Implications of unequal discounting in dynamic contracting*

Ilia Krasikov Rohit Lamba Thomas Mettral

August 2021

Abstract

Dynamic contracts typically allow the principal to relax future incentive constraints by backloading the agent's information rents or asking the agent to post a bond upfront which is liquidated over time. An implicit modeling assumption at play there is that both the principal and agent have equal access to capital, captured by equal discount rates. This paper introduces unequal discounting in a canonical dynamic screening problem where the agent has Markovian private information and limited commitment. The backloading force is tempered by an inter-temporal cost of incentive provision. The optimal contract features cycles with infinite memory. The interaction of Marokvian information and unequal discounting introduces technical challenges by rendering the standard relaxed problem approach invalid for certain parameters. An approximately optimal and simple alternative is provided, where both terms are made formalized.

1 Introduction

Discounting is central to the modeling of dynamic interactions. Other than the direct psychological motivation that humans beings value the present more than the future, the most obvious interpretation of discounting factor is given by interest rates. This is captured by the proverbial identity $\delta = e^{-r} \approx \frac{1}{1+r}$, where δ is the agent's discount factor and r the interest rate in the market.

Unequal discounting then typically attempts to capture one of three forces. In principal-agent settings it allows one party to have better access to capital markets than the other through lower interest rates (eg. Krueger and Uhlig [2006] and Biais, Mariotti, Plantin, and Rochet [2007]). In planning problems, it allows the planner to have longer or shorter time horizon than citizens (eg. Farhi and Werning [2007] and Acemoglu, Golosov, and Tsyvinski

*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*.

[2008]). In pure strategic settings such as bargaining and reputation, it seeks to emulate a power (im)balance between players (eg. Rubinstein [1982] and Fudenberg and Levine [1989]). The goal of this paper is to invoke the first interpretation and under it guise understand how unequal discounting impacts the standard predictions in dynamic mechanism design.

To that end, we focus on a canonical dynamic screening model inspired by Battaglini [2005]. A principal (here a "large" firm) supplies capital that is critical for the production of a final good by an agent (here a "small" firm). A production technology shock is privately observed by the agent, and its follows a two-state Markov process. The principal posts 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. In a departure from the canonical model, the principal is more patient than the agent.

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 delivered a downward distortion, that is an allocation less than than the efficient one.[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 restarts the cycles of previous distortions, and so on. This is in contrast to dynamic mechanism design models with equal discounting that predict eventually vanishing distortions.[3] We provide a brief intuition for this result.

A salient trade-off in much of contract theory is rent-versus-efficiency.[4] How much efficiency to give up in order to reduce the information rent of the agent. Typically, dynamic contracting allows the principal to gradually resolve the tradeoff in favor of the efficiency by *backloading* the agent's payoffs— this reduces shadow price of providing information rent in

[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]Laffont and Martimort [2002] write: "[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 long-run.[5] Unequal discounting limits the extent to which the principal can exploit this instrument by introducing a competing *frontloading* force.

Suppose interest rates faced by the two parties are r_P and r_A respectively. If the principal promises to pay the agent an expected information rent of x tomorrow, limited commitment implies she can extract a maximum of $\frac{x}{1+r_A}$ today. This generates an account of $\frac{x}{1+r_A} - \frac{x}{1+r_P} = -\left(\frac{1}{1+r_P} - \frac{1}{1+r_A}\right)x$ for the principal, which is negative when $r_P < r_A$. We call this the *intertermporal cost of incentive provision*. For any positive value of x, the intertemporal cost creates a new wedge wherein the principal always wants to backload more than what is feasible. So the shadow price for backloading is perennially positive owing to the frontloading push caused by unequal discounting.

These countervailing forces settle on to a compromise culminating in the restart contract. The standard backolading force generates decreasing distortions for consecutive low shocks, which are erased on the arrival of a high shock. However, the interremporal cost of incentive provision never allows the shadow price of providing incentives to ever disappear. So, another low shock after a high one restarts the cycles of distortions.

Now, a vital feature of the model here is role of persistence in the agent's technology. If the production technology is iid over time, the model is relatively uninteresting: the high type gets the efficient allocation as before and the low type gets a static downward distortion that is independent of history. Distortions are still cyclical, but trivially so, with no memory. The interaction of Markovian types with unequal discounting, however, creates richer predictions: a sequence of decreasing distortions with infinite memory, which are cyclical around the arrival of a high shock.[6]

The second result pins 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: Consider a two period version of our model. In the standard model, distortions for the low type disappear once a high shock is realized, so efficient allocation is delivered for history HL, whereas distortions persist for the low shock history, L and LL. With unequal discounting new distortions are introduced for HL, and these can be high enough so that the capital allocation for this history is lower than LL even though the former is better in terms of productivity shocks. This non-monotonicity can violate the incentive constraints ignored in the relaxed problem, and the force is stronger

[5] 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]).

[6] If the model has iid shocks and equal discounting, then it delivers zero distortions from the second period onwards. This is in some sense the reason why much of the dynamic financial contracting literature, which features iid technologies (eg. Clementi and Hopenhayn [2006]) focusses on stronger forms of limited commitment on the agent's side. These take the form of limited liability constraints as opposed to forward looking individual rationality, since the latter do not generate any interesting tension in the the models. More on financial contracting later.

for the infinite horizon model. The optimum in such a case demands an upwardly distorted allocation for the high type, larger than its efficient counterpart.

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" alternative 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) 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.

The purpose of restricting attention to restart contracts is fourfold. First, 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. Second, the idea of restart contracts generated through unequal discounting connects to a sizeable literature in economics, particularly public finance, political economy and sovereign debt.[7] Third, restart contracts arise naturally as the solution to the relaxed problem, which is quite the standard in contract theory. Optimal restart contracts aim to keep that basic structure in tact while ensuring approximate global optimality. And, finally, when the relaxed problem approach fails, even though the optimal contract gets quite complicated, the optimal restart contract continues to be simple in a precise sense that forms the next result fo the paper.

The fourth result formalizes this notion of simplicity for dynamic contracts. Here we take a 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

[7] 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. 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. Aguiar, Amador, and Gopinath [2009] show that in a model of sovereign debt and foreign capital, if the government cannot commit and is more impatient than the market, then long-run cycles emerge in sovereign debt and foreign capital investments. In recent work, Kapon [2021] studies a model of dynamically arriving agents who are screened for crimes they have committed using amnesty programs that feature the restart property; here stochastic arrivals as opposed to unequal discounting creates cycles.

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.

A few observations on the modeling approach and its interpretations are in order. A sizeable 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 the problem of agency in the design of institutions when the principal has some commitment power (see Bergemann and Välimäki [2019] for a recent survey). 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).

Further still, unequal discounting can be thought of as a financial constraint. In the standard model, as in Battaglini [2005], it is costless to move transfers across time. Two modeling ingredients generate this: First limited commitment for the agent captured by the individual rationality constraint imposes a lower bound on current utility plus continuation value. Second, the principal and agent have the same discount factor. A consequence of these assumptions is that transfers are not uniquely pinned down— individual rationality binds only in the first period and is slack forever after. Imposing either a stronger notion of limited commitment or unequal discounting introduces a friction in the model that breaks this irrelevance in timing of payments, while still allowing for some tractability afforded by the quasi-linear framework.

In recent work (Krasikov and Lamba [2021]) we have explored the former route by requiring a stronger form of limited commitment— the stage payoff or current utility of the agent must be non-negative every period.[8] These are hard financial constraints since the agent simply cannot borrow beyond his daily working capital. In addition, persistence in technology shocks makes financial constraints bind for a long time. But eventually, through maximal backloading of incentives, distortions disappear and efficiency is achieved.

In this paper, we take the latter route by requiring the movement of payments across time to be evaluated at different rates for the principal and agent. This too puts restriction on the principal's ability to backload incentives to the extent desirable in the standard model. The financial constraint here is soft since the weaker form of individual rationality allows the agent to borrow beyond what is required for current working capital. However, this borrowing is at a rate worse than the market rate or the rate available to the principal. An inter-temporal cost of incentive provision is then created. It disappears, as in the standard model, when a high shock is realized, but unlike the standard model, the financial constraints imposed by unequal discounting never allows it go away, a low shocks starts the cycle of distortions

[8]See also Clementi and Hopenhayn [2006], Krishna, Lopomo, and Taylor [2013] and Krähmer and Strausz [2015] for related models.

again.[9]

Finally, the impact of unequal discounting has been studied in other related works. In repeated games, the set of equilibrium payoffs expands favoring the patient player (see Lehrer and Pauzner [1999]). 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, as discussed above, the literatures on public finance, political economy and sovereign debt have used unequal discounting to generate allocations that parallel restart contracts either in their cyclicality or persistence of long-run distortion.[10]

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.[11] 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.[12] The principal's stage utility is simply $p - k$.[13] The static surplus is denoted by $s(\theta, k) := \theta R(k) - k$. We consider an infinite horizon setting where the principal and agent discount future utility. However, we *do not restrict them to have the same discount factor*; these are denoted by δ_P and δ_A, respectively, where $\delta_P \geqslant \delta_A$.[14]

[9]In taking financial constraints seriously and breaking the inter-temporal linearity of transfers, while maintaining quasi-linear structure on preferences, both Krasikov and Lamba [2021]) and this paper seek to bring ideas from dynamic financial contracting (see Sannikov [2013] for a survey) to dynamic mechanism design.

[10]Note that there is also a separate literature on hyperbolic discounting, starting with Laibson [1997], that has also been used in the study of contracts (eg. (Gottlieb and Zhang [2021]), political economy (eg. Bisin, Lizzeri, and Yariv [2015]), and more.

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

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

[13]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.

[14]As mentioned in the introduction, 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. 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

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, $\mathbb{P}(\theta_H|\theta_j) = \alpha_j$, 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.

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 $\theta_{\mathbf{j}}^{\mathbf{t}}$.[15] 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.[16]

2.2 Constraints

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

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

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:[17]

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

present." So under this interpretation, different discount factors automatically refer to access to different interest rates.

[15] 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^{∞}.

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

[17] 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.

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

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

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 counterpart; the latter records the fact that 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 into the contractual relationship at any point. 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

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

$\bar{U}_A$ 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 is the notion of Myersonian virtual value (Myerson [1981]). Recollect that the static surplus is given by $s(\theta, k) = \theta R(k) - k$. In our *quasi-linear* environment, define

$$\mathcal{K}_H(\rho) := arg \max_{k \geqslant 0} \ s(\theta_H + \rho\Delta\theta, k), \quad \mathcal{K}_L(\rho) := arg \max_{k \geqslant 0} \ s(\theta_L - \rho\Delta\theta, k).$$

Then, $\theta_H + \rho\Delta\theta$ is the virtual value associated with the high type and $\theta_L - \rho\Delta\theta$ - with the low type. Here $\rho \geq 0$ measures the level of distortion arising out of information asymmetry, and it is pinned down by the set of binding constraints at the optimum. Concavity of R implies that $\mathcal{K}_H$ is an increasing and $\mathcal{K}_L$ a decreasing function of ρ:

$$\mathcal{K}_H(\rho) = (R')^{-1}\left(\frac{1}{\theta_H + \rho\Delta\theta}\right) \text{ and } \mathcal{K}_L(\rho) = (R')^{-1}\left(\frac{1}{\theta_L - \rho\Delta\theta}\right) \text{ for } \rho\Delta\theta < \theta_L, \text{ zero otherwise.}$$

The efficient allocations are given by $k^e_H := \mathcal{K}_H(0)$ and $k^e_L := \mathcal{K}_L(0)$, i.e.,

$$\theta_j R'(k^e_j) = 1 \quad \text{for } j = H, L.$$

3.2 Static problem

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

$$\max_{\langle \mathbf{k}, \mathbf{u} \rangle} \quad \sum_{j=H,L} (s(\theta_j, k(\theta_j)) - u(\theta_j))\mathbb{P}(\theta_j) \quad \text{subject to} \quad \mathbf{k} \geqslant 0 \text{ and } IC_H, IR_H, IC_L, IR_L,$$

where we dropped time subscripts, i.e., k stands for k_1, etc. It is well known that we can look at a relaxed problem where we maximize the objective subject only to $IC_H : u(\theta_H) \geqslant \Delta\theta R(k(\theta_L)) + u(\theta_L)$ and $IR_L : u(\theta_L) \geqslant 0$. Both these constraints hold as equalities, the objective can be re-written as

$$\sum_{j=H,L} s(\theta_j, k(\theta_j))\mathbb{P}(\theta_j) - \Delta\theta R(k(\theta_L))\mathbb{P}(\theta_H) = s(\theta_H, k(\theta_H))\mathbb{P}(\theta_H) + s\left(\theta_L - \Delta\theta\frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)}, k(\theta_L)\right)\mathbb{P}(\theta_L).$$

Using the notation of virtual valuation above, the optimal allocation rule is then given by $k^\star(\theta_H) = k^e_H = \mathcal{K}_H(0)$, and $k^\star(\theta_L) = \mathcal{K}_L(\hat{\rho})$ for $\hat{\rho} = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)} = \frac{\alpha_L}{1-\alpha_H}$. The rent-versus-efficiency 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) 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 downward incentive constraints and individual rationality constraint of the low type:

$$\max_{\langle \mathbf{k}, \mathbf{U} \rangle} \quad \bar{S} - \bar{U}_A - I \quad \text{subject to} \quad \mathbf{k} \geqslant 0 \text{ and } IC_H, IR_L.^{18}$$

We show in the appendix that both IC_H and IR_L bind at the solution to the relaxed problem. These binding constraints can then be used to substitute for U_1 and u_2 in the above expression for $\bar{U}_A$ and I, which now become functions of the primitives and allocation rule, $\mathbf{k}$. Finally, we optimize to obtain the allocations that solve the relaxed problem. The solution is recorded in the following proposition. Recollect that we define $\theta_H R'(k^e_H) = 1$ and $\mathcal{K}_L(\rho) = (R')^{-1}\left(\frac{1}{\theta_L - \rho\Delta\theta}\right)$ for $\rho\Delta\theta < \theta_L$, zero otherwise.

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

a) $k^{\#}_t(\theta^{\mathbf{t-1}}, \theta_H) = k^e_H$ *for all* $\forall \theta^{\mathbf{t-1}}$;

b) $k^{\#}_t(\theta^{\mathbf{t}}_{\mathbf{L}}) = \mathcal{K}_L(\hat{\rho}_t)$ *for* $\hat{\rho}_2 = \frac{\delta_A}{\delta_P}\frac{\alpha_H - \alpha_L}{1-\alpha_L}\hat{\rho}_1 + \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_L}{1-\alpha_L}$, $\hat{\rho}_1 = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_H)}$;

c) $k^{\#}_t(\theta_H, \theta_L) = \mathcal{K}_L(\rho_1)$ *for* $\rho_1 = \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_H}{1-\alpha_H}$.

The high type is always supplied the efficient allocation, and the supply to the low type is distorted downwards. The distortions are pervasive in that $k_t(\theta^{\mathbf{t-1}}, \theta_L) < k^e_L$ for all $\theta^{\mathbf{t-1}}$. We now present the intuition for how these distortions are generated.

Start with the case where $\delta_P = \delta_A$, then there is no intertemporal cost of incentive provision, that is $\bar{U}_P = \bar{U}_A$. The binding IR_L and IC_H constraints give us

$$U_1(\theta_L) = 0,\ U_1(\theta_H) = \Delta\theta R(k_1(\theta_L)) + \delta_A(\alpha_H - \alpha_L)\Delta\theta R(k_2(\theta^{\mathbf{2}}_{\mathbf{L}})) \text{ and } \bar{U}_A = U_1(\theta_H)\mathbb{P}(\theta_H).$$

The only allocations that appear in the rents are $k_1(\theta_L)$ and $k_2(\theta^{\mathbf{2}}_{\mathbf{L}})$, for which the optimal distortions are positive. In the first period, this is captured by the coefficient $\hat{\rho}_1$, which

[18] The statement of the problem is the same as in Section 2.3 with the only difference that $T = \infty$ is replaced by $T = 2$ in the all the summations.

is the same as in the static model. Due to persistence in private information, $\alpha_H \geqslant \alpha_L$, the distortion propagates to the second period along the history of consecutive low shocks: $\hat{\rho}_2 = \frac{\alpha_H - \alpha_L}{1-\alpha_L}\hat{\rho}_1$. It is easy to see that $\frac{\alpha_H-\alpha_L}{1-\alpha_L} < 1$, thus the distortions are decreasing along the low history. Note that if $\alpha_H = \alpha_L$, so the model is iid, there are no distortions in the second period. Hence, persistence is critical for the propagation of distortions along the history of consecutive low shocks. However, the distortions are muted after the high shock, i.e., $\rho_1 = 0$ and $k_2(\theta_H, \theta_L) = k^e_L$. 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 fixed allocations $k_1(\theta_L)$ and $k_2(\theta^{\mathbf{2}}_{\mathbf{L}})$ the value of $\bar{U}_A$ goes down. However, now the principal also incurs the intertemporal cost of incentive provision, which through the binding IC_H and IR_L constraints of the second period is given by

$$I = (\delta_P - \delta_A)\left(\Delta\theta R(k(\theta_H, \theta_L)\alpha_H \mathbb{P}(\theta_H) + \Delta\theta R(k_2(\theta^{\mathbf{2}}_{\mathbf{L}}))\alpha_L \mathbb{P}(\theta_L)\right).$$

This interaction of IC_H and IR_L in the second period generates new distortions for the allocations that appear in the expression for I. In Proposition 1, $a_H := \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_H}{1-\alpha_H}$ represents the coefficient of the distortion for allocation $k_2(\theta_H, \theta_L)$, and $a_L := \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_L}{1-\alpha_L}$ represents the added distortion for allocation $k_2(\theta^{\mathbf{2}}_{\mathbf{L}})$.

To summarize, the low type in the first period is delivered a static distortion $\hat{\rho}_1$. 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 $b := \frac{\delta_A}{\delta_P}\frac{\alpha_H - \alpha_L}{1-\alpha_L}$, resulting in $\hat{\rho}_2 = b\hat{\rho}_1 + a_L$. 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 $\rho_1 = a_H$. In the next section we will see that these four terms— *starter* $\hat{\rho}_1$, *propagator* b, *adder* a_L and *seed* a_H— define all distortions for the infinite horizon model. But before that, two final thoughts for the two-period problem.

Once $\mathbf{k}^{\#}$ is determined by Proposition 1, the utilities, $\mathbf{U}^{\#}$, 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. Moreover, it is possible that the seed a_H is large enough so that at the first-order optimum, $k^{\#}_2(\theta_H, \theta_L) \ll k^{\#}_2(\theta^{\mathbf{2}}_{\mathbf{L}})$. Then, even though (θ_H, θ_L) is a "better" history than $\theta^{\mathbf{2}}_{\mathbf{L}}$ 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. The resulting inequality delivers a condition on the primitives.

4 The relaxed problem

4.1 Optimal contract

As we did in the two-period problem, we start with the standard relaxed problem (or first-order) approach, 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 an account. The goal now is to express the profit solely as a function of allocations by substituting away transfers from the set of binding constraints.

Start by rewriting IC_H as follows:

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

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:

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

Equation (1) gives the expression for the principal's expected payment in terms of the allocations:

$$\bar{U}_A + I = \mathbb{E}[U_1(\theta_1)] + (\delta_P - \delta_A) \sum_{t=2}^{\infty} \delta_P^{t-2} \mathbb{E}\left[U_t(\theta^{\mathbf{t}})\right] = \tag{2}$$

$$= \sum_{t=1}^{\infty} \delta_P^{t-1} \cdot \hat{\rho}_t \cdot \Delta\theta R\left(k_t(\theta_{\mathbf{L}}^{\mathbf{t}})\right) \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}) + \sum_{\theta^{\mathbf{t-1}}} \sum_{s=1}^{\infty} \delta_P^{t-1+s} \cdot \rho_s \cdot \Delta\theta R\left(k_{t+s}(\theta^{\mathbf{t-1}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s}})\right) \mathbb{P}(\theta^{\mathbf{t-1}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s}}),$$

where $\{\hat{\rho}_t\}$ and $\{\rho_t\}$ are (deterministic) measures of agent's information rents, respectively for the lowest history where no high type is ever realized and all other histories where at least one high type occured at some point. Equation (2) implies an important property of the sequential structure of distortions: Any two histories with the same time since the last high shock are isomorphic. So, the distortion for both histories $(\theta_L, \theta_H, \theta_L)$ and $(\theta_L, \theta_H, \theta_H, \theta_L)$ is given by ρ_1, and more generally for any arbitrary history $(\theta^{\mathbf{t-1}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s}})$ it is given by ρ_s.

Recall that we define $\theta_H R'(k_H^e) = 1$ and $\mathcal{K}_L(\rho) = (R')^{-1}\left(\frac{1}{\theta_L - \rho\Delta\theta}\right)$ for $\rho\Delta\theta < \theta_L$, zero otherwise. Optimizing the objective over the set of allocations, we have the following result.

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

a) $k_t^{\#}(\theta^{\mathbf{t-1}}, \theta_H) = k_H^e$ *for all* $\forall \theta^{\mathbf{t-1}}$*;*

b) $k_t^{\#}(\theta_{\mathbf{L}}^{\mathbf{t}}) = \mathcal{K}_L(\hat{\rho}_t)$ *for* $\hat{\rho}_t = b\hat{\rho}_{t-1} + a_L$*,* $\hat{\rho}_1 = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_H)}$*;*

c) $k_{t+s}^{\#}(\theta^{\mathbf{t-1}}\theta_H, \theta_{\mathbf{L}}^{\mathbf{s}}) = \mathcal{K}_L(\rho_s)$ *for all* $\forall \theta^{\mathbf{t-1}}$ *for* $\rho_t = b\rho_{t-1} + a_L$*,* $\rho_1 = a_H$*, 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, whereas the low type allocations are distorted by $\{\hat{\rho}_t\}$ along the lowest history and by $\{\rho_t\}$ if at least one high shock arrived in the past. The closed form expressions for $\{\hat{\rho}_t\}$ and $\{\rho_t\}$ are provided in Theorem 1. 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 $\hat{\rho}_t = b\hat{\rho}_{t-1} + a_L$. Once a high shock arrives all previous distortions are erased. The realization of a new low type leads to a new *seed* distortion $\rho_1 = a_H$, which is then propagated and added to as before for consecutive low shocks: $\rho_t = b\rho_{t-1} + a_L$. We now study the properties of the dynamic contract more closely, in particular the cyclicality of optimal distortions.

4.2 Properties of the optimal contract

The *seed*, *adder*, and *propagator* lend a cyclical pattern to the optimal distortions, a property which we will term *restart*, and the contract that adopts this property to be a *restart contract*. Figure 1 illustrates graphically the workings of a general restart contract. The contract starts in the white circle on the left. 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 that 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 $\{k_t\}$. The first-order optimal allocations, identified in Theorem 1, constitute a restart contract. The high type allocation is efficient, i.e., $k_H^{\#} := k_H^e$, whereas the low type allocations are given by two sequences: $\hat{k}_t^{\#} := \mathcal{K}_L(\hat{\rho}_t)$ and $k_t^{\#} := \mathcal{K}_L(\rho_t)$.

The *starter* and the *propagator* are present in the standard equal discounting model, whereas the *adder* and *seed* are created by unequal discounting. This 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 had 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] calls these properties *generalized no distortion at the top* and *vanishing*

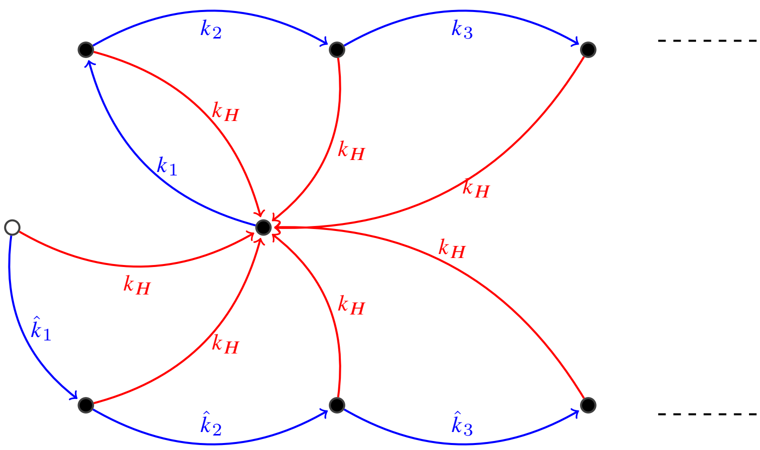

Figure 1: The evolution of allocation in a restart contract. A red arrow indicates a transition, because of a high report, and blue a transition because of a low report.

distortions at the bottom, respectively. In terms of Theorem 1, these respectively imply $k_H^\# = k_H^e$ and $\hat{k}_{t+1}^\# \geqslant \hat{k}_t^\#$ with $\hat{k}_t^\# \to k_L^e$.

When $\delta_P > \delta_A$, we have $\bar{U}_P = \bar{U}_A + I$, where $\bar{U}_A$ is the net present value of standard information rent and I is the intertemporal costs of incentive provision. While $\bar{U}_A$ is now reduced from the perspective of the principal, I is an added cost which ensures that IR_H and IR_L bind permanently. Put differently, for any fixed sequence of allocations, $\bar{U}_A$ is decreasing in δ_A and I is increasing in δ_A. These competing forces interact to endogenously determine the optimal level of allocative distortions presented in Theorem 1. The main conceptual departure from the standard model is that backloading is now constrained by how much of the dynamic information rent the principal can extract upfront since the agent values the future less.

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 [2021], we have explored this model, under persistence. In contrast, here we allow for the more permissive individual rationality constraint $\mathbf{U} \geqslant 0$, so that movement of transfers across time is feasible, but it 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 1. *A contract* $\langle \mathbf{k}, \mathbf{U} \rangle$ *is said to be* **shutdown** *if* $\liminf_{t\to\infty} \mathbb{P}\left(k_t(\theta^{\mathbf{t-1}}, \theta_L) = 0\right) \in (0, 1]$, *and it is said to be* **permanently shutdown** *if* $\liminf_{t\to\infty} \mathbb{P}\left(k_t(\theta^{\mathbf{t-1}}, \theta_L) = 0\right) = 1$.

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, i.e., $k_t^{\#} = 0$ *for some* t*, whenever* $\theta_L \le \rho_1 \Delta\theta$*;*

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

The evolution of distortions here is distinct than both the equal discounting model without financial constraints (eg. Battaglini [2005] and Pavan, Segal, and Toikka [2014]), and the equal discounting model with hard financial constraints (eg. Krishna, Lopomo, and Taylor [2013] and Krasikov and Lamba [2021]). In the former case, depending on the generality of the model, distortions are monotonically decreasing and the efficient allocation is reached in the limit, either along every history, almost surely, or at least on average. In the latter case the distortions are monotonically *increasing* for consecutive bad (or low) shocks, but the contract still does converge almost surely to the efficient allocation. Thus, in their pervasiveness, the distortion dynamics here is distinct from both cases, and decreasing distortions for successive low shocks is reminiscent of the former case.

4.3 Validity of the relaxed problem approach

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, which is documented in the following simple result.

Corollary 2. *The first-order optimal payments are as follows:*

a) $U_t^{\#}(\theta^{\mathbf{t-1}}, \theta_L) = 0$ *for all* $\theta^{\mathbf{t-1}}$*;*

b) $U_t^{\#}(\theta_{\mathbf{L}}^{\mathbf{t-1}}, \theta_H) = \hat{U}_t$ *for* $\hat{U}_t^{\#} := \Delta\theta \sum_{s=t}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-t} R(\hat{k}_s^{\#})$*;*

c) $U_{t+s}^{\#}(\theta^{\mathbf{t-1}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s-1}}, \theta_H) = U_s^{\#}$ *for all* $\theta^{\mathbf{t-1}}$ *for* $U_t^{\#} := \Delta\theta \sum_{s=t}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-t} R(k_s^{\#})$, *where* $\hat{k}_t^{\#} = \mathcal{K}_L(\hat{\rho}_t)$ *and* $k_t^{\#} = \mathcal{K}_L(\rho_t)$*,* $\{\hat{\rho}_t\}$ *and* $\{\rho_t\}$ *are as defined in Theorem 1.*

The low type always gets zero expected payoff. Similar to the optimal distortions, the high type's payoff is determined by two sequences, $\{\hat{U}_t\}$ and $\{U_t\}$: The first determines utility along the history of consecutive low shocks, and the second, as a function of the number of low shocks since the last high shock. We use Corollary 2 to understand when the first-order optimum satisfies IC_L, which can be rewritten as follows:

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

According to Corollary 2, the number of periods since the last high shock is a sufficient statistics for the agent's utilities. As a result, the upward incentive constraint can be succinctly rewritten as

$$\overbrace{\max\{\hat{U}_t^{\#}, U_t^{\#}\}}^{\substack{\text{maximal payoff for } \theta_H \\ \text{after } t \text{ consecutive low shocks}}} \qquad \leqslant \qquad \overbrace{\underbrace{\Delta\theta R(k_H^e)}_{\text{static}} + \underbrace{\delta_A(\alpha_H - \alpha_L)U_1^{\#}}_{\text{dynamic}}}^{\theta_L' s \text{ payoff from misreporting.}}.$$

Since the first-order optimal distortions are monotonically decreasing in the number of low shocks (see Corollary 1) the tightest upward incentive constraint is one at "infinity". The following corollary makes this statement precise.

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

$$\lim_{t\to\infty} \hat{U}_t^{\#} = \lim_{t\to\infty} U_t^{\#} \leqslant \Delta\theta R(k_H^e) + \delta_A(\alpha_H - \alpha_L)U_1^{\#}.$$

The intuition for this result can be expressed as follows. Take two histories, the lowest one $\theta_{\mathbf{L}}^{\mathbf{t+s}}$ and one where one high shock has been realized $(\theta_{\mathbf{L}}^{\mathbf{t}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s-1}})$. In period $t+1$, for the first history, the allocation is $\hat{k}_{t+1}^{\#}$ and for the second history it is $k_1^{\#}$. If the parameters are such that the seed distortion ρ_1 is high, then for a large enough t, since $\hat{k}_t^{\#}$ is increasing in t, we have a situation where $k_{t+s}^{\#}(\theta_{\mathbf{L}}^{\mathbf{t+s}}) \gg k_{t+s}^{\#}(\theta_{\mathbf{L}}^{\mathbf{t}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s-1}})$. In other words, the allocation along the history $\theta_{\mathbf{L}}^{\mathbf{t+s}}$ is much larger than along $(\theta_{\mathbf{L}}^{\mathbf{t}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s-1}})$ even though the latter is "better" in terms of the sequence of productivity shocks. When this force is strong, the inequality in Corollary 3 is violated.

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 $\{\hat{U}_t^{\#}\}$ and $\{U_t^{\#}\}$ 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* $0 < \delta_A \leqslant \delta_P < 1$*. Then for any Markov process* $1 > \alpha_H \geqslant \alpha_L > 0$ *that* *satisfies*

$$(1-\alpha_L)(\alpha_H-\alpha_L)\left(\frac{\alpha_H}{1-\alpha_H}-\frac{\alpha_L}{1-\alpha_L}\right)\geqslant\frac{1}{\delta_A\left(1-\frac{\delta_A}{\delta_P}\right)},$$

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

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

$$\alpha(2\alpha-1)\left(\frac{\alpha}{1-\alpha}-\frac{1-\alpha}{\alpha}\right)\geqslant\frac{1}{\delta_A\left(1-\frac{\delta_A}{\delta_P}\right)}. \tag{3}$$

Inequality (3) can be used to derive some intuition about the (in)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 (3) 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, (3) 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 (3) explodes, so the sufficient condition is satisfied and the relaxed problem approach is violated. Finally, as $\alpha \to \frac{1}{2}$, the Markov process becomes iid and the right-hand side of (3) converges to zero, and we know for the iid model, the relaxed problem approach is valid.

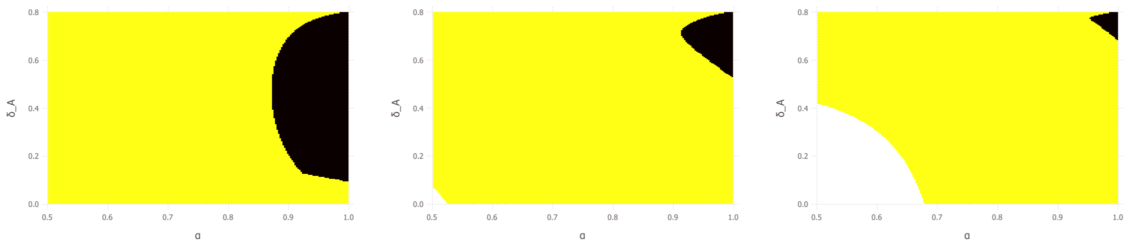

(a) $\Delta\theta = 0.1$ (b) $\Delta\theta = 1.1$ (c) $\Delta\theta = 2.1$

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. Here, $\alpha_H = 1-\alpha_L = \alpha$ 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 $\Delta\theta$. A larger value of $\Delta\theta$ signifies greater ex ante asymmetric information, hence a high distortion for the low type in the first place. So, the non-monotonicity desired above for Corollary 3, viz., $\hat{k}^{\#}_{t+s} \gg k^{\#}_s$ for large enough t, no longer holds, or at least not strongly enough for IC_L to be violated. In fact, for large enough value of $\Delta\theta$ as in Figure 2c, the low-type can be shutdown, which is the highest possible distortion.

5 Restart contract

What can the principal do if she faces parameters for which the relaxed problem approach is not valid? For starters, she can brute force her way to determine the optimal contract despite the large number of binding constraints. We provide this solution in the recursive format—the problem and its solution are detailed in the appendix. However, an alternative, and to us a more appealing, solution is to look for the optimum within the restrictive class of restart contracts that are incentive compatible, simple and approximately optimal.[19]

5.1 The optimal restart contract

A restart contract, described informally above, can be defined formally as follows.

Definition 2. *A contract* $\langle \mathbf{k}, \mathbf{U} \rangle$ *is called* **restart** *if there exists a number* k_H *and a sequence* $\{k_t\}$ *such that for all* $\theta^{\mathbf{t-1}}$*, we have*

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

This definition is depicted in Figure 1. It essentially requires 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.[20] It is immediately clear from Theorem 1 that the first-order optimal contract is indeed a restart contract.

In this section, we search for the optimum in a restrictive class of contracts that are required to be restart and satisfy the full set of constraints; moreover IC_H must hold as an equality:

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

[19] As mentioned in the introduction, we find restart contracts appealing for four reasons: First, is an intrinsic normative appeal to the idea of restartness in the form of "let bygones be bygones". Second, these contracts generated through unequal discounting connects to a sizeable literature in economics, particularly public finance, political economy and sovereign debt. Third, restart contracts arise naturally as the solution to the relaxed problem, which is quite the standard in contract theory. Fourth, these contracts are simple in the sense that will be described in Section 6.

[20] The second sequence is left out in Definition 2 for simplicity, because it is implicit that since the lowest history is the only one which cannot be written in the form $(\theta^{\mathbf{t-1}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s}})$, it will have its own sequence of allocations.

We will denote the solution of this problem by $\langle \mathbf{k}^{\mathbf{R}}, \mathbf{U}^{\mathbf{R}} \rangle$, and refer to it as the *restart optimum*.[21], [22]

When the optimal contract is restart, there is no loss from this extra restriction. It is easy to see that Problems ($\star$), ($\#$) and (R) all have the same objective, but they are nested by the set of constraints imposed on that objective in the following way: $(\#) \subset (\star) \subset (R)$. Hence, $\Pi^R \leqslant \Pi^\star \leqslant \Pi^\#$, and $\Pi^R = \Pi^\star = \Pi^\#$, whenever the relaxed problem approach is valid. In what follows we describe the restart optimum, $\langle \mathbf{k}^{\mathbf{R}}, \mathbf{U}^{\mathbf{R}} \rangle$, and then provide a theoretical bound to precisely capture the gap in profit generated by $\langle \mathbf{k}^\star, \mathbf{U}^\star \rangle$ and $\langle \mathbf{k}^{\mathbf{R}}, \mathbf{U}^{\mathbf{R}} \rangle$.

Theorem 2. *There exists $\bar{\gamma} \geqslant 0$ such that the restart optimum is as follows:*

a) $k_H^R \geqslant k_H^\# = k_H^e$*;*

b) $\hat{k}_t^R = \mathcal{K}_L(\hat{\gamma}_t)$ *for* $\hat{\gamma}_t = \max\{\bar{\gamma}, b\hat{\gamma}_{t-1} + a_L\}$ *for some* $\hat{\gamma}_1 \geqslant \hat{\rho}_1 = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)}$*;*

c) $k_t^R = \mathcal{K}_L(\gamma_t)$ *for* $\gamma_t = \max\{\bar{\gamma}, b\gamma_{t-1} + a_L\}$ *for some* $\gamma_1 \leqslant \rho_1 = a_H$*,*

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* $j = H, L$*.*

Theorem 2 describes the optimal distortions for the low type along the two classes of histories, first where no high type is realized and second where at least one high type has been realized; these are given by $\{\hat{\gamma}_t\}$ and $\{\gamma_t\}$, respectively. The contract here is analogous to the first-order optimal contract (Theorem 1), with three key differences: First, the high type allocation is (potentially) distorted upwards. Second, 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 $\gamma_1 \leqslant \rho_1$. Third, there is a floor $\bar{\gamma}$ on distortions, i.e., $\hat{\gamma}_t, \gamma_t \geqslant \bar{\gamma}$ for all t. Since distortions are still decreasing this means if the floor binds for $\hat{\gamma}_t$ or γ_t, it binds for $\hat{\gamma}_{t+s}$ or γ_{t+s} as well, and thence the contract has finite memory.[23]

Further, note that the *propagator* b and the *adder* a_L are the same as before. The "initial" allocation is determined by three numbers k_H^R, γ_1 and $\hat{\gamma}_1$. These 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).

5.2 Approximate global optimality

How well does the optimal restart contract perform? By construction, $\Pi^R \leqslant \Pi^\star$. Unfortunately, the gap between the two is very hard to theoretically compute when the upward

[21] In general, the optimal restart contract does not have to satisfy all the downward constraints as equality. We require 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.

[22] 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.

[23] 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.

constraints bind, because it requires us to solve a system of non-linear equations defining the values of $\{\bar{\gamma}, k_H^R, \hat{\gamma}_1, \gamma_1\}$ as described in Theorem 2. However, we can still bound the loss by using the expression for the first-order optimal contract, $\Pi^{\#}$, which is calculable in a closed form. Since $\Pi^{\star} \leqslant \Pi^{\#}$, we must have $\Pi^{\star} - \Pi^R \leqslant \Pi^{\#} - \Pi^R$. So, our goal here is to bound $\Pi^{\#} - \Pi^R$, which then provides an upper bound for $\Pi^{\star} - \Pi^R$ as well. We estimate the former gap using sensitivity analysis. We first describe the general mathematical argument and then show how it can be applied to our context.

Remark 1. *Consider the problem of maximizing a smooth concave function $f : X \to \mathbb{R}$ subject to a set of linear inequality constraints: $Ax \geqslant 0$, where X is a closed subset of $\mathbb{R}^n_+$ and A is an $m \times n$ matrix of the coefficients. The set X here incorporates all additional constraints, i.e., linear equality constraints, that can be required.*

Suppose that the function f admits a maximizer on X, say x^*. The goal is to assess the gap between $f(x^*)$ and $\max_{x \in X} f(x)$ subject to $Ax \geqslant 0$. In order to estimate the gap, it is useful to introduce the following auxiliary problem parametrized by $\varepsilon \geqslant 0$:

$$\Pi(\varepsilon) := \max_{x \geqslant 0} f(x) \quad \text{subject to} \quad Ax \geqslant \varepsilon \min\{0, Ax^*\}.$$

The auxiliary problem admits a solution for every $\varepsilon \geqslant 0$, because the unconstrained problem does. Note that $\Pi(1) = f(x^*)$, because x^* is feasible whenever $\varepsilon = 1$, and $\Pi(0)$ corresponds to the optimal value in the original problem.

Since the problem is concave and bounded, strong duality holds, and thus we can express the value of the auxiliary problem as

$$\Pi(\varepsilon) = \min_{\lambda \geqslant 0} \max_{x \geqslant 0} f(x) + \lambda \cdot (Ax - \varepsilon \min\{0, Ax^*\}).$$

It follows that $\Pi(1) \leqslant \max_{x \geqslant 0} f(x) + \lambda(0) \cdot (Ax - \min\{0, Ax^*\})$, where $\lambda(0)$ is the dual variable associated with the constraint for $\varepsilon = 0$. Combining the inequality with the definition of $\Pi(0)$, we finally obtain the following estimate:

$$\Pi(1) - \Pi(0) \leqslant \lambda(0) \cdot \max\{0, -Ax^*\}.$$

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. Similarly to Corollary 2, the low type's utilities are zero at all dates and the high type's utilities feature restarts:

a) $U_t(\theta^{\mathbf{t-1}}, \theta_L) = 0$ for all $\theta^{\mathbf{t-1}}$;

b) $U_t(\theta_{\mathbf{L}}^{\mathbf{t-1}}, \theta_H) = \hat{U}_t$ for $\hat{U}_t := \Delta\theta \sum_{s=t}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-t} R(\hat{k}_s)$;

c) $U_{t+s}(\theta^{\mathbf{t-1}}, \theta_H, \theta_{\mathbf{L}}^{\mathbf{s-1}}, \theta_H) = U_s$ for all $\theta^{\mathbf{t-1}}$ for $U_t := \Delta\theta \sum_{s=t}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-t} R(k_s)$, Moreover, we require the upward incentive constraints (IC_L) to hold, these form the set of linear inequality constraints:

$$\hat{U}_t \leqslant \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)U_1, \quad U_t \leqslant \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)U_1. \tag{4}$$

The first-order optimum solves the problem when (4) is ignored yielding the minimal slack, then our estimate of loss combines this slack and Lagrange multipliers when (4) is imposed.

Attach a Lagrange multiplier to each upward incentive constraint, say $\hat{\eta}_t$ and η_t, and evaluate the multipliers at the restart optimum.[24] Quantify how much slack needs to be added to these constraints so that the solution then coincides with the first-order optimum.[25] Remark 1 implies the following estimate:

$$\Pi^{\#} - \Pi^{R} \leqslant \sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}})\hat{\eta}_t \cdot \hat{\epsilon}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \sum_{t=1}^{\infty} \delta_P^t \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}|\theta_{\mathbf{H}})\eta_t \cdot \epsilon_t =: B,$$

where $\hat{\epsilon}_t$ and ϵ_t are slack variables measuring the extent of violation of the upward incentive constraint by the first-order optimum:

$$\hat{\epsilon}_t := \max\left\{0, \hat{U}_t^{\#} - \Delta\theta R(k_H^e) - \delta_A(\alpha_H - \alpha_L)U_1^{\#}\right\}, \quad \epsilon_t := \max\left\{0, U_t^{\#} - \Delta\theta R(k_H^e) - \delta_A(\alpha_H - \alpha_L)U_1^{\#}\right\}.$$

The general estimate B makes it clear that there is no loss from the restriction to restart contracts whenever the upward incentive constraints are satisfied for the first-order optimum, i.e., $\hat{\epsilon}_t = \epsilon_t = 0$ at all dates. And, even if some of these constraints are actually violated, the loss from using the restart contracts is at most linear in $\hat{\epsilon}_t$ and ϵ_t. In addition, the loss is expected to be relatively small whenever the violation of upward incentive constraint is not severe. Indeed, in the appendix we formally show that the Lagrange multipliers $\{\hat{\eta}_t\}$ and $\{\eta_t\}$ can themselves be bounded using the first-order conditions of Problem (R). This gives an easily computable bound on the gap between the global optimum and restart optimum.. Here we sketch the construction.

We now provide two different ways to bound the gap B as a function of primitives. First, note that Corollaries 1 and 2 jointly imply that both sequences of utilities $\{\hat{U}_t^{\#}\}$ and $\{U_t^{\#}\}$ are increasing to the same limit, i.e., $\lim_{t\to\infty} \hat{U}_t^{\#} = \lim_{t\to\infty} U_t^{\#}$. It follows that the slack sequences $\{\hat{\epsilon}_t\}$ and $\{\epsilon_t\}$ are increasing as well, again with the same limit, thus the gap can be bounded as follows:

$$\Pi^{\#} - \Pi^{R} \leqslant B \leqslant \left(\sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}})\hat{\eta}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \sum_{t=1}^{\infty} \delta_P^t \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}|\theta_{\mathbf{H}})\eta_t \right) \cdot \lim_{t\to\infty} \hat{\epsilon}_t.$$

[24] Formally, the multipliers are defined as $\delta_P^{t-1}\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t-1}})\hat{\eta}_t$ and $\frac{\mathbf{P}(\theta_H)}{1-\delta_P}\delta_P^t\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}|\theta_{\mathbf{H}})\eta_t$, which is merely a normalization.

[25] 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.

We note that the term is in the brackets is the "aggregate" shadow price of the upward incentive constraint. Since the right-hand of each upward incentive constraint is the same, i.e., $\Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)U_1$, the aggregate shadow price is determined in a way that the coefficient in the Lagrangian in front of U_1 is zero. In other words, the aggregate shadow price matches exactly the marginal benefit of adjusting U_1, which turns out to be proportional to the difference in the seeds at the restart and first-order optimums: $\rho_1 - \gamma_1 \geqslant 0$. In the appendix we show that

$$\sum_{t=1}^{\infty} \delta_P^{t-1}\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}})\hat{\eta}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P}\sum_{t=1}^{\infty} \delta_P^t\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}|\theta_{\mathbf{H}})\eta_t = \frac{P(\theta_H)}{1-\delta_P}\frac{\delta_P}{\delta_A}\frac{1-\alpha_H}{1-\alpha_L}(\rho_1 - \gamma_1) \leqslant$$
$$\leqslant \frac{P(\theta_H)}{1-\delta_P}\frac{\delta_P}{\delta_A}\frac{1-\alpha_H}{\alpha_H-\alpha_L}\left(\rho_1 - \lim_{t\to\infty}\rho_t\right),$$

where the last inequality follows from the fact $\{\gamma_t\}$ obeys the same dynamics as $\{\rho_t\}$, i.e., it is monotonically decreasing to $\lim_{t\to\infty}\rho_t$, but is bounded from below by the floor (see Theorem 2). So, we conclude that the general bound B satisfies the following inequality:

$$\Pi^\# - \Pi^R \leqslant B \leqslant \frac{P(\theta_H)}{1-\delta_P}\frac{\delta_P}{\delta_A}\frac{1-\alpha_H}{\alpha_H-\alpha_L}\left(\rho_1 - \lim_{t\to\infty}\rho_t\right) =: B_a^1. \tag{5}$$

We now sketch the second approach to bound B in terms of fundamentals. Using the fact the optimal restart distortions $\{\hat{\gamma}_t\}$ and $\{\gamma_t\}$ are bounded from below by the floor distortion $\bar{\gamma}$ (see Theorem 2), we show in the appendix that $\hat{\eta}_{t+1}, \eta_{t+1} \leqslant (1-b)\,(\rho_1 - \lim_{t\to\infty}\rho_t)$ for all $t \geqslant 2$, and $\hat{\eta}_1 = (\hat{\rho}_1 - \lim_{t\to\infty}\rho_t)$ and $\eta_1 = 0$. Substituting these into the expression for B, we arrive at the second bound:

$$\Pi^\# - \Pi^R \leqslant B \leqslant \mathbb{P}(\theta_L)\,(\rho - \hat{\rho}_1)^+\,\hat{\epsilon}_1 +$$
$$+ (1-b)\left(\rho_1 - \lim_{t\to\infty}\rho_t\right)\sum_{t=2}^{\infty}(\delta_P(1-\alpha_L))^{t-1}\left(\mathbb{P}(\theta_L)\hat{\epsilon}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P}\delta_P(1-\alpha_H)\epsilon_t\right) =: B_a^2. \tag{6}$$

Combining both estimates we obtain that $B \leqslant \min\{B_a^1, B_a^2\}$, and the value of $\min\{B_a^1, B_a^2\}$ is a known function of primitives. In the appendix we actually make the construction even a bit tighter combining $\min\{B_a^1, B_a^2\}$ with the loss from using the best static contract. This adjustment ensures that the relative loss stays bounded as well for all constellations of parameters:

$$B_a := \min\{B_a^1, B_a^2, \Pi^\# - \Pi^S\}, \text{ where } \Pi^S \text{ is the best static contract.} \tag{7}$$

Corollary 5. *There exists two bounds, B_a and B_r, functions of primitives, such that $\Pi^\star - \Pi^R \leqslant B_a$ and $1 - \frac{\Pi^R}{\Pi^\star} \leqslant B_r$, where B_a is defined by equations (5), (6) and (7); $B_r = \frac{B_a}{\Pi^\#}$; and $B_a = B_r = 0$ whenever the optimal contract is restart.*

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 problem approach. Moreover, it is easy to check that as $\Delta\theta \to 0$, as required by our sufficient condition for the invalidity of the first-order approach (Corollary 4), both slack variables $\{\hat{\epsilon}_t\}$ and $\{\epsilon_t\}$ converge to zero. Thus, the additive bound converges to zero as well, and the restart optimum has no loss. 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 % and the actual loss is never more than 4 %.[26]

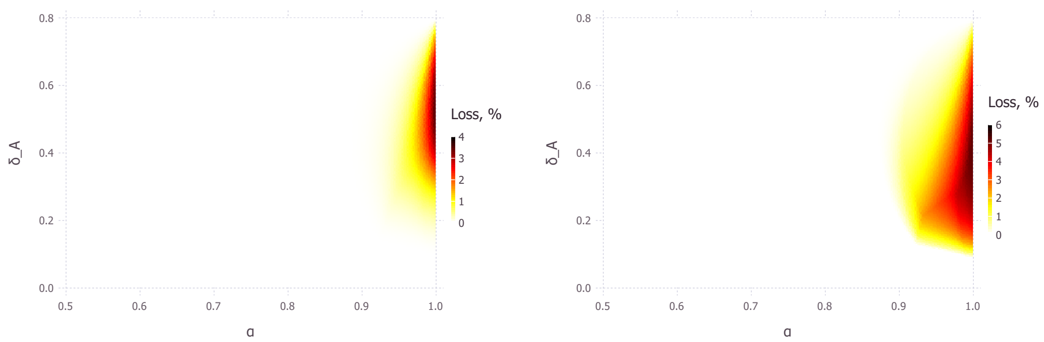

(a) actual loss (b) analytical loss

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 $\Delta\theta = 0.1$.

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 fixes an 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 relaxed approach is valid.

6 Simplicity through recursivity

We now define a notion of simplicity under which any restart contract is simple. In addition, when the optimal contract here is not restart, i.e., the relaxed problem approach is not valid,

[26]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, B_r, for which no optimization is required, it is simply a function of the fundamentals of the model.

then the state space required to encode it is quite rich, which makes the contract not simple. The argument follows the recursive approach to contract design which is described in detail in the appendix. An intuitive sketch follows.

Typically "continuation utility" is the chosen state variable in dynamic games and contracts. With asymmetric information, the number of state variable has to be enriched to be the cardinality of the agent's type space– one continuation utility for each type (see Fernandes and Phelan [2000]). Here the problem can be simplified since IR_L always binds, i.e., $U(\theta^{t-1}, \theta_L) = 0$. The recursive contract thus takes the expected payoff to the high type as a state variable, and then optimizes over the allocation and expected continuation utility to be given to the high type in the next period. That is, taking $w = U(\theta^{\mathbf{t-1}}, \theta_H)$, the optimal recursive contract pins down $(\mathbf{k}, \mathbf{z})$, where $\mathbf{k} = (k_H(w), k_L(w))$ is the allocation and $\mathbf{z} = (z_H(w), z_L(w))$ is the expected continuation utility promised to the high type after either high or low realizations. These are chosen to maximize the principal's profit subject to the recursive versions of incentive compatibility and individual rationality.

It is also well known that a recursive strategy (or contract) can be thought as an automaton. For our model, the automaton starts from some initial state, and then upon the announcement of θ_H or θ_L, it supplies capital and determines the expected continuation utility to the high type agent, which in turn becomes the next state of the automaton, and so on. 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 number of distinct allocations supplied by the automaton "machine".[27]

Unfortunately, in our infinite horizon contracting 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 the set $\left\{k | \exists \theta^{\mathbf{t}} : k = k_t(\theta^{\mathbf{t}}), t \leqslant T\right\}$ to grow exponentially with T. We use an intermediate notion 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*,*

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

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.

[27] This notion was first studied by Moore [1956], and it is often referred to as the Moore-machine. See also Chatterjee and Sabourian [2009] for a survey on the study of simplicity/complexity in strategic settings.

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

The proof of Theorem 3 involves first restating the contract design problem using the recursive approach and then using its properties to determine the cardinality of the optimal mechanism. The optimal recursive contract is given by $\langle U_1^\star(\theta_H), \mathbf{k}(\cdot), \mathbf{z}(\cdot)\rangle$, where $U_1^\star(\theta_H)$ is the initial state, and $\mathbf{k}$ and $\mathbf{z}$ are mappings that take the current state as given and spit out the allocation for this period and the value of the future state. In fact, the optimal sequential contract $\langle \mathbf{k}^\star, \mathbf{U}^\star\rangle$ can be constructed from the optimal recursive contract $\langle U_1^\star(\theta_H), \mathbf{k}(\cdot), \mathbf{z}(\cdot)\rangle$ in the standard inductive fashion. 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, which in turn is the expected continuation utility promised to the high type.

Conceptually, the idea behind Theorem 3 is as follows. When the optimal contract is restart, the cardinality of its support 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 the state space needed to encode it has to keep track of two (new) expected utilities in every period because now $U^\star$ is completely history dependent. Thus, the contract space grows at a rate that is at least as large as 2^T, which render it not simple or complex. This provides a natural normative criterion for why optimal restart contracts might be easier to use than the optimal contracts when the relaxed problem approach is violated.

7 Comparative Statics

We provide two types of comparative statics results: 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. Both results are formally stated in the appendix, here we provided an intuitive exposition.

A standard folk theorem type of result has been reported in the dynamic mechanism design literature– it states that sum of average payoffs across all parties converges to the average surplus maximizing payoff as the discount factor converges to one (see Battaglini [2005] and Athey and Segal [2013]). In Corollary 6, we show that as the principal becomes arbitrarily patient the efficient surplus can be achieved if and only if the agent too becomes arbitrarily patient.[28] The familiar intuition carries through– for an arbitrarily patient principal, even though the standard backloading force becomes almost free of cost, the intertemporal costs

[28]This result is also related 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.

of incentive provision are forever positive as long as the agent's discount factor is bounded away from one.

The second comparative static is with respect to the principal's profit as a function of the agent's discount factor. Assume δ_P to be fixed and less than 1. Does the principal favor an impatient agent or patient agent, and what determines the ranking if there exists any? Due to the complexity of competing forces, there is no easily expressible global comparative static here. A theoretical result can be stated for the limit cases (iid and perfect persistence, see Corollary 7), and numerical arguments explored for the intermediate ones.

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 $\alpha_H = 1-\alpha_L = \alpha$ (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 the limit cases– the principal prefers the patient agent ($\delta_A = \delta_P$) for α sufficiently close to $\frac{1}{2}$, and she prefers the myopic agent ($\delta_A = 0$) for α sufficiently close to 1. 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 those with intermediate values of δ_A.

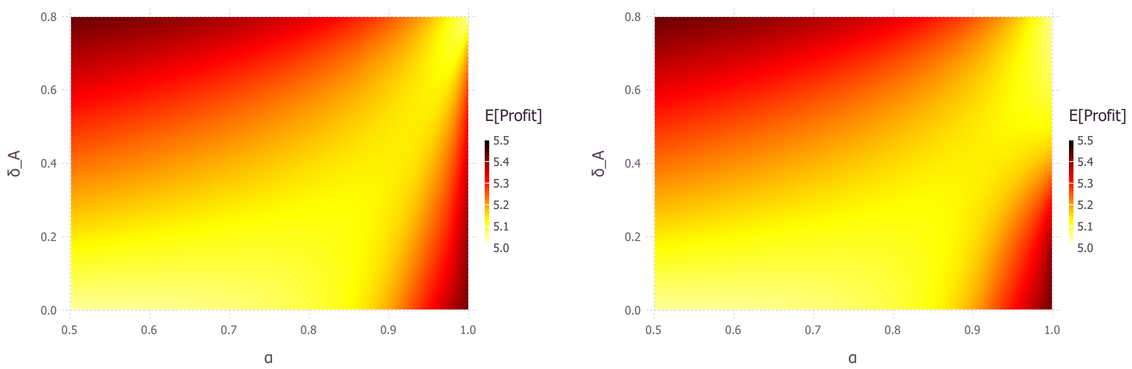

(a) First-order optimal contract (b) Optimal restart contract

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 $\Delta\theta = 0.1$.

The limit cases can be easily understood: When persistence is very high, the principal has to pay a large information rent pointwise, for every history. Thus, in order to bring down the net present value of this cost she prefers a myopic agent, even though this increases the inter-temporal cost. On the other hand, with very low persistence, the pointwise value of information rent is small. So, the principal prefers a forward looking agent because it erases the inter temporal cost and whatever increase in the net present value of information rent accrues, it can be extracted upfront since backloading is not constrained anymore by unequal discounting.

8 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.

Their interaction produces a tradeoff for the principal: backloading agent's payoffs as much as possible to relax 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 both simplicity and approximate optimality are formally defined.

The nature of dynamic distortions poses a question to the literature on dynamic (Myersonian) mechanism design— moving away from the standard equal discounting model changes the structure of distortions in that they become pervasive and cyclical. With equal discounting, Besanko [1985] and Battaglini [2005] show that ex post distortions steadily decrease to zero in the long run for the AR(1) and two type Markov models respectively. Garrett, Pavan, and Toikka [2018] show that distortions steadily decrease 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 upfront. 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 [2021]). 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 of 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".[29] 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.

9 Appendix

9.1 Sequential characterization

9.1.1 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 and Lemma 2 proves that IC_H binds in the relaxed problem. We use the terminology that a constraint is respected when it holds as a weak inequality.

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 the same as in the original mechanism. To see that IR_H holds, inductively expand

[29] 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].

IC_H along the persistent history of θ_L's:

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

where the last inequality follows from positivity of output.

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 a result, the ex-ante profit of $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ is strictly 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*

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

Then, there exists another mechanism $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ that satisfies IC_H, IR_L and yields a strictly 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 both IR_L and IC_H bind in the new mechanism.

As in Lemma 1, $\tilde{\mathbf{U}} \leqslant \mathbf{U}$ and $\tilde{\mathbf{U}} \neq \mathbf{U}$, thus the altered mechanism is cheaper for the principal. It follows that the ex-ante profit of $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ is strictly higher than of $\langle \mathbf{k}, \mathbf{U} \rangle$. □

9.1.2 Relaxed problem approach

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

Proof of Theorem 1. First, we will derive Equation 2 and two sequences of distortions $\{\hat{\rho}_t\}$ and $\{\rho_t\}$, which are described in the statement of the theorem. Then, we will show how to construct the optimal allocations.

Lemmas 1 and 2 jointly imply that 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 for all $\theta^{\mathbf{t-1}}$ we have that $U_t(\theta^{\mathbf{t-1}}, \theta_L) = 0$ and

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

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:

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

where $b = \frac{\delta_A}{\delta_P} \frac{\alpha_H - \alpha_L}{1 - \alpha_L}$ is the propagator as described in Section 4.1.

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

$$I = (\delta_P - \delta_A)\sum_{t=2}^{\infty}\delta_P^{t-2}\mathbb{E}\left[U_t(\theta^{\mathbf{t}})\right] = (\delta_P - \delta_A)\sum_{\theta^{\mathbf{t-1}}:t\geqslant 2}\delta_P^{t-2}\mathbb{P}(\theta^{\mathbf{t-1}},\theta_H)U_t(\theta^{\mathbf{t-1}},\theta_H) =$$

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

To make further progress, we expand the intertemporal cost of incentive provision separately along the lowest history of θ_L's and in the restart phase where $\theta^{\mathbf{t-1}}$ contains at least one θ_H.

In the former case, we have

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

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

where $a_L = \frac{\delta_P-\delta_A}{\delta_P}\frac{\alpha_L}{1-\alpha_L}$ is the adder, which was introduced in Section 4.1. The distortions along the lowest history are then given by

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

In the latter case, we have

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

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

where $a_H = \frac{\delta_P-\delta_A}{\delta_P}\frac{\alpha_H}{1-\alpha_H}$ is the see as defined in Section 4.1. The total distortion in the restart phase can concisely be written as

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

It remains only to find the optimal allocations. Substituting the expressions for the agent's ex ante utility and intertemporal cost of incentive provision into the seller's profit, we obtain the following representation:

$$\bar{S} - \bar{U}_A - I = \sum_{\theta^{t-1}} \delta_P^{t-1} s(\theta_H, k_t(\theta^{t-1}, \theta_H))\mathbb{P}(\theta^{t-1}, \theta_H) + \sum_{t=1}^{\infty} \delta_P^{t-1} s\left(\theta_L - \Delta\theta\hat{\rho}_t, k_t(\theta_L^{t-1}, \theta_L)\right)\mathbb{P}(\theta_L^t) +$$
$$+ \sum_{\theta^{t-1}} \sum_{s=1}^{\infty} \delta_P^{t-1+s} s\left(\theta_L - \Delta\theta\rho_t, k_{t+s}(\theta^{t-1}, \theta_H, \theta_L^s)\right)\mathbb{P}(\theta^{t-1}, \theta_H, \theta_L^s).$$

The reader can verify that pointwise optimization of the above objective yields the first-order optimal allocation rule $\mathbf{k}^{\#}$ as described in Theorem 1. □

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 $\hat{\rho}_{t+1} = f(\hat{\rho}_t)$.

It is routine to 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)* 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. □

9.1.3 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 utilities as a function of primitives, whereas the latter identifies the tightest possible upward incentive constraint in the whole set of IC_L.

Proofs of Corollaries 2 and 3. Corollary 2 follows simply from Equation (1).

As for Corollary 3, we showed in Section 4.3 that the upward incentive constraint for the first-order optimal contract can be expressed as

$$\max\{\hat{U}_t^{\#}, U_t^{\#}\} \leqslant \Delta\theta R(k_H^e) + \delta_A(\alpha_H - \alpha_L)U_1^{\#},$$

where the sequences of utilities $\{\hat{U}_t\}$ and $\{U_t\}$ are defined in Corollary 2. By Corollary 1, both sequences are increasing to the same value, i.e.,

$$\lim_{t\to\infty} \hat{U}_t^{\#} = \lim_{t\to\infty} U_t^{\#} = \frac{\Delta\theta(R \circ \mathcal{K}_L)\left(\lim_{t\to\infty} \rho_t\right)}{1 - \delta_A(\alpha_H - \alpha_L)}.$$

It follows that the "tightest" incentive constraint is one at the "infinity". □

Next, we provide a proof of Corollary 4, which gives a condition for invalidity of the first-order approach.

Proof of Corollary 4. Define a number ϵ to be the slack in the tightest possible upward incentive constraint, that is

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

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

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

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

We now show that ϵ 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, that is

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

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 ϵ is differentiable in $\Delta\theta$ at zero. Taking the common factor outside the brackets, we can express the derivative of ϵ with respect to $\Delta\theta$ as

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

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

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

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

Substituting these into the second term of the above expression, we arrive at

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

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 $b < \delta_A/\delta_P$, thus $\frac{a_L}{1-b} < \frac{\alpha_L}{1-\alpha_L}$. It follows that

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

The assumption of Corollary 4 implies that

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

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

9.1.4 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 binds at every history, i.e., $U_t(\theta^{\mathbf{t-1}}, \theta_L) = 0$. We further restrict a contract space to be the 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 $\theta^{\mathbf{t-1}}$, 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 functions of the allocation rule by the binding downward constraint:

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

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 incentive constraint along the lowest history, the latter corresponds to the restart phase.

It is convenient to rewrite the objective in terms of the aforementioned 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:

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

The term $\frac{\mathbb{P}(\theta_H)}{1-\delta_P}$ is the 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 $\bar{U}_A = \hat{U}_1 \mathbb{P}(\theta_H)$, whereas the intertemporal costs of incentive provision can be factored as

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

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

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

$$\max_{k_H, \{\hat{k}_t\}, \{k_t\}, \{\hat{U}_t\}, \{U_t\}} \quad \bar{S} - \bar{U}_A - I \quad \text{subject to} \quad k_H \geqslant 0, \ \forall t \quad \hat{k}_t, k_t, \hat{U}_t, U_t \geqslant 0, \ \text{and}$$

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

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 $\delta_P^{t-1} \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}) \hat{\gamma}_t$, whereas the upward incentive constraints are associated with dual variables $\delta_P^{t-1} \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}) \hat{\eta}_t$. Similarly, in the restart phase multipliers are $\frac{\mathbb{P}(\theta_H)}{1-\delta_P}\delta_P^t\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}|\theta_{\mathbf{H}})\gamma_t$ and $\frac{\mathbf{P}(\theta_H)}{1-\delta_P}\delta_P^t\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}|\theta_{\mathbf{H}})\eta_t$ for the downward and upward incentive constraints, respectively. Then, the Lagrangian is as follows:

$$\frac{\mathbb{P}(\theta_H)}{1-\delta_P}s\left(\theta_H+\Delta\theta\frac{(1-\delta_P)\kappa}{\mathbb{P}(\theta_H)},k_H\right)+\sum_{t=1}^{\infty}\delta_P^{t-1}s(\theta_L-\Delta\theta\hat{\gamma}_t,\hat{k}_t)\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}})+\sum_{t=1}^{\infty}\delta_P^t s(\theta_L-\Delta\theta\gamma_t,k_t)\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}|\theta_{\mathbf{H}})+$$
$$+\Big(-\mathbb{P}(\theta_H)+\mathbb{P}(\theta_L)(\hat{\gamma}_1-\hat{\eta}_1)\Big)\hat{U}_1+\frac{\mathbb{P}(\theta_H)}{1-\delta_P}\left(-(\delta_P-\delta_A)\alpha_H+\delta_P(1-\alpha_H)(\gamma_1-\eta_1)+\delta_A(\alpha_H-\alpha_L)\frac{(1-\delta_P)\kappa}{\mathbb{P}(\theta_H)}\right)U_1+$$
$$+\sum_{t=1}^{\infty}\delta_P^{t-1}\Big(-(\delta_P-\delta_A)\alpha_L+\delta_P(1-\alpha_L)(\hat{\gamma}_{t+1}-\hat{\rho}_{t+1})-\delta_A(\alpha_H-\alpha_L)\hat{\gamma}_t\Big)\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}})\hat{U}_{t+1}+$$
$$+\frac{\mathbb{P}(\theta_H)}{1-\delta_P}\sum_{t=1}^{\infty}\delta_P^{t-1}\Big(-(\delta_P-\delta_A)\alpha_L+\delta_P(1-\alpha_L)(\gamma_{t+1}-\rho_{t+1})-\delta_A(\alpha_H-\alpha_L)\gamma_t\Big)\mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}|\theta_H)U_{t+1},$$

where the number κ measures of the total shadow price of IC_L, and it is defined by

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

The reader can verify that the allocations are uniquely determined by the set of first-order conditions as a function of the Lagrange multipliers:

$$\hat{k}_t=\mathcal{K}_L(\hat{\gamma}_t),\quad k_t=\mathcal{K}_L(\gamma_t)\quad\text{and}\quad k_H=\mathcal{K}_H\left(\frac{(1-\delta_P)\kappa}{\mathbb{P}(\theta_H)}\right)\geqslant k_H^e.$$

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. We shall use our standard shorthand notations: $b=\frac{\delta_A}{\delta_P}\frac{\alpha_H-\alpha_L}{1-\alpha_L}$, $a_j=\left(1-\frac{\delta_A}{\delta_P}\right)\frac{\alpha_j}{1-\alpha_j}$ for $j=H,L$. We will also use two sequences of first-order optimal distortions $\{\hat{\rho}_t\}$ and $\{\rho_t\}$, which are defined in Theorem 1 as $\hat{\rho}_t=b\hat{\rho}_{t-1}+a_L$, $\hat{\rho}_1=\frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)}$ and $\rho_t=b\rho_{t-1}+a_L$, $\rho_1=a_H$.

To begin, fix $\bar{\gamma}\geqslant 0$ and $\gamma_1\in\left[\lim_{t\to\infty}\rho_t,\rho_1\right]$, where $\lim_{t\to\infty}\rho_t=\frac{a_L}{1-b}$ and $\rho_1=a_H$, and define $\{\hat{\gamma}_t\}$, $\{\gamma_{t+1}\}$ by

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

Then, let $\eta_1:=0$, $\hat{\eta}_1:=(\bar{\gamma}_1-\hat{\rho}_1)^+$ 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 $(\bar{\gamma},\gamma_1)$. 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

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

Note that $\kappa=0$ whenever $\bar{\gamma}$ is sufficiently small, moreover, it is strictly increasing in $\bar{\gamma}$ without bound. Therefore, for any $\gamma_1 \in \left[\lim_{t\to\infty} \rho_t, \rho_1\right)$, there exists a unique value of $\bar{\gamma}$ which makes the aforementioned coefficient equal to zero. For $\gamma_1 = \rho_1$, any $\bar{\gamma} \leqslant \min\left\{\lim_{t\to\infty} \rho_t, \hat{\rho}_1\right\} = \lim_{t\to\infty} \rho_t = \frac{a_L}{1-b}$ 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, $\gamma_1 = \rho_1$ will work. Since the distortions $\{\hat{\gamma}_t\}$, $\{\gamma_t\}$ 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 the "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 $\gamma_1 = \rho_1$, 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 $\gamma_1 = \lim_{t\to\infty} \rho_t$. To see it more formally, let $\gamma_1 = \lim_{t\to\infty} \rho_t$. Then, we have $\bar{\gamma} > \min\left\{\lim_{t\to\infty} \rho_t, \hat{\rho}_1\right\} = \lim_{t\to\infty} \rho_t$. Taking two observations together, by continuity, there exists a value of $\gamma_1 \in \left(\lim_{t\to\infty} \rho_t, \rho_1\right)$ 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 perturbation argument, which is discussed in details in Remark 1.

Consider the first-order optimum $\langle \mathbf{k}^\#, \mathbf{U}^\# \rangle$ and define the slack in the upward incentive constraints as in Section 5.2, that is

$$\epsilon_t := \left(\hat{U}_t^\# - \Delta R(k^e(\theta_H)) - \delta_A(\alpha_H - \alpha_L)U_1^\#\right)^+, \quad \hat{\epsilon}_t := \left(\hat{U}_t^\# - \Delta R(k^e(\theta_H)) - \delta_A(\alpha_H - \alpha_L)U_1^\#\right)^+.$$

Then, Remark 1 implies the following bound on the profit gap $\Pi^\star - \Pi^R$:

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

Our goal is to build an upper bound on the right hand side of this expression in terms of the primitives. We shall use our standard shorthand notations: $b = \frac{\delta_A}{\delta_P}\frac{\alpha_H - \alpha_L}{1-\alpha_L}$, $a_i = \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_j}{1-\alpha_j}$ for $j = H, L$. We will also use two sequences of first-order optimal distortions $\{\hat{\rho}_t\}$ and $\{\rho_t\}$, which are defined in Theorem 1 as $\hat{\rho}_t = b\hat{\rho}_{t-1} + a_L$, $\hat{\rho}_1 = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)}$ and $\rho_t = b\rho_{t-1} + a_L$, $\rho_1 = a_H$.

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. According to Corollary 1, the first-order optimal distortion are increasing to the same limit, thus $\hat{\epsilon}_t, \epsilon_t \leqslant \lim_{s\to\infty} \epsilon_s$ for all t. It follows that B

satisfies the following inequality:

$$B \leqslant \left(\sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}) \hat{\eta}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \sum_{t=1}^{\infty} \delta_P^t \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}} | \theta_{\mathbf{H}}) \eta_t \right) \lim_{t\to\infty} \epsilon_t.$$

In the proof of Theorem 2, we showed that the term in the brackets is such that the coefficient in the Lagrangian in front of U_1 is zero. Using this result, we can rewrite the right-hand side as follows:

$$\left(\sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}) \hat{\eta}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \sum_{t=1}^{\infty} \delta_P^t \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}} | \theta_{\mathbf{H}}) \eta_t \right) \lim_{t\to\infty} \epsilon_t = \frac{\delta_P(1-\alpha_H)}{\delta_A(\alpha_H-\alpha_L)} \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \left(\gamma_1 - \lim_{t\to\infty} \rho_t \right) \lim_{t\to\infty} \epsilon_t.$$

Finally, as was shown in Theorem 2, we have that $\gamma_1 \leqslant \rho_1$, thus the following estimate:

$$\Pi^{\#} - \Pi^R \leqslant B \leqslant \frac{\delta_P(1-\alpha_H)}{\delta_A(\alpha_H-\alpha_L)} \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \left(\rho_1 - \lim_{t\to\infty} \rho_t \right) \lim_{t\to\infty} \epsilon_t =: B_a^1.$$

Our second estimate directly bounds each dual variable. Using the construction used in the proof of Theorem 2, the reader can verify that $\bar{\gamma} \leqslant \gamma_1$, thus $\gamma_{t+1} - a_L - b\gamma_t = \eta_{t+1} \leqslant \bar{\gamma}(1-b) - a_L \leqslant (1-b)\left(\rho_1 - \lim_{t\to\infty} \rho_t \right)$. Similarly, we have $\hat{\eta}_{t+1} \leqslant (1-b)\left(\rho_1 - \lim_{t\to\infty} \rho_t \right)$. At the initial date, $\eta_1 = 0$ and $\hat{\eta}_1 \leqslant (\rho - \hat{\rho}_1)^+$. Combining all pieces together, we obtain the second upper bound B_a^2 on the profit gap:

$$\begin{aligned} \Pi^{\#} - \Pi^R \leqslant & \mathbb{P}(\theta_L)\,(\rho - \hat{\rho}_1)^+\,\hat{\epsilon}_1 + \\ & + (1-b)\left(\rho_1 - \lim_{t\to\infty} \rho_t \right) \sum_{t=2}^{\infty} (\delta_P(1-\alpha_L))^{t-1} \left(\mathbb{P}(\theta_L)\hat{\epsilon}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \delta_P(1-\alpha_H)\epsilon_t \right) =: B_a^2. \end{aligned}$$

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 $k_L^S := K_L(x)$ 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, which resembles the expression for $\Pi^{\#}$, that is

$$\Pi^S := \frac{\mathbb{P}(\theta_H)}{1-\delta_P} s(\theta_H, k_H^e) + \sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}}) s(\theta_L - \Delta\hat{\rho}_t, k_L^S) + \frac{\mathbb{P}(\theta_H)}{1-\delta_P} \sum_{t=1}^{\infty} \delta_P^t \mathbb{P}(\theta_{\mathbf{L}}^{\mathbf{t}} | \theta_H) s(\theta_L - \Delta\rho_t, k_L^S).$$

Then, we have $\Pi^{\#} - \Pi^R \leq \Pi^{\#} - \Pi^S$.

So, we arrive at the following analytical bounds:

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

The former is absolute, whereas the latter is relative.

□

9.2 Recursive characterization

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, $U^\star(\theta^{\mathbf{t-1}}, \theta_L) = 0$ at all histories. Thus, even though the agent's types follow a two state Markov process, a one dimensional state variable, viz. $U(\theta^{\mathbf{t-1}}, \theta_H) = w \in \mathbb{R}_+$, suffices to encode all the required history dependence.

The following recursive formulation is equivalent to the sequence problem described in (★). From the second period onwards, for an expected promised utility of w to the high type and 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 $U^\star(\theta^{\mathbf{t-1}}, \theta_L) = 0$ 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:

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

Denote the optimal recursive contract by $\langle w^\star, \mathbf{k}(\cdot), \mathbf{z}(\cdot)\rangle$ where $(\mathbf{k}(w), \mathbf{z}(w))$ solves $(\mathcal{RP})$ for the given promise $w \geq 0$ and $(w^\star, \mathbf{k}(w^\star), \mathbf{z}(w^\star))$ solves $(\diamond)$.[30] 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.

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 $(\diamond)$.

9.2.1 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 $k_t(\theta^t) = U_t(\theta^t) = 0$ for all $\theta^t \neq \theta_H$. □

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 9.2. 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 $k_{t+1}(\theta_j, \theta^{t-1}, \theta_i) \in \mathbf{K}_i(U_{t+1}(\theta_j, \theta^{t-1}, \theta_H))$ and $U_{t+2}(\theta_j, \theta^{t-1}, \theta_i, \theta_H) \in \mathbf{Z}_i(U_{t+1}(\theta_j, \theta^{t-1}, \theta_H))$ for $i, j = H, L$ and for all θ^{t-1}.

The next claim formally connects the sequential and recursive formulations.

[30] 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).

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 hemi-continuous.*

(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 $U_{s+1}^\star(\theta_j, \boldsymbol{\theta}_{\mathbf{L}}^{\mathbf{s-1}}, \theta_H)$. 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 8, 9 and 10 below, it will continue to bind along the sequence of θ_L's. Then, we have

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

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

$$S_H'\big(z_H(w)\big) - \alpha_H \frac{\delta_P - \delta_A}{\delta_P} + (\alpha_H - \alpha_L)\frac{\delta_A}{\delta_P}\rho_H(w) \begin{cases} = 0 \text{ if } z_H(w) > 0, \\ \leqslant 0 \text{ if } z_H(w) = 0, \end{cases} \tag{8}$$

$$S_L'\big(z_L(w)\big) - \alpha_L \frac{\delta_P - \delta_A}{\delta_P} - (\alpha_H - \alpha_L)\frac{\delta_A}{\delta_P}\rho_L(w) \begin{cases} = 0 \text{ if } z_L(w) > 0, \\ \leqslant 0 \text{ if } z_L(w) = 0. \end{cases} \tag{9}$$

In addition, the Envelope theorem gives:

$$S_j'(w) = (1 - \alpha_j)\rho_L(w) - \alpha_j \rho_H(w) \text{ for } j = H, L. \tag{10}$$

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 $k_t^\star(\theta^{\mathbf{t-2}}, \theta_j, \theta_H) > 0$ is uniquely defined by strict concavity of the objective. It follows from Claim 1 that $\rho_H(w) = \mathcal{K}_H^{-1}\left(k_t^\star(\theta^{\mathbf{t-2}}, \theta_j, \theta_H)\right)$, 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.

9.2.2 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

($\diamond$).[31] 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 $z_H(w) = z_H^e$ 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

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

[31] 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.

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 10) 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.

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,

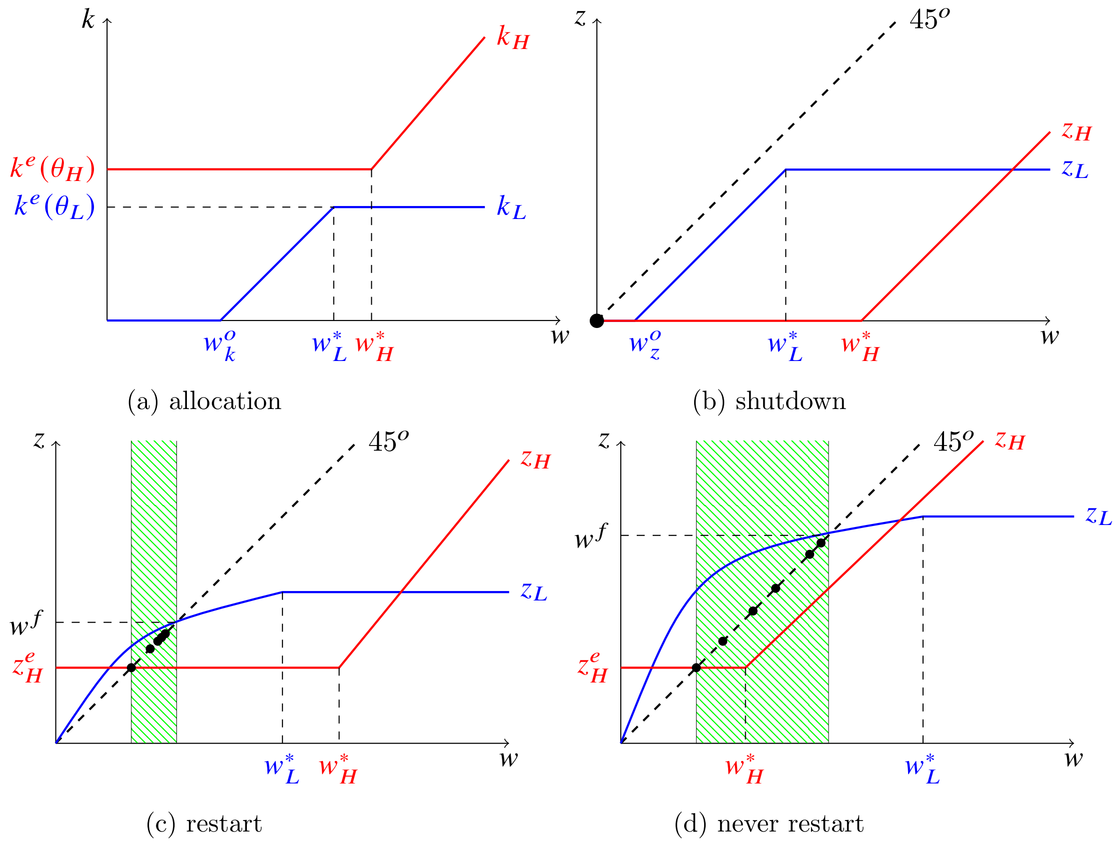

Figure 5: Optimal recursive contract

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.[32] 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* $w \geqslant w_L^*$,

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

$z_L(\cdot)$ 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 $\theta_L > \frac{\alpha_L}{1-b}\Delta\theta$.

(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 10:

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

In fact, $z_L^e \neq z_H^e$ if and only if $S_L'(0) > \alpha_L \frac{\delta_P-\delta_A}{\delta_P}$.

We showed above that $z_j^e \in [0, w_j^*)$ for $j = H, L$. Then, monotonicity of $\rho_L(\cdot)$ and $\rho_H(\cdot)$, as shown in Proposition 2, combined with Equations 8 and 9 translates into monotonicity of both $z_L(\cdot)$ and $z_H(\cdot)$. Finally, we set $w_z^o := \max\{w \geqslant 0 : z_L(w) = 0\}$ 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 $z_H^e < w_H^*$ that implies that z_H^e 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 $w \neq z_H^e > 0$, 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 $\rho_H(w) > 0$:

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

To obtain a contradiction, combine Equations 8 and 10:

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

Since $\rho_L(w) > 0$, the downward constraint binds this period, and it will keep binding a sequence of θ_L's. Formally, let $z_L^s(w)$ be a result of s successive applications of $z_L(\cdot)$ to w, that is $z_L^s(w) := z_L\big(z_L^{s-1}(w)\big)$ with $z_L^0(w) = w$. By Equation 9, $\rho(z_L^s(w)) > 0$ for any s. Then, iterating along this sequence, we arrive at the following equation:

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

which clearly contradicts the premise. As a result, z_H^e 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, $z_L(w) = z_L^e < w_L^* \leqslant w$ whenever $\rho_L(w) = 0$, therefore it must be the case that $w < z_L^e$ 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 9 and 10, 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 $w^f < z_L^e$. 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 $z_j^e < w_j^*$ and $z_H^e \leqslant z_L^e$. So, it remains to establish that $z_H^e \leqslant w^f$. Of course, it is vacuously true whenever $z_H^e = 0$, thus it is without loss of generality, to assume that $z_H^e > 0$. By contradiction, suppose that $w^f < z_H^e$. Since $\rho_H(\cdot)$ is monotone and $z_H^e \leqslant w_H^*$, 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 $\rho_L(\cdot)$, $z_H^e < w^f$ implies $\rho_L(w^f) > \rho_L(z_H^e) = a_H$. As a result, $\frac{a_L}{1-b} > a_H$ that is a clear contradiction. Conclude that $z_H^e \leqslant w^f$.

Part (e). At the initial date, the first-order conditions with respect to $\mathbf{z}(\cdot)$ coincide with Equations 8 and 9. 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 $w^* \in [z_H^e, w^f]$. By contradiction, suppose that $w^* < z_H^e$. 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 $\rho_H(z_H^e) \geqslant \rho_H(w^*) > 0$ that is a contradiction. Conclude that $w^* \geqslant z_H^e$.

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.

9.2.3 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 $\left\{k | \exists \theta^t : k = k_t(\theta^t), t \leqslant T\right\}$ 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

$$\left| \left\{ U | \exists \theta^{t-1} : U = U_t(\theta^{t-1}, \theta_H), t \leqslant T \right\} \right| \geqslant K 2^T.$$

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 $K = 2^{-\tau}$. □

9.3 Comparative statics

Here we present two results which are intuitively sketched in Section 7.

Corollary 6. *Let $\{(\delta_P^n, \delta_A^n)\} \in (0,1)^2$ be a convergent sequence of discount factors such that $\delta_P^n \geqslant \delta_A^n$ for all n. Let $\Pi^{\star,n}$ be the profit at the optimal contract for (δ_P^n, δ_A^n). Then, $(1-\delta_P^n)\Pi^{\star,n} \to s^e$ if and only if $\delta_A^n \to 1$.*

Proof. First of all, note that the optimal contract only depends on the agent's relative patience $\frac{\delta_A}{\delta_P}$ and absolute patience δ_A. By the theorem of maximum, the optimal contract is a continuous function of (δ_A, δ_P) on $0 < \delta_A \leqslant \delta_P < 1$.

Take the convergent sequence $\{(\delta_A^n, \delta_P^n)\}$ as described in the corollary. Suppose that $\delta_A^n \to 1$, then δ_t^P also converges to 1, because $\delta_A^n \leqslant \delta_P^n$ for all elements. It follows that the optimal contract convergences to the first-order optimum, because the latter is always incentive compatible for $\delta_A = \delta_P$. 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 $\delta_P^n \to 1$, 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 $\delta_A^n \to 1$ is necessary for the full surplus extraction. By construction, the value of the first-order program $\Pi^{\#,n}$ is an upper bound on $\Pi^{\star,n}$. Since $\delta_A^n \leqslant \delta_P^n \leqslant 1$ for all elements, the ratio $\frac{\delta_A^n}{\delta_P^n}$ is also convergent. There are two cases two consider.

First, suppose that $\lim_{n\to\infty} \delta_P^n = 1$. Then, $\lim_{n\to\infty} \frac{\delta_A^n}{\delta_P^n} < 1$, therefore the first-order optimal distortions are positive in the restart phase, i.e., $\rho_t > 0$. It follows that the limit of $(1-\delta_P^n)\Pi^{\#,n}$ is well defined, and $s^e > \lim_{n\to\infty}(1-\delta_P^n)\Pi^{\#,n} \geqslant \lim_{n\to\infty}(1-\delta_P^n)\Pi^{\#,n}$.

Finally, suppose that $\lim_{n\to\infty} \delta_P^n < 1$. In this case the low type distortion in the first period is positive, i.e., $\hat{\rho}_1 > 0$. As a result, the principal's profit is strictly less than the surplus: $s^e > \Pi^{\#} \geqslant \Pi^*$. □

Corollary 7. *Fix $\delta_P \in (0,1)$. 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 follows:*

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

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

Proof. 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 cost 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 $\frac{1}{2}$.

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.

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.

A. Bisin, A. Lizzeri, and L. Yariv. Government policy with time inconsistent voters. *American Economic Review*, 105(6):1711–1737, 2015.

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

K. Chatterjee and H. Sabourian. Game theory and strategic complexity. In R. A. Meyers, editor, *Encyclopedia of Complexity and Systems Science*, pages 13–30. Springer, 2009.

G. L. Clementi and H. A. Hopenhayn. A theory of financing constraints and firm dynamics. *Quarterly Journal of Economics*, 121:229–265, 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. Fudenberg and D. K. Levine. Reputation and equilibrium selection in games with a patient player. *Econometrica*, 57(4):759–778, 1989.

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

D. Gottlieb and X. Zhang. Long-term contracting with time-inconsistent agents. *Econometrica*, 89(2):793–824, 2021.

S. Kapon. Dynamic amnesty programs. New York University, 2021.

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. *Journal of Economic Theory*, 193(105196), 2021.

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.

D. Laibson. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112 (2):443–477, 1997.

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.

E. F. Moore. Gedanken-experiments on sequential machines. In C. E. Shannon and J. McCarthy, editors, *Automata Studies, Annals of Mathematics Studies Volume 34*, pages 129–153. 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.

A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

Y. Sannikov. Dynamic security design and corporate financing. In G. M. Constantinides, M. Harris, and R. M. Stulz, editors, *Handbook of the Economics of Finance*, volume 2A, chapter 2, pages 71–122. Elsevier, 2013.

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.