Implications of unequal discounting in dynamic contracting*

Ilia Krasikov Rohit Lamba Thomas Mettral

September 2019

Abstract

This paper studies a canonical dynamic screening model where the principal's discount factor is larger than the agent, the agent has limited commitment and payoff relevant private information that follows a Markov process. The interaction of unequal discounting and limited commitment with persistent agency frictions produces a novel tradeoff: (i) new *intertemporal costs of incentive provision* emerge, and (ii) the net present value of the *standard information rent* decreases. The former ensure that the shadow price of incentive constraints are permanently positive, and the latter contributes towards decreasing distortions since principal and agent evaluate future payoffs differently.

The optimal contract mostly exhibits a rather simple cyclical form that we term *restart*: (i) distortions decrease monotonically in the consecutive number of low shocks; (ii) a high shock erases all previous history of distortions, and then (iii) for every consecutive low shock, distortions follow the same path as before. Invoking an automaton inspired definition, restart contracts are shown to be *simple*. The optimal restart contract is (globally) optimal when the relaxed approach works, and approximately optimal otherwise.

The setup admits a host of applications where one party is "financially bigger" and the other is armed with some private information. Examples include a venture capitalist-entrepreneur relationship, loan contracts between the International Monetary Fund and emerging markets, and governments redistributing amongst heterogeneous citizens.

1 Introduction

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

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

*Krasikov: Penn State University, izk113@psu.edu; Lamba: Penn State University, rlamba@psu.edu; Mettral: Allianz Life Insurance, t.mettral@gmx.de. We are indebted to Leeat Yariv for conversations 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 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*.

The objective of this paper is to understand how the aforementioned tradeoff evolves when the principal and agent contract over time and the principal faces more favorable interest rates than the agent. It studies the interaction of three *forces*– (i) the agent has payoff relevant private information that follows a Markov process, (ii) the principal has commitment power while agent has limited commitment, and (iii) the principal is more patient than the agent. Taking away the third force, novel to this paper, would put us in the rubric of standard dynamic mechanism design models (see Bergemann and Välimäki [2019]), and in addition, assuming that the principal lacks commitment would make our setup akin to a stochastic game (see Hörner et al. [2011]).

Our setup admits a host of applications of principal-agent scenarios, where one party is "financially bigger" and the other is armed with private information, for example: a venture-capitalist plush with cash and an entrepreneur with hidden information about the viability of a project, and an large lending agency such as the International Monetary Fund and a emerging market looking for some capital infusion. In the "macro" interpretation of the model, it can also be applied to taxation environments where a government wants to redistribute amongst citizens with heterogenous shocks in labor productivity.[1]

Think through the following simple example. Suppose the principal and agent interact for two periods. They can make payments in both periods, but trade happens only in the second period. The agent's type, that is his value of the object for trade, can assume one of two numbers with some probabilities. No information is generated in the first period, and the agent privately learns the type in the second period. If the principal and agent sign a contract in the first period, the principal can essentially extract all the surplus from this dynamic contract reducing the agent's payoff to his reservation value, irrespective of his type realization in the second period. The principal is still paying information rent to the agent in the second period, but all of it is extracted in the form of upfront payments in the first period. The *shadow price* of incentives is thus zero– the principal implements the efficient contract and takes the efficient value of the total surplus as profit.[2] Allowing for dynamic contracts relaxes distortions characteristic in the static model.[3]

A crucial assumption in reaching the above conclusion is the absence of any financial or collateral constraints. The agent has enough funds to pay upfront for the whole value of the "firm" or can raise external capital at the same rate as the principal. In the presence of financial constraints, the principal's ability to recover future information rents upfront is restricted. As a result, in the context of the above example, the agent's limited commitment constraint in the second period starts binding, which in turn implies that the shadow price of providing incentives in the second period is no longer zero. We are therefore back to allocative distortions that determine the rent versus efficiency tradeoff, as in the static model.[4]

---

[1] In each of these cases the principal is mentioned first (venture capitalist, IMF, government) followed by the agent (entrepreneur, emerging market, citizen). Moreover, in each case, it is reasonable to assume that the principal has deeper financial capacity than the agent, the agent has some private information, and the agent has limited commitment.

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

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

[4] Allocative distortions mean the wedge between the optimal allocation and the efficient allocation.

Previous work has analyzed hard financial constraints wherein the agent does not have deep pockets and cannot raise external capital (see, for example, Krishna et al. [2013], Krähmer and Strausz [2015], and Krasikov and Lamba [2018]). In this paper, we look at soft(er) financial constraints: the agent can obtain financing but at a higher cost, while maintaining the assumption that agent has limited commitment. These two qualitative assumptions are modeled through unequal discounting, $\delta_P > \delta_A$, and a participation constraint on the agent's lifetime utility in each period.

Under this setup a novel tradeoff emerges. Think back to the two period example. Suppose the agent has to paid an information rent worth x in the second period. The principal can extract a maximal amount of $\delta_A \times x$ in the first period because any higher price will not be accepted by the agent. Therefore, an added cost of delivering an information rent of x in the second period is introduced which is given by $-(\delta_P - \delta_A) \times x$. We call this the *intertemporal cost of incentive provision*, distinct from the *standard information rent* that has to be paid in the second period. It pushes in the direction of more binding incentive constraints, and hence greater allocative distortions.[5]

Now, consider the multi period version of the stated example with trade in every period. Here for a fixed value of per period surplus (say $s_1, s_2, s_3, \ldots$) and standard information rent (say $x_1, x_2, x_3, \ldots$), the principal has a higher net present value of the former ($\sum_t \delta_P^{i-1} s_t$) and smaller value of the latter ($\sum_t \delta_A^{i-1} x_t$), since these are calculated under δ_P and δ_A respectively. This forms the benefit side of the aforementioed tradeoff. It pushes to reduce the shadow price of incentives, and hence decrease allocative distortions.

In describing the above tradeoff, we assumed a fixed value of dynamic information rent; at the optimum, these forces interact to endogenously pin down the optimal value of economic surplus and information rents, which in turn characterizes the extent of inefficiencies in our environment. In what follows we provide a brief description of the model, the forces at play in special cases of our model, and an intuitive characterization of dynamic distortions with all three key ingredients: Markov types, limited commitment and unequal discounting.

The formal model entails a "small" firm (agent) with a production technology whose total factor productivity (TFP) changes periodically according to a two state Markov process, and a "large" supplier (principal) of capital that is critical for production. The principal is more patient than the agent, and the realization of the TFP shock is privately observed by the agent. A contract here is a dynamic menu of capital allocations to the agent in return for payments to the principal. We solve for the profit maximizing contract of the principal subject to incentive compatibility and individual rationality constraints for the agent, where the latter constraint captures limited commitment on part of the agent.

Suppose the two TFP types are given by θ_H and θ_L, where $\theta_H > \theta_L$. The usual *no distortion at the top principle* implies that θ_H is always allocated the efficient capital: $k_H^t = k_H^e$ for all t. The challenge is to solve for the optimal dynamic distortion of the capital allocation to the low type.

[5]Note that the limited commitment assumption is important here, in the absence of which the principal is guaranteed an arbitrage opportunity with no extra costs. Suppose the principal lends x to the agent, and demands it back with interest in the second period. Then, the payment to the principal is $-x + {}^{\delta_P}/_{\delta_A} x > 0$ and agent breaks even in payments across the two periods, x-$\delta_A \times {}^{1}/_{\delta_A} x = 0$. Obviously, the principal will want to maximize the gains from this channel by lending an arbitrarily large amount of money to the agent and demanding it back in the next period. The limited commitment assumption, that the agent needs to break even in the second period as well, precludes this arbitrage channel.

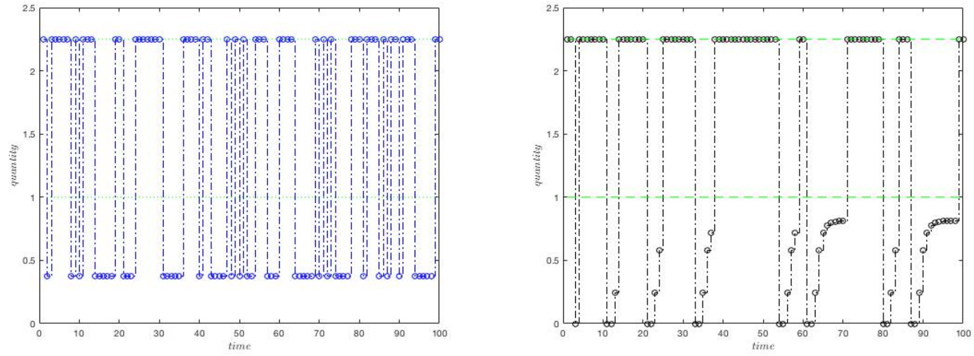

(a) IID types with unequal discounting (b) Correlated types with unequal discounting

Figure 1: Sample of allocations across time

As a first step towards understanding the structure of distortions, consider a special case of our model where the agent's type follow an iid process. Heuristically speaking, distortions in capital allocation take a very simple form:

$$k_L^1 = k_L^e - (\text{static distortion})\,, \quad \text{and} \quad k_L^t = k_L^e - \left(1 - \frac{\delta_A}{\delta_P}\right) \cdot (\text{static distortion}) \;\text{ for } t > 1.$$ [6]

where k_L^t is the capital allocation to the low type in the t-th period.

In the iid model, every binding incentive constraint produces distortions that have a one-period memory. With equal discounting therefore, only the first period allocation is distorted, the shadow price of all future incentive constraints is reduced to zero by paying for the information rent up-front. However, with unequal discounting, this shadow price is now permanently positive through the novel *intertemporal cost of incentive provision*. The distortions, though, do not propagate into the future.[7] Figure 1a depicts the permanent cycle of static distortions for the iid case. The "high" type gets the efficient allocation, and the "low" type is served a distorted allocation where the extent of the distortion is its distance from the dotted horizontal line.[8]

The main applications of this theoretical framework would realistically demand the agent's private information be correlated across time. So the main focus of our paper is to characterize how dynamic distortions would evolve over time with Markovian agency frictions. As before, the high type gets the efficient allocation. However, the distortions for the low type now have infinite memory that assume a very special structure. Again, writing heuristically:

[6]Here k_H^e and k_L^e refer to the efficient capital allocations that would be supplied if the TFP realizations (θ_H and θ_L respectively) were publicly observed. Hence a distortion is the wedge between the efficient and the optimal allocation. In the quasi-linear framework that we follow, this wedge is essentially pinned down by the Myerson virtual value (see Myerson [1981] and Pavan et al. [2014]).

[7]In the iid model, the second (benefits) channel mentioned above does not play a role. Since the *standard information rent* each period itself is static, there is no benefit to the principal from the agent valuing future utils less than her.

[8]In Figure 1 the x-axis represents time and y represents the allocation corresponding to the time for a path of realized private types of the agent. The two dotted horizontal lines at 2.3 and 1 represent the efficient allocation for the high and low types.

$$k_L^t = k_L^e - \left[\frac{\delta_A}{\delta_P} \cdot (\text{standard dynamic distortion})_\tau + \left(1 - \frac{\delta_A}{\delta_P}\right) \cdot (\text{new intemporal distortion})_\tau\right]$$

where $t - \tau \leq t$ is the most recent high type realization in the t-period history of TFP shocks, that is, the history of shocks can written as

$$\theta^t = (\theta^{t-\tau-1}, \theta_H, \underbrace{\theta_L, ..., \theta_L}_{\tau-\text{times}}),$$

so that the exact composition of $\theta^{t-\tau-1}$ is irrelevant, and the distortions are determined as if the contract is τ periods old.

The boxed equation above is the backbone of the optimal contract, and we briefly explain the key economic forces that lead to that exact structure. It can straightway be noted that dynamic distortions are simply a function of the number of consecutive low shocks since the last high shock. Each time a high shock arrives, the memory of distortions is completely erased and future distortions form a repetitive chain. The evolution of distortions is represented clearly in Figure 1b. We call this cyclical pattern of optimal distortions *restart*.

To understand the pattern of distortions, consider the special case of our model with equal discounting ($\delta_A = \delta_P$) and Markov types. There the impact of consecutive low shocks is permanent but gradually weakens over time, and the realization of a high shock ends distortions forever; Battaglini [2005] calls these "vanishing distortions at the bottom" and "generalized no distortion at the top", respectively. The main economic force behind both the properties is *backloading* of the agent's payoffs, which helps minimize the shadow price of incentives. With unequal discounting backloading is costly: for every x that is backloaded only $\delta_A x$ can be recouped upfront.

Two main differences arise in the structure of optimal distortions with unequal discounting: First, the arrival of a high shocks erases the legacy of past binding incentive constraints, however, the shadow price of future incentive constraints still remains positive. Second, akin to the iid model with unequal discounting, the first low shock seeds distortions anew, but due to the Markovian nature of shocks, these distortions now propagate along the sequence of consecutive low shock realizations. Old distortions carry on, weakening over time, and new seeds are added very period due to repeatedly binding incentive constraints. The monotonic nature of these distortions are analytically pinned down: they decreases over time and converge to a fixed positive value.

What happens if the ex ante agency problem, as measure by the magnitude of $\theta_H - \theta_L$, is particularly acute? As can be inferred from Figure 1b, for consecutive low shocks the optimal allocation first falls and then rises to converge to a fixed value. In the figure this convergent value lies above zero. However, if the agency problem is acute, the distortions do not decrease enough for the allocation to converge to a positive number. In such a situation the optimal contract excludes the low type, it gets zero supply across time; we call this feature *shutdown*.

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

As noted earlier, financial constraints in dynamic mechanisms have been modeled as a restriction on the magnitude of per-period transfers or limited liability constraints for the agent, but maintaing equal discounting. In such an environment, distortions actually increases with consecutive low shocks and at any given point a consecutive number of high shocks would make the shadow price of incentives zero. In fact the contract would reach the efficient allocation almost surely (see Krasikov and Lamba [2018]). Thus, the agent can still win his way towards provision of liquidity, and the optimal allocation becomes efficient; however, with financial constraints modeled here as a constant asymmetry in the interest faced by the two parties, inefficiency is permanent, and more specifically it is cyclical.[9]

We also make two other conceptual contributions, on simplicity and approximate optimality in dynamic contracts. Unequal discounting leads to the downward and upward incentive constraints binding simultaneously for certain parameters, even though we operate in a type space that has a cardinality of two. The optimal contract then loses the restart feature and can have a very complicated form. We do two things. First we develop a notion of simplicity of the contractual space in our model using ideas analogous to automaton in repeated games. In this realm, any restart contract is simple, and the optimal contract is simple if and only if it is restart. Second, when, the optimal contract is indeed not restart, we propose the optimal restart contract, that is we optimize over the space of restart contracts, and show, by constructing a theoretical bound, that the loss in profit from this restriction is minimal.

The paper ends with comparative statics on discounting and persistence of Markovian shocks, and an overview of literature and potentially promising ideas for future work.

2 Model

2.1 Primitives

A firm (agent) with access to a production technology approaches a supplier (principal) of a key input; the former is a "small player" while the latter is a "big player" in the market.[10] The total factor productivity (TFP) 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 total factor productivity, and R is a concave production function that satisfies Inada

[9]We discuss in more detail in Section 6, the connections to other dynamic models, especially in public finance and financial contracting, where unequal discounting generates analogous cycles in optimal allocation.

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

conditions.[11] The principal's stage utility is simply $p - k$.[12]

TFP or technology "shocks" can take values in $\Theta = \{\theta_H, \theta_L\}$, where $\theta_L > 0$ and $\theta_H - \theta_L = \Delta\theta > 0$. We will often refer to it as the agent's type. The types follow a Markov chain with transition probabilities $\mathbb{P}(\theta_H|\theta_i) = \alpha_i$, which satisfies first-order stochastic dominance: $\alpha_H \geqslant \alpha_L$. To simplify calculations, we assume that the prior distribution coincides with the invariant distribution of Markov process, that is $\mathbb{P}(\theta_H) = \frac{\alpha_L}{1-\alpha_H+\alpha_L}$. The principal does not observe the output, and therein lies the asymmetric information or agency friction.

We consider an infinite horizon setting where the principal and agent discount future utility. However, critically, we *do not restrict them to have the same discount factor*; these are denoted by δ_P and δ_A, respectively, where $\delta_P \geqslant \delta_A$. The concept of discounting or time preference 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. Krueger and Uhlig [2006] write that different discount factors in principal-agent models can be interpreted as "the gross real interest rate or the return to some storage technology the principal has access to."

The principal can commit to a long-term contract. Then, invoking the revelation principle, it is without loss 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 the history with t consecutive reports of type θ_j by $\theta_{\mathbf{j}}^{\mathbf{t}}$.[13] The principal's objective is to maximize her profit subject to incentive compatibility and participation constraints for the agent. For a fixed mechanism, the agent faces a dynamic decision problem in which her strategy is simply a function that maps his private history into an announcement every period.[14]

2.2 Constraints

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

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

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

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

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

[12] 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 àla Laffont and Tirole [1993] where the principal and agent have preferences $V(k) - p$ and $p - \theta k$ respectively, or the monopolistic screening model àla Mussa and Rosen [1978] where the principal and agent have preferences $p - k^2/2$ and $\theta k - p$, respectively.

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

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

as:[15]

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

$\forall \theta^{\mathbf{t}}, \hat{\theta}_t, \forall t$. Equivalently, incentive compatibility can be expressed directly in terms of $\langle \mathbf{k}, \mathbf{U} \rangle$:

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

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 essential records the fact with Markovian shocks, knowing his type today also gives some information to the agent about his types in the future. It is useful to partition the set of incentive compatibility constraints into *"downward"* (IC_H) corresponding to $\theta_t = \theta_H$ and $\hat{\theta}_t = \theta_L$; *"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 $U_t(\theta^{\mathbf{t}}) \geqslant 0\ \forall \theta^{\mathbf{t}}$. 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 or the agent– he cannot be forced to be in the contractual relationship. The set of participation constraints are analogously partitioned into IR_H for $\theta_t = \theta_H$ and IR_L for $\theta_t = \theta_L$.

2.3 Optimization problem

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

The static surplus (under truthful revelatin) 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

$$\bar{U}_P = \sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{E}\left[u(\tilde{\theta}_t|\tilde{h}^{t-1})\right] \quad \text{and} \quad \bar{U}_A = \sum_{t=1}^{\infty} \delta_A^{t-1} \mathbb{E}\left[u(\tilde{\theta}_t|\tilde{h}^{t-1})\right]$$

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

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

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

$\bar{U}_A$ 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:

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

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

2.4 Virtual value

To conceptualize the solution to problem ($\star$), we shall introduce the notion of Myersonian virtual value (Myerson [1981]). In our *quasi-linear* environment, define

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

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

$$\mathcal{K}_H(x) = arg\max_k \quad \phi_H(x)R(k) - k, \quad \mathcal{K}_L(x) = arg\max_k \quad \phi_L(x)R(k) - k.$$

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

To fix ideas, note that in the static model, the solution would provide the efficient capital to the high type, and a distorted allocation to the low type given by $\mathcal{K}_L(x)$ for $x = \frac{\mathbb{P}(\theta_H)}{\mathbb{P}(\theta_L)} = \frac{\alpha_L}{1-\alpha_H}$. In our setup, this value of x evolves over time as a function of the level of asymmetric information captured by the Markov process, and the fact that the principal and agent discount future payoffs at different rates.

3 Restart contracts

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

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

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

The restart property is modeled as a measurability restriction on the allocation rule: all relevant history dependence is encoded in the number of consecutive low shocks since the last high realization. The allocation is completely characterized by the number k_H and two sequences $\{k_t\}$ and $\{\hat{k}_t\}$. The first sequence, $\{k_t\}$, defines the allocation for consecutive low shocks after a high shock has been realized, and the second sequence, $\{\hat{k}_t\}$, defines the allocation to the low type along the "lowest" history, where the high type has never been realized in the past.[16]

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

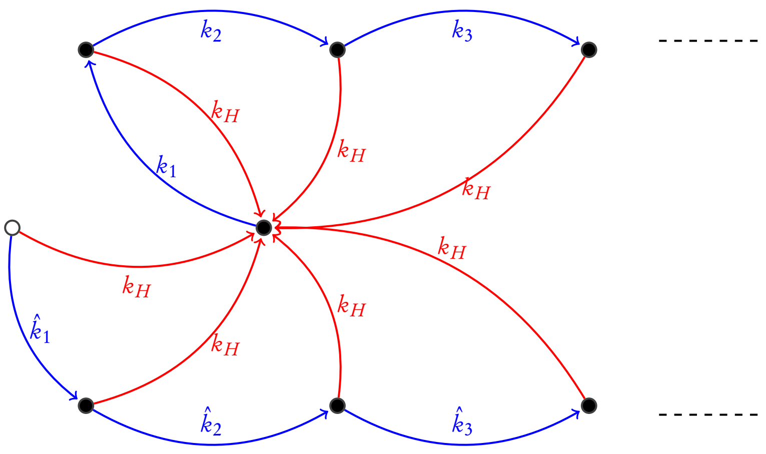

Figure 2: The evolution of allocation in a restart contract. A red/blue arrow indicates a transition, because of a high/low report.

Figure 2 exposits the evolution of restart contracts. The contract starts in the white circle, and then evolves dynamically. If the agent reports θ_H, then k_H is supplied irrespective of the previous history. If θ_L is reported in the first period then allocation is $\hat{k}_1$, followed by $\hat{k}_t$ for every subsequent announcement of θ_L. If θ_L is reported immediately after θ_H, then k_1 is allocated, followed by k_t for every subsequent announcement of θ_L. The restart feature is captured by the fact the allocation always returns to k_H on the realization of a high shock, and remains there until a low shock is realized, which triggers the sequence k_t once again.

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

3.1 Relaxed problem

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

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

We will denote the solution to this problem by $\langle \mathbf{k}^{\#}, \mathbf{U}^{\#} \rangle$ and its profit by $\Pi^{\#}$. This is often referred to in the literature as the *first-order optimum*, because it only takes the "first-order constraints" into account. Interestingly, $k^{\#}$ satisfies the restart property, and by construction it provides an upper bound for the optimal profit, that is $\Pi^{\star} \leqslant \Pi^{\#}$.[17] In what follows we illustrate how to obtain the first-order optimum and provide a closed-form solution.

Start by rewriting IC_H as follows:

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

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

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

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

Equation (1) gives the expression for the agent's expected utility in terms of the allocation:

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

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

where $\{\hat{\rho}_t\}$ and $\{\rho_t\}$ are measures of agent's information rents, respectively for the lowest history where no high type is ever realized and the *restart phase* where at least one high type has been realized at some point. Recall the definitions of $\mathcal{K}_H$ and $\mathcal{K}_L$ from Section 2.4.

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

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

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

The high type allocations are always efficient, the low type allocations starts with the seed distortion $\hat{\rho}_1 = \frac{\alpha_L}{1-\alpha_H}$, which is equal to the optimal distortion in the static model. From then on, every successive low type carries over the previous distortion with a *multiplicative factor* $b = \frac{\delta_A}{\delta_P}\frac{\alpha_H-\alpha_L}{1-\alpha_L}$ and adds to it a new *L-seed*, $a_L = \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_L}{1-\alpha_L}$, culminating in $\hat{\rho}_t = b\hat{\rho}_{t-1} + a_L$. The moment a high shock arrives all previous distortions are erased. Now, the realization of a new low type leads to an *H-seed* distortion $a_H = \left(1 - \frac{\delta_A}{\delta_P}\right)\frac{\alpha_H}{1-\alpha_H}$, different than a_L since the last period type was high. Again, for every successive low type a multiplicative factor b carries over the previous distortion and adds to it a new seed a_L, culminating in $\rho_t = b\rho_{t-1} + a_L$. The key difference in the structure of distortions from the equal discounting model is that now distortions do not disappear completely once a high shock arrives, they simply loose memory; this is because new distortions are seeded again on the arrival of a low shock.

That seed distortions a_H and a_L are created by unequal discounting can be readily seen by observing $\lim_{\delta_A \to \delta_P} a_H = \lim_{\delta_A \to \delta_P} a_L = 0$. For every unit of information rent that has to be paid to the agent, unequal discounting creates the three novel economic forces: (i) backloading is costly, hence IR_L permanently binds, (ii) new intertemporal costs of incentive constraints are introduced, hence IC_H binds permanently, and (iii) net present value of standard information rent goes down. These interact to endogenously pin down the optimal level of allocative distortions.

A seed distortion is created every time an individual rationality constraint binds, and it is propagated through multiplicative distortion created by the linking of binding incentive compatibility constraints. Note that propagation through b is driven by the Markovian nature of shocks; $\lim_{\alpha_H \to \alpha_L} b = 0$. Therefore for iid shocks, captured by $\alpha_H = \alpha_L$, distortions are permanently seeded but have no memory, since there is no propagation.

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

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

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 satisfies the following properties:*

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

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

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

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

The distortions dynamics here are distinct than both the equal discounting model without financial constraints (eg. Battaglini [2005] and Pavan et al. [2014]), and the equal discounting model with hard financial constraints (eg. Krishna et al. [2013] and Krasikov and Lamba [2018]). In the former case, depending on the generality of model, the distortions are monotonically decreasing and the contract converges to the efficient allocation in the limit, either along every history, almost surely, or at least on average. In the latter case (modeled as $\mathbf{u} \geqslant 0$ as opposed to $\mathbf{U} \geqslant 0$) the distortions are monotonically increasing in the bad (or low) shocks, but the contract still does converge almost surely to the efficient allocation. Thus, a soft but permanent financial constraint the form of $\delta_A < \delta_P$ can create greater (long-term) inefficiency than a hard financial constraint in the form of $\mathbf{u} \geqslant 0$; the latter can eventually be overcome in a long enough relationship if the interest rates faced by both parties are the same.

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

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

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

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

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

Recollect that distortions in the restart phase depend only on the number of low shocks since the last high shock (Theorem 1), and they are monotonically decreasing along consecutive low cost realizations (Corollary 1(a)). So, the tightest upward incentive constraint in the restart phase is the "one at infinity". Moreover, distortions along the lowest history and in the restart phase converge to the same value (Corollary 1(b)). Putting these together we get the following simple result.

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

$$\lim_{t\to\infty} U_t^{\#}(\theta_L^{t-1},\theta_H) \leqslant \Delta\theta R(k_H^e) + \delta_A(\alpha_H - \alpha_L)U_2^{\#}(\theta_H^2).$$

Figure 3 partitions the parameter space along the set of binding constraints for a specific example. White and yellow regions represent the validity of the relaxed problem approach, where the optimal contract is restart; the dark region is the space where the optimal contract is never restart and upward incentive constraints bind infinitely often. The white portion in the southwest corner also represents the case of (permanent) shutdown, no capital is supplied to the low type.

For larger values of $\Delta\theta$, signifying greater ex ante asymmetric information, it is easier to separate the two types, and hence the upward constraint does not bind often, culminating in the optimal contract being restart in most of the parameter space. But for larger value of $\Delta\theta$, the gain from serving both types also decreases, so we see more shutdowns.[18]

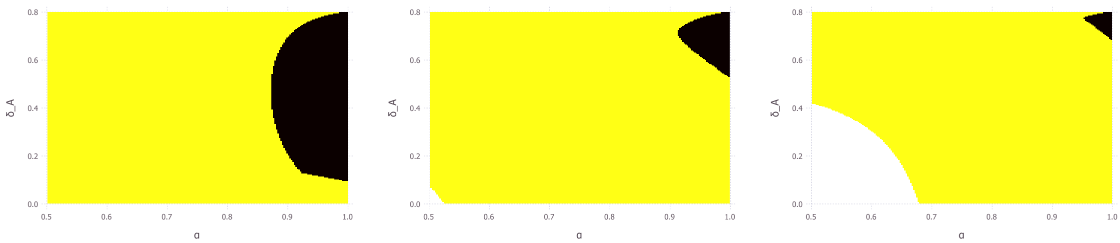

(a) $\Delta\theta = 0.1$ (b) $\Delta\theta = 1.1$ (c) $\Delta\theta = 2.1$

Figure 3: Partitioning parameter space into set of binding constraints. White & yellow: first-order approach works and optimal contract is restart. White: low type is shutdown. Black: upward constraint binds ad infinitum where $\alpha_H = 1 - \alpha_L = \alpha$ on the x-axis, δ_A on the y-axis; $\delta_P = 0.8$, $R(k) = 2\sqrt{k}$, $\theta_L = 1$.

To summarize the discussion on primitives, the first-order optimal contract satisfies the restart property such that the highest distortion occurs with the first low shock after which distortions

[18] In terms of the primitives, note that for $\delta_A = 0$ the optimally contract is trivially restart, the same is true for the iid model ($\alpha_H = 1 - \alpha_L$) and perfect persistence ($\alpha_H = 1 - \alpha_L = 1$). In the intermediate cases, discounting and persistence interact in a non-linear fashion. Technically speaking there is a "discontinuity" in the limit for both discounting and persistence. When the agent is almost as patient as the principal and the types are highly persistent then upward constraints start binding.

progressively decrease to some constant and positive value (seen clearly in Figure 1 from the introduction). As a result, $U^{\#}_{t+1}(\theta^{t-1}, \theta^2_H) = U^{\#}_2(\theta^2_H)$ is the lowest element of the expected utility vector for the high type, and $U^{\#}_{t+s}(\theta^{t-1}, \theta_H, \theta^{s-1}_L, \theta_H)$ is a decreasing function of s. Therefore, for a large enough s, it is possible that the non-monotonicity emanating from the fact that $U^{\#}_{t+s}(\theta^{t-1}, \theta^s_L, \theta_H) > U^{\#}_{t+s}(\theta^{t-1}, \theta_H, \theta^{s-1}_L, \theta_H)$ can violate upward incentive constraints.

3.2 Restart optimum

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

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

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

There are three reasons for focusing on the class of restart contracts: (i) it is a fairly intuitive criterion and simple to describe, (ii) the (global) optimal contract falls within this class for a large measure of parameters, and (iii) even when the optimal contract is not restart, it continues to be close to a restart contract in spirit and especially in the magnitude of loss. 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 demand incentive compatibility.

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

Theorem 2. *There exists $\bar{\gamma}$ such that the restart optimum is as follows: $k^R_H \geqslant k^e_H$, and*

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

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

The optimal distortions along the two class of histories, $\{\hat{\gamma}_t\}$ and $\{\gamma_t\}$, are given in Theorem 2. These are obviously analogous to their counterparts from the first-order optimal contract (Theorem 1), but there are three key differences: (i) the high type allocation is (potentially) distorted upwards, (ii) the initial seed at the lowest history is (weakly) higher and that in the restart phase is lower, and (iii) there is a floor on distortions, so if the floor binds, the contract has a finite memory along consecutive low TFP shocks.[20] The initial allocations, determined by three numbers k^R_H,

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

[20] 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.

γ_1 and $\hat{\gamma}_1$, are picked using the first-order conditions presented in the appendix. Finally, the floor $\bar{\gamma}$, is uniquely determined according to the complementary slackness of the corresponding upward incentive constraints.

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

We estimate the gap using sensitivity analysis. Attach a Lagrange multiplier to each upward incentive constraint and evaluate the multipliers at the restart optimum. Quantify how much slack needs to be added to these constraints so that the solution then coincides with the first-order optimum.[21, 22] The estimate can then be written as

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

Corollary 4. *There exists two bounds, B_a and B_r, functions of primitives, such that $\Pi^\star - \Pi^R \leqslant B_a$ and $1 - \frac{\Pi^R}{\Pi^\star} \leqslant B_r$, and $B_a = B_r = 0$ when the optimal contract is restart.*

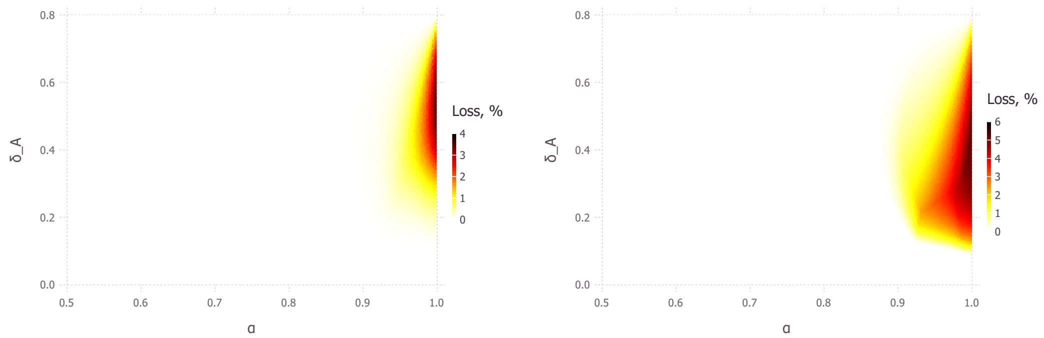

(a) actual loss (b) analytical loss

Figure 4: 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 $\Delta\theta = 0.1$.

One is an additive bound, and the other is a bound on the ratio. In the appendix we provide

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

[22] 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.

closed form expressions in terms of fundamentals. Figure 4 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 coincides with the first-order optimum. When the relaxed problem approach is not valid the analytical bound never exceeds 6 percent and the actual loss is never more than 4 percent.[23]

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

4 Simplicity through recursivity

In this section we characterize the optimum recursively, and document that when the optimal contract is not restart, the state space required to encode it is significantly richer, and thus the optimum is not "simple".

A recursive contract can be thought as an automaton which supplies capital advances to the agent conditional on an announcement of θ_H/θ_L. In such a scenario, one potential notion of simplicity is due to Abreu and Rubinstein [1988]; it counts the number of states or equivalently a number of distinct allocations supplied by the "machine".[24] Unfortunately, in our infinite horizon setting finite state machines are too restrictive, a prospective alternative notion of simplicity is to let the set of allocations $\bigcup_{t=1}^{\infty}\{k_t(\theta^t) : \text{for some } \theta^{\mathbf{t}}\}$ be countable. However, this notion of simplicity is too permissive, specifically it allows the cardinality of $\bigcup_{t=1}^{T}\{k_t(\theta^t) : \text{for some } \theta^{\mathbf{t}}\}$, to grow exponentially with T. We use a more demanding notion that 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 $\forall\, T$:*

$$\frac{1}{T}\left|\bigcup_{t=1}^{T}\left\{k_t(\theta^t) : \textit{for some } \theta^{\mathbf{t}}\right\}\right| \leqslant C.$$

When a contract is not simple, it is termed *complex*. Clearly, any restart contract is simple. We show that the optimal contract is simple if and only if the optimum is restart.

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

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

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

In what follows we briefly describe the recursive problem, which allows us to operationalize the notion of simplicity.. It is easy to show that IR_L will always bind for the optimal contract, hence, $U_t^\star(\theta^{t-1}, \theta_L) = 0$ at all dates. Thus, even though agent's type follows a two state Markov process, a one dimensional state variable, viz. $U(\theta^{t-1}, \theta_H) = w \in \mathbb{R}_+$, will suffice to encode all the required history dependence. From the second period onwards, for a promised expected utility of w to the high type and last period type j, define the objective as follows:

$$(\mathcal{RP}) \qquad S_j(w) = \max_{(\mathbf{k},\mathbf{z})\in\mathbb{R}_+^4} \alpha_j[s(k_H,\theta_H) - (\delta_P - \delta_A)\alpha_H z_H + \delta_P S_H(z_H)]+$$
$$+ (1-\alpha_j)[s(k_L,\theta_L) - (\delta_P - \delta_A)\alpha_L z_L + \delta_P S_L(z_L)] \quad \text{subject to}$$
$$\begin{cases} w \geqslant \Delta\theta R(k_L) + \delta_A(\alpha_H - \alpha_L)z_L \\ w \leqslant \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)z_H \end{cases}$$

The objective is to maximize the surplus when expected utility promised to the agent is fixed at $(w, 0)$ or $\alpha_j w + (1-\alpha_j)0$ in expectation. There are four choice variables: working capital advances $\mathbf{k} = (k_H, k_L)$ and expected utilities $\mathbf{z} = (z_H, z_L)$; note that z_i represents the utility promised to the high TFP 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 a continuation value of z_i. The two constraints are downward and upward incentive constraints, respectively. The participation constraint of θ_H type 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, $w = U(\theta_H) - U(\theta_L)$ must be chosen. The problem reads as follows

$$(\diamond) \quad \Pi^\star = \max_{(w,\mathbf{z},\mathbf{k})\in\mathbb{R}_+^5} -\mathbb{P}(\theta_H)w + \mathbb{P}(\theta_H)[s(k_H,\theta_H) - (\delta_P - \delta_A)\alpha_H z_H + \delta_P S_H(z_H)]+$$
$$+ \mathbb{P}(\theta_L)[s(k_L,\theta_L) - (\delta_P - \delta_A)\alpha_L z_L + \delta_P S_L(z_L)] \quad \text{subject to}$$
$$\begin{cases} w \geqslant \Delta\theta R(k_L) + \delta_A(\alpha_H - \alpha_L)z_L \\ w \leqslant \Delta\theta R(k_H) + \delta_A(\alpha_H - \alpha_L)z_H \end{cases}$$

Denote the optimal recursive contract by $\langle U_1^\star(\theta_H), \mathbf{k}(\cdot), \mathbf{z}(\cdot)\rangle$ where $(\mathbf{k}(w), \mathbf{z}(w))$ solves $(\mathcal{RP})$ for each $w \geqslant 0$. As is standard, the recursive optimal contract can be used to generate the optimum as follows:

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

We establish that the allocations $\mathbf{k}(\cdot)$ are monotone, thus complexity of the optimum is completely determined by richness of the state space used to encode it. The formal details are provided in the appendix, where we first completely characterize the global optimum in its recursive form, and then show that the optimal contract converges to its invariant (or stationary) distribution in finite time. Thus, in order evaluate the simplicity of the optimum, we only need to explore whether the set of allocations (or promised utilities) in the support of the stationary distribution satisfy simplicity.

Next, we first briefly define the stationary distribution of promised utilities and then the idea of simplicity associated with it.

Combined with the process of TFP shocks, the recursive contract induces a Markov process over $\Theta \times \mathbb{R}_+$. Define the joint probability of the event that (i) the expected utility promised to the agent lies in some Borel measurable set $A \subseteq \mathbb{R}_+$ and (ii) type realized today is θ_i; given that the current expected utility and last period's shock are w and θ_j to be:

$$F_{i|j}(A|w) = \mathbb{1}(z_i(w) \in A)\mathbb{P}(\theta_i|\theta_j).$$

By standard arguments F can be shown to have a unique invariant distribution (see Theorem 12.12 of Stokey et al. [1989]). Denote by $supp(F) \subset \mathbb{R}_+$ the projection of support of this unique invariant distribution onto the space of promised utilities.

The set $supp(F^*)$ is a strict subset of the recursive domain, and its cardinality captures the amount of information needed to describe the optimal contract. When the optimal contract is restart, $supp(F^*)$ is relatively small and hence the set of the allocations attained in the stationary distribution satisfies Definition 3. When then optimal contract is not restart, $supp(F^*)$ is exponentially large, violating our notion of simplicity. This captures the spirit of Theorem 3.

5 Comparative Statics

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

5.1 A "folk theorem"

Let $\beta = {}^{\delta_A}/_{\delta_P}$, and define the average (ex-ante) profit of the principal and the average payoff of the agent at any time t to be as follows:

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

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

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

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

Corollary 5. $\Pi^* \to s^e$ *if and only if* $\delta_P\beta \to 1$.

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

In the second case, where $\delta_P \beta < 1$, that, is at least of the player's discount factor is bounded below unity, the total surplus is bounded away from efficiency. Here either the intertemporal costs of incentive provision are forever positive ($\beta < 1$) or the agent's rents do not vanish and the principal distorts allocations along the lowest history ($\delta_P < 1$).

5.2 Patient versus impatient agent

Does the principal favor the impatient agent or the patient agent and what determines the ranking if there exists any? Recollect, that the principal's cost of providing incentives is given by $\bar{U}_P = \bar{U}_A + I$; for a fixed allocation, $\bar{U}_A$ is increasing in δ_A and I is decreasing in δ_A. The aggregate effect therefore depends on the level of asymmetric information in the model as measured by the persistence of the agent's type. Now, $\bar{U}_A$ is increasing in the persistence of the agent's types, and I is not monotonic in persistence. The principal's benefit depends on the agent's discount factor through of channel of information rent that has to be paid in the future, which in turn can be loaned to the agent so the principal can demand it back with an interest without violating the limited commitment requirement.

The complexity of these competing forces does not allow for a global comparative static, but a theoretical result can be stated for the limit cases and numerical arguments explored for the intermediate ones.

Corollary 6. *Let $\alpha_H = 1 - \alpha_L = \alpha$. Principal's ex ante payoff in the first-order optimal, optimal and optimal restart contracts varies with δ_A as follows:*

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

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

Figure 5 plots principal's profit in the first-order optimal contract and the optimal restart contract. It presents a "heat map" where each point in the box represents the expected profit of the principal as a function of α (on the x-axis) and δ_A (on the y-axis), wherein darker shades mean higher values. The northwest and southeast corners of the parametric spaces correspond to cases (a) and (b) of Corollary 6. In the intermediate range it is clear that the for each value of α the principal's profit changes non-linearly as a function of δ_A. For example at $\alpha = 0.9$, the principal prefers either a completely myopic agent ($\delta_A = 0$) or completely forward looking one ($\delta_A = \delta_P$), but not goldilocks.

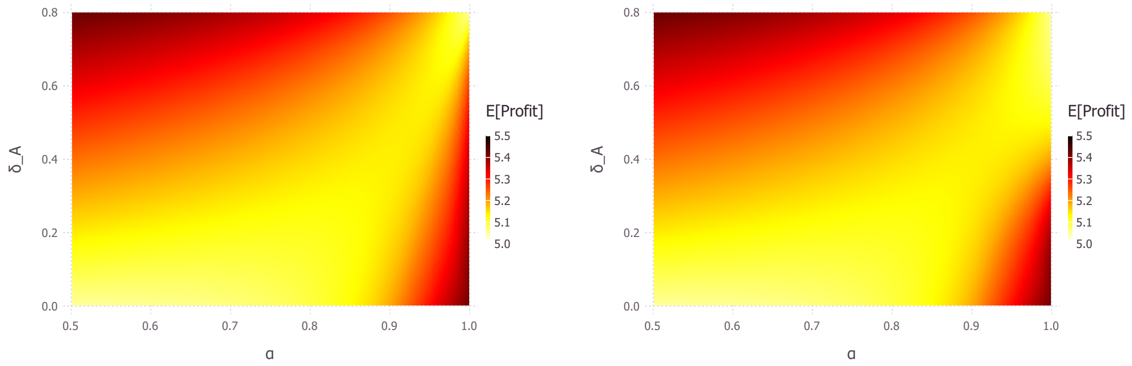

(a) First-order optimal contract

(b) Optimal restart contract

Figure 5: 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 $\Delta\theta = 0.1$.

At low levels of persistence the standard information rent the principal has to pay is quite low, she extracts a large part of the surplus as profit, and does not find it worthwhile to pay the extra intertemporal cost of incentive provision to benefit from differential interest rates. Hence profit increasing in δ_A. As persistence increases, on the cost side, the standard information rent goes up which a lower δ_A helps reduce, but a lower δ_A also increases the intertemporal cost of incentive provision. These competing forces work to cancel each other until persistence becomes very high, then the increased standard information rent dominates all other effects and the principal strictly prefers a myopic agent to minimize the total information rent.

6 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 novel tradeoff: the principal has to incur intertemporal costs of incentive provision, but benefits from having a higher net present value of total surplus and lower value of the standard information rent. These forces interact to to 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 the history dependence. In the latter case, we characterize the optimal restart contract which provides a simpler and approximately optimal alternative.

The nature of dynamic distortions poses a question to the literature on dynamic (Myersonian) mechanism design– a slight perturbation of the standard model of equal discounting renders long-term efficiency unachievable, distortions are pervasive. With equal discounting, Besanko [1985] and Battaglini [2005] show that ex post distortions converge to zero in the long run for the AR(1) and two type Markov models respectively. Garrett, Pavan, and Toikka [2018] show that distortions converge to zero on average for more general types' processes.[25] 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 opposed 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 demands all the information rent upfront and relaxes future incentive constraints. In the presence of hard financial constraints in the form of limited liability the principal binds the limited liability constraints for as long as information rent to be paid out to the agent is recouped, and then eventually implements the efficient contract (see Krishna et al. [2013] and Krasikov and Lamba [2018]). However, a permanent difference in "access to capital" creates a permanent cost in generating the requisite room to relax future incentive constraints, which culminates in cyclical and non-vanishing distortions.

Unequal discounting has been explored to varying degrees in dynamic games and contracts. It is well known that in repeated games with differential rate of time preference the set of equilibrium payoffs expands favoring the patient player (see the classic Lehrer and Pauzner [1999]). Opp and Zhu [2015] analyze the general relational contracting model of Ray [2002] with unequal discounting. There is no private information or unobservable actions. It is a two-sided limited commitment problem, and incentive constraints therein are the equivalent of punishment phase in repeated games, a resort to autarky on deviation from the prescribed plan. The threat of autarky generates backloading of payments and unequal discounting does the frontloading, leading to a cyclical pattern similar to our paper.

In another exposition of the implication of differential discounting, Krueger and Uhlig [2006] study a risk sharing model with a risk averse agent and competing risk neutral principals. For equal discounting the model generates full risk-sharing in the long-run, for moderate differences in discounting partial insurance is the outcome, and when difference in discount factors is very large autarky results are obtained. These results are analogous to our results on long-term efficiency for equal discounting, pervasive distortions for unequal discounting, and shutdowns for large difference in discount rates and large agency frictions. The important difference is that the underlying feature of their model is risk sharing whereas in our setting it is Markovian private information.

Biais et al. [2007] incorporate unequal discounting in a dynamic model of moral hazard with an i.i.d technology, limited liability constraints, and the possibility of liquidation. There exists a reflective boundary below the efficient level that pushes the optimal contract back towards the liquidation region, and the contract is liquidated almost surely in the long-run. The propagation of distortions is sustained in our model through persistence in agency frictions whereas the same is done in their framework by limited liability and the threat of liquidation.[26]

Our paper is also related to the political economy and public finance literature that uses unequal

[25] See also Bergemann and Strack [2015] for the evolution of dynamic distortions in the continuous time setting.

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

discounting as a motivation for long-run distortions. Acemoglu, Golosov, and Tsyvinski [2008] show that when politicians are less patient than the citizens, positive aggregate labor and capital taxes are charged forever to correct for political economy distortions. Farhi and Werning [2007] find that in an Atkeson and Lucas [1992] style risk sharing model with taste shocks, when the social discount factor is higher than the private one, consumption exhibits mean reversion with no immiseration.[27] While the former contains the long-run inefficiency flavor of our results, the latter shows dynamics similar to the optimality of restart.

A limitation of our model is the 'permanency' of the differential interest rates. A more detailed analysis would allow the agent to save his way towards the market rate. There are many plausible ways on introducing this added dimension to our model. One tractable way could perhaps be to allow the discount factor of the agent to depend on the level of equity of the "firm structure".[28] 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 but still 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 also 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 the model as stated is then not "compact"; the lack of an upper bound on transfers that the principal can pay means that the agent will lend the principal an unbounded amount of money in a hope to claw it back in the future. 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.

7 Appendix

7.1 Sequential characterization

First, we establish the set of binding constraints in problems $(\star)$, $(\#)$ and (R). Lemma 1 shows that IR_L binds in all three problems, Lemma 2 states that IC_H binds in the relaxed problem.

Lemma 1. *Let $\langle \mathbf{k}, \mathbf{U} \rangle$ satisfy IC_H and IR_L, then there exists another mechanism $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ which respects IC_H and IR_L, moreover, it yields a weakly higher ex-ante profit. In addition, if $\langle \mathbf{k}, \mathbf{U} \rangle$ satisfies IC_L and IR_H, then $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ can be chosen to respect these constraints as well.*

Proof. Define $\tilde{U}_t(\theta^{t-1}, \theta_H) = 0$ and $\tilde{U}_t(\theta^{t-1}, \theta_H) = U_t(\theta^{t-1}, \theta_H) - U_t(\theta^{t-1}, \theta_L) \geqslant 0$, because the original contract respects IC_H. Note that the incentive compatibility constraints are unaltered and $\tilde{\mathbf{U}} \leqslant \mathbf{U}$, thus $\tilde{U}_P \leqslant U_P$. Conclude that the profit of $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ is higher than of $\langle \mathbf{k}, \mathbf{U} \rangle$. □

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

[28] 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].

Lemma 2. *Let $\langle \mathbf{k}, \mathbf{U} \rangle$ satisfy IC_H and IR_L, then there exists another mechanism $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ in which both constraints hold as equalities, moreover, it yields a weakly higher ex-ante profit.*

Proof. By Lemma 1, it is without loss to assume that $U_t(\theta^{\mathbf{t}-1}, \theta_L) = 0$ for all $\theta^{\mathbf{t}-1}$. Define $\tilde{\mathbf{U}}$ by $\tilde{U}_t(\theta^{\mathbf{t}-1}, \theta_L) = 0$ and $\tilde{U}_t(\theta^{\mathbf{t}-1}, \theta_H) = \sum_{s=1}^{\infty} (\delta_A(\alpha_H - \alpha_L))^{s-1} \Delta\theta R \left(k_{t-1+s}(\theta^{\mathbf{t}-1}, \theta_{\mathbf{L}}^{\mathbf{s}})\right)$. Note that IC_H binds and $\tilde{\mathbf{U}} \leqslant \mathbf{U}$, thus $\tilde{U}_P \leqslant U_P$. Conclude that the profit of $\langle \mathbf{k}, \tilde{\mathbf{U}} \rangle$ is higher than of $\langle \mathbf{k}, \mathbf{U} \rangle$. □

7.1.1 Relaxed problem approach

Now, we are in position to complete the proof of Theorem 1 and Corollary 1.

Proof of Theorem 1. The goal is to obtain Equation 2 and derive the distortions $\hat{\rho}$ and ρ as described in the theorem. By Lemmata 1 and 2 both IC_H and IR_L bind in (#), thus $U_t(\theta^{\mathbf{t}-1}, \theta_H) = 0$, and

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

First, we solve for the agent's ex ante utility:

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

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

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

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

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

Evaluate the sum for the lowest history and restart phase separately, starting with the lowest history:

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

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

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

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

Finally, we look at the restart phase and compute the costs of incentive provision:

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

where $a_H = \frac{\delta_P - \delta_A}{\delta_P} \frac{\alpha_H}{1-\alpha_H}$ is the H-seed. Conclude that total distortions in the restart phase take the following form:

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

□

Proof of Corollary 1. Consider $f(x) = bx + a_L$ 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 distortions satisfy $\rho_{t+1} = f(\rho_t)$ and $\hat{\rho}_{t+1} = f(\hat{\rho}_t)$ for all t.

It is easy to see that f has only one 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 the fixed point is globally stable and the distortions converge to it monotonically. It is easy to see that $\frac{a_L}{1-b} < \frac{\alpha_L}{1-\alpha_H} < a_H$ which immediately implies *(a)* and *(b)*. To see *(c)* and *(d)*, recall the definition of $\mathcal{K}_L(x) = (R')^{-1}\left(\frac{1}{\theta_L - x\Delta\theta}\right)$ for $x\Delta\theta < \theta_L$ and zero otherwise. □

7.1.2 Restart optimum

In this section we characterize the restart optimum (Theorem 2) and derive its profit guarantee (Corollary 4). By Lemma 1, there is no loss to assume that IR_L always binds, $U_t(\theta^{\mathbf{t-1}}, \theta_L) = 0$ for any $\theta^{\mathbf{t-1}}$. Our restrictions on the contract space imply that the agent's expected utilities are pinned down by binding downward incentive constraints, moreover, they also feature restarts. More formally, there exists two sequences $\{U_t\}$ and $\{\hat{U}_t\}$ such that

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

These sequences are determined as a function of allocation in the following manner: $\hat{U}_t = \Delta\theta R(\hat{k}_t) + \delta_A(\alpha_H - \alpha_L)\hat{U}_{t+1}$ and $U_t = \Delta\theta R(k_t) + \delta_A(\alpha_H - \alpha_L)U_{t+1}$. It follows that IC_L is equivalent to

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

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

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

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

The term $\frac{\mathbb{P}(\theta_H)}{1-\delta_P}$ is a discounted probability of θ_H, that is $\sum_{t=1}^{\infty} \delta_P^{t-1} \mathbb{P}(\theta_t = \theta_H)$.

Next, notice that the agent's expected payoff is simply $\mathbb{E}[U_1(\theta_1)] = \mathbb{P}(\theta_H)\hat{U}_1$, and the costs of incentive provision are given by

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

As before, the former term captures the costs along the lowest history, whereas the latter - the restart phase.

To sum up, problem (R) can be equivalently written as

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

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

Now, we are in position to prove Theorem 2 and derive the bound described in Corollary 4, see Figure 4b for a visualization.

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

First order conditions for the allocation rule yield $\hat{k}_t = \mathcal{K}_L(\hat{\gamma}_t)$, $k_t = \mathcal{K}_L(\gamma_t)$ $\forall t$ and $k_H = \mathcal{K}_H(\kappa) \geqslant k_H^e$ where

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

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

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

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

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

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

It is routine to verify that $\hat{\eta}$ and η are both non-negative and continuous in $(\bar{\gamma}, \mu_1)$ on a relevant domain. By construction, coefficients in the Lagrangian in front of $\{\hat{U}_t\}$ and $\{U_{t+1}\}$ are identically zero. In addition, a coefficient in front of U_1 is proportional to

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

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

To conclude the proof, we only need to show that γ_1 can be chosen to satisfy complimentary slackness. The only non-trivial case is when the first-order optimum is not restart, otherwise $\gamma_1 = a_H$. Since distortions are monotone and stays at the same value once upward incentive compatibility starts to bind, it is sufficient to only verify complimentary slackness "at infinity":

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

This condition holds as ">" inequality whenever $\gamma_1 = a_H$ provided that the first order optimum is not restart. On other hand, this condition holds as "<" whenever $\gamma_1 = \frac{a_L}{1-b}$. To see it more formally, let $\gamma_1 = \frac{a_L}{1-b}$, then we must have $\bar{\gamma} > \min\left\{\frac{\alpha_L}{1-\alpha_H}, \frac{a_L}{1-b}\right\} = \frac{a_L}{1-b}$, because $\mu_1 < a_H$. Continuity Then implies that there exists some $\gamma_1 \in \left(\frac{a_L}{1-b}, a_H\right)$ for which complimentary slackness is satisfied. □

Proof of Corollary 4. Define slack variables for upward incentive constraints by $\epsilon_t = \left(\hat{U}_t^\# - \Delta R(k^e(\theta_H)) - \delta_A(\alpha_H - \alpha_L)U_1^\#\right)^+$ and $\hat{\epsilon}_t = \left(\hat{U}_t^\# - \Delta R(k^e(\theta_H)) - \delta_A(\alpha_H - \alpha_L)U_1^\#\right)^+$.

By the standard perturbation argument,

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

Our goal is to evaluate the left hand side of this expression in two different ways.

First, recall that distortions are monotone, thus $\hat{\epsilon}_t, \epsilon_t \leqslant \lim_{t\to\infty} \epsilon_t$ for all t. Using the first-order condition for U_1 and $\frac{a_L}{1-b} \leqslant \gamma_1$:

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

Second, we bound η_t and $\hat{\eta}_t$. It is easy to see that $\bar{\gamma} \leqslant \gamma_1$, thus $\gamma_{t+1} - a_L - b\gamma_t = \eta_{t+1} \leqslant$

$\bar{\gamma}(1-b) - a_L \leq (1-b)\left(a_H - \frac{a_L}{1-b}\right)$ and the same is true for $\hat{\eta}_{t+1}$ for any $t \geqslant 2$. For $t = 1$, $\eta_1 = 0$ and $\hat{\eta}_1 \leqslant \left(a_H - \frac{\alpha_L}{1-\alpha_H}\right)^+$, thus

$$\Pi^\# - \Pi^R \leqslant \mathbb{P}(\theta_L)\left(a_H - \frac{\alpha_L}{1-\alpha_H}\right)^+ \hat{\epsilon}_1 + \\ + (1-b)\left(a_H - \frac{a_L}{1-b}\right)\sum_{t=2}^{\infty}(\delta_P(1-\alpha_L))^{t-1}\left[\mathbb{P}(\theta_L)\hat{\epsilon}_t + \frac{\mathbb{P}(\theta_H)}{1-\delta_P}\delta(1-\alpha_H)\epsilon_t\right] =: B_a^2$$

To make sure that a relative loss does not explode, we also compute the loss from using the optimal static contract which a restart contract with $\hat{k}_t = k_t\ \forall t$. It is easy to show that the optimal static contract supplies the efficient quantity to the high type and $K_L(x)$ 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, it has a clear closed form representation, then we have $\Pi^\# - \Pi^R \leq \Pi^\# - \Pi^S$.

Taking all parts together, we arrive at the following analytical bounds:

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

□

7.2 Recursive characterization

In this section we study the recursive problem introduced in the main text; the formulation follows in other recursive characterizations in dynamic contracts, for example Fernandes and Phelan [2000]. By Lemma 1, it is without loss to assume that IR_L binds at all dates, thus a one-dimensional state suffices.[29]

Let W be the largest set of w such that there exists an incentive compatible and individually rational contract which delivers $U_1(\theta_H) = w$ and $U_1(\theta_L) = 0$. W is the familiar recursive domain described in Spear and Srivastava [1987] and it has a very simple structure.

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

Proof. First of all, $w \geqslant 0$ by IR_H. On the other hand, any $w \geqslant 0$ can be provided by choosing $k_t(\theta^t) = U_t(\theta^t) = 0\ \forall \theta^t$, but $U_1(\theta_1) = w$ and $k_1(\theta_H) = R^{-1}\left(\frac{w}{\Delta\theta}\right)$. □

Using Lemma 3, we can write the recursive problem as $(\mathcal{RP})$ from the second period onwards, and as $(\diamond)$ in the first period, explicitly stated in Section 4.

It is easy to see that $(\star)$ and $(\diamond)$ admit the same solution. To formally show equivalence between the sequential and recursive formulations, we need to introduce auxiliary definitions. The policy correspondence is a correspondence which maps w into $(\mathbf{Z}(w), \mathbf{K}(w))$, that is the set of optimal choices in $(\mathcal{RP})$. We say that a contract is generated from the policy correspondence

[29] Technically, the state space also includes the last period type, thus it is two-dimensional.

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 $\forall \theta^{t-1}$.

Claim 1.

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

(b) The policy correspondence is non-empty, compact-valued and upper hemicontinuous.

(c) A contract is generated from the policy correspondence if and only if it solves $(\mathcal{RP})$ *with* $w = U(\theta_H|\theta_j)$ *for* $j = H, L$*.*

(d) Value functions in $(\mathcal{SP})$ *and* $(\mathcal{RP})$*, as well as in* $(\star)$ *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 1, 2.

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

(a) Each S_j *is concave.*

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

(c) Each S_j *is locally strictly concave at* 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. Indeed, the constraints set is convex and $s(\theta, \cdot)$ is concave. So, concavity is preserved by the Bellman operator. Since the set of concave functions is closed in the space of continuous bounded functions, 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 S_j is concave, thus it is subdifferentiable. Take $\langle \mathbf{k}^\star, \mathbf{U}^\star \rangle$ which solves $(\mathcal{SP})$ 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 some finite time s such that the value function is differentiable at $U_{s+1}^\star(\theta_j, \theta_{\mathrm{L}}^{s-1}, \theta_H)$. 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 a solution. In particular, we show that Slater's condition holds which is sufficient to guarantee that the first-order approach with Lagrange multipliers in l^1 is valid in $(\mathcal{SP})$, because of concavity and boundedness of these problems (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 exists $k_H > k_L > 0$ satisfying:

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

Take $k_{t+1}(\theta_j, \theta^{\mathbf{t-1}}, \theta_H) = k_H$, $k_{t+1}(\theta_j, \theta^{\mathbf{t-1}}, \theta_L) = k_L$ and $U_{t+1}(\theta_j, \theta^{\mathbf{t-1}}, \theta_H) = w\ \forall\theta^{\mathbf{t-1}}$.

Now, we are in a position to show that S_j is continuously differentiable. Let $\langle \mathbf{k}^\star, \mathbf{U}^\star\rangle$ be a solution to $(\mathcal{SP})$ at $t = 2$. It is clear that the capital supplied to θ_H can be distorted only upwards, thus $k^\star_{t+1}(\theta_j, \theta^{\mathbf{t-1}}, \theta_H) > 0$ is uniquely defined $\forall\theta^{\mathbf{t-1}}$ by strict concavity of the objective. In addition, if $k^\star_{t+1}(\theta_j, \theta^{\mathbf{t-1}}, \theta_L) > 0$, then it is unique by strict concavity of the objective.

Next, consider $(\mathcal{RP})$, its solution exists and coincides with one found in $(\mathcal{SP})$. Since S_j is concave, its superdifferential at $w > 0$ is well-defined and it 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.

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 a solution, whereas $\rho_j^-(w)/\rho_L^+(w)$ be the highest/smallest such Lagrange multiplier. Finally, denote by $(\mathbf{z}(w), \mathbf{k}(w))$ 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, $\mathbf{K}_H(w)$ is a singleton and $\rho_H^+(w) = \rho_H^-(w) = \rho_H(w)$ for any w. In addition, if $k_L(w) > 0$, then $\mathbf{K}_L(w)$ is a singleton and $\rho_L^+(w) = \rho_L^-(w) = \rho_L(w)$. So, for $w > 0$, the Lagrange multipliers might be not unique only if there exists some $\rho_L(w) \geqslant \theta_L/\Delta\theta > 0$. Given this $\rho_L(w) > 0$, the downward incentive constraint binds and we have that $z_L(w) = \frac{w}{\delta_A(\alpha_H-\alpha_L)} > w > 0$ is uniquely defined.

Then, the envelope theorem gives $S_j^-(w) - S_j^+(w) = (1-\alpha_j)(\rho_L^-(w) - \rho_L^+(w))$. It is immediate that S_j is differentiable at w if and only if $\rho_L(w)$ is unique. 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^∞.

Finally, continuous differentiability of S_j is implied by differentiability and concavity.

Part (c). 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 $(\mathcal{SP})$ at w and $w+\varepsilon$, 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 each case. By the first-order and envelope conditions, see Equations 3, 4 and 5, it will continue to bind along the sequence of θ_L's, thus

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

The last assertion is a clear contradiction. The similar argument establishes that $S_j'(w - \epsilon) > S_j'(w)$. □

Next, we derive optimality conditions which are useful for our characterization of the optimal contract. Let $(1-\alpha_j)\rho_H$ and $\alpha_j \rho_L$ be Lagrange multipliers on the constraints in $(\mathcal{RP})$. And, let $\mathbb{P}(\theta_H)\rho_H$ and $\mathbb{P}(\theta_L)\rho_L$ be Lagrange multipliers on the constraints in $(\diamond)$. We denote by $(\mathbf{z}(w), \mathbf{k}(w))$ some selection from the optimal correspondence and by $\rho(w)$ some corresponding Lagrange multipliers. So, the first-order conditions are $k_i(w) = \mathcal{K}_i(\rho_i(w))$ for $i = H, L$ and

$$S_H'(z_H(w)) - \alpha_H \frac{\delta_P - \delta_A}{\delta_P} + (\alpha_H - \alpha_L)\frac{\delta_A}{\delta_P}\rho_H(w) \begin{cases} = 0 \text{ if } z_H(w) > 0 \\ \leqslant 0 \text{ if } z_H(w) = 0 \end{cases} \tag{3}$$

$$S_L'(z_L(w)) - \alpha_L \frac{\delta_P - \delta_A}{\delta_P} - (\alpha_H - \alpha_L)\frac{\delta_A}{\delta_P}\rho_L(w) \begin{cases} = 0 \text{ if } z_L(w) > 0 \\ \leqslant 0 \text{ if } z_L(w) = 0 \end{cases} \tag{4}$$

In addition, the Envelope theorem gives:

$$S_j'(w) = (1-\alpha_j)\rho_L(w) - \alpha_j \rho_H(w) \text{ for } j = H, L \tag{5}$$

We proceed by characterizing properties of the recursive optimum. Although, S_j might be not globally strictly concave, we are able to identify next period promised utilities when the incentive constraints do not bind. To be specific, $z_L(w) = z_L^e$ if the downward constraint is slack and $z_H(w) = z_H^e$ if the upward constraint is slack. By part (c) of Claim 1, there exists unique z_j^e satisfying $z_j^e > 0$ and $S_j'(z_j^e) = \alpha_j \frac{\delta_P - \delta_A}{\delta_P}$ or $z_j^e = 0$ and $S_j'(0) \leqslant \alpha_j \frac{\delta_P - \delta_A}{\delta_P}$. Then, define two thresholds $w_j^* = \Delta\theta R(k^e(\theta_j)) + \delta_A(\alpha_H - \alpha_L)z_j^e > 0$.

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

7.2.1 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 $(\diamond)$.[30] We

[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 only z_H could fail to be unique at a single point. The details are provided in Claim 3.

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 TFP 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 6a plots the optimal allocation as the function of agent's expected utility. We have the following simple result.

Proposition 1. *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 1. It suffices to characterize ρ, because its properties translate into $\mathbf{k}$ by the first-order condition $k_i(w) = \mathcal{K}_i(\rho(w))$ for $i = H, L$.

Part (a). If there is no upward incentive constraint, then $k_H(w) = k_H^e$ and $z_H = z_H^e$ by the first-order conditions and definition of z_H^e. Since this choice is feasible if and only if $w \geqslant w_H^*$, the result for ρ_H follows. To see monotonicity of $\rho_H(\cdot)$, take $w' > w \geqslant w_H^*$ and suppose that $\rho_H(w) \geqslant \rho_H(w')$. Concavity and first-order conditions imply that $z_H(w) \geqslant z_H(w')$ which contradicts to $\Delta\theta(R \circ \mathcal{K}_H)(\rho_H(w)) + \delta_A(\alpha_H - \alpha_L)z_H(w) = w < w' = \Delta\theta(R \circ \mathcal{K}_H)(\rho_H(w')) + \delta_A(\alpha_H - \alpha_L)z_H(w')$.

Part (b). By the similar argument to part (a), $\rho_L(\cdot)$ is strictly decreasing on $[0, w_L^*]$, and it is zero afterwards. Finally, since the only feasible choice at $w = 0$ is $k_L(0) = 0$, $w_k^o = \sup\{w \in W : k_L(w) = 0\}$ is well-defined. □

Now, we turn our attention to $\mathbf{z}$ and start by pointing out uniqueness of transfers.

Claim 3. $\mathbf{Z}_L$ *is single-valued, and* $\exists$ *unique* $\bar{w}$ *such that* $\mathbf{Z}_H$ *is single-valued whenever* $w_L^* \geqslant w_H^*$ *or* $w \neq \bar{w}$.

Proof. $z_L(\cdot)$ is unique which follows from the last part of Claim 2, whereas z_H might fail to be unique. Intuitively, z_H could be not unique only when there are multiple z_H with $\rho_L(z_H(w)) = \rho_H(z_H(w)) = 0$.

Define $\bar{w}$ by $(\alpha_H - \alpha_L)\delta_A \rho_H(\bar{w}) = \alpha_H(\delta_P - \delta_A)$. Clearly, it exists and it is unique, because of monotonicity of ρ_H as shown in the proof of Proposition 1.

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 $\mathbb{R}_+$. So, Z_H is single-valued by strict concavity of S_j.

If $w^*_L < w^*_H$, then the envelope conditions (Equation 5) 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 $w \in [w^*_L, w^*_H]$. Therefore, $\mathbf{Z}_H$ is single-valued on $[0, \bar{w})$ by the last part of Claim 2, and $\mathbf{Z}_H(\bar{w}) = [w^*_L, w^*_H]$ by construction. To see that $\mathbf{Z}_H$ is single-valued on $(\bar{w}, +\infty)$, notice 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, by $z_H(\cdot)$ we mean an arbitrary selection from $\mathbf{Z}_H(\cdot)$.

Now, the dynamics of promised expected utility are described in Figure 6. 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^e_H at which $z_H(\cdot)$ intersects the 45° line constitutes a bold dot. Each time a high shock arrives it is possible for the optimal contract to stay at the same expected utility, and it surely does so if the upward constraint is not binding.

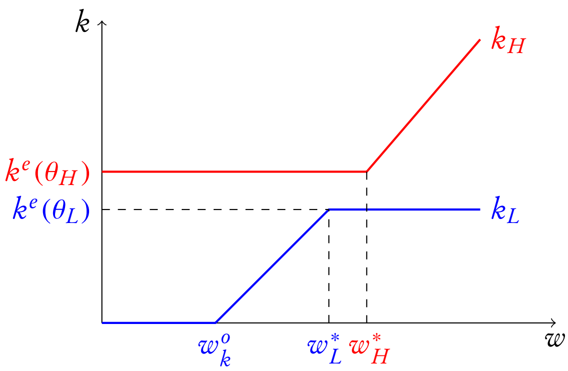

(a) allocation

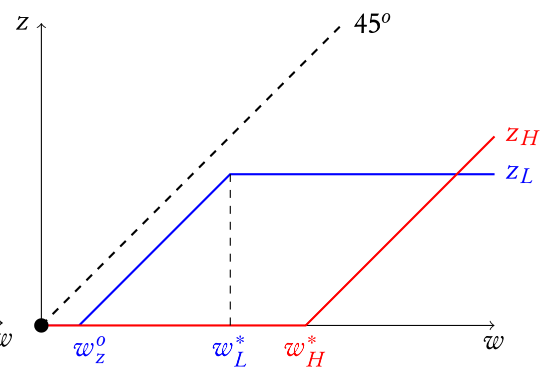

(b) shutdown

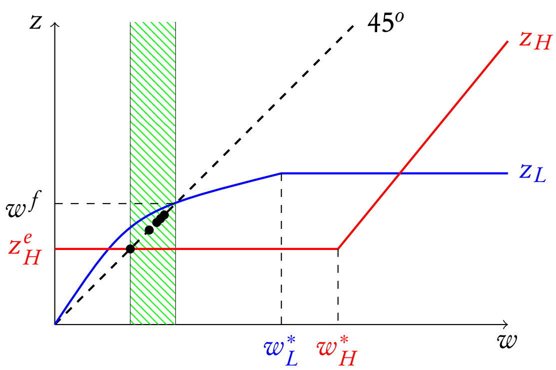

(c) restart

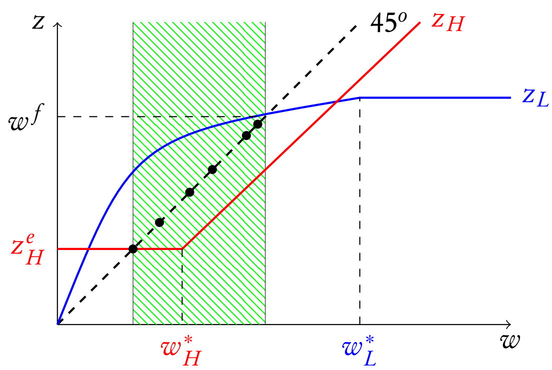

(d) never restart

Figure 6: Optimal recursive contract

Consider first the situation depicted in Figure 6b. Here $z^e_H = 0$. 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 TFP type. In Figures 6c and 6d, we portray the optimal restart contract which does not feature shutdowns. The realization of a high shock pushes the expected utility towards z^e_H. On the realization of a low shock, promised expected utility above w^f contracts, and below w^f it expands. The key condition that characterizes Figure 6c is $w^f \leq w^*_H$. It implies that the upward incentive constraint does not bind in the interval $[z^e_H, w^f]$, and the invariant distribution of the promised expected utility rests therein.[31] In contrast, Figure 6d 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 $w^\star \in [z^e_H, w^f]$. 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 2. *Expected utility of the agent in the optimal recursive contract satisfies the following:*

(a) $\exists w^o_z, z^e_L$ *such that* $z_L(w) = 0$ *if and only if* $w \leqslant w^o_z$, $z_L(w) = z^e_L$ *if and only if* $w \geqslant w^*_L$, $z_L(\cdot)$ *is strictly increasing on* $[w^o_z, w^*_L]$.

(b) $\exists z^e_H$ *such that* $z_H(w) = z^e_H$ *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* $w^f \in [z^e_H, z^*_L]$, *and* z_H *has a unique fixed point* z^e_H *which is positive if and only if* $\theta_L > \frac{a_L}{1-b}\Delta\theta$.

(d) The thresholds satisfy $z^e_H \leqslant w^f \leqslant z^e_L < w^*_L$, $z^e_H < w^*_H$, *and* $z^e_L \neq z^e_H$ *if and only if* $z^e_L > 0$.

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

Proof of Proposition 2. Part (d). Equation 5 says that $S'_H(w)/\alpha_H - S'_L(w)/\alpha_L = \frac{\alpha_L-\alpha_H}{\alpha_H\alpha_L}\rho_L(w) \leqslant 0$. Therefore, $z^e_H \leqslant z^e_L$ with $z^e_L \neq z^e_H$ if and only if $S'_L(0) > \alpha_L\frac{\delta_P-\delta_A}{\delta_P}$ by their definitions and part (c) of Claim 2. For $z^e_L = 0$, $w^*_L > z^e_L$ is trivially satisfied. Suppose that $z^e_L > 0$, then $S'_j(w^*_L) = -\alpha_j\rho_H(w^*_L) \leqslant 0 < S'_j(z^e_L)$, thus $w^*_L > z^e_L$.

Moreover, notice that $w^*_H = \Delta\theta R(k^e(\theta_H)) + \delta_A(\alpha_H - \alpha_L)z^e_H \leqslant z^e_H$ if and only if $z^e_H \geqslant \frac{\Delta\theta}{1-\delta_A(\alpha_H-\alpha_L)}R(k^e(\theta_H))$. On the other hand, $z^e_H < \frac{\Delta\theta}{1-\delta_A(\alpha_H-\alpha_L)}R(k^e(\theta_L))$, because of $z^e_H \leqslant z^e_L < w^*_L$. So, we can not have $z^e_H \geqslant w^*_H$.

It remains to establish that $z^e_H \leqslant w^f$. Of course, it is vacuously true whenever $z^e_H = 0$. So, suppose that $z^e_H > 0$. In this case, $z^e_H \leqslant w^f$ whenever $\frac{a_L}{1-b} \leqslant a_H$. To see this, notice that $\rho_L(w^f) \geqslant \frac{a_L}{1-b}$ with an equality if and only if $\rho_H(w^f) = 0$, as shown in part (c). Suppose that $z^e_H < w^f$ which is equivalent to $\rho_L(w^f) > \rho_L(z^e_H)$ by monotonicity of $\rho_L(\cdot)$. Since $z^e_H < w^*_H$, $\rho_H(w^f) = \frac{a_L}{1-b}$, which contradicts to $\rho_L(w^f) > \rho_L(z^e_H) > 0$.

Recall that $\frac{a_L}{1-b} \leqslant a_H$ if and only if $\frac{\alpha_L}{1-\alpha_L} \leqslant \frac{\alpha_H}{1-\alpha_H}\left(1 - \frac{\delta_A}{\delta_P}\frac{\alpha_H-\alpha_L}{1-\alpha_L}\right)$ which is always satisfied.

[31] To find the support, we repeatedly apply $z_L(\cdot)$ to z^e_H, the bold points in Figure 6c depict this set.

Parts (a) and (b). We established above that $z_j^e \in [0, w_j^*]$ for $j = H, L$. Monotonicity of $\rho(\cdot)$ as shown in Proposition 1 combined with Equations 3 and 4 yields the result of parts (a) and (b).

Part (c). First, we study fixed points of $Z_H(\cdot)$. In the previous part, we showed that $z_H^e < w_H^*$ which implies that z_H^e is a fixed point of $Z_H(\cdot)$. Suppose that there exists $w \neq z_H^e > 0$ with $w \in \mathbf{Z}_H(w)$. By definition, it must be the case that $\rho_H(w) > 0$.

Consider the equation $w = \frac{\Delta\theta}{1-\delta_A(\alpha_H-\alpha_L)}(R \circ \mathcal{K}_H)(\rho_H(w)) > \frac{\Delta\theta}{1-\delta_A(\alpha_H-\alpha_L)}R(k_H^e)$ which is necessary for $w \in \mathbf{Z}_H(w) > 0$ with $\rho_H(w) > 0$. Equation 3 and 5 imply that $(1-\alpha_H)\delta_P\rho_L(w) = \alpha_H(\delta_P - \delta_A) + (\alpha_H\delta_P - (\alpha_H - \alpha_L)\delta_A)\rho_H(w) > 0$.

Since $\rho_L(w) > 0$, the downward constraint binds this period and it will keep binding along the sequence of θ_L's. Formally, let $z_L^s(w)$ be defined by $z_L^s(w) = z_L(z_L^{s-1}(w))$ with $z_L^0(w) = w$. By Equation 4, $\rho(z_L^s(w)) > 0$ for any s. Then, iterating along this sequence, we arrive at the following contradiction:

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

So, z_H^e is the unique fixed point of $\mathbf{Z}_H(\cdot)$.

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. Suppose there exists $0 < w = z_L(w)$. 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 $\rho_L(w) > 0$.

Consider the equation $w = \frac{\Delta\theta}{1-\delta_A(\alpha_H-\alpha_L)}(R\circ\mathcal{K}_L)(\rho_L(w))$ which is necessary when $w = z_L(w) > 0$ with $\rho_L(w) > 0$. One more necessary condition, due to the Equations 4 and 5, is that $((1-\alpha_L)\delta_P - \delta_A(\alpha_H - \alpha_L))\rho_L(w) = \alpha_L(\delta_P - \delta_A) + \alpha_L\delta_P\rho_H(w) > 0$. By monotonicity of ρ (shown in Proposition 1), these two 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 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 (e). At the initial date, the first-order conditions with respect to $\mathbf{z}$ coincide with Equations 3 and 4. 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 $w > (=)0$. Existence and uniqueness directly follows from monotonicity of ρ, see proof of Proposition 1.

Next, we show $w^* \in [z_H^e, w^f]$. By the way of 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 ρ_H is non-decreasing, thus $\rho_H(z_H^e) \geqslant \rho_H(w^*) > 0$ which is a contradiction. Conclude that $w^* \geqslant z_H^e$.

Again, by the way of 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 ρ_H and ρ_L, $\rho_H(w^*) > \rho_H(w^f) > 0$ and $\rho_L(w^*) \leqslant \rho_L(w^f)$ where

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

It is routine to verify that all these conditions can not be satisfied simultaneously, thus $w^* \leqslant w^f$. □

Propositions 1 and 2 precisely characterize the optimal contract. Starting at w^*, each subsequent realization of the agent's type determines the optimal allocation according to Proposition 1 and the optimal expected utility for the next period, the state variable, according to Proposition 2.

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 for the upward incentive constraint to start binding.

7.2.2 Simplicity

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

Proof of Theorem 3. It is easy to see that any restart contract is simple, because a number of possible distinct allocations by time T is at most $2T$. To be concrete, the set $\bigcup_{t=1}^{T} \{k_t(\theta^{\mathrm{t}}) : \text{ for some } \theta^{\mathrm{t}}\}$ contains $\{\hat{k}_t\}_{t=1}^{T}$, $\{k_t\}_{t=1}^{T-1}$ and k_H

Suppose that the optimal contract is not restart. In term of our recursive notations, it means $z_H(w^f) \neq z_H^e$. According to Proposition 2, a necessary condition is that there are no shutdowns, that is $z_H^e > 0$, thus the high type promised utility stays strictly positive. Therefore, it is sufficient to show that the set of utilities promised to θ_H grows at an exponential rate, formally there exists a number K such that

$$\left|\bigcup_{t=1}^{T} \left\{U_t(\theta^{\mathrm{t-1}}, \theta_H) : \text{ for some } \theta^{\mathrm{t-1}}\right\}\right| \geqslant K2^T$$

First of all, notice 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 2 implies that for any $w, w' \in [z_H^e, w^f)$ with $w \neq w'$ we have $z_H(z_L^\tau(w)) \neq z_H(z_L^\tau(w')) \neq z_H^e$. In other words, the number of states is doubled every τ periods, thus we the state space expands exponentially with $K = 2^{-\tau}$. □

7.3 Comparative statics

Proof of Corollary 5. Note that the optimal contract is independent from the level of δ_P. Moreover, by the theorem of maximum, this contract is a continuous function of the agent's relative patience. If $\beta \to 1$, then the optimum convergence to the first-order optimum, because the latter is always incentive compatible when $\beta = 1$. This contract is described in 1 and it has distortions only along the lowest history, that is $\rho_t = 0\ \forall t$. Next, as $\delta_P \to 1$, non-stationary payoffs in Π^* vanish, thus the principal's achieves the maximal surplus.

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

□

Proof of Corollary 6. We start by looking at the first-order optimal contract. The first-order optimal contract is essentially static for $\alpha = \frac{1}{2}$, see Theorem 1. Formally, $\rho_t = \frac{\delta_P - \delta_A}{\delta_P}$ for any t, $\hat{\rho}_t = \frac{\delta_P - \delta_A}{\delta_P}$ for $t \geqslant 2$. Importantly, $\bar{U}_A$ is independent of δ_A, so $\delta_A = \delta_P$ uniquely maximizes the surplus and minimizes the cost of incentive provision at the same time. Since the profit in the first-order optimal contract is continuous with respect α and $\delta_A = \delta_P$ is a strict maximizer for $\alpha = 1$, it is still a maximizer for $\alpha \approx \frac{1}{2}$.

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

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

References

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–âĂŞ453, 1992.

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

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

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

D. Bergemann and P. Strack. Dynamic revenue maximization: A continuous time approach. *Journal of Economic Theory*, 159:819–853, 2015.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

E. F. Moore. Gedanken-experiments on sequential machines. In C. E. Shannon and J. McCarthy, editors, *Automata Studies, Annals of Mathematics Studies Volume 34*, page 129Ð153. Princeton University Press, 1956.

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

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

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

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

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

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

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

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

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

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