

OPTIMAL ASSET MANAGEMENT CONTRACTS WITH HIDDEN SAVINGS

SEBASTIAN DI TELLA

Graduate School of Business, Stanford University

YULIY SANNIKOV

Graduate School of Business, Stanford University

We characterize optimal asset management contracts in a classic portfolio-investment setting. When the agent has access to hidden savings, his incentives to misbehave depend on his precautionary saving motive. The contract dynamically distorts the agent's access to capital to manipulate his precautionary saving motive and reduce incentives for misbehavior. We provide a sufficient condition for the validity of the first-order approach, which holds in the optimal contract: global incentive compatibility is ensured if the agent's precautionary saving motive weakens after bad outcomes. We extend our results to incorporate market risk, hidden investment, and renegotiation.

KEYWORDS: Dynamic contracts, hidden savings, precautionary savings.

1. INTRODUCTION

DELEGATED ASSET MANAGEMENT is ubiquitous in modern economies, from fund managers investing in financial assets to CEOs or entrepreneurs managing real capital assets. To align incentives, agents must retain a risky stake in their investment activity. However, hidden savings pose a significant challenge. Agents can save to self-insure against bad outcomes, undoing the incentive scheme. In this paper, we characterize optimal dynamic asset management contracts when the agent has access to hidden savings.

This paper has two main contributions. First, we build a model of delegated asset management based on a classic environment of portfolio investment. An agent with constant relative risk aversion (CRRA) continuously invests in risky assets, but can secretly divert returns and has access to hidden savings. These assumptions lead to a tractable characterization of the optimal contract and the dynamic distortions induced by hidden savings. Second, we provide a general verification theorem for the validity of the first-order approach that is valid for a wide class of contracts. Global incentive compatibility is ensured as long as the agent's precautionary saving motive weakens after bad outcomes. This condition holds in the candidate optimal contract, proving the validity of the first-order approach in our setting.

When the agent has access to hidden savings, his incentives to divert funds depend on his precautionary motive. If the agent expects a risky consumption stream in the future, hidden savings that allow him to self-insure are very valuable. This makes fund diversion to build up savings attractive. The optimal contract must therefore manage the agent's precautionary saving motive by committing to limit his future risk exposure, especially after bad outcomes. This is accomplished by restricting the agent's access to capital, since giving the agent capital to manage requires exposing him to risk in order to align incentives. By promising a future amount of capital that is *ex post* inefficiently low, the principal makes fund diversion less attractive today and reduces the cost of giving capital to

Sebastian Di Tella: sditella@stanford.edu

Yuliy Sannikov: sannikov@gmail.com

We thank Alex Bloedel and Erik Madsen for outstanding research assistance.

the agent up front. This dynamic trade-off leads to history-dependent distortions in the agent's access to capital and a skewed compensation scheme. After good outcomes his access to capital improves, which allows him to keep growing rapidly. After bad outcomes, he gets starved for capital and stagnates. The flip side is that the agent's consumption is somewhat insured on the downside, and he is punished instead with lower consumption growth.

One of the main methodological contributions of this paper is to provide an analytical verification of the validity of the first-order approach. This is an old, hard problem in the theory of dynamic contracts because the agent has such a rich action space.¹ The first-order approach, introduced to the problem of hidden savings by Werning (2001), incorporates the agent's Euler equation as a constraint in the contract design problem. This ensures the agent does not have a profitable deviation with savings alone. However, the agent may find it attractive to divert funds and save for when he is punished for poor performance in the future. We prove the validity of the first-order approach analytically by establishing an upper bound on the agent's continuation utility off-path, for any deviation and after any history. Global incentive compatibility is ensured if the agent's precautionary saving motive becomes weaker after bad outcomes. Intuitively, if the contract becomes less risky after bad outcomes, hidden savings become less valuable to the agent after diverting funds. This condition is guaranteed to hold in the candidate optimal contract, proving the validity of the first-order approach. But it is worth stressing that this verification argument applies beyond the optimal contract to any contract in which the precautionary saving motive weakens after bad outcomes.

Since the principal uses dynamic distortions to the agent's access to capital to provide incentives, it is natural to ask how *hidden investment* and *renegotiation* affect results. Fortunately, both issues can be easily handled. If the agent is able to secretly invest in risky capital, the principal is limited in his ability to restrict his access to capital. The constraint becomes binding only after sufficiently bad outcomes, when the principal would like to reduce capital managed by the agent so much that the agent would find it attractive to invest on his own (i.e., without sharing risk with the principal). We can add to our model observable *market risk* that commands a premium, and allow the agent to also invest his hidden savings in the market. Our agency model can therefore be fully embedded within the standard setting of continuous-time portfolio choice.

A second concern is that the optimal contract requires commitment. The principal relaxes the agent's precautionary saving motive by promising an inefficiently low amount of capital in the future. It is therefore tempting at that point to renegotiate and start over. To address this issue, we also characterize the optimal renegotiation-proof contract. This leads to a stationary contract that restricts the agent's access to capital in a uniform way.

Finally, we map our optimal contract into a classic portfolio-consumption problem with a dynamic leverage constraint that limits the agent's risk exposure. This constraint reduces the agent's precautionary motive and allows the agent's incentive constraints to hold with a lower retained equity stake, improving risk sharing. Hidden savings therefore link retained equity constraints and leverage constraints, which are widely used in applied macro-finance work. The logic is completely different than that of models with limited commitment such as Hart and Moore (1994) and Kiyotaki and Moore (1997). The leverage constraint exists not because the agent can walk away, but because it limits risk and, therefore, helps relax the equity constraint.

¹Kocherlakota (2004) provides a well-known example in which double deviations are profitable (and the first-order approach fails) when cost of effort is linear. To establish the validity of the first-order approach, the existing literature has often pursued the numerical option (e.g., see Farhi and Werning (2013)).

Literature Review. This paper fits within the literature on dynamic agency problems. There is extensive literature that uses recursive methods to characterize optimal contracts, including Spear and Srivastava (1987), Phelan and Townsend (1991), Sannikov (2008), He (2011), Biais, Mariotti, Plantin, and Rochet (2007), and Clementi and Hopenhayn (2006). Our paper builds upon these standard recursive techniques, but adds the problem of persistent private information. In our case, about the agent's hidden savings.

Our environment is based on the classic portfolio-investment problem, widely used in macroeconomic and financial applications. Two salient features of this environment are concave preferences (CRRA in our case), and endogenous capital under management. Here, we would like to explain the role of each feature and put them in context within the literature, keeping in mind that the focus of our paper is the role of hidden savings.

First, concave preferences create a role for hidden savings. The agency problem we study is one of cash flow diversion, as in DeMarzo and Fishman (2007), DeMarzo and Sannikov (2006), and DeMarzo, Fishman, He, and Wang (2012), but unlike their models we have CRRA rather than risk neutral preferences. With risk-neutral preferences, the optimal contracts with and without hidden savings are the same. Once concave preferences are introduced, hidden savings become binding. Rogerson (1985) shows that the inverse Euler equation, which characterizes the optimal contract without hidden savings, implies that the agent would save if he could. This opens the door to potential distortions to control the agent's incentives to save, but also presents the problem of double deviations. In some settings, distortions do not arise, for example, the CARA settings, such as He (2011) and Williams (2008), where the ratio of the agent's current utility to continuation utility is invariant to contract design.² However, when distortions do arise, it is difficult to characterize the specific form they take, and the first-order approach may fail. We are able to provide a sharp characterization of the optimal contract and the distortions generated by the presence of hidden savings, and provide a verification theorem for global incentive compatibility, which is valid for a wide class of contracts.

Second, the classic portfolio-investment problem where capital can be costlessly adjusted affords us a degree of scale invariance. In our setting, after good performance, the agent does not need to be retired nor outgrow moral hazard as in Sannikov (2008) or Clementi and Hopenhayn (2006), respectively. Likewise, since the project can be scaled down, neither will the agent retire after sufficiently bad outcomes as in DeMarzo and Sannikov (2006). Rather, the optimal contract dynamically scales the size of the agent's fund with performance, taking into account his precautionary saving motive.

To understand the link between concave preferences and endogenous capital better, it is useful to contrast our paper to DeMarzo et al. (2012), which is closely connected to ours. In that paper, agents are risk neutral, so hidden savings are not binding (the Euler and inverse Euler equations coincide). With concave preferences, hidden savings are binding. But concave preferences play another important role. Since the principal can give the agent an unboundedly large amount of capital to manage, if the agent is able to obtain an excess return over the market, he can obtain an unboundedly large utility if he is risk neutral. He effectively has access to an arbitrage opportunity. DeMarzo et al. (2012) introduced adjustment costs to capital to eliminate this issue while maintaining risk neutral preferences. Our environment does not have adjustment costs to capital, but the agent's risk aversion eliminates the arbitrage opportunity and makes the problem well-defined.

²Likewise, the dynamic incentive accounts of Edmans, Gabaix, Sadzik, and Sannikov (2011) exhibit no distortions either, as hidden action enters multiplicatively and project size is fixed.

Our model provides a unified account of equity and leverage constraints, which are two of the most commonly used financial frictions in the macro-finance literature in the tradition of [Bernanke and Gertler \(1989\)](#) and [Kiyotaki and Moore \(1997\)](#).³[Di Tella \(2019\)](#) adopted a version of our setting without hidden savings to study optimal financial regulation policy in a general equilibrium environment. Our paper is also related to models of incomplete idiosyncratic risk sharing, such as [Aiyagari \(1994\)](#) and [Krusell and Smith \(1998\)](#). Here, the focus is on risky capital income, as in [Angeletos \(2007\)](#) or [Christiano, Motto, and Rostagno \(2014\)](#), rather than risky labor income.

Access to capital provides the principal with an important incentive tool. He is able to relax the incentive constraints and improve risk sharing by committing to distort project size below optimum over time and after bad performance. This result stands in contrast to [Cole and Kocherlakota \(2001\)](#), where project scale is fixed and the optimal contract is risk-free debt. We recover the result of [Cole and Kocherlakota \(2001\)](#) only in the special case when the agent can secretly invest on his own just as efficiently as through the principal, so the principal cannot control the scale of investment at all.

Our paper is related to the literature on persistent private information, since the agent has private information about savings. The growing literature in this area includes the fundamental approach of [Fernandes and Phelan \(2000\)](#), who propose to keep track of the agent's entire off-equilibrium value function, and the first-order approach, such as [He, Wei, Yu, and Gao \(2017\)](#) and [DeMarzo and Sannikov \(2016\)](#) who use a recursive structure that includes the agent's "information rent," that is, the derivative of the agent's payoff with respect to private information. In all cases, in designing the contract we have to evaluate the agent's payoff off-path, after actions, which the agent is not supposed to take, in order to ensure that the agent's incentives to refrain from those actions are designed properly. In our case, since we want to control the agent's incentives to save secretly, the agent's information rent is marginal utility of consumption (of an extra unit of savings). Key common issues are (1) the way that information rents enter the incentive constraint, (2) distortions that arise from this interaction, and (3) forces that affect the validity of the first-order approach. In our case, we verify the validity of the first-order approach by characterizing an analytic upper bound, related to the CRRA utility function, on the agent's payoff after deviations. The bound coincides with the agent's utility on path (hence the first-order approach is valid), and its derivative with respect to hidden savings is the agent's "information rent." [Farhi and Werning \(2013\)](#) have verified the first-order approach numerically by computing the agent's value function after deviations explicitly, in the context of insurance with unobservable skill shocks. Other papers that study the problems of persistent private information via a recursive structure that includes information rents include [Garrett and Pavan \(2015\)](#), [Cisternas \(2018\)](#), [Kapička \(2013\)](#), and [Williams \(2011\)](#).

2. THE MODEL

We consider a setting in which an agent manages risky capital and contracts with a set of outside investors to raise funds and share risk. We may think of these investors collectively

³See [Bernanke and Gertler \(1989\)](#), [Kiyotaki and Moore \(1997\)](#), [Bernanke, Gertler, and Gilchrist \(1999\)](#), [Iacoviello \(2005\)](#), [Lorenzoni \(2008\)](#), [Gertler et al. \(2010\)](#), [Gertler and Karadi \(2011\)](#), [Adrian and Boyarchenko \(2012\)](#), [Brunnermeier and Sannikov \(2014\)](#), [He and Krishnamurthy \(2012\)](#), [He and Krishnamurthy \(2013\)](#), [Di Tella \(2017\)](#), [Moll \(2014\)](#).

as “the principal,” whose objective is to maximize the market value of their payoff, so we may also refer to them as “the market.”⁴

The agent manages an amount of capital $k_t \geq 0$, determined by the contract, to obtain a risky return per dollar invested of

$$dR_t = (r + \alpha - a_t) dt + \sigma dZ_t. \quad (1)$$

In this expression, $r > 0$ is the risk-free rate, which is also the required market return on investment with idiosyncratic risk σdZ , $\alpha > 0$ is the excess return, and $a_t \geq 0$ is a hidden action the agent takes to divert returns for private monetary benefit. The Brownian motion Z , defined on a complete probability space (Ω, P, \mathcal{F}) with filtration $\mathbb{F} = \{\mathcal{F}_t\}$, represents agent-specific idiosyncratic risk. If the agent is a fund manager, Z represents the outcome of his particular investment/trading activity. If he is an entrepreneur, Z represents the outcome of his particular project.

The principal can commit to any fully history-dependent contract $\mathcal{C} = (c, k)$, which specifies payments to the agent $c_t > 0$ and capital $k_t \geq 0$ as a function of the observed history of returns R up to time t . Payments c_t may differ from the agent’s true consumption $\tilde{c}_t > 0$, which is also a hidden action. The principal does not observe hidden actions a_t or \tilde{c}_t , and so the contract \mathcal{C} depends only on observed returns R . After signing the contract \mathcal{C} , the agent chooses a strategy (\tilde{c}, a) , that specifies \tilde{c}_t and a_t , also as functions of the history of returns R .

Diversion of $a_t \geq 0$ gives the agent a flow of funds $\phi a_t k_t$. For each stolen dollar, the agent keeps only a fraction $\phi \in (0, 1)$. Given the agent’s payments c_t , consumption \tilde{c}_t and diversion action a_t , his hidden savings h_t evolve according to

$$dh_t = (rh_t + c_t - \tilde{c}_t + \phi k_t a_t) dt, \quad (2)$$

and must remain nonnegative, $h_t \geq 0$.⁵ The agent invests his hidden savings at the risk-free rate r . Later we will introduce hidden investment, and allow the agent to also invest his hidden savings in risky capital.

The agent has CRRA preferences. Given contract \mathcal{C} , under strategy (\tilde{c}, a) the agent gets utility

$$U_0^{\tilde{c}, a} = \mathbb{E} \left[\int_0^\infty e^{-\rho t} \frac{\tilde{c}_t^{1-\gamma}}{1-\gamma} dt \right], \quad (3)$$

with ρ and γ strictly positive and $\gamma \neq 1$.⁶ Given contract \mathcal{C} , we say a strategy (\tilde{c}, a) is *feasible* if (1) utility $U_0^{\tilde{c}, a}$ is finite and (2) $h_t \geq 0$ always. Let $\mathbb{S}(\mathcal{C})$ be the set of feasible strategies (\tilde{c}, a) given contract \mathcal{C} .

The principal pays for the agent’s consumption, but keeps the excess return α on the capital that the agent manages. The principal can access a complete financial market with equivalent martingale measure Q where idiosyncratic risk is not priced, so $Q = P$. In

⁴Even if there is not a single principal, the agent himself would like to commit to a contract to then sell shares to dispersed outside investors.

⁵We do not allow negative hidden savings, $h_t \geq 0$. If the agent can secretly borrow up to some limit, $h_t \geq \underline{h} \in (-\infty, 0]$, the principal can first ask the agent to borrow up to the limit and transfer the funds to the principal to invest on his behalf, so we are back to no hidden borrowing. Also, $a_t \geq 0$ means that the agent cannot secretly use his hidden savings to improve observed returns. Section 8 discusses the role of this assumption in detail.

⁶As long as $\gamma \neq 1$, utility in (3) is always well-defined, but could take values ∞ or $-\infty$. We do not explicitly analyze the log utility case with $\gamma = 1$, but it can be treated as the limit $\gamma \rightarrow 1$.

the Online Appendix (Di Tella and Sannikov (2021)), we allow for aggregate risk with a market price, such that $Q \neq P$. The principal tries to minimize the cost of delivering utility u_0 to the agent

$$J_0 = \mathbb{E}^Q \left[\int_0^\infty e^{-rt} (c_t - k_t \alpha) dt \right]. \tag{4}$$

A standard argument in this setting implies that the optimal contract must implement no stealing, that is, $a = 0$. In addition, without loss of generality and for analytic convenience, we can restrict attention to contracts in which $h = 0$ and $\tilde{c} = c$, that is, the principal saves for the agent.⁷ Of course, the optimal contract has many equivalent and more natural forms, in which the agent maintains savings, but all these forms can be deduced easily from the optimal contract with $\tilde{c} = c$.

We say a contract $\mathcal{C} = (c, k)$ is *admissible* if (1) utility $U_0^{c,0}$ is finite, and (2)⁸

$$\mathbb{E}^Q \left[\int_0^\infty e^{-rt} c_t dt \right] < \infty, \quad \mathbb{E}^Q \left[\int_0^\infty e^{-rt} k_t dt \right] < \infty. \tag{5}$$

We say an admissible contract \mathcal{C} is *incentive compatible* if

$$(c, 0) \in \arg \max_{(\tilde{c}, a) \in \mathbb{S}(\mathcal{C})} U_0^{\tilde{c}, a}.$$

Let \mathbb{IC} be the set of incentive compatible contracts. For an initial utility u_0 for the agent, an incentive compatible contract is *optimal* if it minimizes the cost of delivering utility u_0 to the agent

$$\begin{aligned} v_0 &= \min_{(c, k)} J_0 \\ \text{s.t. } & U_0^{c,0} \geq u_0 \\ & (c, k) \in \mathbb{IC}. \end{aligned} \tag{6}$$

We need several parameter restrictions to make the problem well-defined and avoid infinite profits/utility. We assume throughout that

$$\rho > r(1 - \gamma). \tag{7}$$

If this condition fails, which can happen only if $\gamma < 1$, the agent can obtain infinite utility simply by investing at the risk-free rate $r > 0$. Also, if α is too large, the principal’s profit can be infinite. We assume the following necessary condition throughout:

$$\alpha \leq \bar{\alpha} \equiv \phi \sigma \gamma \sqrt{\frac{2}{1 + \gamma}} \sqrt{\frac{\rho - r(1 - \gamma)}{\gamma}}. \tag{8}$$

If this fails, the principal can get an infinite profit through simple stationary contracts described in Section 3.6. The exact bound on α to guarantee that the profit from the optimal

⁷Lemma O.1 in the Online Appendix (Di Tella and Sannikov (2021)) establishes this in a more general setting with both aggregate risk and hidden investment.

⁸This assumption ensures the principal’s objective function is well-defined.

dynamic contract is finite can be found only numerically, but the following conditions are sufficient:⁹

$$\alpha \leq \underline{\alpha} \equiv \begin{cases} \phi \sigma \gamma \sqrt{\frac{1}{2\gamma}} \sqrt{\frac{\rho - r(1 - \gamma)}{\gamma}} & \text{if } \gamma \geq 1/2, \\ \phi \sigma \gamma \sqrt{2(1 - \gamma)} \sqrt{\frac{\rho - r(1 - \gamma)}{\gamma}} & \text{if } \gamma \leq 1/2, \end{cases} \quad (9)$$

where $\underline{\alpha} < \bar{\alpha}$ for all γ .

REMARK: Notice that we build the model with the aim of tractability, and to highlight our main object of interest, hidden savings. CRRA preferences together with a scalable investment technology give us scale invariance. In particular, if contract (c, k) gives utility u_0 to the agent and has cost v_0 to the principal, then the scaled version $(\lambda c, \lambda k)$ has utility $\lambda^{1-\gamma} u_0$ and cost λv_0 . If the former contract is optimal for utility u_0 , then the latter is optimal for utility $\lambda^{1-\gamma} u_0$. The scale invariance property reduces the problem by one dimension—the utility dimension—and allows us to highlight the dimension central to our interest—the dimension of the precautionary savings motive, as the analysis of the next section shows.

3. SOLVING THE MODEL

We solve the model as follows. We first derive necessary first-order incentive-compatibility conditions for the agent's fund diversion and savings choices, using two appropriate state variables: the agent's continuation utility and consumption level. We introduce a change of variables that takes advantage of the homothetic structure of the problem, and provide a sufficient condition for global incentive compatibility in terms of these transformed variables. It is a condition on how these variables depend on returns under the agent's recommended strategy $(c, 0)$, that is, it is an on-path condition, but it allows us to bound the agent's payoffs off-path, after arbitrary deviations.

We then solve the principal's *relaxed problem*, minimizing cost subject to only first-order conditions, as a control problem, and characterize the solution with an HJB equation. We show that the solution to the relaxed problem satisfies the sufficient condition for global incentive compatibility, so we obtain the optimal contract. More generally, the sufficient condition identifies a whole class of globally incentive compatible contracts, and is useful in a broader context.

3.1. Local Incentive Compatibility

We can express the local incentive compatibility conditions in terms of two state variables: the agent's continuation utility and his consumption. Continuation utility is defined as

$$U_t^{c,0} = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} \frac{c_s^{1-\gamma}}{1-\gamma} ds \right].$$

First, we obtain the law of motion for the agent's continuation utility.

⁹See Lemma O.8 and Lemma O.5 in the Online Appendix (Di Tella and Sannikov (2021)) for details.

LEMMA 1: For any admissible contract $\mathcal{C} = (c, k)$, the agent’s continuation utility $U^{c,0}$ satisfies

$$dU_t^{c,0} = \left(\rho U_t^{c,0} - \frac{c_t^{1-\gamma}}{1-\gamma} \right) dt + \Delta_t \underbrace{(dR_t - (\alpha + r) dt)}_{\substack{\sigma dZ_t \\ \text{if } a_t=0}} \tag{10}$$

for some stochastic process Δ .

Faced with this contract, the agent might consider stealing and immediately consuming the proceeds, that is, following a strategy $(c + \phi k a, a)$ for some a , which results in savings $h = 0$. The agent adds $\phi k_t a_t$ to his consumption, but reduces the observed returns dR_t , and therefore his continuation utility $U_t^{c,0}$, by $\Delta_t a_t$. For $a_t = 0$ to be optimal, we need

$$0 \in \arg \max_{\tilde{a} \geq 0} \frac{(c_t + \phi k_t \tilde{a})^{1-\gamma}}{1-\gamma} - \Delta_t \tilde{a} \tag{11}$$

everywhere. Taking the first-order condition yields

$$\Delta_t \geq c_t^{-\gamma} \phi k_t. \tag{12}$$

We need to give the agent some “skin in game” by exposing him to risk. This is costly because the principal is risk-neutral with respect to Z so he would like to provide full insurance to the agent.

For hidden savings, the first-order condition to make sure the agent does not want to save is the Euler equation: the condition that the agent’s discounted marginal utility $e^{(r-\rho)t} c_t^{-\gamma}$ is a supermartingale.¹⁰

The following lemma summarizes the necessary local incentive conditions, expressed in terms of variables $U_t^{c,0}$ and c_t . We present sufficient incentive compatibility conditions in Section 3.3.¹¹

LEMMA 2: If $\mathcal{C} = (c, k)$ is an incentive compatible contract, then (12) must hold and the agent’s consumption must satisfy

$$\frac{dc_t}{c_t} = \left(\frac{r - \rho}{\gamma} + \frac{1 + \gamma}{2} (\sigma_t^c)^2 \right) dt + \sigma_t^c \underbrace{\frac{1}{\sigma} (dR_t - (\alpha + r) dt)}_{\substack{dZ_t \\ \text{if } a_t=0}} + dL_t \tag{13}$$

for some stochastic process σ^c and a weakly increasing stochastic process L .

Equation (13) gives a lower bound on the growth rate of the agent’s consumption. The first term $\frac{r-\rho}{\gamma}$ captures the benefit of postponing consumption without risk, given by the

¹⁰Notice also that if $c_t = 0$ were allowed, then $c_t^{-\gamma} = \infty$; hence this possibility can only arise at the beginning of the contract until a stopping time, and only with zero capital according to (12). Thus, giving the agent $c_t = 0$ at the beginning is equivalent to delaying the contract, which can never be optimal. Hence, our assumption that $c_t > 0$ at all times is without loss of generality.

¹¹It can be shown that condition (12) is necessary and sufficient to deter any deviations (c, a) with just fund diversion, and condition (13) is necessary and sufficient to deter any deviation $(\tilde{c}, 0)$ with hidden savings. We could make these statements formal, but this still leaves us concerned about double deviations, which we can rule out with an additional sufficient condition.

risk-free rate r , the discount rate ρ , and the elasticity of intertemporal substitution $1/\gamma$. The second term $\frac{1+\gamma}{2}(\sigma_t^c)^2$ captures the agent's precautionary saving motive. A risky consumption profile induces the agent to postpone consumption to self-insure, resulting in a steeper expected consumption profile. The process L in equation (13) allows consumption to grow at a higher rate than the one that makes $e^{(r-\rho)t}c_t^{-\gamma}$ a martingale. This is consistent with incentive compatibility because the agent cannot borrow against future consumption. However, we shall see that in the optimal contract $dL = 0$, so the Euler equation holds with equality.

There is an important interaction between the precautionary motive and incentives to divert funds. Notice how the private benefit of fund diversion depends on the marginal utility of consumption, $c_t^{-\gamma}$. The agent's consumption, in turn, depends on future risk. Exposing the agent to risk today though Δ_t helps provide incentives against diversion. But exposing the agent to risk in his consumption in the future, σ_{t+s}^c , creates a precautionary motive that reduces c_t , which tightens (12) and makes fund diversion today more attractive. The intuition is that when faced with future risk, the agent places a large value on hidden savings he can use to self-insure. Important dynamic properties of the optimal contract arise out of the principal's attempt to manipulate the agent's precautionary motive in order to relax the IC constraints.

Without hidden savings, the principal could control the agent's consumption without worrying about the agent's precautionary motive. In this case, c_t would be the principal's control rather than a state. The principal would choose c_t to minimize the cost of compensating the agent, also taking into account that at any time higher consumption relaxes the constraint (12). The optimal consumption path without hidden savings would satisfy the inverse Euler equation, that is, $e^{(\rho-r)t}c_t^\gamma$ would be a martingale, and c_t would have a lower drift than that required by condition (13).

3.2. A Change of Variables

It is convenient to work with the following transformation of variables that exploits the setting's homothetic structure. For any admissible contract, we can compute

$$x_t = ((1 - \gamma)U_t^{c,0})^{\frac{1}{1-\gamma}} > 0,$$

$$\hat{c}_t = \frac{c_t}{x_t} > 0.$$

Variable x_t is just a monotone transformation of continuation utility, but it is measured in consumption units (up to a constant). As a result, \hat{c}_t measures how front loaded the agent's consumption is. The state \hat{c}_t is related to the agent's precautionary saving motive. If the agent faces risk looking forward, he will want to postpone consumption in an attempt to self-insure, leading to lower \hat{c}_t . While x_t can take any positive value, \hat{c}_t has an upper bound.

LEMMA 3: *For any incentive compatible contract $\mathcal{C} = (c, k)$, we have for all t*

$$\hat{c}_t \leq \hat{c}_h \equiv \left(\frac{\rho - r(1 - \gamma)}{\gamma} \right)^{\frac{1}{1-\gamma}} > 0. \quad (14)$$

If ever $\hat{c}_t = \hat{c}_h$, then the continuation contract satisfies $k_{t+s} = 0$ and $\hat{c}_{t+s} = \hat{c}_h$ at all future times $t + s$ and gives the agent a unique deterministic consumption path with growth $(r - \rho)/\gamma$. The continuation contract has cost $\hat{v}_h x_t$ to the principal, where $\hat{v}_h \equiv \hat{c}_h^\gamma$.

If in addition σ^c is bounded, we have

$$\hat{c}_t \leq \mathbb{E}_t^{\tilde{P}} \left[\int_t^\infty e^{-\int_t^s (\frac{\rho-r(1-\gamma)}{\gamma} - \frac{1-\gamma}{2}(\sigma_u^c)^2) du} ds \right]^{-\frac{1}{1-\gamma}}, \tag{15}$$

with equality if $L_t = 0$ always, where \tilde{P} is an equivalent measure such that $Z_t - \int_0^t (1-\gamma)\sigma_s^c ds$ is a \tilde{P} -martingale.

REMARK: In the optimal contract $L_t = 0$ and (15) holds as an equality. This expression captures exactly how future risk exposure translates to precautionary motive and to a level of \hat{c}_t below \hat{c}_h . In the special case where $\sigma_t^c = \sigma^c$ is constant, we have $\hat{c}_t = (\frac{\rho-r(1-\gamma)}{\gamma} - \frac{1-\gamma}{2}(\sigma^c)^2)^{\frac{1}{1-\gamma}}$ which is decreasing in σ^c .

The upper bound \hat{c}_h corresponds to autarky, where the agent saves and consumes on his own without managing any capital. It can also be interpreted as the point of termination or the agent’s retirement. The corresponding contract minimizes the agent’s precautionary motive because it is fully safe. At the same time, this contract is very costly because it does not give any capital to the agent and so does not take advantage of the excess return $\alpha > 0$. A value $\hat{c}_t < \hat{c}_h$ means the agent expects to manage capital and be exposed to risk in the future, and we may think of \hat{c}_t as indexing how risky the continuation contract is for the agent.

Using Ito’s lemma we can obtain the laws of motion for x_t and \hat{c}_t from (10) and (13). Using the normalization $\Delta_t \sigma / U_t^{c,0} = (1-\gamma)\sigma_t^x$, we obtain

$$\frac{dx_t}{x_t} = \underbrace{\left(\frac{\rho - \hat{c}_t^{1-\gamma}}{1-\gamma} + \frac{\gamma}{2}(\sigma_t^x)^2 \right)}_{\mu_t^x} dt + \sigma_t^x dZ_t \tag{16}$$

and

$$\begin{aligned} \frac{d\hat{c}_t}{\hat{c}_t} = & \underbrace{\left(\frac{\hat{c}_t^{1-\gamma} - \hat{c}_h^{1-\gamma}}{1-\gamma} + \frac{(\sigma_t^x)^2}{2} + \gamma\sigma_t^x\sigma_t^{\hat{c}} + \frac{1+\gamma}{2}(\sigma_t^{\hat{c}})^2 \right)}_{\mu_t^{\hat{c}}} dt \\ & + \sigma_t^{\hat{c}} dZ_t + dL_t, \quad dL_t \geq 0, \end{aligned} \tag{17}$$

with $\sigma_t^{\hat{c}} = \sigma_t^c - \sigma_t^x$. The constraint (12) can be rewritten as

$$\sigma_t^x \geq \hat{c}_t^{-\gamma} \hat{k}_t \phi \sigma, \tag{18}$$

where $\hat{k}_t = \frac{k_t}{x_t}$. It will always be binding because conditional on σ_t^x it is always better to give the agent more capital to manage. The IC constraint (18) establishes a link between the agent’s exposure to risk σ_t^x and the amount of capital he manages \hat{k}_t , mediated by the marginal utility $\hat{c}_t^{-\gamma}$, which captures the precautionary savings motive.

In terms of variables x_t , \hat{c}_t , and k_t , we call any admissible contract *locally incentive compatible* if (17) and (18) hold and $\hat{c}_t \leq \hat{c}_h$. Notice that equation (16) follows automatically from the definition of x_t .

In the other direction, given a pair of processes $x_t > 0$ and $\hat{c}_t > 0$ satisfying (16) and (17) with $\hat{c}_t \leq \hat{c}_h$ and $\sigma_t^x \geq 0$, we can build a locally incentive compatible contract (c, k) , with $c_t = \hat{c}_t x_t > 0$ and $k_t = x_t \sigma_t^x \hat{c}_t^\gamma / (\phi \sigma) \geq 0$. Under technical conditions, the contract (c, k) is admissible and delivers utility $U_t^{c,0} = \frac{x_t^{1-\gamma}}{1-\gamma}$ under good behavior. It is therefore locally incentive compatible.

LEMMA 4: *Let $x > 0$ and $\hat{c} > 0$ be stochastic processes satisfying (16) and (17) with bounded volatility $\sigma^x \geq 0$, and with \hat{c} bounded away from zero and above by \hat{c}_h .*

Then the contract $\mathcal{C} = (c, k)$ with $c_t = \hat{c}_t x_t > 0$ and $k_t = x_t \sigma_t^x \hat{c}_t^\gamma / (\phi \sigma) \geq 0$ delivers utility $U_t^{c,0} = \frac{x_t^{1-\gamma}}{1-\gamma}$ if the agent follows strategy $(c, 0)$. The contract \mathcal{C} is admissible and, therefore, locally incentive compatible, if and only if $\mathbb{E}[\int_0^\infty e^{-rt} x_t dt] < \infty$.

3.3. Sufficient Conditions for Global Incentive Compatibility

Every incentive compatible contract is locally incentive compatible. Here, we provide sufficient conditions for incentive compatibility for any locally incentive compatible contract, that is, an admissible contract that satisfies the local IC constraints (17), (18), and $\hat{c}_t \leq \hat{c}_h$. We shall see that these sufficient conditions are guaranteed to hold in the relaxed optimal contract, which proves the validity of the first-order approach, but are more general than that and can be used to check incentive compatibility of many other contracts of interest.

While (17) and (18) ensure that neither stealing and immediately consuming, nor secretly saving without stealing are attractive on their own, they leave open the possibility that a double deviation—stealing and saving the proceeds for later—could be attractive to the agent. To see how this can happen, notice that since stealing makes bad outcomes more likely, it increases the expected marginal utility of consumption in the future $\mathbb{E}_t^a[e^{(r-\rho)u} c_{t+u}^{-\gamma}]$.¹² Saving the stolen funds for consumption later could therefore be very attractive. However, hidden savings have decreasing marginal value (the first dollar yields $c_t^{-\gamma}$, the second one less than that), which depends on the agent's precautionary saving motive. This observation allows us to derive a sufficient condition to rule out profitable double deviations.

THEOREM 1: *Let $\mathcal{C} = (c, k)$ be a locally incentive compatible contract with \hat{c} bounded away from zero and bounded volatilities σ^x and $\sigma^{\hat{c}}$. Suppose that the contract satisfies the following property:*

$$\sigma_t^{\hat{c}} \leq 0. \quad (19)$$

Then for any feasible strategy (\tilde{c}, a) , with associated hidden savings h , we have the following upper bound on the agent's utility, after any history

$$U_t^{\tilde{c},a} \leq \left(1 + \frac{h_t}{x_t} \hat{c}_t^{-\gamma}\right)^{1-\gamma} U_t^{c,0}. \quad (20)$$

In particular, since $h_0 = 0$, for any feasible strategy $U_0^{\tilde{c},a} \leq U_0^{c,0}$, and the contract \mathcal{C} is therefore incentive compatible.

¹²In other words, even if $e^{(r-\rho)t} c_t^{-\gamma}$ is a martingale under P , it might be a submartingale under P^a .

Theorem 1 shows that condition (19) is sufficient for global incentive compatibility by providing a closed-form explicit upper bound (20) on the agent’s off-equilibrium payoff for any savings level $h_t \geq 0$. According to (20), if the agent does not have any hidden savings, $h_t = 0$, the most utility he could get is $U_t^{c,0}$, that is, the utility level he obtains from “good behavior” $(c, 0)$. Hence, the contract is incentive compatible. However, if the agent had somehow accumulated hidden savings in the past, he would want to deviate from $(c, 0)$ in the future, at the very least to increase his consumption and attain a greater utility. Inequality (20) bounds the utility the agent can get, and the bound tightens as \hat{c}_t rises and the agent’s precautionary saving motive decreases. The bound is consistent with the intuition that the value of hidden savings becomes lower as the agent’s precautionary saving motive decreases (however, remember that this is just an upper bound on achievable utility).

The sufficient condition $\sigma_t^{\hat{c}} \leq 0$ can be understood as follows. Hidden savings become more valuable when the agent faces more risk, that is, has a higher precautionary saving motive. With $\sigma_t^{\hat{c}} \leq 0$, the contract becomes less risky for the agent after bad outcomes (hidden savings become less valuable). Since stealing makes bad outcomes more likely, if the agent steals and saves for later, he expects to have a hidden dollar when it is least valuable to him. This makes double deviations unprofitable.

As we show below, the sufficient condition $\sigma_t^{\hat{c}} \leq 0$ is guaranteed to hold in the relaxed optimal contract. The principal wants to constrain the agent’s precautionary saving motive and the most efficient way to do this is by reducing the agent’s risk exposure after bad outcomes. As it happens, this property is sufficient for global incentive compatibility, which shows the validity of the first-order approach. But we would like to emphasize that Theorem 1 identifies $\sigma_t^{\hat{c}} \leq 0$ as a general sufficient condition for incentive compatibility, without even assuming that the contract is recursive in variables x and \hat{c} . These variables are well-defined for any admissible contract, and their laws of motion do not need to be Markovian for Theorem 1 to apply. Condition $\sigma_t^{\hat{c}} \leq 0$ can be used to verify global incentive compatibility of suboptimal locally incentive compatible contracts of interest. For example, it implies that *stationary* contracts that we discuss below are also globally incentive compatible.

3.4. The Relaxed Problem

The *relaxed problem* consists of minimizing cost within the class of locally incentive compatible contracts, and we call a solution to the relaxed problem a *relaxed optimal contract*.

We approach the relaxed problem as an optimal control problem with states $x_t > 0$ and $\hat{c}_t > 0$ satisfying (16), (17), and $\hat{c}_t \leq \hat{c}_h$, controls $\sigma_t^x \geq 0$, $\sigma_t^{\hat{c}}$, and $dL_t \geq 0$, and initial condition $x_0 = ((1 - \gamma)u_0)^{\frac{1}{1-\gamma}}$. The objective function is

$$\mathbb{E}^Q \left[\int_0^\infty e^{-rt} x_t \left(\hat{c}_t - \hat{c}_t^\gamma \frac{\alpha}{\phi \sigma} \sigma_t^x \right) dt \right],$$

where the cost flow already incorporates the binding IC constraint (18).

To understand the relaxed problem, notice that the initial value of x_0 is determined by the initial utility u_0 , but there is no corresponding constraint for \hat{c}_0 —the principal can set \hat{c}_0 freely to minimize the cost. A low \hat{c}_t means the agent expects a very risky consumption in the future, as captured by (15), so \hat{c}_t determines the principal’s “budget for risk exposure”. If \hat{c}_t ever reached the upper bound \hat{c}_h , the principal would have to give the agent

the perfectly safe autarky contract without any capital, which is a costly way to deliver utility. On the other hand, a low \hat{c}_t means the agent has large incentives to divert funds, as captured by the IC constraint (18). Hence, the choice of \hat{c}_0 takes into account the cost of committing to lower risk in the future and the benefit of improved incentives now.

After \hat{c}_0 is set the principal chooses σ_t^x and $\sigma_t^{\hat{c}}$ dynamically. Higher σ_t^x today, that is, exposing the agent to more risk, allows the principal to give more capital to the agent and earn the excess return α . However, this is costly because (a) the agent is risk averse and must be compensated with more utility in the future—the drift of x_t is increasing in σ_t^x —and (b) it eats up the budget for risk exposure in the future—the term $(\sigma_t^x)^2/2$ raises the drift of \hat{c}_t .

The principal can choose $\sigma_t^{\hat{c}}$ to mitigate the latter effect: when $\sigma_t^x > 0$, setting $\sigma_t^{\hat{c}} < 0$ reduces the drift of \hat{c}_t , preserving the budget for risk exposure. The intuition is that a lower volatility of consumption, $\sigma_t^c = \sigma_t^x + \sigma_t^{\hat{c}}$, weakens the agent's precautionary motive. This is the first of two forces that guide the dynamic choice of $\sigma_t^{\hat{c}}$. The second force is that the benefit of a larger budget of risk exposure (lower \hat{c}_t) is proportional to x_t due to scale invariance, so the principal prefers to have a lower \hat{c}_t when x_t goes up. This also implies $\sigma_t^{\hat{c}} < 0$.

As we show below analytically, $\sigma_t^{\hat{c}} < 0$ is a key property of the relaxed optimal contract: it implies that the contract gets safer after poor outcomes. Theorem 1 shows that $\sigma_t^{\hat{c}} \leq 0$ is a sufficient condition that ensures global incentive compatibility, so the solution of the relaxed problem is indeed the optimal contract.

3.5. The Solution to the Relaxed Problem Gives the Optimal Contract

In this subsection, we characterize the relaxed optimal contract, and use Theorem 1 to show it is globally incentive compatible and, therefore, an optimal contract.

Because preferences are homothetic and the principal's objective is linear, we know the principal's cost function in the relaxed problem takes the form $v(x, \hat{c}) = \hat{v}(\hat{c})x$. We will sometimes write $\hat{v}_t = \hat{v}(\hat{c}_t)$, and \hat{v} instead of $\hat{v}(\hat{c})$.

Notice that since we could always raise \hat{c}_t using $dL_t > 0$, we know that $\hat{v}(\hat{c})$ must be weakly increasing. In fact, we show below that the cost function $\hat{v}(\hat{c})$ has a flat region $(0, \hat{c}_l)$ where the optimal contract does not spend any time, followed by a strictly increasing region (\hat{c}_l, \hat{c}_h) with $dL_t = 0$ in which \hat{c}_t stays over the course of the optimal contract.

The HJB equation in the strictly increasing region with $\hat{v}'(\hat{c}) > 0$ is

$$\begin{aligned} r\hat{v} = & \min_{\sigma^x \geq 0, \sigma^{\hat{c}}} \hat{c} - \sigma^x \hat{c}^\gamma \frac{\alpha}{\phi \sigma} + \hat{v} \left(\frac{\rho - \hat{c}^{1-\gamma}}{1-\gamma} + \frac{\gamma}{2} (\sigma^x)^2 \right) \\ & + \hat{v}' \hat{c} \left(\frac{\hat{c}^{1-\gamma} - \hat{c}_h^{1-\gamma}}{1-\gamma} + \frac{(\sigma^x)^2}{2} + (1+\gamma)\sigma^x \sigma^{\hat{c}} + \frac{1+\gamma}{2} (\sigma^{\hat{c}})^2 \right) + \frac{\hat{v}''}{2} \hat{c}^2 (\sigma^{\hat{c}})^2. \end{aligned} \quad (21)$$

Even though we have two state variables, \hat{c} and x , the HJB equation boils down to a second-order ODE in \hat{c} . This is a feature of homothetic preferences and linear technology that makes the problem more tractable.

In the flat region with $\hat{v}'(\hat{c}) = 0$, the HJB equation (21) must hold as a \leq inequality,

$$A(\hat{c}, \hat{v}) \equiv \min_{\sigma^x \geq 0} \hat{c} - \sigma^x \hat{c}^\gamma \frac{\alpha}{\phi \sigma} - r\hat{v} + \hat{v} \left(\frac{\rho - \hat{c}^{1-\gamma}}{1-\gamma} + \frac{\gamma}{2} (\sigma^x)^2 \right) > 0. \quad (22)$$

The optimal contract does not spend any time in this region.

The following theorem characterizes the solution to the relaxed problem and shows that it is globally incentive compatible and, therefore, an optimal contract.

THEOREM 2: *The relaxed problem has the following properties:*

1. *The cost function $\hat{v}(\hat{c})$ has a flat portion on $(0, \hat{c}_l)$ and a strictly increasing C^2 portion on (\hat{c}_l, \hat{c}_h) , for some $\hat{c}_l \in (0, \hat{c}_h)$. The HJB equation (21) holds with equality for $\hat{c} \in (\hat{c}_l, \hat{c}_h)$. For $\hat{c} < \hat{c}_l$, we have $\hat{v}(\hat{c}) = \hat{v}(\hat{c}_l) \equiv \hat{v}_l$ and the HJB holds as an inequality, $A(\hat{c}, \hat{v}_l) > 0$.*
2. *At \hat{c}_l , we have $\hat{v}'(\hat{c}_l) = 0$, $\hat{v}''_+(\hat{c}_l) > 0$, and $A(\hat{c}_l, \hat{v}_l) = 0$. The cost function satisfies $\hat{v}(\hat{c}) < \hat{c}^\gamma$ for all $\hat{c} \in [\hat{c}_l, \hat{c}_h)$, with $\hat{v}(\hat{c}_h) = \hat{v}_h = \hat{c}_h^\gamma$.*
3. *The state variables x_t and \hat{c}_t follow the laws of motion (16) and (17) with bounded volatilities $\sigma_t^x > 0$ and $\sigma_t^{\hat{c}} < 0$ for all $t > 0$, and $dL_t = 0$ always, so the Euler equation holds as an equality. The state \hat{c}_t starts at $\hat{c}_0 = \hat{c}_l$, with $\mu_0^{\hat{c}} > 0$ and $\sigma_0^{\hat{c}} = 0$, and immediately moves into the interior of the domain never reaching either boundary, that is, $\hat{c}_t \in (\hat{c}_l, \hat{c}_h)$ for all $t > 0$.*
4. *\hat{c}_t does not have a stationary distribution. In the long-run, \hat{c}_t spends almost all the time near \hat{c}_h ,*

$$\frac{1}{t} \int_0^t 1_{\{\hat{c}_s > \hat{c}_h - \epsilon\}}(\hat{c}_s) ds \rightarrow 1 \quad \text{a.s. } \forall \epsilon > 0.$$

However, \hat{c}_h is never reached, even in infinite time, $P\{\hat{c}_t \rightarrow \hat{c}_h\} = 0$.

5. *Since the relaxed optimal contract satisfies the sufficient condition in Theorem 1, it is incentive compatible and, therefore, an optimal contract.*

Let us now highlight some features of the cost function and the relaxed optimal contract from Theorem 2. Figures 1 and 2 illustrate these properties in a numerical example, discussed below in Section 7. The shape of the cost function, with a flat portion up to \hat{c}_l followed by an increasing portion, reflects the optimal choice of $\hat{c}_0 = \hat{c}_l$. It would be sub-optimal to run the contract with $\hat{c}_l < \hat{c}_l$ that makes incentive provision excessively difficult. Yet, since the principal can raise \hat{c} to the optimal level of \hat{c}_l using the process $dL_t \geq 0$, the




FIGURE 1.—The cost function $\hat{v}(\hat{c})$ of the optimal contract. The starting point of the optimal contract is indicated by a dot. Parameters: $\rho = r = 5\%$, $\alpha = 1.7\%$, $\gamma = 1/3$, $\phi\sigma = 0.2$.




FIGURE 2.—The drift, $\mu^{\hat{c}}$ and μ^x , and volatility, $\sigma^{\hat{c}}$ and σ^x , of the state variables \hat{c} and x under the optimal contract. The starting point of the optimal contract is indicated by a dot. Parameters: $\rho = r = 5\%$, $\alpha = 1.7\%$, $\gamma = 1/3$, $\phi\sigma = 0.2$.

cost function is not decreasing but flat over $(0, \hat{c}_l)$. As \hat{c}_l moves above \hat{c}_l , the contract must give the agent an inefficiently low risk exposure, with the end point of \hat{c}_h that corresponds to the safe autarky contract with no capital. As a result, the cost function $\hat{v}(\hat{c})$ is increasing over (\hat{c}_l, \hat{c}_h) . The marginal cost of delivering utility, $\partial v(x, \hat{c})/\partial U = \hat{v}(\hat{c})x^\gamma$, is below the inverse marginal utility of consumption, c^γ , that is, $\hat{v} < \hat{c}^\gamma$, because higher \hat{c} not only delivers utility to the agent, but also relaxes the IC constraint.

An important feature of the relaxed optimal contract is that $\sigma_t^{\hat{c}} < 0$, which can be seen in the bottom left panel of Figure 2. After bad performance \hat{c}_t rises, so the agent expects a less risky consumption in the future. There are two reasons to set $\sigma_t^{\hat{c}} < 0$. First, it reduces the volatility of the agent’s consumption, $\sigma_t^c = \sigma_t^x + \sigma_t^{\hat{c}}$ and, therefore, helps relax the precautionary motive. Second, since the cost is proportional to promised utility, $\hat{v}(\hat{c}_t)x_t$, the principal prefers to use the least costly continuation contracts with low \hat{c}_t after good shocks when he must deliver more utility to the agent. This also implies $\sigma_t^{\hat{c}} < 0$. It follows, using Theorem 1, that the relaxed optimal contract is fully incentive compatible and, therefore, an optimal contract.

The top left panel of Figure 2 shows the drift of \hat{c} , which is positive at \hat{c}_l . This means that while the contract starts at the optimal $\hat{c}_0 = \hat{c}_l$, it immediately moves into the interior of the domain. The reason for this is that, since the Euler equation is forward-looking, promising an inefficiently safe contract at a future time t (higher \hat{c}_t) relaxes the IC constraint for every period $s < t$, so it makes sense to backload distortions. This intertemporal trade-off allows the principal to relax the IC constraints today at the cost of an inefficiently safe contract in the future.

The optimal contract starts at $\hat{c}_0 = \hat{c}_l$ and moves up over time and after bad returns. Given the behavior of $\mu^{\hat{c}}$, one may suspect that there is a stationary distribution of \hat{c} concentrated near the point of zero drift. It turns out that this is not the case. Even though we can show analytically that the drift of \hat{c} is negative near \hat{c}_h , the asymptotic properties of $\mu^{\hat{c}}$ and $\sigma^{\hat{c}}$ near \hat{c}_h are such that the process \hat{c} slows down but never reaches \hat{c}_h . As a result, there is no stationary distribution. The principal backloads distortions in order to relax the agent’s IC constraints, so the contract starts at \hat{c}_l and in the long-run spends most of the time near the upper boundary corresponding to autarky \hat{c}_h . However, $P\{\hat{c}_t \rightarrow \hat{c}_h\} = 0$. No

matter how close any path gets to \hat{c}_h , it will always leave that neighborhood at some point in the future, but this happens progressively less and less often for any path. Part (4) of Theorem 2 formalizes this asymptotic result.

Finally, to understand the choice of σ^x , look at the FOC for σ^x . There is a myopic and a dynamic motive for exposing the agent to risk through σ^x ,

$$\sigma^x = \underbrace{\frac{\alpha}{\gamma\phi\sigma} \frac{\hat{c}^\gamma}{\hat{v}}}_{\text{myopic}} - \underbrace{\frac{\hat{v}'}{\hat{v}} \frac{\hat{c}}{\gamma}}_{\text{dynamic}} ((1 + \gamma)\sigma^{\hat{c}} + \sigma^x). \tag{23}$$

The myopic motive trades off the excess return on capital, $\alpha > 0$, against the higher future utility required to compensate the agent for the risk. The dynamic component takes into account that higher σ^x eats up the budget for risk exposure (increases the drift of \hat{c}). This is costly because the cost $\hat{v}(\hat{c})$ is increasing in \hat{c} . At $t = 0$, we know that $\hat{v}' = 0$ because \hat{c}_1 is chosen optimally, so the dynamic motive disappears and we choose σ^x myopically without taking into account the effect on the agent’s precautionary motive.

What happens to the drift and volatility of x , σ^x and μ^x , when \hat{c} increases? The top and bottom right panels of Figure 2 show their behavior in the numerical solution. While we do not have a proof, the following properties hold across all solutions we computed, and help understand the behavior of the optimal contract. The agent’s risk exposure σ^x declines with \hat{c} . This is intuitive: the principal exposes the agent to less risk when the “budget for risk” exposure gets depleted. Lower risk σ^x , in combination with higher consumption \hat{c} , imply a lower expected growth in the agent’s utility x_t , so μ^x declines in \hat{c} . Thus, since $\sigma^{\hat{c}} < 0$, the agent is punished after bad outcomes in part by slower growth.

According to our original definition, the contract has to specify the pair (c_t, k_t) for any history of output. We have $c_t = \hat{c}_t x_t$ and $k_t = \sigma_t^x \frac{\hat{c}_t^\gamma}{\phi\sigma} x_t$, so \hat{c}_t determines the drift and volatility of these variables. These are shown in the bottom panels of Figure 3. Variable \hat{c}_t reflects the principal’s commitment to make the contract more secure, and the top left

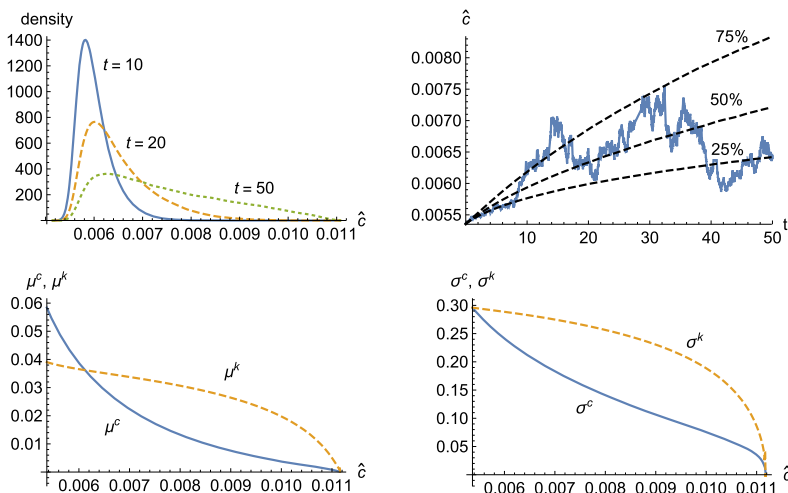


FIGURE 3.—The distribution of \hat{c}_t at different time points (top left panel), a sample path of \hat{c}_t (top right), drifts $\mu^c(\hat{c})$ and $\mu^k(\hat{c})$, and volatilities $\sigma^c(\hat{c})$ and $\sigma^k(\hat{c})$ (bottom panels). Parameters: $\rho = r = 5\%$, $\alpha = 1.7\%$, $\gamma = 1/3$, $\phi\sigma = 0.2$.

panel shows the distributions of values of \hat{c}_t at three different time points. The top right panel shows a sample path of \hat{c}_t starting from \hat{c}_l , relative to the whole probability distribution. Initially, consumption and capital grow fast, and the agent faces a large amount of risk in consumption and capital. Over time, expected growth and risk decline as \hat{c}_t moves up. The volatility of consumption is lower than that of capital. After poor returns, capital declines faster than consumption, because the agent faces a safer contract going forward.

We close this section with a verification theorem. Theorem 2 shows that the principal’s cost function satisfies the HJB equation, but how do we know the converse? The following theorem shows that if an appropriate solution to the HJB has been found, for example, numerically, then it must be the true cost function and we can use it to build a globally incentive compatible optimal contract.

Given an appropriate solution to the HJB equation, we can identify controls σ^x and $\sigma^{\hat{c}}$ as functions of \hat{c} , and use those to build a candidate optimal contract C^* . Specifically, let x^* and \hat{c}^* be solutions to (16) and (17) with $\sigma_t^{x^*} = \sigma^x(\hat{c}_t^*)$, $\sigma_t^{\hat{c}^*} = \sigma^{\hat{c}}(\hat{c}_t^*)$ and $dL_t = 0$, starting from initial values $x_0^* = ((1 - \gamma)u_0)^{\frac{1}{1-\gamma}}$ and $\hat{c}_0^* = \hat{c}_l$. We then construct the candidate contract $C^* = (c^*, k^*)$ with $c^* = \hat{c}^*x^*$ and $k^* = \sigma^{x^*} \frac{(\hat{c}^*)^\gamma}{\phi\sigma} x^*$.

THEOREM 3—Verification theorem: *Let $\hat{v}(\hat{c}) : [\hat{c}_l, \hat{c}_h] \rightarrow [\hat{v}_l, \hat{v}_h]$ be a strictly increasing C^2 solution to the HJB equation (21) for some $\hat{c}_l \in (0, \hat{c}_h)$, such that $\hat{v}_l \equiv \hat{v}(\hat{c}_l) \in (0, \hat{v}_h]$, $\hat{v}'(\hat{c}_l) = 0$, $\hat{v}'(\hat{c}_l) > 0$ and $\hat{v}(\hat{c}_h) = \hat{v}_h$. Assume that for $\hat{c} < \hat{c}_l$ the HJB equation holds as an inequality, $A(\hat{c}, \hat{v}_l) > 0$. Then:*

1. *For any locally incentive compatible contract $C = (c, k)$ that delivers at least utility u_0 to the agent, we have $\hat{v}(\hat{c}_l)((1 - \gamma)u_0)^{\frac{1}{1-\gamma}} \leq J_0(C)$.*
2. *Let C^* be a candidate contract generated by the policy functions of the HJB as described above. If C^* is admissible, then C^* is an optimal contract with cost $J_0(C^*) = \hat{v}(\hat{c}_l)((1 - \gamma)u_0)^{\frac{1}{1-\gamma}}$.*

REMARK: Part (2) requires checking that the candidate optimal contract is admissible. Lemma 4 shows that $\mathbb{E}[\int_0^\infty e^{-rt} x_t dt] < \infty$ is a necessary and sufficient condition.

3.6. Discussion: Hidden Savings and Dynamic Distortions

In the optimal contract, the principal commits to a level of safety summarized by \hat{c}_0 , and then allocates the agent’s risk across future states to minimize the cost of compensation. Initially, \hat{c}_0 minimizes the cost of compensating the agent, and so $\hat{v}'(\hat{c}_0) = 0$. We can say that at time 0, there are no distortions. Afterwards, $\hat{c}_t > \hat{c}_0$ and $\hat{v}'(\hat{c}_t) > 0$, reflecting the principal’s costly commitment to make the contract safer. This commitment benefits incentives before time t , and distorts the principal’s choice of risk after time t .

The optimal contract resolves two trade-offs. First, the standard risk-return trade-off: the principal earns α at the cost of exposing the agent to risk. This is also the trade-off of the standard portfolio choice problem. In addition to this trade-off, after time 0 we have additional distortions because $\hat{v}'(\hat{c}_t) > 0$. When exposing the agent to risk, to respect the law of motion (17), the principal has to raise \hat{c}_t . This carries the marginal cost of $\hat{v}'(\hat{c}_t) > 0$. To sum up, we have the risk-return trade-off as well as distortions to manage the agent’s precautionary motive after time 0.

It is useful to compare the optimal contract with hidden savings to a series of benchmark contracts to better understand these trade-offs. The benchmark contracts are also interesting in their own right.

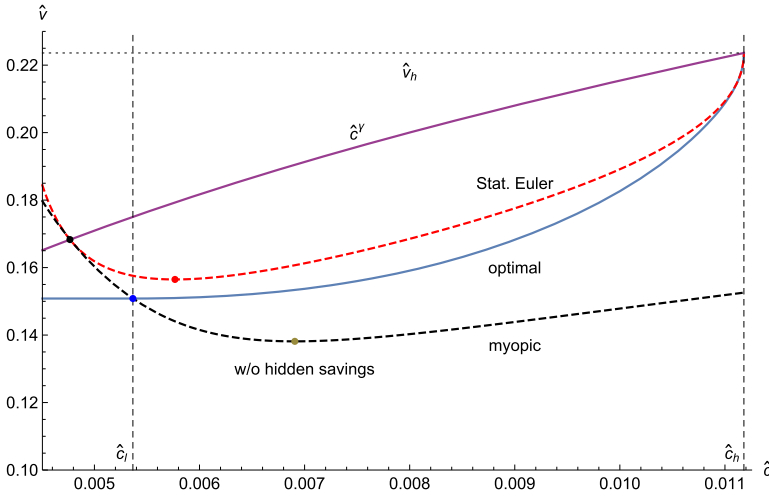


FIGURE 4.—The cost function for the optimal contract $\hat{v}(\hat{c})$ (solid). For reference, dashed lines show the cost of stationary contracts, $\hat{v}_s(\hat{c})$, and contracts with myopic optimization over σ^x , $\hat{v}_m(\hat{c})$. Parameters: $\rho = r = 5\%$, $\alpha = 1.7\%$, $\gamma = 1/3$, $\phi\sigma = 0.2$.

First, the reader may wonder what happens without hidden savings. In this case, the principal determines σ^x purely from the risk-return trade-off, as the choice of \hat{c} is free and unconstrained by the agent’s precautionary motive. For a fixed level of \hat{c} , optimizing over σ^x subject to the law of motion of the agent’s utility (16) and the IC constraint (18), we obtain the cost function $\hat{v}_m(\hat{c})$ that satisfies the HJB equation

$$r\hat{v}_m(\hat{c}) = \min_{\sigma^x \geq 0} \hat{c} - \sigma^x \hat{c}^\gamma \frac{\alpha}{\phi\sigma} + \hat{v}_m(\hat{c}) \left(\frac{\rho - \hat{c}^{1-\gamma}}{1 - \gamma} + \frac{\gamma}{2} (\sigma^x)^2 \right), \tag{24}$$

The optimal choice of σ^x satisfies the first-order condition

$$\sigma^x = \frac{\alpha}{\gamma\phi\sigma} \frac{\hat{c}^\gamma}{\hat{v}_m(\hat{c})}.$$

This condition reflects the pure risk-return trade-off, taking into account how the cost of risk impacts the principal’s value function (hence the ratio $\hat{c}^\gamma/\hat{v}_m(\hat{c})$). We call such choice of σ^x *myopic*, that is, without taking into account the impact on the agent’s precautionary motive (which is irrelevant without hidden savings). $\hat{v}_m(\hat{c})$ is shown as a dashed line in Figure 4.

The optimal contract without hidden savings corresponds to the minimum point of $\hat{v}_m(\hat{c})$, indicated by a dot in Figure 4.¹³ The choice of \hat{c} without hidden savings satisfies the inverse Euler equation, that is, $e^{(\rho-r)t} c_t^\gamma$ must be a martingale, which implies that the agent would like to save if he could (see Lemma O.6).

Once we introduce hidden savings, we are constrained by the Euler equation. For a sharper contrast to the case without hidden savings, it is useful to consider *stationary contracts* with a fixed choice of \hat{c} , with c_t , k_t , and x_t following geometric Brownian motions,

¹³The optimal contract without hidden savings only exists if $\gamma < 1/2$. If $\gamma > 1/2$ the principal can obtain infinite profits for any $\alpha > 0$. See Di Tella (2019). Figure 4 shows results for $\gamma < 1/2$, but hidden savings constrain the principal and make the problem well-defined for any γ .

and $dL_t = 0$. To keep \hat{c} constant while respecting the Euler equation, we must set $\sigma^{\hat{c}} = 0$ and pick σ^x to satisfy

$$\sigma^x = \sigma_s^x(\hat{c}) \equiv \sqrt{2} \sqrt{\frac{\hat{c}_h^{1-\gamma} - \hat{c}^{1-\gamma}}{1-\gamma}}, \tag{25}$$

so that $\mu^{\hat{c}} = 0$ in (17). According to (25), if we want to convince the agent to consume more (i.e., choose higher \hat{c}), we need to expose him to less risk (i.e., lower σ^x) in order to weaken the precautionary motive. The cost of stationary contracts $\hat{v}_s(\hat{c})$ is shown as a dashed line in Figure 4, and satisfies the HJB equation

$$r\hat{v}_s(\hat{c}) = \hat{c} - \sigma_s^x(\hat{c})\hat{c}^\gamma \frac{\alpha}{\phi\sigma} + \hat{v}_s(\hat{c}) \left(\frac{\rho - \hat{c}^{1-\gamma}}{1-\gamma} + \frac{\gamma}{2} \sigma_s^x(\hat{c})^2 \right). \tag{26}$$

The optimal stationary contract is indicated by a dot at \hat{c}_r , with cost $\hat{v}_r = \hat{v}_s(\hat{c}_r)$.

It is interesting that the curves \hat{v}_m and \hat{v}_s meet at a point where the Euler equation holds and σ^x satisfies the risk-return trade-off condition (25), indicated by a dot in Figure 4. These are the two conditions for the classic consumption and portfolio choice problem (with Sharpe ratio $\alpha/(\sigma\phi)$). Here, we do not take into account the effect of risk exposure on the agent’s precautionary motive, and how a reduction in risk can raise \hat{c} and relax the agent’s incentive constraint. Hence, the minimal incentive-compatible share of risk the agent must retain is given by parameter ϕ , and we get the Sharpe ratio of $\alpha/(\sigma\phi)$. (See also Section 6, where we discuss in more detail the division of risk between the principal and agent through fund capital structure).

For the record, the solution of the classic consumption-portfolio choice problem, in which the agent retains a fraction ϕ of the risk, is given by the well-known formula¹⁴

$$\hat{c}_p = \left(\hat{c}_h^{1-\gamma} - \frac{1-\gamma}{2} \left(\frac{\alpha}{\gamma\phi\sigma} \right)^2 \right)^{\frac{1}{1-\gamma}}, \quad \sigma_p^x = \frac{\alpha}{\gamma\phi\sigma}, \quad \hat{v}_p = \hat{c}_p^\gamma. \tag{27}$$

The curves \hat{v}_m and \hat{v}_s are tangent at point \hat{c}_p because for any given \hat{c} the cost \hat{v}_m from the optimal unconstrained choice of σ^x must be weakly below the cost \hat{v}_s that is constrained by the Euler equation.

Relative to the classic portfolio solution, the optimal stationary contract has lower risk exposure, that is, $\sigma_r^x < \sigma_p^x$ and $\hat{c}_r > \hat{c}_p$. To see why, suppose we start at the classic portfolio solution and reduce σ^x in a uniform way, that is, within the family of stationary contracts that respect the Euler equation. This change raises \hat{c} , which creates a first-order benefit relaxing the IC constraint (18). The optimal stationary contract takes into account this effect of the risk exposure σ^x on the stationary level of the agent’s precautionary motive.

The optimal dynamic contract with hidden savings does even better, by taking into account how risk exposure σ_t^x at any time t affects the agent’s precautionary motive and incentives along the entire path leading to time t . At time 0, the contract features myopic risk-return optimization, as the optimal contract without hidden savings. Since we chose $\hat{c}_0 = \hat{c}_l$ optimally, we have $\hat{v}'(\hat{c}_l) = 0$, so the HJB equation (21) reduces to (24) and the

¹⁴Formally, the problem is

$$\max_{(c,k)} U(c) \quad \text{s.t.} \quad dw_t = (rw_t + \alpha k_t - c_t) dt + k_t(\phi\sigma) dZ_t, \quad w_t \geq 0$$

for a given wealth w_0 .

optimal contract therefore starts on the curve \hat{v}_m . In contrast the choice of σ_t^x at time $t > 0$ has to be made taking into account the precautionary motive and the tightening of the incentive constraints that this choice creates leading to time t . In the optimal contract, the benefits that a reduction in risk σ_t^x would have on past incentives must equal the cost of distortions going forward, captured by the derivative $\hat{v}'(\hat{c}_t)$. Distortions will be greater, that is, \hat{c}_t and $\hat{v}'(\hat{c}_t)$ must be higher, when the history prior to t is longer (larger t) and when fund size prior to time t had been higher. Hence, \hat{c}_t tends to rise in the optimal contract over time and after poor outcomes. As distortions build up, the contract looks more and more like a stationary contract with $\hat{c}_t = \hat{c}_h$ (see also Section 4, where the right boundary \hat{c}_h can be any other point on the stationary contract curve \hat{v}_s).

4. HIDDEN INVESTMENT

In this section, we consider the possibility that the agent can secretly save not only in the risk-free asset, but also in his private technology. This may not be a concern when the agent is already exposed to significant risk through the contract: the excess return the agent earns from his technology would not justify the incremental risk. However, hidden investment becomes a concern in future states, in which the principal would like to significantly reduce the agent’s risk exposure to control ex ante precautionary motive. In other words, the possibility of hidden investment may raise the agent’s incentives to save if in some future states he can benefit by investing these savings in his private technology. In this section, we analyze exactly how this plays out. To deal with hidden investment, the principal has to give the agent some minimal level of risk exposure. This constraint binds in some future states. Here, we focus on the main economic insights. The Online Appendix (Di Tella and Sannikov (2021)) has all the formal results and also extends the setting to incorporate aggregate risk.

If the agent can secretly invest in his private technology, his hidden savings follow the law of motion:

$$dh_t = (rh_t + z_t h_t \alpha + c_t - \tilde{c}_t + \phi k_t a_t) dt + z_t h_t \sigma dZ_t. \tag{28}$$

The new term $z_t \geq 0$ is the portfolio weight on his private technology in his hidden savings. If he invests in his private technology, he takes on risk to earn the excess return α . An interpretation for hidden investment is that the principal can give the agent an amount of capital k_t , but the agent can secretly invest more.

A contract $\mathcal{C} = (c, k)$ specifies the contractible payments c and capital k , contingent on returns R . After signing the contract, the agent can choose a strategy (\tilde{c}, a, z) to maximize his utility. The agent’s utility and the principal’s objective function are still given by (3) and (4). As in the baseline setting, it is without loss of generality to look for a contract where the agent does not steal, has no hidden savings, and no hidden investment. A contract is therefore incentive compatible if the agent’s optimal strategy is $(c, 0, 0)$.

The laws of motion of the state variables x_t and \hat{c}_t , (16) and (17), as well as the “skin in the game” constraint (18) remain unchanged. The no-savings constraint is that the agent’s discounted marginal utility $e^{-\rho t} c_t^{-\gamma}$ times the return of any trading strategy available to the agent must be a supermartingale. Otherwise, the agent can benefit by saving in some trading strategy to consume later. Condition $dL_t \geq 0$ in the constraint (17) ensures that the agent cannot benefit through risk-free savings, and we have a new IC constraint for the agent’s private technology,

$$\sigma_t^x + \sigma_t^{\hat{c}} \geq \frac{\alpha}{\gamma \sigma}. \tag{29}$$

The interpretation (29) is simple. The agent can obtain a premium $\frac{\alpha}{\sigma}$ for his idiosyncratic risk. If his exposure to risk were $\sigma_t^c < \frac{\alpha}{\gamma\sigma}$, he would benefit from deviating by secretly investing and taking on risk. The excess return α would more than compensate for the extra risk exposure at the margin.

Hidden investment restricts the principal's ability to provide incentives. In particular, the principal cannot promise to give the agent a perfectly deterministic consumption stream. If the principal tried to do this, the agent would just secretly take on risk by investing on his own. This is costly because the principal would like to promise future safety to relax the agent's precautionary saving motive. The extra constraint implies a lower upper bound \hat{c}_h . In the baseline where the agent cannot invest in his private technology, we have $\hat{c}_h = \left(\frac{\rho - r(1-\gamma)}{\gamma}\right)^{\frac{1}{1-\gamma}}$, corresponding to the \hat{c} in autarky where the agent cannot invest in capital. If the agent can invest in his private technology, the upper bound is

$$\hat{c}_h = \left(\frac{\rho - r(1-\gamma)}{\gamma} - \frac{1-\gamma}{2} \left(\frac{\alpha}{\sigma\gamma} \right)^2 \right)^{\frac{1}{1-\gamma}},$$

which is lower and corresponds to the \hat{c} in autarky where the agent invests in his private technology on his own (so he cannot get any risk sharing).¹⁵

The HJB equation is the same as in the baseline setting, but we have an extra constraint. In the region where $\hat{v}(\hat{c})$ is increasing, the HJB equation is

$$\begin{aligned} 0 = \min_{\sigma^x, \sigma^{\hat{c}}} & \hat{c} - r\hat{v} - \sigma^x \hat{c}^\gamma \frac{\alpha}{\phi\sigma} + \hat{v} \left(\frac{\rho - \hat{c}^{1-\gamma}}{1-\gamma} + \frac{\gamma}{2} (\sigma^x)^2 \right) \\ & + \hat{v}' \hat{c} \left(\frac{\hat{c}^{1-\gamma} - \hat{c}_h^{1-\gamma}}{1-\gamma} + \frac{(\sigma^x)^2}{2} + (1+\gamma)\sigma^x\sigma^{\hat{c}} + \frac{1+\gamma}{2} (\sigma^{\hat{c}})^2 \right) + \frac{\hat{v}''}{2} \hat{c}^2 (\sigma^{\hat{c}})^2 \end{aligned} \quad (30)$$

subject to $\sigma^x \geq 0$ and (29).

Figure 5 shows the optimal contract with hidden investment. As in the baseline setting, the contract starts at some \hat{c}_l and then moves immediately into the interior of the domain (\hat{c}_l, \hat{c}_h) . While the new IC constraint (29) is not binding, the FOC for σ^x and $\sigma^{\hat{c}}$ are the same as in the baseline setting. But with hidden investment the IC constraint (29) can be binding in some region of the state space. In Figure 6, the IC constraint (29) is binding near the upper bound \hat{c}_h . The property $\sigma_t^{\hat{c}} < 0$ holds for all $t > 0$ both when the hidden investment constraint (29) binds, and when it does not. When the constraint binds, $\sigma^x + \sigma^{\hat{c}} = \frac{\alpha}{\sigma\gamma}$ with $\sigma_t^x > \frac{\alpha}{\sigma\gamma}$. Contract dynamics are therefore qualitatively the same as in the baseline. The contract starts at \hat{c}_l with myopic optimization over σ^x and becomes less risky over time and after bad outcomes, never revisiting \hat{c}_l or reaching autarky \hat{c}_h .¹⁶

Just as in the baseline setting without hidden savings, the marginal cost of the agent's utility is lower than the inverse of the marginal utility of consumption, that is, $\hat{v}(\hat{c})x^\gamma < c^\gamma \Leftrightarrow \hat{v}(\hat{c}) < \hat{c}^\gamma$, and the optimal contract optimizes myopically over σ^x at $t = 0$. After that, the principal distorts the agent's access to capital to manipulate his precautionary motive. How can the principal do this, if the agent has access to hidden investment? The answer

¹⁵The parameter restrictions (7) and (8) ensure \hat{c}_h is positive in both cases.

¹⁶There is one subtle difference. The optimal contract without hidden investment does not have a stationary distribution, but instead limits in the long-run to autarky in the sense of Theorem 2. With hidden investment, there is a stationary distribution $\psi(\hat{c})$. See Theorem O.2 in the Online Appendix (Di Tella and Sannikov (2021)).

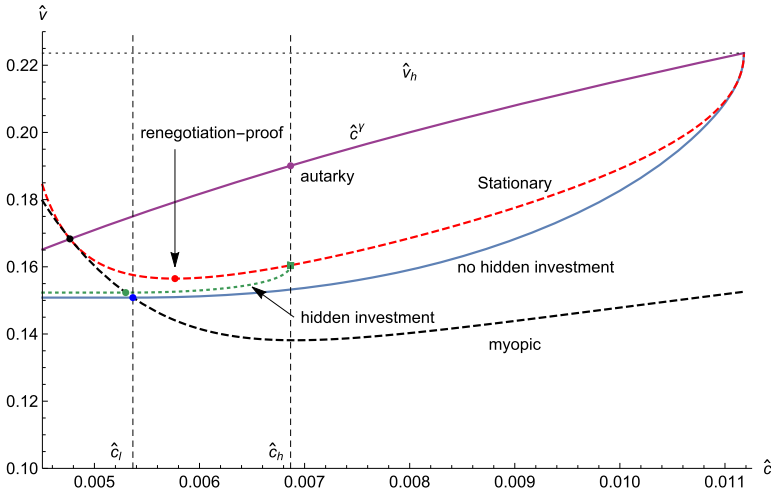


FIGURE 5.—The cost function $\hat{v}(\hat{c})$ without hidden investment (solid), and with hidden investment (dotted). For reference, dashed lines show the cost of stationary contracts and contracts with myopic optimization over σ^x . Parameters: $\rho = r = 5\%$, $\alpha = 1.7\%$, $\gamma = 1/3$, $\phi = 0.8$, $\sigma = 0.2/0.8$.

is that if the agent invests on his own he is the residual claimant and must bear the whole risk, while if he invests through the principal he gets to share some risk because $\phi < 1$. The principal is really restricting the amount of capital he is willing to share risk on. The agent would like to invest more if he could share the risk on the extra capital with the principal, but not if he must bear the whole risk on his own. For the same reason, although \hat{c}_h corresponds to the consumption profile under autarky, there are still gains from trade at that point. Since $\phi < 1$, the principal can give the agent the same consumption process that he would get in autarky, but more capital. As a result, the cost to the principal is




FIGURE 6.—The drift, $\mu^{\hat{c}}$ and μ^x , and volatility, $\sigma^{\hat{c}}$ and σ^x , of the state variables \hat{c} and x with hidden investment. The dotted part is where the hidden investment IC constraint is binding. The dashed lines indicate \hat{c}_l and \hat{c}_h with hidden investment. Parameters: $\rho = r = 5\%$, $\alpha = 1.7\%$, $\gamma = 1/3$, $\phi = 0.8$, $\sigma = 0.2/0.8$.

lower than what the agent could get in autarky, $\hat{v}(\hat{c}_h) = \hat{v}_s(\hat{c}_h) < \hat{c}_h^\gamma$. The cost of autarky is indicated by a dot in Figure 5.

It is useful to ask under what conditions the gains from trade are completely exhausted, and the optimal contract corresponds to autarky. Lemma O.9 in the Online Appendix (Di Tella and Sannikov (2021)) shows this is true in the special (but salient) case with hidden investment and $\phi = 1$. In this case, the optimal contract, the myopic contract, and autarky coincide. Without hidden investment there would still be gains from trade. While the agent can save on his own, the principal can still control his access to capital to provide incentives, and can therefore provide some risk sharing. However, with hidden investment, the optimal contract coincides with autarky. Intuitively, the agent can both save and invest on his own, so the principal cannot provide any risk sharing in an incentive compatible way. We can see this case as a limit in Figure 5. If we let $\phi \rightarrow 1$, while also adjusting σ so that $\phi\sigma = \varsigma$ is constant, all the curves corresponding to the case with no hidden investment remain unchanged. The optimal contract with hidden investment, however, becomes progressively worse (the dotted curve shifts up), because the agent finds investing on his own more attractive, and the upper bound \hat{c}_h shrinks. In the limit as $\phi \rightarrow 1$ while $\phi\sigma = \phi\sigma = \varsigma$ remains constant, the upper bound \hat{c}_h converges to $\hat{c}_p = \left(\frac{\rho-r(1-\gamma)}{\gamma} - \frac{1-\gamma}{2} \left(\frac{\alpha}{\gamma\varsigma}\right)^2\right)$ from above.

We close this section by considering the validity of the first-order approach with hidden investment. The agent's ability to invest his hidden savings makes the verification of global incentive compatibility potentially more difficult. The agent has a greater space of potential deviations with his savings. Fortunately, we can extend the results of Theorem 1 to deal with hidden investment. Theorem O.1 in the Online Appendix (Di Tella and Sannikov (2021)) shows that $\sigma^{\hat{c}} \leq 0$ is still sufficient to ensure global incentive compatibility even when the agent has hidden investment in both his private technology and aggregate risk. Since $\sigma^{\hat{c}} \leq 0$ holds in the optimal contract with hidden investment, the first-order approach remains valid. This is formalized in Theorem O.2 in the Online Appendix (Di Tella and Sannikov (2021)).

5. RENEGOTIATION

The optimal contract requires commitment. The principal can relax the agent's precautionary saving motive by promising an inefficiently low amount of capital and risk over time and after bad outcomes. This suggests that the agent and principal could be tempted to renegotiate the contract and “start over,” undoing the whole incentive scheme. Here, we formalize a notion of renegotiation, and characterize the optimal renegotiation-proof contract. As it turns out, the optimal renegotiation-proof contract is the best stationary contract, so even without commitment we are able to do better than under myopic optimization. This result is consistent with the presence of hidden investment and aggregate risk introduced in Section 4 and the Online Appendix (Di Tella and Sannikov (2021)).

After signing an incentive compatible contract $\mathcal{C} = (c, k)$, the principal can at any time offer a new continuation contract that leaves the agent at least as well off (the offer is “take it or leave it”). The question is, what kind of contracts can he offer—what is a valid “challenger” to the original contract? Here, we use a notion of internal consistency. If \mathcal{C} is renegotiation proof, then surely appropriately scaled parts of it should be a valid challenger. At any point in time, $\hat{v}_t = J_t/x_t$ is the continuation cost per unit of promised utility x_t . Define $\hat{v} = \inf \hat{v}(\omega, t)$, where the minimization is across both $\omega \in \Omega$ and $t \in \mathbb{R}_+$. At any stopping time τ , the principal can renegotiate and get a continuation cost

$x_\tau \times \hat{v}$ while delivering utility x_τ to the agent. With this in mind, we say that an incentive compatible contract is *renegotiation-proof (RP)* if

$$\infty \in \arg \min_{\tau} \mathbb{E}^Q \left[\int_0^\tau e^{-rt} (c_t - k_t \alpha) dt + e^{-r\tau} x_\tau \hat{v} \right].$$

The optimal contract with hidden savings is not renegotiation proof, because after any history at any $t > 0$, $\hat{v}_t > \hat{v}_l = \hat{v}$, so the principal is always tempted to “start over.”

It is easy to see that an incentive compatible contract is renegotiation-proof if and only if the continuation cost \hat{v}_t is constant. In Section 3.6, we define a class of stationary contracts in which \hat{c}_t is constant and c_t , k_t and x_t follow proportional geometric Brownian motions and $dL_t = 0$. These contracts have constant \hat{v}_t ; hence they are renegotiation proof. However, these are not the only renegotiation-proof contracts. There are contracts in which $dL_t > 0$, and c_t , k_t and x_t follow proportional geometric Brownian motions; they are also renegotiation proof. In addition, there could be nonstationary contracts with a constant cost \hat{v}_t .

However, Theorem O.4 in the Online Appendix (Di Tella and Sannikov (2021)) shows that the optimal renegotiation-proof contract is the optimal stationary contract, with cost \hat{v}_r , indicated in Figure 5 by a dot. This result shows that even without full commitment, the optimal contract with hidden savings distorts how much capital the agent receives. That is, the principal can do better than the myopic contract given by (27).

REMARK: It is possible that $\hat{c}_r = \hat{c}_h$ if the agent can invest his hidden savings and ϕ is close enough to 1. In the special case with hidden investment and $\phi = 1$, we have $\hat{c}_r = \hat{c}_p = \hat{c}_h$, as shown in Lemma O.9.

6. FUND CAPITAL STRUCTURE

In this section, we show how dynamic distortions become reflected in fund leverage, payout rate and the division of risk between the agent and principal (i.e., outside investors). This perspective is useful because these quantities are of applied interest. Also, they connect our results to the classic consumption-portfolio problem.

As a starting point, notice that without a long-term contract, the agent could simply shift fraction $1 - \phi$ of risk to the outside market and allocate his wealth between the risk-free asset and capital, as in the standard optimal consumption and portfolio choice problem. This risk sharing is incentive compatible because insurance does not cover more than the cost of cash diversion, hence the agent does not want to steal. The resulting contract is the myopic contract in (27), which satisfies the Euler equation and myopic optimization over risky capital (taking into account that only a fraction ϕ of the risk is retained by the agent).

Relative to this simple benchmark, distortions in the optimal contract take the form of dynamic deviations from myopic optimization for capital: the contract restricts the funds invested in capital (i.e., the ratio k_t/x_t), particularly after bad outcomes. That is, if the agent could increase the investment in capital, while passing on some risk to the principal in an incentive compatible way, he would. It is important that the optimal contract ties the agent’s hands dynamically with respect to capital investment k_t relative to x_t .

These distortions allow the principal to improve risk sharing. That is, the agent can shift a greater fraction of risk than $1 - \phi$ to the principal and still respect incentive compatibility. The agent would like to invest more but the contract limits the investment, so an

extra dollar of returns relaxes the limit. Thus, the restricted access to capital makes the marginal value of funds invested in capital greater than the marginal utility of consumption, which relaxes the incentive constraint. As we move away from the simple optimal portfolio benchmark, distorting the allocation to capital has a second-order negative effect on welfare, but the incentive and improved risk sharing effect of this distortion has a first-order positive effect.

To see this in more detail, let us be more explicit about how the optimal contract deviates from the myopic consumption-portfolio solution. At any time, there is capital k_t in the fund. Since the cost of compensating the agent is $J_t = \mathbb{E}_t[\int_t^\infty e^{-r(s-t)}(c_s - \alpha k_s) ds]$, the value of the principal’s stake in the fund is $k_t - J_t$. This gives us a specific division of the value of capital k_t : the principal’s stake is $k_t - J_t$ and the agent’s stake is J_t . Recall that the agent needs to contribute J_0 at time 0 in order for the principal to break even.

The principal has to earn the required return of r on his stake in the fund. This means the agent’s stake J_t earns the α and some portion of risk, that is, J_t follows the law of motion:

$$dJ_t = (rJ_t - c_t + \alpha k_t) dt + \tilde{\phi}_t \sigma k_t dZ_t \tag{31}$$

for some stochastic process $\tilde{\phi}_t$. This equation is a classic dynamic budget equation, with the allocation of funds to the risk-free asset, capital (of which risk $1 - \tilde{\phi}_t$ is insured) and consumption c_t . We call $\tilde{\phi}_t$ the agent’s *retained equity stake*. Ideally, we would set $\tilde{\phi}_t = 0$ to obtain perfect risk-sharing, but the retained equity stake is important to give the agent some “skin in the game” and ensure incentive compatibility.

For any contract that is recursive in \hat{c}_t and scale invariant in x_t , with a cost function $J_t = \hat{v}(\hat{c}_t)x_t$, we can match terms to obtain

$$k_t/J_t = \frac{\sigma_t^x}{\hat{v}_t \hat{c}_t^{-\gamma} \phi \sigma}, \quad \tilde{\phi}_t = (1 + \sigma_t^{\hat{v}}/\sigma_t^x)(\hat{v}_t \hat{c}_t^{-\gamma}) \phi, \quad c_t/J_t = \frac{\hat{c}_t}{\hat{v}_t}, \tag{32}$$

where $\sigma^{\hat{v}} = \hat{v}'(\hat{c}_t)/\hat{v}(\hat{c}_t)\sigma_t^{\hat{c}}\hat{c}_t$ is the volatility of \hat{v}_t .

The optimal contract distorts the agent’s portfolio-weight on capital, or *leverage*, k/J , to weaken the agent’s incentives to steal. These distortions imply a lower incentive compatible retained equity stake

$$\tilde{\phi}_t = \underbrace{(1 + \sigma_t^{\hat{v}}/\sigma_t^x)}_{<1} \underbrace{(\hat{v}_t \hat{c}_t^{-\gamma})}_{<1} \phi < \phi.$$

The retained equity stake $\tilde{\phi}_t$ is below ϕ for two reasons. First, the marginal value of consumption is below the marginal value of funds that can be invested in capital, that is, $\hat{v}_t \hat{c}_t^{-\gamma} < 1$. Second, after bad outcomes the agent is punished not only with less wealth, but also with tighter distortions, $\sigma_t^{\hat{v}}/\sigma_t^x < 0$. Facing less risk, the agent’s consumption rate out of wealth, c_t/J_t , is above the level corresponding to the simple benchmark consumption-portfolio problem (the myopic contract).

Figure 7 shows these quantities in our optimal contract. Fund leverage k_t/J_t is decreasing with \hat{c}_t , so the agent’s margin tightens after bad outcomes. The agent’s share of risk $\tilde{\phi}_t$ is below ϕ , and the consumption rate c_t/J_t is above the level of the optimal portfolio benchmark. Notice that the optimal contract uses future distortions to relax the retained equity stake, $\tilde{\phi}_0 < \phi$, but optimizes myopically over capital at $t = 0$, so fund leverage is higher at $t = 0$ than what it would have been under the simple portfolio benchmark.




FIGURE 7.—Fund leverage k/J , consumption rate c/J , and retained equity $\tilde{\phi}$ for the optimal contract. Dashed line indicates the simple benchmark portfolio (myopic contract). Parameters: $\rho = r = 5\%$, $\alpha = 1.7\%$, $\gamma = 1/3$, $\phi = 0.8$, $\sigma = 0.2/0.8$.

7. NUMERICAL APPLICATION

We computed a numerical example to illustrate our results and show how our theorems are applied. The interest rate and impatience rate of the agent are $r = \rho = 5\%$, and risk aversion $\gamma = 1/3$. The excess return the agent obtains from investing in capital is $\alpha = 1.7\%$. The efficiency of stealing ϕ and the idiosyncratic risk σ enter together, and we set $\phi\sigma = 20\%$. These parameter values seem reasonable and satisfy the constraints in Section 2, but there is not a specific calibration target.

For the baseline case without hidden investment or renegotiation, we solved the HJB numerically as an ODE, plugging in the FOCs. The resulting μ^x , σ^x , $\mu^{\hat{c}}$, and $\sigma^{\hat{c}}$ are bounded and $\mu^x(\hat{c})$ is bounded below away from r , so Lemma 4 ensures that the resulting contract is admissible and delivers utility u_0 to the agent. The verification Theorem 3 shows we have an optimal contract. Figures 1 and 2 show the cost function and the laws of motion of the states, and Figure 3 shows a simplified simulated path for the main variables of interest.

Lemma 4 and Theorem 1 can also be used to verify local and global incentive compatibility for the suboptimal contracts introduced in Section 3.6. We can get a quantitative sense of the importance of the incentive mechanism in the optimal contract by comparing its cost with the cost of the myopic stationary contract, which corresponds to a simple portfolio problem where the agent invests and consumes on his own, subject only to retaining fraction ϕ of the equity. The myopic contract has a cost of $\hat{v}_p = 0.168$, compared to $\hat{v}(\hat{c}_l) = 0.151$ for the optimal contract. This means that switching from the myopic stationary contract to the optimal contract reduces the cost of delivering utility x_0 by roughly 10%. Equivalently, switching to the optimal contract allows us to leave the principal indifferent and improve the agent’s utility by an amount equivalent to a proportional 10% increase in consumption. To put this in context, the optimal contract without hidden savings has a cost $\hat{v}_n = 0.138$. Eliminating the agent’s access to hidden savings would therefore reduce costs by a further 8.4%. Switching from the myopic stationary contract to the optimal contract with hidden savings closes roughly half the gap in cost between the myopic stationary contract and the optimal contract without hidden savings.

We can also quantitatively explore the cost of introducing renegotiation and hidden investment into the environment. The optimal contract with renegotiation corresponds to the best stationary contract, with cost $\hat{v}_r = 0.156$. This is roughly 3.75% more costly than the optimal contract without renegotiation. Renegotiation eliminates roughly a third of the gain from going from the myopic stationary contract (the simple portfolio problem) to the optimal contract with full commitment.

We solve for the optimal contract with hidden investment using $\phi = 0.8$ and keeping $\phi\sigma = 20\%$ unchanged. We solve the appropriate HJB equation (O.20) in the Online Appendix (Di Tella and Sannikov (2021)) numerically, and we use the more general

Lemma O.4 and Theorem O.3 to verify we have constructed an optimal contract. Figure 5 shows the optimal contract with hidden investment. The cost is $\hat{v}(\hat{c}_t) = 0.152$, or roughly 1% more costly than the optimal contract without hidden investment. Introducing hidden investment eliminates 1/10 of the gain from going from the myopic stationary contract to the optimal contract without hidden investment. We know, however, that as $\phi \rightarrow 1$, while $\phi\sigma$ remains fixed, the cost of the optimal contract with hidden investment converges to the cost of the myopic stationary contract, while the cost of the optimal contract without hidden investment is unchanged. In that limit introducing hidden investment eliminates all of the gains of the optimal contract over the myopic stationary contract.

8. DISCUSSION: DYNAMIC PRIVATE INFORMATION

Our model is related to a whole class of problems with the broad title “dynamic adverse selection.” These are dynamic environments where the agent has private information. In our case, information about savings. In these environments the agent’s deviation payoff plays a central role, also called his information rent. This term originates from the static auction environment of Myerson (1981), where agents whose type (valuation) is higher can mimic lower types, and their higher valuation earns them higher utility or rents. Although the two problems may appear unrelated, at the core of both lie the distortions that the principal uses to control the payoff of a deviating type.

In the static Myerson (1981) setting, the distortion can take the form of a reserve price, which reduces the rents of high valuation bidders and helps the principal screen. In dynamic models of adverse selection distortions are more complex.¹⁷ In our setting, the principal distorts the agent’s access to capital ex post to reduce his precautionary motive. This lowers the value of hidden savings and improves ex ante incentives against stealing.

Generally, the agent’s incentives depend on his entire off-path value function. This is the approach in Fernandes and Phelan (2000), which is challenging due to an infinite-dimensional state space. Instead, we use the first-order approach, which focuses on local incentives, and verify global incentive compatibility using an upper bound of the agent’s off-path value function. In terms of transformed utility, this bound takes the form $x_t + \hat{c}_t^{-\gamma} h_t$. The agent’s deviation payoff would take this exact form if we had $\phi = 1$ and the agent can choose any real a_t (so the agent can move funds between returns and hidden savings without friction), and if we took a contract with $\sigma_t^{\hat{c}} = 0$. The stationary contracts of Section 3.6 have this property, as well as contracts in which the path for \hat{c}_t is deterministic. Indeed, the proof of Theorem 1, which shows that $x_t + \hat{c}_t^{-\gamma} h_t$ is an upper bound on the agent’s off-path value function in our setting, extends to the possibility that $a_t \leq 0$ if also $\sigma_t^{\hat{c}} = 0$. To see why in this case the bound is attainable, consider the utility he would get if he immediately put all his hidden savings back into returns. In general, this utility is

$$x_t + \int_0^h \hat{c}(h')^{-\gamma} dh', \quad (33)$$

where $\hat{c}(h')$ is the level of \hat{c} that results after the agent uses h' of savings to boost returns, because $\hat{c}^{-\gamma}$ is the marginal benefit of a dollar of returns from incentive compatibility. This expression resembles the bound used in DeMarzo and Sannikov (2016) and Pavan, Segal, and Toikka (2014). With deterministic contracts, $\hat{c}(h') = \hat{c}_t$ for all h' because \hat{c}_t does not depend on the history of observed returns, so the agent obtains utility $x_t + \hat{c}_t^{-\gamma} h_t$, achieving

¹⁷Pavan, Segal, and Toikka (2014), Battaglini (2005), DeMarzo and Sannikov (2016), He et al. (2017).

the upper bound. So in this special case we actually know exactly what the agent’s off-path utility is, and we can therefore show that contracts with $\sigma_t^{\hat{c}} = 0$ are globally incentive compatible even if we allow $a_t \leq 0$.

In our model, $x_t + \hat{c}_t^{-\gamma} h_t$ is only an upper bound. The assumption that $a_t \geq 0$ is necessary for this result, since otherwise the agent can obtain a higher utility of (33) by immediately using all his hidden savings to boost returns. Indeed, when \hat{c}_t declines with reported returns (i.e., when $\sigma_t^{\hat{c}} < 0$), this utility is greater than the bound, $x_t + \int_0^h \hat{c}(h')^{-\gamma} dh' > x_t + \hat{c}_t^{-\gamma} h_t$. The restriction $a_t \geq 0$ implies that $x_t + \hat{c}_t^{-\gamma} h_t$ is an upper bound on the agent’s deviation payoff, which allows us to prove global incentive compatibility for all contracts with $\sigma_t^{\hat{c}} \leq 0$.

It is interesting what would happen if we allowed $a_t \leq 0$. We do not know in this case if the solution to the relaxed problem remains incentive compatible, that is, whether the first-order approach is valid. We know for sure that in that contract the agent’s deviation payoff is above $x_t + \hat{c}_t^{-\gamma} h_t$: it is at least (33), and if it is higher than that, the first-order approach fails. If this is so, the principal can always reduce the agent’s off-equilibrium utility to our bound, $x_t + \hat{c}_t^{-\gamma} h_t$, by using deterministic contracts. This means that solving the relaxed problem with the extra constraint $\sigma_t^{\hat{c}} = 0$ provides an upper bound on the cost of the optimal contract, while the solution to the relaxed problem with free $\sigma_t^{\hat{c}}$ provides a lower bound. The optimal contract may be in between, and may use dynamic distortions to bound the agent’s deviation payoff at savings levels $h_t > 0$.

9. CONCLUSIONS

We study the role of hidden savings in a classic portfolio-investment problem with fund diversion. The agent’s precautionary saving motive plays a central role in his incentives to divert funds. If the agent expects a large exposure to risk in the future, he places a large value on hidden savings that he can use to self-insure. As a result, the principal must manipulate the agent’s precautionary saving motive by committing to limit his exposure to risk in the future, especially after bad outcomes. Since giving capital to the agent requires exposing him to risk to align incentives, this leads to dynamic distortions in the agent’s access to capital and a skewed compensation scheme. After good outcomes, the agent’s access to capital improves, allowing his fund to keep growing rapidly. After bad outcomes, his access to capital is restricted and he stagnates. In exchange, his consumption is somewhat insured on the downside, and he is punished instead with lower growth. We also extend our environment to incorporate hidden investment and renegotiation, and show that the optimal contract can be mapped into a consumption-portfolio allocation with a retained equity constraint and a leverage constraint.

An important methodological contribution is to provide a sufficient analytical condition for the validity of the first-order approach. If the agent’s precautionary saving motive is weaker after bad outcomes, the contract is globally incentive compatible. This condition holds in the optimal contract and in a wider class of contracts beyond the optimal one. In fact, the sufficient condition does not even require a recursive structure, and it is valid even with aggregate risk and hidden investment.

APPENDIX: OMITTED PROOFS

PROOF OF LEMMA 1: Consider the strategy $(c, 0)$. Define

$$Y_t = \mathbb{E}_t \left[\int_0^\infty e^{-\rho s} \frac{c_s^{1-\gamma}}{1-\gamma} ds \right] = \int_0^t e^{-\rho s} \frac{c_s^{1-\gamma}}{1-\gamma} ds + e^{-\rho t} U_t^{c,0}.$$

Since Y is a P -martingale adapted to the filtration generated by Brownian motion Z , we can write

$$dY_t = e^{-\rho t} \frac{c_t^{1-\gamma}}{1-\gamma} dt + e^{-\rho t} dU_t^{c,0} - \rho e^{-\rho t} U_t^{c,0} dt = e^{-\rho t} \Delta_t \sigma dZ_t$$

for some stochastic process Δ , also adapted to the same filtration. Dividing by $e^{-\rho t}$ and rearranging we get (10). Q.E.D.

PROOF OF LEMMA 2: We will show that if (11) fails then the agent has a profitable deviation with stealing. Equation (12) is simply the first-order condition necessary for (11). After that, we show that (13) follows from the standard Euler equation, which must necessarily hold to rule out deviations with savings (without stealing).

First, suppose that (11) fails on a set of positive measure. Consider the stealing process a where a_t is the action that achieves the maximum of (11) on $[0, \bar{a}]$, and \bar{a} is an arbitrary bound. Since the objective in (11) is concave, we have that $a_t > 0$ whenever (11) fails. Consider the resulting strategy $(c + \phi k a, a)$, which implies zero savings, $h_t = 0$.

Define the process

$$V_t \equiv \int_0^t e^{-\rho s} \frac{(c_s + \phi k_s a_s)^{1-\gamma}}{1-\gamma} ds + e^{-\rho t} U_t^{c,0},$$

which corresponds to the utility of following this strategy until time t and then reverting to $(c, 0)$. From the representation (10), under the strategy $(c + \phi k a, a)$ the process V_t has drift

$$e^{-\rho t} \left(\frac{(c_t + \phi k_t a_t)^{1-\gamma}}{1-\gamma} - \rho U_t^{c,0} + \rho U_t^{c,0} - \frac{c_t^{1-\gamma}}{1-\gamma} - \Delta_t a_t \right),$$

since $dR_t - (\alpha + r) dt$ has drift $-a_t$. By our construction of a , the drift is nonnegative and positive on a set of positive measure, so V_t is a local submartingale, and we can pick a large enough stopping time τ such that

$$U_0^{c,0} = V_0 < \mathbb{E}_0^a[V_\tau] = \mathbb{E}_0^a \left[\int_0^\tau e^{-\rho s} \frac{(c_s + \phi k_s a_s)^{1-\gamma}}{1-\gamma} ds + e^{-\rho \tau} U_\tau^{c,0} \right],$$

The last expectation is the agent's payoff from following $(c + \phi k a, a)$ until time τ and then reverting to $(c, 0)$, and this strategy is a profitable deviation over following $(c, 0)$ throughout.

For (13), notice that $e^{-(\rho-r)t} c_t^{-\gamma}$ must be a supermartingale to ensure that the agent does not want to save. Then we can use the Doob–Meyer decomposition to write $e^{-(\rho-r)t} c_t^{-\gamma} = M_t - A_t$, where M_t is a local martingale and A a weakly increasing process. Since M is adapted to the filtration generated by Z , we can write $M_t = \int_0^t \sigma_t^M dZ_t$ for some process σ_t^M . Define σ_t^c by $\sigma_t^M = -\gamma \sigma_t^c e^{-(\rho-r)t} c_t^{-\gamma}$. Then using Ito's lemma we obtain (13) with a weakly increasing process L . Q.E.D.

PROOF OF LEMMA 3: To establish the bound $\hat{c}_t \leq \hat{c}_h$, use the fact that in an incentive compatible contract $m_t = e^{(r-\rho)t} c_t^{-\gamma}$ is a supermartingale. Let $y_t = \frac{c_t^{1-\gamma}}{1-\gamma} = e^{(\frac{\gamma-1}{\gamma}(\rho-r)t} \frac{m_t^{\frac{\gamma-1}{\gamma}}}{1-\gamma}$. We can write, for $s > t$,

$$\mathbb{E}_t[y_s] \geq e^{(\frac{\gamma-1}{\gamma}(\rho-r)(s-t)} \frac{\mathbb{E}_t[m_s]^{\frac{\gamma-1}{\gamma}}}{1-\gamma} \geq e^{(\frac{\gamma-1}{\gamma}(\rho-r)(s-t)} e^{(r-\rho)(\frac{\gamma-1}{\gamma}t)} c_t^{1-\gamma},$$

where the first inequality follows from Jensen’s inequality (with equality only if c_t is deterministic), and the second inequality from m_t being a supermartingale (with equality if it is a martingale). Now we can write

$$\begin{aligned}
 U_t &= \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} y_s ds \right] \geq \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} e^{(\frac{\gamma-1}{\gamma}(\rho-r))s} \frac{e^{(r-\rho)(\frac{\gamma-1}{\gamma})t} c_t^{1-\gamma}}{1-\gamma} ds \right] \\
 &= \frac{c_t^{1-\gamma}}{1-\gamma} \mathbb{E}_t \left[\int_t^\infty e^{-\frac{\rho-r(1-\gamma)}{\gamma}(s-t)} ds \right] = \frac{c_t^{1-\gamma}}{1-\gamma} \left(\frac{\rho-r(1-\gamma)}{\gamma} \right)^{-1}.
 \end{aligned}$$

It follows that

$$\begin{aligned}
 x_t &= ((1-\gamma)U_t)^{\frac{1}{1-\gamma}} \geq c_t \left(\frac{\rho-r(1-\gamma)}{\gamma} \right)^{-\frac{1}{1-\gamma}} \\
 \implies \hat{c}_t &= \frac{c_t}{x_t} \leq \left(\frac{\rho-r(1-\gamma)}{\gamma} \right)^{\frac{1}{1-\gamma}} = \hat{c}_h.
 \end{aligned}$$

In addition, this upper bound can only be achieved with deterministic consumption, so $U_t^{c,0}$ must be deterministic too in that case. This implies that both c_t and x_t grow at rate $\frac{r-\rho}{\gamma}$, so \hat{c}_h is an absorbing state. In light of (12), we must have $k_{t+u} = 0$ in the continuation contract, so we have the autarky contract with cost $\hat{v}_h x_t$.

For expression (15), use the law of motion of c_t in equation (13) to compute

$$\begin{aligned}
 dy_t &= (1-\gamma)y_t \left(\frac{r-\rho}{\gamma} + \frac{1+\gamma}{2}(\sigma_t^c)^2 - \frac{\gamma}{2}(\sigma_t^c)^2 \right) dt + (1-\gamma)y_t \sigma_t^c dZ_t + dA_t, \\
 dy_t &= y_t(1-\gamma) \left(\frac{r-\rho}{\gamma} + \frac{1}{2}(\sigma_t^c)^2 \right) dt + (1-\gamma)y_t \sigma_t^c dZ_t + dA_t,
 \end{aligned}$$

where A is a weakly increasing process coming from the term dL_t in (13). This is a linear SDE with solution

$$y_s \geq y_t e^{\int_t^s (1-\gamma)(\frac{r-\rho}{\gamma} + \frac{1}{2}(\sigma_u^c)^2) du} e^{\int_t^s ((1-\gamma)\sigma_u^c dZ_u - \frac{1}{2}(1-\gamma)^2(\sigma_u^c)^2 du)} \tag{34}$$

for $s > t$, with equality if $L_t = 0$.

Suppose that $\exp(\int_0^t (1-\gamma)\sigma_u^c dZ_u - \int_0^t \frac{1}{2}(1-\gamma)^2(\sigma_u^c)^2 du)$ is a proper martingale. A sufficient condition for this is that σ^c be bounded. It then defines a probability measure $\tilde{P}(\sigma^c)$ such that $\tilde{Z}_t = Z_t - \int_0^t (1-\gamma)\sigma_s^c ds$ is a \tilde{P} -martingale. Now write

$$\begin{aligned}
 U_t &= \mathbb{E}_t \left[\int_t^\infty e^{-\rho s} y_s ds \right] \\
 &\geq y_t \mathbb{E}_t \left[\int_t^\infty e^{-\rho s} e^{\int_t^s (1-\gamma)(\frac{r-\rho}{\gamma} + \frac{1}{2}(\sigma_u^c)^2) du} e^{\int_t^s (1-\gamma)\sigma_u^c dZ_u - \frac{1}{2}(1-\gamma)^2(\sigma_u^c)^2 du} ds \right] \\
 &= y_t \mathbb{E}_t \left[\int_t^\infty e^{-\int_t^s (\frac{\rho-r(1-\gamma)}{\gamma} - \frac{1-\gamma}{2}(\sigma_u^c)^2) du} e^{\int_t^s (1-\gamma)\sigma_u^c dZ_u - \frac{1}{2}(1-\gamma)^2(\sigma_u^c)^2 du} ds \right],
 \end{aligned}$$

and, therefore,

$$U_t \geq y_t \mathbb{E}_t^{\tilde{P}} \left[\int_t^\infty e^{-\int_t^s (\frac{\rho-r(1-\gamma)}{\gamma} - \frac{1-\gamma}{2} (\sigma_s^c)^2) du} ds \right].$$

Now compute as before

$$\begin{aligned} x_t &= ((1-\gamma)U_t)^{\frac{1}{1-\gamma}} \geq c_t \mathbb{E}_t^{\tilde{P}} \left[\int_t^\infty e^{-\int_t^s (\frac{\rho-r(1-\gamma)}{\gamma} - \frac{1-\gamma}{2} (\sigma_s^c)^2) du} ds \right]^{\frac{1}{1-\gamma}} \\ \implies \frac{c_t}{x_t} &\leq \mathbb{E}_t^{\tilde{P}} \left[\int_t^\infty e^{-\int_t^s (\frac{\rho-r(1-\gamma)}{\gamma} - \frac{1-\gamma}{2} (\sigma_s^c)^2) du} ds \right]^{-\frac{1}{1-\gamma}}. \end{aligned} \tag{35}$$

With $L_t = 0$, we have equality in (34), so we have equality in (35). Q.E.D.

PROOF OF LEMMA 4: First, we show the contract delivers utility $\frac{x_t^{1-\gamma}}{1-\gamma} < \infty$ if the agent follows strategy $(c, 0)$. Let $Y_t = \frac{x_t^{1-\gamma}}{1-\gamma}$, and using the law of motion of x , (16), we get

$$dY_t = \underbrace{Y_t(1-\gamma) \left(\mu_t^x - \frac{\gamma}{2} (\sigma_t^x)^2 \right)}_{\rho Y_t - \frac{c_t^{1-\gamma}}{1-\gamma}} dt + Y_t(1-\gamma) \sigma_t^x dZ_t. \tag{36}$$

Integrating we obtain

$$Y_0 = \mathbb{E} \left[\int_0^{\tau^n} e^{-\rho s} \frac{c_s^{1-\gamma}}{1-\gamma} ds + e^{-\rho \tau^n} Y_{\tau^n} \right]$$

for an increasing sequence of bounded stopping times with $\tau^n \rightarrow \infty$ a.s. Take the limit $n \rightarrow \infty$, using the monotone convergence theorem on the first term to get

$$Y_0 = \underbrace{\mathbb{E} \left[\int_0^\infty e^{-\rho s} \frac{c_s^{1-\gamma}}{1-\gamma} ds \right]}_{U_0^{c,0}} + \lim_{n \rightarrow \infty} \mathbb{E} [e^{-\rho \tau^n} Y_{\tau^n}].$$

We will now show that the last term is zero,

$$\lim_{n \rightarrow \infty} \mathbb{E} [e^{-\rho \tau^n} Y_{\tau^n}] = 0.$$

Since σ^x is bounded and \hat{c} bounded away from zero and above by \hat{c}_h , μ^x is bounded, too, and so is therefore the growth rate, $(1-\gamma)(\mu_t^x - \frac{\gamma}{2}(\sigma_t^x)^2)$, and volatility, $(1-\gamma)\sigma_t^x$, of Y_t in (36). Furthermore, the growth rate of Y_t in (36), is bounded away below ρ ,

$$(1-\gamma) \left(\mu_t^x - \frac{\gamma}{2} (\sigma_t^x)^2 \right) - \rho = -\hat{c}_t^{1-\gamma} \leq \max \{ -\hat{c}^{1-\gamma}, -\hat{c}_h^{1-\gamma} \} < 0,$$

where \hat{c} is a lower bound on \hat{c}_t . We then get that $\lim_{n \rightarrow \infty} \mathbb{E} [e^{-\rho \tau^n} Y_{\tau^n}] = 0$ and, therefore, $U_0^{c,0} = Y_0 = \frac{x_0^{1-\gamma}}{1-\gamma} < \infty$. The same reasoning yields $U_t^{c,0} = \frac{x_t^{1-\gamma}}{1-\gamma} < \infty$ for all t .

Now we show that the resulting contract \mathcal{C} is admissible if and only if $\mathbb{E}^Q[\int_0^\infty e^{-rt} x_t dt] < \infty$. To show sufficiency, notice that since \hat{c} is bounded above by \hat{c}_h and σ^x bounded, we can write

$$\mathbb{E}^Q \left[\int_0^\infty e^{-rt} (c_t + k_t) dt \right] \leq 2 \max \left\{ \hat{c}_h, \frac{\bar{\sigma}^x \hat{c}_h^\gamma}{\phi \sigma} \right\} \mathbb{E}^Q \left[\int_0^\infty e^{-rt} x_t dt \right] < \infty,$$

where $\bar{\sigma}^x$ is an upper bound on σ_t^x . For necessity, since \hat{c}_t is bounded away from zero, $\hat{c}_t \geq \underline{\hat{c}} > 0$,

$$\infty > \mathbb{E}^Q \left[\int_0^\infty e^{-rt} c_t dt \right] \geq \underline{\hat{c}} \times \mathbb{E}^Q \left[\int_0^\infty e^{-rt} x_t dt \right].$$

Since the contract is admissible and satisfies (17) and (18), and $\hat{c}_t \leq \hat{c}_h$, it is locally incentive compatible. Q.E.D.

PROOF OF THEOREM 1: We will do the proof for $dL_t = 0$ to keep it simple; it can be easily generalized for $dL_t \neq 0$. First, write the bound $(1 + \hat{h}_t \hat{c}_t^{-\gamma})^{1-\gamma} U_t^{c,0} = \frac{\bar{x}_t^{1-\gamma}}{1-\gamma}$, where $\bar{x}_t = x_t + h_t \hat{c}_t^{-\gamma}$, and define $\check{c}_t = \tilde{c}_t / \bar{x}_t$.

For any feasible strategy (\tilde{c}, a) , we can write the difference between the utility of that strategy, $U_t^{\tilde{c},a}$, and the bound $\frac{\bar{x}_t^{1-\gamma}}{1-\gamma}$ as

$$\begin{aligned} & e^{-\rho t} \left(U_t^{\tilde{c},a} - \frac{\bar{x}_t^{1-\gamma}}{1-\gamma} \right) \\ &= \mathbb{E}_t^a \left[\int_t^{\tau^n} e^{-\rho u} \frac{\tilde{c}_u^{1-\gamma}}{1-\gamma} du + \int_t^{\tau^n} d \left(e^{-\rho u} \frac{\bar{x}_u^{1-\gamma}}{1-\gamma} \right) + e^{-\rho \tau^n} \left(U_{\tau^n}^{\tilde{c},a} - \frac{\bar{x}_{\tau^n}^{1-\gamma}}{1-\gamma} \right) \right], \end{aligned}$$

for a localizing sequence $\{\tau^n\}$ with $\tau^n \rightarrow \infty$ a.s. We would like to show that this is always nonpositive.

First, take the first two terms on the rhs and write them:¹⁸

$$\begin{aligned} & \mathbb{E}_t^a \left[\int_t^{\tau^n} e^{-\rho u} \frac{\tilde{c}_u^{1-\gamma}}{1-\gamma} du + \int_t^{\tau^n} d \left(e^{-\rho u} \frac{\bar{x}_u^{1-\gamma}}{1-\gamma} \right) \right] \\ &= \mathbb{E}_t^a \left[\int_t^{\tau^n} e^{-\rho u} \bar{x}_u^{-\gamma} \left(\frac{\check{c}_u^{1-\gamma} - \rho}{1-\gamma} \bar{x}_u + \bar{x}_u \mu_u^{\bar{x}} - \bar{x}_u \frac{\gamma}{2} (\sigma_u^{\bar{x}})^2 \right) du \right], \end{aligned} \tag{37}$$

where $\mu^{\bar{x}}$ and $\sigma^{\bar{x}}$ are the geometric drift and volatility of \bar{x} , and satisfy

$$\begin{aligned} \bar{x}_t \mu_t^{\bar{x}} &= x_t \left(\frac{\rho - \hat{c}_t^{1-\gamma}}{1-\gamma} + \frac{\gamma}{2} (\sigma_t^x)^2 - \frac{\sigma_t^x}{\sigma} a_t \right) + \hat{c}_t^{-\gamma} (r h_t + c_t - \tilde{c}_t + \phi k_t a_t) \\ &\quad + h_t \hat{c}_t^{-\gamma} \left(\rho - r - \gamma \frac{\hat{c}_t^{1-\gamma} - \rho}{1-\gamma} - \frac{\gamma}{2} (\sigma_t^x)^2 - \gamma^2 \sigma_t^x \sigma_t^{\hat{c}} + \gamma \sigma_t^{\hat{c}} a_t \right), \\ \bar{x}_t \sigma_t^{\bar{x}} &= \sigma_t^x x_t - h_t \hat{c}_t^{-\gamma} \gamma \sigma_t^{\hat{c}}. \end{aligned}$$

¹⁸If $dL_t \neq 0$, we would get an extra negative term on the rhs; $dL_t > 0$ pushes \bar{x}_t down because $h_t \geq 0$. Since we want to prove that the rhs is nonpositive, we can take $dL_t = 0$ wlog.

We will show that the rhs of (37) is nonpositive. To do this, we will split the integrand into three parts,

$$\frac{\check{c}_u^{1-\gamma} - \rho}{1-\gamma} \bar{x}_u + \bar{x}_u \mu_u^{\bar{x}} - \bar{x}_u \frac{\gamma}{2} (\sigma_u^{\bar{x}})^2 = A_u a_u + B_u + C_u,$$

and we will show that A_t , B_t , and C_t are always nonpositive.

The term A_t collects the terms that multiply a_t :

$$A_t = \underbrace{\hat{c}_t^{-\gamma} \phi k_t - x_t \frac{\sigma_t^x}{\sigma}}_{\leq 0 \text{ by IC}} + \underbrace{h_t \hat{c}_t^{-\gamma} \gamma \sigma_t^{\hat{c}}}_{\leq 0} \leq 0,$$

where the second term is nonpositive because $h_t \geq 0$ and $\sigma_t^{\hat{c}} \leq 0$. Since $a_t \geq 0$, we can wlog take $a_t = 0$.¹⁹

The term B_t collects the remaining terms that do not have volatilities

$$B_t = x_t \frac{\rho - \hat{c}_t^{1-\gamma}}{1-\gamma} + \hat{c}_t^{-\gamma} (r h_t + \hat{c}_t x_t - \check{c}_t \bar{x}_t) + h_t \hat{c}_t^{-\gamma} \left(\rho - r - \gamma \frac{\hat{c}_t^{1-\gamma} - \rho}{1-\gamma} \right) + \bar{x}_t \frac{\check{c}_u^{1-\gamma} - \rho}{1-\gamma}.$$

This expression is maximized for $\check{c}_t = \hat{c}_t$, and then it simplifies to

$$B_t \leq \hat{c}_t^{-\gamma} (-\hat{c}_t h_t \hat{c}_t^{-\gamma}) + h_t \hat{c}_t^{-\gamma} \left(-\gamma \frac{\hat{c}_t^{1-\gamma}}{1-\gamma} \right) + h_t \hat{c}_t^{-\gamma} \frac{\hat{c}_t^{1-\gamma}}{1-\gamma} = 0.$$

Finally, C_t collects the terms involving volatilities

$$C_t = x_t \frac{\gamma}{2} (\sigma_t^x)^2 + h_t \hat{c}_t^{-\gamma} \left(-\frac{\gamma}{2} (\sigma_t^x)^2 - \gamma^2 \sigma_t^x \sigma_t^{\hat{c}} \right) - \bar{x}_t \frac{\gamma}{2} (\sigma_t^{\bar{x}})^2,$$

$$C_t = \frac{\gamma}{2} \left(x_t (\sigma_t^x)^2 - h_t \hat{c}_t^{-\gamma} (\sigma_t^x)^2 - h_t \hat{c}_t^{-\gamma} 2\gamma \sigma_t^x \sigma_t^{\hat{c}} - \frac{1}{\bar{x}_t} (\bar{x}_t \sigma_t^{\bar{x}})^2 \right),$$

$$C_t = \frac{\gamma}{2} \left(x_t (\sigma_t^x)^2 - h_t \hat{c}_t^{-\gamma} (\sigma_t^x)^2 - h_t \hat{c}_t^{-\gamma} 2\gamma \sigma_t^x \sigma_t^{\hat{c}} \right. \\ \left. - \frac{1}{\bar{x}_t} (x_t^2 (\sigma_t^x)^2 + (h_t \hat{c}_t^{-\gamma} \gamma \sigma_t^{\hat{c}})^2 - x_t \sigma_t^x h_t \hat{c}_t^{-\gamma} \gamma \sigma_t^{\hat{c}}) \right),$$

$$C_t = \frac{\gamma}{2} \frac{1}{\bar{x}_t} (x_t^2 (\sigma_t^x)^2 + h_t \hat{c}_t^{-\gamma} x_t (\sigma_t^x)^2 - x_t h_t \hat{c}_t^{-\gamma} (\sigma_t^x)^2 - (h_t \hat{c}_t^{-\gamma})^2 (\sigma_t^x)^2 - x_t h_t \hat{c}_t^{-\gamma} 2\gamma \sigma_t^x \sigma_t^{\hat{c}} \\ - (h_t \hat{c}_t^{-\gamma})^2 2\gamma \sigma_t^x \sigma_t^{\hat{c}} - (x_t^2 (\sigma_t^x)^2 + (h_t \hat{c}_t^{-\gamma} \gamma \sigma_t^{\hat{c}})^2 - 2x_t \sigma_t^x h_t \hat{c}_t^{-\gamma} \gamma \sigma_t^{\hat{c}})),$$

$$C_t = \frac{\gamma}{2} \frac{1}{\bar{x}_t} \left(-(h_t \hat{c}_t^{-\gamma})^2 (\sigma_t^x)^2 - (h_t \hat{c}_t^{-\gamma})^2 2\gamma \sigma_t^x \sigma_t^{\hat{c}} - (h_t \hat{c}_t^{-\gamma} \gamma \sigma_t^{\hat{c}})^2 \right),$$

$$C_t = -\frac{\gamma}{2} \frac{(h_t \hat{c}_t^{-\gamma})^2}{\bar{x}_t} \left((\sigma_t^x)^2 + 2\gamma \sigma_t^x \sigma_t^{\hat{c}} + (\gamma \sigma_t^{\hat{c}})^2 \right) = -\frac{\gamma}{2} \frac{(h_t \hat{c}_t^{-\gamma})^2}{\bar{x}_t} (\sigma_t^x + \gamma \sigma_t^{\hat{c}})^2 \leq 0.$$

¹⁹Notice that if $\phi = 1$ the expression for A_t is valid also if we allow $a_t \in \mathbb{R}$. As a result, in this case contracts with $\sigma_t^{\hat{c}} = 0$ are IC even if we allow $a_t \in \mathbb{R}$.

Putting this together and plugging into (37), we get

$$e^{-\rho t} \left(U_t^{\tilde{c},a} - \frac{\bar{x}_t^{1-\gamma}}{1-\gamma} \right) \leq \mathbb{E}_t^a \left[e^{-\rho \tau^n} \left(U_{\tau^n}^{\tilde{c},a} - \frac{\bar{x}_{\tau^n}^{1-\gamma}}{1-\gamma} \right) \right].$$

Taking the limit $n \rightarrow \infty$, it only remains to show that the tail term

$$\lim_{n \rightarrow \infty} \mathbb{E}_t^a \left[e^{-\rho \tau^n} \left(U_{\tau^n}^{\tilde{c},a} - \frac{\bar{x}_{\tau^n}^{1-\gamma}}{1-\gamma} \right) \right] \leq 0.$$

Since the agent’s strategy (\tilde{c}, a) is feasible, we have that

$$\lim_{n \rightarrow \infty} \mathbb{E}_t^a [e^{-\rho \tau^n} U_{\tau^n}^{\tilde{c},a}] = 0.$$

For $\gamma < 1$, we have $\frac{\bar{x}_t^{1-\gamma}}{1-\gamma} \geq 0$, so it follows that

$$\lim_{n \rightarrow \infty} \mathbb{E}_t^a \left[e^{-\rho \tau^n} \frac{\bar{x}_{\tau^n}^{1-\gamma}}{1-\gamma} \right] \geq 0.$$

As a result, when we take $n \rightarrow \infty$ we get $\lim_{n \rightarrow \infty} \mathbb{E}_t^a [e^{-\rho \tau^n} (U_{\tau^n}^{\tilde{c},a} - \frac{\bar{x}_{\tau^n}^{1-\gamma}}{1-\gamma})] \leq 0$ and, therefore, $U_t^{\tilde{c},a} \leq \frac{\bar{x}_t^{1-\gamma}}{1-\gamma}$ as desired.

For $\gamma > 1$, if $\lim_{n \rightarrow \infty} \mathbb{E}_t^a [e^{-\rho \tau^n} \frac{\bar{x}_{\tau^n}^{1-\gamma}}{1-\gamma}] = 0$ for any feasible strategy (\tilde{c}, a) , then we are done. To show this, notice that the law of motion of \bar{x} satisfies

$$d\bar{x}_t \leq (\lambda_1 \bar{x}_t - \lambda_2 \tilde{c}_t) dt + \sigma_t^{\bar{x}} \bar{x}_t dZ_t^a,$$

where $\lambda_2 = \hat{c}_h^{-\gamma} > 0$, and $\lambda_1 = |\bar{\mu}^x| + r + \hat{c}^{1-\gamma} + \gamma |\bar{\mu}^{\hat{c}}| + \frac{1}{2} \gamma ((1 + \gamma)(\bar{\sigma}^{\hat{c}})^2)$. Here is where we use the assumption that σ^x and $\sigma^{\hat{c}}$ are bounded and \hat{c} bounded below \hat{c}_h and away above zero. They imply that the drifts μ^x and $\mu^{\hat{c}}$ are also bounded, and we let $\bar{\mu}^x$, $\bar{\mu}^{\hat{c}}$, $\bar{\sigma}^{\hat{c}}$, and \hat{c} be appropriate bounds. Notice that stealing only reduces the drift of \bar{x}_t , since the change in the drift of x_t and h_t cancel out, and it increases the drift of \hat{c} . Since $\bar{x}_t > 0$ always, Lemma O.14 and its corollary ensure the desired limit. Q.E.D.

PROOF OF THEOREM 2: The proof is split into parts (which do not correspond to the numbers in the statement of the theorem).

(1) The cost function must be bounded above by \hat{v}_h since we can always just give consumption to the agent without any capital, and obtain cost \hat{v}_h . It must be strictly positive because if $\hat{v}(\hat{c}) = 0$ for any $\hat{c} \in [0, \hat{c}_h]$, then we can scale up the contract and give infinite utility to the agent at zero cost, or else achieve infinite profits. Because we can always move \hat{c} up using dL_t , we know that \hat{v} must be weakly increasing.

(2) Lemma O.11 in the Online Appendix (Di Tella and Sannikov (2021)) shows that \hat{v} has (1) a flat portion on $(0, \hat{c}_l)$ for some $\hat{c}_l \in (0, \hat{c}_h)$, where the HJB holds as an inequality; and (2) a strictly increasing, C^2 portion (\hat{c}_l, \hat{c}_h) with $\hat{v}'(\hat{c}) > 0$ where the HJB equation holds. At \hat{c}_l we have the smooth pasting condition $\hat{v}'(\hat{c}_l) = 0$. To understand this, it is useful to write the function $A(\hat{c}; \hat{v})$ from (22) as

$$A(\hat{c}, \hat{v}) \equiv \hat{c} - r\hat{v} - \frac{1}{2} \frac{\left(\frac{\alpha \hat{c}^\gamma}{\phi \sigma} \right)^2}{\hat{v} \gamma} + \hat{v} \frac{\rho - \hat{c}^{1-\gamma}}{1 - \gamma}, \tag{38}$$

which is the HJB equation when $\hat{v}(\hat{c})$ is flat. Lemma O.11 says that in the flat region $(0, \hat{c}_l)$ we have $A(\hat{c}, \hat{v}(\hat{c})) > 0$. This means that any contract that spends time in this region has cost strictly greater than one that immediately jumps up to \hat{c}_l and attains cost $\hat{v}(\hat{c}) = \hat{v}(\hat{c}_l)$. If the smooth pasting condition did not hold, $\hat{v}'(\hat{c}_l) > 0$, we would get a kink at \hat{c}_l , which Lemma O.11 rules out.

(3) Now let us show that $A(\hat{c}_l, \hat{v}(\hat{c}_l)) = 0$. To see this, it is useful to use the FOC for σ^x conditional on $\sigma^{\hat{c}}$ to obtain

$$\sigma^x = \frac{\frac{\alpha \hat{c}^\gamma}{\phi \sigma} - \hat{v}' \hat{c} (1 + \gamma) \sigma^{\hat{c}}}{\hat{v} \gamma + \hat{v}' \hat{c}}, \tag{39}$$

and plug it into the HJB equation. We can then rewrite the HJB

$$0 = \min_{\sigma^{\hat{c}}} \bar{A} + \bar{B} \sigma^{\hat{c}} + \frac{1}{2} \bar{C} (\sigma^{\hat{c}})^2,$$

with

$$\bar{A} = \hat{c} - r\hat{v} - \frac{1}{2} \left(\frac{\alpha \hat{c}^\gamma}{\phi \sigma} \right)^2 + \hat{v} \left(\frac{\rho - \hat{c}^{1-\gamma}}{1 - \gamma} \right) - \hat{v}' \hat{c} \left(\frac{\rho - r}{\gamma} + \frac{\rho - \hat{c}^{1-\gamma}}{1 - \gamma} \right), \tag{40}$$

$$\bar{B} = \hat{v}' \hat{c} (1 + \gamma) \frac{\frac{\alpha \hat{c}^\gamma}{\phi \sigma}}{\hat{v} \gamma + \hat{v}' \hat{c}} > 0, \tag{41}$$

$$\bar{C} = \gamma \hat{v}' \hat{c} (1 + \gamma) \frac{\hat{v} - \hat{v}' \hat{c}}{\hat{v} \gamma + \hat{v}' \hat{c}} + \hat{v}'' \hat{c}^2. \tag{42}$$

For the HJB to have a minimum, it must be that $\bar{C} > 0$, and hence

$$\sigma^{\hat{c}} = -\frac{\bar{B}}{\bar{C}} < 0 \tag{43}$$

for all $\hat{c} \in (\hat{c}_l, \hat{c}_h)$. Since $\hat{v}'(\hat{c}_l + \epsilon) > 0$ and $\hat{v}'(\hat{c}_l) = 0$, we must have $\hat{v}''(\hat{c}_l + \epsilon) \geq 0$. We can then show that $\frac{\bar{C}}{\bar{B}^2} \rightarrow \infty$ as $\hat{c} \searrow \hat{c}_l$, which implies $\frac{\bar{B}^2}{\bar{C}} \rightarrow 0$ and, therefore, $\bar{A} \rightarrow 0$. So we get that $A(\hat{c}, \hat{v}(\hat{c})) \rightarrow 0$ as $\hat{c} \searrow \hat{c}_l$, as desired. Now since at \hat{c}_l we have $A(\hat{c}, \hat{v}(\hat{c})) = 0$, this is a root of $A(\hat{c}, \hat{v}(\hat{c}_l))$. For $\hat{c} < \hat{c}_l$ we have $A(\hat{c}, \hat{v}(\hat{c}_l)) \geq 0$. From Lemma O.13, we know that this can only be the case if \hat{c}_l is the first root of $A(\hat{c}, \hat{v}(\hat{c}_l))$.

(4) Now we want to show that $\hat{v}''_+(\hat{c}_l) > 0$. Since we know \hat{c}_l is the first root of $A(\hat{c}, \hat{v}(\hat{c}_l))$, we know that $A'_1(\hat{c}_l, \hat{v}_l) \leq 0$. Consider the first-order ODE

$$\underbrace{\hat{c} - r\hat{f} - \sigma^x \hat{c}^\gamma \frac{\alpha}{\phi \sigma} + f \left(\frac{\rho - \hat{c}^{1-\gamma}}{1 - \gamma} + \frac{\gamma}{2} (\sigma^x)^2 \right)}_{A(\hat{c}, f)} + f' \hat{c} \left(\frac{\hat{c}^{1-\gamma} - \hat{c}_h^{1-\gamma}}{1 - \gamma} + \frac{(\sigma^x)^2}{2} \right) = 0$$

that fixes $\sigma^{\hat{c}} = 0$ and $\sigma^x = \frac{\alpha \hat{c}^\gamma}{\sigma \gamma \hat{v} \phi}$ in the HJB equation. Consider the solution with boundary condition $f(\hat{c}_l) = \hat{v}(\hat{c}_l)$. Since we already know that $A(\hat{c}_l, \hat{v}(\hat{c}_l)) = 0$, we must have

$f'(\hat{c}_l) = 0$ because the term in parenthesis is the drift $\mu^{\hat{c}}$ if $\sigma^{\hat{c}} = 0$, and Lemma O.15 shows this is strictly positive under these conditions. Furthermore, we must have $f''(\hat{c}_l) \leq \hat{v}''_+(\hat{c}_l)$. To see this, if $f''(\hat{c}_l) > \hat{v}''_+(\hat{c}_l) \geq 0$, then $f(\hat{c}_l + \epsilon) > \hat{v}(\hat{c}_l + \epsilon)$, since both have equal first derivative. We can then slightly lower $f(\hat{c}_l) < \hat{v}(\hat{c}_l)$. Since the ODE is locally Lipschitz continuous, the solution f changes continuously so it must intersect \hat{v} above \hat{c}_l , and the drift $\mu^{\hat{c}}$ is still strictly positive. The cost $f(\hat{c}_l) < \hat{v}(\hat{c}_l)$ is therefore attainable with such a deterministic contract that starts at \hat{c}_l , drifts up, and then reverts to the optimal contract. This cannot be, so we must have $f''(\hat{c}_l) \leq \hat{v}''_+(\hat{c}_l)$.

Differentiating the first-order ODE with respect to \hat{c} we obtain

$$0 = A'_1(\hat{c}_l, \hat{v}_l) + f''(\hat{c}_l)\hat{c}_l \underbrace{\left(\frac{\hat{c}^{1-\gamma} - \hat{c}_h^{1-\gamma}}{1-\gamma} + \frac{(\sigma^x)^2}{2} \right)}_{>0},$$

where we have used $f'(\hat{c}_l) = 0$ and the envelope theorem to compute the derivative $A'_1(\hat{c}_l, \hat{v}_l)$. If $A'_1(\hat{c}_l, \hat{v}_l) < 0$, it follows that $\hat{v}''_+(\hat{c}_l) \geq f''(\hat{c}_l) > 0$. It only remains to rule out $A'_1(\hat{c}_l, \hat{v}(\hat{c}_l)) = 0$ which means $A(\hat{c}, \hat{v}(\hat{c}_l)) > 0$ for all $\hat{c} \neq \hat{c}_l$ (from Lemma O.13). Since $A(\hat{c}, \hat{v}) = 0$ is the HJB equation if \hat{v} is flat, by a martingale verification argument as in Theorem 3 we obtain that the cost of any incentive compatible contract \mathcal{C} is strictly larger than $\hat{v}(\hat{c}_l)$, unless it has $\hat{c}_t = \hat{c}_l$ always so that $A(\hat{c}_t, \hat{v}_t) = 0$ always and $\sigma^x = \frac{\alpha}{\sigma\gamma} \frac{\hat{c}_l^\gamma}{\hat{v}_l\phi}$.

This requires $\mu^{\hat{c}} = \sigma^{\hat{c}} = 0$. But Lemma O.15 shows that for $\sigma^{\hat{c}} = 0$ and $\sigma^x = \frac{\alpha}{\sigma\gamma} \frac{\hat{c}_l^\gamma}{\hat{v}_l\phi}$, and $A(\hat{c}_l, \hat{v}_l) = 0$ we get $\mu^{\hat{c}}(\hat{c}_l) > 0$. Since the optimal contract must be incentive compatible and achieve cost $\hat{v}(\hat{c}_l)$, $A'_1(\hat{c}_l, \hat{v}(\hat{c}_l)) = 0$ cannot be. So we have $\hat{v}''_+(\hat{c}_l) > 0$.

(5) Now we can study the properties of $\sigma^{\hat{c}}$, σ^x , and $\mu^{\hat{c}}$. First, $\hat{v}''_+(\hat{c}_l) > 0$ and $\hat{v}'(\hat{c}_l) = 0$ implies $\sigma^{\hat{c}}(\hat{c}_l) = 0$ and, therefore, with Lemma O.15, we have $\mu^{\hat{c}}(\hat{c}_l) > 0$. From (43), we get that for $\hat{c} \in (\hat{c}_l, \hat{c}_h)$, we have $\sigma^{\hat{c}}(\hat{c}) < 0$, and (39) implies $\sigma^x(\hat{c}) > 0$. Since \hat{v} is C^2 , $\sigma^{\hat{c}}(\hat{c})$ and $\sigma^x(\hat{c})$ are bounded in any compact subset of (\hat{c}_l, \hat{c}_h) . The argument in part (7) below shows they are also bounded near the boundaries \hat{c}_l and \hat{c}_h , which are inaccessible from the interior, so $\sigma_t^{\hat{c}}$ and σ_t^x are bounded.

(6) Now we can show that $\hat{v} < \hat{c}^\gamma$ for all $\hat{c} \in [\hat{c}_l, \hat{c}_h)$, with $\hat{v}(\hat{c}_h) = \hat{c}_h^\gamma$. We already know from Lemma O.15 that at \hat{c}_l we have $\hat{v}(\hat{c}_l) < \hat{c}_l^\gamma$. We also know that $\hat{v}(\hat{c}) \leq \hat{v}_s(\hat{c})$ for $\hat{c} > \hat{c}_a$ (defined in (O.31)), so from Lemma O.16 if ever $\hat{v}(\hat{c}) = \hat{c}^\gamma$ for some $\hat{c} > \hat{c}_l$, it must be either that $\hat{c} = \hat{c}_h$; or that $\hat{c} \leq \hat{c}_p$ and, therefore, $A(\hat{c}, \hat{c}^\gamma) = A(\hat{c}, \hat{v}(\hat{c})) \geq 0$ and $\partial_1 A(\hat{c}, \hat{v}(\hat{c})) < 0$ (because $\hat{c} \geq \hat{c}_l > 0$). From Lemma O.13, we know that $A(\hat{c}, \hat{v})$ is positive near 0 and either has one root in \hat{c} if $\gamma \geq 1/2$, or is convex with at most two roots if $\hat{c} \leq 1/2$. This means that $A(\hat{c} - \delta, \hat{c}^\gamma) > 0$ for all $\delta \in (0, \hat{c}]$.

Now consider a distortion of the problem, changing α . We can pick an $\alpha' < \alpha$ so that $\hat{v}_{\alpha'}(\hat{c}_l^{\alpha'}) = \hat{v}(\hat{c})$, and $\hat{c}_l^{\alpha'} < \hat{c}$, because $\hat{v}_{\alpha'}(\hat{c})$ is decreasing in α' . However, $A_{\alpha'}(\hat{c}, \hat{c}^\gamma)$ is decreasing in α , so we get $A_{\alpha'}(\hat{c}_l^{\alpha'}, \hat{c}^\gamma) > A(\hat{c}_l^{\alpha'}, \hat{c}^\gamma) > 0$, which contradicts $A_{\alpha'}(\hat{c}_l^{\alpha'}, \hat{v}_{\alpha'}(\hat{c}_l^{\alpha'})) = 0$.

We have $\hat{v}(\hat{c}_h) = \hat{c}_h^\gamma$. The only way to avoid reaching \hat{c}_h and becoming absorbed there is to set $\sigma^{\hat{c}} \approx 0$ and $\sigma^x \approx 0$ near \hat{c}_h , so $\mu^{\hat{c}} \approx 0$. This means the contract spends an arbitrarily long time near \hat{c}_h with k_t arbitrarily close to zero, so the cost is $\hat{v}(\hat{c}_h) \geq \hat{v}_h$. So we conclude that $\hat{v}(\hat{c}) < \hat{c}^\gamma$ for all $\hat{c} \in [\hat{c}_l, \hat{c}_h)$ and $\hat{v}(\hat{c}_h) = \hat{c}_h^\gamma$.

(7) We want to show that \hat{c}_l and \hat{c}_h are inaccessible from (\hat{c}_l, \hat{c}_h) , so that $\hat{c}_t \in (\hat{c}_l, \hat{c}_h)$ for all $t > 0$, $L_t = 0$ always, and the Euler equation holds as an equality. First, note that we already have proven that $\sigma^{\hat{c}}(\hat{c}) < 0$ in the interior, so \hat{c}_t is a regular diffusion. For \hat{c}_l , use

(43) and replace $\hat{v}' \approx \hat{v}'_+ \epsilon$ to obtain $\sigma^{\hat{c}}(\hat{c}_l + \epsilon) \approx K\epsilon$, with $K = -\sigma^x(\hat{c}_l)(1 + \gamma)\hat{c}_l^{-1} < 0$, and recall $\mu^{\hat{c}}(\hat{c}_l)\hat{c}_l = \bar{\mu} > 0$. Compute the scale function²⁰

$$S(\hat{c}) = \int^{\hat{c}} \exp\left(-\int^y \frac{2\mu^{\hat{c}}(z)z}{(\sigma^{\hat{c}}(z)z)^2} dz\right) dy.$$

Using the approximation near \hat{c}_l , we obtain $S(\hat{c}) = -\infty$, so \hat{c}_l is inaccessible, $P\{\tau_{\hat{c}_l} < \infty\} = 0$, and nonattracting, $P\{\hat{c}_t \rightarrow \hat{c}_l\} = 0$. The speed function is

$$m(\hat{c}) = \frac{1}{(\sigma^{\hat{c}}(\hat{c})\hat{c})^2} \exp\left(\int^{\hat{c}} \frac{2\mu^{\hat{c}}(z)z}{(\sigma^{\hat{c}}(z)z)^2} dz\right),$$

and using the approximation we evaluate

$$\int_{\hat{c}_l} S(\hat{c})m(\hat{c}) d\hat{c} < \infty.$$

So we conclude that \hat{c}_l is an entrance boundary: \hat{c}_t starts at \hat{c}_l and immediately moves into the interior and never returns.

To show that \hat{c}_h is inaccessible, we use the approximation in Lemma O.17, $\mu^{\hat{c}}\hat{c} = (4\gamma - 6(1 + \gamma)^2)\hat{c}_h^{-\gamma}(\hat{c}_h - \hat{c})^2$ and $(\sigma^{\hat{c}}\hat{c})^2 = 8(1 + \gamma)^2\hat{c}_h^{-\gamma}(\hat{c}_h - \hat{c})^3$. Plug into the expression for the scale function to obtain $S(\hat{c}) = \text{const} \times (\frac{-1}{K+1})(\hat{c}_h - \hat{c})^{K+1}$, where $K = \frac{\gamma}{(1+\gamma)^2} - \frac{3}{2} < -1$ for any $\gamma > 0$. We take $\hat{c} \rightarrow \hat{c}_h$ and obtain $S(\hat{c}_h) = \infty$. This means \hat{c}_h is inaccessible and nonattracting, too. In fact, we can also show that \hat{c}_h is a natural boundary because

$$\int^{\hat{c}_h} S(\hat{c})m(\hat{c}) d\hat{c} = \infty.$$

(8) Now we show that the optimal contract of the relaxed problem does not have a stationary distribution. In the long-run, it spends almost all the time near \hat{c}_h ,

$$\frac{1}{t} \int_0^t 1_{\{\hat{c} > \hat{c}_h - \epsilon\}}(\hat{c}_s) ds \rightarrow 1 \quad \text{a.s. } \forall \epsilon > 0.$$

Using the same approximation near \hat{c}_h in Lemma O.17,

$$\begin{aligned} \mu^{\hat{c}}\hat{c} &= (4\gamma - 6(1 + \gamma)^2)\hat{c}_h^{-\gamma}(\hat{c}_h - \hat{c})^2, \\ (\sigma^{\hat{c}}\hat{c})^2 &= 8(1 + \gamma)^2\hat{c}_h^{-\gamma}(\hat{c}_h - \hat{c})^3, \end{aligned}$$

we compute the speed measure

$$m(\hat{c}) = (\sigma^{\hat{c}}(\hat{c})\hat{c})^{-2} \exp\left(\int^{\hat{c}} \frac{2\mu^{\hat{c}}(z)z}{(\sigma^{\hat{c}}(z)z)^2} dz\right) = \text{const} \times (\hat{c}_h - \hat{c})^{-K-3},$$

²⁰See Karlin and Taylor (1981, Chapter 15, part 6).

where $K = \frac{\gamma}{(1+\gamma)^2} - \frac{3}{2} < -1$, so that $-(K + 3) = -(\frac{\gamma}{(1+\gamma)^2} + \frac{3}{2}) < -3/2$. If there is a stationary distribution it must be proportional to $m(\hat{c})$ and integrate to 1. But

$$\int^{\hat{c}} m(y) dy = \frac{1}{-K - 2} (\hat{c}_h - \hat{c})^{-K-2}.$$

When we take $\hat{c} \rightarrow \hat{c}_h$ we find that $\int^{\hat{c}_h} m(y) dy = \infty$ because $-(K + 2) < 0$, which means there cannot be a stationary distribution.

The same computation near \hat{c}_l shows that $\int_{\hat{c}_l}^{\hat{c}_l+\epsilon} m(\hat{c}) d\hat{c} < \infty$. This means that

$$\frac{1}{t} \int_0^t 1_{\{\hat{c} < \hat{c}_h - \epsilon\}}(\hat{c}_s) ds \rightarrow \frac{\int_{\hat{c}_l}^{\hat{c}_h - \epsilon} m(\hat{c}) d\hat{c}}{\int_{\hat{c}_l}^{\hat{c}_h} m(\hat{c}) d\hat{c}} = 0 \quad \text{a.s. } \forall \epsilon > 0,$$

which in turn implies that

$$\frac{1}{t} \int_0^t 1_{\{\hat{c} > \hat{c}_h - \epsilon\}}(\hat{c}_s) ds \rightarrow 1 \quad \text{a.s. } \forall \epsilon > 0.$$

$P\{\hat{c}_t \rightarrow \hat{c}_h\} = 0$ follows from $S(\hat{c}_h) = \infty$.

(9) Finally, we show the relaxed optimal contract is globally incentive compatible and therefore an optimal contract. We know $0 < \hat{c}_l \leq \hat{c}_t \leq \hat{c}_h$ for all t , and σ^x and $\sigma^{\hat{c}}$ are bounded. Most importantly, $\sigma_t^{\hat{c}} \leq 0$ always, so we can use Theorem 1 to obtain our result. *Q.E.D.*

PROOF OF THEOREM 3: Consider any locally incentive compatible contract $\mathcal{C} = (c, k)$ that delivers utility of at least u_0 to the agent, with associated state variables x and \hat{c} . Because $\hat{v}'(\hat{c}_l) = 0$ we can use Ito's lemma²¹ and the HJB equation to obtain

$$\begin{aligned} e^{-r\tau^n} \hat{v}(\hat{c}_{\tau^n}) x_{\tau^n} &\geq \hat{v}(\hat{c}_0) x_0 - \int_0^{\tau^n} e^{-rt} (\hat{c}_t - \hat{k}_t \alpha) x_t dt \\ &\quad + \int_0^{\tau^n} e^{-rt} \hat{v}(\hat{c}_t) x_t \left(\frac{\hat{v}'(\hat{c}_t)}{\hat{v}(\hat{c}_t)} \hat{c}_t \sigma_t^{\hat{c}} + \sigma_t^x \right) dZ_t, \end{aligned}$$

for localizing sequence of stopping times $\{\tau^n\} \rightarrow \infty$ a.s. such that the last term has expectation zero. Take expectations to obtain

$$\mathbb{E}_0^Q [e^{-r\tau^n} \hat{v}(\hat{c}_{\tau^n}) x_{\tau^n}] \geq \hat{v}(\hat{c}_0) x_0 - \mathbb{E}_0^Q \left[\int_0^{\tau^n} e^{-rt} (c_t - k_t \alpha) dt \right]. \tag{44}$$

Now we would like to take the limit $n \rightarrow \infty$. For the second term on the right, we can use the dominated convergence theorem, using the fact that locally incentive compatible contracts are admissible.

²¹Notice \hat{v}' is discontinuous at \hat{c}_l , but this does not change Ito's formula. See Proposition 4.12 in Harrison (2013).

Also, for an admissible contract we have

$$0 \leq \lim_{n \rightarrow \infty} \mathbb{E}_0^Q [e^{-r\tau^n} \hat{v}(\hat{c}_{\tau^n}) x_{\tau^n}] \leq \lim_{n \rightarrow \infty} \mathbb{E}_0^Q [e^{-r\tau^n} \hat{v}_h x_{\tau^n}] = 0.$$

To see why the last equality holds, notice that since $\hat{v}_h x$ is the cheapest way of delivering utility to the agent without capital, the cost of consumption on the contract is

$$\infty > \mathbb{E}^Q \left[\int_0^\infty e^{-rt} c_t dt \right] \geq \mathbb{E}^Q \left[\int_0^{\tau^n} e^{-rt} c_t dt + e^{-r\tau^n} \hat{v}_h x_{\tau^n} \right].$$

Taking the limit $n \rightarrow \infty$ and using the dominated convergence theorem, we obtain $0 \leq \lim_{n \rightarrow \infty} \mathbb{E}^Q [e^{-r\tau^n} \hat{v}_h x_{\tau^n}] \leq 0$.

Upon taking the limit $n \rightarrow \infty$ in (44), we obtain

$$\mathbb{E}_0^Q \left[\int_0^\infty e^{-rt} (c_t - k_t \alpha) dt \right] \geq \hat{v}(\hat{c}_0) x_0.$$

Using $\hat{v}(\hat{c}_t) \leq \hat{v}(\hat{c}_0)$ and $x_0 \geq ((1 - \gamma)u_0)^{\frac{1}{1-\gamma}}$, we obtain the first result.

For the second part, first we use the same steps as in the proof of Theorem 2 to show that $\hat{c}_t^* \in (\hat{c}_l, \hat{c}_h)$ for all $t > 0$, and that $\sigma_t^{\hat{c}^*} < 0$, $\sigma_t^{x^*} > 0$, $\mu_t^{x^*}$, and $\mu_t^{\hat{c}^*}$ are all bounded. We also know that, given $x_0 > 0$, any solution x^* to (16) is strictly positive. Lemma 4 then ensures C^* delivers utility $U_t^{c^*,0} = (x_t^*)^{1-\gamma}/(1-\gamma)$, with $U_0^{c^*,0} = u_0$. Since C^* is admissible, it is locally incentive compatible. The same argument as in the first part shows the cost of C^* is $\hat{v}_h x_0^*$. C^* is therefore a relaxed optimal contract, which implies it is also an optimal contract. *Q.E.D.*

REFERENCES

- ADRIAN, T., AND N. BOYARCHENKO (2012): "Intermediary Leverage Cycles and Financial Stability," Federal Reserve Bank of New York Staff Reports, 567. [1102]
- AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 659–684. [1102]
- ANGELETOS, M. (2007): "Uninsured Idiosyncratic Investment Risk and Aggregate Saving," *Review of Economic Dynamics*, 10, 1–30. [1102]
- BATTAGLINI, M. (2005): "Long-Term Contracting With Markovian Consumers," *American Economic Review*, 95, 637–658. [1125]
- BERNANKE, B., AND M. GERTLER (1989): "Agency Costs, Net Worth, and Business Fluctuations," *The American Economic Review*, 79, 14–31. [1102]
- BERNANKE, B., M. GERTLER, AND S. GILCHRIST (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework," in *Handbook of Macroeconomics*, Vol. 1, ed. by J. Taylor and M. Woodford. Elsevier, 1341–1393, Chapter 21. [1102]
- BIAIS, B., T. MARIOTTI, G. PLANTIN, AND J. ROCHET (2007): "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications," *Review of Economic Studies*, 74, 345–390. [1101]
- BRUNNERMEIER, M., AND Y. SANNIKOV (2014): "A Macroeconomic Model With a Financial Sector," *American Economic Review*, 104, 379–421. [1102]
- CHRISTIANO, L., R. MOTTO, AND M. ROSTAGNO (2014): "Risk Shocks," *American Economic Review*, 104. [1102]
- CISTERNAS, G. (2018): "Two-Sided Learning and the Ratchet Principle," *Review of Economic Studies*, 85, 307–351. [1102]
- CLEMENTI, G., AND H. HOPENHAYN (2006): "A Theory of Financing Constraint and Firm Dynamics," *The Quarterly Journal of Economics*, 121, 229–265. [1101]
- COLE, H. L., AND N. R. KOCHERLAKOTA (2001): "Efficient Allocations With Hidden Income and Hidden Storage," *The Review of Economic Studies*, 68, 523–542. [1102]

- DEMARZO, P., AND M. FISHMAN (2007): "Optimal Long Term Financial Contracting," *Review of Financial Studies*, 20, 2079–2127. [1101]
- DEMARZO, P., AND Y. SANNIKOV (2006): "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," *The Journal of Finance*, 61, 2681–2724. [1101]
- (2016): "Learning, Termination, and Payout Policy in Dynamic Incentive Contracts," *The Review of Economic Studies*, 84, 182–236. [1102,1125]
- DEMARZO, P., M. FISHMAN, Z. HE, AND N. WANG (2012): "Dynamic Agency and the Q-Theory of Investment," *The Journal of Finance*, 67, 2295–2340. [1101]
- DI TELLA, S. (2017): "Uncertainty Shocks and Balance Sheet Recessions," *Journal of Political Economy*, 125, 2038–2081. [1102]
- (2019): "Optimal Regulation of Financial Intermediaries," *American Economic Review*, 109, 271–313. [1102,1116]
- DI TELLA, S., AND Y. SANNIKOV (2021): "Supplement to 'Optimal Asset Management Contracts With Hidden Savings,'" *Econometrica Supplemental Material*, 89, <https://doi.org/10.3982/ECTA14929>. [1104,1105,1118,1119,1121,1122,1124,1132]
- EDMANS, A., X. GABAIX, T. SADZIK, AND Y. SANNIKOV (2011): "Dynamic CEO Compensation," *The Journal of Finance*, 67, 1603–1647. [1101]
- FARHI, E., AND I. WERNING (2013): "Insurance and Taxation Over the Life Cycle," *The Review of Economic Studies*, 80, 596–635. [1100,1102]
- FERNANDES, A., AND C. PHELAN (2000): "A Recursive Formulation for Repeated Agency With History Dependence," *Journal of Economic Theory*, 91, 223–247. [1102,1125]
- GARRETT, D. F., AND A. PAVAN (2015): "Dynamic Managerial Compensation: A Variational Approach," *Journal of Economic Theory*, 159, 775–818. [1102]
- GERTLER, M., AND P. KARADI (2011): "A Model of Unconventional Monetary Policy," *Journal of Monetary Economics*, 58, 17–34. [1102]
- GERTLER, M., N. KIYOTAKI et al. (2010): "Financial Intermediation and Credit Policy in Business Cycle Analysis," *Handbook of Monetary Economics*, 3, 547–599. [1102]
- HARRISON, J. M. (2013): *Brownian Models of Performance and Control*. Cambridge University Press. [1136]
- HART, O., AND J. MOORE (1994): "A Theory of Debt Based on the Inalienability of Human Capital," *The Quarterly Journal of Economics*, 109, 841–879. [1100]
- HE, Z. (2011): "Dynamic Compensation Contracts With Private Savings," *Review of Financial Studies*. [1101]
- HE, Z., AND A. KRISHNAMURTHY (2012): "A Model of Capital and Crises," *Review of Economic Studies*, 79, 735–777. [1102]
- (2013): "Intermediary Asset Pricing," *The American Economic Review*, 103, 732–770. [1102]
- HE, Z., B. WEI, J. YU, AND F. GAO (2017): "Optimal Long-Term Contracting With Learning," *The Review of Financial Studies*, 30, 2006–2065. [1102,1125]
- IACOVIELLO, M. (2005): "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *The American Economic Review*, 95, 739–764. [1102]
- KAPIČKA, M. (2013): "Efficient Allocations in Dynamic Private Information Economies With Persistent Shocks: A First-Order Approach," *The Review of Economic Studies*, rds045. [1102]
- KARLIN, S., AND H. E. TAYLOR (1981): *A Second Course in Stochastic Processes*. Elsevier. [1135]
- KIYOTAKI, N., AND J. MOORE (1997): "Credit Cycles," *Journal of Political Economy*, 105, 211–248. [1100,1102]
- KOCHERLAKOTA, N. (2004): "Figuring out the Impact of Hidden Savings on Optimal Unemployment Insurance," *Review of Economic Dynamics*, 7, 541–554. [1100]
- KRUSELL, P., AND A. A. JR. SMITH (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896. [1102]
- LORENZONI, G. (2008): "Inefficient Credit Booms," *Review of Economic Studies*, 75 (3), 809–833. [1102]
- MOLL, B. (2014): "Productivity Losses From Financial Frictions: Can Self-Financing Undo Capital Misallocation?" *The American Economic Review*, 104, 3186–3221. [1102]
- MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73. [1125]
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 82, 601–653. [1125]
- PHELAN, C., AND R. TOWNSEND (1991): "Computing Multi-Period, Information-Constrained Optima," *Review of Economic Studies*, 58, 853–881. [1101]
- ROGERSON, W. P. (1985): "Repeated Moral Hazard," *Econometrica: Journal of the Econometric Society*, 69–76. [1101]
- SANNIKOV, Y. (2008): "A Continuous-Time Version of the Principal: Agent Problem," *The Review of Economic Studies*, 75, 957–984. [1101]

- SPEAR, S., AND S. SRIVASTAVA (1987): "On Repeated Moral Hazard With Discounting," *The Review of Economic Studies*, 54, 599–617. [1101]
- WERNING, I. (2001): "Repeated Moral Hazard With Unmonitored Wealth: A Recursive First-Order Approach," Working paper, MIT. [1100]
- WILLIAMS, N. (2008): "On Dynamic Principal-Agent Problems in Continuous Time," Working paper, U of Wisconsin–Madison. [1101]
- (2011): "Persistent Private Information," *Econometrica*, 79, 1233–1275. [1102]

Co-editor Dirk Bergemann handled this manuscript.

Manuscript received 13 December, 2016; final version accepted 15 November, 2020; available online 3 December, 2020.