_L^{t}) + \sum_{\theta^{t-1}}\sum_{s=1}^{\infty}\delta_P^{t-1+s}\cdot\rho_s\cdot\Delta\theta R\left(k_{t+s}(\theta^{t-1},\theta_H,\theta_L^{s})\right)\mathbb{P}(\theta^{t-1},\theta_H,\theta_L^{s}).\end{aligned} \tag{4}$$

where $\{\hat{\rho}_t\}$ and $\{\rho_t\}$ are measures of agent's information rents, respectively for the lowest history where no high type is ever realized and the *restart phase* where at least one high type has been realized at some point. Equation (4) is essentially the infinite horizon version of Equations (1) and (2) that pinned down information rents in the two-period model. Recall the definition of $\mathcal{K}_L$ from Section 3.1: $\mathcal{K}_L(\rho) = (R')^{-1}\left(\frac{1}{\theta_i - \rho\Delta\theta}\right)$ for $\rho\Delta\theta < \theta_L$, zero otherwise. We have the following result.

Theorem 1. *The first-order optimum (solution to problem (#)) is a restart contract with* $k_H^{\#} = k_H^{e}$, *and*

$$\begin{cases} \hat{k}_t^{\#} = \mathcal{K}_L(\hat{\rho}_t) & \textit{for} \quad \hat{\rho}_t = b\hat{\rho}_{t-1} + a_L, \quad \hat{\rho}_1 = \frac{\alpha_L}{1-\alpha_H}, \\ k_t^{\#} = \mathcal{K}_L(\rho_t) & \textit{for} \quad \rho_t = b\rho_{t-1} + a_L, \quad \rho_1 = a_H \end{cases}$$

[18] The first-order optimum solves the original problem, thus $\Pi^{\star} = \Pi^{\#}$, if and only if $\langle \mathbf{k}^{\#}, \mathbf{U}^{\#}\rangle$ satisfies the remaining constraints, namely IC_L and IR_H.

where $b = \frac{\delta_A}{\delta_P}\frac{\alpha_H - \alpha_L}{1-\alpha_L}$ *and* $a_j = \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_j}{1-\alpha_j}$ *for* $j = H, L$.

The high type allocations are always efficient, the low type allocations in the first period, $\hat{\rho}_1$, is the same as in the static model. We call $\hat{\rho}_1$ the *starter* since it is how distortions begin. From then on, every successive low type carries over the previous distortion with the *propagator* b, and bolsters it with a_L, which we term the *adder*. This culminates into the identity $\hat{\rho}_t = b\hat{\rho}_{t-1} + a_L$. Further, the moment a high shock arrives all previous distortions are erased. Now, the realization of a new low type leads to a new *seed* of distortion $\rho_1 = a_H$, which is then propagated and added to as before for consecutive low shocks; this is captured by the identity: $\rho_t = b\rho_{t-1} + a_L$.

As promised in the two-period model, the four characters– *starter*, *propagator*, *adder*, and *seed*– pin down the solution to the relaxed problem. A simple way to think about the dynamics is the following: In each period IC_H and IR_L bind to produce a new distortion that propagates through persistence in private information. The first time these constraints bind is at the inception, the standard static distortion, $\hat{\rho}_1$, is induced. If a low shock comes at the backdrop of previous low shocks, the added distortion is a_L, and if the low shock comes immediately after a high shock, then the new distortion is a_H. Finally, each successive low shocks propagates previous distortions through b, and the arrival of a high shock stops the propagation and in fact erases the history.

While the *starter* and the *propagator* are grounded in the standard screening model with Markovian shocks and limited commitment; that the *adder* and *seed* are created by unequal discounting can be readily seen by observing $\lim_{\delta_A \to \delta_P} a_H = \lim_{\delta_A \to \delta_P} a_L = 0$. When $\delta_A = \delta_P$, the arrival of a high shock permanently removes all distortions– the principal is still paying the information rent generated by the efficient allocation, but this has been extracted through the upfront payment at the start of the contract, and hence the shadow price of all these incentives is zero. Moreover, along the consecutive low shock history, distortions propagate and eventually converge to zero. Battaglini [2005] terms these generalized no distortion at the top and vanishing distortions at the bottom.

When $\delta_P > \delta_A$, for every unit of information rent that has to be paid to the agent, unequal discounting creates the three novel economic forces: (i) backloading is costly, hence IR_L permanently binds, (ii) new intertemporal costs of incentive constraints are introduced, hence IC_H binds permanently, and (iii) net present value of standard information rent goes down (since NPV evaluated under δ_A is lower than under δ_P). These interact to endogenously determine the optimal level of allocative distortions presented in Theorem 1.

This is also a good place to make a comparative observation on dynamic models of agency. If we operated in the iid model, so that $\alpha_H = \alpha_L$, distortions are periodically renewed, $a_H = a_L > 0$, but they are completely static. Since there is no propagation, $b = 0$, there is no memory. In addition, if discounting is equal, then $a_H = a_L = 0$– there are no distortions at all beyond the first period. Therefore, to make the analysis empirically relevant, the iid models of agency (such as Clementi and Hopenhayn [2006]) and Biais et al. [2007]) invoke limited liability as a *natural* constraint that introduces history dependent distortions. That modeling choice would imply strengthening individual rationality from $\mathbf{U} \geqslant 0$ to $\mathbf{u} \geqslant 0$. In earlier work, Krasikov and Lamba [2018], we have explored this model, under persistence. In contrast, here we allow for the more permissive $\mathbf{U} \geqslant 0$, so that movement of transfers across time is feasible, but this is constrained by unequal discounting.

The magnitude of distortions can be more precisely described. The allocation for consecutive low shocks is monotonically increasing. Two things can happen in the time limit: either the limit allocation is positive, or even in the limit the distortions are not small enough to make the allocation positive. In the latter case the principal permanently shuts down the market for the low type agent. More generally, we can define *shutdown* as follows.

Definition 2. *A contract* $\langle \mathbf{k}, \mathbf{U} \rangle$ *is said to be* **shutdown** *if* $\lim_{t\to\infty} \mathbb{P}\left(k_t(\theta^{\mathbf{t-1}}, \theta_L) = 0\right) \in (0, 1]$, *and it is said to be* **permanently shutdown** *if* $\lim_{t\to\infty} \mathbb{P}\left(k_t(\theta^{\mathbf{t-1}}, \theta_L) = 0\right) = 1$.

The following list consolidates the key properties exhibited by the dynamic distortions of the first-order optimal contract.

Corollary 1. *The first-order optimal contract (solution to Problem (#)) satisfies the following properties:*

(a) distortions are monotonically decreasing: $\hat{\rho}_t > \hat{\rho}_{t+1}$ *and* $\rho_t > \rho_{t+1}$ *for all* t*;*

(b) distortions are pervasive: $\lim_{t\to\infty} \hat{\rho}_t = \lim_{t\to\infty} \rho_t = \frac{a_L}{1-b} > 0$*;*

(c) there are shutdowns in the restart phase: $k_t^{\#} = 0$ *for some* t *whenever* $\theta_L \leq \rho_1 \Delta\theta$*;*

(d) shutdowns are permanent: $k_t^{\#} = 0$ *for all* t *whenever* $\theta_L \leq \lim_{t\to\infty} \rho_t \Delta\theta$.

The evolution of distortion dynamics here are distinct than both the equal discounting model without financial constraints (eg. Battaglini [2005] and Pavan et al. [2014]), and the equal discounting model with hard financial constraints (eg. Krishna et al. [2013] and Krasikov and Lamba [2018]). In the former case, depending on the generality of model, the distortions are monotonically decreasing and the contract converges to the efficient allocation in the limit, either along every history, almost surely, or at least on average. In the latter case the distortions are monotonically increasing in the bad (or low) shocks, but the contract still does converge almost surely to the efficient allocation. Thus, in their pervasiveness, distortion dynamics here are distinct from both cases, and decreasing distortions for successive low shocks is reminiscent of the former case.

Finally, we identify the set of primitives for which the first-order optimum is globally optimal, that is when all upward incentive constraints are slack. Observe that the binding IC_H and IR_L uniquely pin down transfers as a function of allocation, thus transfers inherit the restart property, which is documented in the following simple result.

Corollary 2. *The first-order optimal payments are as it follows:* $U_t^{\#}(\theta^{\mathbf{t-1}}, \theta_L) = 0$ *and*

$$\begin{cases} U_t^{\#}(\theta_L^{t-1}, \theta_H) = \Delta\theta \sum_{s=t}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-t} (R \circ \mathcal{K}_L)(\hat{\rho}_s), \\ U_{t+s}^{\#}(\theta^{t-1}, \theta_H, \theta_L^{s-1}, \theta_H) = \Delta\theta \sum_{r=s}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{r-s} (R \circ \mathcal{K}_L)(\rho_r) \quad \forall s. \end{cases}$$

We use Corollary 2 to understand when the first-order optimum satisfies IC_L, which can be rewritten as follows:

$$U_t(\theta^{t-1}, \theta_H) - U_t(\theta^{t-1}, \theta_L) \leqslant \Delta\theta R\left(k_H^e\right) + \delta_A(\alpha_H - \alpha_L)\left(U_{t+1}(\theta^{t-1}, \theta_H^2) - U_{t+1}(\theta^{t-1}, \theta_H, \theta_L)\right).$$

Recollect that distortions depend only on the number of low shocks since the last high shock (Theorem 1), and they are monotonically decreasing along consecutive low cost realizations (Corollary 1(a)). So, the tightest upward incentive constraint is "after consecutive infinite" low cost realizations. Moreover, since $\{\rho_t\}$ and $\{\hat{\rho}_t\}$ converge to the same value (Corollary 1(b)), we can just choose the lowest history to express the tightest upward incentive constraint. Putting these together we get the following result.

Corollary 3. *The first-order optimum is globally optimal if and only if the following holds:*

$$\overbrace{\lim_{t\to\infty} U_t^{\#}(\theta_L^{t-1}, \theta_H)}^{\theta'_H s \text{ payoff at the end of the lowest history}} \leqslant \overbrace{\underbrace{\Delta\theta R(k_H^e)}_{\text{static}} + \underbrace{\delta_A(\alpha_H - \alpha_L)U_2^{\#}(\theta_H^2)}_{\text{dynamic}}}^{\theta'_L s \text{ payoff from misreporting.}}.$$

Essentially, the structure of primitives can be such that the distortions for a history of shocks LHL is much higher than the distortion LLL because the former contains the seed a_H and latter does not. The condition in Corollary 3 ensures that this does not happen for the tightest possible IC_L, which as argued above, is the one "after infinite" low cost realizations.

It can be noted that Corollary 3 is a necessary and sufficient condition *on the primitives* of the environment. This is because Corollary 2 pins down the formula for $\mathbf{U}^{\#}$ in the terms of the parameters. Since the condition is tight there is no obvious way of simplifying it. In the next result, we provide a stronger sufficient condition for the invalidity of the first-order approach that has a clearer intuitive appeal.

Corollary 4. *Fix $\delta_A \in (0, \delta_P)$ and $\kappa = \frac{2}{\delta_A}\left(1 - \frac{\delta_A}{\delta_P}\right)^{-1} \in (0, \infty)$. Then for any Markov process (α_H, α_L) that satisfies*

$$(1 - \alpha_L)(\alpha_H - \alpha_L)\left(\frac{\alpha_H}{1 - \alpha_H} - \frac{\alpha_L}{1 - \alpha_L}\right) > \kappa,$$

there exists $\Delta\theta$ small enough so that the first-order optimum is not incentive compatible.

To simplify condition stated above, assume a symmetric Markov process: $\alpha_H = 1 - \alpha_L = \alpha$, so α is the persistence. Then, the inequality can be re-written as:

$$\alpha \cdot (2\alpha - 1)\left(\frac{\alpha}{1 - \alpha} - \frac{1 - \alpha}{\alpha}\right) > \frac{2}{\delta_A\left(1 - \frac{\delta_A}{\delta_P}\right)}. \tag{5}$$

Inequality (5) can be used to derive some intuition about the validity of the relaxed problem approach. Figure 2 partitions the parameter space along the set of binding constraints– α on the x-axis and δ_A on the y-axis and three plots for different values of $\Delta\theta$. White and yellow regions represent the validity of the relaxed problem approach, the dark region is the space where the upward incentive constraints bind. The white portion in the southwest corner also represents the case of (permanent) shutdown, no capital is supplied to the low type.

Note that the right-hand side of (5) is inversely quadratic in δ_A, the term explodes as $\delta_A \to 0$ and $\delta_A \to \delta_P$. In both cases, for any fixed Markov process, (5) is not satisfied, and numerically we can see in Figure 2 that the relaxed problem approach is valid. In contrast, for fixed discounting, as $\alpha \to 1$ the left-hand side of (5) explodes, so the sufficient condition is satisfied and the relaxed problem approach is violated. Finally, as $\alpha \to 1/2$, the Markov process becomes iid and the right-hand side of (5) converges to zero, and we know for the iid model, the relaxed problem approach is valid.

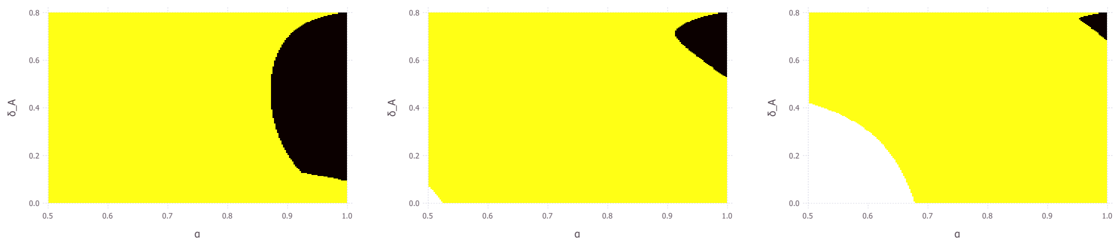

(a) $\Delta\theta = 0.1$ (b) $\Delta\theta = 1.1$ (c) $\Delta\theta = 2.1$

Figure 2: Partitioning parameter space into set of binding constraints. White & yellow: first-order approach works and optimal contract is restart. White: low type is shutdown. Black: upward constraint binds ad infinitum where $\alpha_H = 1 - \alpha_L = \alpha$ on the x-axis, δ_A on the y-axis; $\delta_P = 0.8$, $R(k) = 2\sqrt{k}$, $\theta_L = 1$.

The requirement of the smallness of $\Delta\theta$ for the sufficiency condition in Corollary 4 is also depicted in the shrinking region of binding upward incentive constraints in Figure 2 as we increase the value of $\Delta\theta$. The intuition for this is as follows: Larger values of $\Delta\theta$ signify greater ex ante asymmetric information, hence a high distortion for the low type in the first place. So LHL cannot produce a distortion significantly higher than consecutive LLL, and upward incentives are preserved. In fact for large enough value of $\Delta\theta$ as in Figure 2c, the low-type is excluded for low levels of discounting and persistence, which is the highest possible distortion.

To summarize the discussion on primitives, the first-order optimal contract satisfies the restart property such that the highest distortion occurs with the first low shock after which distortions progressively decrease to some constant and positive value. As a result $U^{\#}_{t+s}(\theta^{t-1}, \theta_H, \theta_L^{s-1}, \theta_H)$ is a decreasing function of s and so for a large enough s, it is possible that the non-monotonicity emanating from the fact that $U^{\#}_{t+s}(\theta^{t-1}, \theta_L^{s}, \theta_H) > U^{\#}_{t+s}(\theta^{t-1}, \theta_H, \theta_L^{s-1}, \theta_H)$ can violate the upward incentive constraints. Corollary 3 provides a necessary and sufficiency condition and Corollary 4 provides an easily interpretable sufficient condition for the (in)validity of the relaxed problem approach.

4.2 Restart optimum

In this section, we consider a more restrictive problem where the class of contracts is required to be restart and satisfy the full set of constraints; moreover IC_H must hold as an equality:

$$(R) \quad \Pi^R = \max_{\langle \mathbf{k},\mathbf{U}\rangle : \langle \mathbf{k},\mathbf{U}\rangle \text{ is restart, } IC_H \text{ binds}} \bar{S} - \bar{U}_A - I \quad \text{subject to } \mathbf{k} \geq 0 \text{ and } IC_H, IC_L, IR_H, IR_L.$$

We will denote the solution of this problem as $\langle \mathbf{k^R}, \mathbf{U^R} \rangle$, and refer to it as the *restart optimum*. It is easy to see that $\Pi^R \leq \Pi^\star$. When the optimal contract is restart, there is no loss from this extra restriction, that is $\Pi^R = \Pi^\star$.[19, 20]

As we mentioned in the introduction, there are three reasons for focusing on the class of restart contracts: (i) the (global) optimal contract falls within this class for a large measure of parameters, (ii) there is a normative appeal in its specific simplicity, letting bygones be bygones with a good outcomes, but introducing new and propagating old distortions for each successive bad shock, and (iii) the cyclicality of distortions and its pervasiveness connects to a large literature across fields, in public finance, sovereign debt analysis, risk-sharing and financial contracting.

In what follows we describe the restart optimum and then provide a theoretical bound to precisely capture the gap in profit between $\langle \mathbf{k^\star}, \mathbf{U^\star} \rangle$ and $\langle \mathbf{k^R}, \mathbf{U^R} \rangle$.

Theorem 2. *There exists a floor $\bar{\gamma}$ such that the restart optimum is as follows: $k_H^R \geqslant k_H^e$, and*

$$\begin{cases} \hat{k}_t^R = \mathcal{K}_L(\hat{\gamma}_t) & \text{for} \quad \hat{\gamma}_t = \max\{\bar{\gamma}, b\hat{\gamma}_{t-1} + a_L\} \text{ for some } \hat{\gamma}_1 \geqslant \frac{\alpha_L}{1-\alpha_H}, \\ k_t^R = \mathcal{K}_L(\gamma_t) & \text{for} \quad \gamma_t = \max\{\bar{\gamma}, b\gamma_{t-1} + a_L\} \text{ for some } \gamma_1 \leqslant a_H \end{cases}$$

where $b = \frac{\delta_A}{\delta_P} \frac{\alpha_H - \alpha_L}{1-\alpha_L}$ and $a_j = \left(1 - \frac{\delta_A}{\delta_P}\right) \frac{\alpha_j}{1-\alpha_j}$ for $j = H, L$.

The optimal distortions along the two class of histories, $\{\hat{\gamma}_t\}$ and $\{\gamma_t\}$, are given in Theorem 2. These are obviously analogous to their counterparts from the first-order optimal contract (Theorem 1), but there are three key differences: (i) the high type allocation is (potentially) distorted upwards, (ii) the *starter* is (weakly) higher and the *seed* (weakly) lower than its first-order optimum counterpart, i.e. $\hat{\gamma}_1 \geqslant \hat{\rho}_1$ and $\gamma_1 \leqslant \rho_1$, and (iii) there is a floor on distortions, so if the floor binds, the contract has a finite memory along consecutive low shocks.[21] Note that the *propagator* b and the *adder* a_L are the same as before. The "initial" allocations, determined by three numbers k_H^R, γ_1 and $\hat{\gamma}_1$, are picked using the first-order conditions presented in the appendix. Finally, the floor $\bar{\gamma}$, is uniquely determined according to the complementary slackness of the corresponding upward incentive constraints (IC_L).

[19] In general, the optimal restart contract does not have to satisfy all the downward constraints as equality. We require the IC_H to bind to reduce complexity of the problem, and the difference in profits is very small by not having this added restriction. Both the notion of complexity and bound on profits will be made precise.

[20] Technically, our approach here is somewhat analogous to Chassang [2013] in that it emphasizes the search for approximately optimal contracts by constraining the instruments available to the principal, but it is also different in that we do still operate within the Bayesian paradigm and demand incentive compatibility.

[21] However, it must be noted that the optimal restart contract has positive memory in that it is not the same as the static optimum, it does strictly better than the repetition of the static optimum.

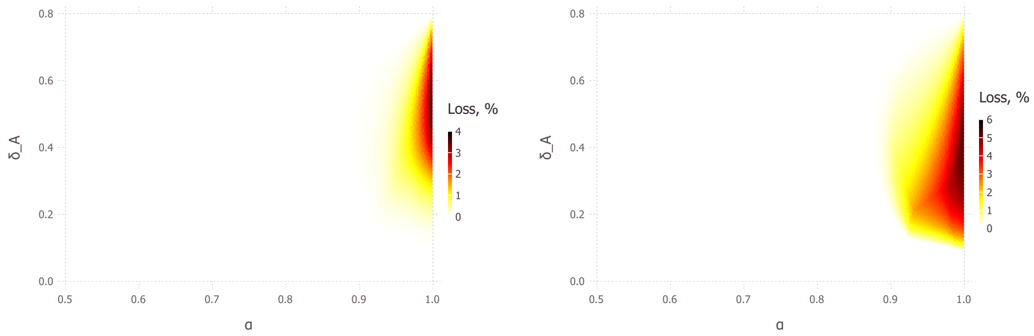

(a) actual loss

(b) analytical loss

Figure 3: Percentage loss, $\left(1 - \frac{\Pi^R}{\Pi^*}\right) * 100$ where $\alpha_H = 1 - \alpha_L = \alpha$ on the x-axis, δ_A on the y-axis; $\delta_P = 0.8$, $R(k) = 2\sqrt{k}$, $\theta_L = 1$ and $\Delta\theta = 0.1$.

How well does the optimal restart contract perform? By definition, the principal's profit from the optimal restart contract is lower than the optimal contract, $\Pi^R \leqslant \Pi^\star$. Unfortunately, the gap between the two is very hard to theoretically compute when the upward constraints bind. However, we can still bound the loss by using the expression for the first-order optimal contract, $\Pi^\#$, which is calculable in closed form. Since $\Pi^\star \leqslant \Pi^\#$, we must have $\Pi^\star - \Pi^R \leqslant \Pi^\# - \Pi^R$.

We estimate the gap using sensitivity analysis. Attach a Lagrange multiplier to each upward incentive constraint and evaluate the multipliers at the restart optimum. Quantify how much slack needs to be added to these constraints so that the solution then coincides with the first-order optimum.[22, 23] The estimate can then be written as

$$\Pi^\# - \Pi^R \leqslant \text{Lagrange multipliers} \cdot \text{Slack}.$$

Corollary 5. *There exists two bounds, B_a and B_r, functions of primitives, such that $\Pi^\star - \Pi^R \leqslant B_a$ and $1 - \frac{\Pi^R}{\Pi^\star} \leqslant B_r$, and $B_a = B_r = 0$ when the optimal contract is restart.*

One is an additive bound, and the other is a bound on the ratio. In the appendix we provide closed form expressions in terms of fundamentals. Some limit cases can be quickly registered. For the equal discounting case $\delta_A = \delta_P$, iid case $\alpha_H = \alpha_L$, and more generally when the relaxed problem approach is valid, the bound is zero, showing that it is tight with the validity of the relaxed

[22] We first describe a mathematical argument and then show how it can be applied to our setting. Consider a problem of maximizing a smooth concave function $f : \mathbb{R}^n_+ \to \mathbb{R}$ subject to a set of linear inequality constraints: $Ax \geqslant 0$. Denote the solution to this problem when the *constraints are ignored* by x^*, and consider an auxiliary problem parametrized by $\varepsilon \geqslant 0$: $\Pi(\varepsilon) = \max_{x \geqslant 0} f(x)$ subject to $Ax \geqslant \varepsilon \min\{0, Ax^*\}$. By the strong duality (assume it holds), $\Pi(\varepsilon) = \min_{\lambda \geqslant 0} \max_{x \geqslant 0} f(x) + \lambda \cdot (Ax - \varepsilon \min\{0, Ax^*\})$, thus we can estimate $\Pi(1) \leqslant \max_{x \geqslant 0} f(x) + \lambda(0) \cdot (Ax - \min\{0, Ax^*\})$ where $\lambda(0)$ is the dual variable associated with $\varepsilon = 0$. Conclude, that $\Pi(0) - \Pi(1) \leqslant \lambda(0) \cdot \max\{0, -Ax^*\}$. This argument can be easily extended to allow for linear equality constraints. In our setting: we maximize the seller's profit over $\langle \mathbf{k}, \mathbf{U} \rangle$ which is restart and satisfies (IC_H) as an equality. Moreover, we require the upward incentive constraints (IC_L) to hold, these forms the set of linear inequality constraints. The first-order optimum solves the problem when (IC_L) is ignored yielding the minimal slack, then our estimate of loss combines this slack and Lagrange multipliers when (IC_L) is imposed.

[23] Our approach of slacking upward incentive constraints and quantifying the loss associated from the exercise has a flavor of Madarász and Prat [2017] where a robust approach to multidimensional screening entails the principal giving up profits in order to relax global incentive constraints.

problem approach. Figure 3 depicts the loss from using the optimal restart contract for a specific example. As before we set $\theta_L = 1$, $\delta_P = 0.8$ and $R(k) = 2\sqrt{k}$. The unshaded region represents the validity of the relaxed problem approach so the optimal restart contract is in fact the global optimum. When the relaxed problem approach is not valid the analytical bound never exceeds 6 percent and the actual loss is never more than 4 percent.[24]

To summarize, when upward constraints (IC_L) bind at the optimum, the optimal contract can take a complicated sequential form, which is hard to pin down in a closed form. This is because both high and low type allocations are now distorted in a history dependent fashion. To generate tractable predictions, we look instead at the optimal restart contract. Restart contract kills history dependence in the allocation for the high type, and encodes all history dependence in the allocation for the low type through the number of consecutive low shocks since the last high one. This allows us to write down a simple contract that is approximately optimal in general and exactly optimal when the optimal contract is itself restart.

5 Simplicity through recursivity

In this section we first describe the recursive problem and then use its structure to formulate the idea of simplicity in contracts.

5.1 Recursive formulation

We now characterize the optimal contract through the recursive approach. We present the statement of the problem in the main text and the complete solution, since it involves additional notations, is presented in the appendix. Note that the solution to the recursive problem is globally valid in that it delivers the optimal recursive contract irrespective of whether the relaxed problem works or not.

It is well known that in order to *recursify* a dynamic contracting sequence problem where the agent's types follow an N-state Markov chain, the state variable of promised utility has to be N-dimensional (Fernandes and Phelan [2000]). In our model, it is easy to show that IR_L will always bind for the optimal contract, hence, $U^*(\theta^{t-1}, \theta_L) = 0$ at all histories. Thus, even though the agent's types follow a two state Markov process, a one dimensional state variable, viz. $U(\theta^{t-1}, \theta_H) = w \in \mathbb{R}_+$, suffices to encode all the required history dependence. The following recursive formulation can be shown to be equivalent to the sequence problem described in ($\star$), detailed proofs of this are provided in the appendix.

From the second period onwards, for an expected promised utility of w to the high type and

[24] By actual loss, we mean the exact numerical value of the loss associated with using the optimal restart contract as opposed to the first-order optimal contract, and by analytical loss we mean the value of the theoretical bound, $\min\{B_a, B_r\}$, for which no optimization is required, it is simply a function of the fundamentals of the model.

last period type j, define the objective as follows:

$$(\mathcal{RP}) \qquad S_j(w) = \max_{(\mathbf{k},\mathbf{z})\in\mathbb{R}^4_+} \alpha_j\big(s(k_H,\theta_H) - (\delta_P - \delta_A)\alpha_H z_H + \delta_P S_H(z_H)\big) +$$
$$+ (1-\alpha_j)\big(s(k_L,\theta_L) - (\delta_P - \delta_A)\alpha_L z_L + \delta_P S_L(z_L)\big) \text{ subject to}$$
$$w \geq \Delta\theta R(k_L) + \delta_A(\alpha_H - \alpha_L)z_L,$$
$$w \leq \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)z_H.$$

The objective is to maximize the surplus, $S_j(w)$, when expected utility promised to the agent is fixed at w for the high type and 0 for the low type, or $\alpha_j w + (1-\alpha_j)0$ in expectation. There are four choice variables: capital advances $\mathbf{k} = (k_H, k_L)$ and expected continuation utilities $\mathbf{z} = (z_H, z_L)$; note that z_i represents the continuation utility of the high productivity type next period if the current type is θ_i. The term $(\delta_P - \delta_A)\alpha_i z_i$ captures the intertemporal cost of incentive provision incurred by the principal in providing the continuation value of z_i. The two constraints are the downward and upward incentive constraints, IC_H and IC_L, respectively. The reader can verify that these simply re-write the constraint from Section 2.2, with an additional substitution $U^*(\theta^{t-1}, \theta_L) = 0$ since IR_L binds at the optimum. Finally, note that the participation constraint IR_H is subsumed in the recursive domain.

At date $t = 1$, the problem is different for two reasons: the belief equals the prior and contract has not yet been initialized. To initialize the contract, $w = U(\theta_H) - U(\theta_L) \geqslant 0$ must be chosen. The problem reads as follows:

$$(\diamond) \qquad \Pi^* = \max_{(w,\mathbf{z},\mathbf{k})\in\mathbb{R}^5_+} -\mu_H w + \mu_H\big[s(k_H,\theta_H) - (\delta_P - \delta_A)\alpha_H z_H + \delta_P S_H(z_H)\big] +$$
$$+ \mu_L\big[s(k_L,\theta_L) - (\delta_P - \delta_A)\alpha_L z_L + \delta_P S_L(z_L)\big] \text{ subject to}$$
$$w \geq \Delta\theta R(k_L) + \delta_A(\alpha_H - \alpha_L)z_L,$$
$$w \leq \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)z_H.$$

Denote the optimal recursive contract by $\langle w^*, \mathbf{k}(\cdot), \mathbf{z}(\cdot)\rangle$ where $(\mathbf{k}(w), \mathbf{z}(w))$ solves $(\mathcal{RP})$ for the given promise $w \geq 0$ and $(w^*, \mathbf{k}(w^*), \mathbf{z}(w^*))$ solves $(\diamond)$.[25] In the appendix we present the complete characterization of the optimal recursive contract. In what follows we use this recursive formulation to define a notion of simplicity for dynamic contracts.

5.2 Simplicity

We now show that when the optimal contract is not restart, the state space required to encode it is quite rich, and thus the optimum is not simple.

A recursive contract can be thought as an automaton which supplies capital advances to the agent conditional on an announcement of θ_H/θ_L. In such a scenario, one potential notion of simplicity is due to Abreu and Rubinstein [1988]; it counts the number of states or equivalently the

[25] As in the sequential first-order optimal contract, the allocation and transfers are uniquely pinned down. To be precise, we formally show in the appendix that $\mathbf{k}$ is unique and $\mathbf{z}$ is almost surely unique (Claim 3).

number of distinct allocations supplied by a "machine".[26] Unfortunately, in our infinite horizon setting, finite state machines are intractable and also too restrictive for they do not even allow a contract to be time dependent. A prospective alternative notion of simplicity is to let the set of allocations $\left\{k|\exists\theta^{\mathbf{t}} : k = k_t(\theta^{\mathbf{t}})\right\}$ be countable. However, this notion of simplicity is too permissive, specifically, it allows the cardinality of $\left\{k|\exists\theta^{\mathbf{t}} : k = k_t(\theta^{\mathbf{t}}), t \leqslant T\right\}$, to grow exponentially with T. We use an intermediate notion notion that is richer than finiteness but does not allow the state space to grow too fast.

Definition 3. *A contract $\langle \mathbf{k}, \mathbf{U}\rangle$ is said to be **simple** if there exists a number C such that for all T,*

$$\frac{1}{T}\left\{k|\exists\theta^{\mathbf{t}} : k = k_t(\theta^{\mathbf{t}}), t \leqslant T\right\}\Big| \leqslant C.$$

This definition allows the space of allocations to grow linearly. To the best of our knowledge, this is the first such notion of simplicity for dynamic contracts or mechanisms. When a contract is not simple, it is termed *complex*. Clearly, any restart contract is simple. We show that the optimal contract is simple if and only if the optimum is restart.

Theorem 3. *Any restart contract is simple. Moreover, the optimal contract is simple iff it is restart.*

The proof of Theorem 3 uses the recursive approach, we provide a brief sketch of the main argument here. Let $\langle U_1^\star(\theta_H), \mathbf{k}(\cdot), \mathbf{z}(\cdot)\rangle$ be the optimal recursive contract. Then, we can generate the optimum $\langle \mathbf{k}^\star, \mathbf{U}^\star\rangle$ by the following iterative procedure:

$$k_t^\star(\theta^{\mathbf{t-1}}, \theta_j) = k_j\left(U_t^\star(\theta^{\mathbf{t-1}}, \theta_j)\right), \quad U_{t+1}^\star(\theta^{\mathbf{t-1}}, \theta_j, \theta_H) = z_j\left(U_t^\star(\theta^{\mathbf{t-1}}, \theta_H)\right).$$

In the appendix, we show that the optimal recursive allocation $\mathbf{k}(\cdot)$ is a monotone function, thus complexity of the optimum is completely determined by richness of the state space used to encode it. Further, we establish that the optimal sequential contract $\langle \mathbf{k}^\star, \mathbf{U}^\star\rangle$ converges to its invariant (or stationary) distribution almost surely in finite time. Thus, in order to evaluate the simplicity of the optimum, we only need to explore whether the set of allocations (or promised utilities) in the support of the stationary distribution satisfy simplicity.

Combined with the process of productivity shocks (or types), the recursive contract induces a Markov process over $\Theta \times \mathbb{R}_+$, the two dimensional vector of type realization and continuation expected utility. For promised expected utility w and last period's shock θ_j, define the joint probability of the event that (i) the continuation expected utility lies in a Borel measurable set $A \subseteq \mathbb{R}_+$ and (ii) type realized today is θ_i as

$$F(\theta_i, A|\theta_j, w) := \mathbb{1}(z_i(w) \in A)\mathbb{P}(\theta_i|\theta_j).$$

By standard arguments, F admits a unique invariant distribution say μ (see Theorem 12.12 of Stokey et al. [1989]). Denote by μ_2 its marginal over the second argument, that is the set of promised utilities, and by $supp(\mu_2) \subseteq \mathbb{R}_+$ the support of this unique invariant distribution onto the space of continuation expected utilities.

[26] This notion was first studied by Moore [1956], and it is often referred to as the Moore-machine.

The set $supp(\mu_2)$ is a strict subset of the recursive domain, and its cardinality captures the amount of information needed to describe the optimal contract. When the optimal contract is restart, $supp(F)$ has the same "size" as the flow of time, for it is completely captured by a sequence of allocations for consecutive low shocks since the last high shock. When the optimal contract is not restart, $supp(F)$ is exponentially large for optimal distortions are completely history dependent, this violates are our notion of simplicity.

6 Comparative Statics

We provide two types of comparative statics results here: a folk theorem type of result when the principal is infinitely patient and a comparison of patient versus impatient agent from the perspective of the principal.

6.1 A folk theorem

Let $\beta = \frac{\delta_A}{\delta_P}$. Slightly abusing notations, define the average (ex-ante) profit of the principal and the average payoff of the agent at any time t to be as it follows:

$$\Pi := (1-\delta_P)\sum_{t=1}^{\infty}\delta_P^{t-1}\mathbb{E}[p_t - k_t], \quad U_t := (1-\beta\delta_P)\sum_{s=t}^{\infty}(\beta\delta_P)^{s-1}\mathbb{E}\left[\theta_s R(k_s) - p_s\right].$$

As before, we abbreviate the principal's profit at the optimum as $\Pi^\star$.

We consider the principal's profit as she becomes infinitely patient. Define s^e to be expected efficient surplus under the stationary distribution, that is

$$s^e := \mathbb{P}(\theta_H)s(\theta_H, k_H^e) + \mathbb{P}(\theta_L)s(\theta_L, k_L^e),$$

where, of course, $(\mathbb{P}(\theta_H), \mathbb{P}(\theta_L))$ is the prior, which is assumed to be the stationary distribution of the two-state Markov chain. Thus, we have the following "folk theorem".

Corollary 6. *$\Pi^\star \to s^e$ if and only if $\delta_P\beta \to 1$.*

This result can be classified into two cases. In the first case, as $\delta_P\beta \to 1$, both players are equally infinitely patient and the principal guarantees himself the total economic surplus. For imperfectly correlated types, the agent's type in the long-run is (almost) symmetrically unknown. Since the principal only cares about long-run payoffs, the information rent payable initially forms a negligible part of it. So, the principal can implement the efficient contract in the long-run and extract the associated information rent upfront. This corresponds to the standard long-term efficiency result from dynamic mechanism design for patient players (see Battaglini [2005] and Athey and Segal [2013]), and to the folk theorem in repeated games with differential discounting (Sugaya [2015]). In the folk theorem, difference between the rate of convergence of discount factor for the two players matters for the equilibrium payoff set, but the "best" achievable equilibrium does not depend on the rate, only on the limit, which is true here as well for the commitment payoff.

In the second case, where $\delta_P \beta < 1$, at least one of the player's discount factor is bounded below unity, the total surplus is bounded away from efficiency. Here either the intertemporal costs of incentive provision are forever positive ($\beta < 1$) or the standard distortions along with lowest history do not vanish fast enough ($\delta_P < 1$), or both.

6.2 Patient versus impatient agent

Does the principal favor an impatient agent or patient agent, and what determines the ranking if there exists any? Recollect that the principal's cost of providing incentives is given by $\bar{U}_P = \bar{U}_A + I$. For a fixed allocation, $\bar{U}_A$ is increasing in δ_A and I is decreasing in δ_A. The aggregate effect depends on other parameters, in particular, the level of asymmetric information as measured by the persistence of the agent's type. It can also be noted that $\bar{U}_A$ is increasing in the persistence of the agent's types, and I is not monotonic in persistence. The complexity of these competing forces does not allow for a global comparative static, but a theoretical result can be stated for the limit cases and numerical arguments explored for the intermediate ones.

Corollary 7. *Consider a symmetric Markov chain with $\alpha_H = 1 - \alpha_L = \alpha$. The principal's ex ante payoff in the first-order optimal, optimal and optimal restart contracts varies with δ_A as it follows:*

(a) principal prefers patient agent ($\delta_A = \delta_P$) for α sufficiently close to $\frac{1}{2}$.

(b) principal prefers myopic agent ($\delta_A = 0$) for α sufficiently close to 1.

Figure 4 plots principal's profit in the first-order optimal contract and the optimal restart contract. It presents a "heat map" where each point in the box represents the expected profit of the principal as a function of α (on the x-axis) and δ_A (on the y-axis), where darker shades mean higher values. The northwest and southeast corners of the parametric spaces correspond to cases (a) and (b) of Corollary 7. In the intermediate range it is clear that the for each value of α the principal's profit changes non-linearly as a function of δ_A. For example at $\alpha = 0.9$, the principal prefers either a completely myopic agent ($\delta_A = 0$) or completely forward looking one ($\delta_A = \delta_P$), but not goldilocks.

When persistence is very high, principal has to pay a large information rent. Thus, in order to bring down this cost she prefers a myopic agent. On the other hand, with very low persistence, since the information rent is small the principal can generate a large surplus and extract most of it upfront, even though it involves incurring the inter-temporal cost of incentive provision. So, in this case, she prefers a myopic agent.

7 Final remarks

Many long-term contractual situations involve one party that is financially bigger or more integrated in capital markets and the other endowed with private information. What kind of contracts do we expect to observe in such environments? Pursuing such a framework, we analyzed a dynamic principal-agent model with three ingredients: persistent private information, limited commitment and unequal discounting.

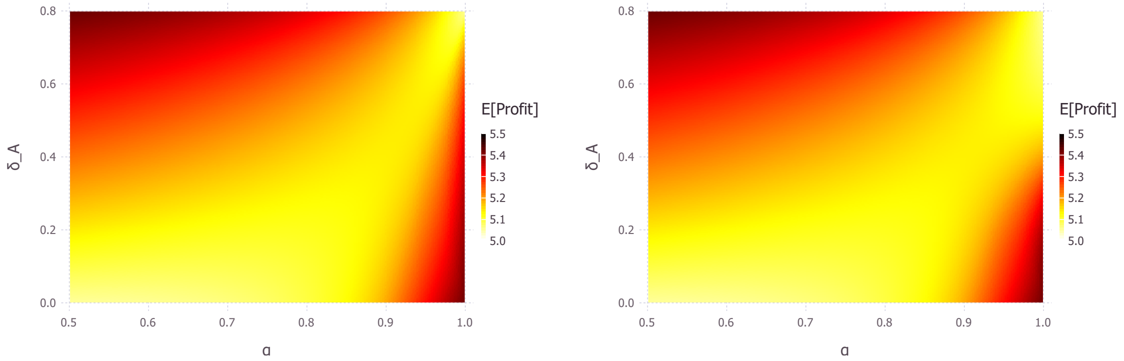

(a) First-order optimal contract

(b) Optimal restart contract

Figure 4: Principal's profit where $\alpha_H = 1 - \alpha_L = \alpha$ on the x-axis, δ_A on the y-axis; $\delta_P = 0.8$, $R(k) = 2\sqrt{k}$, $\theta_L = 1$ and $\Delta\theta = 0.1$.

Their interaction produces a tradeoff for the principal: backloading agent's payoffs as much as possible to lax future incentive constraints, and front-loading them to minimize the inter-temporal cost of incentive provision. This constant tussle between the two forces produce a cyclical structure of allocative distortions that we term *restart*. The optimal contract is completely characterized– sequentially for the relaxed problem and recursively for the global optimum. When the relaxed problem approach is valid, the optimal contract is restart, and when it is not valid, the optimum requires an exponentially growing state space to encode all relevant history dependence. In the latter case, we characterize the optimal restart contract that provides a simpler and approximately optimal alternative, where the both simplicity and approximate optimality are formally defined.

The nature of dynamic distortions poses a question to the literature on dynamic (Myersonian) mechanism design– a slight perturbation of the standard model of equal discounting renders long-term efficiency unachievable, distortions are pervasive. With equal discounting, Besanko [1985] and Battaglini [2005] show that ex post distortions converge to zero in the long run for the AR(1) and two type Markov models respectively. Garrett, Pavan, and Toikka [2018] show that distortions converge to zero on average for more general types' processes. Our results make clear that these predictions will not hold for unequal discounting.

The modeling of financial constraints as differential interest rates through unequal discounting and limited commitment as compared to limited liability constraints is a departure from standard dynamic financial contracting literature. We term this as soft versus hard financial constraints. In the absence of financial constraints the principal extracts maximal possible information rent up-front. In the presence of hard financial constraints in the form of limited liability, the principal binds the limited liability constraints for as long as information rent to be paid out to the agent is recouped, and then eventually implements the efficient contract (see Krishna et al. [2013] and Krasikov and Lamba [2018]). However, a permanent difference in access to capital creates a permanent cost in generating the requisite room to relax future incentive constraints, which culminates in cyclical and non-vanishing distortions.

The paper also discussed the connection of our modeling approach to a sizeable literature in macroeconomics, public finance and political economy, which uses unequal discounting to understand forces as disparate as debt dynamics, societal altruism for future generations and evolution of capital taxes. In each cases, some mechanism resembling the restart contract emerges.

A limitation of our model is 'permanency' of the differential interest rates. A more detailed analysis would allow the agent to save his way towards the market rate. There are many plausible ways on introducing this added dimension to our model. One tractable way could perhaps be to allow the discount factor of the agent to depend on the level of equity of the "firm structure".[27] So, as the agent's share in total surplus increases, the interest rate he faces also converges to the one faced by the principal. It would be a reduced form yet an endogenous way of allowing for the effects of financial constraints to be mitigated in the long-run. This seems to us a fruitful question for future research.

Finally, one can ask the question- what if the agent is more patient than the principal? Though most of our applications fit the patient principal model, this is an interesting theoretical question in its own right. It turns out that the model as stated is then not compact; the lack of an upper bound on transfers means that the principal will borrow or demand an unbounded amount of money hoping to create a Ponzi scheme. Imposing an upper bound rectifies the problem- the optimal allocation rule in the equal discounting case continues to be the optimum for the model with $\delta_A > \delta_P$, and transfers are uniquely pinned down through the upper bound.

8 Appendix

8.1 Sequential characterization

8.1.1 Binding constraints

First, we establish the set of binding constraints in problems ($\star$), ($\#$) and (R): Lemma 1 shows that IR_L binds in all three problems, Lemma 2 states that IC_H binds in the relaxed problem. We use the terminology that a constraint is respected when it holds as a weak inequality- may or may not bind but is satisfied.

Lemma 1. *Consider a mechanism $\langle \mathbf{k}, \mathbf{U} \rangle$ that respects IC_H, IR_L such that $U_t(\tilde{\theta}^{\mathbf{t-1}}, \theta_L) > 0$ for some history $\tilde{\theta}^{\mathbf{t-1}}$. Then, there exists another mechanism $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ that satisfies IC_H, IR_L, IR_H and yields a higher ex-ante profit. In addition, if $\langle \mathbf{k}, \mathbf{U} \rangle$ respects IC_L, then $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ can be chosen to do this as well.*

Proof. Define $\tilde{\mathbf{U}}$ as $\tilde{U}_t(\theta^{\mathbf{t-1}}, \theta_L) = 0$ and $\tilde{U}_t(\theta^{\mathbf{t-1}}, \theta_H) = U_t(\theta^{\mathbf{t-1}}, \theta_H) - U_t(\theta^{\mathbf{t-1}}, \theta_L)$. By construction, the new mechanism satisfies IR_L, moreover, its incentive compatibility constraints are exactly identical the same as in the original mechanism. To see that IR_H holds, inductively expand IC_H along a persistent history of θ_L's:

$$\tilde{U}_t(\theta^{\mathbf{t-1}}, \theta_H) \geqslant \sum_{s=1}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-1} \Delta\theta R(k_{t-1+s}(\theta^{\mathbf{t-1}}, \theta_{\mathbf{L}}^{\mathbf{s}})) \geqslant 0,$$

[27] In dynamic contracting models with agency frictions, the share of the principal can be regarded as the debt and the share of the agent as equity, and the sum of two as the total value of the firm that is born out of the contractual relationship between the two, see for example Clementi and Hopenhayn [2006].

where the last inequality follows from non-negativity of production.

Finally, note that $\tilde{\mathbf{U}} \leqslant \mathbf{U}$ and $\tilde{\mathbf{U}} \neq \mathbf{U}$, thus the altered mechanism is cheaper for the principal as $\tilde{U}_P < U_P$. Conclude that the ex-ante profit of $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ is higher than of $\langle \mathbf{k}, \mathbf{U} \rangle$. □

Lemma 2. *Consider a mechanism $\langle \mathbf{k}, \mathbf{U} \rangle$ that respects IC_H, satisfies IR_L as an equality, but*

$$U_t(\tilde{\theta}^{\mathbf{t-1}}, \theta_H) > \sum_{s=1}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-1} \, \Delta\theta R(k_{t-1+s}(\tilde{\theta}^{\mathbf{t-1}}, \theta_{\mathbf{L}}^{\mathbf{s}})) \quad \textit{for some history} \quad \tilde{\theta}^{\mathbf{t-1}}.$$

Then, there exists another mechanism $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ that satisfies IC_H, IR_L and yields a higher ex-ante profit.

Proof. Define $\tilde{\mathbf{U}}$ as $\tilde{U}_t(\theta^{\mathbf{t-1}}, \theta_L) = 0$ and $\tilde{U}_t(\theta^{\mathbf{t-1}}, \theta_H) = \sum_{s=1}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-1} \, \Delta\theta R(k_{t-1+s}(\theta^{\mathbf{t-1}}, \theta_{\mathbf{L}}^{\mathbf{s}}))$. The reader can verify that IR_L and IC_H bind in the new mechanism.

Note that $\tilde{\mathbf{U}} \leqslant \mathbf{U}$ and $\tilde{\mathbf{U}} \neq \mathbf{U}$, thus the altered mechanism is cheaper for the principal as $\tilde{U}_P < U_P$. Conclude that the ex-ante profit of $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ is higher than of $\langle \mathbf{k}, \mathbf{U} \rangle$. □

8.1.2 Relaxed problem approach

We now complete the proof of Theorem 1 and Corollary 1.

Proof of Theorem 1. The goal is to obtain Equation 4 and derive two sequences of distortions $\{\hat{\rho}_t\}$ and $\{\rho_t\}$, which are described in the statement of the theorem.

By Lemmata 1 and 2, it is without loss of generality to focus on mechanisms in which both IC_H and IR_L bind at every history. So, consider a mechanism $\langle \mathbf{k}, \mathbf{U} \rangle$ satisfying these two properties, that is $U_t(\theta^{\mathbf{t-1}}, \theta_L) = 0$ and

$$U_t(\theta^{\mathbf{t-1}}, \theta_H) = \sum_{s=1}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-1} \, \Delta\theta R\left(k_{t-1+s}(\theta^{\mathbf{t-1}}, \theta_{\mathbf{L}}^{\mathbf{s}})\right).$$

We now use the set of binding constraints to rewrite the principal's profit as a function of allocations. First, we solve for the agent's ex ante utility:

$$\mathbb{E}\left[U_1(\theta_1)\right] = \sum_{t=1}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{t-1} \mathbb{P}(\theta_H) \Delta\theta R\left(k_t(\theta_{\mathbf{L}}^{\mathbf{t}})\right) = \sum_{t=1}^{\infty} (\delta_P b)^{t-1} \frac{\alpha_L}{1-\alpha_H} \Delta\theta R\left(k_t(\theta_{\mathbf{L}}^{\mathbf{t}})\right) \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}).$$

where $b = \frac{\delta_A}{\delta_P} \frac{\alpha_H - \alpha_L}{1-\alpha_L}$ is the multiplicative distortion.

Next, we solve for the intertemporal costs of incentive provision:

$$\begin{aligned} I &= (\delta_P - \delta_A) \sum_{t=2}^{\infty} \delta_P^{t-2} \mathbb{E}\left[U_t(\theta^{\mathbf{t}})\right] = (\delta_P - \delta_A) \sum_{\theta^{\mathbf{t-1}}: t \geqslant 2} \delta_P^{t-2} \mathbb{P}(\theta^{\mathbf{t-1}}, \theta_H) U_t(\theta^{\mathbf{t-1}}, \theta_H) = \\ &= (\delta_P - \delta_A) \sum_{\theta^{\mathbf{t-1}}: t \geqslant 2} \sum_{s=1}^{\infty} \delta_P^{t-2} \mathbb{P}(\theta^{\mathbf{t-1}}, \theta_H) (\delta_A(\alpha_H - \alpha_L))^{s-1} \Delta\theta R\left(k_{t-1+s}(\theta^{\mathbf{t-1}}, \theta_{\mathbf{L}}^{\mathbf{s}})\right). \end{aligned}$$

Expand the intertemporal costs of incentive provision separately at the lowest history of θ_L's and the restart phase where $\theta^{\mathbf{t-1}}$ contains at least one θ_H.

In the former case, we have

$$(\delta_P - \delta_A)\sum_{t=2}^{\infty}\sum_{s=1}^{\infty}\delta_P^{t-2}\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t-1}},\theta_H)(\delta_A(\alpha_H-\alpha_L))^{s-1}\Delta\theta R\left(k_{t-1+s}(\theta_{\mathbf{L}}^{\mathbf{t-1+s}})\right) =$$

$$= a_L\sum_{t=2}^{\infty}\delta_P^{t-1}\left(\sum_{s=1}^{t-1}b^{s-1}\right)\Delta\theta R\left(k_t(\theta_{\mathbf{L}}^{\mathbf{t}})\right)\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}})$$

where $a_L = \frac{\delta_P-\delta_A}{\delta_P}\frac{\alpha_L}{1-\alpha_L}$ is the L-seed. The distortions along the lowest history are then given by

$$\hat{\rho}_t = b^{t-1}\frac{\alpha_L}{1-\alpha_H} + a_L\left(\sum_{s=1}^{t-1}b^{s-1}\right) = b\hat{\rho}_{t-1} + a_L.$$

In the latter case, we have

$$(\delta_P - \delta_A)\sum_{\theta^{\mathbf{t-1}}:\theta^{\mathbf{t-1}}\neq\theta_{\mathbf{L}}^{\mathbf{t-1}},t\geqslant 2}\sum_{s=1}^{\infty}\delta_P^{t-2}\mathbb{P}(\theta^{\mathbf{t-1}},\theta_H)(\delta_A(\alpha_H-\alpha_L))^{s-1}\Delta\theta R\left(k_{t-1+s}(\theta^{\mathbf{t-1}},\theta_{\mathbf{L}}^{\mathbf{s}})\right) =$$

$$= a_H\sum_{\theta^{\mathbf{t-1}}}\sum_{s=1}^{\infty}\delta_P^{t-1+s}\left(\sum_{r=1}^{s}b^{r-1}\right)\Delta\theta R\left(k_{t+s}(\theta^{\mathbf{t-1}},\theta_H,\theta_{\mathbf{L}}^{\mathbf{s}})\right)\mathbb{P}(\theta^{\mathbf{t-1}},\theta_H,\theta_{\mathbf{L}}^{\mathbf{s}})$$

where $a_H = \frac{\delta_P-\delta_A}{\delta_P}\frac{\alpha_H}{1-\alpha_H}$ is the H-seed. The total distortions in the restart phase can concisely be written as

$$\rho_t = b^{t-1}a_H + a_L\left(\sum_{s=1}^{t-1}b^{s-1}\right) = b\rho_{t-1} + a_L.$$

□

Proof of Corollary 1. Consider a function $f(x) = bx + a_L$ with $b = \frac{\delta_A}{\delta_P}\frac{\alpha_H-\alpha_L}{1-\alpha_L}$ and $a_j = \frac{\delta_P-\delta_A}{\delta_P}\frac{\alpha_j}{1-\alpha_j}$ for $j = H, L$. By Theorem 1, the first-order optimum is characterized by two sequences of distortions $\{\rho_t\}$ and $\{\hat{\rho}_t\}$ that satisfy $\rho_{t+1} = f(\rho_t)$ and $\hat{\rho}_{t+1} = f(\hat{\rho}_t)$.

The reader can verify that the function f has a unique non-zero fixed point, that is $\frac{a_L}{1-b}$. Moreover, $f(x) \gtrless x$ whenever $x \lessgtr \frac{a_L}{1-b}$. Thus, this fixed point is globally stable and both sequences of distortions converge to it monotonically.

Parts *(a)* and *(b)* immediately follow from the following set of inequalities:

$$\frac{a_L}{1-b} < \frac{\alpha_L}{1-\alpha_H} < a_H.$$

Parts *(c)* and *(d)* clearly follow from the definition of $\mathcal{K}_L(x)$, that is $\mathcal{K}_L(x) = (R')^{-1}\left(\frac{1}{\theta_L - x\Delta\theta}\right)$ for $x\Delta\theta < \theta_L$ and zero otherwise. □

8.1.3 Validity of the relaxed problem approach

Corollaries 2 and 3 provide the necessary and sufficient condition for the validity of the relaxed problem approach. The former expresses expected utility as a function of primitives, whereas the latter identifies the tightest possible upward incentive constraint in the whole set of IC_L.

Proof of Corollaries 2 and 3. Corollary 2 follows simply from Equation (3). For Corollary 3, note that xx. □

Next, we provide a proof of Corollary 4, which gives a stronger sufficient condition that is more intuitive.

Proof of Corollary 4. Define a number Γ as a slack in the tightest possible upward incentive constraint, that is

$$\Gamma := \lim_{t\to\infty} U_t^{\#}(\theta_{\mathrm{L}}^{\mathrm{t-1}}, \theta_H)/\Delta\theta - R(k_H^e) - \delta_A(\alpha_H - \alpha_L)U_2^{\#}(\theta_H^2)/\Delta\theta.$$

According to Corollary 3, the first-order optimum is globally optimal if and only if $\Gamma \leqslant 0$. It is useful to rewrite Γ only in terms of the optimal distortions $\{\rho_t\}$, which are identified in Theorem 1:

$$\Gamma = \sum_{t=0}^{\infty}[\delta_A(\alpha_H - \alpha_L)]^t \, (R \circ \mathcal{K}_L)\left(\lim_{s\to\infty}\rho_s\right) - R(k_H^e) - \sum_{t=1}^{\infty}[\delta_A(\alpha_H - \alpha_L)]^t \, (R \circ \mathcal{K}_L)\,(\rho_t)\,.$$

The reader can verify that the value of Γ at $\Delta\theta = 0$ is zero. Therefore, to prove the claim it is sufficient to establish that Γ is increasing in $\Delta\theta$ in a neighborhood of zero.

We now show that Γ is increasing in $\Delta\theta$ for values that are sufficiently close to zero. First of all, $\mathcal{K}_L(x)$, which is defined as a solution to $\max_{k\geqslant 0}(\theta_L - x\Delta\theta)R(k) - k$, is positive for $\Delta\theta$ that is sufficiently close to zero:

$$1/(\theta_L - x\Delta\theta) = R'(\mathcal{K}_L(x)).$$

By the implicit function theorem, $\mathcal{K}_L$ is differentiable in $\Delta\theta$ at zero, moreover, its derivate is proportional to the value of x, that is

$$\left.\frac{\partial \mathcal{K}_L(x)}{\partial \Delta\theta}\right|_{\Delta\theta=0} = x\frac{1}{(\theta_L)^2 R''(k_L^e)}.$$

Note also that $k_H^e = \mathcal{K}_H(0) = \mathcal{K}_L(-1)$, as a result Γ is differentiable in $\Delta\theta$ at zero. Taking the common factor outside the brackets, we can express the derivative of Γ with respect to $\Delta\theta$ as

$$\left.\frac{\partial \Gamma}{\partial \Delta\theta}\right|_{\Delta\theta=0} = \underbrace{\frac{1}{(\theta_L)^2 R''(k_L^e)}}_{<0}\left(\lim_{s\to\infty}\rho_s + 1 + \sum_{t=1}^{\infty}(\delta_A(\alpha_H - \alpha_L))^t(\lim_{s\to\infty}\rho_s - \rho_t)\right).$$

The first term is negative due to strict concavity of R. We claim that under the premise of corollary the second term is also negative, thus $\left.\frac{\partial\Gamma}{\partial\Delta\theta}\right|_{\Delta\theta=0} > 0$.

We now compute the second term of the above expression using the notations introduced in Theorem 1, i.e., $b = \frac{\delta_A}{\delta_P}\frac{\alpha_H - \alpha_L}{1-\alpha_L}$ and $a_j = \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_j}{1-\alpha_j}$ for $j = H, L$. Recall that the optimal distortions are defined as $\rho_t = (1 - b^{t-1})\frac{a_L}{1-b} + b^{t-1} a_H$, therefore

$$\lim_{s \to \infty} \rho_s - \rho_t = b^{t-1}\left(\frac{a_L}{1-b} - a_H\right).$$

Substitute these into the second term of the above expression:

$$\frac{a_L}{1-b} + 1 + \frac{\delta_A(\alpha_H - \alpha_L)}{1 - b\delta_A(\alpha_H - \alpha_L)}\left(\frac{a_L}{1-b} - a_H\right) =: \zeta.$$

To complete the proof, we need to show that $\zeta < 0$ under the assumption of Corollary 4. To see it formally, multiply the left hand side by $(1-\alpha_L)[1 - b\delta_A(\alpha_H - \alpha_L)]$ and rearrange to obtain that $\zeta < 0$ if and only if

$$\delta_A\left(1 - \frac{\delta_A}{\delta_P}\right)(1-\alpha_L)(\alpha_H - \alpha_L)\left(\frac{\alpha_H}{1-\alpha_H} - \frac{\alpha_L}{1-\alpha_L}\right) > (1-\alpha_L)\left(\frac{a_L}{1-b} + 1 - b\delta_A(\alpha_H - \alpha_L)\right).$$

Note that $b < \delta_A/\delta_P$ and $\alpha_H \geqslant \alpha_L \geqslant 2$, therefore the right hand side is not higher than 2. Then, the premise of Corollary 4 implies that

$$(1-\alpha_L)(\alpha_H - \alpha_L)\left(\frac{\alpha_H}{1-\alpha_H} - \frac{\alpha_L}{1-\alpha_L}\right) > 2/\delta_A\left(1 - \frac{\delta_A}{\delta_P}\right)^{-1}.$$

As a result, $\zeta < 0$, thus $\frac{\partial \Gamma}{\partial \Delta\theta} > 0\Big|_{\Delta\theta = 0} > 0$. By continuity of Γ, the first-order optimum is not incentive compatible for $\Delta\theta > 0$ that is close to zero.

□

8.1.4 Restart optimum

We now characterize the restart optimum (Theorem 2) and derive its profit guarantee (Corollary 5).

By Lemma 1, it is without loss of generality to focus on mechanisms such that IR_L bind at every history, i.e., $U_t(\theta^{\mathbf{t-1}}, \theta_L) = 0$. We further restrict the contract space to be a set of mechanisms satisfying IC_H as an equality at every history. Our restriction on the contract space implies that the agent's expected utilities are pinned down by the binding downward incentive constraints, moreover, they also feature restarts. In other words, for any permissible mechanism $\langle \mathbf{k}, \mathbf{U} \rangle$ there exist two sequences $\{U_t\}$ and $\{\hat{U}_t\}$ such that for all $\theta^{\mathbf{t-1}}$ and s, we have

$$U_t(\theta_{\mathbf{L}}^{\mathbf{t-1}}, \theta_H) = \hat{U}_t, \quad U_{t+s}(\theta^{t-1}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s-1}}, \theta_H) = U_s.$$

These two sequences are determined as a function of the allocation rule as

$$\hat{U}_t = \Delta\theta R(\hat{k}_t) + \delta_A(\alpha_H - \alpha_L)\hat{U}_{t+1}, \quad U_t = \Delta\theta R(k_t) + \delta_A(\alpha_H - \alpha_L)U_{t+1}.$$

It follows that IC_L is equivalent to the following system of inequalities:

$$\hat{U}_t \leqslant \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)U_1, \quad U_t \leqslant \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)U_1.$$

The former is the upward constraint along the lowest history, the latter corresponds to the restart phase.

It is convenient to rewrite the objective in terms of the sequences of allocations and utilities. First, we decompose the expected surplus into three terms: the high type surplus, the surplus along the lowest history and the surplus in the restart phase:

$$\begin{aligned}\bar{S} &= \sum_{t=1}^{\infty} \delta_P^{t-1}\mathbb{E}\left[s\left(\theta_t|k_t(\theta^{\mathrm{t}})\right)\right] = \\ &= \frac{\mathbb{P}(\theta_H)}{1-\delta_P}s(\theta_H, k_H) + \sum_{t=1}^{\infty}\delta_P^{t-1}\mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t}})s(\theta_L, \hat{k}_t) + \frac{\mathbb{P}(\theta_H)}{1-\delta_P}\sum_{t=1}^{\infty}\delta_P^t\mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t}}|\theta_H)s(\theta_L, k_t).\end{aligned}$$

The term $\frac{\mathbb{P}(\theta_H)}{1-\delta_P}$ is a discounted probability of θ_H, that is $\sum_{t=1}^{\infty}\delta_P^{t-1}\mathbb{P}(\theta_t = \theta_H)$. Next, note that the agent's expected payoff is simply $\mathbb{E}[U_1(\theta_1)] = \mathbb{P}(\theta_H)\hat{U}_1$, whereas the costs of incentive provision can be factored as

$$\begin{aligned}I &= (\delta_P - \delta_A)\sum_{\theta^{\mathrm{t-1}}:t\geqslant 2}\delta_P^{t-2}\mathbb{P}(\theta^{\mathrm{t-1}}, \theta_H)U_t(\theta^{\mathrm{t-1}}, \theta_H) = \\ &= (\delta_P - \delta_A)\sum_{t=2}^{\infty}\delta_P^{t-2}\mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t-1}}, \theta_H)\hat{U}_t + (\delta_P - \delta_A)\frac{\mathbb{P}(\theta_H)}{1-\delta_P}\sum_{t=1}^{\infty}\delta_P^{t-1}\mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t-1}}, \theta_H|\theta_H)U_t.\end{aligned}$$

The former term captures the costs along the lowest history, and the latter reflects the costs in the restart phase.

Taking all pieces together, Problem (R) can be equivalently written as

$$\max_{k_H, \{\hat{k}_t\}, \{k_t\}, \{\hat{U}_t\}, \{U_t\}} \quad \bar{S} - \mathbb{E}[U_1(\theta_1)] - I \quad \text{subject to} \quad k_H \geqslant 0, \ \forall t \quad \hat{k}_t, k_t, \hat{U}_t, U_t \geqslant 0, \text{ and}$$

$$\begin{aligned}\hat{U}_t &= \Delta\theta R(\hat{k}_t) + \delta_A(\alpha_H - \alpha_L)\hat{U}_{t+1} \leqslant \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)U_1, \\ U_t &= \Delta\theta R(\hat{k}_t) + \delta_A(\alpha_H - \alpha_L)U_{t+1} \leqslant \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)U_1.\end{aligned}$$

We are now in position to prove Theorem 2 and derive the bound described in Corollary 5 (see Figure 3b for a visualization).

Proof of Theorem 2. Problem (R) is strictly concave and bounded, thus the restart optimum can be characterized using the Lagrangian method. We first build the Lagrangian by attaching a multiplier to each constraint. Specifically, the downward incentive constraints along the lowest history are associated with dual variables $\delta_P^{t-1}\mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t}})\hat{\gamma}_t$, whereas the upward incentive constraints are associated with dual variables $\delta_P^{t-1}\mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t}})\hat{\eta}_t$. Similarly, in the restart phase multipliers are $\frac{\mathbb{P}(\theta_H)}{1-\delta_P}\delta_P^t\mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t}}|\theta_{\mathrm{H}})\gamma_t$ and $\frac{\mathrm{P}(\theta_H)}{1-\delta_P}\delta_P^t\mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t}}|\theta_{\mathrm{H}})\eta_t$ for the downward and upward incentive constraints, respectively.

We now derive the set of first-order conditions, these are $\hat{k}_t = \mathcal{K}_L(\hat{\gamma}_t)$, $k_t = \mathcal{K}_L(\gamma_t)$ and $k_H = \mathcal{K}_H(\kappa) \geqslant k_H^e$ where

$$\kappa := \sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t}}) \hat{\eta}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \sum_{t=1}^{\infty} \delta_P^t \mathbb{P}(\theta_{\mathrm{L}}^{\mathrm{t}} | \theta_{\mathrm{H}}) \eta_t.$$

In what follows we establish existence of the set of dual variables satisfying the properties outlined in Theorem 2, moreover, we show that there is no duality gap.

To begin, fix $\bar{\gamma} \geqslant 0$, $\mu_1 \in \left[\frac{a_L}{1-b}, a_H\right]$, and define $\{\hat{\gamma}_t\}$, $\{\gamma_{t+1}\}$ by

$$\hat{\gamma}_t := \max\left\{\bar{\gamma}, b^{t-1}\hat{\gamma}_1 + (1-b^{t-1})\frac{a_L}{1-b}\right\}, \quad \gamma_{t+1} := \max\left\{\bar{\gamma}, b^t \gamma_1 + (1-b^t)\frac{a_L}{1-b}\right\}.$$

Then, let $\eta_1 := 0$, $\hat{\eta}_1 := \left(\hat{\gamma}_1 - \frac{\alpha_L}{\alpha_H}\right)^+$ and

$$\hat{\eta}_{t+1} := \hat{\gamma}_{t+1} - b\hat{\gamma}_t - a_L, \quad \eta_{t+1} := \gamma_{t+1} - b\gamma_t - a_L.$$

The reader can verify that $\{\eta_t\}$ and $\{\eta_t\}$ are both non-negative and continuous in $(\bar{\gamma}, \mu_1)$. By construction, the coefficients in the Lagrangian in front of $\{\hat{U}_t\}$ and $\{U_{t+1}\}$ are identically zero. In addition, the coefficient in front of U_1 is proportional to

$$(a_H - \gamma_1)\frac{\delta_A}{\delta_P}\frac{1-\alpha_H}{\alpha_H - \alpha_L}\frac{\mu_H}{1-\delta_P} - \kappa.$$

Note that $\kappa = 0$ whenever $\bar{\gamma}$ is sufficiently small, moreover, it is strictly increasing in $\bar{\gamma}$ without bound. Therefore, for any $\gamma_1 \in \left[\frac{a_L}{1-b}, a_H\right)$ there exists a unique value of $\bar{\gamma}$ which makes the aforementioned coefficient equal to zero. For $\gamma_1 = a_H$, any $\bar{\gamma} \leqslant \min\left\{\frac{\mu_H}{\mu_L}, \frac{a_L}{1-b}\right\} = \frac{a_L}{1-b}$ will do.

To conclude the proof, we need to show that there exists a value of γ_1 such that the complimentary slackness is satisfied at all histories. The only non-trivial case is when the first-order optimum is not incentive compatible, otherwise, $\gamma_1 = a_H$ will work. Since the distortions $\gamma, \hat{\gamma}$ are monotone and stay at the same value once upward incentive compatibility starts to bind, it is sufficient to only verify the complimentary slackness "at infinity", that is

$$\lim_{t\to\infty} U_t = \Delta R(k_H) + \delta_A(\alpha_H - \alpha_L)U_1.$$

The reader can check that the left hand side is larger than the right hand side for $\gamma_1 = a_H$, provided that the first-order optimum is not incentive compatible. On the other hand, the left hand side is smaller than the right hand side for $\gamma_1 = \frac{a_L}{1-b}$. To see it more formally, let $\gamma_1 = \frac{a_L}{1-b}$. Then, we have $\bar{\gamma} > \min\left\{\frac{\alpha_L}{1-\alpha_H}, \frac{a_L}{1-b}\right\} = \frac{a_L}{1-b}$, because $\mu_1 < a_H$. Taking two observations together, by continuity, there exists a value of $\gamma_1 \in \left(\frac{a_L}{1-b}, a_H\right)$ for which the complimentary slackness is satisfied. □

Proof of Corollary 5. Since the only difference between problems (R) and (#) is the set of upward incentive constraints, the difference in ex-ante profits of these two problems can be assessed using the standard perturbation argument.

Consider the first-order optimum $\langle \mathbf{k}^{\#}, \mathbf{U}^{\#}\rangle$ and define the slack in the upward incentive constraints as

$$\epsilon_t := \left(\hat{U}^{\#}_t - \Delta R(k^e(\theta_H)) - \delta_A(\alpha_H - \alpha_L)U^{\#}_1\right)^+, \quad \hat{\epsilon}_t := \left(\hat{U}^{\#}_t - \Delta R(k^e(\theta_H)) - \delta_A(\alpha_H - \alpha_L)U^{\#}_1\right)^+.$$

Then, we have

$$\Pi^{\#} - \Pi^R \leqslant \sum_{t=1}^{\infty} \delta_P^{t-1}\mathbb{P}(\theta^{\mathrm{t}}_{\mathrm{L}})\hat{\eta}_t \cdot \epsilon_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P}\sum_{t=1}^{\infty}\delta_P^t \mathbb{P}(\theta^{\mathrm{t}}_{\mathrm{L}}|\theta_{\mathrm{H}})\eta_t \cdot \epsilon_t.$$

Our goal is to build an upper bound on the right hand side of this expression in terms of the primitives. We now provide two different ways to measure the right hand side.

Our first bound is based on the fact that the slack variables are monotone, thus we can substitute the largest slack variable into the right hand side. Recall that the distortions in the first-order optimum are monotone, thus $\hat{\epsilon}_t, \epsilon_t \leqslant \lim_{s\to\infty} \epsilon_s$ for all t. Using the first-order condition for U_1 in Problem (R) and $\frac{a_L}{1-b} \leqslant \gamma_1$, we obtain the first upper bound B^1_a on the profit gap:

$$\Pi^{\#} - \Pi^R \leqslant \frac{\delta_P(1-\alpha_H)}{\delta_A(\alpha_H - \alpha_L)}\left(a_H - \frac{a_L}{1-b}\right)\frac{\mathbb{P}(\theta_H)}{1-\delta_P}\lim_{t\to\infty}\epsilon_t =: B^1_a.$$

Our second bound is based on assessing the dual variables $\{\eta_t\}$ and $\{\hat{\eta}_t\}$. The reader can verify that $\bar{\gamma} \leqslant \gamma_1$, thus $\gamma_{t+1} - a_L - b\gamma_t = \eta_{t+1} \leqslant \bar{\gamma}(1-b) - a_L \leqslant (1-b)\left(a_H - \frac{a_L}{1-b}\right)$. Similarly, we have $\hat{\eta}_{t+1} \leqslant (1-b)\left(a_H - \frac{a_L}{1-b}\right)$. At the initial date, $\eta_1 = 0$ and $\hat{\eta}_1 \leqslant \left(a_H - \frac{\alpha_L}{1-\alpha_H}\right)^+$. Combining all pieces together, we obtain the second upper bound B^1_a on the profit gap:

$$\begin{aligned}\Pi^{\#} - \Pi^R \leqslant\ & \mathbb{P}(\theta_L)\left(a_H - \frac{\alpha_L}{1-\alpha_H}\right)^+\hat{\epsilon}_1 + \\ & + (1-b)\left(a_H - \frac{a_L}{1-b}\right)\sum_{t=2}^{\infty}(\delta_P(1-\alpha_L))^{t-1}\left(\mathbb{P}(\theta_L)\hat{\epsilon}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P}\delta(1-\alpha_H)\epsilon_t\right) =: B^2_a.\end{aligned}$$

We now construct an upper bound on the relative profit loss. To make sure it does not explode, we also compute the loss from using the optimal static contract, which specifies a history independent allocation to θ_L. The optimal static contract supplies the efficient quantity to the high type and $k_L = K_L(x)$ to the low type where

$$x = \frac{1-\delta_A}{1-\delta_A(\alpha_H - \alpha_L)}\frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)}.$$

Denote the profit from using this static contract by Π^s, of course, it has a clear closed form representation. Then, we have $\Pi^{\#} - \Pi^R \leq \Pi^{\#} - \Pi^S$.

To sum up, we arrive at the following analytical bounds:

$$\Pi^* - \Pi^R \leqslant \min\{B^1_a, B^2_a, \Pi^{\#} - \Pi^S\} =: B_a \quad \text{and} \quad 1 - \frac{\Pi^R}{\Pi^*} \leqslant B_a/\Pi^{\#} =: B_r.$$

The former is absolute, whereas the latter is relative.

□

8.2 Recursive characterization

In this section we study the recursive problem introduced in the main text, and then use it to prove the result on simplicity. In what follows we first completely characterize the solutions to the problem jointly defined by ($\mathcal{RP}$) and ($\diamond$), see Section 5.1.

8.2.1 Preliminary results

Let W be the largest set of promised utilities $w \in \mathbb{R}$ such that there exists an incentive compatible and individually rational contract which delivers $U_1(\theta_H) = w$ and $U_1(\theta_L) = 0$. The set W is a familiar recursive domain, which was introduced in Spear and Srivastava [1987]. In our setting the recursive domain has a very simple structure as shown in the following lemma.

Lemma 3 (Recursive domain)**.** $W = \mathbb{R}_+$.

Proof. First of all, every $w \in W$ must be such that $w \geqslant 0$ by IR_H. On the other hand, any $w \geqslant 0$ can be implemented, for example, the following mechanism $\langle \mathbf{k}, \mathbf{U} \rangle$ will do: $U_1(\theta_1) = w, k_1(\theta_H) = R^{-1}\left(\frac{w}{\Delta\theta}\right)$ and $k_t(\theta^{\mathbf{t}}) = U_t(\theta^{\mathbf{t}}) = 0$ for all $\theta^{\mathbf{t}} \neq \theta_H$. □

Using Lemma 3, we can express the recursive problem as ($\mathcal{RP}$) from the second period onwards, and as ($\diamond$) in the first period, explicitly stated in Section 5.2. The reader can verify that the sequential problem and its recursive counterpart admit the same solution. To formally show equivalence between the sequential and recursive formulations, we need to introduce auxiliary notations.

A policy correspondence $w \mapsto (\mathbf{K}(w), \mathbf{Z}(w))$ specifies a set of optimal solutions in ($\mathcal{RP}$) for every $w \in \mathbb{R}_+$. We say that a mechanism $\langle \mathbf{k}, \mathbf{U} \rangle$ is generated from the policy correspondence $(\mathbf{K}(\cdot), \mathbf{Z}(\cdot))$ if $k_{t+1}(\theta_j, \theta^{\mathbf{t-1}}, \theta_i) \in \mathbf{K}_i(U_{t+1}(\theta_j, \theta^{\mathbf{t-1}}, \theta_H))$ and $U_{t+2}(\theta_j, \theta^{\mathbf{t-1}}, \theta_i, \theta_H) \in \mathbf{Z}_i(U_{t+1}(\theta_j, \theta^{\mathbf{t-1}}, \theta_H))$ for $i, j = H, L$ and for all $\theta^{\mathbf{t-1}}$.

The next claim formally connects the sequential and recursive formulations.

Claim 1.

(a) There exists a unique continuous bounded function $S_j(w)$ satisfying the Bellman equation in ($\mathcal{RP}$).

(b) The policy correspondence $(\mathbf{K}(\cdot), \mathbf{Z}(\cdot))$ is non-empty, compact-valued and upper hemicontinuous.

(c) A contract $\langle \mathbf{k}, \mathbf{U} \rangle$ is optimal if and only if it is generated from the policy correspondence $(\mathbf{K}(\cdot), \mathbf{Z}(\cdot))$, and $\langle (U_1(\theta_H), (k_1(\theta_H), k_1(\theta_L)), (U_2(\theta_H, \theta_H), U_2(\theta_L, \theta_H)) \rangle$ solves Problem ($\diamond$).

(d) The value of problems ($\star$) and ($\diamond$) coincide.

Proof. The result follows from Exercises 9.4, 9.5 in Stokey et al. [1989]. □

In the rest of the section we outline several standard properties of the value function (Claim 2), establish uniqueness of transfers (Claim 3) and prove Propositions 2, 3.

Claim 2 (Properties of the value function)**.**

(a) S_j *is concave.*

(b) S_j *is continuously differentiable on* $\mathbb{R}_{++}$*.*

(c) S_j *is locally strictly concave at every* w *satisfying* $S_j'(w) > 0$*.*

Proof.

Part (a). The argument is standard, we need to show that the Bellman operator, implicitly defined in ($\mathcal{RP}$), preserves concavity. Note that the constraints set is convex and $s(\theta, \cdot)$ is concave. Then, the result follows from Theorem 3.1 and its Corollary 1 in Stokey et al. [1989].

Part (b). We established concavity of the value function using the standard argument. As for differentiability, the standard argument of Benveniste and Scheinkman [1979] is not applicable in our context, because it might not to be possible to change $\mathbf{k}$ keeping $\mathbf{z}$ constant. We give a different argument that is close to Rincón-Zapatero and Santos [2009] in its spirit. We shall use the fact that S_j is concave, thus it is subdifferentiable.

Take $\langle \mathbf{k}^\star, \mathbf{U}^\star \rangle$ which solves the sequence problem starting from the second period with $U_2^\star(\theta_j, \theta_H) = w$. Using the generalized first-order and envelope conditions for ($\mathcal{RP}$), we argue that there exists a finite time s such that the value function is differentiable at $U_{s+1}^\star(\theta_j, \theta_{\mathbf{L}}^{\mathbf{s}-1}, \theta_H)$. Then, the value function turns out to be differentiable at the original point, w.

Before we show differentiability, we shall validate that the first-order conditions are sufficient to characterize the solution. In particular, we show that the Slater's condition holds, which is sufficient to guarantee that the first-order approach with Lagrange multipliers in l^1 is valid in the sequence problem, because of concavity and boundedness (see Morand and Reffett [2015]). We claim that for any $w > 0$, there exists a feasible point such that the constraint map is uniformly bounded away from 0. The argument is constructive: since $w > 0$, there exist two numbers $k_H > k_L > 0$ satisfying

$$\frac{\Delta\theta}{1 - \delta_A(\alpha_H - \alpha_L)} R(k_L) < w < \frac{\Delta\theta}{1 - \delta_A(\alpha_H - \alpha_L)} R(k_H).$$

Then, define a contract $\langle \mathbf{k}, \mathbf{U} \rangle$ as $k_{t+1}(\theta_j, \theta^{\mathbf{t}-1}, \theta_H) = k_H$, $k_{t+1}(\theta_j, \theta^{\mathbf{t}-1}, \theta_L) = k_L$ and $U_{t+1}(\theta_j, \theta^{\mathbf{t}-1}, \theta_H) = w$.

We are now in a position to show that S_j is continuously differentiable. Recall that $\langle \mathbf{k}^\star, \mathbf{U}^\star \rangle$ is the solution to the sequence problem at $t = 2$ with $U_2^\star(\theta_j, \theta_H) = w$. The reader can verify that the capital supplied to θ_H can be distorted only upwards, thus $k_{t+1}^\star(\theta_j, \theta^{\mathbf{t}-1}, \theta_H) > 0$ is uniquely defined at all histories by strict concavity of the objective. In addition, if $k_{t+1}^\star(\theta_j, \theta^{\mathbf{t}-1}, \theta_L) > 0$, then it is unique by strict concavity of the objective.

Next, consider the recursive problem ($\mathcal{RP}$), its solution exists and coincides with one found above. Since S_j is concave, its superdifferential at $w > 0$ is well-defined and equals to $\partial S_j(w) =$

$[S_j^+(w), S_j^-(w)]$, and at $w = 0$ it is $S_j^+(0)$ where a plus/minus denotes a right/left subderivative. The goal is to establish that the right and left subderivatives coincide.

Let $\alpha_j \rho_H$ and $(1-\alpha_j)\rho_L$ be Lagrange multipliers for the upward and downward incentive constraints, respectively. And, $\rho_j(w)$ be some Lagrange multiplier supporting the solution , whereas $\rho_j^-(w)/\rho_j^+(w)$ be the highest/smallest such Lagrange multiplier. Finally, denote by $(\mathbf{k}(w), \mathbf{z}(w))$ some point in the optimal correspondence. The first-order conditions with respect to $\mathbf{k}$ are $k_i(w) = \mathcal{K}_i(\rho_i(w))$ for $i = H, L$. By the above argument, $K_H(w)$ is a singleton and $\rho_H^+(w) = \rho_H^-(w) = \rho_H(w)$ for every w. In addition, if $k_L(w) > 0$, then $K_L(w)$ is a singleton and $\rho_L^+(w) = \rho_L^-(w) = \rho_L(w)$. So, for $w > 0$, there might be multiple multipliers only if $\rho_L^-(w) \geqslant \theta_L/\Delta\theta > 0$. It follows that the downward incentive constraint must bind, and we have that $z_L(w) = \frac{w}{\delta_A(\alpha_H - \alpha_L)} > w > 0$ is uniquely defined.

Then, the envelope theorem reads $S_j^-(w) - S_j^+(w) = (1-\alpha_j)(\rho_L^-(w) - \rho_L^+(w))$. It is immediate that S_j is differentiable at w if and only if $\rho_L^-(w) = \rho_L^+(w)$. The first-order condition with respect to z_L when $z_L(w) > 0$ reads as follows:

$$\delta_P S_L^-(z_L(w)) \geqslant \alpha_L(\delta_P - \delta_A) + (\alpha_H - \alpha_L)\delta_A \rho_L(w) \geqslant \delta_P S_L^+(z_L(w)).$$

If $\rho_L(z_L(w))$ is unique, then $\rho_L(w)$ is so and S_j is differentiable at w. Now, define recursively $z_L^s = z_L(z_L^{s-1})$ with $z_L^0 = w > 0$ for some selection from Z_L. There are two potential cases, namely $\rho_L(z_L^s)$ is unique for some s or it is not for all s. In the former case, S_j is differentiable at w by our previous argument. In the latter case, $z_L^s = \frac{w}{\delta_A^s(\alpha_H - \alpha_L)^s} \to \infty$ as $s \to \infty$ which is impossible, because any solution must be in l^∞. To complete the proof, note that continuous differentiability of S_j is implied by differentiability and concavity.

Part (c). The proof is by contradiction. Suppose that $S_j'(w) = S_j'(w+\epsilon) > 0$ for some $w, \epsilon > 0$. Consider $\langle \mathbf{k}^\star, \mathbf{U}^\star \rangle$ and $\langle \mathbf{k}^\varepsilon, \mathbf{U}^\varepsilon \rangle$ solving the sequence problem at $t = 2$ with $U_2^\star(\theta_j, \theta_H) = w$ and $U_2^\varepsilon(\theta_j, \theta_H) = w + \varepsilon$, respectively. Since $s(\theta, \cdot)$ is strictly concave, it must be that $\mathbf{k}^\star = \mathbf{k}^\varepsilon$. Otherwise, we would have $S_j'(w) < S_j'(w+\varepsilon)$.

Now, since $S_j'(w) = S_j'(w+\varepsilon) > 0$, the envelope theorem implies that the downward incentive constraint binds in both cases. By the first-order and envelope conditions, see Equations 6, 7 and 8 below, it will continue to bind along the sequence of θ_L's. Then, we have

$$w = \Delta\theta \sum_{s=1}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-1} R\left(k_{t+s-1}^\star(\theta^{t-2}, \theta_j, \theta_L^s)\right) = w + \varepsilon.$$

The last assertion is a clear contradiction. By the same argument, $S_j'(w-\epsilon) > S_j'(w)$. □

We now derive the set of optimality conditions that is useful for our characterization of the optimal contract. Let $(1-\alpha_j)\rho_H$ and $\alpha_j \rho_L$ be Lagrange multipliers attached to the constraints in $(\mathcal{RP})$. And, let $\mathbb{P}(\theta_H)\rho_H$ and $\mathbb{P}(\theta_L)\rho_L$ be Lagrange multipliers attached to the constraints in $(\diamond)$. We denote by $(\mathbf{k}(\cdot), \mathbf{z}(\cdot))$ a selection from the optimal correspondence and by $\rho(\cdot)$ the corresponding

Lagrange multipliers. So, the first-order conditions are $k_i(w) = \mathcal{K}_i(\rho_i(w))$ for $i = H, L$ and

$$S'_H(z_H(w)) - \alpha_H \frac{\delta_P - \delta_A}{\delta_P} + (\alpha_H - \alpha_L)\frac{\delta_A}{\delta_P}\rho_H(w) \begin{cases} = 0 \text{ if } z_H(w) > 0, \\ \leqslant 0 \text{ if } z_H(w) = 0, \end{cases} \tag{6}$$

$$S'_L(z_L(w)) - \alpha_L \frac{\delta_P - \delta_A}{\delta_P} - (\alpha_H - \alpha_L)\frac{\delta_A}{\delta_P}\rho_L(w) \begin{cases} = 0 \text{ if } z_L(w) > 0, \\ \leqslant 0 \text{ if } z_L(w) = 0. \end{cases} \tag{7}$$

In addition, the Envelope theorem gives:

$$S'_j(w) = (1 - \alpha_j)\rho_L(w) - \alpha_j \rho_H(w) \text{ for } j = H, L. \tag{8}$$

Finally, we argue that the Lagrange multipliers are unique. Let $\langle \mathbf{k}^\star, \mathbf{U}^\star \rangle$ be the solution to the sequence problem at $t = 2$ with $U_2^\star(\theta_j, \theta_H) = w$. Since the capital supplied to θ_H can be distorted only upwards, thus $k_t^\star(\theta^{t-2}, \theta_j, \theta_H) > 0$ is uniquely defined by strict concavity of the objective. It follows from Claim 1 that $\rho_H(w) = \mathcal{K}_H^{-1}\left(k_t^\star(\theta^{t-2}, \theta_j, \theta_H)\right)$, and ρ_H is continuous, because $\langle \mathbf{k}^\star, \mathbf{U}^\star \rangle$ changes continuously with w. It remains to select $\rho_L(w)$ to satisfy the envelope condition.

8.2.2 Optimal recursive contract

In this section, we exposit the properties of the optimal recursive contract, $\langle w^*, \mathbf{k}(\cdot), \mathbf{z}(\cdot) \rangle$ where $w^* = U_1^\star(\theta_H)$ and $(\mathbf{k}(w), \mathbf{z}(w))$ solves $(\mathcal{RP})$ for each $w \geqslant 0$; $(w^*, \mathbf{k}(w^*), \mathbf{z}(w^*))$ solves $(\diamond)$.[28] We start with registering the monotonicity of allocation with respect to expected utility given to the high type.

For the optimal recursive contract, allocations for the high and low productivity shocks are increasing in the state variable, w. Intuitively speaking, the downward incentive constraint binds only for low values of w. In this case, the allocation and promised expected utility upon announcing the low type (that is, k_L and $\alpha_L z_L$) must be distorted downwards to prevent the high type from misreporting. Indeed, there exists a critical value w_L^* so that the downward incentive constraint binds only for $w \leqslant w_L^*$. The incentive problem is more severe for low values of w, there exists another threshold w_k^o below which the contract does not supply θ_L.

By the similar reasoning, the allocation and promised expected utility upon announcing the high type (that is, k_H and $\alpha_H z_H$) must be distorted upwards if the upward incentive constraint binds. And, there exists a critical value w_H^* such that this constraint binds if and only if $w \geqslant w_H^*$. Figure 5a plots the optimal allocation as the function of agent's expected utility.

For the latter references, it is useful to construct these threshold formally. By Part (c) of Claim 2, there exists a unique number z_L^e such that $z_L(w) = z_L^e$ whenever the downward incentive constraint is slack. By the same token, there exists a unique number z_H^e such that $z_H(w) = z_H^e$

[28] As in the sequential first-order optimal contract, the allocation and transfers are uniquely pinned down. To be precise, we formally show in the appendix that only z_H could fail to be unique at a single point. The details are provided in Claim 3.

whenever the upward incentive constraint is slack. The reader can verify that each number satisfies $z_j^e > 0$ and $S_j'(z_j^e) = \alpha_j \frac{\delta_P - \delta_A}{\delta_P}$ or $z_j = 0$ and $S_j'(0) \leqslant \alpha_j \frac{\delta_P - \delta_A}{\delta_P}$. Then, the critical thresholds are defined as $w_j^* := \Delta\theta R(k^e(\theta_j)) + \delta_A(\alpha_H - \alpha_L)z_j^e > 0$.

We then have the following simple result.

Proposition 2. *The allocation in the optimal recursive contract satisfies the following:*

(a) $\exists w_H^*$ *such that* $k_H(w) = k_H^e$ *if and only if* $w \leqslant w_H^*$, $k_H(\cdot)$ *is strictly increasing on* $[w_H^*, \infty)$.

(b) $\exists w_k^o, w_L^*$ *such that* $k_L(w) = 0$ *if and only if* $w \leqslant w_k^o$, $k_L(w) = k_L^e$ *if and only if* $w \geqslant w_L^*$, $k_L(\cdot)$ *is strictly increasing on* $[w_k^o, w_L^*]$.

Proof of Proposition 2. It suffices to characterize the optimal distortions $\rho_L(\cdot)$ and $\rho_H(\cdot)$, because their properties translate into $\mathbf{k}(\cdot)$ by the first-order condition $k_i(w) = \mathcal{K}_i(\rho(w))$ for $i = H, L$.

Part (a). If the upward incentive constraint is slack, then, by definition, $k_H(w) = k_H^e$ and $z_H = z_H^e$. Since this choice is feasible if and only if $w \geqslant w_H^*$, the result for $\rho_H(\cdot)$ follows.

We now establish monotonicity of $\rho_H(\cdot)$. Take $w' > w \geqslant w_H^*$ and suppose, by contradiction, that $\rho_H(w) \geqslant \rho_H(w')$. Concavity and the first-order conditions jointly imply that $z_H(w) \geqslant z_H(w')$ which contradicts to

$$\Delta\theta(R \circ \mathcal{K}_H)(\rho_H(w)) + \delta_A(\alpha_H - \alpha_L)z_H(w) = w < w' = \Delta\theta(R \circ \mathcal{K}_H)(\rho_H(w')) + \delta_A(\alpha_H - \alpha_L)z_H(w').$$

Part (b). By the same argument as in Part (a), $\rho_L(\cdot)$ is strictly decreasing on $[0, w_L^*]$, and it is zero afterwards. To complete the proof, let $w_k^o := \max\{w \in W : k_L(w) = 0\}$. The threshold w_k^o is well-defined, because $k_L(\cdot)$ is a continuous function (Claim 1) and $k_L = 0$ is feasible for all values of w.

□

We now turn our attention to $\mathbf{z}(\cdot)$. Our first result establishes uniqueness of transfers, the second completely characterizes the shape of optimal policy.

Claim 3. *$Z_L(\cdot)$ is single-valued. Moreover, if $w_L^* \geqslant w_H^*$, then $Z_H(\cdot)$ is single-valued, otherwise, there exists a number $\bar{w}$ such that $Z_H(w)$ is a singleton if and only if $w \neq \bar{w}$.*

Proof. Uniqueness of $z_L(\cdot)$ is directly implied by the last part of Claim 2. In contrast, $z_H(\cdot)$ might fail to be unique. We now establish the second part of the claim.

First, suppose that $w_L^* \geqslant w_H^*$. Then, $S_j'(w) = (1 - \alpha_j)\rho_L(w) - \alpha_j \rho_H(w)$ is strictly decreasing on $\mathbb{R}_+$. As a result, $Z_H(\cdot)$ is single-valued by strict concavity of S_j.

Second, suppose that $w_L^* < w_H^*$, then the envelope conditions (Equation 8) imply that $S_j'(w) > 0$ on $[0, w_L^*]$, $S_j'(w) < 0$ on $[w_H^*, +\infty)$ and $S_j'(w) = 0$ for any $w \in [w_L^*, w_H^*]$. Define $\bar{w}$ by

$$(\alpha_H - \alpha_L)\delta_A \rho_H(\bar{w}) = \alpha_H(\delta_P - \delta_A).$$

The reader can verify that such value exists, and it is unique, due to of monotonicity of $\rho_H(\cdot)$, which was established in Proposition 2. As a result, $Z_H(\cdot)$ is single-valued on $[0, \bar{w})$ by the last part of Claim 2, and $Z_H(\bar{w}) = [w_L^*, w_H^*]$ by construction. To see that $Z_H(\cdot)$ is single-valued on $(\bar{w}, +\infty)$, note that $w = \Delta\theta(R \circ \mathcal{K}_H)(\rho_H(w)) + \delta_A(\alpha_H - \alpha_L)z_H(w)$ whenever $\rho_H(w) > 0$. Since $\rho_H(w) > 0$ for any $w > \bar{w}$, $z_H(w)$ could be uniquely identified from the upward incentive constraint. □

To sum up, $z_H(w)$ is almost surely unique, it is *not* unique only when $w_L^* < w_H^*$ and $w = \bar{w}$. In what follows, $z_H(\cdot)$ stand for an arbitrary selection from $Z_H(\cdot)$.

Now, the dynamics of promised expected utility are described in Figure 5. In each case z_H and z_L are plotted as functions of w. The 45° line partitions the quadrant into regions where expected utility increases or decreases in the next period. w_H^* and w_L^* are the thresholds as defined above. And the bold dots represent some points in the support of the invariant distribution of the optimal contract. For example, in all the figures the point z_H^e at which $z_H(\cdot)$ intersects the 45° line constitutes a bold dot. Each time a high shock arrives it is possible for the optimal contract to stay at the same expected utility, and it surely does so if the upward constraint is not binding.

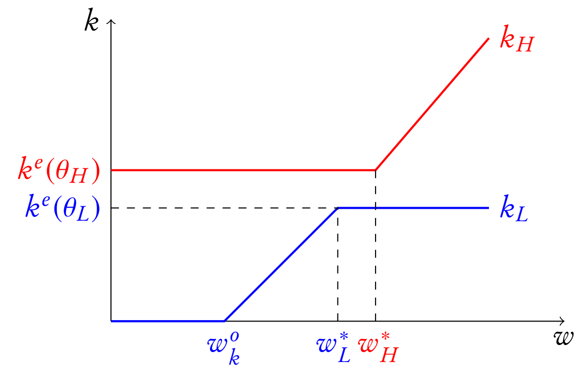

(a) allocation

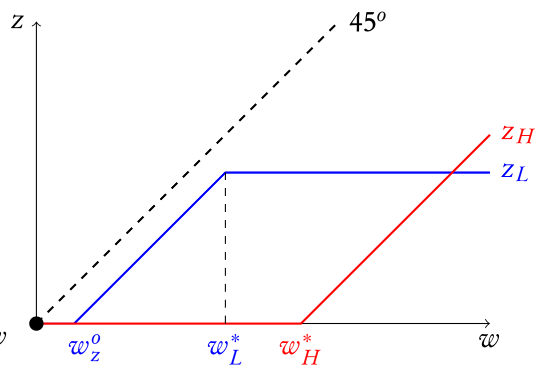

(b) shutdown

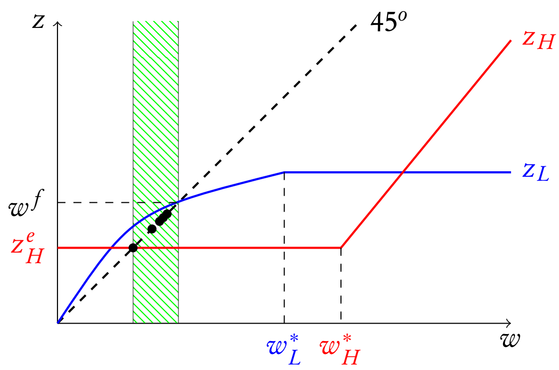

(c) restart

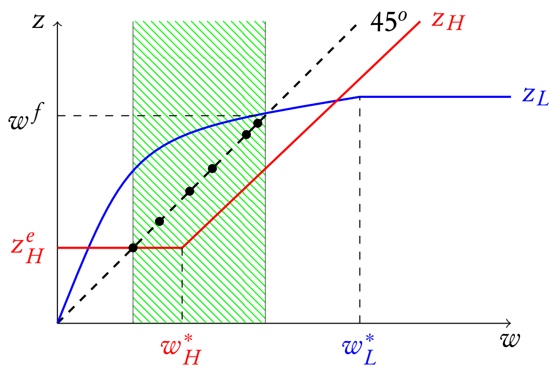

(d) never restart

Figure 5: Optimal recursive contract

Consider first the situation depicted in Figure 5b. Here $z_H^e = 0$. Since both curves lie below the 45°, the recursive contract continually shrinks in expected value. It quickly converges to, most often immediately, to the bold point at zero which implies an expected utility of zero and a complete shutdown of the low productivity type. In Figures 5c and 5d, we portray the optimal restart contract which does not feature shutdowns. The realization of a high shock pushes the expected utility towards z_H^e. On the realization of a low shock, promised expected utility above

w^f, which is the largest fixed point of $z_L(\cdot)$, contracts, and below w^f it expands. The key condition that characterizes Figure 5c is $w^f \leq w_H^*$. It implies that the upward incentive constraint does not bind in the interval $[z_H^e, w^f]$, and the invariant distribution of the promised expected utility rests therein.[29] In contrast, Figure 5d exposits the case with perennial binding of the upward incentive constraint which is captured by the condition $w^f > w_H^*$.

Finally, the only missing piece is initialization- where does the optimal recursive contract start? We show that the recursive contract is initialized at a unique point $w^\star \in [z_H^e, w^f]$. Therefore, at the inception the downward incentive constraint always binds, while the upward constraint may or may not bind. The next proposition summarizes the evolution of expected utility in the optimal recursive contract.

Proposition 3. *The expected utility of the agent in the optimal recursive contract satisfies the following:*

(a) $\exists w_z^o, z_L^e$ *such that* $z_L(w) = 0$ *if and only if* $w \leqslant w_z^o$, $z_L(w) = z_L^e$ *if and only if* $w \geqslant w_L^*$, $z_L(\cdot)$ *is strictly increasing on* $[w_z^o, w_L^*]$.

(b) $\exists z_H^e$ *such that* $z_H(w) = z_H^e$ *if and only if* $w \leqslant w_H^*$, $z_H(\cdot)$ *is strictly increasing on* $[w_H^*, \infty)$.

(c) $z_L(\cdot)$ *has a unique globally stable fixed point* $w^f \in [z_H^e, z_L^*]$, *and* z_H *has a unique fixed point* z_H^e *which is positive if and only if* $\theta_L > \frac{a_L}{1-b}\Delta\theta$.

(d) The thresholds satisfy $z_H^e \leqslant w^f \leqslant z_L^e < w_L^*$, $z_H^e < w_H^*$, *and* $z_L^e \neq z_H^e$ *if and only if* $z_L^e > 0$.

(e) $\exists w^* \in [z_H^e, w^f]$ *such that the optimal contract starts at this point, and it always stays within* $[z_H^e, w^f]$.

Proof of Proposition 3.

Parts (a) and (b). First, we show that $z_L^e < w_L^*$. The claim is vacuously true whenever $z_L^e = 0$, because $w_L^* = \Delta\theta R(k_L^e) > 0$. Consider the alternative case with $z_L^e > 0$. Then, by definition, w_L^* satisfies $S_j'(w_L^*) = -\alpha_j \rho_H(w_L^*) \leqslant 0$, and z_L^e satisfies $S_j'(z_L^e) > 0$. Strict concavity of S_j, which was shown in Part (c) of Claim 2, implies $w_L^* > z_L^e$.

Next, we establish that $z_H^e < w_H^*$. By contradiction, suppose that $w_H^* \leqslant z_H^e$, equivalently we have

$$\frac{\Delta\theta}{1-\delta_A(\alpha_H-\alpha_L)}R(k^e(\theta_H)) \leqslant z_H^e.$$

We claim that $z_H^e \leqslant z_L^e$, therefore $z_H^e < w_L^*$ that implies

$$z_H^e < \frac{\Delta\theta}{1-\delta_A(\alpha_H-\alpha_L)}R(k^e(\theta_L))$$

contradicting to the inequality above as $k_L^e < k_H^e$. To complete the argument, we need to establish that $z_H^e \leqslant z_L^e$. This clearly follows from Equation 8:

$$S_H'(w)/\alpha_H - S_L'(w)/\alpha_L = \frac{\alpha_L-\alpha_H}{\alpha_H\alpha_L}\rho_L(w) \leqslant 0,$$

[29] To find the support, we repeatedly apply $z_L(\cdot)$ to z_H^e, the bold points in Figure 5c depict this set.

In fact, $z^e_L \neq z^e_H$ if and only if $S'_L(0) > \alpha_L \frac{\delta_P - \delta_A}{\delta_P}$.

We showed above that $z^e_j \in [0, w^*_j)$ for $j = H, L$. Then, monotonicity of $\rho_L(\cdot)$ and $\rho_H(\cdot)$, as shown in Proposition 2, combined with Equations 6 and 7 translates into monotonicity of both $z_L(\cdot)$ and $z_H(\cdot)$. Finally, we set $w^o_z := \max\{w \geqslant 0 : z_L(w) = 0\}$ that is uniquely-defined, because $z_L(\cdot)$ is a continuous function with $z_L(0) = 0$.

Part (c). We begin with fixed points of $Z_H(\cdot)$. In the previous part, we showed that $z^e_H < w^*_H$ that implies that z^e_H is a fixed point of $Z_H(\cdot)$. We now show that there are no other fixed points. By contradiction, suppose that there exists another fixed point $w \neq z^e_H > 0$, it must be the case that $\rho_H(w) > 0$. The following equation is necessary for existence of such $w \in \mathbf{Z}_H(w) > 0$ with $\rho_H(w) > 0$:

$$w = \frac{\Delta\theta}{1 - \delta_A(\alpha_H - \alpha_L)} \big(R \circ \mathcal{K}_H\big)\big(\rho_H(w)\big) > \frac{\Delta\theta}{1 - \delta_A(\alpha_H - \alpha_L)} R(k^e_H)$$

To obtain a contradiction, combine Equations 6 and 8:

$$(1 - \alpha_H)\delta_P \rho_L(w) = \alpha_H(\delta_P - \delta_A) + \big(\alpha_H \delta_P - (\alpha_H - \alpha_L)\delta_A\big)\rho_H(w) > 0.$$

Since $\rho_L(w) > 0$, the downward constraint binds this period, and it will keep binding a sequence of θ_L's. Formally, let $z^s_L(w)$ be a result of s successive applications of $z_L(\cdot)$ to w, that is $z^s_L(w) := z_L\big(z^{s-1}_L(w)\big)$ with $z^0_L(w) = w$. By Equation 7, $\rho(z^s_L(w)) > 0$ for any s. Then, iterating along this sequence, we arrive at the following equation:

$$w = \Delta\theta \sum_{\tau=0}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{\tau} \big(R \circ \mathcal{K}\big)\big(\rho_L(z^{\tau}_L(w))\big) < \frac{\Delta\theta}{1 - \delta_A(\alpha_H - \alpha_L)} R(k^e_L).$$

which clearly contradicts the premise. As a result, z^e_H is the unique fixed point of $Z_H(\cdot)$.

Now, we turn our attention to fixed points of $z_L(\cdot)$. Of course, 0 is always a fixed point, and our goal is to identify a positive fixed point, that is $0 < w = z_L(w)$. First of all, $z_L(w) = z^e_L < w^*_L \leqslant w$ whenever $\rho_L(w) = 0$, therefore it must be the case that $w < z^e_L$ and $\rho_L(w) > 0$. The following equation is necessary for existence of a fixed point with $\rho_L(w) > 0$:

$$w = \frac{\Delta\theta}{1 - \delta_A(\alpha_H - \alpha_L)} (R \circ \mathcal{K}_L)(\rho_L(w)).$$

The other necessary condition, due to the Equations 7 and 8, is that

$$\big((1 - \alpha_L)\delta_P - \delta_A(\alpha_H - \alpha_L)\big)\rho_L(w) = \alpha_L(\delta_P - \delta_A) + \alpha_L \delta_P \rho_H(w) > 0.$$

Moreover, the reader can verify that these two equations are jointly sufficient for w to be a positive fixed point of $z_L(\cdot)$. By monotonicity of both ρ_L and ρ_H (shown in Proposition 2), the equations have a root if and only if $\theta_L > \frac{a_L}{1-b}\Delta\theta$. And, if such a root exists, then it is unique.

Let w^f be the largest fixed point, i.e., it is the root of the aforementioned equations for $\theta_L > \frac{a_L}{1-b}\Delta\theta$, and $w^f = 0$, otherwise. For $\theta_L > \frac{a_L}{1-b}\Delta\theta$, global stability follows from $z_L(\cdot)$ crossing the

45-degree line only once and from above, because $w^f < z^e_L$. For $\theta_L/\Delta\theta \leqslant \frac{a_L}{1-b}$, global stability is trivial, because 0 is the unique fixed point.

Part (d). We have already established in Parts (a) and (b) that each $z^e_j < w^*_j$ and $z^e_H \leqslant z^e_L$. So, it remains to establish that $z^e_H \leqslant w^f$. Of course, it is vacuously true whenever $z^e_H = 0$, thus it is without loss of generality, to assume that $z^e_H > 0$. By contradiction, suppose that $w^f < z^e_H$. Since $\rho_H(\cdot)$ is monotone and $z^e_H \leqslant w^*_H$, we have $\rho_H(w^f) = 0$ that implies $\rho_L(w^f) = \frac{a_L}{1-b}$. On the other hand, by monotonicity of $\rho_L(\cdot)$, $z^e_H < w^f$ implies $\rho_L(w^f) > \rho_L(z^e_H) = a_H$. As a result, $\frac{a_L}{1-b} > a_H$ that is a clear contradiction. Conclude that $z^e_H \leqslant w^f$.

Part (e). At the initial date, the first-order conditions with respect to $\mathbf{z}(\cdot)$ coincide with Equations 6 and 7. The extra first condition is $\mathbb{P}(\theta_L)\rho_L(w) - \mathbb{P}(\theta_H)\rho_H(w) = (\leqslant)\mathbb{P}(\theta_H)$ whenever $w > (=)0$. Then, existence and uniqueness directly follows from monotonicity of both $\rho_L(\cdot)$ and ρ_H, see proof of Proposition 2.

We now show $w^* \in [z^e_H, w^f]$. By contradiction, suppose that $w^* < z^e_H$. Since $\frac{P(\theta_H)}{P(\theta_L)} \leqslant a_H$, we must have $\rho_H(w^*) > 0$. Recall that $\rho_H(\cdot)$ is non-decreasing, thus $\rho_H(z^e_H) \geqslant \rho_H(w^*) > 0$ that is a contradiction. Conclude that $w^* \geqslant z^e_H$.

Again, by contradiction, suppose that $w^* > w^f$. Since $\frac{P(\theta_H)}{P(\theta_L)} \geqslant \frac{a_L}{1-b}$, we must have $\rho_H(w^f) > 0$. By monotonicity of $\rho_H(\cdot)$ and $\rho_L(\cdot)$, $\rho_H(w^*) > \rho_H(w^f) > 0$ and $\rho_L(w^*) \leqslant \rho_L(w^f)$ where

$$\rho_L(w^f) = \frac{a_L}{1-b}\left(1 + \frac{1}{1-\delta_A/\delta_P}\rho_H(w^f)\right), \quad \rho_L(w^*) = \frac{P(\theta_H)}{P(\theta_L)}\left(1 + \rho_H(w^*)\right).$$

The reader can verify that these conditions cannot be satisfied simultaneously, as a result we have $w^* \leqslant w^f$. □

Propositions 2 and 3 precisely characterize the optimal contract. Starting at w^*, each subsequent realization of the agent's type determines the optimal allocation according to Proposition 2 and the optimal expected utility for the next period, the state variable, according to Proposition 3.

There is, of course, a one-to-one relationship between the optimal recursive contract, and the sequential optimum. First of all, the downward incentive constraints always bind, and the low type always gets the promised utility of zero. The high type allocation can be distorted only upwards, whereas the low type allocation is always distorted downwards.

Moreover, the realization of each θ_H decreases the promised utility offered to the high type in the next period which reduces distortion for the high type allocation, but increases a distortion in the low type. It takes an endogenous number of consecutive θ_H for the upward incentive constraint to stop binding. θ_L always increases the promised utility offered to the high type in the next period which tightens the distortion for the high type allocation, but relaxes distortions for the low type allocation. It takes an endogenous number of consecutive θ_L shocks for the upward incentive constraint to start binding.

8.2.3 Simplicity

Here the characterization of the optimal recursive contract to is used to establish Theorem 3.

Proof of Theorem 3. First of all, any restart contract is simple, because a number of possible distinct allocations by time T is at most $2T$. Indeed, the set $\{k|\exists\theta^t : k = k_t(\theta^t), t \leqslant T\}$ is a union of $\{\hat{k}_t\}_{t=1}^{T}$, $\{k_t\}_{t=1}^{T-1}$ and k_H. As a result, if the optimum is restart, then it is simple.

We now show the converse. Suppose that the optimal contract is not restart. In term of our recursive notations, this means that $z_H(w^f) \neq z_H^e$. According to Proposition 3, it must be the case that there are no shutdowns, that is $z_H^e > 0$. Thus, the allocation supplied to the low type and the promised utility of the high type are both strictly positive. Since $\mathbf{k}(\cdot)$ is monotone, it suffices to show that the set of utilities promised to θ_H grows at an exponential rate. Formally, we claim that there exists a number K such that

$$\left|\left\{U|\exists\theta^{t-1} : U = U_t(\theta^{t-1}, \theta_H), t \leqslant T\right\}\right| \geqslant K2^T.$$

First of all, note that z_H^e is reached after sufficiently many consecutive high shocks. Since $z_H(w^f) \neq z_H^e$, there exists a natural number τ such that $z_H(z_L^\tau(z_H^e)) \neq z_H^e$. Moreover, Proposition 3 implies that for any $w, w' \in [z_H^e, w^f)$ with $w \neq w'$, we have that $z_H(z_L^\tau(w)) \neq z_H(z_L^\tau(w')) \neq z_H^e$. In other words, the number of states is doubled every τ periods. As a result, the state space expands exponentially with the constant $K = 2^{-\tau}$. □

8.3 Comparative statics

Proof of Corollary 6. First of all, note that the optimal contract only depends on the agent's relative patience β and absolute patience $\beta\delta_P$. By the theorem of maximum, the contract is a continuous function of the agent's patience. If $\beta \to 1$, then the optimum convergence to the first-order optimum, because the latter is always incentive compatible for $\beta = 1$. Moreover, by Theorem 1, the first-order optimum exhibits distortions only along the lowest history, that is $\rho_t = 0$ at all dates. Next, as $\delta_P \to 1$, the weight on the payoffs along the lowest history goes to zero. As a result, the principal's achieves the maximal surplus.

We now show that $\delta_P\beta \to 1$ is necessary for the full surplus extraction. By construction, the value of the first-order program in an upper bound on $\Pi^\star$. Since $\delta_P\beta < 1$, the distortions along the lowest history are positive. Thus, the principal's profit is strictly less than the surplus: $s^e > \Pi^\# \geqslant \Pi^*$. □

Proof of Corollary 7. We first consider the first-order optimum. This contract is essentially static for $\alpha = \frac{1}{2}$, see Theorem 1: $\rho_t = \frac{\delta_P - \delta_A}{\delta_P}$ for all t, $\hat{\rho}_t = \frac{\delta_P - \delta_A}{\delta_P}$ for all $t \geqslant 2$. Importantly, $\bar{U}_A$ is independent of δ_A. Since the surplus and the costs of incentive provision are both strictly increasing in δ_A, $\delta_A = \delta_P$ uniquely maximizes the principal's profit. By continuity of the profit in the first-order optimal contract with respect to α, $\delta_A = \delta_P$ is still a maximizer for any α in a sufficiently small neighborhood of 1/2.

If $\alpha \to 1$, then $\hat{\rho}_t = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)}\left(\frac{\delta_A}{\delta_P}\right)^{t-1}$ for all t, thus the intertemporal cost of incentive provision goes to zero. As a result, $\lim_{\alpha\to 1} \bar{U}_P = \lim_{\alpha\to 1} \bar{U}_A$, and the limit is strictly increasing in δ_A. By continuity, $\delta_A = 0$ is a maximizer for any α in a sufficiently small neighborhood of 1.

Recall that the first-order optimal contract is incentive compatible for either iid or constant types, see Corollary 3. Therefore, the proposition is true for the optimum and the restart optimum as well. □

References

D. Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56(2): 383–396, 1988.

D. Abreu and A. Rubinstein. The structure of Nash equilibrium in repeated games with finite automata. *Econometrica*, 56(6):1259–1281, 1988.

D. Acemoglu, M. Golosov, and A. Tsyvinski. Political economy of mechanisms. *Econometrica*, 76 (3):619–641, 2008.

M. Aguiar, M. Amador, and G. Gopinath. Investment cycles and sovereign debt overhang. *Review of Economic Studies*, 76(1):1–31, 2009.

S. Athey and I. Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.

A. Atkeson and R. E. Lucas. On efficient distribution with private information. *Review of Economic Studies*, 59(3):427–453, 1992.

D. P. Baron and D. Besanko. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267–302, 1984.

M. Battaglini. Long-term contracting with markovian consumers. *American Economic Review*, 95 (3):637–658, 2005.

L. M. Benveniste and J. A. Scheinkman. On the differentiability of the value function in dynamic models of economics. *Econometrica*, 47:727–732, 1979.

D. Bergemann and J. Välimäki. Dynamic mechanism design: an introduction. *Journal of Economic Literature*, forthcoming, 2019.

D. Besanko. Multi-period contracts between principal and agent with adverse selection. *Economics Letters*, 17(1-2):33–37, 1985.

B. Biais, T. Mariotti, G. Plantin, and J.-C. Rochet. Dynamic security design: Convergence to continuous time and asset pricing implications. *Review of Economic Studies*, 74(2):345–390, 2007.

S. Chassang. Calibrated incentive contracts. *Econometrica*, 81(5):1935–1971, 2013.

G. L. Clementi and H. A. Hopenhayn. A theory of financing constraints and firm dynamics. *Quarterly Journal of Economics*, 121:229–265, 2006.

P. M. DeMarzo and Y. Sannikov. Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance*, 61(6):2681–2724, 2006.

A. K. Dixit. *Optimization in economic theory*. Oxford University Press, 1990.

E. Farhi and I. Werning. Inequality and social discounting. *Journal of Political Economy*, 115(3): 365–402, 2007.

A. Fernandes and C. Phelan. A recursive formulation for repeated agency with history dependence. *Journal of Economic Theory*, 91(2):223–247, 2000.

D. Garrett, A. Pavan, and J. Toikka. Robust predictions of dynamic optimal contracts. Toulouse School of Economics, Northwestern University and MIT, 2018.

J. Hörner, T. Sugaya, S. Takahashi, and N. Vieille. Recursive methods in discounted stochastic games: An algorithm for $\delta \to 1$ and a folk theorem. *Econometrica*, 79(4):1277–1318, 2011.

D. Krähmer and R. Strausz. Optimal sales contracts with withdrawal rights. *Review of Economic Studies*, 82(2):762–790, 2015.

I. Krasikov and R. Lamba. A theory of dynamic contracting with financial constraints. Pennsylvania State University, 2018.

R. V. Krishna, G. Lopomo, and C. Taylor. Stairway to heaven or highway to hell: Liquidity, sweat equity, and the uncertain path to ownership. *RAND Journal of Economics*, 44(1):104–127, 2013.

D. Krueger and H. Uhlig. Competitive risk sharing contracts with one-sided commitment. *Journal of Monetary Economics*, 53(7):1661–1691, 2006.

J.-J. Laffont and D. Martimort. *The theory of incentives: the principal-agent model*. Princeton university press, 2002.

J.-J. Laffont and J. Tirole. *A theory of incentives in procurement and regulation*. MIT press, 1993.

J.-J. Laffont and J. Tirole. Pollution permits and compliance strategies. *Journal of Public Economics*, 62(1-2):85–125, 1996.

E. Lehrer and A. Pauzner. Repeated games with differential time preferences. *Econometrica*, 67(2): 393–412, 1999.

K. Madarász and A. Prat. Sellers with misspecified models. *Review of Economic Studies*, 84(2): 790–815, 2017.

F. Modigliani and M. H. Miller. The cost of capital, corporation finance and the theory of investment. *American Economic Review*, 48(3):261–297, 1958.

E. F. Moore. Gedanken-experiments on sequential machines. In C. E. Shannon and J. McCarthy, editors, *Automata Studies, Annals of Mathematics Studies Volume 34*, page 129ÃŘ153. Princeton University Press, 1956.

O. Morand and K. Reffett. Lagrange multipliers in convex programs with applications to classical and nonoptimal stochastic one-sector growth models. University of Connecticut and Arizona State University, 2015.

M. Mussa and S. Rosen. Monopoly and product quality. *Journal of Economic Theory*, 18(2):301–317, 1978.

R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

M. M. Opp and J. Y. Zhu. Impatience versus incentives. *Econometrica*, 83(4):1601–1617, 2015.

A. Pavan, I. Segal, and J. Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014.

D. Ray. The time structure of self-enforcing agreements. *Econometrica*, 70(2):547–582, 2002.

J. P. Rincón-Zapatero and M. S. Santos. Differentiability of the value function without interiority assumptions. *Journal of Economic Theory*, 144(5):1948–1964, 2009.

S. E. Spear and S. Srivastava. On repeated moral hazard with discounting. *Review of Economic Studies*, 54(4):599–617, 1987.

N. L. Stokey, R. E. Lucas Jr, and E. Prescott. *Recursive methods in economic dynamics*. Harvard University Press, 1989.

T. Sugaya. Characterizing the limit set of perfect and public equilibrium payoffs with unequal discounting. *Theoretical Economics*, 10(3):691–717, 2015.