Firm Size and Compensation Dynamics with Risk Aversion and Persistent Private Information*

Gerard Maideu-Morera †

February 10, 2023

Abstract

I study a dynamic cash flow diversion model between a risk neutral lender and a *risk averse* entrepreneur that has *persistent private information* about the firm's productivity. Risk aversion decouples firm size and compensation dynamics. In the optimal contract, firm size drifts downwards, and the entrepreneur's compensation is smoothed, but the cross-sectional variance increases without bound. These results contrast equivalent models with risk neutrality, where firm size tends to increase over time, and the entrepreneur is compensated once the optimal size is reached. I use numerical simulations to study a (quasi-)implementation with simpler contracts. Risk aversion requires separately keeping track of the entrepreneur's wealth and equity share in the firm, which should decrease over time if shocks are persistent. The implementation provides an intuitive explanation for the opposite firm size dynamics.

JEL codes:

Keywords:

*I am grateful to Christian Hellwig and Nicolas Werquin for their advice and guidance. I have benefited very much from advice and discussions from Charles Brendon and Fabrice Collard. I would also like to thank Martin Beraja, Jonas Gathen, Eugenia Gonzalez-Aguado, Johannes Hörner, David Martimort, Alessandro Pavan, Jean-Charles Rochet, Andreas Schaab, Stéphane Villeneuve and seminar participants at the TSE Macro Workshop, EEA, the Naples PhD workshop and the European Winter Meeting of the Econometric Society.

†Toulouse School of Economics. Email: gerard.maideumorera@tse-fr.eu

1 Introduction

Financing constraints slow down firms' growth over their lifecycle. Dynamic contracting models have proved useful in understanding the underlying agency frictions that generate financing constraints and prevent, particularly young firms, from operating at their optimal size. The canonical setting in this literature is the cash flow diversion model: at each period, an entrepreneur needs funds from a lender to operate a project, but only the entrepreneur observes the project's cash flows and can secretly divert them for consumption. A regular outcome of this class of models is that, in the optimal contract, the firm size drifts upwards, and the entrepreneur is compensated once the undistorted first best size is reached (Clementi and Hopenhayn (2006))[1].

The literature has typically assumed that the entrepreneur is risk neutral and that the shocks to the firm's cash flows are i.i.d[2] (Clementi and Hopenhayn (2006), DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007a), Biais *et al.* (2007)). However, by making these two assumptions, these models may abstract from first-order concerns for designing financial contracts. First, with risk neutrality, there is no need to smooth the entrepreneur's consumption, so the compensation can be backloaded at little (or no) cost. Second, the i.i.d assumption restricts the extent to which the entrepreneur may have more information about the firm's future profitability. Thus limiting the potential gains of misreporting and the extent to which the entrepreneur's preferences for future contract arrangements depend on the current productivity.

In this paper, I study a dynamic contracting problem between a risk neutral lender and a *risk averse* entrepreneur that has *persistent private information* about the firm's productivity. I solve for the optimal contract and analyze the implied firm size and compensation dynamics. Together, risk aversion and persistence lead to remarkably different dynamics than models previously studied. Firm size (i.e. working capital invested) drifts downwards. The entrepreneur's compensation is smoothed, but the variance of consumption is permanently increasing (as in Thomas and Worrall (1990)). Moreover, compared to the risk neutral case, where compensation and firm size are linked one to one (Clementi and Hopenhayn (2006)), with risk aversion, there is an almost complete separation between these two variables. For

[1] Empirically, it is not obvious that the firms' financing constraints are eventually relaxed as they age. For instance, in developing economies, where financing constraints are more stringent, we observe that old firms are relatively smaller than in developed economies (Hsieh and Klenow (2014)).

[2] A notable exception is Fu and Krishna (2019), who study a similar cash-flow diversion model with risk neutrality but with persistent shocks. However, as I will show, the role of persistence on firm dynamics crucially depends on the entrepreneur's risk aversion.

instance, after several periods, the size of the firm can be distorted downwards, but the entrepreneur receives a high compensation. These dynamics are shown theoretically but also illustrated with numerical simulations.

With private information, the firm dynamics can be characterized with return-dependent investment wedges, which lower the implicit expected marginal product of capital. The size of the investment wedges depends on three terms: (i) the elasticity of the marginal product of capital with respect to productivity, intuitively, lending is more costly when it benefits relative more the productive entrepreneurs (it increases relatively more their ability to divert funds); (ii) the upper Pareto tail of the distribution of productivity shocks; and (iii) a normalized shadow cost of insurance which captures the amount of insurance that the lender wants (or has promised) to provide to the entrepreneur. With risk aversion and persistent private information, the lender gains by promising to provide more insurance in future periods as it helps screen types[3]. Consequently, the investment wedges tend to increase over time, and firm size tends to decrease[4]. However, with i.i.d productivity shocks, there is no gain of promising higher future insurance, so wedges and firm size are stationary.

The entrepreneur's consumption process satisfies a Generalized Inverse Euler Equation (GIEE) similar to Hellwig (2021). As expected, with risk aversion, the lender smooths the entrepreneur's compensation. In the GIEE, a savings wedge captures the incentive costs of savings at period t. Higher savings lower information rents at $t+1$, and the pass-through as an incentive cost at t is increasing on the persistence of the process. Hence, higher persistence lowers the savings wedge, so savings are less discouraged on the margin. The cross-sectional variance of the entrepreneur's consumption grows over time without bound. After a history of high (low) productivity shocks, the entrepreneur is rewarded with high (low) consumption. However, the numerical simulations show that this dispersion in compensation does not translate to firm size distortions. The investment wedges are essentially uncorrelated with compensation. Basically, the entrepreneur is compensated with higher future consumption, not with low future firm size distortions. Firm size distortions are instead driven by the gains of promising high future insurance. This finding also contrasts the model with risk

[3]More concretely, with risk aversion and persistent private information, different types θ_t have different preferences for the expected variation in utilities (i.e. the insurance) of contracts offered at $t+1$. In particular, types $\theta' > \theta_t$ prefer contracts where there is less insurance at period $t+1$ than θ_t. Therefore, the principal can commit to increase the insurance provided at $t+1$ to reduce the cost of screening types at t. This is the same reason why the labor wedges tend to increase over time in dynamic Mirrlees models (see Farhi and Werning (2013) and Makris and Pavan (2020)).

[4]Interestingly, Clementi *et al.* (2010) also find that the optimal contract implies that firm size tends to decrease over time. But the setting and rationale for firm decline are different. I discuss in more detail the differences between the two models in the literature review section.

neutrality, where, with i.i.d types, there is a one-to-one mapping between promised utility and firm size distortions (Clementi and Hopenhayn (2006)).

To further understand the compensation dynamics, I use numerical simulations and analyze a (quasi-)implementation with simpler contracts. The implementation is also useful for understanding the drivers of the different firm size dynamics with risk neutrality and risk aversion. With i.i.d shocks, the following simple contract gets very close to the optimal allocation[5]. The lender gives the entrepreneur a constant equity share on the firm's reported cash flows. Then the entrepreneur can pledge her shares as collateral and borrow to smooth consumption given his implied wealth. Pledging shares is a common practice (Fabisik (2019)); this implementation shows how it can be rationalized as part of a nearly optimal contract[6].

With persistent private information, the principal's problem contains an extra state variable that captures the insurance promised to the agent. This state variable naturally maps to the equity share given to the entrepreneur. Intuitively, types $\theta'' > \theta'$ know they are expected to obtain higher cash flows at $t+1$ than θ', so it is less attractive for them to give up equity. Therefore, when the lender buys equity at $t+1$ to some type θ', it discourages the diversion of funds for types $\theta'' > \theta'$. That is, the lender optimal lowers the equity share at period $t+1$ (to an inefficient level once at $t+1$) because it helps screen types at t. In any case, in the numerical simulations, the contract with a constant equity share, augmented to account for capital gains, still delivers small losses to the lender compared to the optimal allocation.

The implementation clarifies the discrepancy of the firm size dynamics with risk neutrality and risk aversion. With risk neutrality, it is optimal to reward the entrepreneur solely through a higher stake in the project to minimize diversion incentives. Therefore, the promised utility can be mapped to the entrepreneur's equity (Clementi and Hopenhayn (2006)). This is no longer the case with risk aversion. As I show, promised utility better maps to the entrepreneur's wealth, and the promised insurance maps to the entrepreneur's equity share. Consequently, both models obtain a positive relation between equity and firm size. Intuitively, regardless of the entrepreneur's preferences, a high equity share disincentivizes the diversion of funds, so the lender is willing to provide more capital. However, with risk neutrality, the equity share drifts upwards, but with risk aversion and persistence, it drifts

[5]More concretely, for every draw of a montecarlo simulation, I compute the distance between the induced consumption allocation under the optimal contract and the implementation at all periods.

[6]Pledging shares aligns the entrepreneur's consumption with the firm's value, but without having to sell shares and independently of dividend payout policies. In this model selling shares may not be optimal, as lowering the entrepreneur's stake on the firms increases his incentives to divert funds. This rationale is consistent with the primary motive for pledging shares estimated in Fabisik (2019): obtain liquidity while maintaining ownership.

downwards[7]. The equity share dynamics of the risk neutral model may be at odds with what we observe in the data. For example, in the venture capital industry, the founder's equity share gets diluted over the financing rounds as the firm grows (Azevedo *et al.* (2023)). Therefore, it may be challenging for this class of models to generate realistic firm size and equity dynamics simultaneously.

I explore three extensions to the main model: (i) limited commitment of the entrepreneur as in (Albuquerque and Hopenhayn (2004)), (ii) a model where the entrepreneur can choose the fraction of funds invested and diverted, and (iii) allowing the lender to terminate the contract. Although termination may be optimal, it does not affect the equations characterizing the optimal contract presented throughout the paper. Moreover, I also discuss that, contrary to the model with risk neutrality (Clementi and Hopenhayn (2006)), termination probabilities should tend to increase over time. In a simplified version of the model, I show that if termination is optimal, termination probabilities increase with the persistence of the process; the intuition is similar to that of the equity purchases.

I use two tools from the dynamic public finance literature to solve and characterize the optimal allocation. The first is the first-order approach (FOA) as in Kapička (2013), Farhi and Werning (2013), Pavan, Segal and Toikka (2014) and Golosov *et al.* (2016a). It consists of solving a relaxed problem with the local incentive compatibility constraints. The FOA is popular in dynamic public finance, but it is also used more broadly in dynamic mechanism design. The FOA allows solving the model with persistent private information. A priori, global incentive compatibility constraints may bind. Following the procedure in Kapička (2013) and Farhi and Werning (2013), I verify ex-post that this is not the case in all the numerical simulations.

The second tool allows for deriving analytical characterizations of the optimal allocation with risk aversion. This is the change of measure used in Hellwig (2021) for a Mirrlees taxation problem with non-separable preferences between consumption, leisure and type. The challenge is that, with risk aversion, the marginal information rents depend on consumption. So if the lender redistributes consumption around some type θ, the slope of the profile of information rents for types $\theta' > \theta$ changes. Following Hellwig (2021), incentive-adjusted

[7] Another interesting implication concerns the role of capital structure on the firm's value (i.e. the Modigliani-Miller theorem). With risk neutrality, the firm's value does depend on the promised utility given to the entrepreneur (Clementi and Hopenhayn (2006)). Instead, with risk aversion, firm size, and so firm value, is approximately independent of promised utility, but they are decreasing on the amount of insurance promised. This observation corroborates the idea that, with risk aversion, promised utility maps to the entrepreneur's private wealth and is unrelated to the firm's capital structure.

measures can be employed to reweight the density of types such that the lender's evaluation of allocations accounts for the changes in information rents, and therefore, incentive compatibility is preserved.

Related literature: This paper contributes to the dynamic financial contracting literature. Important early work on this class of models includes Clementi and Hopenhayn (2006), Albuquerque and Hopenhayn (2004), Biais *et al.* (2007), Biais *et al.* (2010), DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007a), DeMarzo and Fishman (2007b) and DeMarzo *et al.* (2012). In particular, I contribute to the literature by studying a workhorse dynamic cash flow diversion model with risk aversion and persistent private information.

That firm size decreasing over time can be the outcome of an optimal contract has also been shown in Clementi *et al.* (2010). They study a dynamic moral hazard model where the firm's productivity distribution depends on the entrepreneur's costly effort exerted. In their model, the entrepreneur becomes wealthier over time which lowers the effort and, consequently, firm size. Thus, these dynamics are driven by a wealth effect on incentives, which is different from the increase in insurance resulting from the persistent private information.

Models with persistence have been recently analyzed in DeMarzo and Sannikov (2016), Fu and Krishna (2019) and Krasikov and Lamba (2021), but all these papers assume a risk neutral entrepreneur. To my knowledge, this is the first paper in the dynamic financial contracting literature with both persistent private information and risk aversion. As I will show, some key effects of persistence on the optimal allocation, especially the downwards drift in firm size, are only present with risk aversion. Fu and Krishna (2019) and Krasikov and Lamba (2021) show some interesting role of persistence on the dynamics of distortions. However, as in the i.i.d risk neutral models, they still find that distortions eventually disappear[8]. Models with risk aversion have been studied in He (2012) and Di Tella and Sannikov (2021). Both papers study a hidden savings problem, so the entrepreneur has persistent private information about his savings. I do not allow for hidden savings but allow for persistent private information about the firm's productivity.

Throughout the paper, I use tools and insights from the dynamic Mirrlees literature[9]. As discussed, I use the FOA and set up the principal's problem recursively as in Kapička (2013),

[8]As shown in Makris and Pavan (2020), with risk neutrality wedge dynamics only depend on the impulse response of the initial period, and so wedges converge to zero.

[9]For a review of the literature see Stantcheva (2020). In some aspects, the model also resembles the setting of the dynamic taxation problems in Stantcheva (2017) and Brendon (2022).

Farhi and Werning (2013) or Golosov *et al.* (2016a). I also use incentive-adjusted probability measures as in Hellwig (2021) to derive analytical characterizations of the optimal contract. The finding that firm size drifts downwards follows from the insight of the Dynamic Mirrlees literature that labor wedges tend to increase over time (Farhi and Werning (2013), Makris and Pavan (2020)).

Finally, this paper is also related to the literature on insurance with persistent private information (Williams (2011), Bloedel *et al.* (2018) and Bloedel *et al.* (2020)). With fixed capital, the cash flow diversion model studied in this paper is equivalent to the hidden endowment model used in this literature. Their focus is on the role of persistent private information for the long-run distribution of consumption and whether or not it features immiseration (Thomas and Worrall (1990) and Atkeson and Lucas (1992)). In the paper, I present some results and discussion on the long-run consumption dynamics. Nevertheless, numerical simulations show that in this model, immiseration is a very long-run phenomenon so that it may be irrelevant for the usual lifespan of a firm.

Outline: The rest of the paper is organized as follows. Section 2 describes the model, sets up the relaxed planning problem, and presents the first best allocation. Section 3 presents the main results on the firm size and consumption dynamics, and section 4 illustrates them with numerical simulations. Section 5 studies the quasi-implementation. Section 6 discusses the differences in models with risk neutrality and risk aversion and their implications. Finally, section 7 briefly summarizes the extensions to the main model, and section 8 concludes.

2 Model

Time is discrete and indexed by $t = 0, 1, ..., \infty$. Every period an entrepreneur (the agent, "he") needs funds k_t from a lender (the principal, "she") to operate a project. Both the entrepreneur and the lender are long-lived. At period t, the project generates a cash flow $f(k_t, \theta_t)$, where $\theta_t \in [\underline{\theta}, \overline{\theta}]$ is the entrepreneur's productivity type. The agent's type history is denoted by $\theta^t = \{\theta_0, ..., \theta_t\}$ and is the agent's private information. θ_t follows a Markov process with conditional density $\varphi_t(\theta_t|\theta^{t-1})$.

The lender cannot observe the returns and instead relies on the entrepreneur's report.The entrepreneur can misreport and divert a fraction of the cash flow for his consumption. There is a deadweight loss $(1 - \iota) \in [0, 1)$ on diverted funds. That is, for every dollar of funds diverted, the entrepreneur only gets to consume a fraction $(1-\iota)$. After the entrepreneur reports returns $f(k_t, \widetilde{\theta}_t)$, the lender asks for a repayment $b_t(\widetilde{\theta}_t)$ and advances funds $k_{t+1}(\widetilde{\theta}_t)$ for the next period. The entrepreneur cannot privately save[10], so the entrepreneur's period t consumption if the true cash flow is $f(k_t, \theta_t)$ but he reports $f(k_t, \widetilde{\theta}_t)$ is

$$c_t = f(k_t, \theta_t) - (1-\iota)\Big(f(k_t, \theta_t) - f(k_t, \widetilde{\theta}_t)\Big) - b_t(\widetilde{\theta}_t) \tag{1}$$

In particular, if the entrepreneur does not misreport returns he consumes $c_t = f(k_t, \theta_t) - b_t(\theta_t)$. As is common, I further assume that the agent cannot overreport his returns. That is, reports are restricted to $\widetilde{\theta}_t \leq \theta_t$. This assumption is motivated by the restriction that the entrepreneur cannot save outside the contract with the lender. The entrepreneur is risk averse, derives utility $u(c_t)$ from consumption, and discounts the future at rate β. Throughout the paper, I will use the following notation for the derivatives of the production function

$$f_k(k_t, \theta_t) \equiv \frac{\partial f(k_t, \theta_t)}{\partial k_t} \quad f_\theta(k_t, \theta_t) \equiv \frac{\partial f(k_t, \theta_t)}{\partial \theta_t} \quad f_{\theta k}(k_t, \theta_t) \equiv \frac{\partial^2 f(k_t, \theta_t)}{\partial \theta_t \partial k_t}$$

Below I summarize all the assumptions on the productivity process and the functions f and u.

Assumptions

A1: The conditional density $\varphi_t(\theta_t|\theta^{t-1})$ is differentiable with respect to the second argument and persistent, i.e.

$$\mathcal{E}(\theta_t, \theta_{t-1}) \equiv \frac{\frac{\partial \varphi_t(\theta_t|\theta^{t-1})}{\partial \theta_{t-1}}}{\varphi_t(\theta_t|\theta^{t-1})}$$

is non-decreasing in θ_t.

A2: The production function is twice differentiable and satisfies $f_{kk} < 0 < f_k$, $f_\theta > 0$, the Inada conditions $\lim_{k\to 0} f_k(k, \theta) = \infty$ and $\lim_{k\to\infty} f_k(k, \theta) = 0$, and $f_{\theta k} > 0$

A3: The utility function satisfies $u'' < 0 < u'$, and the Inada conditions $\lim_{c\to 0} u'(c) = \infty$ and $\lim_{c\to\infty} u'(c) = 0$

[10] If the agent's savings are observable, it is without loss to have the lender do all the savings for the entrepreneur. In this case, allowing the agent to also save by himself is straightforward. Let d_t be dividend payments, w_t the agent's net worth and B_t the funds advanced by the principal. Then we would have $c_t = d_t$, a LOM for the entrepreneur's net worth $w_{t+1} = f(k_t, \theta_t) - b_t - d_t + w_t$ and investment equal to $k_{t+1} = w_{t+1} + B_{t+1}$.

The first assumption (A1) requires that type process has either positive persistence or is independent over time, in which case $\frac{\partial \varphi_t(\theta_t|\theta^{t-1})}{\partial \theta_{t-1}} = 0$. The process is allowed to be time-dependent. Differentiability will be needed to use the envelope condition for the local incentive constraint. For future use, it is useful to define:

$$\rho_t(\theta^t) \equiv \frac{1-\Phi_t(\theta_t|\theta^{t-1})}{\varphi_t(\theta_t|\theta^{t-1})}\mathbb{E}\left[\mathcal{E}(\theta',\theta^{t-1})|\theta' \geq \theta_t, \theta^{t-1}\right] = \frac{\frac{\partial}{\partial \theta_{t-1}}\left(1-\Phi(\theta_t|\theta^{t-1})\right)}{\varphi_t(\theta_t|\theta^{t-1})} \tag{2}$$

This is the impulse response of θ_t to θ_{t-1} as defined in Pavan, Segal and Toikka (2014). It is a measure of the persistence of the process. If the type process follows an AR(1) with autoregressive parameter ρ, then $\mathcal{E}(\theta_t,\theta^{t-1}) = -\rho\frac{\partial \varphi_t(\theta_t|\theta^{t-1})}{\partial \theta_t}/\varphi_t(\theta_t|\theta^{t-1})$ and $\rho_t(\theta^t) = \rho$. Assumption A2 states that there is decreasing marginal product of investment, higher types obtain higher returns and have a higher marginal product. This last assumption ($f_{\theta k} > 0$) is key as it will imply that higher capital increases information rents.

2.1 Lender's problem

The lender is risk neutral and discounts the future at rate q. By the revelation principle, it is without loss to focus on direct mechanisms. At any history, the entrepreneur sends a report $r \in [\underline{\theta},\theta_t]$ about θ_t to the lender. Define a reporting strategy by $\sigma = \{\sigma_t(\theta^t)\}$, it implies a history of reports $\sigma^t(\theta^t) = \{\sigma_1(\theta_0), ..., \sigma_t(\theta^t)\}$. Let $\Sigma = \{\sigma|\sigma_t(\theta^t) \leq \theta_t \quad \forall \theta^t \in [\underline{\theta},\overline{\theta}]^t\}$ be the set of feasible reporting strategies. The entrepeneur's continuation utility with truth-telling can be written recursively as

$$w_t(\theta^t) = u(c(\theta^t)) + \beta\int w_{t+1}(\theta^t,\theta_{t+1})\varphi_{t+1}(\theta_{t+1}|\theta^t)d\theta_{t+1} \tag{3}$$

where $c(\theta^t) = f(k_t(\theta^{t-1}),\theta_t) - b_t(\theta^t)$. The continuation utility of type θ^t with reporting strategy σ is

$$w_t^\sigma(\theta^t) = u(c(\theta_t,\sigma^t(\theta^t))) + \beta\int w_{t+1}^\sigma(\theta^t,\theta_{t+1})\varphi_{t+1}(\theta_{t+1}|\theta^t)d\theta_{t+1} \tag{4}$$

where

$$c(\theta_t,\sigma^t(\theta^t)) = \iota f(k_t(\sigma^{t-1}(\theta^{t-1})),\theta_t) + (1-\iota)f(k_t(\sigma^{t-1}(\theta^{t-1})),\sigma_t(\theta_t)) - b_t(\sigma^t(\theta^t)) \tag{5}$$

The lender problem consists of choosing an allocation $\{k_{t+1}(\theta^t), b_t(\theta^t)\}$ to minimize the cost of providing expected utility v_0[11] subject to the incentive compatibility constraints:

$$\begin{aligned} K(v_0) \quad &= \quad \min_{\{k_{t+1}(\theta^t), b_t(\theta^t)\}} \mathbb{E}_0 \left[\sum_{t=1}^{\infty} q^t \left(k_{t+1}(\theta^t) - b_t(\theta^t) \right) \right] \\ & \quad\quad s.t \quad \mathbb{E}_0 \left[w_1(\theta^1) \right] \geq v_0 \qquad (PK) \\ & \quad\quad w_t(\theta^t) \geq w_t^{\sigma}(\theta^t) \quad \forall \theta^t \in [\underline{\theta}, \overline{\theta}]^t \text{ and } \ \sigma \in \Sigma \quad (IC) \end{aligned} \tag{6}$$

2.2 Relaxed problem

With Markov shocks, it is sufficient to consider only the temporary incencentive compatibility constraint (Fernandes and Phelan (2000), Kapička (2013))

$$w_t(\theta^t) = \max_{r \in [\underline{\theta}, \theta_t]} u(c(\theta_t, r)) + \beta \int w_{t+1}(\theta^{t-1}, r, \theta_{t+1}) \varphi_{t+1}(\theta_{t+1} | \theta^t) d\theta_{t+1} \tag{7}$$

where type's θ_t consumption if he reports r, $c(\theta_t, r)$, is given by equation (5). This allows us to solve a recursive problem. Write entrepreneur's continuation utility under truth-telling as

$$w_t(\theta^t) = u(c(\theta^t)) + \beta v_t(\theta^t) \tag{8}$$

$$v_t(\theta^t) = \int w_{t+1}(\theta^{t+1}) \varphi_{t+1}(\theta_{t+1} | \theta^t) d\theta_{t+1} \tag{9}$$

Following Kapička (2013), Farhi and Werning (2013) and Pavan, Segal and Toikka (2014), I use the first-order approach. That is, I solve a relaxed problem with the local IC constraint[12]. The envelope condition of the temporary IC (7) gives

$$\frac{\partial}{\partial \theta_t} w_t(\theta^t) = \underbrace{u'(c(\theta^t)) \iota f_\theta(k_t(\theta^{t-1}), \theta_t)}_{\text{Static marginal info rent}} + \underbrace{\beta \Delta_t(\theta^t)}_{\text{Markov marginal info rent}} \tag{10}$$

$$\Delta_t(\theta^t) = \int w_{t+1}(\theta^{t+1}) \frac{\partial \varphi_{t+1}(\theta_{t+1} | \theta^t)}{\partial \theta^t} d\theta_{t+1} \tag{11}$$

With persistent private information, the marginal information rents depend on two terms.

[11] The constant v_0 may correspond to the entrepreneur's outside option, or it can be pinned down by the lender's break-even condition, i.e. $K(v_0) = 0$.

[12] Following Kapička (2013) and Farhi and Werning (2013), I verify numerically that the global IC constraints do not bind. More details can be found in section 4 and Appendix E.

The static component captures how much the agent can gain by marginally misreporting returns in the current period. The Markov information rent, which can be rewritten as $\Delta_t(\theta^t) = \mathbb{E}\left[\rho(\theta^{t+1})\frac{\partial w(\theta^{t+1})}{\partial \theta_{t+1}}|\theta^t\right]$, captures the information rent that the agent obtains by having more information about future types than the principal. If types are i.i.d we have $\Delta_t(\theta^t) = 0$.

If the entrepreneur is risk averse the static marginal information rent $(u'(c(\theta^t))\iota f_\theta(k_t(\theta^{t-1}), \theta_t))$ depends on the entrepreneur's consumption. Intuitively, if the entrepreneur's productivity increases by $d\theta_t$, he generates an extra return of $f_\theta(k_t(\theta^{t-1}), \theta_t)d\theta_t$. The entrepreneur can then decide to mimick the returns of the type right below him and divert the extra funds, he can then obtain $\iota f_\theta(k_t(\theta^{t-1}), \theta_t)d\theta_t$ extra consumption units. This extra information rent has to be transformed into utils by multiplying by $u'(c(\theta^t))$. The fact that information rents depend on the entrepreneur's consumption poses a challenge for characterizing the solution to this problem. If the principal increases consumption of type θ_t, then this type's information marginal rent changes. But then the information rents of all types $\theta' > \theta_t$ must be adjusted non-linearly in order to preserve incentive compatibility. In section 3.1, I will discuss how the incentive-adjusted probability measures developed in Hellwig (2021) can be used to take into account these changes in information rents.

The principal solves a dynamic optimization program where within every period, there is an optimal control problem. The relaxed problem is

$$\begin{aligned}
K_t(v_{t-1}, \Delta_{t-1}, \theta_{t-1}, k_t) = \min \int \big(k_{t+1}(\theta^t) - b_t(\theta^t) + qK_{t+1}(v_t(\theta^t), \Delta_t(\theta^t), \theta_t, k_{t+1}(\theta^t))\big)\, \varphi_t(\theta_t|\theta^{t-1})d\theta_t \\
s.t \quad (PK) \quad & w_t(\theta^t) = u(c(\theta^t)) + \beta v_t(\theta^t) \qquad [\varphi_t(\theta_t|\theta^{t-1})\xi_t(\theta^t)] \\
& v_{t-1} = \int w_t(\theta^t)\varphi_t(\theta_t|\theta^{t-1})d\theta_t \qquad [\varphi_t(\theta_t|\theta^{t-1})\lambda_t] \\
(IC) \quad & \dot{w}_t(\theta) = u'(c(\theta^t))\iota f_\theta(k_t, \theta_t) + \beta\Delta_t(\theta^t) \qquad [\mu_t(\theta^t)] \\
& \Delta_{t-1} = \int w_t(\theta^t)\frac{\partial \varphi_t(\theta_t|\theta^{t-1})}{\partial \theta^{t-1}}d\theta_t \qquad [\varphi_t(\theta_t|\theta^{t-1})\gamma_t] \\
(Feasibility) \quad & c(\theta^t) = f(k_t, \theta_t) - b_t(\theta^t)
\end{aligned} \tag{12}$$

Along with the promised utility, v_{t-1}, the previous period's type, θ_{t-1}, and the funds advanced at $t-1$, k_t, the past Markov information rents, Δ_{t-1}, become an extra state variable of the problem. Intuitively, the principal can lower Markov information rents by promising to reduce information rents in future periods. That is, she promises she will increase insurance (or equivalently distortions) in the future to lower current Markov information rents. Because the past promises must be satisfied, Δ_{t-1} has to be added as an extra state variable of the problem. Throughout the paper, I will directly refer to this state variable as the promised insurance.

The co-state variable of the within period Hamiltonian is $\mu_t(\theta^t)$. This co-state variable will become key for the dynamics later one. We will refer to it as the shadow cost of insurance. Intuitively, $\mu_t(\theta^t)$ captures the resource gain from redistributing consumption around θ^t, while preserving incentive compatibility, promised expected utility (v_{t-1}) and prior promised insurance (Δ_{t-1}). Note that I write inside square brackets the multipliers associated with each constraint. The Hamiltonian of this problem and the derivation of the optimality conditions can be found in Appendix B. To economize notation, I will often write directly $u(\theta^t)$ and $f(\theta^t)$ instead of $u(c(\theta^t))$ and $f(k_t(\theta^{t-1}), \theta_t)$. The sequential problem (6) can be recovered by treating Δ_0 and k_1 as free variables, $K(v_0) = \min_{\Delta_0, k_1} K(v_0, \Delta_0, \theta_0, k_1)$.

2.3 First Best

To gain intuition on the model, it is useful to first look at the first best allocation, i.e. with no private information. The results are summarized in the following proposition

Proposition 1. *In the First Best, at any history θ^t, there is*

1. *No diversion of funds*

$$f(k_t, \widetilde{\theta}_t) = f(k_t, \theta_t) \tag{13}$$

2. *No distortion of the firm's size*

$$\frac{1}{q} = \mathbb{E}\left[f_k(k_{t+1}(\theta^t), \theta_{t+1}) | \theta^t\right] \tag{14}$$

3. *Full insurance and intertemporal consumption smoothing*

$$u'(c(\theta^t)) = \frac{\beta}{q} u'(c(\theta^{t+1})) \tag{15}$$

Because diverting funds is inefficient and by the revelation principle, in the second best, there will also be no diversion of funds. However, points 2. and 3. of the proposition do not hold in the second best allocation. In particular, firm size is distorted downwards, and the entrepreneur is exposed to risk. In the following section, we will study how with private information, the firm size and compensation dynamics differ from the first best benchmark.

3 Optimal allocation

In this section, I present the two main results on the dynamics of the optimal allocation. I start with the firm size dynamics (section (3.1)). First, I show that they are driven by the dynamics of the normalized shadow cost of insurance ($\widetilde{\mu}_t$). Second, I introduce the incentive-adjusted probability measures as in Hellwig (2021) to characterize $\widetilde{\mu}_t$ and discuss why it tends to increase over time with persistent private information. Then, I turn to the compensation dynamics (section (3.2)). I again use incentive-adjusted measures to characterize the entrepreneur's consumption process and discuss the implications.

The approach to characterizing the dynamics is different from the risk neutral case. With risk neutrality, there is a one-to-one relation between promised utility (compensation) and firm size (Clementi and Hopenhayn (2006)). Risk aversion breaks this tight relation. Although there can be some interactions, the firm size and compensation dynamics are characterized independently. This separation will also be illustrated with the numerical simulations (section 4).

3.1 Firm size dynamics

The firm dynamics implied by this cash flow diversion model with risk neutrality are well understood. On average, firm size tends to increase over time until it converges to the first best (Clementi and Hopenhayn (2006)). This is true regardless of whether the shocks are i.i.d or persistent (Fu and Krishna (2019)). However, as I will show in this section, the firm dynamics are remarkably different when we allow the entrepreneur to be risk averse. With both risk aversion and persistent private information, the firm's size tends to decrease over time.

Following the Public Finance tradition, it is helpful to describe the optimal allocation in terms of implicit wedges, i.e. distortions in the second best allocation relative to the first best. We define the investment wedges as the type-dependent distortion in the marginal product of capital faced by the lender

$$\frac{1}{q} = \mathbb{E}\left[f_k(k_{t+1}(\theta^t), \theta_{t+1})(1 - \tau^k(\theta^{t+1}))|\theta^t\right] \tag{16}$$

Besides the direct effect of the productivity process $\{\theta^t\}$, now the dynamics of the firm's size $(k_{t+1}(\theta^t))$ also depend on the dynamics of the investment wedge. Therefore, to characterize the firm size dynamics in the second best, it is sufficient to focus on the dynamics of the investment wedge. The following proposition shows that the investment wedge satisfies an ABC-type formula (Diamond (1998)).

Proposition 2. *At any history θ^{t+1}, the investment wedge $\tau^k(\theta^{t+1})$ satisfies*

$$\tau^k(\theta^{t+1}) \equiv \iota\mathcal{E}^f(\theta^{t+1}) \times \widetilde{\mu}_{t+1}(\theta^{t+1}) \times \Psi(\theta^{t+1}) \geq 0$$

where

$$\mathcal{E}^f(\theta^{t+1}) = \frac{\theta_{t+1} f_{\theta k}(\theta^{t+1})}{f_k(\theta^{t+1})} > 0$$

$$\widetilde{\mu}_{t+1}(\theta^{t+1}) = \frac{\mu_{t+1}(\theta^{t+1})}{1-\Phi_{t+1}(\theta_{t+1}|\theta^t)} u'(\theta^{t+1}) \geq 0$$

$$\Psi(\theta^{t+1}) = \frac{1-\Phi_t(\theta_{t+1}|\theta_t)}{\theta_{t+1}\varphi_{t+1}(\theta_{t+1}|\theta_t)} \geq 0$$

Because $\tau^k(\theta^{t+1}) \geq 0$, the investment wedges lower the implicit marginal product of capital, and so $k^{SB}_{t+1}(\theta^t) \leq k^{FB}_{t+1}(\theta^t)$. The first term equals the elasticity of the marginal product of capital with respect to productivity, $\mathcal{E}^f(\theta^{t+1})$, times the ability to consume diverted funds, ι. Intuitively, because $f_{\theta k} > 0$, increasing capital increases relatively more the returns of higher types. Therefore, their ability to divert funds increases, i.e. the higher types' information rents (in consumption units) increase by more, which is costly for the lender. This cost is proportional to the normalized shadow cost of insurance $\widetilde{\mu}(\theta^{t+1})$. This term increases when the lender wants (or has promised) to provide more insurance to the entrepreneur. So when $\widetilde{\mu}(\theta^{t+1})$ is high, an increase in information rents is more costly. The last term, $\Psi(\theta^{t+1})$, measures the thickness of the right tail of the distribution of productivity shocks. The wedges increase in $\Psi(\theta^{t+1})$ because the cost of increasing information rents is higher when there is a higher mass of types above θ^{t+1}.

For production functions of the form $f(k_{t+1}, \theta_t) = \theta_{t+1} f(k_{t+1})$, the elasticity simplifies to $\mathcal{E}^f(\theta^{t+1}) = 1$, and for log-additive functions (i.e. $f(k_{t+1}, \theta_t) = g(\theta_{t+1}) f(k_{t+1})$) $\mathcal{E}^f(\theta^{t+1})$ is only a function of θ_{t+1}. Consequently, the terms $\iota\mathcal{E}^f(\theta^{t+1})$ and $\Psi(\theta^{t+1})$ usually depend only on exogenous variables, and they are stationary as long as the type process is also stationary. Therefore, the normalized shadow cost of insurance ($\widetilde{\mu}(\theta^{t+1})$) typically drives all the wedge dynamics.

It has been shown in dynamic screening models that with risk aversion and persistent private information, the shadow costs of insurance (and so wedges) tend to increase over time (Farhi and Werning (2013), Makris and Pavan (2020)[13]). This is also true in this model, which implies that if types are persistent, firm size will tend to decrease over time. In what follows, we first characterize the shadow costs and then show they tend to increase over time.

Characterization of $\widetilde{\mu}_{t+1}$: As discussed, the main challenge for characterizing the optimal allocation in this problem is that the static marginal information rents, $u'(c(\theta^t))\iota f_\theta(k(\theta^t), \theta_t)$, depend on consumption. The literature on dynamic mechanism design with risk aversion typically analyses models with separable preferences of the form

$$U(\theta, y, c) = u(c) - \psi(y, \theta) \tag{17}$$

where y can represent the agent's income or effort. As discussed in Makris and Pavan (2020), this includes dynamic public finance models with separable preferences but also models of managerial compensation, among others. In these settings, the static marginal information rents are $\psi_\theta(y, \theta)$, which do not depend on consumption. Therefore, there is a complete separation between insurance (or redistribution) and marginal information rents. Imagine that at history θ^{t-1} the principal redistributes consumption from types all types $\theta_t' < \theta_t$ to all $\theta_t'' > \theta_t$. It is well known that in settings as (17), it is sufficient to increase utility uniformly to all types $\theta_t'' > \theta_t$ to preserve incentive compatibility, so consumption has to be redistributed in proportion to $\frac{1}{u'(\theta'')}$. In our setting, this perturbation in consumption changes marginal information rents $(u'(c(\theta^t))\iota f_\theta(k(\theta^t), \theta_t))$. Incentive compatibility now requires redistributing consumption according to $M(\theta^t) = \frac{1}{u'(\theta^t)} e^{-\int_{\theta_t}^{\overline{\theta}} \frac{u''(\theta^{t-1},\theta')\iota f_\theta(\theta^{t-1},\theta')}{u'(\theta^{t-1},\theta')} d\theta'}$. This is the same challenge one encounters in models with arbitrary non-separable preferences $U(\theta, y, c)$ (see Hellwig (2021), Hellwig and Werquin (2022)). Following Hellwig (2021), the factor $m(\theta^t) = u(\theta^t)M(\theta^t)$ can be interpreted as a reweighting of the type distribution (or pareto weights). Accordingly, we define the incentive-adjusted probability measures as

$$\hat{\varphi}_t(\theta_t|\theta^{t-1}) \equiv \frac{\varphi_t(\theta_t|\theta^{t-1})m(\theta^t)}{\mathbb{E}[\varphi_t(\theta_t|\theta^{t-1})m(\theta^t)|\theta^{t-1}]} \tag{18}$$

[13]The early papers attributed these wedge dynamics to the variance of the types increasing over time, which is the case if the type process follows a random walk. However, Makris and Pavan (2020) have clarified why this intuition is incomplete, as wedges can increase even if the variance of the types decreases over time.

Figure 1: Incentive-adjusted probability measure

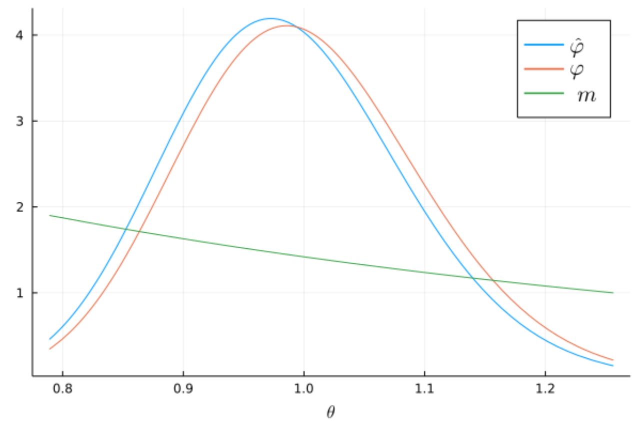

Note: The plot is computed with the same calibration as the main simulations in section 4 for i.i.d types. Observe that m is monotonically decreasing, and the incentive-adjusted measure $\hat{\varphi}$ puts more weight on the lower type realizations.

Therefore, the new measure $\hat{\varphi}$ reweights the density of types such that these perturbations in consumption preserve incentive compatibility. Because $\frac{m'(\theta^t)}{m(\theta^t)} = \frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)} < 0$, the function $m(\theta^t) = e^{-\int_{\theta_t}^{\overline{\theta}} \frac{u''(\theta', \theta^{t-1})\iota f_\theta(\theta', \theta^{t-1})}{u'(\theta', \theta^{t-1})} d\theta'}$ is decreasing in θ_t. So, $\Phi_t(\cdot|\theta^{t-1})$ first-order stochastically dominates $\hat{\Phi}_t(\cdot|\theta^{t-1})$. That is, incentive compatibility requires evaluating allocations as if the principal puts more weight on lower types, see Figure 1. Intuitively, lower types have lower returns, so they collect lower information rents in consumption units. Hence, their marginal utility is higher. When the principal redistributes consumption, information rents change more for types with high marginal utility. Therefore, the incentive-adjusted measure that guarantees incentive compatibility has to put more weight on lower types. The following proposition uses the incentive-adjusted measure to characterize $\widetilde{\mu}_t$.

Proposition 3. *(Hellwig (2021)) The normalized shadow cost of insurance $\widetilde{\mu}_{t+1}(\theta^{t+1})$ satisfies*

$$\widetilde{\mu}_{t+1}(\theta^{t+1}) = \hat{MB}(\theta^{t+1}) + \hat{\rho}(\theta^{t+1})\frac{\beta}{q}\Psi(\theta^t)\frac{\theta_t}{u'(\theta^t)}\widetilde{\mu}_t(\theta^t) \tag{19}$$

With

$$\hat{MB}(\theta^{t+1}) = \frac{\mathbb{E}\left(m(\theta^t, \theta')|\theta' \geq \theta_{t+1}, \theta^t\right)}{M(\theta^{t+1})}\left\{\hat{\mathbb{E}}\left[\frac{1}{u'\left(\theta', \theta^t\right)} \mid \theta' \geq \theta_{t+1}, \theta^t\right] - \hat{\mathbb{E}}\left[\frac{1}{u'\left(\theta^{t+1}\right)}|\theta^t\right]\right\} \tag{20}$$

$$\hat{\rho}(\theta^{t+1}) \equiv \frac{\mathbb{E}\left(m(\theta^t, \theta')|\theta' \geq \theta_{t+1}, \theta^t\right)}{M(\theta^{t+1})}\left\{\hat{\mathbb{E}}\left[\mathcal{E}(\theta', \theta^t) \mid \theta' \geq \theta_{t+1}, \theta^t\right] - \hat{\mathbb{E}}\left[\mathcal{E}(\theta_{t+1}, \theta^t)|\theta^t\right]\right\} \tag{21}$$

Note that the operator $\hat{\mathbb{E}}$ denotes expectations under the measure $\hat{\varphi}$. The proposition shows that the normalized shadow cost of insurance is a function of two terms. The first, $\hat{MB}(\theta^{t+1})$, is a static marginal benefit of redistributing consumption from types $\theta' > \theta_{t+1}$ to $\theta'' < \theta_{t+1}$. The second is a backward-looking term that accounts for past insurance promises. With utility functions as in (17), one can derive the same characterization but under the original measure φ and with $m(\theta^t) = 1$ (Makris and Pavan (2020), Brendon (2013), Hellwig (2021)).

The normalized shadow costs $\widetilde{\mu}$ (and so investment wedges τ^k) increase over time: Proposition 3 shows that the shadow costs $\widetilde{\mu}_{t+1}$ are persistent. Iterating backward on equation (19) we get

$$\widetilde{\mu}_{t+1}(\theta^{t+1}) = \sum_{\tau=0}^{t} \left(\frac{\beta}{q}\right)^{\tau} \prod_{s=0}^{\tau-1} \left(\hat{\rho}_{t+1-s}(\theta^{t+1-s}) \Psi(\theta^{t-s}) \frac{\theta_{t-s}}{u'(\theta^{t-s})} \right) \hat{MB}_{t-\tau}(\theta^{t+1-\tau}) \tag{22}$$

The formula shows that the current shadow costs of insurance are a function of current and past marginal benefits of insurance $\hat{MB}_{t-\tau}(\theta^{t+1-\tau})$. Notice that at the initial period $\widetilde{\mu}_1(\theta^1) = \hat{MB}_1(\theta^1)$, and the passthrough of past $\hat{MB}_{t-\tau}(\theta^{t+1-\tau})$ to $\widetilde{\mu}_{t+1}(\theta^{t+1})$ is always positive. Therefore, $\widetilde{\mu}_{t+1}(\theta^{t+1})$ and $\tau^k(\theta^{t+1})$ will tend to grow with the distance from the starting period. That is, it is optimal for the principal to increase the insurance provided to the agent over time. The intuition is the following. With persistent private information, different types θ_t have different preferences for period $t+1$ contracts. In particular, higher types value relatively less contracts with high insurance at $t+1$, as they know they are expected to be more productive then. The principal can use this to lower the resource cost of screening types at every period. More concretely, if the principal promises to provide more insurance to type (θ^{t-1}, θ') (i.e lowers $\Delta_t(\theta^{t-1}, \theta')$) this relaxes the incentive constraints of types (θ^{t-1}, θ') with $\theta'' > \theta'$. Because every period the principal can gain by promising to provide more insurance in the future, the shadow costs $\widetilde{\mu}_t$ will tend to increase over time. However, as will be shown in the numerical simulations, the wedges may, over time, converge to a stationary distribution. I use this intuition to explain why the lender may want to use equity purchases in the implementation (see section 5.2)[14].

[14] Alternatively, imagine the principal increases consumption of all types $(\theta^{t-1}, \widetilde{\theta}_t)$ with $\widetilde{\theta}_t > \theta_t$. To preserve incentive compatibility, the principal needs to adjust the information rent of all types $(\theta^{t-2}, \theta'_{t-1})$ with $\theta'_{t-1} > \theta_{t-1}$. Because if types are persistent (i.e $\rho_t(\theta^t) > 0$), types $\theta'_{t-1} > \theta_{t-1}$ have a higher probability of being type $\widetilde{\theta}_t$ at period t. This adjustment has to be done for all types $(\theta^{\tau-1}, \theta'_{\tau-1})$ with $\theta'_{\tau-1} > \theta_{\tau-1}$ at all periods $\tau < t$. Therefore, these costs will tend to increase over time if types are persistent. For a clearer and more detailed intuition on this, see Makris and Pavan (2020).

It is important to stress that for every type θ_t firm size $(k_{t+1}(\theta_t))$ is never larger than in the initial period. The reason is that the principal initializes the contract by setting Δ_0 freely. So Δ_0 is set to not have any "extra" promised insurance. Consequently, for every $\theta_t \in [\underline{\theta}, \overline{\theta}]$, the wedges will not be smaller than in the initial period.

Moreover, both risk aversion and persistence are necessary to have investment wedges increasing over time. If the agent is risk neutral we have $\hat{MB}_{t+1}(\theta^{t+1}) = 0$, which implies $\widetilde{\mu}_{t+1}(\theta^{t+1}) = 0$. If the type process is not persistent we have $\hat{\rho}_t(\theta^t) = \rho_t(\theta^t) = 0$ and

$$\widetilde{\mu}_{t+1}(\theta^{t+1}) = \hat{MB}_{t+1}(\theta^{t+1}) \tag{23}$$

so past marginal benefits of insurance do not affect the current shadow costs. As I will show in the numerical simulations, wedges are stationary with i.i.d types.

In a separable model with utility as in (17), one also obtains the same formula as (22) but with ρ, i.e. with impulse responses under the original type measure. The change of measure can amplify or dampen the persistence of the wedges. The (unormalized) persistence is $\frac{u'\theta^t)\varphi(\theta_t|\theta^{t-1})}{1-\Phi(\theta_t|\theta^{t-1})}\hat{\rho}_t(\theta^t) \gtreqless \rho_t(\theta^t)$ if $\rho_t(\theta, \theta^{t-1})\frac{u''(\theta,\theta^{t-1})f_\theta(\theta,\theta^{t-1})}{u'(\theta,\theta^{t-1})}$ is increasing/constant/decreasing in θ (see proposition 3 in Hellwig (2021)). If the type process is (log) AR(1) with autoregressive parameter ρ (i.e $\frac{\partial\varphi_t(\theta_t|\theta^{t-1})}{\partial\theta^{t-1}} = -\rho\frac{\partial\varphi_t(\theta_t|\theta^{t-1})}{\partial\theta_t}$ and $\rho_t(\theta^t) = \rho$) and the production function is linear in the type (i.e $f_{\theta\theta} = 0$), then we have $\hat{\rho}_t(\theta^t) = \rho$ if the agent has CARA utility. However, if the utility features decreasing absolute risk aversion (DARA)[15], the persistence of the wedges is amplified, i.e. $\hat{\rho}_t(\theta^t) > \rho$.

In the data, we consistently observe a strong lifecycle component in firm dynamics (Evans (1987)). Young firms are usually small and face strong financing constraints. Over time, the firm size tends to increase and financing constraints are relaxed. A cash flow diversion with a risk neutral agent and limited liability (Clementi and Hopenhayn (2006), Fu and Krishna (2019)) can qualitatively replicate the dynamics observed in the data. However, this is no longer the case once we introduce risk aversion and persistent private information. The opposite dynamics emerge, the firm size tends to decrease over time, and the first best size is never reached.

In Appendix C.1, I study a model where the entrepreneur has limited commitment. The firm dynamics induced by this type of model do not change in any meaningful way once risk aversion and persistent private information are introduced. So this type of friction can

[15] Note CRRA utility functions belong to the DARA class.

still generate dynamics where firm size increases over time (as found in Albuquerque and Hopenhayn (2004)[16]). In section 6, I discuss in more detail why models with risk neutrality generate different firm dynamics and their implications.

More generally, the firm lifecycle dynamics are driven by many different frictions. This model could generate more consistent firm dynamics in a straightforward manner by allowing for a drift in the productivity process $\{\theta_t\}$. Then, this type of friction may act as a constraint on the size that firms can eventually reach rather than on the growth of young firms. This may then help explain other empirical facts. For instance, we often observe fewer large firms (Hsieh and Klenow (2014)) in developing economies, where financing frictions are more stringent.

3.2 Compensation dynamics

We now turn to the compensation dynamics. As in all dynamic insurance models, at the optimum, the principal equalizes the cost of increasing the agent's utility at periods t and $t+1$ in an incentive-compatible manner, i.e.

$$\lambda_{t+1}(\theta^t) = \frac{\beta}{q}\xi(\theta^t)$$

where $\lambda_{t+1}(\theta^t)$ is the multiplier on the period $t+1$ promise keeping constraint and $\xi(\theta^t)$ is the multiplier on the type's θ^t period t continuation utility constraint. With separable preferences as in (17), this leads to the well know Inverse Euler Equation $\frac{1}{u'(c(\theta^t))} = \frac{\beta}{q}\mathbb{E}\left[\frac{1}{u'(c(\theta^{t+1}))}|\theta^t\right]$. One cannot derive this tight characterization in all other settings studied in the literature: this includes models with taste shocks (as in Atkeson and Lucas (1992)), hidden endowment (as in Thomas and Worrall (1990)), Mirrlees with non-separable preferences and also this model. For this reason, results on the agent's consumption process are usually derived from the principal's marginal cost martingale (see Golosov *et al.* (2016b)).

The challenge is the same we encountered for characterizing $\widetilde{\mu}_t$. When the principal promises to increase utilities at period $t+1$, this changes all marginal information rents (as they depend on consumption). So utility has to be distributed non-linearly to preserve incentive compatibility. Hellwig (2021) shows how incentive-adjusted measures can be used to derive

[16]Unfortunately, the limited commitment friction alone also does not generate firm dynamics consistent with the data. The issue is that, with i.i.d shocks, firm size is always weakly increasing, so the firm cannot be downsized.

a Generalized Inverse Euler Equation (GIEE). Not surprisingly, we can derive a similar characterization in this model. The GIEE gives an intuitive representation that clarifies what effects drive consumption dynamics and allows to perform direct comparative statics.

Proposition 4. *In the optimal allocation, the following Generalized Inverse Euler Equation holds at any history* θ^t

$$\frac{q}{\beta}\hat{\mathbb{E}}\left[\frac{1}{u'(\theta^{t+1})}|\theta^t\right] = \frac{1}{u'(\theta^t)}(1+s(\theta^t)) \tag{24}$$

where

$$s(\theta^t) = \left(\frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)} - \hat{\mathbb{E}}\left[\rho_{t+1}(\theta^{t+1})\frac{u''(\theta^{t+1})\iota f(\theta^{t+1})}{u'(\theta^{t+1})}|\theta^t\right]\right)\frac{\theta_t}{\mathcal{E}^f(\theta^t)}\tau^k(\theta^t) \tag{25}$$

As in the standard Inverse Euler Equation, the principal arbitrages between the costs of increasing utility at period t and $t+1$, which are proportional to $\frac{1}{u'(\theta^t)}$ and $\frac{1}{u'(\theta^{t+1})}$. However, expectations are taken with respect to the incentive-adjusted probability measure because utility at $t+1$ has to be redistributed non-uniformly to preserve incentive compatibility. Moreover, an extra wedge emerges that captures how savings decisions affect marginal information rents at periods t and $t+1$. Changes in marginal information rents at $t+1$ are passed as a cost at period t at rate $\rho_{t+1}(\theta^{t+1})$. Therefore, the size and sign of the savings wedge depends on the persistence of the process. Intuitively, when the persistence is higher, increasing consumption at $t+1$ lowers the cost of incentive provision at t by more, and so the principal wants relatively higher savings. In general, if persistence (i.e. $\rho_{t+1}(\theta^{t+1})$) is not too high, we will have $s(\theta^t) < 0$, and so savings are on the margin discouraged. The savings wedge is then scaled by the cost of insurance provision at period t.

The savings wedge takes a particularly simple form with CARA utility $u(c) = -e^{-\sigma c}$ with $\sigma > 0$. Assume also an autoregressive process $\rho_t(\theta^t) = \rho$ and $f(k,\theta) = \theta f(k)$, then

$$\begin{aligned} s(\theta^t) &= -\sigma\iota\theta_t \times \tau^k(\theta^t) \times \big(f(k_t) - \rho f(k_{t+1}(\theta^t))\big) \\ &= -\sigma\iota \times \tau^k(\theta^t) \times \big(f(k_t,\theta_t) - \mathbb{E}\left[f(k_{t+1}(\theta^t),\theta_{t+1})|\theta_t\right]\big) \end{aligned}$$

Because $-\sigma\iota\theta_t \times \tau^k(\theta^t) \le 0$, we have $s(\theta^t) \le 0$ if $\rho \le \frac{f(k_t)}{f(k_{t+1}(\theta^t))}$. Moreover, savings are, on the margin, more discouraged when the agent is more risk-averse (higher σ), the costs of diverting funds are small (high ι), and the costs of incentive provision are high (high $\tau^k(\theta^t)$). With fixed capital ($k_t = k$) and $\iota = 1$, this model nests a hidden endowment model[17]. In

[17]With CARA utility, it is also equivalent to a taste shocks model as in Atkeson and Lucas (1992).

this case, $s(\theta^t) = 0$ and we can use the following result.

Proposition 5. *Assume $\frac{q}{\beta} \leq 1$, if $s(\theta^t) \geq 0$ marginal utility follows a super-martingale*

$$u'(\theta^t) \geq \mathbb{E}\left[u'(\theta^{t+1})|\theta^t\right]$$

Moreover, if $s(\theta^t) \geq 0$ for all θ^t, $u' \to 0$ almost surely.

The proposition shows that when $s(\theta^t) \geq 0$ the marginal utility dynamics are preserved under the original measure. Therefore, in a hidden endowment model with a unit root process ($\rho = 1$) there is no immiseration (Thomas and Worrall (1990) and Atkeson and Lucas (1992)), and the contract send the agent to bliss, which is consistent with the results in Bloedel *et al.* (2018) and Bloedel *et al.* (2020)[18]. When $s(\theta^t) < 0$, we do not have direct implications for the dynamics under the original measure. The numerical simulations indicate, as expected, that consumption converges to zero, and so there is immiseration. However, the convergence is very slow, hence these results may be irrelevant for the usual lifespan of a firm.

Compared to a hidden endowment model, time-varying capital generates an extra motive to increase the variance in compensation over time. For the parametric specification above, as long as $\frac{f(k_t)}{f(k_{t+1}(\theta^t))}$ is decreasing in θ_t[19], given some high enough ρ, there can exist a $\widetilde{\theta}_t$ such that $s(\theta^t) < 0$ if $\theta_t \leq \widetilde{\theta}_t$ and $s(\theta^t) \geq 0$ otherwise. Intuitively, because $f_{\theta k} > 0$, higher capital increases information rents. If lower types will have less capital at $t+1$, their incentive constraints will be less tight. Hence, the benefit of increasing consumption at $t+1$ to lower information rents is smaller for lower types.

In sum, the lender minimizes the cost of compensating the agent across periods in an incentive-compatible manner. For this reason, it is optimal to smooth the entrepreneur's compensation over time. Moreover, because the entrepreneur always needs to be compensated for reporting a high cash flow, the variance of consumption will grow without bound.

[18]Bloedel *et al.* (2018) and Bloedel *et al.* (2020) have corrected the findings in Williams (2011) and shown (with more general utility functions and processes) that there is immiseration whenever there is some mean-reversion in the type process.

[19]This would not be the case, if for some types $\theta'_t > \theta''_t$, the effect of higher wedges at $t+1$ for type θ'_t is stronger than from the higher expected productivity. The numerical simulations verify that $k_{t+1}(\theta^t)$ is indeed increasing in θ_t, see Figure 6 in Appendix A.

4 Numerical simulations

In this section, I numerically solve and simulate the model. This will help us better understand the results in the previous section and allow us to quantify the effect of persistent private information on firm size and compensation dynamics. The numerical simulations will also be used to guide the implementation in section 5. I assume the agent has CRRA utility

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}$$

and the production function is given by

$$f(k, \theta) = z\theta k^{\alpha}$$

where $\alpha \in (0, 1)$ and z is a positive constant used to scale up the problem. The agent's productivity follows a geometric AR(1) process

$$\theta_t = \theta_{t-1}^{\rho} \varepsilon_t$$

where $\log(\varepsilon_t) \sim N(\mu, \sigma_\varepsilon^2)$. I set $\alpha = 3/4$, $1 - \iota = 0.05$, $\sigma = 2$ and assume the lender and the entrepreneur have the same discount rate $\beta = q = 0.95$. For the productivity process, I set $\mu = 1$ and $\sigma_\varepsilon^2 = 0.01$. The comparative statics of this section focus on the effect of the persistence ρ, the model is solved with $\rho = 0$ (i.i.d types) and $\rho = 0.7$. I also solve the model with different parametrizations of the utility function (log utility ($\sigma = 1$) and CARA), qualitatively, the results are the same (see Appendix A). Details on the solution method, algorithm and the procedure to check global incentive compatibility can be found in Appendix E. After solving for the value functions (K, v and Δ), the policy functions (c_t, λ_{t+1}, γ_{t+1} and k_{t+1}) and the costate (μ_t), I run a Monte Carlo simulation with 10^6 draws over 25 periods each.

Figure 2 illustrates the evolution of the mean and standard deviation of consumption along the cross-section over time with $\rho = 0$ and $\rho = 0.7$. As expected, the variance of consumption is permanently increasing in both cases. With i.i.d types, average consumption is approximately constant. With persistence, there is also a slight increase in average consumption in the initial periods. Since the savings wedge $s(\theta^t)$ is proportional to the investment wedge $\tau^k(\theta^t)$, this is consistent with the initial increase in the investment wedge that we will observe (see Figure 4). Moreover, because the agent is risk averse, average marginal utility tends to increase over time.

To visualize the immiseration dynamics, in Figure 5 in Appendix A I plot the median and quantiles of the distribution of consumption over a long time horizon. The median consumption monotonically decreases, indicating that consumption will converge to its lower bound. However, the decrease is very slow, so it may be irrelevant for the usual lifespan of a firm.

Figure 2: Consumption dynamics

(a) i.i.d ($\rho = 0$) (b) Persistence ($\rho = 0.7$)

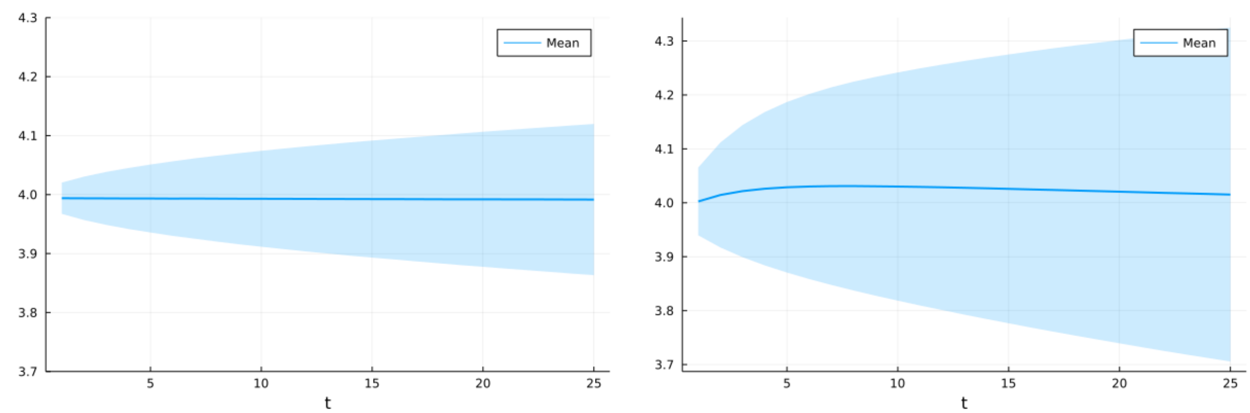

Note: For both panels, I initialize the simulation by setting $\lambda_0 = 15.9$. The initial types, θ_1, are drawn from the ergodic distribution of each process. For each period, the blue line is the mean consumption along the cross-section, and the shaded blue area is one standard deviation.

As discussed, with risk aversion there is a separation between consumption and firm size (or wedge) dynamics. The large differences in consumption resulting from the immiseration dynamics do not translate into differences in firm size distortions. This separation can be illustrated very clearly with the numerical simulations. Figure 3 shows the relation between the promised utility and the investment wedge at age 20. There appears to be some positive association between the two variables, but they are not linked one to one[20]. We can observe that there is some probability that at age 20, the entrepreneur receives a high compensation (high v_t) but that the firm is financially constrained (high τ^k). The converse is also possible, the compensation is low, but the constraints are also low. This is not the case in a model with risk neutrality (Clementi and Hopenhayn (2006)), where the promised utility is linked one-to-one with the distortions on firm size.

[20]This is also the case with i.i.d shocks (trivially because wedges are stationary) and with the other parametrizations of the utility function

Figure 3: Investment wedge and promised utility at $t = 20$ ($\rho = 0.7$)

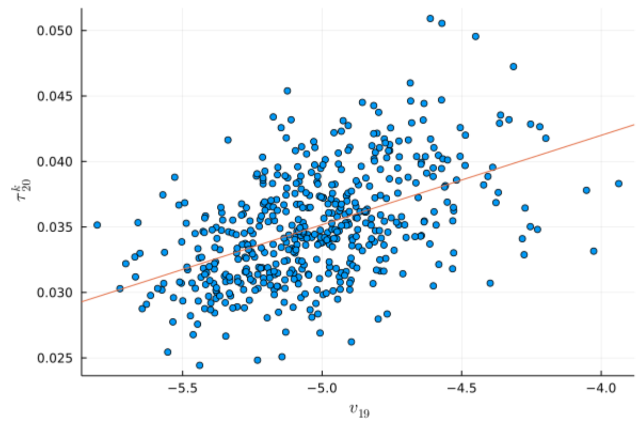

Note: Each dot is a random realization of the investment wedge and promised utility at period 20. The red line is a linear regression line on the 500 draws ploted.

Figure 4 shows the firm size and investment wedge dynamics. In both cases, the firm size closely follows the dynamics of the investment wedge. With i.i.d shocks, the wedges are stationary, so firm size is constant (Panels 4a and 4c). So the firm size dynamics are essentially independent of the consumption dynamics. This is because the lender compensates the entrepreneur by permanently increasing his consumption, not by lending more capital to the firm[21]. The firm's size instead depends on how costly it is to increase information rents. Because the investment wedge is proportional to $\widetilde{\mu}(\theta^t) = \frac{\mu(\theta^t)}{1-\Phi_t(\theta_t|\theta^{t-1})} u'(\theta^t)$, a priori, the marginal utility process could have some effects on wedge dynamics. However, I find that these effects, if any, are negligible. This is also the case with the other parametrizations of the utility function[22]. Moreover, the wedges are small, so firm size is also very close to the first best level.

For the persistent case, at the first best, all variation is driven only by differences in expected returns. Moreover, because the type process is mean-reverting, firm size is stationary. At the second best, on average, the wedges tend to increase over time and firm size tends to decrease (Panels 4b and 4d). However, the wedges do not increase indefinitely. Over time it appears they converge to a stationary distribution, and so does firm size. With log utility (lower risk aversion), the wedges and the decrease in firm size are smaller (see Figure 10 in

[21]This observation also helps explain why, in the implementation, promised utility is linked to the entrepreneur's wealth and unrelated to the firm's capital structure.

[22]This finding also motivates using the normalized shadow cost of insurance ($\widetilde{\mu}$) as the relevant object for the characterization of the wedge dynamics in section 3.1

Appendix A). Overall the decrease in firm size will be larger the higher the risk aversion and persistence.

Figure 4: Firm size and investment wedge dynamics

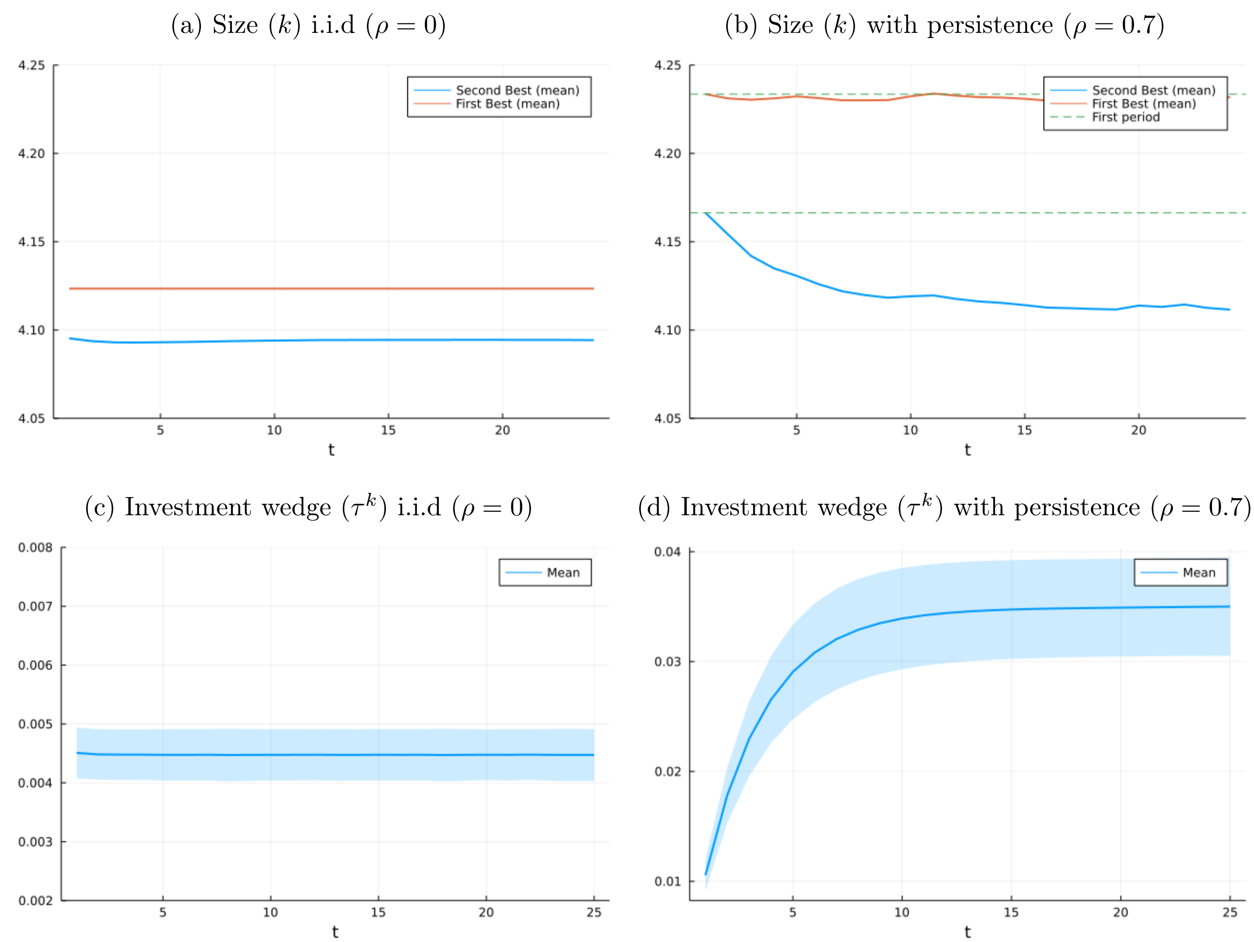

(a) Size (k) i.i.d ($\rho = 0$)

(b) Size (k) with persistence ($\rho = 0.7$)

(c) Investment wedge (τ^k) i.i.d ($\rho = 0$)

(d) Investment wedge (τ^k) with persistence ($\rho = 0.7$)

Note: Panel (a): The red line is the size at the first best (constant). The blue line is the average size at the second best; it is the same for almost all realizations as expected wedges are approximately constant. Panel (b): The red (blue) line is the average size in the first (second) best. The dashed green lines are the average size in the first period. Panels (c) and (d): The blue lines are the average investment wedges, and the shaded blue areas are one standard deviation.

5 Quasi-implementation

The optimal contract studied thus far may a priori be complex, which limits the insights we can derive from the problem. Therefore, it is helpful to study implementations of the optimal contract. The implementation will also provide intuitions on what drives the different firm size dynamics in models with risk neutrality and risk aversion. A full implementation of the optimal contract is challenging and left for future work. For this reason, I analyze two simpler problems. In this section, I use numerical simulations to study a quasi-implementation with simpler contracts. Then, in Appendix D, I study a full implementation in a simplified two-period version of the model.

The two implementations are also different in the sense that they are based on two different margins of distortion for the lender. In particular, we consider two ways in which the principle can lower Markov information rents to better screen types. In the first one, the lender dynamically distorts the entrepreneur's compensation. Given the distortions to compensation, she may then optimally increase or decrease the firm size (k_{t+1})[23]. In the second, the lender directly distorts firm size by discouraging the entrepreneur from reinvesting profits.

The approach to deriving the quasi-implementation with the numerical simulations will be the following. First, I use regressions with the model simulated data to better understand the compensation dynamics. Then, I propose a simple contract and use the simulated data and regression estimates to calibrate the parameters of the contract. Finally, I solve the entrepreneur's problem under the simple contract and compare the induced consumption dynamics with the optimal contract. With i.i.d types, firm size is constant, so I also fix capital to be constant in the implementation. For simplicity, I will also keep capital fixed for the persistent case. Therefore, we will focus solely on the compensation dynamics.

5.1 i.i.d types

I use the simulated data from section 4 to run regressions of consumption on returns and promised utility. The regression results are in Table 1; we make three observations:

1. Variation in returns at any period $t - k$ has the same effect as returns at t on consumption at t (column 2). Relatedly, consumption follows a random walk (column 5). Suggests that compensation is perfectly smoothed across periods.

2. The sensitivity of compensation to returns does not depend on current promised utility. Note the interaction $returns_t \times v_{t-1}$ is close to 0 in column 3.

[23] In this sense, the cash flow diversion model is equivalent to a hidden endowment model (as in Thomas and Worrall (1990)) where the principal has some control over the agent's income process.

3. The sensitivity of compensation to returns is close to linear. Note in column $returns_t^2$ is close to 0 in column 4.

Table 1: Regressions with i.i.d type process

	(1)	(2)	(3)	(4)	(5)
	c_t	c_t	c_t	c_t	c_t
$returns_t$	0.0489***	0.0511***	0.0508***	0.0704***	
	(10558.17)	(399.15)	(108.58)	(926.10)	
v_{t-1}	0.792***		0.790***	0.792***	
	(14968.42)		(1459.73)	(15010.68)	
$returns_{t-5}$		0.0490***			
		(382.92)			
$returns_t * v_{t-1}$			0.000386***		
			(4.13)		
$returns_t^2$				-0.00185***	
				(-283.17)	
c_{t-1}					0.998***
					(5300.13)
N	2400000	1900000	2400000	2400000	2300000
R^2	0.999	0.139	0.999	0.999	0.924

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Points 2. and 3. suggest that a constant equity share can be a good approximation. If the promised utility (v_{t-1}) were related to the equity share, we would observe that it affects the sensitivity of consumption to returns, even if the entrepreneur is smoothing consumption[24]. Point 1. indicates that in the implementation, the entrepreneur's implicit wealth can be used to perfectly smooth consumption intertemporally. As is known, the promised utility can be naturally mapped to the agent's wealth (Atkeson and Lucas (1992), Brendon (2022)). Let W_t denote the agent's wealth and χ the (inside) equity share, i.e. the portion of cash flows accruing to the entrepreneur. Let $\overline{f}(k_t) = \mathbb{E}\left[f(k_t, \theta_t)\right]$ denote the expected returns if capital is k_t. I fix capital to the optimum in the second best k_{SB}. The entrepreneur also receives

[24]This observation is key to understanding the differences with risk neutrality, where the promised utility maps to the value of equity (Clementi and Hopenhayn (2006)). I discuss this distinction in more detail in section 6.

initial cash W_0[25]. Therefore, at period 1, the entrepreneur's wealth is

$$W_1 = W_0 + \frac{\chi \overline{f}(k_{SB})}{1-q}$$

At every period, after returns realized, if the entrepreneur does not misreport, his wealth changes by $\chi\left(f(k_{SB},\theta_t) - \overline{f}(k_{SB})\right)$. So the LOM of the entrepreneur's wealth follows

$$c_t + W_{t+1} = \frac{1}{q}W_t + \chi\left(f(k_{SB},\theta_t) - \overline{f}(k_{SB})\right) \equiv C(W_t,\theta_t) \tag{26}$$

Given the entrepreneur's wealth, savings can be chosen to smooth consumption. Therefore, this contract is equivalent to allowing the entrepreneur to pledge his shares as collateral and borrow to consume. This practice is prevalent; Fabisik (2019) reports that between 2007 and 2016, 7.6% of CEOs of US public companies had pledged shares. Moreover, she estimates that 90.5% of CEOs use it to obtain liquidity while maintaining ownership. This motive is consistent with this implementation. Pledging shares aligns the entrepreneur's consumption with the firm's value but without having to sell shares, which is costly as it reduces the entrepreneur's incentives. Moreover, the implementation is independent of dividend payout policies. Notice that it is equivalent if the extra returns $\left(f(k_{SB},\theta_t) - \overline{f}(k_{SB})\right)$ are paid as dividends or are kept as savings inside the firm, and the entrepreneur and the firm face the same interest rate q.

The next step for the numerical implementation is to obtain a value for χ. I back out this value from the regressions on model simulated data. For an entrepreneur that does not misreport and is allowed to save by himself, to a first order approximation, we have

$$\frac{dc_t}{df(k_t,\theta_t)} \approx (1-q)\chi$$

So χ can be identified from the regressions as $\hat{\chi} = \frac{\beta_{returns}}{(1-q)} = \frac{0.0488}{0.05} \approx \iota$, where $\beta_{returns}$ is the regression coefficient on returns in column (1) of Table 1. So I set directly $\hat{\chi} = \iota$[26]. Then,

[25] This is just a free variable used to match the chosen initial promised utility in the second best, so we may have also have $W_0 < 0$ if the entrepreneur initially transfers funds to the lender.

[26] It is a regular result in cash flow diversion models (especially in static versions) that the equity share is linked to the deadweight loss of diverting funds.

given $\hat{\chi}$, the entrepreneur recursive problem with wealth W_t and productivity θ_t is

$$\begin{aligned} \mathcal{W}(W_t, \theta_t) &= \max_{\tilde{\theta} \leq \theta} u(\tilde{c}_t) + \beta \mathcal{V}(W_{t+1}) \\ s.t \qquad W_{t+1} &= qC(W_t, \tilde{\theta}_t) \\ c_t &= (1-q) C(W_t, \tilde{\theta}_t) \\ \tilde{c}_t &= c_t + \iota(f(k_{SB}, \theta) - f(k_{SB}, \tilde{\theta})) \end{aligned} \tag{27}$$

Where $\mathcal{V}(W_{t+1}) = \mathbb{E}\left[\mathcal{W}(W_{t+1}, \theta_{t+1})\right]$, $C(W_t, \theta_t) = \frac{1}{q} W_t + \hat{\chi}\left(f(k_{SB}, \tilde{\theta}_t) - \overline{f}(k_{SB})\right)$ and W_0 is chosen such that $\mathcal{V}(W_1) = v_1$, i.e. the promised utility under the direct mechanism. Throughout the paper, I have assumed that the entrepreneur cannot secretly save. So in the implementation, there is a double deviation problem if the entrepreneur is allowed to save freely. That is, the entrepreneur deviates by misreporting funds and saving more. For this reason, I assume that the lender directly assigns a consumption/savings level given the entrepreneur's report and wealth $(W_t, \tilde{\theta}_t)$[27]. Equivalently, we can imagine that the entrepreneur is penalized if the lender observes that his savings choices are not optimal given the reported type and wealth.

I solve numerically for the policy functions $\tilde{\theta}(W_t, \theta_t)$ in the entrepreneur's problem (27). Then, I run the same Monte Carlo simulation as for the optimal allocation and compare the results[28]. Figure 9 in Appendix A shows that the consumption paths are very close to the optimal allocation and that this contract induces minimal diversion of funds. Not surprisingly, this simple contract also reaps most of the benefits of the optimal allocation (see Table 2). Given a fixed initial promised utility (v_0), we can decompose the lender's loss

[27] I assign the consumption to be $c_t = (1-q)C(W_t, \tilde{\theta}_t)$ because I observe that average consumption is approximately constant in the numerical simulations. But this is not the optimal savings level of the entrepreneur, as he would save more for precautionary motives. To relax this restriction, we could introduce an extra wedge (or tax) on the entrepreneur's returns on savings to exactly counteract the precautionary motive.

[28] To have accurate comparisons, in the Monte Carlo simulation, for each realization of the shock process $\{\varepsilon_t\}_{t=1}^{25}$ I compute consumption for both the optimal allocation and the implementation. Then for each realization and period, I compute the distance and average across all draws. That is, I compute for every period $\overline{c}_t^{dist} = \sum_i \sqrt{\left(c_t^{SB}(\{\varepsilon_{i.\tau}\}_{\tau=1}^t) - c_t^I(\{\varepsilon_{i.\tau}\}_{\tau=1}^t)\right)^2}$, where c^{SB} is the consumption under the optimal allocation and c^I under the implementation.

Table 2: Welfare comparisons i.i.d

	Total Welfare	Deadweight loss diversion of funds	Risk premium (relative to SB)
Optimal contract (SB)	-55.88	0	0
Quasi-Implementation ($\chi = 0.95$)	-56.11 (-0.4% loss)	5.7e-8	0.22

from using the simple contract

$$K^I(v_0,\theta_0) - K^{SB}(v_0,\theta_0) =$$
$$\underbrace{(1-\iota)\mathbb{E}\left[\sum_{t=0}^{\infty} q^t \left(f(k_{SB},\theta^t) - f(k_{SB},\widetilde{\theta}_t(\theta^t))\right)|\theta_0\right]}_{\geq 0\text{,Deadweight loss diversion of funds}} + \underbrace{\mathbb{E}\left[\sum_{t=0}^{\infty} q^t \left(c^I(\theta^t) - c^{SB}(\theta^t)\right)|\theta_0\right]}_{\text{Risk premium,}>0\text{ if less risk in SB}} \quad (28)$$

where the superscript I is used to denoting allocations under the implementation. As shown in Table 2, most of the losses from the simple contract result from exposing the entrepreneur to more risk, but the differences are negligible. The implementation performs even better with log utility, see Figure 9 in Appendix A[29].

5.2 Persistent types

With persistent private information, there is an extra state variable in the recursive planning problem (12), Δ_{t-1}. Intuitively, this state variable captures how much insurance (or incentives) the principal has promised to provide to the agent, as equation (11) can be written as

$$\Delta_{t-1} = \mathbb{E}\left[\rho(\theta^t)\frac{\partial w(\theta^t)}{\partial \theta_t}|\theta^{t-1}\right] \quad (29)$$

For a given level of persistence, a lower Δ_{t-1} implies more insurance is provided to the agent. We can also verify this in the regressions with model simulated data, where we obtain

$$c_t = -0.2^{***}\theta_t - 3.327^{***}\Delta_{t-1} + 1.011^{***}\theta_t \times \Delta_{t-1} + 0.652^{***}\theta_{t-1} + 0.382^{***}v_t$$

The coefficient on the interaction term $\Delta_{t-1} \times \theta_t$ is positive. So when the lender has promised high insurance (i.e low Δ_{t-1}), the entrepreneur's compensation is less sensitive to the type

[29]With log utility average consumption in the optimal contract is exactly constant, so the savings choices imposed in the implementation give a better approximation.

realization. In this implementation, the level of insurance provided to the entrepreneur is controlled by the equity share. To see this, notice that for the i.i.d case (problem (27)), if it is optimal for the entrepreneur to not divert funds, we have

$$\frac{\partial \mathcal{W}(W_t, \theta_t)}{\partial \theta_t} = \chi \times u'(c_t) f_\theta(k_{SB}, \theta_t)$$

Thus, a full implementation of the optimal contract would need to allow for a time-varying equity share. In general, lowering the entrepreneur's equity is beneficial as it increases insurance, but it also comes at the cost of increasing the entrepreneur's incentives to misreport funds. If types are persistent, there is an extra gain of lowering the equity share at period $t+1$ because it helps screen types.

Why does buying equity help screen types? Imagine that, at period t, the lender offers to buy some equity from type (θ^{t-1}, θ'). Assume also that the lender offers to pay him a price $P_{\Delta\chi}((\theta^{t-1}, \theta'))$ such that he is indifferent between accepting the offer or rejecting it. If returns are persistent, types (θ^{t-1}, θ'') with $\theta'' > \theta'$ have higher expected returns at period $t+1$. So it is not attractive for them to sell equity at price $P_{\Delta\chi}((\theta^{t-1}, \theta'))$. Therefore, the lender can use equity purchases, which inefficiently lower the equity share, to better screen types.

More formally, this intuition is related to the Atkinson and Stiglitz (1976) result for commodity taxation. With i.i.d shocks, less productive entrepreneurs are also more willing to sell equity as they have higher marginal utility. But in this case, the willingness to sell equity does not reveal any information to the lender that is not already contained in reported returns. With persistence, lower types would be more willing to sell equity even if they had the same marginal utility as higher types[30]. So the lender optimally distorts the equity share as it directly reveals information about the entrepreneur's productivity.

An implementation with a time-varying equity share is substantially more challenging[31]. However, I find that the contract with a constant equity share still delivers small welfare losses relative to the optimal contract[32]. Compared to the i.i.d case, we now only have to make

[30]With risk neutrality, there is no insurance motive in the contract. Because there is no cost of exposing the agent to more risk, the entrepreneur's future compensation is provided only through equity. So the equity share is only constrained by the current promised utility Clementi *et al.* (2010).

[31]Now it is more challenging to infer the equity share from the regressions directly. Moreover, it may follow a complicated stochastic process. As it would be persistent but also because there is no distortion at the top $(\overline{\theta})$ and bottom $(\underline{\theta})$ in the promised insurance.

[32]I have experimented with contracts where the equity share is uniformly decreased over time for all types. The idea is that when the agent underreports at t, he experiences a capital loss but expects to recover it at

one modification. Notice that when the entrepreneur's period t returns increase, the net present value of the firm's cash flows also increases. So the firm's value increases and the entrepreneur experiences a capital gain. Define the value of the firm by

$$\overline{f}_{t+1}(k_{SB}, \theta_t) \equiv \mathbb{E}\left[\sum_{\tau=1}^{\infty} q^{\tau-1} f(k_{SB}, \theta_{t+\tau})|\theta_t\right]$$

Recall capital is fixed to the same level k_{SB} as in the i.i.d case. Then, the entrepreneur's cash on hand at period t if he reports type $\widetilde{\theta}_t$ and his past type report was $\widetilde{\theta}_{t-1}$ is

$$\begin{aligned} C(W_t, \widetilde{\theta}_t, \widetilde{\theta}_{t-1}) &= \frac{1}{q} W_t + \chi\left(f(k_{SB}, \widetilde{\theta}_t) + q\overline{f}_{t+1}(k_{SB}, \widetilde{\theta}_t) - \overline{f}_t(k_{SB}, \widetilde{\theta}_{t-1})\right) \\ &= \frac{1}{q} W_t + \chi\left(\underbrace{f(k_{SB}, \widetilde{\theta}_t) - \mathbb{E}\left[f(k_{SB}, \widetilde{\theta}_t)|\widetilde{\theta}_{t-1}\right]}_{\text{Innovation returns}} + \underbrace{q\left(\overline{f}_{t+1}(k_{SB}, \widetilde{\theta}_t) - \overline{f_{t+1}}(k_{SB}, \widetilde{\theta}_{t-1})\right)}_{\text{Capital gain}}\right) \end{aligned} \tag{30}$$

The entrepreneur's problem is the same as in (27) but with the cash on hand given (30). Because the entrepreneur can borrow using his shares as collateral, the capital gains also increase the entrepreneur's consumption. Moreover, we also need to keep track of the past report $\widetilde{\theta}_{t-1}$ as an extra state variable because it affects the expected returns. Table 3 contains the welfare comparison and decompositions with the optimal contract. The risk premium is higher than for the i.i.d case, but the welfare losses from the simple contract continue to be small[33].

Fu and Krishna (2019) study a similar model with persistent private information and a risk neutral entrepreneur. They show that as persistence increases, the convexity of the entrepreneur's compensation also increases. In their implementation, this implies that the entrepreneur is compensated more with stock options and less with equity. A priori, a full implementation of the optimal contract could require the use of stock options. But only with

$t+1$ with returns that are higher than expected. However, if his equity share is lower at $t+1$, he cannot fully recover the capital loss. However, I have not found any gains from this type of contracts.

[33] In general, the optimal contract studied is not renegotiation-proof. When the principal lowers Δ_t at period t, this is inefficient from the period $t+1$ perspective. Because these effects are not present with i.i.d types, the optimal contracts are renegotiation-proof. I conjecture that this quasi-implementation with constant equity could be a good approximation to the optimal renegotiation-proof contract with persistent private information. The constant equity implies constant sensitivity as the optimal renegotiation proof contracts in Strulovici (2022). Moreover, as in Strulovici (2022) the sensitivity is increasing in the level of persistence because capital gains are larger. Because the constant equity contract delivers small welfare losses compared to the optimal, I further conjecture that, at least for the calibration used, the losses from restricting to renegotiation-proof contracts may be small.

Table 3: Welfare comparison with persistence

	Total Welfare	Deadweight loss diversion of funds	Risk premium (relative to SB)
Optimal contract (SB)	-58.49	0	0
Quasi-Implementation ($\chi = 0.95$)	-59.12 (-1.07% loss)	3e-3	0.572

equity the entrepreneur's compensation is already convexified. Because the entrepreneur experiences a capital gain and can borrow using his shares as collateral, his compensation increases more than linearly in returns. I conjecture that accounting for capital gains could be sufficient to generate the sensitivity of consumption to returns required to implement the optimal allocation. In fact, by expanding the dynamic (or Markov) information rent term (equation (29)), we can write

$$\Delta_t = \iota \times \mathbb{E}\left[\sum_{j=1}^{\infty} I_t^{t+j}(\theta^{t+j})\beta^{j-1}u'(c(\theta^{t+j}))f_\theta(k_{sb},\theta_{t+j})|\theta^t\right]$$

where $I_t^{t+j}(\theta^{t+j}) \equiv \rho(\theta^t) \times ... \times \rho(\theta^{t+j})$ are the impulse response functions as defined in Pavan *et al.* (2014). With $\chi = \iota$ and fixed capital, this is actually the capital gain of the entrepreneur from an increase in productivity $d\theta_t$ if the firm is priced with the entrepreneur's stochastic discount factor.

6 Comparison with risk neutral and equity dynamics

The quasi-implementation helps understand the different firm size dynamics with risk neutrality and risk aversion. With risk neutrality, as long as the limited liability constraint is satisfied, increasing the agent's exposure to risk bears no cost. After high returns, it is optimal to compensate the entrepreneur with a higher stake in the project, i.e. by increasing his equity share. Therefore, with risk neutrality, the entrepreneur's promised utility maps to the value of equity, as shown in Clementi and Hopenhayn (2006).

If the entrepreneur is risk averse, increasing his exposure to risk through a higher equity share is costly. In the numerical simulations, we have seen that the entrepreneur's exposure to returns is independent of his promised utility. So with i.i.d types, a constant equity share and mapping the entrepreneur's promised utility to his private wealth gives a good approximation to the optimal allocation. With persistent types, the equity share should also be time-varying as in the risk neutral model, but the driving forces are different. With persistence, the lender has an incentive to lower equity below the efficient level at $t+1$ as it helps screen types at period t. Hence, over time, the equity share of the entrepreneur tends to decrease. Then, when the equity share is low, the entrepreneur has more incentives to divert funds, so the lender is less willing to advance capital.

Consequently, both models obtain a positive relation between equity and firm size. Lower equity always increases the implicit lending costs because the incentives to divert funds are higher. However, equity drifts in opposite directions. With risk neutrality, equity drifts upwards, but with risk aversion and persistence, equity drifts downwards. With risk neutrality and i.i.d types, firm size converges to the first best level only because the entrepreneur's equity share goes to one (Clementi and Hopenhayn (2006)). That is, he becomes the sole owner of the firm, and the value of debt and outside equity go to zero. With persistent types and risk neutrality, the equity share does not necessarily have to converge to one for the firm's size to reach the first best (Fu and Krishna (2019)). However, the equity share still grows over time as the firm size grows. These equity dynamics may be inconsistent with what is observed in the data. For example, in the venture capital industry, the founder's ownership is typically diluted over time as the firm's capital grows through multiple financing rounds (Azevedo *et al.* (2023)). Accordingly, to simultaneously explain firm size and equity dynamics, it may be necessary to break the tight link between equity and firm size that these models generate.

Modigliani-Miller and promised utility: The distinction between wealth and equity also helps understand other differences with the risk neutral case. One of them concerns how the firm's value depends on the capital structure. With risk neutrality, the firm's value depends on the value of equity (or promised utility), so the Modigliani-Miller theorem does not hold (Clementi and Hopenhayn (2006)). Interestingly, we have found in the numerical simulations that, with risk aversion, the lending policy functions k_{t+1} are approximately independent of v_t. Therefore, the firm's value is approximately independent of the initial promised utility given to the entrepreneur. So in this sense, Modigliani-Miller does hold "over" promised utility. However, because lending depends on promised insurance, the firm's value will vary with promised insurance. This observation corroborates the idea of the implementation with risk aversion presented in the previous section. The promised utility does not map properly to the entrepreneur's equity; instead, it maps to the entrepreneur's private wealth. Moreover, it is promised insurance what maps to the entrepreneur's equity share and so what affects the firm's capital structure.

7 Extensions

The problem studied throughout the paper is the simplest version of a cash flow diversion model with persistent private information and risk aversion. To focus on the role of persistence and risk aversion, I have imposed some assumptions and abstracted from other interesting margins. I explore three extensions in Appendix C.

Limited commitment: In Appendix C.1, I relax the assumption of full commitment of the entrepreneur. At every period, the entrepreneur can steal all the capital advanced by the lender and leave the contract. This friction also lowers firm size but adds an incentive to have promised utilities increase over time. So when the limited commitment is binding, it can generate dynamics where firm size increases over time, as in Albuquerque and Hopenhayn (2004), even with risk aversion and persistence.

Endogenous termination: I have assumed that the lender does not have the option to terminate the project. I discuss endogenous termination in Appendix C.2. Although termination may sometimes be optimal, I show that it does not affect any of the results presented. I also discuss what inefficiencies may cause termination with risk aversion and persistence: too high promised insurance or a combination of private information and limited commitment, similar to Dovis (2019). In both cases, the termination probabilities should tend to increase over time. This is the opposite of the model with risk neutrality, where the exit rates tend to decrease over time.

Moreover, if at some point terminating with some probability $\alpha_t(\theta^t) \in (0, 1)$ is optimal, $\alpha_t(\theta^t)$ should increase with the persistence of the shocks. I show this in a simplified two-period and two-type version of the model. The intuition is similar to the equity purchases. Imagine that the principal increases the termination probability of type θ^t while increasing his payment after termination such that his ex-ante utility is held constant. With persistence, types (θ^{t-1}, θ') with $\theta' > \theta_t$ know they have higher expected returns at $t+1$ than θ_t, so a higher termination probability is relatively less attractive for them. Hence, a higher termination probability can discourage misreporting for types $\theta' > \theta_t$, and so it lowers the cost of screening types.

Divert funds before investing (screening): Finally, in Appendix C.3, I study a model where instead of diverting cash flows, the entrepreneur can choose the fraction of available funds invested in the firm and divert the rest. Then the lender can observe the project returns but not invested funds. The investment wedge can now be defined as the distortion to invested and diverted funds relative to the first best. Moreover, this model yields the same characterizations of the shadow cost of insurance, the GIEE, and the firm size dynamics.

8 Conclusion

This paper studied a dynamic cash flow diversion model with a risk averse entrepreneur that has persistent private information about the firm's productivity. The firm size and compensation dynamics differ significantly from models with risk neutrality. Most notably, firm size tends to decrease over time. The implementation helps understand the opposite size dynamics. Regardless of the entrepreneur's preferences, capital is increasing in the equity share. However, equity drifts upwards with risk neutrality and downwards with risk aversion and persistence. These findings suggest that it may be challenging for these type of models to generate realistic firm size and equity dynamics. With risk neutrality and i.i.d types (Clementi and Hopenhayn (2006)), firm size converges to the first best only because the entrepreneur's equity share goes to one. As discussed, these dynamics may be at odds with what we observe empirically (Azevedo *et al.* (2023)).

References

Albuquerque, R. and Hopenhayn, H. A. (2004). Optimal lending contracts and firm dynamics. *The Review of Economic Studies*, **71** (2), 285–315.

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

Atkinson, A. B. and Stiglitz, J. E. (1976). The design of tax structure: direct versus indirect taxation. *Journal of public Economics*, **6** (1-2), 55–75.

Azevedo, E., Scheuer, F. and Smetters, K. (2023). The returns to high-powered entrepreneurship.

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

—, —, Rochet, J.-C. and Villeneuve, S. (2010). Large risks, limited liability, and dynamic moral hazard. *Econometrica*, **78** (1), 73–118.

Bloedel, A., Krishna, R. and Strulovici, B. (2020). *Persistent private information revisited*. Tech. rep., Technical report, Stanford University.

Bloedel, A. W., Krishna, R. V. and Leukhina, O. (2018). Insurance and inequality with persistent private information. *Available at SSRN 3091457*.

Brendon, C. (2013). Efficiency, equity, and optimal income taxation.

— (2022). Optimal nonlinear savings taxation.

Clementi, G. L., Cooley, T. F. and Di Giannatale, S. (2010). A theory of firm decline. *Review of Economic Dynamics*, **13** (4), 861–885.

— and Hopenhayn, H. A. (2006). A Theory of Financing Constraints and Firm Dynamics. *The Quarterly Journal of Economics*, **121** (1), 229–265.

DeMarzo, P. M. and Fishman, M. J. (2007a). Agency and optimal investment dynamics. *The Review of Financial Studies*, **20** (1), 151–188.

— and — (2007b). Optimal long-term financial contracting. *The Review of Financial Studies*, **20** (6), 2079–2128.

—, —, He, Z. and Wang, N. (2012). Dynamic agency and the q theory of investment. *The Journal of Finance*, **67** (6), 2295–2340.

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

— and — (2016). Learning, termination, and payout policy in dynamic incentive contracts. *The Review of Economic Studies*, **84** (1), 182–236.

Di Tella, S. and Sannikov, Y. (2021). Optimal asset management contracts with hidden savings. *Econometrica*, **89** (3), 1099–1139.

DIAMOND, P. A. (1998). Optimal income taxation: an example with a u-shaped pattern of optimal marginal tax rates. *American Economic Review*, pp. 83–95.

DOVIS, A. (2019). Efficient sovereign default. *The Review of Economic Studies*, **86** (1), 282–312.

EVANS, D. S. (1987). The relationship between firm growth, size, and age: Estimates for 100 manufacturing industries. *The journal of industrial economics*, pp. 567–581.

FABISIK, K. (2019). Why do us ceos pledge their own company's stock? *Swiss Finance Institute Research Paper*, pp. 19–60.

FARHI, E. and WERNING, I. (2013). Insurance and taxation over the life cycle. *Review of Economic Studies*, **80** (2), 596–635.

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

FU, S. and KRISHNA, R. V. (2019). Dynamic financial contracting with persistent private information. *The RAND Journal of Economics*, **50** (2), 418–452.

GOLOSOV, M., TROSHKIN, M. and TSYVINSKI, A. (2016a). Redistribution and social insurance. *American Economic Review*, **106** (2), 359–86.

—, TSYVINSKI, A. and WERQUIN, N. (2016b). Recursive contracts and endogenously incomplete markets. In *Handbook of Macroeconomics*, vol. 2, Elsevier, pp. 725–841.

HE, Z. (2012). Dynamic compensation contracts with private savings. *The Review of Financial Studies*, **25** (5), 1494–1549.

HELLWIG, C. (2021). Static and dynamic mirrleesian taxation with non-separable preferences: A unified approach.

— and WERQUIN, N. (2022). A fair day's pay for a fair day's work: Optimal tax design as redistributional arbitrage.

HSIEH, C.-T. and KLENOW, P. J. (2014). The life cycle of plants in india and mexico. *The Quarterly Journal of Economics*, **129** (3), 1035–1084.

KAPIČKA, M. (2013). Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach. *Review of Economic Studies*, **80** (3), 1027–1054.

Krasikov, I. and Lamba, R. (2021). A theory of dynamic contracting with financial constraints. *Journal of Economic Theory*, **193**, 105196.

Makris, M. and Pavan, A. (2020). Wedge dynamics with evolving private information.

Ndiaye, A. (2020). *Flexible Retirement and Optimal Taxation*. Tech. rep.

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

Stantcheva, S. (2017). Optimal taxation and human capital policies over the life cycle. *Journal of Political Economy*, **125** (6).

— (2020). Dynamic taxation. *Annual Review of Economics*, **12**, 801–831.

Strulovici, B. (2022). Renegotiation-proof contracts with persistent states.

Thomas, J. and Worrall, T. (1990). Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory*, **51** (2), 367–390.

Williams, N. (2011). Persistent private information. *Econometrica*, **79** (4), 1233–1275.

Appendix

A Additional tables and figures

Figure 5: Immiseration in the long run

(a) Marginal utility i.i.d ($\rho = 0$)

(b) Marginal utility with persistence ($\rho = 0.7$)

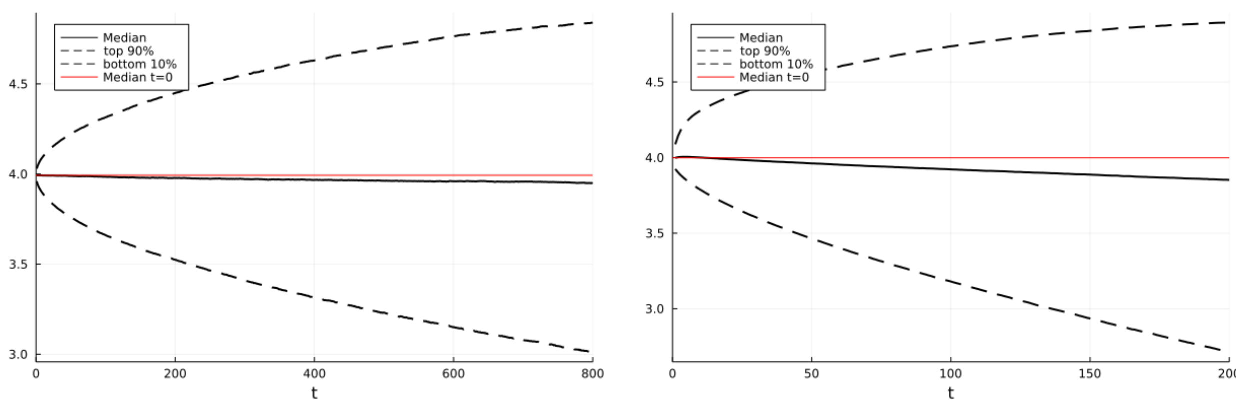

Note: The figures show the median, 10%, and 90% quantiles of the distribution of consumption at every period. For reference, the red line displays the median at period $t = 0$. The median monotonically decreases and the growth of the 90% quantile decreases over time. This implies that consumption will converge to its' lower bound. However, we also observe that this convergence is very slow.

Figure 6: Relation $k_{t+1}(\theta_t)$ and θ_t

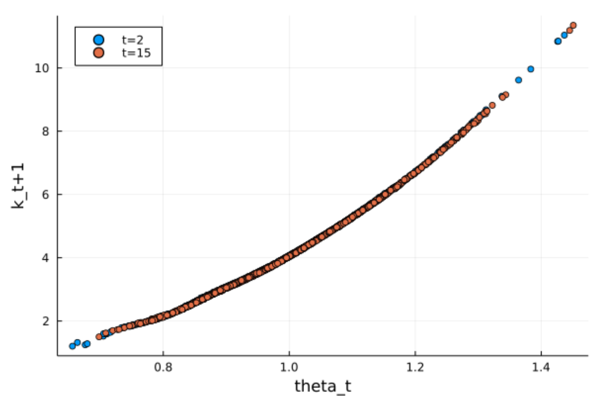

Note: For a random subsample of 1000 realizations, the plot shows the policy functions of k_{t+1} as a function of θ_t. The blue dots are policies at period $t = 2$, and the red dots at $t = 15$.

Figure 7: Shadow cost insurance μ at different γ

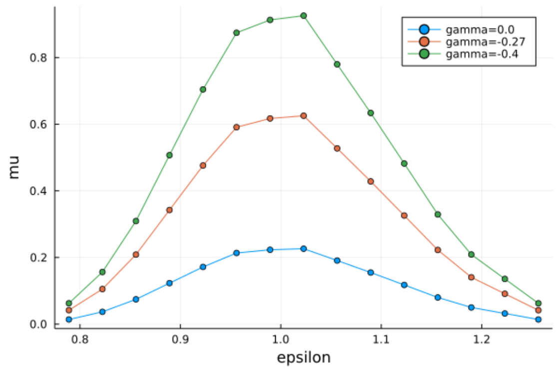

Note: For a fixed (λ_-, k, θ_-), the figure shows the shadow cost of insurance μ as a function of the shock ε for different γ_-. The promised insurance Δ_- is increasing in γ_-. So when the agent is promised more insurance (low Δ_- and γ_-), the shadow costs μ are higher. The increase is more pronounced for the types in the middle.

Figure 8: Simulations implementation i.i.d

(a) Average distance consumption SB and implementation

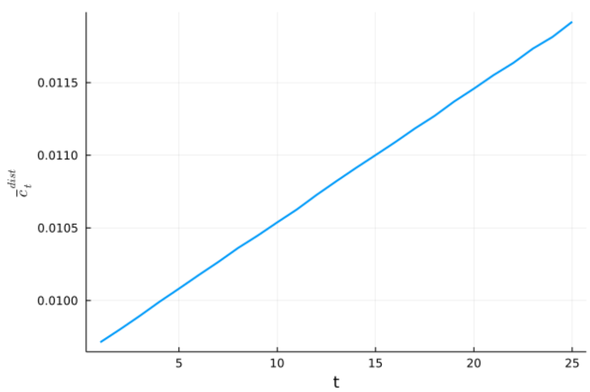

(b) Average fraction of diverted funds

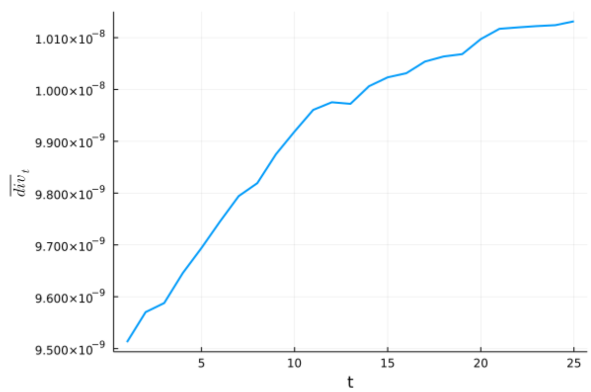

Note: The left figure shows, for every period, the average distance between consumption in the optimal contract (c^{SB}) and the implementation (c^I), i.e. $\bar{c}_t^{dist} = \frac{1}{N}\sum_i \sqrt{(c_t^{SB}(\{\varepsilon_{i,\tau}\}_{\tau=1}^t) - c_t^I(\{\varepsilon_{i,\tau}\}_{\tau=1}^t))^2}$. The right figure shows the average of the diverted funds as a fraction of total returns, i.e. $\overline{div}_t = \frac{1}{N}\sum_i \frac{f(k_{SB},\theta_i^t) - f(k_{SB},\widetilde{\theta}_t(\theta_i^t))}{f(k_{SB},\theta_i^t)}$

Figure 9: Simulations implementation i.i.d with log utility ($\sigma = 1$)

(a) Average distance consumption SB and implementation

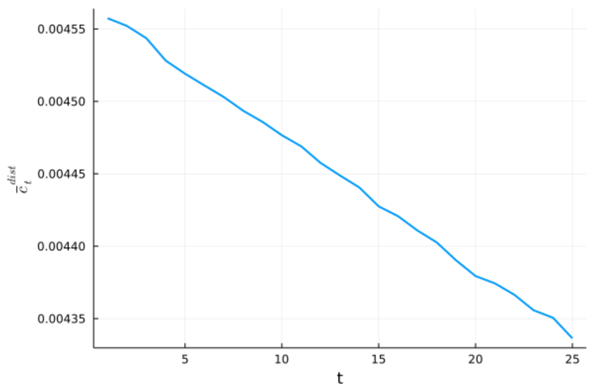

(b) Average fraction of diverted funds

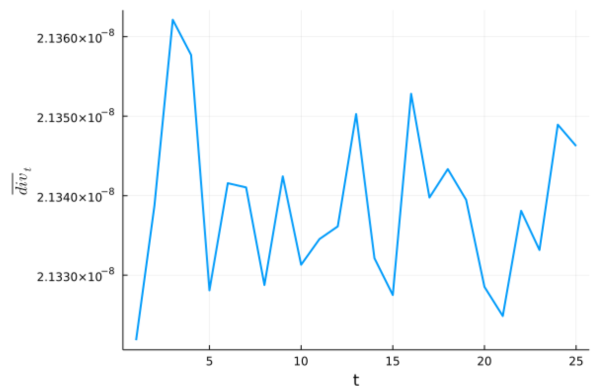

Figure 10: Log utility ($\sigma = 1$)

(a) Consumption i.i.d ($\rho = 0$)

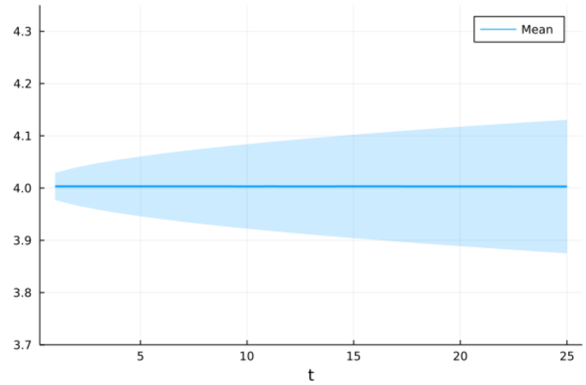

(b) Consumption with persistence ($\rho = 0.7$)

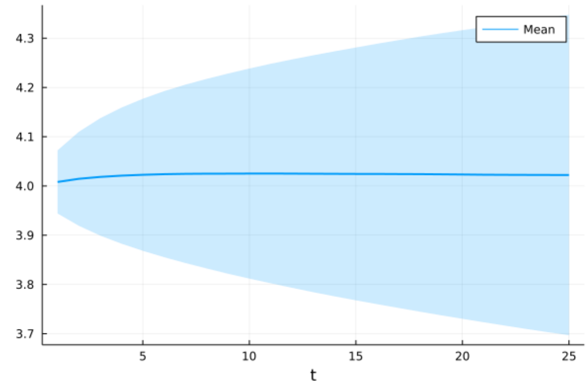

(c) Firm size i.i.d ($\rho = 0$)

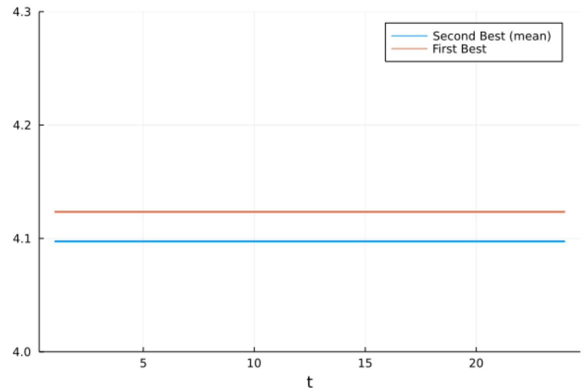

(d) Firm size with persistence ($\rho = 0.7$)

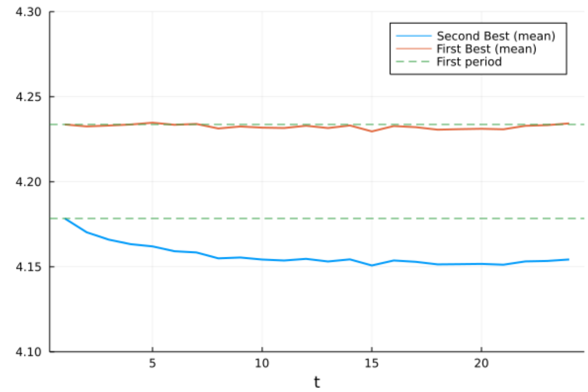

(e) Investment wedge i.i.d ($\rho = 0$)

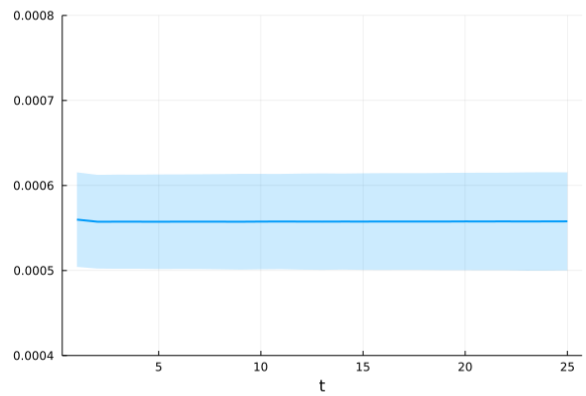

(f) Investment wedge with persistence ($\rho = 0.7$)

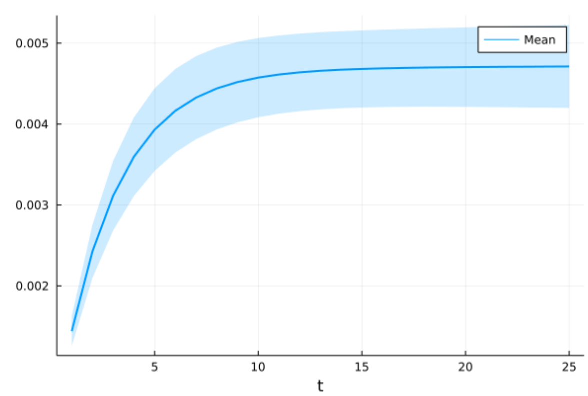

Figure 11: CARA utility

(a) Consumption i.i.d ($\rho = 0$) (b) Firm size i.i.d ($\rho = 0$)

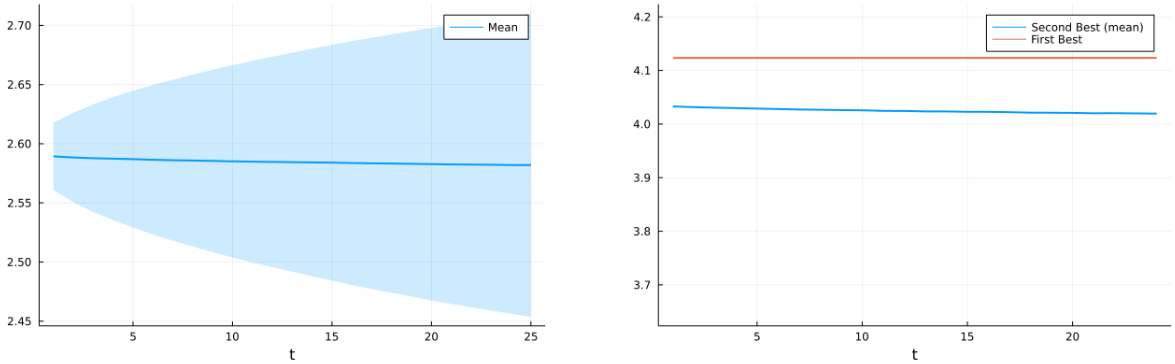

B Derivations and proofs

The Hamiltonian of the recursive principal's problem is

$$\begin{aligned} \mathcal{H} = \left[k_{t+1}(\theta^t) - b_t(\theta^t) + qK_{t+1}(v_t(\theta^t), \Delta_t(\theta^t), \theta_t, k_{t+1}(\theta^t))\right] \varphi_t(\theta_t|\theta^{t-1}) \\ -\lambda_t \varphi_t(\theta_t|\theta^{t-1}) \left[w_t(\theta^t) - v_{t-1}\right] - \gamma_t \varphi_t(\theta_t|\theta^{t-1}) \left[w_t(\theta^t)\mathcal{E}(\theta_t, \theta^{t-1}) - \Delta_{t-1}\right] \\ +\mu_t(\theta^t) \left[u'(f(k_t, \theta_t) - b_t(\theta^t))\iota f_\theta(k_t, \theta_t) + \beta\Delta_t(\theta^t)\right] \\ +\xi_t(\theta^t)\varphi_t(\theta_t|\theta^{t-1}) \left[w_t(\theta^t) - u(f(k_t, \theta_t) - b_t(\theta^t)) - \beta v_t(\theta^t)\right] \end{aligned}$$

The optimality conditions are

$b_t(\theta^t)$:

$$\xi_t(\theta^t) = \frac{1}{u'(\theta^t)} \left[1 + \frac{\mu_t(\theta^t)}{\varphi_t(\theta_t|\theta^{t-1})} \iota f_\theta(\theta^t) u''(\theta^t)\right] \tag{31}$$

The envelope conditions are

$$\frac{\partial K_{t+1}}{\partial v_t(\theta^t)} = \lambda_{t+1}(\theta^t) \tag{32}$$

$$\frac{\partial K_{t+1}}{\partial \Delta_t(\theta^t)} = \gamma_t(\theta^t) \tag{33}$$

$$\frac{\partial K_{t+1}}{\partial k_{t+1}(\theta^t)} = \mathbb{E}\left[-\xi_{t+1}(\theta^{t+1})u'(\theta^{t+1})f_k(\theta^{t+1})|\theta^t\right] + \mathbb{E}\left[\frac{\mu_{t+1}(\theta^{t+1})}{\varphi_{t+1}(\theta_{t+1}|\theta^t)}\left(u''(\theta^{t+1})\iota f_\theta(\theta^{t+1})f_k(\theta^{t+1}) + u'(\theta^{t+1})\iota f_{\theta k}(\theta^{t+1})\right)|\theta^t\right] \tag{34}$$

Using the envelope conditions (32) and (33) we get

$v_t(\theta^t)$:

$$\lambda_{t+1}(\theta^t) = \frac{\beta}{q}\xi_t(\theta^t) \tag{35}$$

$\Delta_t(\theta^t)$:

$$\gamma_{t+1}(\theta^t) = -\frac{\beta}{q}\frac{\mu_t(\theta^t)}{\varphi_t(\theta_t|\theta^{t-1})} \tag{36}$$

Substituting (31) and (34) into the FOC for $k_{t+1}(\theta^t)$ we get

$$\frac{1}{q} = \mathbb{E}\left[f_k(\theta^{t+1}) - \frac{\mu_{t+1}(\theta^{t+1})}{\varphi_{t+1}(\theta_{t+1}|\theta^t)}u'(\theta^{t+1})\iota f_{\theta k}(\theta^{t+1})|\theta^t\right] \tag{37}$$

Finally the law of motion for the co-state is

$$\dot{\mu}_t(\theta^t) = -\left[\xi_t(\theta^t) - \lambda_t - \gamma_t\mathcal{E}(\theta_t,\theta^{t-1})\right]\varphi_t(\theta_t|\theta^{t-1}) \tag{38}$$

Proof Proposition 1 Set $\mu_t(\theta^t) = 0$ for all θ^t, then from equation (37) we obtain point 3. For point 2, note that with $\mu_t(\theta^t) = 0$, equation (38) becomes

$$\xi_t(\theta^t) = \lambda_t$$

From equation (31),

$$\frac{1}{u'(\theta^t)} = \xi_t(\theta^t)$$

and using (35) gives point 2. Point 1 holds in the first best and second best allocations.

Proof Proposition 2 From the FOC for $k_{t+1}(\theta^t)$ (equation (37)), multiplying the second term inside the expectation by $\frac{f_k(\theta^{t+1})}{f_k(\theta^{t+1})}$ and using the definition of the investment wedge, we get

$$\tau^k(\theta^{t+1}) = \frac{\mu_{t+1}(\theta^{t+1})}{\varphi_{t+1}(\theta_{t+1}|\theta^t)} u'(\theta^{t+1}) \iota \frac{f_{\theta k}(\theta^{t+1})}{f_k(\theta^{t+1})}$$

Multiplying by $\frac{\theta_{t+1}}{\theta_{t+1}}$ and $\frac{1-\Phi_{t+1}(\theta_{t+1}|\theta^t)}{1-\Phi_{t+1}(\theta_{t+1}|\theta^t)}$ and rearranging terms

$$\tau^k(\theta^{t+1}) = \iota \frac{\theta_{t+1} f_{\theta k}(\theta^{t+1})}{f_k(\theta^{t+1})} \times \frac{\mu_{t+1}(\theta^{t+1})}{1-\Phi_{t+1}(\theta_{t+1}|\theta^t)} u'(\theta^{t+1}) \times \frac{1-\Phi_{t+1}(\theta_{t+1}|\theta^t)}{\theta_{t+1}\varphi_{t+1}(\theta_{t+1}|\theta^t)}$$

Proof Proposition 3 These are the same steps as proposition 1 in Hellwig (2021). Substitue $\xi_t(\theta^t)$ in the LOM of the co-state (38):

$$\dot{\mu}_t(\theta^t) + \mu_t(\theta^t) \frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)} = -\left[\frac{1}{u'(\theta^t)} - \lambda_t - \gamma_t \mathcal{E}(\theta_t, \theta^{t-1})\right] \varphi_t(\theta_t|\theta^{t-1})$$

substitute $\frac{m'(\theta^t)}{m(\theta^t)} = \frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)}$, using the boundary conditions $\mu_t(\underline{\theta}) = 0$ and $\mu_t(\overline{\theta}) = 0$ and integrating upwards

$$\mu_t(\theta^t) m(\theta^t) = \int_{\theta_t}^{\overline{\theta}} \left[\lambda_t + \gamma_t \mathcal{E}(\theta', \theta^{t-1}) - \frac{1}{u'(\theta', \theta^{t-1})}\right] \varphi_t(\theta'|\theta^{t-1}) m(\theta^t) d\theta'$$

Using the definition of the incentive-adjusted measure

$$\frac{\mu_t(\theta^t)}{\varphi_t(\theta_t|\theta^{t-1})} = \frac{1-\hat{\Phi}_t(\theta_t|\theta^{t-1})}{\hat{\varphi}_t(\theta_t|\theta^{t-1})} \left\{ \hat{\mathbb{E}}\left[\frac{1}{u'(\theta', \theta^{t-1})} \mid \theta' \geq \theta_t, \theta^{t-1}\right] - \gamma_t \hat{\mathbb{E}}\left[\mathcal{E}(\theta', \theta^{t-1}) \mid \theta' \geq \theta_t, \theta^{t-1}\right] - \lambda_t \right\} \tag{39}$$

To get λ_t, note that using the boundary conditions we have

$$0 = \int_{\underline{\theta}}^{\overline{\theta}} \left[\lambda_t + \gamma_t \mathcal{E}(\theta', \theta^{t-1}) - \frac{1}{u'(\theta', \theta^{t-1})}\right] \varphi_t(\theta_t|\theta^{t-1}) m(\theta^t) d\theta'$$

Or

$$\lambda_t = \hat{\mathbb{E}}\left[\frac{1}{u'(\theta^t)} \mid \theta^{t-1}\right] - \gamma_t \hat{\mathbb{E}}\left[\mathcal{E}(\theta_t, \theta^{t-1}) \mid \theta^{t-1}\right]$$

Substiuting back λ_t into equation (39) and using the definition of $\hat{\rho}(\theta^t)$ (equation (21)) we get the solution.

Proof Proposition 4 This proof also follows similar steps to Theorem 1 in Hellwig (2021). Using the characterization of λ_t in Proposition 3 and substitute the multipliers $\lambda_{t+1}(\theta^t)$ and $\gamma_{t+1}(\theta^t)$ from the optimality conditions (35) and (36), and using equation (31) to substitute for ξ_t:

$$\frac{1}{u'(\theta^t)} + \frac{\mu_t(\theta^t)}{\varphi_t(\theta_t|\theta^{t-1})}\frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)} = \frac{q}{\beta}\hat{\mathbb{E}}\left[\frac{1}{u'(\theta^{t+1})}|\theta^t\right] + \frac{\mu(\theta^t)}{\varphi_t(\theta_t|\theta^{t-1})}\hat{\mathbb{E}}\left(\mathcal{E}(\theta_{t+1},\theta^t)|\theta^t\right) \tag{40}$$

where we can rewrite

$$\hat{\mathbb{E}}\left[\mathcal{E}(\theta_{t+1}|\theta^t)|\theta^t\right] = \hat{\mathbb{E}}\left[\rho(\theta^{t+1})\frac{u''(\theta^{t+1})\iota f_\theta(\theta^{t+1})}{u'(\theta^{t+1})}|\theta^t\right]$$

To show this, note that

$$\begin{aligned}\hat{\mathbb{E}}\left[\mathcal{E}(\theta_{t+1},\theta^t)|\theta^t\right] &= \int_{\underline{\theta}}^{\overline{\theta}} \mathcal{E}(\theta_{t+1},\theta^t)\frac{\varphi(\theta_{t+1}|\theta^t)m(\theta^{t+1})}{\mathbb{E}\left[m(\theta^{t+1})|\theta^t\right]}d\theta_{t+1}\\ &= \frac{1}{\mathbb{E}\left[m(\theta^{t+1})|\theta^t\right]}\int_{\underline{\theta}}^{\overline{\theta}}\left(-\int_{\theta_{t+1}}^{\overline{\theta}}\mathcal{E}(\theta',\theta^t)\varphi(\theta'|\theta^t)d\theta'\right)' m(\theta^{t+1})d\theta_{t+1}\end{aligned}$$

Integrate by parts and use $\mathbb{E}\left[\mathcal{E}(\theta_{t+1},\theta^t)|\theta^t\right] = \int_{\underline{\theta}}^{\overline{\theta}}\frac{\partial\varphi(\theta_{t+1}|\theta^t)}{\partial\theta_t}d\theta_{t+1} = 0$. Then using the definition of $\rho(\theta^{t+1})$ and $\frac{m'(\theta^t)}{m(\theta^t)} = \frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)}$

$$\begin{aligned}\hat{\mathbb{E}}\left[\mathcal{E}(\theta_{t+1},\theta^t)|\theta^t\right] &= \int_{\underline{\theta}}^{\overline{\theta}}\int_{\theta_{t+1}}^{\overline{\theta}}\mathcal{E}(\theta_{t+1},\theta^t)\varphi_{t+1}(\theta'|\theta^t)d\theta'\frac{m'(\theta^{t+1})}{\mathbb{E}\left[m(\theta^{t+1})|\theta^t\right]}d\theta_{t+1}\\ &= \int_{\underline{\theta}}^{\overline{\theta}}\frac{1}{\varphi_{t+1}(\theta_{t+1}|\theta^t)}\int_{\theta_{t+1}}^{\overline{\theta}}\mathcal{E}(\theta',\theta^t)\varphi_{t+1}(\theta'|\theta^t)d\theta'\frac{m'(\theta^{t+1})}{m(\theta^{t+1})}\frac{m(\theta^{t+1})}{\mathbb{E}\left[m(\theta^{t+1})|\theta^t\right]}\varphi_{t+1}(\theta_{t+1}|\theta^t)d\theta_{t+1}\\ &= \int_{\underline{\theta}}^{\overline{\theta}}\rho(\theta^{t+1})\frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)}\hat{\varphi}_{t+1}(\theta_{t+1}|\theta^t)d\theta_{t+1}\end{aligned}$$

Substitute back and use the definition of the investment wedge to substitute $\frac{\mu_t(\theta^t)}{\varphi(\theta_t|\theta^{t-1})}$

$$\frac{1}{u'(\theta^t)}+\frac{f_k(\theta^t)}{f_{\theta k}(\theta^t)}\frac{\tau^k(\theta^t)}{u'(\theta^t)}\frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)} = \frac{q}{\beta}\hat{\mathbb{E}}\left[\frac{1}{u'(\theta^{t+1})}|\theta^t\right]+\frac{f_k(\theta^t)}{f_{\theta k}(\theta^t)}\frac{\tau^k(\theta^t)}{u'(\theta^t)}\hat{\mathbb{E}}\left[\rho(\theta^{t+1})\frac{u''(\theta^{t+1})\iota f_\theta(\theta^{t+1})}{u'(\theta^{t+1})}|\theta^t\right]$$

$$\frac{1}{u'(\theta^t)}+\frac{\theta_t}{\mathcal{E}^f(\theta^t)}\frac{\tau^k(\theta^t)}{u'(\theta^t)}\frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)}-\frac{\theta_t}{\mathcal{E}^f(\theta^t)}\frac{\tau^k(\theta^t)}{u'(\theta^t)}\hat{\mathbb{E}}\left[\rho(\theta^{t+1})\frac{u''(\theta^{t+1})\iota f_\theta(\theta^{t+1})}{u'(\theta^{t+1})}|\theta^t\right] = \frac{q}{\beta}\hat{\mathbb{E}}\left[\frac{1}{u'(\theta^{t+1})}|\theta^t\right]$$

Rearranging we get

$$\frac{q}{\beta}\hat{\mathbb{E}}\left[\frac{1}{u'(\theta^{t+1})}|\theta^t\right] = \left\{1 + \underbrace{\left[\frac{u''(\theta^t)\iota f_\theta(\theta^t)}{u'(\theta^t)} - \hat{\mathbb{E}}\left[\rho(\theta^{t+1})\frac{u''(\theta^{t+1})\iota f_\theta(\theta^{t+1})}{u'(\theta^{t+1})}|\theta^t\right]\right]\frac{\theta_t}{\mathcal{E}^f(\theta^t)}\tau^k(\theta^t)}_{\equiv s(\theta^t)}\right\}\frac{1}{u'(\theta^t)}$$

Proof Proposition 5 If $s(\theta^t) \geq 0$,

$$\frac{1}{u'(\theta^t)} \leq \hat{\mathbb{E}}\left[\frac{1}{u'(\theta^{t+1})}|\theta^t\right] = \frac{\mathbb{E}\left[M(\theta^{t+1})|\theta^t\right]}{\mathbb{E}\left[u'(\theta^{t+1})M(\theta^{t+1})|\theta^t\right]}$$

where $M(\theta^{t+1}) = \frac{m(\theta^{t+1})}{u'(\theta^{t+1})}$, rearranging

$$\frac{\mathbb{E}\left[u'(\theta^{t+1})M(\theta^{t+1})|\theta^t\right]}{\mathbb{E}\left[M(\theta^{t+1})|\theta^t\right]} \leq u'(\theta^t)$$

Because $u'(\theta^{t+1})$ is decreasing, to show $\mathbb{E}\left[u'(\theta^{t+1})|\theta^t\right] \leq u'(\theta^t)$ we only need to show that $M(\theta^{t+1})$ is weakly decreasing. Differentiating

$$\frac{d}{d\theta_{t+1}}\left(M(\theta^{t+1})\right) = \underbrace{M(\theta^{t+1})\frac{u''(\theta^{t+1})}{u'(\theta^{t+1})}}_{\leq 0}\left(\iota f_\theta(\theta^{t+1}) - c'(\theta^{t+1})\right)$$

from the the local IC constraint $\frac{dw(\theta^{t+1})}{d\theta_{t+1}} = \frac{\partial w(\theta^{t+1})}{\partial \theta_{t+1}}$ we have $c'(\theta^{t+1}) + \beta\frac{\mathbb{E}\left(\frac{dw(\theta^{t+2})}{d\theta_{t+1}}|\theta_{t+1}\right)}{u'(\theta^{t+1})} = \iota f_\theta(\theta^{t+1})$. Incentive compatibility requires that $\mathbb{E}\left(\frac{dw(\theta^{t+2})}{d\theta_{t+1}}|\theta_{t+1}\right) \geq 0$, which implies $\iota f_\theta(\theta^{t+1}) - c'(\theta^{t+1}) \geq 0$, and so $M(\theta^{t+1})$ is weakly decreasing. For the second part, assume $s(\theta^t) \geq 0$ for all θ^t, then u' follows a non-negative super-martingale. By Doob's super-martingale convergence theorem u' converges almost surely to a finite limit. By contradiction, assume u' converges to a positive limit $\overline{u'} > 0$. Then, almost sure convergence implies that for some τ we have $u'(\theta^\tau) = u'(\theta^\tau, \theta_{\tau+1}) = ... = \overline{u'}$, which would violate incentive compatibility. Hence we must have $u' \to 0$ almost surely.

C Extensions

C.1 Limited commitment

In this section, I relax the assumption of full commitment of the entrepreneur. Limited commitment leads to very different firm size and compensation dynamics than the private information friction. The limited commitment works as follows. At every period, before knowing the realization of his productivity, the entrepreneur can divert and consume all the funds advanced by the lender and terminate the project. In this case, I assume the entrepreneur would obtain utility $h(k_{t+1}(\theta^t))$[34], where h is increasing and concave. Therefore, the agent will not terminate the project at period $t+1$ if $h(k_{t+1}(\theta^t)) \leq v_t(\theta^t)$. This limited commitment constraint can be added directly to the planning problem (12). Because the limited commitment constraint does not affect the within-period insurance and incentives trade-off, the characterization of the shadow cost of insurance (Proposition 3) is not affected by the limited commitment assumption.

However, the limited commitment constraint does modify the consumption and firm size dynamics. Let $\eta_t(\theta^t)$ be the multiplier on the limited commitment constraint. Then the GIEE is given by

$$\frac{q}{\beta}\hat{\mathbb{E}}\left[\frac{1}{u'(\theta^{t+1})}|\theta^t\right] = \frac{1}{u'(\theta^t)}(1+s(\theta^t)) + \frac{\eta_t(\theta^t)}{\beta}$$

Because $\eta_t(\theta^t) \geq 0$, the limited commitment gives a force to have a downward drift in marginal utilities. As is well known, in models with only limited commitment, the agent's consumption is backloaded, and consumption follows a sub-martingale. Therefore, the private information and limited commitment frictions will generally have opposite effects on consumption dynamics.

The investment wedge is now given by

$$\tau^{k,LC}(\theta^{t+1}) = \tau^k(\theta^t) + \eta_t(\theta^t)\frac{h'(k_{t+1}(\theta^t))}{f_k(k_{t+1}(\theta^t),\theta_t)} \geq 0$$

where $\tau^k(\theta^t)$ is the wedge from the private information friction in proposition 2. Because $\eta_t(\theta^t)h'(k_{t+1}(\theta^t)) \geq 0$, the limited commitment friction also lowers firm size relative to the first. However, if promised utility increases over time, the limited commitment constraint

[34] A natural specification of the function h is $\frac{u((1-\iota)(1-q)k_{t+1})}{1-\beta}$, this is the value that the agent would obtain if he could keep a fraction $(1-\iota)$ of the capital and then save outside the contract at rate $\frac{1}{q}$.

will eventually not bind ($\eta_t(\theta^t) = 0$). Therefore, this friction still gives an incentive to have firm size increasing over time.

C.2 Endogenous termination

In this section, I show how the model can be extended to allow for endogenous termination of the contract. As is well known, in regions of the state space where the contract becomes very inefficient, the principal may be better of terminating the project or randomizing between terminating and continuing the contract at an efficient point. I assume that after termination, the lender receives a scrap value S. At period t, based on θ^t, the lender can choose a probability $\alpha_{t+1}(\theta^t)$ of termination at $t+1$. In that event, the principal can also give the entrepreneur a compensation of $Q_{t+1}(\theta^t)$. In case of no termination at period t the objective of the principal is

$$\int \left[-b(\theta^t) + \alpha_{t+1}(\theta^t) q \left(S - Q_{t+1}(\theta^t)\right) + \left(1 - \alpha_{t+1}(\theta^t)\right) \left(k_{t+1}(\theta^t) + qK_{t+1}(v_t(\theta^t), \Delta_t(\theta^t), \theta^t, k_{t+1}(\theta^t))\right)\right] \varphi_t(\theta_t|\theta^{t-1}) d\theta_t$$

I assume that after terminating the contract, the entrepreneur can freely save $Q_{t+1}(\theta^t)$ and obtains a per period gross return $\frac{1}{q}$. Then, his value after terminating the contract is $\frac{u\left((1-q)Q_{t+1}(\theta^t)\right)}{(1-q)}$. The continuation utility now becomes

$$w_t(\theta^t) = u(c(\theta^t)) + \beta \left[\alpha_{t+1}(\theta^t) \frac{u\left((1-q)Q_{t+1}(\theta^t)\right)}{(1-q)} + (1 - \alpha_{t+1}(\theta^t)) v_t(\theta^t)\right]$$

And the local IC

$$\dot{w}_t(\theta^t) = u'(c(\theta^t)) \iota f_\theta(k_t, \theta_t) + \beta(1 - \alpha_{t+1}(\theta^t)) \Delta_t(\theta^t)$$

It is then easy to see that the optimality condtions for $b(\theta^t)$, $k_{t+1}(\theta^t)$,$v_t(\theta^t)$, $\Delta_t(\theta^t)$ and $w_t(\theta^t)$ are the same as in the main model. Therefore, although it may be optimal to terminate the contract, the characterizations of the optimal allocation presented in the paper do not rely on the assumption of no termination.

It is interesting to understand when and where (in the state space) termination may occur in this model and how it compares with the risk neutral case. With risk neutrality (Clementi *et al.* (2010)) the Pareto frontier is increasing in the promised utility. Consequently, termination occurs when the promised utility is low. Moreover, because utilities drift upward, the termination probability tends to decrease over time. With risk aversion, the Pareto frontier is decreasing in promised utility (see figure 12a). So the motives for termination differ from the risk neutral case.

Two other sources of inefficiencies could motivate the termination of the firm. First, high promised insurance (i.e. low Δ_0) is inefficient. Panel 12b shows that the Pareto frontier is increasing in Δ_0. Over time, if the lender has promised too much insurance to the entrepreneur, termination could potentially become optimal. Second, introducing a limited commitment constraint as in C.1 could also generate endogenous termination. If $v_t(\theta^t)$ decreases sufficiently, the limited commitment constraint may require that $k_{t+1}(\theta^t) \to 0$. Then, as shown in Dovis (2019), the Inada condition $\lim_{k\to 0} f_k(k, \theta) = \infty$ implies that in this region, the Pareto frontier is increasing in v. When the frontier is increasing, there may be a range of scrap values S where it is optimal for the lender to randomize between termination and continuing at a higher v. Because both insurance and the variance of promised utility tend to increase, both inefficiencies should imply that the termination probabilities tend to increase over time. Again, these are the opposite dynamics of what is found with risk neutrality.

Figure 12: Pareto frontier

(a) As a function of promised utility v_0 (b) As a function of promised insurance Δ_0

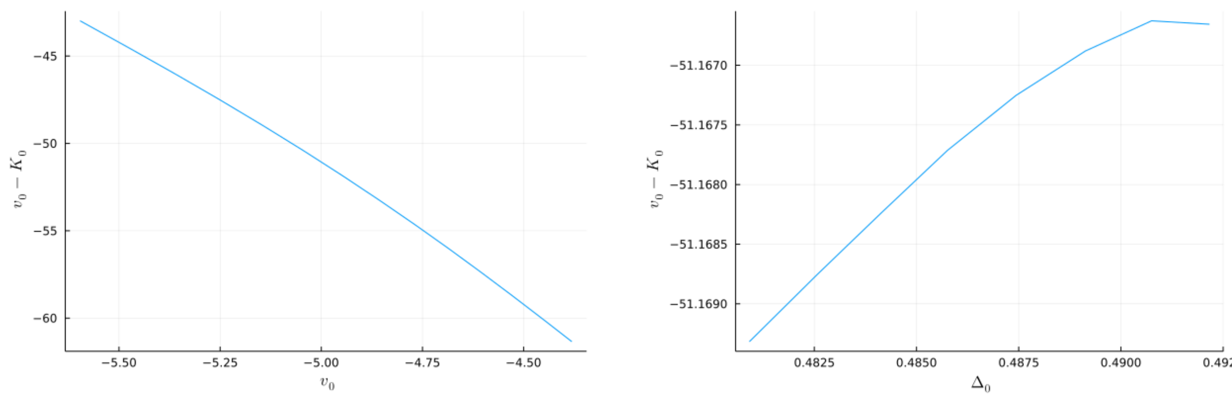

Note: For Panel (a), the promised insurance is set at the optimal level at $t = 0$, i.e. I set $\gamma_0 = 0$. For Panel (b), λ_0 is adjusted for every value of Δ_0 so that v_0 is kept fixed.

Persistent shocks and optimal termination probabilities: Another interesting observation is that whenever termination is optimal, the lender may have more incentives to increase the termination probabilities when the persistence of the shocks is higher. The reason is that a higher termination probability decreases the Markov information rents. The intuition is similar to the equity purchases and the distortions in firm size. Imagine that, at history θ^{t-1}, the lender increases the termination probability of type θ_t and compensates him by increasing $Q(\theta^t)$ such that his ex-ante continuation utility is kept constant. Types $\theta' > \theta_t$ know they are expected to obtain higher returns at $t+1$, so they have a relatively higher preference for continuing to operate the firm. Therefore, the increase in $\alpha(\theta^t)$ makes deviations less attractive for $\theta' > \theta_t$, and so it lowers the cost of screening types.

I show this intuition more formally in a simplified two-period and two-type version of the model. Assume the entrepreneur's productivity can take values $\{\theta^H, \theta^L\}$ with $\theta^H > \theta^L$. In the first period, $P(\theta_1 = \theta^H) = p^1$, and in the second one, $P(\theta_2 = \theta^H | \theta_1 = \theta^H) = p^H$ and $P(\theta_2 = \theta^H | \theta_1 = \theta^L) = p^L$. Let $\rho = p^H - p^L \geq 0$, if $\rho > 0$ we say types are persistent. We assume the production function is of the form $f(\theta, k) = \theta$, so we can abstract away from the choice of firm size. In the second (and last) period, we assume that the entrepreneur consumes all its endowment, so there is no repayment. The principal's objective is

$$K(v) = p^1 \left[-b^H + q\alpha^H(S - Q^H)\right] + (1 - p^1)\left[-b^L + q\alpha^L(S - Q^L)\right]$$

The values of the high and low types are

$$\begin{aligned} w^H &= u(\theta^H - b^H) + \beta\left[\alpha^H u(Q^H) + (1 - \alpha^H)\mathbb{E}^H\left(u(\theta_2)\right)\right] \\ w^L &= u(\theta^L - b^L) + \beta\left[\alpha^L u(Q^L) + (1 - \alpha^L)\mathbb{E}^L\left(u(\theta_2)\right)\right] \end{aligned}$$

where for $j \in \{H, L\}$, $\mathbb{E}^j\left(u(\theta_2)\right) = p^j u(\theta^H) + (1 - p^j)u(\theta^L)$. The participation constraint is

$$p^1 w^H + (1 - p^1) w^L = v$$

and the IC constraint can be written as

$$w^H = w^L + \underbrace{u(\theta^H - b^L) - u(\theta^L - b^L)}_{\text{static info rent}} + \underbrace{(1 - \alpha^L)\beta\rho(u(\theta^H) - u(\theta^L))}_{\text{Markov info rent}}$$

Notice that the Markov information rent is increasing in ρ and decreasing in α^L. I directly assume that the parameters are such that $\alpha^L \in (0, 1)$ is optimal and show that the principal increases the termination probability when the persistence increases.

Proposition 6. *If $\alpha^L \in (0, 1)$ is optimal, the optimal contract is such that $\frac{\partial \alpha^L}{\partial \rho} > 0$.*

Proof. The proof is as follows. Starting from the optimal contract, we consider a perturbation where we increase α^L while preserving the IC and PK constraints and show that the resource gains are increasing in ρ . To this end, let $\Delta\alpha^L = \varepsilon > 0$, for ε small. We perturb the allocation along the low type's indifference curve, so to keep w^L constant, we increase Q^L by

$$Q^L = \frac{\left[u(Q^L) - \mathbb{E}^L\left(u(\theta_2)\right)\right]}{\alpha^L u'(Q^L)}\varepsilon^L$$

The perturbation lowers the Markov information rents, so it relaxes the IC constraint. This allows us to lower the high type's period one utility by

$$\Delta u^H = -\beta\rho(u(\theta^H) - u(\theta^L))\varepsilon$$

Because w^L is kept fixed, this changes the ex-ante utility by $p^1\Delta u^H$. Then, to satisfy the PK constraint, we increase the period one utility of both types in an incentive-compatible manner. Because information rents depend on consumption, increasing utilities uniformly would not be incentive compatible. If we increase the low type's utility by Δu^L, the IC constraint requires increasing the utility of the high type by

$$\Delta u^{H,IC} = \frac{u'(\theta^H - b^L)}{u'(\theta^L - b^L)}\Delta u^L$$

The ex-ante utility is kept fixed if

$$p^1\frac{u'(\theta^H - b^L)}{u'(\theta^L - b^L)}\Delta u^L + (1-p^1)\Delta u^L = -p^1\Delta u^H$$

which implies

$$\Delta u^L = -\frac{p^1 u'(\theta^L - b^L)}{p^1 u'(\theta^H - b^L) + (1-p^1)u'(\theta^L - b^L)}\Delta u^H$$

Therefore, the total change in the high type utility is

$$\begin{aligned}\Delta u^{H,TOT} &= \Delta u^H + \Delta u^{H,IC} \\ &= (1-p^1)\frac{u'(\theta^L - b^L)}{p^1 u'(\theta^H - b^L) + (1-p^1)u'(\theta^L - b^L)}\Delta u^H\end{aligned}$$

Finally, he resource gain from this pertubation is

$$\begin{aligned}\frac{\Delta K}{\varepsilon} &\approx p^1\left[\frac{1}{u'(\theta^H - b^H)}\Delta u^{H,TOT}\right] + (1-p^1)\left[\frac{1}{u'(\theta^L - b^L)}\Delta u^L\right] + \Omega \\ &\approx \left[\frac{1}{u'(\theta^L - b^L)} - \frac{1}{u'(\theta^H - b^H)}\right]\frac{u'(\theta^L - b^L)}{p^1 u'(\theta^H - b^L) + (1-p^1)u'(\theta^L - b^L)}(1-p^1)p^1\beta\rho(u(\theta^H) - u(\theta^L)) + \Omega\end{aligned}$$

where Ω collects all the terms that do not depend on ρ. Because the initial allocation is optimal, $\left[\frac{1}{u'(\theta^L-b^L)} - \frac{1}{u'(\theta^H-b^H)}\right] < 0$ and $u(\theta^H) - u(\theta^L) > 0$. Therefore, the principal's resource gain from this perturbation is increasing in ρ, i.e. $\frac{\partial \frac{\Delta K}{\varepsilon}}{\partial \rho} < 0$, which implies that $\frac{\partial \alpha^L}{\partial \rho} > 0$ is optimal. □

C.3 Screening model: divert funds before investing

In this section, I study a screening version of the model where the entrepreneur can choose what fraction of the funds available he invests in the project. The remaining funds are secretly diverted for consumption. Now the lender can observe the entrepreneur's returns but not the entrepreneur's productivity nor invested and diverted funds. In this sense, the investment decision is similar to the labor/leisure choice in the Mirrlees taxation problem. This model yields the same characterization of the shadow costs μ_t, the GIEE, and the firm size dynamics. Moreover, we can directly define the investment wedge $\tau^k(\theta^t)$ as the wedge between invested and diverted funds relative to the first best.

Denote by B_t the funds advanced by the lender. The entrepreneur can use these funds to invest in the project k_t, but he can also divert a portion a_t of the funds for his consumption. Therefore, invested and diverted funds are subject to the flow of funds constraint

$$k_t + a_t \leq B_t \tag{41}$$

The lender now observes returns $f(k_t, \theta_t)$ but not productivity θ_t and how funds are used, i.e. k_t and a_t. Diverted funds are converted into consumption units according to the function $g(a_t)$, with $g'' \leq 0 < g'$, so the entrepreneur's consumption is

$$c_t = f(k_t, \theta_t) - b_t + g(a_t) \tag{42}$$

The principal's within period objective now is $B_t - b_t$. The envelope condition is

$$\frac{\partial}{\partial \theta_t} w_t(\theta^t) = u'(c_t(\theta^t)) f_\theta(k_t(\theta^t), \theta_t) + \beta \Delta_t(\theta^t)$$

Now the investment wedge can be defined explicitly as the distortion in invested and diverted funds relative to the first best (where we would have $f_k(k_t(\theta^t), \theta_t) = g'(a_t(\theta^t))$. Define

$$\tau^k(\theta^t) \equiv 1 - \frac{g'(a(\theta^t))}{f_k(k(\theta^t), \theta_t)}$$

The rest of the planning problem is the same but with the extra flow of funds constraint (41). The optimality condition for diverted funds is

$$\zeta_t(\theta^t) = g'(a_t(\theta^t))$$

where $\zeta_t(\theta^t)$ is the multiplier on the flow of funds constraint (41). The FOC for investment is

$$\zeta_t(\theta^t) = f_k(k_t(\theta^t), \theta_t) - \frac{\mu_t(\theta^t)}{\varphi_t(\theta_t|\theta^{t-1})} u'(\theta^t) f_{\theta k}(k_t(\theta^t), \theta_t)$$

Then combining the two optimality conditions we get

$$\tau^k(\theta^t) = \frac{\mu_t(\theta^t)}{\varphi_t(\theta_t|\theta^{t-1})} \frac{f_{\theta k}(\theta^t)}{f_k(\theta^t)} u'(\theta^t) > 0$$

which is the same as in proposition 2. Because $\tau^k(\theta^t) > 0$, there is more cash diversion than in the first best. This is the standard screening result; the principal distorts effort (here investment k_t) downwards to screen types at a lower cost. When shadow costs ($\mu_t(\theta^t)$) are high, the principal increases distortions to reduce the costs of screening types. Moreover, this wedge also captures the distortions to firm size as in the cash flow diversion model. Combining the FOC for $B_{t+1}(\theta^t)$ and the envelope condition, we get

$$\begin{aligned} \frac{1}{q} &= \mathbb{E}\left[f_k(k_{t+1}(\theta^t), \theta_{t+1}) - \frac{\mu_{t+1}(\theta^{t+1})}{\varphi_{t+1}(\theta_{t+1}|\theta^{t+1})} u'(\theta^{t+1}) f_{\theta k}(k_{t+1}(\theta^t), \theta_{t+1}) |\theta^t \right] \\ &= \mathbb{E}\left[f_k(k_{t+1}(\theta^t), \theta_{t+1}) \left(1 - \tau^k(\theta^{t+1})\right) |\theta^t \right] \end{aligned}$$

which is exactly how the wedges where defined for the cash flow diversion model in equation (16). Then, it is also easy to verify that this model yields the same characterization for the shadow costs $\mu_t(\theta^t)$ and the GIEE.

D Full implementation in two-period model

In this section, I study a full implementation in a simplified two-period version of the model. Now, the entrepreneur can freely choose between paying himself dividends or reinvesting in the firm. This implementation focuses on another way the principal can lower Markov information rents. Instead of buying equity, the principal discourages reinvesting profits to distort firm size. Because lower future capital is less attractive for higher types if shocks are persistent, this margin can be used to screen types. This implementation of the optimal allocation involves either a (nonlinear) subsidy on dividend payouts or a (nonlinear) tax on reinvested funds so that capital is lowered in period 2 relative to the first best. The magnitudes of the marginal subsidy (or tax) are increasing in the persistence of productivity.

There are two periods $t = 1, 2$. I assume the production function is of the form $f(k, \theta) = \theta f(k)$; for the first period, I normalize $f(k_1) = 1$ and assume there is no cost of diverting funds $\iota = 1$. The entrepreneur is risk neutral in the second period and there is no repayment. So the entrepreneur's utility if he is type θ_1 and reports $\widetilde{\theta}_1$ is

$$w(\theta_1, \widetilde{\theta}_1) = u(\theta_1 - b(\widetilde{\theta}_1)) + \beta \mathbb{E}\left[\theta_2|\theta_1\right] f(k_2(\widetilde{\theta}_1))$$

Solving the lenders problem we obtain

$$u'(\theta_1 - b(\theta_1)) = \beta \mathbb{E}\left[\theta_2|\theta_1\right] f'(k(\theta_1)) \left(1 - \tau^k(\theta_1)\right)$$

where

$$\tau^k(\theta_1) = \widetilde{\mu}(\theta_1)\Psi(\theta^1)\theta_1 \left(\frac{\frac{\partial \mathbb{E}[\theta_2|\theta_1]}{\partial \theta_1}}{\mathbb{E}\left[\theta_2|\theta_1\right]} - \frac{u''(\theta_1 - b(\theta_1))}{u'(\theta_1 - b(\theta_1))} \right) > 0$$

The wedge to investment is increasing in the persistence of the process, as with higher persistence, higher types have an even higher preference for future capital. But now also depends on the absolute risk aversion[35]. We now turn to the implementation. The entrepreneur can freely use his returns to pay dividends d or reinvest in the firm I. The optimal can be implemented with either a tax T on investment such that $k_2 = T(I)$ or a subsidy on dividend payments $c = S(d)$. Here I consider only the subsidy on dividends, so $I = k_2$. Then the entrepreneur's problem is

$$\begin{aligned} w(\theta_1) &= \max_{d, k_2} u(S(d)) + \beta \mathbb{E}\left[\theta_2|\theta_1\right] f(k_2) \\ s.t \quad & d + k_2 = \theta_1 \end{aligned}$$

The marginal subsidy on dividend payments that implements the optimum is

$$S'(d(\theta_1)) = \frac{1}{1 - \tau^k(\theta_1)}$$

[35] Because there is no repayment b_2 in the second period, capital also plays a similiar role as savings for the entrepreneur.

Moreover, we have $S'(d(\overline{\theta})) = S'(d(\underline{\theta})) = 1$ and $S'(d(\theta)) > 1$ for $\theta \in (\underline{\theta}, \overline{\theta})$, so the marginal subsidy is inverse U-shaped. Because $\tau^k(\theta_1)$ is increasing in the persistence of productivity, the marginal subsidy will also be increasing in the persistence.

E Details numerical simulations

I follow a similar procedure as Farhi and Werning (2013), Stantcheva (2017) and Ndiaye (2020). In these papers (and in Kapička (2013) and Golosov *et al.* (2016a)), the model is solved with a geometric random walk process. This allows to normalize the principal's optimization problem and drop θ_{t-1} as a state variable. Here, the problem can also be normalized if the production function is assumed to be of the form $f(k, \theta) = z\theta^{1-\alpha}k^{\alpha}$. However, I am interested in performing comparative statics with respect to the persistence of the process (ρ). Therefore, I solve the full problem without renormalizing.

It is convenient to transform the problem to write the Hamiltonian as a function of the current shock ε_t instead of the current productivity θ_t. Denote the density function of the shock by $g_\varepsilon(\varepsilon_t)$, then it follows that

$$\varphi\left(\theta_t \mid \theta_{t-1}\right) = \frac{g_\varepsilon(\varepsilon_t)}{\theta_{t-1}^{\rho}}$$

moreover, we also have that

$$\frac{\partial \varphi\left(\theta_t \mid \theta_{t-1}\right)}{\partial \theta_{t-1}} = -\frac{\rho}{\varepsilon_t \theta_{t-1}^{1+\rho}} \frac{1}{\sigma_\varepsilon \sqrt{2\pi}} \frac{\left(\log \theta_t - \rho \log \theta_{t-1} - \mu\right)}{\sigma_\varepsilon^2} \exp\left\{-\frac{\left(\log \theta_t - \rho \log \theta_{t-1} - \mu\right)^2}{2\sigma_\varepsilon^2}\right\}$$

and

$$\frac{\partial g_\varepsilon(\varepsilon_t)}{\partial \varepsilon_t} = -\frac{1}{\varepsilon_t^2 \sigma_\varepsilon \sqrt{2\pi}} \frac{\left(\log \varepsilon_t - \mu\right)}{\sigma_\varepsilon^2} \exp\left\{-\frac{\left(\log \varepsilon_t - \mu\right)^2}{2\sigma_\varepsilon^2}\right\}$$

therefore,

$$\tilde{g}_\varepsilon(\varepsilon_t) \equiv g_\varepsilon(\varepsilon_t) + \varepsilon \frac{\partial g_\varepsilon(\varepsilon_t)}{\partial \varepsilon_t} = \frac{\theta_{t-1}^{1+\rho}}{\rho} \frac{\partial \varphi\left(\theta_t \mid \theta_{t-1}\right)}{\partial \theta_{t-1}}$$

Then note that $d\theta_t = \theta_{t-1}^{\rho} d\varepsilon_t$ implies

$$\varphi\left(\theta_t \mid \theta_{t-1}\right) d\theta_t = g_\varepsilon(\varepsilon_t) d\varepsilon_t$$

and

$$\frac{\partial\varphi\left(\theta_t \mid \theta_{t-1}\right)}{\partial\theta_{t-1}}d\theta_t = \rho\frac{\tilde{g}_\varepsilon(\varepsilon_t)}{\theta_{t-1}}d\varepsilon_t$$

The planning problem over the shock ε_t is

$$\begin{aligned}
K(v_{t-1},\Delta_{t-1},k_t,\theta_{t-1}) = \min \int \left(k_{t+1}(\varepsilon_t) - b_t(\varepsilon_t) + qK(v_t(\varepsilon_t),\Delta_t(\varepsilon_t),k_{t+1}(\varepsilon_t),\theta_{t-1}^{\rho}\varepsilon_t)\right) g_\varepsilon(\varepsilon_t)d\varepsilon_t \\
s.t \quad (PK) \quad w_t(\varepsilon_t) = u(c_t(\varepsilon_t)) + \beta v_t(\varepsilon_t) \qquad [g_\varepsilon(\varepsilon_t)\xi_t(\varepsilon_t)] \\
v_{t-1} = \int w_t(\varepsilon_t)g_\varepsilon(\varepsilon_t)d\varepsilon_t \qquad [g_\varepsilon(\varepsilon_t)\lambda_{t-1}] \\
(IC) \quad \dot{w}_t(\varepsilon_t) = \theta_{t-1}^{\rho}\left(u'(c(\varepsilon_t))\iota f_\theta(k_t,\theta_{t-1}^{\rho}\varepsilon_t) + \beta\Delta_t(\varepsilon_t)\right) \qquad [\mu_t(\varepsilon_t)] \\
\Delta_{t-1} = \int w_t(\varepsilon_t)\frac{\rho}{\theta_{t-1}}\tilde{g}_\varepsilon(\varepsilon_t)d\varepsilon_t \qquad [g_\varepsilon(\varepsilon_t)\gamma_{t-1}] \\
(Feasibility) \quad c_t(\varepsilon_t) = f(k_t,\theta_{t-1}^{\rho}\varepsilon_t) - b_t(\varepsilon_t)
\end{aligned}$$

The optimality conditions are

$$\frac{q}{\beta}\lambda_t(\varepsilon_t) = \frac{1}{u'(c_t(\varepsilon_t))}\left[1 + \frac{\mu(\varepsilon_t)}{g_\varepsilon(\varepsilon_t)}\theta_{t-1}^{\rho}\iota f_\theta(k_t,\theta_{t-1}^{\rho}\varepsilon_t)u''(c(\varepsilon_t))\right] \tag{43}$$

$$\gamma_t(\varepsilon_t) = -\frac{\beta}{q}\theta_{t-1}^{\rho}\frac{\mu(\varepsilon_t)}{g_\varepsilon(\varepsilon_t)} \tag{44}$$

And the two LOM

$$\dot{\mu}(\varepsilon_t) = -\left[\frac{q}{\beta}\lambda_t(\varepsilon_t) - \lambda_{t-1} + \gamma_{t-1}\frac{\rho}{\theta_{t-1}}\frac{\tilde{g}_\varepsilon(\varepsilon_t)}{g_\varepsilon(\varepsilon_t)}\right]g_\varepsilon(\varepsilon_t) \tag{45}$$

$$\dot{w}_t(\varepsilon_t) = \theta_{t-1}^{\rho}\left(u'(c(\varepsilon_t))\iota f_\theta(k_t,\theta_{t-1}^{\rho}\varepsilon_t) + \beta\Delta_t(\varepsilon_t)\right) \tag{46}$$

I truncate the distribution of ε_t at the 0.01 and 0.99 percentiles, the boundary conditions then need to be adjusted to $\mu(\overline{\varepsilon}) = -\gamma_{t-1}\frac{\rho}{\theta_{t-1}}\overline{\varepsilon}g_\varepsilon(\overline{\varepsilon})$ and $\mu(\underline{\varepsilon}) = -\gamma_{t-1}\frac{\rho}{\theta_{t-1}}\underline{\varepsilon}g_\varepsilon(\underline{\varepsilon})$.

To solve the model, the state space is modificed to $(\lambda_-,\gamma_-,k,\theta_-)$, so the multipliers λ_- and γ_- are used instead of v_- and Δ_-, respectively. I use 14 grid points for λ_-, 8 for γ_-, 20 for k and 10 for θ_-. I interpolate on K, v and Δ with cubic splines and allow to extrapolate. To solve the model with an i.i.d type process, the algorithm is the same but with $\Delta = 0$ and without the state variables γ_- and θ_-.

Algortihm

Step 0: Guess the value function K', promised utility v' and promised marginal utility Δ' on the grid $(\lambda_-, \gamma_-, k, \theta_-)$

Step 1: Compute the policy functions for k_+ on a grid $(\lambda_{pol}, \gamma_{pol}, \theta)$ by minimizing

$$k_+ + qK'(\lambda_{pol}(i), \gamma_{pol}(i), k_+, \theta(i))$$

(Note: k_+ needs to be computed multiple times at every step while solving the ODE. But to improve speed, we can solve before the policies on a dense grid and then interpolate when solve the ode).

Step 2: For each point in $(\lambda_-, \gamma_-, k, \theta_-)$, solve the optimal control problem with a shooting method.

- a) Guess continuation utility of lowest type $w(\underline{\varepsilon}) = \underline{w}$
- b) For each ε, solve $\lambda(\varepsilon)$ in equation (43) and $\gamma(\varepsilon)$ in equation (44). To compute $c(\varepsilon)$, first compute $k_+(\varepsilon)$ by interpolationg the array of policies on $(\lambda(\varepsilon), \gamma(\varepsilon), \theta_-^\rho \varepsilon)$. Then obtain $v(\varepsilon)$ by interpolation of v' on $(\lambda(\varepsilon), \gamma(\varepsilon), k_+(\varepsilon), \theta_-^\rho \varepsilon)$ and solve

  $$c(\varepsilon) = u^{-1}\left(w(\varepsilon) - \beta v(\varepsilon)\right)$$

  With these solutions solve the differential equations (45) and (46). Note when solving (45) also need to interpolate Δ' on $(\lambda(\varepsilon), \gamma(\varepsilon), k_+(\varepsilon), \theta_-^\rho \varepsilon)$.
- c) Check the boundary condition $\mu(\overline{\varepsilon}) = -\gamma_- \frac{\rho}{\theta_-} \overline{\varepsilon} g_\varepsilon(\overline{\varepsilon})$. If it does not satisfy the tolerance, go back to step a).

Step 3: Given the solution $(\mu(\varepsilon), w(\varepsilon))$, repeat step b) to obtain all policy functions on a grid $(\lambda_-, \gamma_-, k, \theta_-, \varepsilon)$, also compute $b(\varepsilon) = f(k, \theta_-^\rho \varepsilon) - c(\varepsilon)$.

Step 4: Compute the lender's value function, promised utility and expected marginal utility at every grid point

$$v(\lambda_-, \gamma_-, k, \theta_-) = \int w(\lambda_-, \gamma_-, k, \theta_-, \varepsilon) g_\varepsilon(\varepsilon_t) d\varepsilon_t$$

$$\Delta(\lambda_-, \gamma_-, k, \theta_-) = \int w(\lambda_-, \gamma_-, k, \theta_-, \varepsilon) \frac{\rho}{\theta_-} \tilde{g}_\varepsilon(\varepsilon_t) d\varepsilon_t$$

$$K(\lambda_-,\gamma_-,k,\theta_-) = \int \left(k_+(\lambda_-,\gamma_-,k,\theta_-,\varepsilon) - b(\lambda_-,\gamma_-,k,\theta_-,\varepsilon) + qK'(\lambda(\varepsilon),\gamma(\varepsilon),k_+(\varepsilon),\theta_-^{\rho}\varepsilon)\right) g_\varepsilon(\varepsilon_t)d\varepsilon_t$$

Calculate the distance with previouss guess of K', v' and Δ' , and repeat from **Step 1** until the convergence criteria is satisfied.

For the Montecarlo simulation, at the starting period, λ_0 and θ_0 can be fixed at arbitrary values. Because Δ_0 is a free variable, we must set $\gamma_0 = 0$. Then k_1 is chosen optimally given (λ_0, θ) and $\gamma_0 = 0$.

E.1 Check global IC constraints

The first-order approach consists in solving a relaxed problem where only the local incentive constraints are considered. A priori global incentive constraints may bind, in which case the solutions of the relaxed program (12) and the full program (6) would not coincide. I follow the approach outlined in Kapička (2013) and Farhi and Werning (2013) to verify ex-post that only the local incentive constraints bind.

The procedure is the following. First, after solving numerically the relaxed problem, we have obtained the policy functions $\lambda(\lambda_-,\gamma_-,k,\theta_-,\varepsilon)$, $\gamma(\lambda_-,\gamma_-,k,\theta_-,\varepsilon)$, $k_+(\lambda_-,\gamma_-,k,\theta_-,\varepsilon)$ and $b(\lambda_-,\gamma_-,k,\theta_-,\varepsilon)$ and the value function $v(\lambda_-,\gamma_-,k,\theta_-)$. Let $\widetilde{\varepsilon}$ denote the agent's report about the innovation to the productivity. Then we consider a problem where the entrepreneur takes as given the policy functions and can report any $\widetilde{\varepsilon} \in [\underline{\varepsilon}, \varepsilon]$ and verify that for every $(\lambda_-,\gamma_-,k,\theta_-,\varepsilon)$

$$\begin{aligned}
\varepsilon = \arg\max_{\widetilde{\varepsilon}\in[\underline{\varepsilon},\varepsilon]} & u(\widetilde{c}(\lambda_-,\gamma_-,k,\theta_-,\varepsilon,\widetilde{\varepsilon})) \\
& + \beta v(\lambda(\lambda_-,\gamma_-,k,\theta_-,\widetilde{\varepsilon}),\gamma(\lambda_-,\gamma_-,k,\theta_-,\widetilde{\varepsilon}),k_+(\lambda_-,\gamma_-,k,\theta_-,\widetilde{\varepsilon}),\theta_-^{\rho}\widetilde{\varepsilon}) \\
s.t \quad & \widetilde{c}(\lambda_-,\gamma_-,k,\theta_-,\varepsilon,\widetilde{\varepsilon}) = \iota f(k,\varepsilon) + (1-\iota) f(k,\widetilde{\varepsilon}) - b(\lambda_-,\gamma_-,k,\theta_-,\widetilde{\varepsilon})
\end{aligned}$$

E.2 Solution implementation

With persistent shocks and constant equity, the entrepreneur's problem in the quasi-implementation is

$$
\begin{aligned}
\mathcal{W}(W_t, \theta_{t-1}, \widetilde{\theta}_{t-1}, \varepsilon_t) = &\max_{\widetilde{\theta}_t \le \theta_{t-1}^{\rho}\varepsilon_t} u(\widetilde{c}_t) + \beta \mathcal{V}(W_{t+1}, \theta_{t-1}^{\rho}\varepsilon_t, \widetilde{\theta}_t) \\
s.t \qquad W_{t+1} &= qC(W_t, \widetilde{\theta}_t, \widetilde{\theta}_{t-1}) \\
c_t &= (1-q)C(W_t, \widetilde{\theta}_t, \widetilde{\theta}_{t-1}) \\
\widetilde{c}_t &= c_t + \iota(f(k_{SB}, \theta_{t-1}^{\rho}\varepsilon_t) - f(k_{SB}, \widetilde{\theta}_t))
\end{aligned}
$$

where

$$
C(W_t, \widetilde{\theta}_t, \widetilde{\theta}_{t-1}) = \frac{1}{q}W_t + \chi(f(k_{SB}, \widetilde{\theta}_t) + q\overline{f}(k_{SB}, \widetilde{\theta}_t) - \overline{f}(k_{SB}, \widetilde{\theta}_{t-1}))
$$

$$
\overline{f}(k_{SB}, \theta_t) = \mathbb{E}\left[\sum_{\tau=1}^{\infty} q^{\tau-1} f(k_{SB}, \theta_{t+\tau}) | \theta_t\right]
$$

and

$$
\mathcal{V}(W_{t+1}, \theta_t, \widetilde{\theta}_t) = \int \mathcal{W}(W_{t+1}, \theta_t, \widetilde{\theta}_t, \varepsilon_{t+1}) g_{\varepsilon}(\varepsilon_{t+1}) d\varepsilon_{t+1}
$$

With i.i.d shocks, the problem is the same but without θ_{t-1} and $\widetilde{\theta}_{t-1}$ as state variables and with $\overline{f}(k_{SB})$ independent of θ_t. The problem is solved with standard value function iteration, and to have the closest comparison with the solutions of the optimal allocation, $\mathcal{V}(W_{t+1}, \theta_t, \widetilde{\theta}_t)$ is computed with numerical integration.