Persistent Private Information Revisited*

Alexander W. Bloedel[†] R. Vijay Krishna[‡] Bruno Strulovici[§]

April 29, 2023

Abstract

This paper revisits Williams' (2011) (henceforth PPI's) continuous-time principal-agent model of optimal dynamic insurance with persistent private information. We identify three independent issues in PPI that implicate its characterizations of incentive compatible and optimal contracts: (i) the agent cannot over-report *increments of* his type, a constraint that does not follow from the common assumption that the agent cannot over-report his type; (ii) the agent's feasible set of reporting strategies does not include standard "no Ponzi" constraints, without which PPI's main analysis of infinite-horizon incentive compatibility is incomplete; and (iii) most importantly, in PPI's main application, which concerns hidden endowments, the contract identified as optimal is *generically strictly suboptimal*. For this application, we address the three issues by analyzing a class of "self-insurance contracts" that can be implemented as consumption-saving problems for the agent, and which includes the contract derived in PPI as a particular case. We characterize the optimal self-insurance contract and show that, generically, it strictly dominates PPI's. Our analysis does not support PPI's main economic finding that immiseration generally fails or its attribution of this failure to continuous time and persistence.

1. Introduction

In an influential paper, Williams (2011) (henceforth PPI) introduces a continuous-time framework to study optimal contracts in dynamic principal-agent settings where the

*First version: November 2020. Parts of this paper build on and subsume portions of the working papers Strulovici (2011) and Bloedel, Krishna, and Leukhina (2018). We thank Jakša Cvitanić, Arash Fahim, Ilya Segal, Neng Wang, and various seminar and conference audiences for helpful comments. We are particularly grateful to the Co-Editor and three anonymous referees for detailed feedback that greatly improved the paper. Bloedel is grateful to Caltech for its hospitality and financial support while part of this research was conducted, and to the B.F. Haley and E.S. Shaw Fellowship for Economics for financial support through a grant to the Stanford Institute for Economic Policy Research.

†UCLA. Email: abloedel@econ.ucla.edu

‡Florida State University. Email: rvk3570@gmail.com

§Northwestern University. Email: b-strulovici@northwestern.edu

agent's private information is persistent. In the present paper, we identify and address three independent issues in PPI's model formulation and analysis, elucidate their consequences for PPI's conclusions, and discuss implications for the broader literature. In doing so, we develop some new results and methods that may be of independent interest and highlight some open questions that we hope will stimulate future work.

Overview of PPI. PPI aims to make two methodological contributions and one economically substantive contribution (see pp. 1233–35).[1]

The methodological contributions are, first, to formulate a first-order approach to incentive compatibility (IC) in contracting models with a persistent private state (§§2–3) and, second, to provide a set of sufficient conditions under which that approach is valid (§4). Many of these methods have been fruitfully applied in the subsequent literature to shed light on new and difficult problems in contract theory.[2]

For the substantive contribution, PPI applies these techniques to study *optimal contracts*, formulating the principal's optimization problem in an abstract setting (§5) and explicitly solving for optimal contracts in a canonical class of dynamic insurance problems in which the agent's private information concerns his endowment (§§6–7) or taste shocks (§8). PPI's analysis appears to overturn common wisdom about the qualitative features of dynamic insurance contracts, which are the subject of longstanding literatures in both micro- and macroeconomic theory.[3] Specifically, in both applications, PPI's key finding is that the optimal contract sends the agent to *bliss*: the agent's consumption and utility converge almost surely to their *upper* bounds. This finding goes against the classic literature's hallmark result that optimal insurance contracts generate *immiseration*: the agent's consumption and utility converge almost surely to their *lower* bounds. As the classic literature focuses on discrete-time models with i.i.d. private information, PPI attributes this discrepancy to fundamental differences in the agent's IC constraints driven by continuous time and persistence (e.g., see p. 1235 and pp. 1256–58).

Three Issues with PPI. PPI's model formulation and analysis contain three issues that implicate most aspects of the paper's latter two contributions (sufficient conditions for IC and analysis of optimal contracts).[4] We summarize these issues below. We list them

[1] Throughout, we use page numbers and the § symbol to refer, respectively, to pages and sections in PPI. We reserve the term "section" for reference to sections in the present paper.

[2] For instance, Chen (2021), Cisternas (2017), DeMarzo and Sannikov (2016), Prat and Jovanovic (2014), Ramos and Sadzik (2019), and Sannikov (2014) adopt versions of PPI's first-order approach and aspects of its proof strategy for establishing full incentive compatibility.

[3] See Green (1987), Thomas and Worrall (1990), Atkeson and Lucas (1992), and Phelan (1998) for classic contributions, Kocherlakota (2010) and Golosov, Tsyvinski, and Werquin (2016) for surveys, and Bloedel, Krishna, and Leukhina (2021) for a recent contribution and additional references.

[4] These issues do not have any notable implications for PPI's first contribution (the first-order approach to IC). Meanwhile, PPI's economic conclusions about the failure of immiseration and the associated roles

in their logically most linear order, but note that the third issue is the most substantive.

1. *Sign restrictions on misreports:* PPI's description of the model includes the standard assumption that the agent can only under-report his type. However, PPI's actual analysis uses the significantly stronger restriction that the agent can only under-report *each increment of* his type, implying that the agent *cannot correct for any past under-reports*.[5] This restriction, which does not follow from the initial under-reporting assumption (Observation 1), is conceptually problematic and has no known analogues in the literature. Without it, however, the class of IC contracts shrinks and the analysis of IC requires different arguments. In particular, the proof of PPI's general sufficient conditions for IC (Theorem 4.1, p. 1247 and §A.2) and the verification of IC in PPI's main application (§6 and §A.3.1) rely on this stronger restriction.
2. *Tail restrictions on misreports:* PPI initially formulates the agent's reporting problem in a finite-horizon setting (§§2–4). However, the entire formulation and analysis of *optimal* contracts takes place in an infinite-horizon setting (§§5–8), informally motivated as a limit of finite-horizon models. In this infinite-horizon setting, PPI omits standard "no Ponzi" restrictions on the asymptotic growth rate of the agent's feasible reporting strategies and does not check whether the agent's value function satisfies the appropriate transversality condition. PPI's main analysis of infinite-horizon IC is incomplete as a result of these omissions. In particular, the contract identified as optimal in PPI's main application (defined below as Contract PPI) generically violates IC under PPI's assumptions (Observation 2).
3. *Generic suboptimality of the contract identified in PPI:* PPI's main application concerns a risk-sharing model in which the agent has a privately observed endowment (§§6–7).[6] Due to an incorrect numerical characterization of the optimal contract's initial condition (see Appendix A for details), the contract that PPI identifies as optimal (Contract PPI) is, in fact, strictly suboptimal in the generic case that the agent's endowment has non-zero mean-reversion (Observation 3). Consequently, PPI's main economic conclusions about the failure of immiseration and the associated roles of continuous time and persistence do not follow from that paper's analysis. Closely related results in the recent literature further suggest that these conclusions are, in fact, incorrect (see Section 7.2).

These three issues are *logically independent* of each other: even if any two are addressed,

of continuous-time and persistence, which we show are not warranted, have been informally echoed in the literature (e.g., Zhang (2009, pp. 637,652); Kapička (2013, p. 1029); Prat and Jovanovic (2014, p. 885)).

[5] In particular, this restriction prevents the agent from engaging in "one-shot deviations" that consist of under-reporting his type for a small amount of time and truthfully reporting his type thereafter.

[6] PPI's other application (to a taste shock model in §8) considers only the special case in which the agent's type has exactly zero mean-reversion and is essentially equivalent (modulo a change of variables) to the analogous special case of the main hidden endowment application. Thus, our analysis effectively covers both of PPI's applications.

the third remains problematic. They are also essentially *unrelated to PPI's continuous-time formulation*: the same observations apply almost verbatim to the natural discrete-time versions of PPI's model and Contract PPI. The generic suboptimality of Contract PPI—which we view as the most important and economically substantive issue—remains an issue regardless of the sign or tail restrictions that one imposes on the agent's feasible set, and independently of whether one formulates the model in continuous or discrete time. We study the first two issues primarily to emphasize this point.

Addressing the Issues. After recalling PPI's general model (Section 2) and main application (Section 3), we describe the three issues summarized above (Section 4). These issues implicate the general analysis in PPI. For our main analysis (Sections 5 and 6), we then turn our attention to PPI's primary application, in which the agent has a privately observed endowment that follows an OU process (i.e., continuous-time version of a Gaussian AR(1) process) and CARA utility over consumption.

In this setting, we address the first two issues by (i) allowing for a range of sign restrictions on the agent's feasible misreports (including arbitrary over- and under-reporting) and (ii) imposing a standard "no Ponzi" constraint on the agent that limits the asymptotic growth rate of his misreports. This allows us to demonstrate that (a) under our "no Ponzi" constraint, Contract PPI is IC given *any* sign restrictions on the agent's misreports, and (b) our "no Ponzi" constraint on the agent is a minimal sufficient condition for Contract PPI to be IC. We then elucidate the third issue by showing that Contract PPI is strictly suboptimal in the generic case that the agent's endowment has non-zero mean-reversion. We approach the suboptimality issue in two steps:

1. We first observe that *the agent's consumption under Contract PPI is characterized by a standard Euler equation*, and therefore coincides with the solution to a standard consumption-saving problem for the agent, in which there is no principal and the agent simply self-insures by investing in a risk-free bond at the ambient market rate (which is equal to the discount rate). Consequently, PPI's main economic results, which concern the long-run properties and comparative statics of Contract PPI, can be interpreted as standard results about precautionary savings in self-insurance problems. Notably, the classic contracting literature has found that optimal contracts typically do *not* coincide with solutions to self-insurance problems (e.g., Allen (1985); Cole and Kocherlakota (2001)).
2. We then introduce a new class of *Self-Insurance Contracts (SI Contracts)* defined in terms of an indirect implementation in which the principal acts as the agent's "bank."[7] Specifically, the principal provides the agent with some initial wealth and

[7] As we discuss in Section 5.1, versions of our solution to the agent's self-insurance problem, which in our setting arises as his best-response within an SI Contract, appear in the self-insurance literature (Caballero 1990; Wang 2003, 2006). To the best of our knowledge, neither our definition nor our analysis of optimality for SI Contracts (Section 5.2–5.4) has close analogues in that literature.

then allows the agent to self-insure at a risk-free interest rate equal to the sum of the ambient market rate and a savings tax that the principal charges the agent. As suggested by Step 1, Contract PPI is the specific SI Contract with *zero taxes*. However, whenever the agent's endowment has non-zero mean-reversion, we show that the *optimal* SI Contract imposes *strictly positive taxes*. We interpret this result in terms of the agent's precautionary savings behavior.

We also derive some additional properties of SI Contracts that may be of independent interest beyond this paper.

In Section 6, we provide two proofs that all SI Contracts, including Contract PPI, are IC when re-formulated as direct mechanisms: one based on an "indirect" revelation principle argument and the other based on a "direct" analysis of the agent's reporting problem.[8] Each approach is independently instructive, requires different analysis than that in PPI, reveals similarities between IC in continuous- and discrete-time models, and may be useful for analyzing IC in other continuous-time contracting models.

(Fully) Optimal Contracts. Our analysis of SI Contracts raises two questions:

1. *Are there conditions under which Contract PPI is, in fact, optimal?* We provide two positive answers. First, it is optimal in an *alternative model* in which the agent can covertly save and borrow at the market rate ("hidden savings"), linking PPI's analysis to classic discrete-time studies of optimal hidden savings contracts (Allen 1985; Cole and Kocherlakota 2001). Second, it is optimal in the non-generic case of PPI's *original model* in which the agent's endowment has exactly zero mean-reversion ("permanent shocks"), confirming PPI's finding in this special case. We thus obtain optimality foundations for Contract PPI, albeit in distinct and rather specific environments.[9]
2. *Is the optimal SI Contract "fully" optimal (i.e., among all IC contracts)?* We offer a negative answer: under regularity conditions, the fully optimal contract strictly dominates the optimal SI Contract whenever the agent's endowment has non-zero mean-reversion ("transient shocks"). While our analysis of SI Contracts has some implications for the fully optimal contract under transient shocks, a full characterization of the latter contract remains an important open problem.

We summarize these analyses in Sections 5.4 and 7 (details are in Appendices H and J).

[8]That is, for the "indirect" approach, we construct an explicit mapping between consumption-saving strategies in the indirect mechanism and reporting strategies in the direct mechanism. One implication of this mapping is that our "no Ponzi" restriction on reporting strategies is equivalent to the standard no Ponzi condition on savings in the agent's self-insurance problem.

[9]The assumption of hidden savings, while realistic in some settings, corresponds to a fundamentally different model than that studied in PPI and much of the social insurance literature, which focuses on the implications of the agent's private information in isolation. The assumption of permanent endowment shocks is knife-edge and constitutes a fundamental departure from the classic literature with i.i.d. types: as we discuss in Section 7, it leads to qualitatively different tradeoffs for the principal than arise under type processes with even arbitrarily slow mean-reversion.

Broader Implications. In Section 7, we discuss implications of our analysis for the interpretation of PPI's results and for the broader literature.

- *Immiseration and Persistence:* PPI's main economic conclusion that immiseration fails under persistent private information is based on the analysis of Contract PPI. We find that this conclusion is warranted *only* in the non-generic case of permanent shocks; it is driven by the *absence* of mean-reversion in the agent's type process—*rather than persistence per se or continuous time*, as asserted in PPI. We further argue that immiseration should, in fact, be expected to hold under the optimal contract in PPI's model when shocks are transient.
- *Continuous vs. Discrete Time:* A central claim in PPI is that its results differ from those in the prior discrete-time literature because IC constraints are qualitatively different in discrete- and continuous-time models. Our analysis casts doubt on this claim: our SI Contracts (including Contract PPI) have precise discrete-time analogues, and our analysis of IC reveals fundamental parallels between incentive constraints in continuous- and discrete-time settings.

2. Model

Section 2.1 introduces PPI's general model. Section 2.2 introduces various possible restrictions on the agent's feasible set of reporting strategies.

2.1. Environment

Time is continuous and runs over an infinite horizon. At $t = 0$, a risk-neutral principal (she) offers an insurance contract to a risk-averse agent (he). Once the contract is signed, neither party may renege at a later date.

Type Process. At each instant t, the agent privately observes his *type*, $b_t \in \mathbb{R}$. The agent's *type process* $b = (b_t)_{t \geqslant 0}$ evolves according to the equation

[**2.1**] $$\mathrm{d}b_t = (\mu - \lambda b_t)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t,$$

where $\sigma > 0$ and $W = (W_t)_{t \geqslant 0}$ is a standard Brownian motion. The initial condition $b_0 \in \mathbb{R}$ is common knowledge. The parameter $\lambda \geqslant 0$ specifies the rate of mean reversion. When $\lambda = 0$, b is a Brownian motion with constant drift. We refer to this as the *permanent shock* case because the time-t shock (the Brownian increment $\mathrm{d}W_t$) has a non-vanishing additive effect on b_T for all $T > t$. When $\lambda > 0$, b is an Ornstein-Uhlenbeck (OU) process. We refer to this as the *transient shock* case because the time-t shock has a vanishing effect on b_T as $T \to \infty$. Note that *smaller* values of λ correspond to *greater* persistence.[10]

[10]Given any realization of b_t and time $T > t$, the solution to [**2.1**] is $b_T = \mu/\lambda + (b_t - \mu/\lambda)e^{-\lambda(T-t)} + e^{-\lambda(T-t)} \int_t^T \sigma e^{\lambda\tau}\,\mathrm{d}W_\tau$. It follows that $\frac{\mathrm{d}b_T}{\mathrm{d}b_t} = e^{-\lambda(T-t)}$. In the language of Pavan, Segal, and Toikka (2014),

We let $\mathbf{P}$ denote the probability measure over paths of b.[11]

Reporting Strategies. At each instant t, the agent reports a type $y_t \in \mathbb{R}$. The process $y = (y_t)_{t \geqslant 0}$ is the agent's *reporting strategy*, and is assumed to be adapted to the filtration of process b (henceforth, b-adapted).[12] The process $m = (m_t)_{t \geqslant 0}$ where $m_t := y_t - b_t$ is the agent's *misreporting strategy*, and is also b-adapted. Thus, $m_t > 0$ corresponds to over-reporting one's type while $m_t < 0$ corresponds to under-reporting it. Clearly, each reporting strategy y uniquely defines a corresponding misreporting strategy m and vice versa. However, it is useful to distinguish between these objects because the principal only observes the realized sample path of y, while the agent observes the realized sample path of b and his own misreporting strategy. Every misreporting strategy m induces a probability measure $\mathbf{P}^m$ over paths of y. We let $y^* := b$ and $m^* := 0$ denote the *truthful* reporting and misreporting strategies, respectively, and let $\mathbf{P}^* := \mathbf{P}^{m^*}$ denote the corresponding measure over report paths. Note that while $\mathbf{P}^*$ coincides with $\mathbf{P}$, they are measures over paths of different processes (y and b, respectively).

As in PPI, we assume that the agent's misreports have absolutely continuous sample paths, i.e., there exists a process $\Delta = (\Delta_t)_{t \geqslant 0}$ such that $m_t \equiv \int_0^t \Delta_s \, \mathrm{d}s$ (where "$\equiv$" denotes that $\mathbf{P}$-a.s., the processes are equal for almost all $t \geqslant 0$). Thus, the agent's report evolves as $\mathrm{d}y_t = \mathrm{d}b_t + \Delta_t \, \mathrm{d}t$, where the drift adjustment Δ_t corresponds to misreporting the *increment* $\mathrm{d}b_t$. Let $\mathscr{M}$ denote the space of such misreporting strategies.

Contracts. A *contract* is a continuous y-adapted process $s = (s_t)_{t \geqslant 0}$ that specifies transfers (of the consumption good) from the principal to the agent as a function of the history of reports.

Agent's Incentives. The agent's type determines his preference over consumption: if the agent's current type is b_t and he receives transfer $s_t \in \mathbb{R}$ from the principal, his flow utility is $v(s_t, b_t)$. The agent may be restricted to a subset of misreporting strategies, his *feasible set* $F \subseteq \mathscr{M}$, that contains the truthful strategy m^*. Given the agent's feasible set

$e^{-\lambda(T-t)}$ is the *impulse response* of b_T to b_t. Smaller values of λ correspond to larger impulse responses; when $\lambda = 0$, the impulse responses are identically 1.

[11]Formally, this requires specifying the measure space of sample paths. Following PPI (p. 1239), we use the space of continuous paths $\mathbf{C}([0, \infty))$ endowed with the standard Borel sigma-algebra (cf. Appendix C).

[12]That is, y_t is measurable with respect to the sigma-algebra generated by the paths $(b_\tau)_{\tau \in [0,t]}$. PPI further assumes that y is *predictable* with respect to the filtration generated by b (i.e., is b-adapted and, moreover, y_t does not depend on the contemporaneous type b_t). However, this distinction is inconsequential under PPI's assumption, which we adopt below, that the misreporting process $m := y - b$ has absolutely continuous sample paths. (When y has continuous paths, it is adapted if and only if it is predictable.) An analogous qualifier is relevant when we define contracts below as being y-adapted, rather than y-predictable. We assume mere adaptedness here because it is the appropriate assumption when, in Section 6, we consider extending the agent's strategy space to allow for discontinuous "jump" reports (see Section 6, Section 7.3, and Appendix I for further discussion).

F, a contract is said to be *F-incentive compatible (F-IC)* if it satisfies[13]

[IC]
$$m^* \in \arg\max_{m \in F} \mathbf{E}_0^m \left[\int_0^\infty e^{-\rho t} v(s_t, b_t)\, \mathrm{d}t \right],$$

where $\rho > 0$ is the agent's discount rate.

Principal's Problem. The principal also has discount rate $\rho > 0$. A standard interpretation, which will be important later, is that ρ represents the interest rate at which the principal finances the contract on a risk-free bond market. Given the agent's feasible set $F \subseteq \mathscr{M}$, the principal chooses a contract to minimize the expected lifetime cost of transfers to the agent (under truthful reporting)

[2.2]
$$\mathbf{E}_0^* \left[\int_0^\infty e^{-\rho t} s_t\, \mathrm{d}t \right]$$

subject to the contract (i) being F-IC and (ii) satisfying the *promise keeping* constraint $q_0 \leqslant \mathbf{E}_0^* \left[\int_0^\infty e^{-\rho t} v(s_t, b_t)\, \mathrm{d}t \right]$, where the initial *promised utility* $q_0 < 0$ is a given parameter. An *F-optimal (full-commitment) contract* is any contract that minimizes the principal's costs subject to these two constraints.[14]

Remark 2.1. The model described above is identical to the model introduced in PPI, except for two distinctions:

(i) We define the incentive compatibility and optimality of a contract as a function of the agent's feasible set $F \subseteq \mathscr{M}$. This allows us to be more explicit than PPI about which restrictions are being imposed on the agent's strategy space (see Section 2.2 below) and their implications for the class of IC contracts (Section 4).

(ii) We formulate the model directly over an infinite horizon. Although PPI initially formulates a finite-horizon version of the model (§2) and analyzes incentive compatibility in that context (§§3–4), the entire analysis of optimal contracts, both in an abstract setting (§5) and in applications (§§6–8), takes place in a informal infinite-horizon limit of that initial finite-horizon model (see Footnote 23 below). We proceed directly to the infinite-horizon formulation because (a) following PPI, it is the relevant setting for studying optimal contracts, and (b) PPI's infinite-horizon model as written is not fully specified because it relies simultaneously on a finite-horizon formulation of the agent's reporting problem and an infinite-horizon formulation of the principal's contracting problem (see Section 4.2 below).

[13] We let $\mathbf{E}_t^m$ denote the expectation with respect to the agent's time-t information under misreporting strategy $m \in \mathscr{M}$. We also let $\mathbf{E}_t^* := \mathbf{E}_t^{m^*}$ for brevity, as in [**2.2**] below.

[14] As usual, we also implicitly restrict attention to the class of contracts for which [**2.2**] is well-defined.

2.2. Restrictions on Reporting Strategies

We introduce various restrictions on the agent's feasible set F, some of which come from PPI and some of which are new. Doing so allows us to clarify the restrictions imposed in PPI, and to modify these restrictions after showing that some of them are problematic. Throughout, it will be useful to keep in mind the following elementary fact:

Fact 1. If the feasible sets $F, F' \subseteq \mathscr{M}$ satisfy $F \subseteq F'$, then:

(i) Every F'-IC contract is also F-IC, i.e., the set of F'-IC contracts is smaller than the set of F-IC contracts.

(ii) The optimal F'-IC contract has a weakly higher cost to the principal (i.e., is no better) than the optimal F-IC contract.

Sign Restrictions. A class of restrictions considered in PPI concerns the sign of the agent's misreports. One such restriction is that the agent can only *under-report his type*, so that his misreporting process is everywhere non-positive. For reasons discussed in Section 4.1 below, we dub this restriction *No Hidden Borrowing (NHB)* and denote the set of misreporting strategies consistent with NHB by

[**NHB**] $$\mathscr{M}_- := \{m \in \mathscr{M} : m \leqslant 0\}.$$

A stronger restriction is that the agent can only under-report the *increments of* his type, so that the Δ process is everywhere non-positive. Because this implies that the misreporting process has non-increasing sample paths, we refer to it as *Increasing Magnitude of Lies (IML)* and denote the set of misreporting strategies consistent with IML by

[**IML**] $$\mathscr{M}_\leqslant := \{m \in \mathscr{M} : \Delta \leqslant 0\}.$$

As explained below in Section 4.1, PPI motivates NHB as part of the model formulation but relies on IML for the formal analysis.

Absolute Continuity of Measures. Another class of restrictions considered in PPI concerns the kinds of misreporting strategies that are (un)detectable by the principal.

Recall that each $m \in \mathscr{M}$ formally induces a measure P^m over sample paths of y, with $\mathsf{P}^* := \mathsf{P}^{m^*}$ the measure induced by truthtelling. PPI assumes, as part of the finite-horizon model formulation (pp. 1240–42), that the agent can only misreport in ways that generate an *absolutely continuous (AC)* change-of-measure: given the finite horizon $[0, T]$ and letting P^m_T denote the marginal of P^m over $[0, T]$-truncated sample paths, PPI assumes that m is feasible only if P^m_T is AC with respect to P^*_T (denoted $\mathsf{P}^m_T \ll \mathsf{P}^*_T$).[15]

[15] To be precise, PPI assumes that $\mathsf{P}^m_T \ll \mathsf{P}^0_T$, where P^0 is the measure induced by the non-truthful strategy under which y is a driftless Brownian motion. For any finite T, this is equivalent to requiring

Economically, this restriction is meant to capture the fact that, in continuous time, there are certain kinds of misreports that the principal can instantaneously detect (with probability one), and therefore deter at zero expected cost by "shooting the agent" upon detection (cf. PPI p. 1240).[16] Technically, it facilitates PPI's use of Girsanov's Theorem to reformulate the agent's reporting problem as one of choosing a "density process" for the change-of-measure (§2.3).

With an infinite horizon, however, there are two standard but distinct notions of AC changes-of-measure. Specifically:

- P^m is *locally AC* with respect to P^* if the finite-horizon marginals satisfy $\mathsf{P}^m_t \ll \mathsf{P}^*_t$ for all $t > 0$. We let

  $$\text{[\textbf{LAC}]} \qquad \mathscr{M}^{\text{LAC}} := \{m \in \mathscr{M} : \mathsf{P}^m_t \ll \mathsf{P}^*_t \text{ for all } t > 0\}$$

  denote the class of strategies inducing locally AC changes-of-measure. Intuitively, no $m \in \mathscr{M}^{\text{LAC}}$ is detectable by the principal *in finite time*.
- P^m is *globally AC* with respect to P^* if $\mathsf{P}^m \ll \mathsf{P}^*$, i.e., the measures over entire infinite-horizon paths are AC. We let

  $$\text{[\textbf{GAC}]} \qquad \mathscr{M}^{\text{GAC}} := \{m \in \mathscr{M} : \mathsf{P}^m \ll \mathsf{P}^*\}$$

  denote the class of strategies inducing globally AC changes-of-measure. Intuitively, no $m \in \mathscr{M}^{\text{GAC}}$ is detectable even *in infinite time* (i.e., "at $t = \infty$", after the entire sample path has been observed).

It is well known that GAC is strictly more demanding than LAC because the former imposes fairly strong restrictions on the asymptotic behavior of the m process as $t \to \infty$ (see Appendix C). As we discuss in Section 4.2, PPI does not explicitly specify which notion is meant to be imposed in the infinite-horizon model.

No Ponzi Constraints. Our analysis will illustrate the importance of imposing tail restrictions that constrain the asymptotic growth rate of the agent's misreports (see

that $\mathsf{P}^m_T \ll \mathsf{P}^*_T$. However, in the infinite-horizon model, requiring that $\mathsf{P}^m \ll \mathsf{P}^0$ would imply that truthful reporting is infeasible, which would be inappropriate (cf. [**GAC**] below).

[16] The AC change-of-measure assumption is stronger than this intuition suggests. Formally, suppose the principal fixes a P^*-null event $N \subset \mathbf{C}[0,\infty)$ of "detectable lies" (e.g., sample paths that have different quadratic variation than b) and "shoots" the agent if and only if he reports a path in N. While this constrains the agent to strategies m for which $\mathsf{P}^m(N) = 0$, he could still set $\mathsf{P}^m(N') > 0$ for a *different* P^*-null set N'. For instance, he could pick any path $\hat{y} \in \mathbf{C}[0,\infty)\backslash N$ and always report it; the principal, who only observes *one* path of reports, would be unable to distinguish this strategy from truthtelling conditional on $\hat{y}$ being the *realized* path of b. Nonetheless, the AC change-of-measure assumption rules out such strategies by preventing the agent from placing positive weight on *any* P^*-null event. Which misreports should be deemed detectable is a subtle open question for the continuous-time contracting literature (cf. Acciaio, Crowell, and Cvitanić (2022)).

Section 4.2). However, we find it convenient to impose tail restrictions that are weaker—and admit an arguably more natural economic interpretation—than those implied by GAC, and that, unlike GAC, can be formulated as pathwise constraints. To this end, we say that misreporting strategy m satisfies the *no Ponzi condition (at rate $r > 0$)* if

$$\text{[NP-}m\text{]} \qquad \lim_{t\to\infty} e^{-rt} \int_0^t m_\tau \mathrm{d}\tau \geqslant 0 \qquad \mathbf{P}\text{-a.s.}$$

Notice that [**NP-*m***] constrains the asymptotic growth rate of the agent's *under-reports* only by ruling out sample paths along which m_t diverges to $-\infty$ too quickly relative to the rate r. As formalized in Section 6, [**NP-*m***] is analogous to the standard no Ponzi constraint $\lim_{t\to\infty} e^{-rt} A_t \geqslant 0$ that arises in a consumption-saving problem with "interest rate" $r > 0$ and "asset process" $A_t \equiv \int_0^t m_\tau \mathrm{d}\tau$. We denote the set of misreporting strategies consistent with this no Ponzi constraint by

$$\text{[2.3]} \qquad \mathscr{M}^r := \{m \in \mathscr{M} : \text{[NP-}m\text{] holds for rate } r\}\,.$$

Clearly, [**NP-*m***] is strictly more permissive for higher rates: $\mathscr{M}^r \subsetneq \mathscr{M}^{r'}$ whenever $r < r'$. Moreover, GAC implies that [**NP-*m***] holds for all rates $r > 0$ (see Appendix C).

Notation. We define $\mathscr{M}^r_- := \mathscr{M}^r \cap \mathscr{M}_-$ and $\mathscr{M}^r_\leqslant := \mathscr{M}^r \cap \mathscr{M}_\leqslant$, with a similar convention for $\mathscr{M}^{\text{LAC}}_-$, $\mathscr{M}^{\text{LAC}}_\leqslant$, and so on. We also adopt the convention that "NHB" stands for "the assumption that the agent's feasible set is $F = \mathscr{M}_-$," and so on.

3. Hidden Endowment Application

In this section, we recall the setting of PPI's main application and the contract derived as optimal therein (§§6–7 and §A.3).

Setting. The agent's time-t type b_t now corresponds to his endowment at that time. The agent's utility of consumption at time t is given by $v(s_t, b_t) = u(s_t + b_t)$, where u takes the CARA form $u(c) = -e^{-\theta c}$ for some $\theta > 0$. Given this structure, it is convenient to re-express contractual variables as follows. Define the agent's y-adapted *(recommended) consumption process* $c = (c_t)_{t\geqslant 0}$ by $c_t := s_t + y_t$ and his y-adapted *(recommended) flow utility process* $u = (u_t)_{t\geqslant 0}$ by $u_t := u(c_t)$. These are the consumption and flow utility processes intended by the contract, presuming that the agent is truthful. Thus, the agent's *actual* consumption is given by the b-adapted process $c^m = (c^m_t)_{t\geqslant 0}$ where $c^m_t := c_t - m_t = s_t + b_t$. In this context, we can interpret m_t as the amount that the agent "diverts" for private consumption.

Contract PPI. To describe the contract identified as optimal in PPI—hereafter *Contract PPI*—we introduce three processes. First, define the y-adapted process $W^y = (W^y_t)_{t\geqslant 0}$ by $\sigma W^y_t := y_t - b_0 - \int_0^t (\mu - \lambda y_\tau)\,\mathrm{d}\tau$. That is, W^y is the shock process that the principal would infer the agent faced if (a) the principal were to assume that the agent is truthful and (b) the agent actually follows strategy y. Note that W^y coincides with the standard Brownian motion W when the agent follows the truthful strategy y^* and is a Brownian motion with drift more generally (viz., whenever $m \in \mathscr{M}^{\mathrm{LAC}}$). Second, define the y-adapted *promised utility* process $q = (q_t)_{t \geqslant 0}$ by

$$q_t := \mathbf{E}_t^* \left[\int_t^\infty e^{-\rho(\tau - t)} u_\tau \, \mathrm{d}\tau \right]. \tag{3.1}$$

Thus, q_t is the agent's lifetime utility from time t onward *under truthful reporting*. Third, define the y-adapted *(negative) marginal promised utility* process $p = (p_t)_{t \geqslant 0}$ by

$$p_t := \mathbf{E}_t^* \left[\int_t^\infty e^{-(\rho+\lambda)(\tau - t)} \theta u_\tau \, \mathrm{d}\tau \right]. \tag{3.2}$$

Recall that $\theta u_\tau \equiv -u'(c_\tau)$ for CARA utility. Therefore, p_t represents the agent's "marginal incentives" for misreporting by a small amount at time t *conditional on having reported truthfully at all dates* $\tau < t$. In particular, [3.2] is the continuous-time version of the dynamic envelope formula (Pavan, Segal, and Toikka (2014, Theorem 1)) and we can informally view p_t as the derivative $\mathrm{d}q_t/\mathrm{d}b_t$, which is a local measure of the agent's on-path information rents. PPI argues that IC contracts can be written recursively with q_t and p_t as state variables (pp. 1244–46).

Definition 3.1 (Contract PPI)**.** The contract identified as being (uniquely) optimal in PPI (henceforth Contract PPI) is that under which promised utility satisfies

$$q_t = q_0 \exp\left(-\tfrac{1}{2}(k_0^*\sigma)^2 t - k_0^*\sigma W_t^y\right), \tag{3.3}$$

marginal promised utility satisfies $p_t = k_0^* q_t$, and recommended consumption is

$$c_t = \bar{c}(q_0, \rho) + \frac{(k_0^*\sigma)^2}{2\theta} t + \frac{k_0^*\sigma}{\theta} W_t^y, \tag{3.4}$$

where $k_0^* := \rho\theta/(\rho+\lambda)$ and $\bar{c}(q, r) := -\log(-rq)/\theta$.

Contract PPI has several striking features that will be important going forward.

Fact 2. Under truthful reporting, Contract PPI satisfies the following properties:

(i) It generates long-run *bliss*: As $t \to \infty$, we have $q_t, u_t \to 0$ and $c_t \to +\infty$ almost surely, i.e., the agent receives maximal utility and consumption in the long-run.[17]

[17] The fact that $c_t \to +\infty$ a.s. follows from the Strong Law of Large Numbers for Brownian Motion applied to [3.4] (with truthful reporting, $y = y^*$), and implies the asserted a.s. convergence results for u_t and q_t.

(ii) It is *Markovian in promised utility*: The continuation contract is a function of q_t alone because $p_t \equiv k_0^* q_t$ and, substituting [3.3] into [3.4], we have $c_t \equiv \bar{c}(q_t, \rho)$.
(iii) Promised utility is a *martingale*: Applying Itô's lemma to [3.3] yields $\mathrm{d}q_t = -\sigma k_0^* q_t \,\mathrm{d}W_t^y$.
(iv) It has a constant *utility delivery rate* of ρ: Formally, $u_t \equiv \rho q_t$, which means the principal delivers promised utility at a constant rate ρ.

PPI emphasizes parts (i)–(iii) of Fact 2 (pp. 1253–54) and notes property (iv) (p. 1253, second-to-last display). Each property in Fact 2 stands in contrast to the literature's findings in closely related settings:

(a) In contrast to Fact 2(i), the literature has found that, *when private information is either (i) i.i.d. or (ii) persistent and mean-reverting, optimal insurance contracts generate immiseration* (which in the current notation means $q_t, u_t, c_t \to -\infty$). See Section 7.2 below for references and further discussion.
(b) In contrast to Fact 2(ii), the literature has found that, *when private information is persistent, optimal contracts typically cannot be written recursively in promised utility alone*. See Section 7.1 below for references and further discussion.
(c) Points (iii) and (iv) of Fact 2, together with the property of CARA utility that $u'(c) = -\theta u(c)$, imply that the agent's marginal utility of consumption $u'(c_t)$ is a martingale, so that his consumption obeys the Euler equation familiar from consumption-saving problems. In contrast, the literature has found that *optimal contracts typically do not induce the agent's Euler equation, unless the agent has access to hidden savings outside of the contract*. See Section 5 below for references and further discussion.

PPI emphasizes feature (a) above (pp. 1235, 1257, 1264) but does not note either (b) or (c).[18] We will see in Section 4.3 that property (c) is key to the observation that Contract PPI is (generically) suboptimal.

4. Issues in PPI

This section formally describes the three issues in PPI's model formulation and analysis (with some details in Appendices A and D). Section 4.1 describes how PPI introduces the NHB assumption that $m \leqslant 0$ as part of the model formulation, but then bases the formal analysis on the strictly stronger IML assumption that $\Delta \leqslant 0$ (Observation 1). Section 4.2 describes how PPI omits tail restrictions (such as GAC or NP-m) on the agent's feasible set of mispeporting strategies, so that, among other things, Contract PPI

[18] PPI does emphasize that the *inverse* Euler equation—i.e., the property that $1/u'(c_t)$ is a martingale, which arises as an optimality condition for the principal in many dynamic Mirrleesian models—is *not* satisfied in either of PPI's applications (pp. 1235-36, 1257-58, 1264). However, this is not relevant for understanding PPI's results. It is well known that the inverse Euler equation holds when the agent's utility is additively separable across his consumption and private type, in both discrete and continuous time (Golosov, Kocherlakota, and Tsyvinski 2003; Zhang 2009; Farhi and Werning 2013; Kapička 2013). Conversely, it typically fails in settings without such separability, including the discrete-time models on which PPI's applications are based (Thomas and Worrall 1990; Atkeson and Lucas 1992).

	Sign Restriction		
Tail restriction	IML	NHB	None
None / LAC	Not IC	Not IC	Not IC
NP-*m* ($r < \rho$)	IC, suboptimal	IC, suboptimal	IC, suboptimal
GAC	IC, suboptimal (only $m^* \equiv 0$ feasible)	IC, suboptimal	IC, suboptimal

Table 1: Properties of Contract PPI under transient shocks ($\lambda > 0$) and different restrictions on reporting strategies.

is generically not IC under PPI's assumptions (Observation 2). Perhaps most importantly, Section 4.3 shows that Contract PPI is generically strictly suboptimal, even after the first two issues have been addressed (Observation 3).

Table 1 summarizes the implications of these issues for PPI's hidden endowment application and Contract PPI, which motivate our analysis in the remainder of the paper. The first two issues also have implications for PPI's abstract analysis of IC and optimal contracts, which we discuss in Sections 4.1 and 4.2 for completeness. Readers interested primarily in the suboptimality of Contract PPI—which we view as the most important and economically substantive issue—and the economic implications thereof may proceed directly to Section 4.3 with little loss of continuity.

4.1. Sign Restrictions on Misreporting Strategies

As part of the model formulation, PPI initially imposes the NHB assumption that only $m \in \mathscr{M}_-$ are feasible for the agent, writing (p. 1239): "To simplify matters, I assume that the agent cannot overreport the true state, so $y_t \leqslant b_t$ [i.e., $m_t \leqslant 0$]." This restriction is natural and common in the literature; as noted in PPI (p. 1239), it captures the ideas that (i) the agent cannot borrow or save outside of the contract and (ii) the agent's endowment is partially verifiable (e.g., the principal can require that he deposit the reported amount in a joint account). However, PPI's formal analysis is based on the stronger IML assumption that only $m \in \mathscr{M}_\leqslant$ are feasible for the agent. PPI writes (on p. 1240): "Since the agent can report (or deposit) at most his entire state [i.e., $m_t \leqslant 0$], we must have $\Delta_t \leqslant 0$." Unfortunately, this assertion is incorrect: *NHB ($m \leqslant 0$) is a strictly weaker assumption than IML ($\Delta \leqslant 0$).* There exist many non-positive functions $t \mapsto m_t$ with locally strictly positive derivatives $\Delta_t = \mathrm{d}m_t/\mathrm{d}t > 0$.[19] Formally, the set of

[19] NHB requires that $\Delta_t \leqslant 0$ if and only if $m_t = 0$ ("on path") and allows for any $\Delta_t \in \mathbb{R}$ when $m_t < 0$ ("off path"). Two points warrant clarification. First, PPI defines m_t as the "stock of lies" (p. 1240).

strategies satisfying IML is strictly smaller than the set satisfying NHB, regardless of the tail restrictions that one imposes.

Observation 1. $\mathcal{M}^{\dagger}_{\leqslant} \subsetneq \mathcal{M}^{\dagger}_{-}$ *for all* $\dagger \in \mathbb{R}_{++} \cup \{LAC, GAC\}$.

Proof. The weak inclusions are trivial, so it suffices to find an $m \in \mathcal{M}^{\mathrm{GAC}}_{-} \backslash \mathcal{M}_{\leqslant}$ (GAC is the strongest tail restriction). Let $m_t := -t$ for $t \in [0, 1)$, $m_t := t - 2$ for $t \in [1, 2)$, and $m_t := 0$ for $t \geqslant 2$. Then $m \in \mathcal{M}_{-} \backslash \mathcal{M}_{\leqslant}$ because $m_t \leqslant 0$ everywhere and $\Delta_t = 1$ for $t \in [1, 2)$, and $m \in \mathcal{M}^{\mathrm{GAC}}$ because Δ is bounded and $m_t \equiv 0$ for $t \geqslant 2$ (cf. Appendix C). □

PPI's reliance on IML is notable for two reasons:[20]

(i) *If PPI intended to assume only NHB, then the analysis pertaining to the verification of IC is incomplete because it relies on IML.* In PPI's general model, the distinction between NHB and IML implicates the sufficient conditions for IC in PPI's Theorem 4.1, the proof of which (in §A.2) would not go through as stated without IML (see Appendix D). In the hidden endowment application, this distinction implicates the attempted verification (in §A.3.2) that Contract PPI is IC in the generic $\lambda > 0$ case, which is not sufficient to show that Contract PPI would be IC under NHB. This attempted verification also involves the derivation of a value function for the agent (restated as V^W in [**4.1**] below) that (a) diverges to $-\infty$ as the current misreport $m_t < 0$ approaches a *finite* value $\underline{M} < 0$ and (b) is different from the agent's value function under NHB, which is everywhere finite (see Remark B.2 in Appendix B). This indicates that IML is a much more stringent constraint on the agent than NHB.[21]

(ii) *If PPI intended to adopt IML as a separate assumption, no economic motivation for it is given. To the best of our knowledge, it has no analogue in the literature.* IML implies that once the agent under-reports his endowment, he can never revert to truthtelling: $m_t < 0$ implies that $m_\tau \leqslant m_t < 0$ for all $\tau \geqslant t$. Unlike NHB, this stronger property does not follow from the principal's ability to partially verify the agent's endowment

However, m_t is a *flow* variable with the same units as b_t. Second, PPI refers to $\Delta_t = 0$ as "truth-telling" (pp. 1267 and 1272) or corresponding to a "truthful current report" (p. 1244) *even at histories where* $m_t < 0$. This terminology blurs the distinction between truthful reporting ($m_t = 0$) and truthful reporting *of increments* ($\Delta_t = 0$), potentially generating confusion about the definition of IC and its implications for the agent's behavior at off-path histories (which we describe in Section 6.4 below).

[20]To our knowledge, references to PPI in the literature only mention the weaker NHB assumption. For instance, Kapička (2013, p. 1029) writes "[PPI] assumes that the agent cannot overstate her true shock [i.e., endowment]," and Battaglini and Lamba (2019, p. 1459, fn. 36) write "[PPI] limits the set of possible deviations available to the agent (who can report only incomes lower or equal to the true income)."

[21]Intuitively, when $m_t < 0$ the principal expects stronger positive mean-reversion than actually occurs, and so punishes the agent for not reporting increments $\mathrm{d}y_t > \mathrm{d}b_t$. Meanwhile, the IML constraint forces the agent to report increments $\mathrm{d}y_t \leqslant \mathrm{d}b_t$, so that he cannot avoid such punishments. As $m \searrow \underline{M}$, the punishments become so severe that the agent's continuation value decreases without bound. As we show in Section 6.4, if the agent were not constrained by IML, he would find it optimal to "immediately" correct for all past under-reports by "jumping" back to $m_t = 0$.

reports. By construction, it also rules out the possibility of "one-shot" deviations, which are central to virtually all dynamic analyses of incentive compatibility, making PPI's model and analysis incomparable to those in the literature.

Our Approach. *Given that IML is used throughout PPI's formal analysis, we henceforth adopt the perspective that IML was PPI's intended assumption.* Even under this stronger assumption, PPI's main analysis of infinite-horizon IC omits necessary tail restrictions on the agent's feasible set and Contract PPI is generically suboptimal (Observation 2 and 3 below). However, *our main analysis (Sections 5–7) imposes neither IML nor NHB*, allowing us to (a) bypass what may be viewed as an economically problematic assumption and (b) demonstrate that Contract PPI remains generically suboptimal under *any* sign restrictions that one might wish to adopt (per Fact 1). To analyze IC without IML, we employ techniques, explained in Section 6, that may prove relevant in other continuous-time contracting models.

4.2. Tail Restrictions on Misreporting Strategies

PPI's definition and analysis of IC contracts proceeds in two steps. First, PPI formulates the model over a finite time horizon $[0, T]$ (§2) and then provides necessary conditions (§3) and sufficient conditions (§4) for finite-horizon IC. Second, PPI formulates the principal's problem of finding *optimal* contracts in an infinite-horizon version of the model, first in an abstract setting (§5) and then in two solved applications (§§6–8).

There are two standard approaches for carrying out the second step, each of which is common in the literature and has distinct advantages:[22]

(i) Use the definition of finite-horizon IC from the first step to solve for optimal contracts over each finite horizon $[0, T]$, and then study the limit contract as $T \to \infty$.

(ii) Define and characterize IC contracts directly over the infinite horizon $[0, \infty)$, and then solve for the optimal such contract.

PPI adopts neither of these approaches. PPI defines IC contracts over the finite horizon $[0, T]$, informally motivates the infinite-horizon model as a limit of the finite-horizon model as $T \to \infty$, and then directly defines the principal's optimization problem, solves for optimal contracts, and verifies that they are "IC" in the infinite-horizon model *without defining the agent's feasible set $F \subseteq \mathcal{M}$ of misreporting strategies in that context or checking transversality conditions on the agent's value function.*[23] We describe these

[22]Approach (i) obviates the need to distinguish between LAC and GAC, but can be difficult to interpret because the limit contract may be neither IC nor optimal in the appropriately defined infinite-horizon model (see, e.g., Prat and Jovanovic (2014, Section 4)). Approach (ii) is typically carried out by imposing GAC (as in Sannikov (2014); DeMarzo and Sannikov (2016); Chen (2021)), which might be viewed as an overly stringent assumption.

[23]PPI describes the infinite-horizon setting as follows (p. 1248): "I now turn to the principal's problem of optimal contract design over an infinite horizon. Formally, I take limits as $T \to \infty$ in the analysis

omissions and their implications below.

Definition of Agent's Strategy Space. As noted in Section 2.2, PPI's formulation of the *finite*-horizon model explicitly assumes that the agent is restricted to strategies inducing AC changes-of-measure over the finite horizon. In the *infinite*-horizon context, PPI *never specifies* analogous restrictions, such as LAC or GAC. It can be shown that GAC is not a viable assumption in conjunction with IML because, in the generic case of transient shocks ($\lambda > 0$), truthtelling is the only feasible strategy satisfying both conditions.[24] *We therefore explore the consequences of LAC combined with IML, which seems to provide the infinite-horizon version of PPI's model that is most faithful to its finite-horizon version, while still providing a fully formulated contracting problem. For brevity, we henceforth refer to this combination of restrictions as "PPI's assumptions."*

Transversality in Agent's Problem. PPI's main analysis of infinite-horizon IC (pp. 1271–72 in §A.3.2) pertains to Contract PPI under transient shocks ($\lambda > 0$). Therein, PPI formulates the agent's reporting problem as a stochastic control problem in which Δ_t is a control variable and the current promised utility q_t and misreport m_t serve as state variables, so that the agent's value function at time t can be written as a function of (q_t, m_t).[25] PPI first conjectures that the agent's value function takes the form

$$V^W(q,m) := \begin{cases} q\exp(\theta m)\cdot\left(\frac{\rho+\lambda}{\rho+\lambda+\theta\lambda m}\right) & \text{for } m \in (\underline{M}, 0] \\ -\infty & \text{for } m \leqslant \underline{M}, \end{cases} \tag{4.1}$$

[of finite-horizon IC contracts] above. Thus we no longer have the terminal conditions for the co-states [promised utility q_T and marginal promised utility p_T] in (11) and (12) [on p. 1244 of PPI]; instead we have the transversality conditions $\lim_{T\to\infty} e^{-\rho T} q_T = 0$ and $\lim_{T\to\infty} e^{-\rho T} p_T = 0$." These transversality (or terminal) conditions should be understood to hold only under truthful reporting (i.e., $\mathbf{P}^*$-a.s. but *not* necessarily $\mathbf{P}^m$-a.s. for other $m \in F$) for two reasons. First, the terminal conditions of the finite-horizon model analyzed in earlier sections of PPI apply only under truthful reporting. Second, more generally, the first-order approach—which PPI derives via the Maximum Principle and other papers derive via the Envelope Theorem (e.g., Pavan, Segal, and Toikka 2014)—only delivers necessary conditions for IC that hold "on path" under truthful reporting. Thus, PPI's transversality conditions make no reference to the agent's feasible set and are not sufficient conditions for IC. To verify that a contract is IC, one should instead use a transversality condition that holds under all feasible strategies to ensure that the agent's value function is sufficiently "continuous at infinity" that he cannot benefit from infinite-length deviations (cf. the proof of Theorem 3 in Pavan, Segal, and Toikka 2014).

[24] See Fact 3(i) in Appendix C, which implies that all contracts are $\mathscr{M}^{\mathrm{GAC}}_{\leqslant}$-IC. Meanwhile, the set $\mathscr{M}^{\mathrm{GAC}}_{\leqslant}$ does contain nontrivial strategies when $\lambda = 0$, and the sets $\mathscr{M}^{\mathrm{GAC}}_{-}$ and $\mathscr{M}^{\mathrm{GAC}}$ contain many nontrivial strategies for all values of $\lambda \geqslant 0$. Thus, GAC is problematic only in conjunction with IML and $\lambda > 0$.

[25] PPI also provides a separate argument based on an informal infinite-horizon adaptation of PPI's Theorem 4.1. That argument is also incorrect for reasons similar to those described here (see Appendix D).

where $\underline{M} := -(\rho+\lambda)/(\theta\lambda)$.[26] This is the agent's lifetime utility from always setting $\Delta = 0$ irrespective of the current misreport m. PPI then shows that V^W satisfies a suitable HJB equation (p. 1272) and concludes from this that (a) V^W is the agent's value function given the feasible set $F = \mathcal{M}_{\leqslant}^{\mathrm{LAC}}$ and (b) Contract PPI is $\mathcal{M}_{\leqslant}^{\mathrm{LAC}}$-IC.

Conclusions (a) and (b) are both unwarranted because the agent's HJB equation may have multiple solutions. Standard "verification theorems" in stochastic control require verifying that V^W satisfies the transversality condition

$$\text{[TVC]} \qquad \lim_{t\to\infty} \mathbf{E}^m \left[e^{-\rho t} V^W(q_t, m_t)\right] = 0 \ \text{ for all } m \in F$$

before concluding that it is the agent's value function, rather than some other solution to the HJB equation.[27] PPI does not carry out this verification step. In fact, it follows from [**4.1**] that V^W *violates* [**TVC**] unless the feasible set F prevents the agent's misreports m_t from approaching $\underline{M}$ too quickly. For instance, $\lim_{t\to\infty} \mathbf{E}^m \left[e^{-\rho t} V^W(q_t, m_t)\right] = -\infty$ under any misreporting strategy that crosses $\underline{M}$ (and thereafter stays below) in finite time. As there are many such strategies in $\mathcal{M}_{\leqslant}^{\mathrm{LAC}}$, one cannot draw either conclusion (a) or (b) under PPI's assumptions using the standard guess-and-verify approach.[28]

Consequences of Omissions. Unfortunately, conclusions (a) and (b) above are also both incorrect. Formally, we show that the agent's true value function is identically zero under PPI's assumptions, meaning that the agent can achieve near-infinite consumption by deviating from truthtelling. Moreover, Contract PPI is not IC if the agent's misreports are permitted to violate NP-m with rate ρ (i.e., if the cumulative misreports $\int_0^T m_\tau \mathrm{d}\tau$ can diverge to $-\infty$ at an exponential rate faster than ρ).

Observation 2. *If $\lambda > 0$, then Contract PPI satisfies the following properties:*

(i) *The agent's value function under feasible set $F = \mathcal{M}_{\leqslant}^{LAC}$ is identically zero.*

(ii) *It is not $\left[\mathcal{M}_{\leqslant}^{LAC} \cap \mathcal{M}^r\right]$-IC for any $r > \rho$.*

Proof. Given any $\varepsilon > 0$ and $\kappa > 0$, define the "Ponzi scheme" misreporting strategy

---

[26] PPI does not mention the constant $\underline{M}$ or specify that $V^W(q,m) = -\infty$ when $m \leqslant \underline{M}$, apparently intending for $V^W(q,m)$ to be defined as in the first line of [**4.1**] for all $m \leqslant 0$. However, the expression in the first line of [**4.1**] is strictly positive whenever $m < \underline{M}$, making it impossible for V^W so-defined to be the agent's lifetime utility under any strategy (as the agent's CARA utility function is strictly negative).

[27] For instance, see Pham (2009, Theorem 3.5.3). The importance of verifying transversality conditions in problems with unbounded returns is also familiar from discrete-time dynamic programming (e.g., Stokey, Lucas, and Prescott 1989). The condition [**TVC**] differs from PPI's transversality condition $\lim_{T\to\infty} e^{-\rho T} q_T = 0$ $\mathbf{P}^*$-a.s. (recall Footnote 23) in two ways. First, [**TVC**] concerns the conjectured value function process V^W rather than the promised utility process q, which do not coincide off-path. Second, [**TVC**] concerns all feasible strategies, rather than just truthful reporting.

[28] The condition [**TVC**] is sufficient but *not* necessary for V^W to be the true value function. By adapting arguments from Appendix I, it can be shown that V^W is, in fact, the agent's true value function under IML and the no Ponzi constraint [**NP-*m***] with $r = \rho$, which allows for many strategies that violate [**TVC**].

$m^{(\varepsilon,\kappa)}$ as follows:

$$m_t^{(\varepsilon,\kappa)} := \begin{cases} -t/\varepsilon & \text{for } t \in [0,\varepsilon], \\ -e^{\kappa(t-\varepsilon)} & \text{for } t > \varepsilon. \end{cases}$$

Under this strategy, the agent under-reports forever at an exponentially growing rate of κ (after initializing his misreport at time $t = \varepsilon$ to $m_\varepsilon = -1$). This strategy is plainly in $\mathscr{M}^r_\leqslant$ for all $r > \kappa$ and, being deterministic and bounded on each finite time horizon $[0,T]$, is also in $\mathscr{M}^{\text{LAC}}$ (cf. Appendix C). Next, notice that [3.4] allows us to write the agent's actual consumption process $c_t^m \equiv c_t - m_t$ under Contract PPI and any strategy $m \in \mathscr{M}$ as

$$c_t^m = c_t^{m*} + \frac{\lambda}{\lambda+\rho} \underbrace{\left(\rho \int_0^t m_\tau d\tau - m_t\right)}_{=: \, \xi_t^m}$$

where c^{m*} is the agent's consumption process under truthful reporting and the ξ^m process is proportional to the agent's "extra consumption" from misreporting. For the strategy $m^{(\varepsilon,\kappa)}$, a simple calculation yields

$$\xi_t^{m^{(\varepsilon,\kappa)}} = \begin{cases} \frac{t}{\varepsilon}\left(1 - \frac{\rho}{2}t\right) & \text{for } t \in [0,\varepsilon), \\ e^{\kappa(t-\varepsilon)}\left(1-\frac{\rho}{\kappa}\right) + \rho\left(\frac{1}{\kappa} - \frac{\varepsilon}{2}\right) & \text{for } t \in [\varepsilon,\infty). \end{cases}$$

It is easy to see that $\xi_t^{m^{(\varepsilon,\kappa)}} > 0$ for all $t > 0$ whenever $\kappa > \rho$ and $\varepsilon \in (0, \overline{\varepsilon}(\rho,\kappa))$, where we let $\overline{\varepsilon}(\rho,\kappa) := \min\{2/\rho, 2/\kappa\}$.

We may now prove Observation 2. Let $\lambda, \rho > 0$ and $r > \rho$ be given. For any $\kappa \in (\rho, r)$ and $\varepsilon \in (0, \overline{\varepsilon}(\rho,\kappa))$, it follows from the above that the agent derives strictly greater consumption at all times $t > 0$ from the strategy $m^{(\varepsilon,\kappa)}$ than from truthtelling. This establishes part (ii). As for part (i), consider the sequence of strategies $m^{(1/n,n)}$, so that $\kappa_n := n \to +\infty$ and $\varepsilon_n := 1/n$ satisfies $\varepsilon_n < \overline{\varepsilon}(\rho,\kappa_n)$ for all $n > \rho/2$. By construction, $\xi_t^{m^{(1/n,n)}} \to +\infty$ for each fixed $t > 0$. This implies that the agent's actual consumption $c_t^{m^{(1/n,n)}} \to +\infty$ for each fixed $t > 0$. It is then easy to see that the agent's lifetime utility under $m^{(1/n,n)}$ converges to its upper bound of zero as $n \to \infty$. As each such strategy is in $\mathscr{M}^{\text{LAC}}_\leqslant$, this establishes part (i) and thereby completes the proof. □

Our Approach & the Necessity of Tail Restrictions. Observation 2(ii) implies that the no Ponzi constraint [**NP-*m***] with some $r \leqslant \rho$ is a *(nearly) necessary* condition for Contract PPI to be IC. More generally, our analysis suggests that the no Ponzi constraint is important for PPI's model to be well-behaved, and we conjecture that it may be necessary for PPI's model to admit IC contracts with nontrivial risk-sharing.[29] *For these*

[29] For instance, we show in Appendix H that *without* this constraint, Contract PPI *cannot* be implemented as a direct mechanism but *can* be indirectly implemented in an alternative model where the agent has

reasons, we henceforth focus on F-IC contracts for feasible sets F satisfying [**NP-*m***].

4.3. Strict Suboptimality of Contract PPI

PPI states (p. 1254 in §6) that Contract PPI is optimal in the hidden endowment model. This statement is problematic for two reasons. First, Observation 2 demonstrates that Contract PPI is not IC under PPI's assumptions. Even if this could be addressed by imposing sufficiently tight tail restrictions on the agent's strategy space, it is *a priori* unclear whether Contract PPI would be IC if IML were relaxed to NHB (or if sign restrictions were dropped altogether). Second, and more importantly, we show that—independently of these issues—Contract PPI is strictly suboptimal whenever the agent's endowment process has non-zero mean reversion. (We describe PPI's derivation of Contract PPI and the source of this discrepancy in Appendix A.) Formally:

Observation 3. *If $\lambda > 0$, then Contract PPI satisfies the following properties:*[30]

(i) *It is $\mathcal{M}^\rho$-IC.*

(ii) *Given any $r \in (0, \rho)$, it is $\mathcal{M}^r$-IC but strictly suboptimal among $\mathcal{M}^r$-IC contracts.*

Per Fact 1, the same is true if either NHB or IML is also imposed.

Proof. This follows from the analysis in Sections 5 and 6 below. Specifically, part (i) is an immediate corollary of Theorem 2 in Section 6, while part (ii) is an immediate corollary of that result in conjunction with Theorem 1 in Section 5. □

Observation 3 has three implications. First, in conjunction with Observation 2, it demonstrates that restricting the agent's strategy space to $\mathcal{M}^\rho$ is a necessary *and sufficient* condition for Contract PPI to be IC. Second, it establishes that Contract PPI is strictly suboptimal under the slightly stronger assumption that the agent's strategy space is $\mathcal{M}^r$ for some $r \in (0, \rho)$. Third, a corollary of these observations and Fact 1 is that the incentive compatibility and strict suboptimality of Contract PPI are both *independent of the sign restrictions* that one imposes on the agent's feasible set.

access to hidden savings (see Section 5.4), which is unreasonable because the agent has access to more deviations in the latter model. Subsequent to our working paper, Acciaio, Crowell, and Cvitanić (2022) show that, without any tail restrictions on the agent's misreports, deterministic contracts—which do not condition on the agent's reports, and hence provide no insurance—are optimal within a class of "linear contracts" that includes the SI Contracts that we study in Sections 5 and 6. (Acciaio, Crowell, and Cvitanić (2022) also verify that, under suitable tail restrictions, the optimal SI Contract that we identify in Theorem 1 is optimal among all linear contracts.) We view this as evidence that, for technical reasons, tail restrictions are needed for PPI's model to deliver a non-degenerate contracting problem.

[30] We state Observation 3 in this two-part manner to emphasize the weakest technical conditions under which we are able to (separately) establish the incentive compatibility and suboptimality of Contract PPI. The stronger, and standard, GAC assumption is sufficient for both parts of Observation 3.

Our Approach & Self-Insurance. The formal analysis underlying Observation 3 is somewhat involved, but the basic idea is simple. Under truthful reporting, the consumption process induced by Contract PPI satisfies the agent's *Euler equation*: plugging [**3.4**] into the agent's marginal utility $u'(c) = -\theta u(c)$ yields

$$u'(c_t) = \mathbf{E}_t\left[u'(c_\tau)\right] \quad \text{for all } \tau \geqslant t, \tag{4.2}$$

so the agent's marginal utility is a martingale under Contract PPI. This familiar equation represents the agent's optimal consumption-saving behavior in a setting where (a) there is no principal and (b) the agent self-insures by investing in a risk-free bond with interest rate $r = \rho$. The prior literature shows that optimal contracts in discrete-time analogues of PPI's model typically do not coincide with solutions to self-insurance problems (Thomas and Worrall 1990), except in an *alternative model* in which the agent also has access to hidden savings (Allen 1985; Cole and Kocherlakota 2001). The crux of our approach is to argue that this classic finding remains true in PPI's model.

5. A Self-Insurance Approach

Section 5.1 presents a standard self-insurance problem for the agent. Section 5.2 uses this problem to construct an indirect implementation for our main class of "self-insurance contracts." Section 5.3 shows that this class contains Contract PPI and characterizes the optimal self-insurance contract. Section 5.4 summarizes additional results.

5.1. Agent's Self-Insurance Problem

Consider the classic *self-insurance problem* faced by the agent with CARA utility $u(c) = -e^{-\theta c}$ when there is no principal to provide insurance. In this problem, the agent receives only (i) his endowment stream b and (ii) some initial asset holdings $A_0 \in \mathbb{R}$, and must self-insure by borrowing and saving in a risk-free bond market at the given interest rate $r > 0$. Formally, the agent solves[31]

$$V^{\mathrm{SI}}(A_0, b_0) := \sup_{\hat{c} \in \mathscr{A}(A_0, b_0)} \mathbf{E}_0\left[\int_0^\infty e^{-\rho t} u(\hat{c}_t)\,\mathrm{d}t\right] \tag{5.1}$$

where $\mathscr{A}(A_0, b_0)$ is the set of (A_0, b_0)-*feasible* consumption strategies $\hat{c} = (\hat{c}_t)_{t\geqslant 0}$, which consists of all consumption processes that are b-adapted and induce an *asset process* $A^{\hat{c}} = \left(A_t^{\hat{c}}\right)_{t\geqslant 0}$ that solves

$$\mathrm{d}A_t^{\hat{c}} = (rA_t^{\hat{c}} + b_t - \hat{c}_t)\,\mathrm{d}t \tag{5.2}$$

[31] Recall that $\mathbf{P}$ denotes the probability measure over paths of b induced by the law of motion [**2.1**], and let $\mathbf{E}_t$ denote the associated conditional expectation operators.

subject to the initial condition $A_0^{\hat{c}} = A_0$ and the *no Ponzi* condition

$$\text{[NP-A]} \qquad \lim_{t\to\infty} e^{-rt} A_t^{\hat{c}} \geqslant 0 \qquad \mathbf{P}\text{-a.s.}$$

A consumption strategy $\hat{c}$ is *optimal in the agent's self-insurance problem* if it attains the supremum in [5.1], where $V^{\mathrm{SI}}(A_0, b_0)$ is the agent's *self-insurance value function*.

Lemma 5.1. The consumption strategy $\hat{c}^*$ defined by

$$\text{[5.3]} \qquad \hat{c}_t^* := \hat{C}(A_0, b_0) + \left(\frac{r - \rho + \sigma^2 f(r;\lambda)^2/2}{\theta}\right) t + \frac{\sigma f(r;\lambda)}{\theta} W_t$$

is optimal in the agent's self-insurance problem,[32] where $f(r;\lambda) := r\theta/(r+\lambda)$ and

$$\text{[5.4]} \qquad \hat{C}(A, b) := rA + \frac{r}{r+\lambda} b - \bar{A}(r;\lambda)$$

$$\text{[5.5]} \qquad \bar{A}(r;\lambda) := \frac{r - \rho + \sigma^2 f(r;\lambda)^2/2}{r\theta} - \frac{\mu}{r+\lambda}.$$

The proof of Lemma 5.1 is in Appendix E. The consumption strategy in [5.3]–[5.5] is the continuous-time limit of the discrete-time self-insurance solutions in Caballero (1990) and Wang (2003), and was previously derived in continuous time by Wang (2004, 2006). Our derivation involves slightly different arguments than in this prior work, allowing us to dispense with some technical conditions imposed therein.[33]

Lemma 5.1 implies that the agent's optimal consumption strategy can be expressed recursively as $\hat{c}_t^* \equiv \hat{C}(A_t^*, b_t)$, where $A_t^* := A_t^{\hat{c}^*}$ is the induced asset process. Specifically, at each time t, the agent consumes a multiple r of his *permanent income*

$$\text{[5.6]} \qquad A_t^* + \mathbf{E}_t\left[\int_t^\infty e^{-r(\tau-t)} b_\tau \,\mathrm{d}\tau\right] = A_t^* + \frac{1}{r}\left[\frac{r}{r+\lambda} b_t + \frac{\mu}{\lambda+r}\right]$$

adjusted by subtraction of a constant term. This facilitates a natural interpretation of the agent's *risk exposure*, viz., the sensitivity $f(r;\lambda)/\theta = r/(r+\lambda)$ of consumption to endowment shocks in [5.3]. Observe that the derivative of the agent's permanent income [5.6] with respect to his current endowment b_t is $1/(r+\lambda)$. Because wealth effects are absent under CARA utility, the agent optimally responds to a marginal increase in his

[32] Furthermore, this strategy is uniquely optimal ($\mathbf{P}$-a.e.) because the set of feasible consumption strategies is convex (see [E.1] in Appendix E) and the agent's objective function is strictly concave.

[33] Wang (2004, 2006) imposes a technical integrability condition on the space of admissible asset processes, as well as a transversality condition on a function that is conjectured to be the agent's value function (and verified to be so under these assumptions). Our argument instead works directly with the agent's true value function and only requires that the asset process satisfies [5.2] and [NP-A]. This allows us to construct an exact mapping between the agent's consumption strategies in the indirect "self-insurance contracts" of this section and his reporting strategies in their direct revelation counterparts (see Section 6).

endowment by permanently shifting up his future consumption by a constant amount of $r/(r+\lambda)$ at all future dates. This increases the present value of consumption by $1/(r+\lambda)$, exactly matching the increase in permanent income.

It will be useful in what follows to note the agent's risk exposure is strictly increasing in r when $\lambda > 0$ but is constant in r when $\lambda = 0$. Intuitively, as the coefficient $f(r;\lambda)$ depends only on the ratio λ/r, that ratio is the endowment's "interest-adjusted rate of mean reversion." For a fixed $\lambda > 0$, the adjusted rate λ/r is decreasing in r, so that increasing r increases the "effective persistence" of endowment shocks. However, when $\lambda = 0$, endowment shocks remain perfectly persistent regardless of r.

Lemma 5.1 has several further implications that will prove useful (details are in Appendix E). First, the agent's optimal strategy satisfies the familiar *Euler equation*

[**5.7**] $$e^{(r-\rho)t}u'(\hat{c}_t^*) = \mathbf{E}_t\left[e^{(r-\rho)\tau}u'(\hat{c}_\tau^*)\right] \qquad \text{for all } \tau > t,$$

which specifies that the agent's discounted marginal utility is a martingale (and reduces to [**4.2**] when $r = \rho$). Second, the agent's continuation value process $V = (V_t)_{t\geqslant 0}$ defined by $V_t := V^{\mathrm{SI}}(A_t^*, b_t)$ satisfies

[**5.8**] $$\frac{u(\hat{c}_t^*)}{r} = V_t$$

[**5.9**] $$= V_0 \exp\left[-\left(r - \rho + \frac{\sigma^2 f(r;\lambda)^2}{2}\right)t - f(r;\lambda)\sigma W_t\right].$$

Display [**5.8**] implies that the agent's self-insurance solution induces a *constant utility delivery rate* of r,[34] while [**5.9**] implies that the agent's continuation value process V_t is a geometric Brownian motion and the *discounted* value process $e^{(r-\rho)t}V_t$ is a martingale. Finally, plugging the optimal strategy [**5.3**] into the agent's objective [**5.1**] lets us express the agent's self-insurance value function as

[**5.10**] $$V^{\mathrm{SI}}(A_0, b_0) = \hat{V}^{\mathrm{SI}} \exp\left[-\theta r\left(A_0 + \frac{b_0}{r+\lambda}\right)\right]$$

where $\hat{V}^{\mathrm{SI}} := -\frac{1}{r}\exp\left(\theta\bar{A}(r;\lambda)\right)$ is a constant.

5.2. Self-Insurance Contracts

We now introduce a class of contracts defined by an indirect implementation in which the principal acts as the agent's "bank" by (i) giving an initial lump-sum (asset) transfer to the agent and then (ii) allowing the agent to self-insure at interest rate $r > 0$, which

[34]This property arises because, under CARA utility, the Euler equation [**5.7**] implies that discounted flow utility $e^{(r-\rho)t}u(\hat{c}_t^*)$ is itself a martingale.

is determined by a constant marginal tax (if $r < \rho$) or subsidy (if $r > \rho$) imposed on the agent's savings. (Recall from Section 2.1 that ρ can be interpreted as the market rate at which the principal finances the contract.) In this implementation, the agent is *only* able to borrow and save via his "account" with the principal—*not* in the ambient market—so that the principal can observe, and hence tax, the agent's asset holdings. The principal does not observe the agent's endowment or actual consumption. There is no communication: the agent does not submit endowment reports.

Definition 5.2. The *Self-Insurance Contract* (SI Contract) (b_0, q_0, r) is the indirect mechanism consisting of the following steps:

(i) The principal gives the agent initial assets

[**5.11**] $$A_0(b_0, q_0, r) := \frac{\bar{c}(q_0, r)}{r} - \frac{b_0}{r + \lambda} + \frac{\bar{A}(r; \lambda)}{r}$$

where $\bar{c}(q, r) := -\log(-rq)/\theta$ and $\bar{A}(r; \lambda)$ is defined in [**5.5**].

(ii) The principal, acting as the agent's bank, allows the agent to solve his self-insurance problem at the rate r (as defined in Section 5.1) by imposing a constant marginal tax (or subsidy) $\tau(r) := 1 - r/\rho$ on the agent's capital gains ρA_t. The principal thereby collects tax revenue $\tau(r)\rho A_t$ at each instant.

The principal's expected lifetime cost is $\Pi(b_0, q_0, r) := A_0 - \mathbf{E}_0\left[\int_0^\infty e^{-\rho t}\tau(r)\rho A_t \mathrm{d}t\right]$.

Clearly, the agent's best-response when faced with the SI Contract (b_0, q_0, r) is to follow the consumption-saving strategy described in Lemma 5.1 for the induced self-insurance problem with initial assets $A_0(b_0, q_0, r)$ and rate r. It is then easy to verify from [**5.8**] that the agent's initial lifetime utility satisfies $V_0 = \frac{1}{r}u\big(\hat{C}(A_0(b_0, q_0, r), b_0)\big) = q_0$, so that the contract does, in fact, deliver the requisite promised utility. Plugging the agent's optimal strategy into the principal's cost function, we find that the principal's cost of the SI Contract (b_0, q_0, r) admits a simple closed-form expression.

Lemma 5.3. The principal's cost $\Pi(b_0, q_0, r)$ of the SI Contract (b_0, q_0, r) satisfies

[**5.12**] $$\begin{aligned} \Pi(b_0, q_0, r) &= J^*(b_0, q_0) + \frac{\log(\rho/r)}{\rho\theta} + \frac{r - \rho}{\theta\rho^2} + \underbrace{\frac{\sigma^2\theta r^2}{2\rho^2(r + \lambda)^2}}_{= \sigma^2 f(r;\lambda)^2/(2\theta\rho^2)} \\ &= \mathbf{E}_0\left[\int_0^\infty e^{-\rho t}\left(\hat{c}_t^* - b_t\right)\mathrm{d}t\right] \end{aligned}$$

where $J^*(q_0, b_0)$ is the principal's *first-best* cost function and $\hat{c}^*$ is defined in [**5.3**].[35]

[35]That is, $J^*(b_0, q_0)$ is the principal's cost of providing full insurance in the full-information problem where the agent's endowment is observable and contractable. It is easy to show, as in PPI (p. 1252), that $J^*(b, q) = [(\rho + \lambda)\bar{c}(q, \rho) - \mu - \rho b]/\rho(\rho + \lambda)$, where $\bar{c}(q, \rho) = -\log(-\rho q)/\theta$ as defined above in [**5.11**].

The proof of Lemma 5.3 is in Appendix F. Display [5.12] states that the principal's cost of the SI Contract (b_0, q_0, r), which is defined in terms of the initial wealth transfer and subsequent tax revenue, coincides with the expected resource cost of the contract, as defined in [2.2] for truthful direct revelation contracts.[36]

5.3. Contract PPI and the Optimal SI Contract

We now establish the main result of this section: Contract PPI can be indirectly implemented as an SI Contract with zero taxes ($r = \rho$), whereas the optimal SI Contract features strictly positive taxes ($r < \rho$) whenever endowment shocks are transient ($\lambda > 0$). Consequently, Contract PPI is generically suboptimal within the class of SI Contracts.

Implementing Contract PPI as an SI Contract. To begin, we observe that the agent's consumption process and the principal's costs are identical under (i) Contract PPI if the agent reports truthfully and (ii) the corresponding SI Contract with rate $r = \rho$ if the agent follows his *optimal* consumption-saving strategy.

To verify this, compare the agent's promised utility process q and recommended consumption process c under Contract PPI (recall [3.3] and [3.4]) to his value function process V and optimal consumption strategy $\hat{c}$ in the self-insurance problem with interest rate $r = \rho$ (recall [5.3] and [5.9]). If the agent follows the truthful reporting strategy m^* in Contract PPI, then these processes are identical: $q = V$ and $c = c^{m^*} = \hat{c}$. By Lemma 5.3, the indirect implementation costs the principal

$$\Pi(b_0, q_0, \rho) = J^*(b_0, q_0) + \frac{\sigma^2 \theta}{2(\rho + \lambda)^2}, \tag{5.13}$$

which coincides with PPI's expression (p. 1253) for the principal's lifetime cost (as defined in [2.2]) under Contract PPI.

In other words, Contract PPI (under truthful reporting) is outcome-equivalent to letting the agent self-insure at the ambient market rate ρ. This equivalence allows us to reinterpret the two main economic findings that PPI derives from Contract PPI in terms of known results from the self-insurance literature. First, as stated in Fact 2(i), PPI finds that Contract PPI leads to long-run bliss: $c_t \to +\infty$ and $q_t \to 0$ almost surely. Correspondingly, in the context of self-insurance, it is known quite generally that when the interest rate satisfies $r \geqslant \rho$, the agent's desire to accumulate precautionary savings leads him to both save and consume without bound: $A_t^*, \hat{c}_t^* \to +\infty$ $\mathbf{P}$-a.s. (Sotomayor (1984); Chamberlain and Wilson (2000); Ljunqvist and Sargent (2000, Ch. 17)). Second, PPI (pp. 1254–55) finds that, under Contract PPI, the quality of risk-sharing degrades as

[36] Furthermore, any SI Contract can be implemented using only flow transfers (as in Section 2) by replacing the lump-sum transfer of A_0 with a deterministic transfer process $\hat{s}$ satisfying $\mathrm{d}\hat{s}_t := (\alpha - \lambda \hat{s}_t)\,\mathrm{d}t$ for suitably chosen parameters $\hat{s}_0, \alpha \in \mathbb{R}$.

the agent's endowment becomes more persistent (i.e., $f(\rho;\lambda)$ is decreasing in λ). This comparative static is known to arise from the agent's optimal consumption-smoothing in the corresponding self-insurance problem (Caballero (1990); Wang (2003)).

The Optimal SI Contract. This discussion suggests a simple way to improve upon Contract PPI. Namely, Lemma 5.3 and a short calculation reveal that

$$\frac{\mathrm{d}}{\mathrm{d}r}\Pi(b_0,q_0,r)|_{r=\rho} \geqslant 0, \text{ with strict inequality if and only if } \lambda > 0.$$

Thus, whenever $\lambda > 0$, some SI Contract (b_0, q_0, r) with $r < \rho$ delivers the same lifetime utility q_0 to the agent at a strictly lower cost to the principal than [5.13]. In other words, Contract PPI is generically improvable by taxing the agent's savings.

Theorem 1. *For any initial condition (b_0, q_0), there exists an optimal SI Contract (b_0, q_0, r^*), where r^* is a minimizer of $\Pi(b_0, q_0, \cdot)$ from [5.12].[37] It satisfies the following:*

(i) *If $\lambda > 0$, then $r^* < \rho$ and the optimal SI Contract has a strictly lower cost than Contract PPI.*

(ii) *If $\lambda = 0$, then $r^* = \rho$ and the optimal SI Contract implements Contract PPI.*

Proof. Let (b_0, q_0) be given. It is easy to see that $\lim_{r\to 0}\Pi(b_0,q_0,r) = \lim_{r\to\infty}\Pi(b_0,q_0,r) = \infty$ and that $\Pi(b_0,q_0,\cdot)$ is continuously differentiable on $\mathbb{R}_{++}$. Consequently, there exists a minimizer $r^* \in \mathbb{R}_{++}$, and any such minimizer satisfies the necessary first-order condition $\frac{\mathrm{d}}{\mathrm{d}r}\Pi(b_0,q_0,r)|_{r=r^*} = 0$, which can be expressed as

$$\text{[5.14]} \qquad \frac{\mathrm{d}}{\mathrm{d}r}\left[-\log(r)+\frac{r}{\rho}\right]\Bigg|_{r=r^*} + \underbrace{\frac{\mathrm{d}}{\mathrm{d}r}\left[\frac{\sigma^2 f^2(r;\lambda)}{2\rho}\right]}_{>0 \text{ if } \lambda>0,\ =0 \text{ if } \lambda=0}\Bigg|_{r=r^*} = 0.$$

When $\lambda > 0$, the second term in [5.14] is strictly positive, implying that the first term is strictly negative. It follows that any minimizer satisfies $r^* < \rho$, delivering part (i). When $\lambda = 0$, the second term in [5.14] is identically zero and the objective $\Pi(b_0, q_0, \cdot)$ is strictly convex. Thus, the unique minimizer is $r^* = \rho$, delivering part (ii). □

To get additional intuition for Theorem 1, observe that under an SI Contract with interest rate r, the agent's continuation utility process V from [5.9] evolves as

$$\text{[5.15]} \qquad \frac{\mathrm{d}V_t}{V_t} = (\rho - r)\mathrm{d}t - f(r;\lambda)\sigma\,\mathrm{d}W_t.$$

Intuitively, [5.15] reflects two distortions away from first-best insurance provision: the "drift distortion" determined by the difference $|\rho - r| > 0$ and the "risk distortion"

[37] Going forward, we slightly abuse terminology by referring to *the* optimal SI Contract. The optimum is always unique when $\lambda = 0$ and, when $\lambda > 0$, is unique for generic specifications of model parameters.

determined by the risk exposure coefficient $f(r;\lambda) > 0$. The optimal SI Contract results from the principal's optimal trade-off between these two distortions. Note that when $r = \rho$, as in Contract PPI, there is non-zero risk exposure but zero drift distortion, corresponding to PPI's observation that the agent's promised utility under Contract PPI is a martingale (Fact 2(iii)). There are two cases to consider:

(a) *Transient endowment shocks ($\lambda > 0$):* In this case, the risk exposure coefficient $f(r;\lambda)$ is strictly increasing in r. Therefore, marginally decreasing r from the value $r = \rho$ has two effects: it creates a non-zero drift distortion while decreasing the agent's risk exposure. The principal's first-order condition [**5.14**] equates the marginal cost of the former effect (first term) with the marginal benefit of the second effect (second term), leading to an optimal rate of $r^* < \rho$.[38] This corresponds to imposing a strictly positive tax on the agent's savings, consistent with the literature's finding that optimal contracts in related insurance settings feature a "savings wedge" that relaxes the agent's IC contraints (e.g., Golosov, Kocherlakota, and Tsyvinski (2003)).

(b) *Permanent endowment shocks ($\lambda = 0$):* In this case, the risk exposure coefficient $f(\cdot;0) \equiv \theta$ is constant in r. Consequently, the principal cannot manipulate the agent's risk exposure by imposing a tax or subsidy on savings, as reflected by the fact that the second term in the principal's first-order condition [**5.14**] vanishes. Moving r away from the value $r = \rho$ therefore only generates the cost of a non-zero drift distortion, without creating any risk-reduction benefit. To eliminate this cost, the principal optimally sets $r = \rho$ and thereby implements Contract PPI.

In Section 6 below, we show that SI Contracts are IC when reformulated as direct-revelation contracts. This analysis, combined with Theorem 1, completes the proof of our main observation that Contract PPI is generically suboptimal (Observation 3).

5.4. Further Analysis

Properties of SI Contracts. In Appendix G, we establish some additional useful facts about SI Contracts. We summarize them here:

(i) The class of SI Contracts is characterized by the property, which we call *Stationarity*, that promised utility and marginal promised utility are proportional, viz., $p_t \equiv k_0 q_t$ for some constant $k_0 > 0$. Per Fact 2(ii), Contract PPI plainly has this property. As noted in Appendix A, the key incorrect step in PPI's derivation of Contract PPI involves arguing that restricting to Stationary Contracts is without loss of optimality. We show that this conclusion is false: Theorem 1 is equivalent to the statement that Contract PPI is optimal *within the class of Stationary Contracts* if and only if $\lambda = 0$, contradicting that Contract PPI is generally optimal among all IC contracts.

[38] It is never optimal to set $r > \rho$ because doing so creates a non-zero drift distortion while also increasing the agent's risk exposure, both of which are costly to the principal.

(ii) The class of SI Contracts coincides with the class of “state-consistent” renegotiation-proof contracts introduced by Strulovici (2022). This equivalence provides a separate foundation for our focus on SI Contracts based on the principal’s limited commitment. It also shows that SI Contracts provide a simple indirect implementation for the full class of renegotiation-proof contracts in PPI’s model.

(iii) The long-run behavior of the optimal SI Contract depends starkly on the persistence of the agent’s endowment: there exists a threshold $\overline{\lambda} > 0$ such that this contract results in immiseration when persistence is sufficiently low ($\lambda > \overline{\lambda}$) and results in bliss when persistence is sufficiently high ($\lambda < \overline{\lambda}$). This implies that the long-run bliss property of Contract PPI is generally not robust to optimization within the class of SI Contracts. This result may also be of independent interest, given the literature’s emphasis on long-run properties of optimal contracts.

Hidden Savings. PPI adopts the standard assumption that the agent cannot save or borrow outside of his relationship with the principal. Our SI Contracts rely on this assumption: if instead the agent could covertly trade at the market rate ρ, he could circumvent any tax imposed by the principal. Since Contract PPI is implemented with zero taxes, this suggests that it may be uniquely robust to such manipulation. We formalize this idea in Appendix H, where we consider an *alternative version* of PPI’s model with hidden savings. In this setting, Theorem 4 shows, in essence, that every IC contract is outcome-equivalent to Contract PPI, which is therefore the optimal contract. This parallels Allen’s (1985) and Cole and Kocherlakota’s (2001) classic discrete-time characterizations of optimal hidden savings contracts and provides a viable foundation for Contract PPI, albeit one that is rather different than suggested in PPI.

6. Incentive Compatibility of SI Contracts

In Section 5, we characterized the agent’s optimal consumption strategy in an *indirect* implementation. We now reformulate SI Contracts as *direct* mechanisms and show that they are IC, i.e., truthtelling is the agent’s optimal reporting strategy. Our analysis addresses the two issues in PPI related to IC (Observations 1 and 2) and introduces some new techniques that may be applicable in other continuous-time models.

6.1. Extended Reporting Problem

We will see below (Section 6.4) that, if the agent’s recent reports are not truthful ($\lim_{\varepsilon \searrow 0} m_{t-\varepsilon} \neq 0$), then he has an incentive to “instantly” revert to truthtelling (set $m_t = 0$). This requires him to submit a discontinuous “jump report.” While PPI’s formulation allows the agent to approximate such behavior (by sending $|\Delta_t| \to \infty$ for a vanishing measure of times), discontinuous report paths are ruled out by construction. Thus, for the purpose of analyzing IC, we find it convenient in this section to “extend”

PPI's model formulation (from Section 2.1) in two ways:[39]

(i) We allow the agent's feasible set F of misreporting strategies to include $m \notin \mathscr{M}$.

(ii) We allow the principal's contract, the y-adapted transfer process s, to respond in any way we choose to y processes corresponding to $m \in F \backslash \mathscr{M}$. To maintain consistency with Section 2.1, we still require y processes corresponding to $m \in \mathscr{M}$ to induce continuous sample paths of s.

This approach is justified by an elementary variant of Fact 1: *if a contract as defined in (ii) above is F-IC for a given feasible set F of b-adapted misreporting strategies (not necessarily satisfying $F \subseteq \mathscr{M}$), then the contract is also $[F \cap \mathscr{M}]$-IC.* If the agent does not have a profitable deviation from truthtelling within the "extended" feasible set F, then he cannot have a profitable deviation within the "actual" (smaller) feasible set $F \cap \mathscr{M}$. Therefore, if the goal is to verify that a contract is $[F \cap \mathscr{M}]$-IC, then we are free to "extend" the contract arbitrarily to paths of y that correspond to "infeasible" paths of $m \in F \backslash \mathscr{M}$, and then verify that the resulting "extended" contract is F-IC.[40,41]

6.2. Direct Revelation SI Contracts

We can recast any SI Contract as a direct-revelation contract as follows:

Definition 6.1. The *Direct-Revelation Self-Insurance Contract* (DR-SIC) (b_0, q_0, r) is the direct revelation mechanism in which:

(i) The principal keeps track of the y-adapted *virtual asset* process A^v defined by

[6.1] $$A_t^v := A_0(b_0, q_0, r) + \bar{A}(r;\lambda)t + \frac{\lambda}{r+\lambda}\int_0^t y_\tau \, \mathrm{d}\tau$$

where $A_0(b_0, q_0, r)$ is defined in [**5.11**] and $\bar{A}(r;\lambda)$ is defined in [**5.5**].

(ii) The agent's *recommended consumption* is the y-adapted process c defined by

[6.2] $$c_t := \hat{C}\left(A_t^v, y_t\right) = rA_t^v + \frac{r}{\lambda + r} y_t - \bar{A}(r;\lambda)$$

[39] This approach builds on ideas from Strulovici (2022). While this approach is convenient, it is *not* necessary: Section 7.3 and Appendix I describe how to analyze IC under PPI's assumption that $F \subseteq \mathscr{M}$.

[40] One possible "extension," described in Footnote 16, involves "shooting" the agent if his strategy $m \in F \backslash \mathscr{M}$ generates an "immediately detectable" deviation, such as a discontinuous sample path of reports $\hat{y} \notin \mathbf{C}[0, \infty)$, with positive probability. While this extension is without loss of generality, it is "trivial" in the sense that it effectively prevents the agent from using $m \notin \mathscr{M}$. In Section 6.2–6.4 below, we find it more useful to consider an extension that "compactifies" the agent's reporting problem by considering a "closure" of the original contract that treats jump reports as the limit of "approximate jump" reports (whereby $|\Delta_\tau| \to \infty$ for a vanishing measure of times).

[41] When the feasible set F permits m with discontinuous sample paths, we can no longer view the agent as choosing a measure $\mathbf{P}^m$ on the space $\mathbf{C}([0, \infty))$ of continuous paths (cf. Footnote 11). This has no adverse consequence for our analysis in Section 6.2–6.4, which does not use any change of measure.

where $\hat{C}(A, y)$ is defined in [**5.4**].

Put simply, in a DR-SIC the principal "saves" on the agent's behalf and recommends that he consume as would be optimal in his self-insurance problem, *assuming that the agent is truthful so that* $y \equiv b$. We emphasize that virtual assets are not a "physical" object, but simply a state variable that the principal uses to track what the agent would have done in his self-insurance problem.[42]

Observe that every DR-SIC: (a) is well-defined under any report process y that induces a well-defined virtual asset process A^v in [**6.1**], and (b) specifies continuous transfers whenever y is continuous. Denote the corresponding set of *extended* misreporting strategies by

[**6.3**] $$\mathscr{M}_{\text{ext}} := \left\{ m : m \text{ is } b\text{-adapted and } \int_0^t |m_\tau| \mathrm{d}\tau < \infty \ \mathbf{P}\text{-a.s. } \forall t \geqslant 0 \right\}.$$

Following our treatment of tail restrictions from Section 2.2, we also define, for every $r > 0$, the subset of extended strategies $\mathscr{M}_{\text{ext}}^r \subsetneq \mathscr{M}_{\text{ext}}$ by

[**6.4**] $$\mathscr{M}_{\text{ext}}^r := \{ m \in \mathscr{M}_{\text{ext}} : \text{ [NP-}m\text{] holds for rate } r \}.$$

We will study the "extended" reporting problem in which the agent's feasible set is $F = \mathscr{M}_{\text{ext}}^r$ (for a suitable value of $r > 0$). Clearly, $\mathscr{M}_{\text{ext}}^r \supsetneq \mathscr{M}^r$ because [**6.3**]–[**6.4**] allow for strategies with discontinuous sample paths. Notably, any DR-SIC responds to a "jump report" at time t (whereby $m_t - \lim_{\varepsilon \searrow 0} m_{t-\varepsilon} = M \neq 0$) as if it were the limit of "approximate jump reports" at time t (whereby $\Delta_\tau = M/\varepsilon$ for $\tau \in [t, t+\varepsilon)$ and $\varepsilon \searrow 0$).[43]

Using this notation, we can state the main result of this section as follows:

Theorem 2. *For any given $r > 0$, every DR-SIC (b_0, q_0, r) is $\mathscr{M}_{\text{ext}}^r$-IC. Thus, such contracts are also F-IC for any smaller strategy space $F \subseteq \mathscr{M}_{\text{ext}}^r$.*

The assumptions in Theorem 2 cannot be substantively relaxed: by mimicking the proof of Observation 2, it is easy to show that no DR-SIC with rate r is $\left[\mathscr{M}_{\leqslant}^{\text{LAC}} \cap \mathscr{M}^{r'}\right]$-IC for any $r' > r$. We conclude that sufficiently tight no Ponzi constraints are essential for DR-SICs to be IC (while sign and absolute continuity restrictions are not).

We provide two independently instructive proofs of Theorem 2 below. Section 6.3 presents an indirect "revelation principle" argument in which we characterize the mapping between reporting strategies in a DR-SIC and consumption-saving strategies in

[42] As defined above, DR-SICs are recursive in the state variables (A_t^v, y_t), viz., the y-adapted transfer process s can be written as $s_t = \hat{C}(A_t^v, y_t) - y_t$. By defining promised utility as $q_t := V^{\text{SI}}(A_t^v, y_t)$ (where V^{SI} is defined in [**5.10**]), we can equivalently write any DR-SIC recursively in the state variables (q_t, y_t), as in PPI's treatment. In this alternative formulation, it is easy to see that a DR-SIC with interest rate ρ coincides with Contract PPI (see Definition 3.1) under any reporting strategy.

[43] See Lemma B.1 and Remark B.2 in Appendix B.

the corresponding indirect SI Contract. This approach clarifies that the no Ponzi condition [**NP-*m***] on the agent's misreports in the direct mechanism is equivalent to the no Ponzi condition [**NP-A**] on the agent's assets in the indirect mechanism. Section 6.4 then presents an alternative argument in which we directly analyze the agent's reporting incentives in a DR-SIC using stochastic control. This approach elucidates the agent's reporting incentives at non-truthful histories.

6.3. Indirect Proof: Revelation Principle

We establish Theorem 2 in two steps. First, we characterize the agent's actual consumption process c^m in a DR-SIC under an arbitrary misreporting strategy $m \in \mathscr{M}_{\text{ext}}$ (Lemma 6.2). Second, we show that an actual consumption process c^m can be induced in a DR-SIC with interest rate r by some strategy $m \in \mathscr{M}^r_{\text{ext}}$ if and only if c^m is a feasible consumption strategy for the agent in the corresponding indirect SI Contract (Lemma 6.3). Theorem 2 is then immediate: the absence of profitable deviations in the indirect implementation implies the absence of profitable deviations in the direct mechanism.

Step 1: Induced Consumption Processes. Let $\hat{c}^*$ denote the agent's optimal consumption strategy in the SI Contract (b_0, q_0, r) and let A^* denote the corresponding asset process, so that $\hat{c}^*_t = \hat{C}(A^*_t, b_t)$ as in Lemma 5.1. The following lemma characterizes how misreports in the direct mechanism correspond to consumption choices in the agent's self-insurance problem.

Lemma 6.2. Given any misreporting strategy $m \in \mathscr{M}_{\text{ext}}$, the agent's actual consumption $c^m_t := c_t - m_t$ in the DR-SIC (b_0, q_0, r) satisfies

$$c^m_t = \hat{C}(A^v_t, b_t) - \frac{\lambda}{r+\lambda} m_t \tag{6.5}$$

$$= \hat{c}^*_t + \frac{\lambda}{r+\lambda}\left(r\int_0^t m_\tau \,\mathrm{d}\tau - m_t\right). \tag{6.6}$$

Proof. Display [**6.5**] follows from recalling the identity $c^m_t \equiv c_t - m_t$ and substituting the identity $y_t \equiv b_t + m_t$ into the expression for c_t in [**6.2**]. To obtain [**6.6**], observe from [**6.1**] that $A^v_t \equiv A^*_t + \frac{\lambda}{r+\lambda}\int_0^t y_\tau \mathrm{d}\tau$ and substitute this identity into [**6.5**]. □

Display [**6.5**] implies that under-reporting by setting $m_t < 0$ in the DR-SIC corresponds to over-consuming by $-\frac{\lambda}{r+\lambda} m_t > 0$ in the indirect SI Contract, *holding fixed the agent's assets at A^v_t*. At the same time, because virtual assets evolve as

$$\frac{\mathrm{d}A^v_t}{\mathrm{d}t} = \frac{\mathrm{d}A^*_t}{\mathrm{d}t} + \frac{\lambda}{r+\lambda} m_t, \tag{6.7}$$

an under-report of $m_t < 0$ in the DR-SIC corresponds to under-saving by $\frac{\lambda}{r+\lambda} m_t < 0$ in the indirect SI Contract. Display [6.6] describes how these virtual savings distortions accumulate over time, so that the agent's actual consumption c_t^m in the DR-SIC may be either higher or lower than his optimal consumption $\hat{c}_t^*$ in the indirect SI Contract.

Step 2: Outcome-Equivalence of Feasible Sets. The following lemma relates the sets of feasible consumption processes in the direct and indirect implementations.

Lemma 6.3. For any b-adapted consumption process $\hat{c}$, the following are equivalent:

(i) It is feasible in the SI Contract (b_0, q_0, r).

(ii) It satisfies $\hat{c} = c^m$ for some misreporting strategy $m \in \mathscr{M}_{\text{ext}}^r$ in the corresponding DR-SIC (b_0, q_0, r).

Proof. That (ii) implies (i) is immediate from the definitions. To see that (i) implies (ii), suppose that $\hat{c}$ is feasible in the SI Contract (b_0, q_0, r) and let $A^{\hat{c}}$ be the corresponding asset process (i.e., solution to [5.2]–[NP-A]). Motivated by [6.5], define the misreporting strategy $m_t^{\hat{c}} := \frac{r+\lambda}{\lambda}\left[\hat{C}(A_t^{\hat{c}}, b_t) - \hat{c}_t\right]$, which is in $\mathscr{M}_{\text{ext}}$ by construction.[44] We claim that $m^{\hat{c}}$ induces the virtual asset process $A^v = A^{\hat{c}}$ and actual consumption process $c^{m^{\hat{c}}} = \hat{c}$ in the DR-SIC (b_0, q_0, r). Substitute $m^{\hat{c}}$ into [6.1], expand the definition of $\hat{C}$ from [5.4], and recall that $A^{\hat{c}}$ (being a solution to [5.2]) satisfies $A_t^{\hat{c}} = A_0(b_0, q_0, r) + \int_0^t \left(rA_\tau^{\hat{c}} + b_\tau - \hat{c}_\tau\right) \mathrm{d}\tau$. After simplification, this yields $A_t^v \equiv A_t^{\hat{c}}$. Since $A^{\hat{c}}$ satisfies [NP-A], it follows that $m^{\hat{c}}$ satisfies [NP-*m*], and thus $m^{\hat{c}} \in \mathscr{M}_{\text{ext}}^r$. Finally, substituting $A_t^v \equiv A_t^{\hat{c}}$ and the definition of $m^{\hat{c}}$ into [6.5] (from Lemma 6.2) yields $c_t^{m^{\hat{c}}} \equiv \hat{c}_t$, completing the proof. □

The key observation underlying Lemma 6.3 is that, in a DR-SIC, the no Ponzi condition [NP-*m*] on misreports is *equivalent* to the no Ponzi condition [NP-A] on his virtual asset process A^v. Formally, given any misreporting strategy $m \in \mathscr{M}_{\text{ext}}$, the limits $\lim_{T\to\infty} e^{-rT} A_T^v$ and $\lim_{T\to\infty} e^{-rT} \int_0^T m_t \,\mathrm{d}t$ are P-a.s. equal.[45] Since it is well-understood that [NP-A] is necessary for the agent's self-insurance problem to be well-posed, this suggests that [NP-*m*] is a natural constraint on misreporting strategies in DR-SICs.

Proof of Theorem 2. By definition, $\hat{c}^*$ is the agent's optimal consumption process in the indirect SI Contract (b_0, q_0, r). Thus, Lemma 6.3 implies that any misreporting strategy $m \in \mathscr{M}_{\text{ext}}^r$ for which $c^m = \hat{c}^*$ is optimal for the agent in the corresponding

[44] Clearly, $m^{\hat{c}}$ is b-adapted. To verify that it is integrable, expand the definition of $\hat{C}$ from [5.4], regroup terms, and use the facts that b solves [2.1] and $A^{\hat{c}}$ solves [5.2] to conclude that $\int_0^t |m_\tau^{\hat{c}}| \mathrm{d}\tau \leqslant \int_0^t |rA_\tau^{\hat{c}} + b_\tau - \hat{c}_\tau| \mathrm{d}\tau + \int_0^t |\bar{A}(r;\lambda)| \mathrm{d}\tau + \frac{\lambda}{r+\lambda} \int_0^t |b_\tau| \mathrm{d}\tau < \infty$ P-a.s.

[45] To see this, integrate the "virtual savings" equation [6.7] to obtain $A_T^v = A_T^* + \lambda(\lambda + r)^{-1} \int_0^T m_t \,\mathrm{d}t$, where A^* is the asset process in the corresponding SI Contract induced by the agent's optimal consumption-saving strategy therein (Lemma 5.1). We have $\lim_{T\to\infty} e^{-rT} A_T^* = 0$ because the agent optimally "leaves no money on the table" in his self-insurance problem (see Lemma E.3 in Appendix E).

DR-SIC (b_0, q_0, r). By [**6.6**] in Lemma 6.2, the truthful strategy $m^* \equiv 0$ is one such strategy. It follows that the DR-SIC (b_0, q_0, r) is $\mathscr{M}^r_{\text{ext}}$-IC. $\square$

6.4. Direct Proof: Dynamic Programming

We now establish Theorem 2 by directly analyzing the agent's reporting incentives.[46]

Proof of Theorem 2. Fix a DR-SIC with interest rate r. By [**6.5**] in Lemma 6.2, the agent's actual consumption process c^m is Markovian in the contemporaneous virtual assets A^v_t, true endowment b_t, and misreport m_t. Thus, the agent's reporting problem (with strategy space $\mathscr{M}^r_{\text{ext}}$) can be cast as one of stochastic control in which (A^v_t, b_t) serve as state variables and m_t is a control variable. Let $V^{\text{DR}}(A^v_t, b_t)$ denote the agent's value function in this control problem. By standard arguments, it can be shown that[47]

$$V^{\text{DR}}(A^v_t, b_t) = \hat{V}^{\text{DR}} \exp\left[-\theta r \left(A^v_t + \frac{b_t}{r+\lambda}\right)\right] \tag{6.8}$$

for some constant $\hat{V}^{\text{DR}} < 0$, and that $V^{\text{DR}}(A^v_t, b_t)$ is a solution to the HJB equation

$$\begin{aligned}\rho V^{\text{DR}}(A^v_t, b_t) = \sup_{m_t \in \mathbb{R}} \Bigg[u\left(\hat{C}(A^v_t, b_t) - \frac{\lambda m_t}{r+\lambda}\right) + \left(\bar{A}(r;\lambda) + \frac{\lambda(b_t + m_t)}{r+\lambda}\right) V^{\text{DR}}_A(A^v_t, b_t)\Bigg] \\ + (\mu - \lambda b_t) V^{\text{DR}}_b(A^v_t, b_t) + \tfrac{1}{2}\sigma^2 V^{\text{DR}}_{bb}(A^v_t, b_t).\end{aligned} \tag{6.9}$$

Routine analysis of [**6.9**] then yields a full characterization of the agent's optimal strategy: since the supremum in [**6.9**] is *uniquely attained by setting* $m_t = 0$, the agent's unique optimal strategy is to truthfully report his current endowment *at every history*.[48] It follows that the given DR-SIC with interest rate r is $\mathscr{M}^r_{\text{ext}}$-IC. $\square$

This proof implies that the agent truthfully reports his current type b_t (by setting

[46] Strulovici (2022) uses a related argument to analyze IC for an equivalent class of renegotiation-proof contracts (see Section 5.4). A notable technical difference is that Strulovici (2022) casts the agent's reporting problem as one of impulse-control, viewing m as a state variable (as in PPI's treatment) and Δ as an extended-real-valued control variable, with infinite values of Δ corresponding to jumps in m. Our approach of viewing m as a control variable parallels our treatment of the agent's self-insurance problem (in which consumption, which has the same units as m, is a control variable), and facilitates connections to discrete-time reporting problems (see the discussion below and in Section 7.3).

[47] The details are analogous to those from our derivation of the agent's self-insurance solution (Lemma 5.1) in Appendix E. The fact that V^{DR} satisfies [**6.8**] with $\hat{V}^{\text{DR}} < 0$ (rather than $\hat{V}^{\text{DR}} = 0$) follows for Lemma E.1, with [**NP-*m***] replacing [**NP-A**]. That V^{DR} satisfies [**6.9**] follows from standard results in stochastic control (Yong and Zhou 1999, Theorem 3.3; Touzi 2018, Propositions 2.4-2.5).

[48] Plugging [**6.8**] into the first-order condition from [**6.9**] yields that the agent's unique maximizer at state (A^v_t, b_t), call it $m^\star(A^v_t, b_t)$, satisfies $\hat{V}^{\text{DR}} = -r^{-1}\exp\left[\theta\left(\bar{A}(r;\lambda) + \frac{\lambda}{r+\lambda} m^\star(A^v_t, b_t)\right)\right]$. Plugging this expression back into [**6.9**] yields that $\hat{V}^{\text{DR}} = -r^{-1}\exp\left(\theta \bar{A}(r;\lambda)\right)$ (which equals $\hat{V}^{\text{SI}}$ from [**5.10**]) and that $m^\star(A^v_t, b_t) = 0$, as desired. One can then verify that the implied strategy is optimal by following the same verification argument from the proof of Lemma 5.1 (see Appendix E).

$m_t = 0$) *even at off-path histories* where he has recently misreported ($\lim_{\varepsilon \searrow 0} m_{t-\varepsilon} \neq 0$), requiring him to submit a discontinuous "jump" report. This implication is natural from two perspectives. First, in terms of the indirect implementation, it simply means that the agent immediately "re-initializes" his actual consumption c_t^m at the optimal level $\hat{C}(A_t^v, b_t)$ prescribed by Lemma 5.1. Second, it is well known that in discrete-time "Markovian" contracting problems, a contract is IC if and only if it satisfies the (seemingly stronger) property that the agent finds it optimal to truthfully report at *all* histories, both on- and off-path.[49] The strategy derived here exhibits the continuous-time version of this property, highlighting a close connection between discrete- and continuous-time reporting problems that we expect is valid beyond PPI's model (see Section 7.3 below).

7. Discussion

We conclude by discussing implications of our analysis for the fully optimal contract in PPI's hidden endowment model, long-run properties of optimal insurance contracts, and the relation between discrete- and continuous-time contracting models.

7.1. Fully Optimal Contracts

Our analysis leaves open an important question: *What is the fully optimal contract in PPI's hidden endowment model?* The answer depends on whether endowment shocks are transient ($\lambda > 0$) or permanent ($\lambda = 0$).

Transient Shocks. When $\lambda > 0$, it is natural to ask whether the optimal SI Contract is fully optimal. Theorem 7 in Appendix J suggests a negative answer: under regularity conditions, the optimal SI Contract is strictly suboptimal within the broader class of "FO-IC contracts," i.e., the class of contracts for which the agent's value function admits an envelope formula.[50] If the optimal FO-IC contract is actually IC (i.e., the first-order approach is valid), it follows that the optimal SI Contract is not fully optimal. This is to be expected: SI Contracts are renegotiation-proof and "stationary" (Section 5.4), whereas prior work on contracting with persistent states has found that fully optimal contracts typically violate both properties (e.g., Fernandes and Phelan 2000).

A full solution to PPI's model with transient shocks remains an important open question. While we do not know the answer, our analysis does have some implications

[49] In the special case of i.i.d. types, this fact forms the basis of the classic recursive formulations of the agent's incentive constraints in terms of promised utility (Green 1987; Thomas and Worrall 1990; Atkeson and Lucas 1992). For settings with persistent private information, see Fernandes and Phelan (2000, Lemma 2.1, Theorem 2.1) for an early treatment and Pavan, Segal, and Toikka (2014, pp. 620-22, 645-46) for a modern treatment and general definition of "Markovian" environments. This property does not appear in PPI because it is ruled out by PPI's IML assumption, as discussed in Section 4.1.

[50] See Appendix A for a formal definition. As discussed there, this class includes all contracts considered in PPI's analysis, and we conjecture that it includes all IC contracts.

for it. We record them in Appendix A with the hope of stimulating work on this problem.

Permanent Shocks. When $\lambda = 0$, Theorem 1 shows that Contract PPI is optimal among SI Contracts. For this special case, Theorem 6 in Appendix J further shows that Contract PPI is optimal among all FO-IC contracts. This confirms PPI's results for the $\lambda = 0$ case (though we offer a simpler proof). The intuition mirrors that of Theorem 1(ii): when $\lambda = 0$, the principal cannot manipulate his risk exposure, and therefore focuses on eliminating drift distortions. What is potentially surprising is that this logic extends to the full class of FO-IC contracts.

In a more general setting, Bloedel, Krishna, and Strulovici (2022) provide an explanation by observing that, under permanent shocks, the agent is necessarily indifferent among *(essentially) all* reporting strategies under *any* FO-IC contract. Roughly speaking, when the agent's type is subject to permanent shocks, his information rents are "so large" that the principal cannot elicit any useful information. By contrast, when the agent's type is mean-reverting—even arbitrarily slowly—nontrivial screening and risk-sharing is possible because the agent's private information is "short-lived."[51] *We conclude that the nature of IC constraints under permanent shocks are, at least in some respects, fundamentally different than under transient shocks (even as $\lambda \searrow 0$).* It is therefore not surprising that optimal contracts differ in these two cases, as well.

7.2. Immiseration and Persistence

Perhaps the most striking claim in PPI is that the classic immiseration property of (fully) optimal insurance contracts breaks down in PPI's model due to (i) the persistence of the agent's information and (ii) fundamental differences between IC constraints in discrete- and continuous-time models. *Our analysis demonstrates that this claim does not follow from PPI's results in the generic case of transient shocks ($\lambda > 0$) because Contract PPI is optimal only when shocks are permanent ($\lambda = 0$).* We address the role of persistence here and turn to the role of continuous time in Section 7.3 below.

Transient Shocks. Theorem 3 in Appendix G shows that the optimal SI Contract generates immiseration when persistence is low ($\lambda > \overline{\lambda}$) and bliss when persistence is high ($\lambda < \overline{\lambda}$). Because SI Contracts are not fully optimal when $\lambda > 0$ (under regularity conditions), this finding does not have direct implications for the fully optimal contract. In a general discrete-time setting, Bloedel, Krishna, and Leukhina (2021) find that fully optimal contracts generate immiseration when the agent's private type follows a finite-state, fully-connected Markov process. As their model permits arbitrarily good approximations of PPI's hidden endowment model, it is natural to expect that the fully

[51]For any $\lambda > 0$, the influence of the agent's time-t type has a vanishing influence on his time-τ type as $\tau \to \infty$ (recall Footnote 10), so that the principal and agent have "almost symmetric" information over the agent's time-τ preferences.

optimal contract in PPI's model may also generate immiseration for all $\lambda > 0$. Further investigation of this question is an important direction for future research.

Permanent Shocks. When $\lambda = 0$, immiseration fails under the optimal FO-IC contract (Contract PPI). *We view this as a knife-edge result specific to $\lambda = 0$.* The literature often attributes immiseration to the principal's manipulation of the agent's risk exposure for cost-smoothing purposes (see BKL for details). Lemma J.1 in Appendix J formalizes a sense in which the principal can manipulate the agent's risk exposure if and only if $\lambda > 0$. Thus, when $\lambda = 0$, the classic rationale for immiseration is shut off by construction.[52]

7.3. Continuous vs. Discrete Time

A central claim in PPI is that (i) there is a fundamental difference between reporting incentives in discrete- and continuous-time models and (ii) this difference is at least partially responsible for the differences between Contract PPI and optimal contracts derived in the prior literature. This claim is based on the observation that, in PPI's model, the agent's time-t choice of Δ_t affects his future misreports m_τ and actual consumption c_τ^m for $\tau > t$, but does *not* affect the current misreport m_t or consumption c_t^m, which apparently stands in contrast to discrete-time models in which the agent can freely choose his current misreport and therefore affect his current consumption.[53] While the latter observation is correct, *our analysis supports neither component of PPI's claim.*

(i) Reporting Incentives. As discussed in Section 6.4, our direct verification of IC for DR-SICs suggests fundamental similarities between reporting incentives in discrete- and continuous-time models. We showed there how to "extend" the agent's strategy space (and the contract's responses) to allow for discontinuous "jump" reports. In the extended problem, as in discrete-time models, the agent directly chooses his misreport m_t at time t and instantaneously affects his time-t consumption c_t^m. We also described why considering such "extensions" is without loss of generality for the purposes of verifying IC, and why the agent's incentive to truthfully report his current type even at off-path histories—familiar from discrete-time models—makes such "extensions" particularly natural in continuous time.

[52]In a more general setting with permanent shocks, Bloedel, Krishna, and Strulovici (2022) show that the optimal long-run properties depend on details of the agent's type process and utility function.

[53]For instance, PPI writes (p. 1235): "[T]hese differences [between Contract PPI and the optimal contract in Thomas and Worrall (1990)] rely at least partly on differences in the environments. In the discrete [time] analogue of my model, when deciding what to report in the current period, the agent trades off current consumption and future promised utility. In my continuous-time formulation, the agent's private state follows a process with continuous paths and the principal knows this. Thus in the current period the agent only influences the future increments of the reported state. Thus current consumption is independent of the current report and all that matters for the reporting choice is how future transfers are affected … [T]he reporting problem and, hence, the incentive constraints become fully forward-looking …" Similar statements are made elsewhere in PPI (pp. 1239, 1257-58, 1263, 1264).

In Appendix I, we show how to directly analyze the agent's reporting problem in DR-SICs while maintaining PPI's original assumption that misreports have absolutely continuous sample paths ($m \in \mathscr{M}$). Under this assumption, the agent can still approximate the effect of a jump report at time t through a sequence of strategies in $\mathscr{M}$ by sending $|\Delta_\tau| \to \infty$ for a vanishing interval of times $\tau > t$. Intuitively, although the agent can no longer *instantaneously* affect his misreport m_t or his consumption c_t^m, he can still affect these variables at *arbitrarily close* times $\tau > t$, and finds it optimal to do so because DR-SICs respond similarly to "jump" reports and such "approximate jump" reports (see Appendix B). While the details of this argument are specific to PPI's model and the class of DR-SICs, we expect that the basic points apply more broadly in other continuous-time contracting models. Consequently, our analysis suggests that the differences between discrete- and continuous-time models identified in PPI (see Footnote 53) are cosmetic, rather than substantive.[54]

(ii) Optimal Contracts. Our results extend to the natural discrete-time analogue of PPI's hidden endowment model in which the agent's endowment follows a Gaussian AR(1) process. Observation 1 plainly extends. For Observations 2–3 and the associated analysis of SI Contracts, recall that our solution to the agent's self-insurance problem (Lemma 5.1) is the continuous-time limit of Caballero's (1991) and Wang's (2003) discrete-time self-insurance solutions. Thus, we can replicate our analysis of SI Contracts and DR-SICs in discrete time, using those papers' solutions to construct the discrete-time SI Contracts and defining the discrete-time version of Contract PPI as the particular one with zero taxes. Discrete-time versions of Observations 2–3 and Theorems 1 and 2 readily follow. Consequently, any differences between Contract PPI and the optimal contracts in the discrete-time models on which PPI's model is based (Thomas and Worrall 1990; Atkeson and Lucas 1992) are not driven by the distinction between discrete and continuous time. Rather, the discrete-time analogue of Contract PPI is also generically suboptimal in the discrete-time version of PPI's model.[55]

Appendix

Appendices A–B are presented here. Appendices C–K are in the Online Appendix.

A. PPI's Derivation of Contract PPI

We begin with a few (sometimes implicit) definitions from PPI, stated in our terminology. By the Martingale Representation Theorem, the agent's promised utility process (defined

[54] Future work might explore more general connections between discrete- and continuous-time screening problems by studying discrete-time models with vanishing period length (cf. Sadzik and Stacchetti 2015).

[55] Our finding that Contract PPI is optimal in an alternative model with hidden savings extends to discrete time by the same logic described above. Our confirmation that Contract PPI is optimal under permanent shocks extends to discrete time per a special case of Bloedel, Krishna, and Strulovici (2022).

in [**3.1**]) under any contract satisfies

$$\text{[A.1]} \qquad \mathrm{d}q_t = (\rho q_t - u_t)\,\mathrm{d}t + \gamma_t \sigma\,\mathrm{d}W_t^y,$$

where W^y is the "inferred shock process" under reporting strategy y (as defined in Section 3) and γ is a y-adapted *sensitivity process* determined by the contract. Recall the agent's marginal promised utility process p from [**3.2**]. The condition

$$\text{[FO-IC]} \qquad \gamma_t + p_t \equiv 0.$$

corresponds to the familiar *envelope formula* for the agent's information rents: it equates the marginal value γ_t of an increase in the agent's *report* and the marginal value $-p_t$ of an increase in the agent's *true type*. Any contract that satisfies [**FO-IC**] is said to be *first-order IC* (FO-IC). The class of FO-IC contracts includes all DR-SICs (which satisfy [**FO-IC**] with $p_t \equiv f(r;\lambda)q_t$) and all other contracts allowed for in PPI's treatment.[56, 57] We say that a contract is an *optimal FO-IC contract* if it minimizes the principal's lifetime cost under truthtelling (as in [**2.2**]) among all FO-IC contracts delivering a pre-specified promised utility level. Note that an FO-IC contract, including the optimal one, need not be IC.

Standing Assumption. *For the remainder of this appendix, we focus on the generic case of PPI's hidden endowment model with transient shocks ($\lambda > 0$).*

PPI's Claims. PPI's claim that Contract PPI is the optimal $\mathscr{M}_{\leqslant}^{\mathrm{LAC}}$-IC contract can be decomposed into three sub-claims:

[56] §3 of PPI claims (in a general finite-horizon setting) that all IC contracts are FO-IC, and we conjecture that this is indeed the case in PPI's hidden endowment model. The basic idea is to show that [**FO-IC**] is implied by [**IC**] by deriving the former as an infinitesimal optimality condition for the agent that rules out "local" deviations from truthtelling. Technically, this requires appealing to the envelope theorem (e.g., Kapička 2013; Pavan, Segal, and Toikka 2014), the stochastic maximum principle (as in PPI, pp. 1243-44, 1264), or other variational arguments to establish that such infinitesimal conditions are well-defined. The usual approach in the literature is to allow for all conceivable contracts while imposing regularity conditions on the agent's type process and preferences to ensure that such arguments are applicable (e.g., Pavan, Segal, and Toikka 2014, Theorem 1). However, to our knowledge, the agent's CARA utility function in the present model does not satisfy regularity conditions found in the literature because it is unbounded below and does not satisfy standard growth conditions. To avoid these purely technical considerations, we simply restrict attention to the class of FO-IC contracts.

[57] PPI initially states a weaker version of [**FO-IC**] requiring only that $\gamma_t + p_t \geqslant 0$ (display (10) on p. 1244), which is the appropriate "one-sided" envelope formula when either NHB ($m \leqslant 0$) or IML ($\Delta \leqslant 0$) is imposed. However, PPI's entire analysis of optimal contracts (§§5-8) is based on the version of [**FO-IC**] stated above. In §5, PPI writes (p. 1250): "We now assume that the incentive constraint . . . binds, so that $\gamma = -p$. We relax this condition in an example below and verify that it holds. However, we also conjecture that it holds more generally." PPI does not carry out the asserted verification, instead relying on $\gamma = -p$ throughout §§5-8. Nonetheless, at least in the special case of permanent shocks ($\lambda = 0$) this restriction is indeed without loss of optimality (Remark J.2 in Appendix J.2).

- **Claim 1:** *Contract PPI is the optimal FO-IC contract.*
- **Claim 2:** *Contract PPI is $\mathcal{M}_{\leqslant}^{LAC}$-IC.*
- **Claim 3:** *The optimal $\mathcal{M}_{\leqslant}^{LAC}$-IC contract is FO-IC.*

Below, we identify issues with each of these claims and where they arise in PPI.[58]

PPI's Argument for Claim 1. §5 of PPI presents an HJB equation for the principal's value function over [**FO-IC**] contracts in the infinite-horizon setting (see display (19) on p. 1249 and the first-order conditions in displays (24)–(25) on p. 1251). PPI then presents the main derivation of Contract PPI in §6.2.1 and §6.2.2 (pp. 1252-54) and §A.3.1 (pp. 1269-71). The main steps of that derivation are as follows:

Step 1. Conjecture that the principal's value function at time t—which in general can be written as a function $J(y_t, q_t, p_t)$—depends on marginal promised utility p_t (as defined in [**3.2**]) only through the ratio $k_t = p_t/q_t$. Also conjecture that the dependence on k_t is additively separable, viz., $J(y_t, q_t, p_t) = \hat{J}(y_t, q_t) + h(k_t)$ for some functions $\hat{J}$ and h (see display (26) on p. 1252).

Step 2. Assume that $h(\cdot)$ is smooth, plug the conjectured form of the value function from Step 1 into the principal's HJB equation. This delivers a second-order ODE for $h(\cdot)$ (see display (A.11) on p. 1269).

Step 3. Use the policy functions derived from the HJB equation, together with Itô's lemma, to conclude that the resulting k process must satisfy a particular law of motion (display (28) on p. 1253).

Step 4. Numerically solve the ODE for $h(\cdot)$ derived in Step 2, plot the solution for specific parameter values, graphically observe that $h(\cdot)$ appears to be minimized at the value $k_0^* = \theta\lambda/(\rho + \lambda)$, and conclude that k_0^* is the optimal initial condition for the k process (pp. 1270-71). (Recall that the initial conditions q_0 and $y_0 = b_0$ are given, while the principal's choice of FO-IC contract determines p_0.)

Step 5. Observe that, given the law of motion from Step 3 and the initial condition k_0^* from Step 4, the k process is necessarily constant. Conclude from this that the policy function from the HJB generates Contract PPI, and therefore that Contract PPI is the optimal [**FO-IC**] contract (see pp. 1253-54).

Issues with Claim 1. In Appendix J (see Remark J.10), we verify that Steps 1–3 and Step 5 above are correct as stated (while also showing that PPI's conjectures in Step 1 can be derived from first principles, and making explicit some technical assumptions that PPI makes implicitly). Meanwhile, Theorem 1 establishes the existence of an FO-IC

[58] Recall from Section 4.2 that PPI does not fully specify the agent's strategy space in the infinite-horizon setting of §§5–8, so we adopt our best understanding of PPI's assumptions, which is that the agent's feasible set is $F = \mathcal{M}_{\leqslant}^{LAC}$. Claim 1 and the issues that we identify with it are *independent of this convention*, as the class of FO-IC contracts is defined independently of the agent's feasible set.

contract that strictly dominates Contract PPI.[59] *It follows that PPI's conclusion in Step 4—that k_0^* is the optimal initial condition—is incorrect.* We also note that PPI also does not specify boundary conditions for the ODE for $h(\cdot)$ in Step 4; it is therefore unclear what ODE system the numerical solution in Step 4 pertains to.

Issues with Claims 2–3. These follow from discussion above:

- In Section 4.2, we describe issues with PPI's argument for Claim 2. We also show that Claim 2 is false by constructing a strategy $m \in \mathcal{M}_{\leqslant}^{\text{LAC}}$ that the agent strictly prefers to truthtelling (Observation 2).
- In Footnotes 56 and 57 above, we describe how PPI (i) initially presents a weaker "one-sided" version of [**FO-IC**], (ii) claims that it is satisfed by every $\mathcal{M}_{\leqslant}^{\text{LAC}}$-IC contract in the finite-horizon setting, (iii) restricts attention to FO-IC contracts (as defined here) when studying optimal contracts in the infinite-horizon setting, and (iv) does not actually verify that this restriction is without loss of optimality. Consequently, PPI does not establish Claim 3. We nonetheless conjecture that Claim 3 is true.

Implications for the Optimal FO-IC Contract. Through Steps 1–3 above, PPI derives two correct properties of the optimal FO-IC contract.[60] Taking b_0 and q_0 as given, let $k_0^\dagger \in \mathbb{R}_{++}$ denote the true optimal initial condition for the process $k_t := p_t/q_t$. First, PPI finds that the principal's optimal lifetime cost is

$$J(b_0, q_0, k_0^\dagger q_0) = J^*(b_0, q_0) + \frac{\sigma^2}{2\rho^2\theta}\left(k_0^\dagger\right)^2, \tag{A.2}$$

where $J^*(b_0, q_0)$ is the first-best value function (recall Footnote 35). Second, PPI finds that the agent's initial recommended consumption is

$$c_0^\dagger = \bar{c}(q_0, \rho), \tag{A.3}$$

where $\bar{c}(q_0, \rho) = -\log(-\rho q_0)/\theta$. Note that Contract PPI specifies the same initial condition for recommended consumption. Taken together, our results and [**A.2**]–[**A.3**] have the following implications for the optimal FO-IC contract:[61]

1. *The initial condition satisfies $k_0^\dagger < k_0^*$, where $k_0^* = \rho\theta/(\rho + \lambda)$ is the initial condition for Contract PPI.* This holds because the principal's cost of Contract PPI (stated in [**5.13**]) can be expressed as $J^*(b_0, q_0) + \frac{\sigma^2}{2\rho^2\theta}(k_0^*)^2$, our Theorem 1 and the fact that SI

[59] Theorem 2 shows that the contract we identify in Theorem 1 is IC in a suitable sense, but this is not relevant for the present discussion of FO-IC contracts.

[60] See the calculations on pp. 1269–71 in §A.3.1 of PPI, which assume that $h(\cdot)$ is twice continuously differentiable and satisfies $h''(k_0^\dagger) > 0$. We maintain these assumptions in the discussion here.

[61] We emphasize that our results, [**A.2**]–[**A.3**], and the implications below are mutually consistent because our DR-SICs correspond to the strict subset of FO-IC contracts that are "stationary" (viz., feature a constant $k_t = p_t/q_t$ process). See Section 5.4, and Lemma G.1 in Appendix G.

Contracts are FO-IC imply that this cost is strictly greater than $J(b_0, q_0, k_0^\dagger q_0)$, the $k_t = p_t/q_t$ process is strictly positive, and hence [A.2] implies $k_0^* > k_0^\dagger$.

2. *At* $t = 0$, *we have* $\mathrm{d}k_t = (\rho + \lambda)\big[k_0^\dagger - k_0^*\big]\mathrm{d}t < 0$. The expression for $\mathrm{d}k_0$ follows from substituting $h'(k_0^\dagger) = 0$ into the expressions for $\hat{c}(k_0^\dagger)$ and $\hat{Q}(k_0^\dagger)$ on pp. 1269–70 of PPI, and then substituting those expressions into display (28) on p. 1253 of PPI. The strict inequality follows from the first implication above.

Further analysis of the optimal FO-IC contract is an exciting direction for future research.

B. On the Agent's Reporting Problem in DR-SICs

When the agent's feasible set is $F \subseteq \mathscr{M}$, his reporting problem can be cast as one of stochastic control with states (A_t^v, y_t, m_t) and control Δ_t, as in PPI. For a DR-SIC with rate $r > 0$, $V^{\mathrm{NJ}}(A_t^v, y_t, m_t)$ denotes the agent's value function under the "no jump" feasible set $F = \mathscr{M}^r$; as in Section 6.4, $V^{\mathrm{DR}}(A_t^v, y_t - m_t)$ denotes that under the extended feasible set $F = \mathscr{M}^r_{\text{ext}}$ (where we have used the identity $b_t \equiv y_t - m_t$). We show that these two value functions coincide, i.e., the agent does not strictly benefit from having access to jump reports (as in the treatment of IC from Section 6) either on- or off-path. (In Appendix I, we directly analyze the "no jump" reporting problem.)

Lemma B.1. For any DR-SIC with rate $r > 0$, the agent's value functions satisfy

[B.1] $$V^{\mathrm{DR}}(A_t^v, y_t - m_t) \equiv V^{\mathrm{NJ}}(A_t^v, y_t, m_t) \equiv q_t \exp\left[f(r;\lambda)m_t\right].$$

Proof. By construction, $q_t \equiv V^{\mathrm{SI}}(A_t^v, y_t)$. Thus, [5.10], [6.8], and the identity $m_t \equiv y_t - b_t$ yield $V^{\mathrm{DR}}(A_t^v, y_t - m_t) \equiv q_t \exp\left[f(r;\lambda)m_t\right]$. Next, $\mathscr{M}^r \subsetneq \mathscr{M}^r_{\text{ext}}$ implies that $V^{\mathrm{NJ}}(A_t^v, y_t, m_t) \leqslant q_t \exp\left[f(r;\lambda)m_t\right]$. If $m_t = 0$, the upper bound is trivially attained; for $m_t \neq 0$, we show its attainment via an approximation argument. Let (A_t^v, y_t, m_t) with $m_t \neq 0$ be a state at time t in the agent's reporting problem with (non-extended) feasible set $\mathscr{M}^r$. For any continuation strategy $m \in \mathscr{M}^r$ and time $T > t$, integrating [A.1] (with $u_\tau \equiv rq_\tau$ and $\gamma_\tau \equiv -f(r;\lambda)q_\tau$) over $\tau \in [t, T]$ delivers

[B.2] $$q_T = \hat{q}_T \cdot \exp\left[f(r;\lambda)\left((m_t - m_T) + \lambda m_t(T - t) - \lambda \int_t^T m_\tau \mathrm{d}\tau\right)\right],$$

where $\hat{q}_T$ is the counterfactual time-T promised utility that would have arisen under the same endowment shocks $(W_\tau)_{\tau\in[t,T]}$ had the agent set $\Delta_\tau \equiv 0$ on $\tau \in [t, T]$.[62]

[62]Fix the path $(m_\tau)_{\tau\in[0,t)}$. Under the continuation strategy m, the "inferred shock process" satisfies $\sigma W_T^y = \sigma W^T + m_T + \lambda \int_0^T m_\tau \mathrm{d}\tau$. Under the continuation strategy $\Delta_\tau \equiv 0$ for all $\tau \in [t, T]$, the corresponding process, call it $\hat{W}^y$, satisfies $\sigma \hat{W}_T^y = \sigma W^T + m_t + \lambda \int_0^t m_\tau \mathrm{d}\tau + \lambda m_t(T - t)$. Integrating [A.1] (with $u_\tau \equiv rq_\tau$ and $\gamma_\tau \equiv -f(r;\lambda)q_\tau$) from the fixed initial condition $q_t = V^{\mathrm{SI}}(A_t^v, y_t - m_t)$ over $\tau \in [t, T]$ under both strategies yields $q_T = \hat{q}_T \exp\left[f(r;\lambda)\sigma\left(\hat{W}_T^y - W_T^y\right)\right]$, which reduces to [B.2].

For each $T > t$, consider the specific continuation strategy $m^T \in \mathcal{M}^r$ induced by $\Delta_\tau^T := -m_t/(T-t) \cdot \mathbf{1}(\tau \in [t,T])$. Under m^T, [**B.2**] reduces to

$$q_T = \hat{q}_T \cdot \exp\left[f(r;\lambda)m_t\right] \cdot \exp\left[-\lambda f(r;\lambda)m_t \cdot \frac{(T-t)}{2}\right]. \tag{B.3}$$

As $T \searrow t$, m^T approximates an instantaneous jump to $m_{t+} = 0$ and truthful reporting for all $\tau > t$ (the optimal continuation strategy from the extended problem in Section 6.4); we have $\hat{q}_T \to \hat{q}_t = q_t$ and thus $q_T \to q_t \exp\left[f(r;\lambda)m_t\right]$ **P**-a.s. Furthermore, because m^T specifies truthful reporting for $\tau \geqslant T$, q_T is also the agent's true time-T continuation utility. It follows that $V^{\mathrm{NJ}}(A_t^v, y_t, m_t) \geqslant q_t \exp\left[f(r;\lambda)m_t\right]$, as desired. □

Remark B.2. Lemma B.1 and its proof have a few notable implications:

(i) Recommended consumption satisfies $c_t \equiv \bar{c}(q_t, r)$, so sending $T \searrow t$ in [**B.3**] also implies that the agent's recommended and actual consumption (hence, also the transfers s_t) converge to those under the optimal extended strategy from Section 6.4.

(ii) Because $V^{\mathrm{NJ}} = V^{\mathrm{DR}}$ and the latter value function is uniquely attained by the optimal "report truthfully at all histories" strategy from Section 6.4—which is *not* in $\mathcal{M}^r$—it follows that *when the agent's feasible set is* $\mathcal{M}^r$*, he does not have a well-defined optimal continuation strategy at off-path histories where* $m_t \neq 0$. At such histories, the only way to attain continuation value $V^{\mathrm{DR}}(A_t^v, y_t - m_t)$ is to *approximate* a jump report back to $m_t = 0$, as in the above proof.

(iii) By restricting attention to $m_t \leqslant 0$, the arguments here and in Section 6.4 apply almost verbatim when NHB is imposed, i.e., the agent's feasible set is $\mathcal{M}_-^r$. Given this feasible set, the agent's value function is therefore $V^{\mathrm{NHB}}(q,m) = q \exp\left[f(r;\lambda)m\right]$ as in [**B.1**], but restricted to the domain where $m \leqslant 0$. As noted in Section 4.1, for Contract PPI ($r = \rho$), this differs from the agent's value function when the stronger [**IML**] constraint is imposed (cf. [**4.1**]).

References

Acciaio, Beatrice, Robert A. Crowell, and Jakša Cvitanić. 2022. "A Note on Persistent Private Information". Working paper.

Allen, Franklin. 1985. "Repeated Principal-Agent Relationships with Hidden Lending and Borrowing". *Economics Letters* 17 (1-2): 27–31.

Atkeson, Andrew, and Robert E. Lucas Jr. 1992. "On Efficient Distribution with Private Information". *Review of Economic Studies* 59 (3): 427–453.

Battaglini, Marco, and Rohit Lamba. 2019. "Optimal Dynamic Contracting: The First-Order Approach and Beyond". *Theoretical Economics* 14 (4): 1435–1482.

Bloedel, Alexander W., R. Vijay Krishna, and Oksana Leukhina. 2018. “Insurance and Inequality with Persistent Private Information”. Working paper, Stanford University and Federal Reserve Bank of St. Louis.

— . 2021. *Insurance and Inequality with Persistent Private Information*. Tech. rep. Federal Reserve Bank of St. Louis.

Bloedel, Alexander W., R. Vijay Krishna, and Bruno Strulovici. 2022. “On the Limits of Dynamic Screening under Permanent Shocks”.

Caballero, Ricardo J. 1990. “Consumption Puzzles and Precautionary Savings”. *Journal of Monetary Economics* 25 (1): 113–136.

— . 1991. “Earnings Uncertainty and Aggregate Wealth Accumulation”. *American Economic Review*: 859–871.

Chamberlain, Gary, and Charles A. Wilson. 2000. “Optimal Intertemporal Consumption under Uncertainty”. *Review of Economic Dynamics* 3 (3): 365–395.

Chen, Yi. 2021. “Dynamic Delegation with a Persistent State”. *Theoretical Economics*, no. forthcoming.

Cisternas, Gonzalo. 2017. “Two-sided Learning and the Ratchet Principle”. *Review of Economic Studies* 85 (1): 307–351.

Cole, Harold L., and Narayana R. Kocherlakota. 2001. “Efficient Allocations with Hidden Income and Hidden Storage”. *Review of Economic Studies* 68 (3): 523–542.

Cvitanić, Jakša, and Jianfeng Zhang. 2012. *Contract Theory in Continuous-Time Models*. Springer Science & Business Media.

DeMarzo, Peter M., and Yuliy Sannikov. 2016. “Learning, termination, and payout policy in dynamic incentive contracts”. *Review of Economic Studies* 84 (1): 182–236.

DeMarzo, Peter M, and Yuliy Sannikov. 2006. “Optimal security design and dynamic capital structure in a continuous-time agency model”. *The journal of Finance* 61 (6): 2681–2724.

Doepke, Matthias, and Robert M Townsend. 2006. “Dynamic Mechanism Design with Hidden Income and Hidden Actions”. *Journal of Economic Theory* 126 (1): 235 –285.

Edlin, Aaron S., and Chris Shannon. 1998. “Strict monotonicity in comparative statics”. *Journal of Economic Theory* 81 (1): 201–219.

Farhi, Emmanuel, and Iván Werning. 2013. “Insurance and Taxation over the Lifecycle”. *Review of Economic Studies* 80 (2): 596–635.

Fernandes, Ana, and Christopher Phelan. 2000. “A Recursive Formulation for Repeated Agency with History Dependence”. *Journal of Economic Theory* 91:223–247.

Golosov, Mikhail, Narayana R. Kocherlakota, and Aleh Tsyvinski. 2003. “Optimal Indirect and Capital Taxation”. *Review of Economic Studies* 70 (3): 569–587.

Golosov, Mikhail, Aleh Tsyvinski, and Nicolas Werquin. 2016. “Recursive Contracts and Endogenously Incomplete Markets”. In *Handbook of Macroeconomics, Volume 2*.

Green, Edward J. 1987. “Lending and the Smoothing of Uninsurable Income”. In *Contractual Agreements for Intertemporal Trade*, ed. by Edward C. Prescott and Neil Wallace. University of Minnesota Press.

He, Zhiguo, Bin Wei, Jianfeng Yu, and Feng Gao. 2017. “Optimal Long-Term Contracting with Learning”. *Review of Financial Studies* 30 (6): 2006–2065.

Kapička, Marek. 2013. “Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach”. *The Review of Economic Studies*: rds045.

Karatzas, Ioannis, and Steven E. Shreve. 1998. *Brownian Motion and Stochastic Calculus*. Second Edition. New York, NY: Springer.

Katzourakis, Nikos. 2015. *An introduction to viscosity solutions for fully nonlinear PDE with applications to calculus of variations in L*. Springer.

Kocherlakota, Narayana R. 2010. *The New Dynamic Public Finance*. Princeton, NJ: Princeton University Press.

Ljunqvist, Lars, and Thomas J. Sargent. 2000. *Recursive Macroeconomic Theory*. 1st ed. Cambridge, Ma: MIT Press.

Mörters, Peter, and Yuval Peres. 2010. *Brownian Motion*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

Oksendal, Bernt, and Agnes Sulem. 2019. *Applied Stochastic Control of Jump Diffusions*. Third. Springer.

Pavan, Alessandro, Ilya Segal, and Juuso Toikka. 2014. “Dynamic Mechanism Design: A Myersonian Approach”. *Econometrica* 82 (2): 601–653.

Pham, Huyên. 2009. *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Berlin: Springer-Verlag.

Phelan, Christopher. 1998. “On the Long Run Implications of Repeated Moral Hazard”. *Journal of Economic Theory* 79:174–191.

Prat, Julien, and Boyan Jovanovic. 2014. “Dynamic contracts when the agent’s quality is unknown”. *Theoretical Economics* 9 (3): 865 –914.

Ramos, Joao, and Tomasz Sadzik. 2019. “Partnership with Persistence”. *Available at SSRN 3308865*.

Sadzik, Tomasz, and Ennio Stacchetti. 2015. “Agency Models with Frequent Actions”. *Econometrica* 83 (1): 193–237.

Sannikov, Yuliy. 2014. *Moral hazard and long-run incentives*. Tech. rep. Princeton University.

Sotomayor, Marilda. 1984. “On Income Fluctuations and Capital Gains”. *Journal of Economic Theory* 32 (1): 14–35.

Stokey, Nancy L., Robert E. Lucas Jr, and Edward C. Prescott. 1989. *Recursive Methods in Economic Dynamics*. Cambridge, Ma: Harvard University Press.

Strulovici, Bruno. 2022. *Renegotiation-Proof Contracts with Persistent States*. Tech. rep. Northwestern University.

—— . 2011. *Renegotiation-proof Contracts with Moral Hazard and Persistent Private Information*. Tech. rep. Center for Mathematical Studies in Economics and Management Science.

Thomas, Jonathan P., and Tim Worrall. 1990. “Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem”. *Journal of Economic Theory* 51:367–390.

Touzi, Nizar. 2018. “Stochastic Control and Application to Finance”. Manuscript, Ecole Polytechnique Paris.

Wang, Neng. 2003. “Caballero meets Bewley: The Permanent-Income Hypothesis in General Equilibrium”. *American Economic Review* 93 (3): 927–936.

—— . 2006. “Generalizing the permanent-income hypothesis: Revisiting Friedman’s conjecture on consumption”. *Journal of Monetary Economics* 53 (4): 737–752.

—— . 2004. “Precautionary saving and partially observed income”. *Journal of Monetary Economics* 51 (8): 1645–1681.

Williams, Noah. 2011. “Persistent Private Information”. *Econometrica* 79 (4): 1233–1275.

Yong, Jiongmin, and Xun Yu Zhou. 1999. *Stochastic Controls: Hamiltonian Systems and HJB Equations*. Vol. 43. Springer Science & Business Media.

Zhang, Yuzhe. 2009. “Dynamic Contracting with Persistent Shocks”. *Journal of Economic Theory* 144:635–675.

Online Appendix
to *Persistent Private Information Revisited*

Appendices C–K present technical proofs and secondary results omitted from the main paper. Most of this material is presented in the same order that it is mentioned in the main text; Appendix K collects auxiliary mathematical facts.

C. Facts about AC Change-of-Measure

Following PPI (p. 1239) and Karatzas and Shreve (1998, p. 59), we view $\mathbf{P}^*$ and $\mathbf{P}^m$ as probability measures on the space $\mathbf{C}[0,\infty)$ of continuous paths. There exists a *density process* $\Gamma_t^m \equiv \mathrm{d}\,\mathbf{P}_t^m/\mathrm{d}\,\mathbf{P}_t^*$ iff $m \in \mathcal{M}^{\text{LAC}}$. By Karatzas and Shreve (1998, p. 191): (a) Γ^m is a continuous $\mathbf{P}^*$-local martingale and can be expressed as[1]

$$\Gamma_t^m \equiv \exp\left[\int_0^t \frac{(\Delta_\tau + \lambda m_\tau)}{\sigma}\,\mathrm{d}W_\tau^y - \frac{1}{2}\int_0^t \frac{(\Delta_\tau + \lambda m_\tau)^2}{\sigma^2}\,\mathrm{d}\tau\right],$$

where W^y is a standard Brownian motion under $\mathbf{P}^*$;[2] (b) $m \in \mathcal{M}^{\text{LAC}}$ iff Γ^m is a $\mathbf{P}^*$-martingale; and (c) $m \in \mathcal{M}^{\text{GAC}}$ iff Γ^m is a uniformly integrable (UI) $\mathbf{P}^*$-martingale, in which case there exists an infinite-horizon density $\Gamma_\infty^m := \mathrm{d}\,\mathbf{P}^m/\mathrm{d}\,\mathbf{P}^*$ and $\Gamma_t^m \to \Gamma_\infty^m$ $\mathbf{P}^*$-a.s. These facts imply that the strategies constructed in the proofs of Observations 1–2 are in the claimed feasible sets. Furthermore, as claimed in Sections 2.2 and 4.2:

Fact 3. The following hold:

(i) If $\lambda > 0$, every $m \in \mathcal{M}_{\leqslant}^{\text{GAC}}$ satisfies $m_t \equiv 0$.

(ii) $\mathcal{M}_{\dagger}^{\text{GAC}} \subseteq \mathcal{M}_{\dagger}^{\text{LAC}} \cap \mathcal{M}_{\dagger}^{r}$ for all $\dagger \in \{\leqslant, -\}$ and $r > 0$.

Proof. Part (i): Suppose there exists $m \in \mathcal{M}_{\leqslant}^{\text{GAC}} \backslash \{m^* \equiv 0\}$. Then $\Gamma_t^m \equiv \exp(X_t - \frac{1}{2}\langle X\rangle_t)$ is a UI $\mathbf{P}^*$-martingale, where $X_t := \int_0^t \frac{(\Delta_\tau+\lambda m_\tau)}{\sigma}\,\mathrm{d}W_\tau^y$ is a martingale and $\langle X\rangle_t = \int_0^t \frac{(\Delta_\tau+\lambda m_\tau)^2}{\sigma^2}\,\mathrm{d}\tau$ is its quadratic variation. On the event $\{m \not\equiv 0\}$, we have $\lim_{t\to\infty}\langle X\rangle_t = \infty$ because m is nonincreasing. Thus, there exists a time-changed Brownian motion B such that $X_t = B_{\langle X\rangle_t}$ (see Theorem 4.6 and Problem 4.7 in Karatzas and Shreve (1998, Ch. 3)). The SLLN for Brownian motion implies that $\lim_{t\to\infty} X_t/\langle X\rangle_t = 0$, so on the event $\{m \not\equiv 0\}$ we have $\Gamma_\infty^m = \lim_{t\to\infty}\exp\left[\langle X\rangle_t\left(X_t/\langle X\rangle_t - \frac{1}{2}\right)\right] = 0$. Hence, $\Gamma_\infty^m \lneq 1$ on $\mathbf{C}([0,\infty))$ and therefore $\mathbf{P}^m(\mathbf{C}([0,\infty))) < 1$, a contradiction.

Part (ii): Trivially, $\mathcal{M}_{\dagger}^{\text{GAC}} \subseteq \mathcal{M}_{\dagger}^{\text{LAC}}$. The inclusions $\mathcal{M}_{\dagger}^{\text{GAC}} \subseteq \mathcal{M}_{\dagger}^{r}$ follow from y being a $\mathbf{P}^*$-OU process, Lemma K.3(iii) in Appendix K, and the definitions [**GAC**] and [**NP-*m***]. $\square$

[1] This formula presumes that $\int_0^t \frac{(\Delta_\tau+\lambda m_\tau)}{\sigma}\,\mathrm{d}W_\tau^y$ is well-defined (e.g., $\Delta + \lambda m \in L^2_{\text{loc}}$).

[2] As in Section 3, $\mathrm{d}W_t^y := \frac{1}{\sigma}[\mathrm{d}y_t - (\mu - \lambda y_t)\mathrm{d}t]$, which is a $\mathbf{P}^*$-Brownian motion by definition of $\mathbf{P}^*$. When using the "weak formulation" of the agent's reporting problem to conduct change-of-measure (see Cvitanić and Zhang (2012)), PPI denotes W^y by W^* and the process W that drives b by $W_t^\Delta \equiv W_t^* - \int_0^t \left(\frac{\Delta_\tau+\lambda m_\tau}{\sigma}\right)\mathrm{d}\tau$. We maintain the (equivalent) notation W^y and W for simplicity.

D. PPI's Sufficient Conditions for IC

Theorem 4.1 in §4 of PPI presents sufficient conditions under which the first-order approach from §3 is valid in the finite-horizon setting from §2. That is, fixing a finite horizon $[0, T]$, PPI's Theorem 4.1 offers conditions under which we can conclude that a given contract is IC, assuming that (i) the agent's feasible set consists of all reporting strategies satisfying IML and finite-horizon AC change-of-measure ($\mathbf{P}_T^m \ll \mathbf{P}_T^*$), and (ii) the contract satisfies the equation [**A.1**] for q, the *downward FO-IC* condition $\gamma_t + p_t \geqslant 0$ (cf. Footnote 57), and the equation [**J.3**] for p (see also pp. 1244-45 in PPI). These conditions take the form of inequalities (displays (17) and (18) on p. 1247) to be satisfied by the y-adapted process $Q = (Q_t)_{t \in [0,T]}$ controlling the volatility of p in [**J.3**].

Issue 1: The proof of Theorem 4.1 requires IML. PPI presents the proof of Theorem 4.1 in §A.2 (pp. 1265-69). Three key steps of the proof rely on IML and would not go through as stated under NHB alone.

(a) On p. 1267, PPI writes: "when $Q_t \leqslant 0$, then we have both $Q_t m_t^2 \leqslant 0$ and $Q_t m_t \Delta_t \leqslant 0$." The final inequality requires that $m_t \Delta_t \geqslant 0$, which is implied by IML but can fail under NHB (which permits $m_t < 0 < \Delta_t$).

(b) Later on p. 1267, PPI writes: "Thus the optimality condition for truthtelling ($\Delta_t = 0$) is $Q_t m_t + \xi_t \geqslant 0$," where the inequality appears as display (A.9) and ξ is a co-state process in the agent's reporting problem. Direct inspection reveals that PPI's derivation of the inequality (A.9) from the preceding display (A.8) requires the IML constraint $\Delta_t \leqslant 0$. Without IML, to conclude that setting $\Delta_t = 0$ in (A.8) is optimal for the agent, one would need to strengthen (A.9) to the *equality* $Q_t m_t + \xi_t = 0$. Furthermore, truthful reporting corresponds to $m_t = 0$ rather than $\Delta_t = 0$; without IML, under many FO-IC contracts—including all DR-SICs—the agent will find it optimal to set $\Delta_t = +\infty$ in (A.8) when $m_t < 0$ (see Sections 6 and 7.3 and Appendix I).

(c) The final two displays on p. 1268 and the first display on p. 1269 invoke the inequality version of (A.9), and therefore also require IML. Without IML but maintaining the *equality* version of (A.9) and PPI's hypothesis that $m_\tau = m_t$ for all $\tau \geqslant t$ (middle of p. 1268), the final two displays on p. 1268, in particular (A.10), would need to hold as *equalities*. This is problematic because the inequality at the top of p. 1269, which yields (17) in the statement of Theorem 4.1, is not sufficient to ensure that the *equality* version of (A.10) holds. Furthermore, as noted in Point (b) above, without IML the equality version of (A.9) and the hypothesis that $m_\tau = m_t$ for all $\tau \geqslant t$ are both too demanding, so a different argument would be needed even if the equality version of (A.10) could be addressed.

In summary, the proof of Theorem 4.1 would not go through as stated if IML were weakened to NHB. We do not know whether the statement of Theorem 4.1 would remain valid, but conjecture that it would not.

Issue 2: The application of Theorem 4.1 to Contract PPI is incomplete. In §A.3.2 (p. 1271), PPI attempts to apply Theorem 4.1 to prove that Contract PPI is $\mathscr{M}^{\text{LAC}}_{\leqslant}$-IC when $\lambda > 0$. This conclusion is false (Observation 2). There are two issues with PPI's argument:

(a) *Theorem 4.1 pertains to the finite-horizon model, while Contract PPI arises in the infinite-horizon model.* PPI attempts to apply *an infinite-horizon variant of* Theorem 4.1 in which the key inequalities for the Q process (displays (17) and (18) on p. 1247) are modified by replacing the terminal time $T < \infty$ with $T = \infty$ wherever the former appears. PPI does not formally state or prove such a variant of Theorem 4.1.

(b) *The proof of Theorem 4.1, as stated, does not extend to the infinite-horizon model unless a sufficiently tight lower bound is imposed on the agent's misreports. In particular, it does not extend when the agent's infinite-horizon feasible set is $\mathscr{M}^{LAC}_{\leqslant}$.* To illustrate, we specialize to Contract PPI in the hidden endowment model. Following PPI's derivation of (A.4) on p. 1265, the agent's lifetime utility gain $V(m) - q_0$ from using strategy m instead of truthfully reporting is

[**D.1**]
$$V(m) - q_0 = \mathbf{E}_0^m \left[\int_0^T e^{-\rho t} \left[u(c_t - m_t) - u(c_t) \right] \mathrm{d}t + \int_0^T e^{-\rho t} \gamma_t \mathrm{d}W_t^y \right] + e^{-\rho T} \mathbf{E}_0^m \left[\int_T^\infty e^{-\rho(t-T)} u(c_t - m_t) \mathrm{d}t - q_T \right],$$

where [**D.1**] is an accounting identity that holds for *all* $T > 0$.[3] For a given $T > 0$, *if the strategy* m *satisfies* $m_t = m_T$ *for all* $t \geqslant T$, then by the same logic as PPI's (A.5) on p. 1266, the second line of [**D.1**] is bounded above by $e^{-\rho T} \mathbf{E}_0^m \left[p_T m_T \right]$.[4] This yields the overall bound: for all $T > 0$,

[**D.2**]
$$V(m) - q_0 \leqslant K(s, m, T) + e^{-\rho T} \mathbf{E}_0^m \left[p_T m_T \right],$$

where $K(s, m, T)$ denotes the first expectation in [**D.1**]. After using PPI's (A.4) to substitute out $e^{-\rho T} p_T m_T$, [**D.2**] reduces to PPI's (A.6), which is the basis for the rest of PPI's (finite-horizon) proof of Theorem 4.1. To adapt those later proof steps to derive the infinite-horizon variant of Theorem 4.1 that PPI applies to Contract PPI (viz., the infinite-horizon variant of (17) invoked on p. 1271), one must consider the

---

[3] [**D.1**] is the same as PPI's (A.4), specialized to the hidden endowment model and with time-T *continuation* payoffs in place of time-T *terminal* payoffs. As noted in Footnote 2, our W^y is PPI's W^*.

[4] PPI's (A.5) concerns $p_T m_T$ at the terminal time T. With an infinite horizon, let $U_T(m_T) := \mathbf{E}_T^{m_T} \left[\int_T^\infty e^{-\rho(t-T)} u(c_t - m_T) \mathrm{d}t \right]$ denote the agent's time-T continuation payoff under the continuation strategy $m_t \equiv m_T$. Then $U_T(\cdot)$ is smooth and concave, with $U'(m_T) = -\mathbf{E}_T^{m_T} \left[\int_T^\infty e^{-(\rho+\lambda)(t-T)} u'(c_t - m_T) \mathrm{d}t \right]$. Concavity yields $U_T(m_T) - U_T(0) \leqslant m_T U_T'(0)$ and [**3.1**]–[**3.2**] imply that $U_T(0) = q_T$ and $U_T'(0) = p_T$. Taking expectations yields the desired upper bound.

$T \to \infty$ limit of [D.2]. This adaptation is meaningful only if[5]

$$\text{[D.3]} \qquad \lim_{T\to\infty} e^{-\rho T}\,\mathbf{E}_0^m\left[p_T m_T\right] < \infty,$$

for otherwise subsequent calculations would involve expectations that are either infinite or ill-defined. *However, there exist strategies in* $\mathscr{M}_{\leqslant}^{LAC}$ *that violate* [D.3]. *In particular, any deterministic strategy* $m \in \mathscr{M}_{\leqslant}^{LAC}$ *that, for some fixed time* $\hat{T} > 0$, *satisfies* $m_t = M < \underline{M} := -(\rho+\lambda)/(\theta\lambda)$ *for all* $t \geqslant \hat{T}$ *violates* [D.3].[6] To see this, recall that $p_t \equiv k_0^* q_t$ under Contract PPI (where $k_0^* = \frac{\rho\theta}{\rho+\lambda}$), and use [3.3] and the identity $W_t^y \equiv W_t + \int_0^t \left(\frac{\lambda m_\tau + \Delta_\tau}{\sigma}\right) \mathrm{d}\tau$ to write

$$p_T m_T = k_0^* q_T^* m_T \exp\left[-k_0^*\left(m_T + \lambda\int_0^T m_t \mathrm{d}t\right)\right],$$

where $q_T^* := q_0 \exp\left(-\frac{1}{2}(k_0^*\sigma)^2 T - k_0^*\sigma W_T\right)$ is promised utility *under truthful reporting*. Thus, we have

$$\begin{aligned}
\lim_{T\to\infty} e^{-\rho T}\,\mathbf{E}_0^m\left[p_T m_T\right] &= p_0 M \cdot \lim_{T\to\infty} \exp\left[-\rho T - k_0^*\left(M + \lambda M(T-\hat{T}) + \lambda\int_0^{\hat{T}} m_t \mathrm{d}t\right)\right] \\
&= p_0 M e^{k_0^*\left(\lambda M(\hat{T}-1) - \lambda\int_0^{\hat{T}} m_t \mathrm{d}t\right)} \cdot \lim_{T\to\infty} \exp\left[-T\cdot(\rho + k_0^*\lambda M)\right] \\
&= \infty,
\end{aligned}$$

where the first equality uses the facts that (a) m is deterministic and satisfies $m_t = M$ for all $t \geqslant \hat{T}$ and (b) $\mathbf{E}_0^m\left[k_0^* q_T^*\right] = k_0^* q_0 = p_0$ because q^* is a martingale, and the last equality uses the fact that $M < \underline{M}$ implies $\rho + k_0^*\lambda M < 0$. We conclude that PPI's argument for the infinite-horizon variant of Theorem 4.1 requires the additional constraint that $m_t \geqslant \underline{M}$ for all $t \geqslant 0$.

E. Proof of Lemma 5.1

Note that $\hat{c}$ satisfies [5.2] and [NP-A] iff it satisfies the intertemporal constraint

$$\text{[E.1]} \qquad \int_0^\infty e^{-rt}\hat{c}_t\,\mathrm{d}t \leqslant \int_0^\infty e^{-rt} b_t\,\mathrm{d}t + A_0 \quad \mathbf{P}\text{-a.s.}$$

Given any $\xi \in \mathbb{R}$, [E.1] implies the following properties: (i) $\hat{c} \in \mathscr{A}(A_0, b_0)$ iff $\hat{c} + r\xi \in \mathscr{A}(A_0 + \xi, b_0)$, and (ii) $\mathscr{A}(A_0, b_0 + \xi) = \mathscr{A}(A_0 + \xi/(r+\lambda), b_0)$. Property (i) is immediate.

---

[5] IML implies $p_T m_T > 0$ for all $T \geqslant 0$.

[6] One such strategy is $m_t \equiv \max\{M, Mt/\hat{T}\}$; all such strategies, being eventually constant, satisfy the hypothesis used above to derive [D.2] from [D.1]. Also, note the parallel to Section 4.2: $\underline{M}$ is the same constant defined below [4.1] and the strategies considered here violate [TVC].

Property (ii) follows from inserting the closed-form solution for b_t (see Footnote 10) into [E.1]. Using these facts, we can characterize the agent's value function V^{SI} from [5.1] up to a parameter, $\gamma \in \mathbb{R}$, to be determined later.

Lemma E.1. Let $\xi \in \mathbb{R}$. The value function $V^{\mathrm{SI}} : \mathbb{R}^2 \to \mathbb{R}_{--}$ satisfies:

(i) $V^{\mathrm{SI}}(A_0 + \xi, b_0) = e^{-\theta r \xi} V^{\mathrm{SI}}(A_0, b_0)$.

(ii) $V^{\mathrm{SI}}(A_0, b_0 + \xi) = e^{-f(r;\lambda)\xi} V^{\mathrm{SI}}(A_0, b_0)$.

(iii) $V^{\mathrm{SI}}(A_0, b_0) = -\exp\big(-\theta r(A_0 + b_0/(r+\lambda) + \gamma)\big)$, where $-e^{-\theta r \gamma} := V^{\mathrm{SI}}(0,0)$.

Proof. Fix $(A_0, b_0) \in \mathbb{R}^2$. Note that $V^{\mathrm{SI}}(A_0, b_0) \in \mathbb{R}_{--}$ is well-defined: $V^{\mathrm{SI}}(A_0, b_0) < 0$ because $u(\cdot) < 0$ and [E.1] renders infeasible consumption processes approximating $\hat{c}_t \equiv +\infty$, and $V^{\mathrm{SI}}(A_0, b_0) > -\infty$ because there exist $\hat{c} \in \mathscr{A}(A_0, b_0)$ that deliver finite lifetime utility to the agent.[7] Point (i) follows from property (i) above, point (ii) follows from property (ii) above and point (i), and point (iii) follows from points (i) and (ii). □

Lemma E.1(iii) implies that $V^{\mathrm{SI}} \in \mathbf{C}^\infty(\mathbb{R}^2)$. Thus, letting (A_t, b_t) denote a generic state, standard arguments imply that V^{SI} is a classical solution to the HJB equation[8]

$$\rho V^{\mathrm{SI}}(A_t, b_t) = \sup_{c_t \in \mathbb{R}} \left[u(c_t) + (rA_t + b_t - c_t) V_A^{\mathrm{SI}}(A_t, b_t)\right] + (\mu - \lambda b_t) V_b^{\mathrm{SI}}(A_t, b_t) + \tfrac{1}{2}\sigma^2 V_{bb}^{\mathrm{SI}}(A_t, b_t). \tag{E.2}$$

This allows us to determine γ and the optimal (Markovian) control.

Lemma E.2. The following hold:

(i) The parameter $\gamma \in \mathbb{R}$ from Lemma E.1(iii) is

$$\gamma = \frac{\mu}{r(\lambda + r)} + \frac{\log(r)}{r\theta} - \left[\frac{r-\rho}{\theta r^2} + \frac{1}{2}\frac{(f(r;\lambda)\sigma)^2}{\theta r^2}\right]. \tag{E.3}$$

(ii) The supremum in [E.2] at state (A_t, b_t) is uniquely attained by

$$\hat{C}(A_t, b_t) := rA_t + \frac{r}{\lambda + r} b_t + \frac{\mu}{\lambda + r} - \left[\frac{r-\rho}{\theta r} + \frac{1}{2}\frac{(f(r;\lambda)\sigma)^2}{\theta r}\right]. \tag{E.4}$$

(iii) At every state (A_t, b_t), we have $u(\hat{C}(A_t, b_t)) = rV^{\mathrm{SI}}(A_t, b_t)$.

Proof. The FOC for [E.2] in state (A_t, b_t) is $u'(c_t) = V_A^{\mathrm{SI}}(A_t, b_t)$. Because u is exponential

[7] For one example, see the consumption process [5.3] constructed below.

[8] For instance, see Yong and Zhou (1999, Theorem 3.3) or Touzi (2018, Propositions 2.4-2.5).

and V^{SI} satisfies Lemma E.1(iii), it follows that $u(c_t) = rV^{\mathrm{SI}}(A_t, b_t)$. Equivalently,

[E.5]
$$c_t = rA_t + \frac{r}{\lambda + r} b_t + r\gamma - \frac{\log r}{\theta}.$$

Substituting [**E.5**] into [**E.2**] and solving for γ yields [**E.3**], hence point (i). Substituting [**E.3**] into [**E.5**] yields [**E.4**], hence point (ii). Point (iii) follows from the above observation that $u(c_t) = rV^{\mathrm{SI}}(A_t, b_t)$ and point (ii). □

Next, plugging [**E.5**] into the flow constraint [**5.2**] yields the following:

Lemma E.3. Define the asset process A^* by

[E.6]
$$\mathrm{d}A_t^* = \left(\frac{r - \rho}{\theta r} + \frac{1}{2} \frac{\sigma^2 f(r;\lambda)^2}{\theta r} + \frac{\lambda b_t - \mu}{\lambda + r} \right) \mathrm{d}t.$$

The consumption process $\hat{c}^*$ defined by $\hat{c}_t^* := \hat{C}(A_t^*, b_t)$ satisfies:

(i) Its induced asset process $A^{\hat{c}^*}$ satisfies $A^{\hat{c}^*} = A^*$.

(ii) It evolves as

[E.7]
$$\mathrm{d}\hat{c}_t^* = \left(\frac{r - \rho + \sigma^2 f(r;\lambda)^2/2}{\theta} \right) \mathrm{d}t + \frac{\sigma f(r;\lambda)}{\theta} \mathrm{d}W_t$$

(iii) It is feasible, i.e., is b-adapted and satisfies [**5.2**] and [**NP-A**].

Proof. Point (i) follows from plugging $\hat{c}_t^* = \hat{C}(A_t^*, b_t)$ into [**5.2**]. Point (ii) follows from noting that $\mathrm{d}\hat{c}_t^* = r\,\mathrm{d}A_t^* + \frac{r}{(r+\lambda)}\,\mathrm{d}b_t$ and plugging in [**E.6**] and [**2.1**]. For point (iii), only [**NP-A**] is nontrivial. Writing [**E.6**] in integrated form yields

[E.8]
$$A_t^* = A_0 + \left(\frac{r - \rho}{\theta r} + \frac{1}{2} \frac{\sigma^2 f(r;\lambda)^2}{\theta r} - \frac{\mu}{\lambda + r} \right) t + \frac{\lambda}{\lambda + r} \int_0^t b_\tau \,\mathrm{d}\tau.$$

When $\lambda = 0$, A^* is deterministic and affine in t, so $\lim_{t\to\infty} e^{-\alpha t} A_t^* = 0$ for every $\alpha > 0$. When $\lambda > 0$, the same conclusion holds because $\lim_{t\to\infty} e^{-\alpha t} \int_0^t b_\tau \,\mathrm{d}\tau = 0$ for all $\alpha > 0$ by Lemma K.3(iii) in Appendix K. □

Lemmas E.2 and E.3 immediately yield:

Corollary E.4. The strategy $\hat{c}^*$ from Lemma E.3 satisfies $u(\hat{c}_t^*) \equiv rV^{\mathrm{SI}}(A_t^{\hat{c}^*}, b_t)$.

The next two lemmas are useful technical facts:

Lemma E.5. Under $\hat{c}^*$ from Lemma E.3, $e^{(r-\rho)t} u'(\hat{c}_t^*)$, $e^{(r-\rho)t} u(\hat{c}_t^*)$, and $e^{(r-\rho)t} V^{\mathrm{SI}}(A_t^*, b_t)$ are martingales.

Proof. Note that $e^{(r-\rho)t}u(\hat{c}_t^*) = -e^{-\theta\hat{c}_0^*}\exp\left(-\frac{1}{2}\sigma^2(f(r;\lambda))^2 - \sigma f(r;\lambda)W_t\right)$ is a martingale, $u'(c) = -\theta u(c)$ by CARA, and $V^{\mathrm{SI}}(A_t^*, b_t) = \frac{1}{r}u(\hat{c}_t^*)$ by Corollary E.4. ☐

Lemma E.6. Under A^* from [**E.6**], $M_t := \int_0^t e^{-\rho\tau}V_b^{\mathrm{SI}}(A_\tau^*, b_\tau)\sigma\,\mathrm{d}W_\tau$ is a martingale.

Proof. It suffices to show $\mathbf{E}_0\left[\int_0^T\left(e^{-\rho t}V_b^{\mathrm{SI}}(A_t^*, b_t)\sigma\right)^2\mathrm{d}t\right] < \infty$ for all $T > 0$. Lemma E.1(iii) implies that $V_b^{\mathrm{SI}}(A_t^*, b_t) \equiv \frac{(-r\theta)}{r+\lambda}V^{\mathrm{SI}}(A_t^*, b_t)$ and Lemma E.2(i) implies that $u(\hat{c}_t^*) = rV^{\mathrm{SI}}(A_t^*, b_t)$. Thus, by Fubini's Theorem, it suffices to show that $\int_0^T e^{-2\rho t}\,\mathbf{E}_0\left[u(\hat{c}_t^*)^2\right]\mathrm{d}t < \infty$. To that end, Lemma E.3(ii) implies that

$$u(\hat{c}_t^*)^2 = \exp(-2\theta\hat{c}_0^*)\exp(-2(r-\rho)t + \sigma^2 f^2(r;\lambda)t)\exp(-2\sigma^2 f^2(r;\lambda)t - 2\sigma f(r;\lambda)W_t).$$

Because $\exp(-2\sigma^2 f^2(r;\lambda)t - 2\sigma f(r;\lambda)W_t)$ is a martingale, we have $e^{-2\rho t}\,\mathbf{E}_0\left[u(\hat{c}_t^*)^2\right] = \exp(-2\theta\hat{c}_0^*)\exp(-2rt + \sigma^2 f^2(r;\lambda)t)$. Thus, $\int_0^T e^{-2\rho t}\,\mathbf{E}_0\left[u(\hat{c}_t^*)^2\right]\mathrm{d}t < \infty$, as desired. ☐

We now are in a position to prove Lemma 5.1 itself:

Proof of Lemma 5.1. We show that, given initial condition (A_0, b_0), the strategy $\hat{c}^*$ from Lemma E.3 attains lifetime utility $V^{\mathrm{SI}}(A_0, b_0)$. The Itô expansion of $e^{-\rho t}V^{\mathrm{SI}}(A_t^*, b_t)$ is

$$\text{[E.9]}\quad \begin{aligned} e^{-\rho T}V^{\mathrm{SI}}(A_t^*, b_t) = V^{\mathrm{SI}}(A_0, b_0) &+ \int_0^T e^{-\rho t}\left[\mathscr{L}^{\hat{c}^*}V^{\mathrm{SI}}(A_t^*, b_t) - \rho V^{\mathrm{SI}}(A_t^*, b_t)\right]\mathrm{d}t \\ &+ \int_0^T e^{-\rho t}V_b^{\mathrm{SI}}(A_t^*, b_t)\sigma\,\mathrm{d}W_t, \end{aligned}$$

where for any $v \in \mathbf{C}^2(\mathbb{R}^2)$, we let $\mathscr{L}^{\hat{c}^*}v(A,b) := (rA + b - \hat{c}^*)\partial_A v(A,b) + \lambda(\mu/\lambda - b)\partial_b v(A,b) + \frac{1}{2}\sigma^2\partial_{bb}v(A,b)$. Lemma E.2(ii) implies that $\mathscr{L}^{\hat{c}^*}V^{\mathrm{SI}}(A_t^*, b_t) - \rho V^{\mathrm{SI}}(A_t^*, b_t) = -u(\hat{c}_t^*)$. Substituting this into [**E.9**] and applying Lemma E.6 yields

$$\text{[E.10]}\quad V^{\mathrm{SI}}(A_0, b_0) = \mathbf{E}_0\left[\int_0^T e^{-\rho t}u(\hat{c}_t^*)\,\mathrm{d}t\right] + \mathbf{E}_0\left[e^{-\rho T}V^{\mathrm{SI}}(A_t^*, b_t)\right].$$

Lemma E.5 shows that $e^{(r-\rho)t}V^{\mathrm{SI}}(A_t^*, b_t)$ is a martingale. Therefore,

$$\mathbf{E}_0\left[e^{-\rho T}V^{\mathrm{SI}}(A_t^*, b_t)\right] = e^{-rT}\,\mathbf{E}_0\left[e^{(r-\rho)T}V^{\mathrm{SI}}(A_t^*, b_t)\right] = e^{-rT}V^{\mathrm{SI}}(A_0, b_0),$$

so that letting $T \to \infty$ in [**E.10**] yields $V^{\mathrm{SI}}(A_0, b_0) = \mathbf{E}_0\left[\int_0^\infty e^{-\rho t}u(\hat{c}_t^*)\,\mathrm{d}t\right]$, as desired. ☐

F. Proof of Lemma 5.3

We must calculate $A_0 - \mathbf{E}_0[\int_0^\infty e^{-\rho t}(\rho - r)A_t^* \,\mathrm{d}t]$, where $A_0 := A_0(b_0, q_0, r)$ and the process A^* is from [E.6]. We claim that

$$\text{[F.1]} \qquad A_0 - (\rho - r)\int_0^\infty e^{-\rho t} A_t^* \,\mathrm{d}t = \hat{\Pi}(b_0, q_0, r) + \frac{(r-\rho)\lambda\sigma}{\rho(r+\lambda)(\rho+\lambda)} \int_0^\infty e^{-\rho t} \,\mathrm{d}W_t,$$

where $\hat{\Pi}(b_0, q_0, r)$ denotes the expression for $\Pi(b_0, q_0, r)$ on the RHS of the first line of [5.12]. To see this, note that $\int_0^\infty e^{-\rho t} A_t^* \,\mathrm{d}t = \frac{A_0}{\rho} + \frac{1}{\rho}\int_0^\infty e^{-\rho t} \,\mathrm{d}A_t^*$ because $\mathrm{d}(e^{-\rho t}A_t^*) = -\rho e^{-\rho t} A_t^* \,\mathrm{d}t + e^{-\rho t}\,\mathrm{d}A_t^*$ and, by (the proof of) Lemma E.3, $\lim_{T\to\infty} e^{-\rho T} A_T^* = 0$. Plugging in [E.6], the expression for $\int_0^\infty e^{-\rho t} b_t \,\mathrm{d}t$ in Lemma K.4, and simplifying yields [F.1]. A similar calculation yields

$$\text{[F.2]} \qquad A_0 - (\rho - r)\int_0^\infty e^{-\rho t} A_t^* \,\mathrm{d}t = \int_0^\infty e^{-\rho t}(\hat{c}_t^* - b_t)\,\mathrm{d}t.$$

Taking expectations in [F.1]–[F.2] and noting that $\mathbf{E}_0[\int_0^\infty e^{-\rho t}\,\mathrm{d}W_t] = 0$ yields [5.12].[9] □

G. Further Properties of SI Contracts

This appendix develops the properties of SI Contract described in Section 5.4.

Stationarity & State-Consistency. We say that a (direct-revelation) contract is *Stationary* if (i) it is FO-IC and (ii) $k_t := p_t/q_t \equiv k_0$ for some constant $k_0 > 0$. Strulovici (2022) defines the class of "state-consistent" renegotiation-proof contracts; Theorem 1 therein states, in our terminology, that a contract is state-consistent iff it is Stationary.

Lemma G.1. Suppose that $\lambda > 0$. A contract is Stationary with constant $k_0 > 0$ if and only if it is a DR-SIC with rate $r > 0$ such that $k_0 = f(r;\lambda)$.

Proof. For the"if" direction, note that under a DR-SIC with rate $r > 0$, the agent's promised utility satisfies $q_t \equiv V^{\mathrm{SI}}(A_t^v, y_t)$ (where V^{SI} is defined in [5.10]) and hence

$$\text{[G.1]} \qquad \mathrm{d}q_t = (\rho - r)q_t \mathrm{d}t - f(r;\lambda)q_t\sigma \mathrm{d}W_t^y.$$

Furthermore, the Euler equation [5.7] and fact that $u'(c) = -\theta u(c)$ imply $p_t \equiv f(r;\lambda)q_t$. Thus, the contract satisfies [FO-IC] and $k_t \equiv f(r;\lambda)$. For the "only if" direction, fix a Stationary contract with $k_0 > 0$. By [A.1] and [FO-IC], promised utility satisfies

$$\text{[G.2]} \qquad \mathrm{d}q_t = (\rho q_t - u_t)\mathrm{d}t - k_0 q_t \sigma \mathrm{d}W_t^y.$$

[9] The process $M_t := \int_0^t e^{-r\tau}\,\mathrm{d}W_\tau$ is a uniformly integrable martingale because $\int_0^\infty (e^{-\alpha\tau})^2\,\mathrm{d}\tau < \infty$ (see Exercise 5.24 of Karatzas and Shreve (1998, p. 38)). Thus, $\mathbf{E}_0[\int_0^\infty e^{-\alpha t}\,\mathrm{d}W_t] = M_0 = 0$.

Comparing [G.1]–[G.2], it suffices to show that $u_t \equiv r(k_0;\lambda)q_t$, where $r(k_0;\lambda) := \frac{\lambda k_0}{\theta - k_0}$ is the inverse of $k_0(r) := f(r;\lambda)$. By the Martingale Representation Theorem, marginal promised utility [3.2] satisfies

$$\mathrm{d}p_t = [(\rho+\lambda)p_t - \theta u_t]\,\mathrm{d}t + Q_t\sigma \mathrm{d}W_t^y \tag{G.3}$$

for some y-adapted process Q. By the Stationarity hypothesis (that $p_t \equiv k_0 q_t$) and the unique decomposition property for Itô processes, the drift (volatility) of p in [G.3] must a.e. equal k_0 times the drift (volatility) of q in [G.2]. Equating the drifts and using the identity $p_t \equiv k_0 q_t$ yields $u_t \equiv r(k_0;\lambda)q_t$, as desired. □

Long-Run Properties. Under an SI Contract, we say that the agent *converges to misery* if $u_t, V_t \to -\infty$ $\mathbf{P}$-a.s. and say that he *converges to bliss* if $u_t, V_t \to 0$ $\mathbf{P}$-a.s., where $u_t := u(\hat{c}_t^*)$ and $V_t := V^{\mathrm{SI}}(A_t^{\hat{c}^*}, b_t)$.

Theorem 3. *Under the optimal SI Contract, the following hold:*

(i) For each $\sigma > 0$, there exists a $\overline{\lambda}(\sigma) > 0$ such that the agent converges to misery if $\lambda > \overline{\lambda}(\sigma)$ and to bliss if $\lambda \in [0, \overline{\lambda}(\sigma))$.[10]

(ii) For each $\lambda > 0$, there exists a $\overline{\sigma}(\lambda) \geqslant 0$ such that the agent converges to misery if $\sigma > \overline{\sigma}(\lambda)$ and to bliss if $\sigma \in (0, \overline{\sigma}(\lambda))$. Furthermore, $\overline{\sigma}(\lambda) = 0$ if and only if $\lambda \geqslant \rho$.

By Lemma 5.3, the principal's optimization over SI Contracts can be written as

$$\inf_{r>0}\left[-\frac{\log(r)}{\theta\rho} + \frac{r - \rho + \sigma^2 f(r;\lambda)^2/2}{\theta\rho^2}\right]. \tag{G.4}$$

Let $(\lambda,\sigma) \mapsto r^*(\lambda,\sigma)$ denote an arbitrary selection from the argmin correspondence of [G.4], which is nonempty. Let $k^*(\lambda,\sigma) := f(r^*(\lambda,\sigma),\lambda)$. We require two lemmas.

Lemma G.2. Under the optimal SI Contract:

(i) For each $\sigma > 0$, $k^*(\cdot,\sigma)$ is strictly decreasing, with $\lim_{\lambda\to 0} k^*(\lambda,\sigma) = k^*(0,\sigma) = \theta$ and $\lim_{\lambda\to\infty} k^*(\lambda,\sigma) = 0$.

(ii) For each $\lambda > 0$, $k^*(\lambda,\cdot)$ is strictly decreasing, with $\lim_{\sigma\to 0} k^*(\lambda,\sigma) = f(\rho;\lambda)$ and $\lim_{\sigma\to\infty} k^*(\lambda,\sigma) = 0$.

Proof. Point (i): Let $\sigma > 0$ be given. For each $\lambda > 0$, let $r(k_0;\lambda) := \frac{\lambda k_0}{\theta - k_0}$ denote the inverse of $k_0(r) := f(r;\lambda)$. Changing variables in [G.4] from r to k and noting that $\frac{\partial^2}{\partial k \partial \lambda} r(k;\lambda) > 0$, Edlin and Shannon (1998, Theorem 1) (for minimization problems) implies that $k^*(\cdot,\sigma)$ is strictly decreasing. Next, recall from the proof of Theorem 1 that $r^*(\lambda,\sigma)$ satisfies the FOC [5.14] and $0 < r^*(\lambda,\sigma) \leqslant \rho$ for all $\lambda \geqslant 0$. (Furthermore,

[10] For $\lambda = \overline{\lambda}(\sigma)$, u_t and V_t are transient, with $\liminf_{t\to\infty} u_t, V_t = -\infty$ and $\limsup_{t\to\infty} u_t, V_t = 0$.

$r^*(0,\sigma) = \rho$, which yields $k^*(0,\sigma) = \theta$.) Multiplying [5.14] through by $\rho r^*(\lambda,\theta) > 0$ and rearranging yields

$$\rho = r^*(\lambda,\sigma) + \sigma^2 \cdot \frac{\lambda\theta^2 r^*(\lambda,\sigma)^2}{(r^*(\lambda,\sigma)+\lambda)^3}. \tag{G.5}$$

Because $r^*(\cdot,\sigma)$ is bounded, the second term in [G.5] goes to 0 as $\lambda \to \infty$. This implies $\lim_{\lambda\to\infty} r^*(\lambda,\sigma) = \rho$, and hence $\lim_{\lambda\to\infty} k^*(\lambda,\sigma) = \lim_{\lambda\to\infty} f(\rho;\lambda) = 0$. Finally, we show that $\lim_{\lambda\to 0} r^*(\lambda,\sigma) = \rho$, which implies $\lim_{\lambda\to 0} k^*(\lambda,\sigma) = \lim_{\lambda\to 0} f(\rho;\lambda) = \theta$. To this end, notice that $\underline{r} := \liminf_{\lambda\to 0} r^*(\lambda,\sigma) > 0$ (if not, [G.4] would explode as $\lambda \to 0$, contradicting the finite upper bound from setting $r = \rho$). Consequently, we have

$$0 \leqslant \liminf_{\lambda\to 0} \frac{\lambda\theta^2 r^*(\lambda,\sigma)^2}{(r^*(\lambda,\sigma)+\lambda)^3} \leqslant \limsup_{\lambda\to 0} \frac{\lambda\theta^2 r^*(\lambda,\sigma)^2}{(r^*(\lambda,\sigma)+\lambda)^3} \leqslant \limsup_{\lambda\to 0} \frac{\lambda\theta^2 r^*(\lambda,\sigma)^2}{\underline{r}^3} = 0,$$

where the equality uses $r^*(\cdot,\sigma) \leqslant \rho$. Display [G.5] then yields $\lim_{\lambda\to 0} r^*(\lambda,\sigma) = \rho$.
Point (ii): Let $\lambda > 0$ be given. Because $\frac{\partial^2}{\partial r \partial\sigma}\left[\sigma^2 f(r;\lambda)\right] > 0$, Edlin and Shannon (1998, Theorem 1) (for minimization problems) applied to [G.4] implies that $r^*(\lambda,\cdot)$ is strictly decreasing, which further implies that $k^*(\lambda,\cdot)$ is strictly decreasing. The limit properties follow from calculations similar to those in point (i) above, which are omitted. □

Lemma G.3. Let $D^* : \mathbb{R}_+ \times \mathbb{R}_{++} \to \mathbb{R}$ be defined by

$$D^*(\lambda,\sigma) := r^*(\lambda,\sigma) - \rho + \frac{\sigma^2}{2} f\left(r^*(\lambda,\sigma);\lambda\right)^2. \tag{G.6}$$

Under the optimal SI Contract, the following hold:
(i) If $D^*(\lambda,\sigma) > 0$, then $\hat{c}_t^* \to \infty$ and $V_t, u_t \to 0$ $\mathbf{P}$-a.s.
(ii) If $D^*(\lambda,\sigma) < 0$, then $\hat{c}_t^* \to -\infty$ and $V_t, u_t \to -\infty$ $\mathbf{P}$-a.s.

Proof. By [5.3], $\hat{c}^*$ is a Brownian motion with drift $D^*(\lambda,\sigma)/\theta$. The long-run properties of $\hat{c}^*$ thus follow from the SLLN for Brownian motion (Lemma K.1 in Appendix K). The long-run properties of u_t and V_t then follow from CARA utility, [5.8], and the Continuous Mapping Theorem. □

Proof of Theorem 3. Recall from the proof of Theorem 1 that $r^*(\lambda,\sigma)$ satisfies the FOC [5.14], which as noted above is equivalent to [G.5]. Plugging [G.5] into [G.6] yields

$$D^*(\lambda,\sigma) = \underbrace{\sigma^2 f\left(r^*(\lambda,\sigma);\lambda\right)^2}_{>0} \cdot \left[\frac{1}{2} - \frac{\lambda}{r^*(\lambda,\sigma)+\lambda}\right]. \tag{G.7}$$

It follows from [G.7] that

$$D^*(\lambda,\sigma) > 0 \iff r^*(\lambda,\sigma) > \lambda \iff k^*(\lambda,\sigma) > \theta/2, \tag{G.8}$$

$$D^*(\lambda,\sigma) < 0 \iff r^*(\lambda,\sigma) < \lambda \iff k^*(\lambda,\sigma) < \theta/2. \tag{G.9}$$

Point (i): Let $\sigma > 0$ be given. By Lemma G.2(i), there exists a unique $\overline{\lambda}(\sigma) > 0$ such that $k^*(\lambda,\sigma) > \theta/2$ iff $\lambda \in [0, \overline{\lambda}(\sigma))$ and $k^*(\lambda,\sigma) < \theta/2$ iff $\lambda > \overline{\lambda}(\sigma)$. The result then follows from [G.8]–[G.9] and Lemma G.3.

Point (ii): Let $\lambda > 0$ be given. By Lemma G.2(ii), there exists a unique $\overline{\sigma}(\lambda) \geqslant 0$ such that $k^*(\lambda,\sigma) > \theta/2$ iff $\sigma \in (0, \overline{\sigma}(\lambda))$ and $k^*(\lambda,\sigma) < \theta/2$ iff $\sigma > \overline{\sigma}(\lambda)$. Furthermore, $\lim_{\sigma\to 0} k^*(\lambda,\sigma) = f(\rho;\lambda)$, and $f(\rho;\lambda) > \theta/2$ iff $\rho > \lambda$. Thus, $\overline{\sigma}(\lambda) > 0$ iff $\rho > \lambda$. The result then follows from [G.8]–[G.9] and Lemma G.3. $\square$

H. Hidden Savings

In this appendix, we consider the *hidden savings* variant of PPI's hidden endowment model in which (a) the agent can directly self-insure via the market at rate ρ, and (b) both the agent's endowment and trading activity are his private information (as in Allen 1985 and Cole and Kocherlakota 2001).

Given a contract s (as in Section 2.1) and a feasible set of (extended) misreporting strategies $F \subseteq \mathscr{M}_{\text{ext}}$ (as in Section 6.1), the agent chooses an $m \in F$ and a b-adapted consumption strategy $\hat{c}$ subject to the constraint that the induced b-adapted *asset process* $A^{m,\hat{c}}$ solves the equation

$$A_t^{m,\hat{c}} = \big(\rho A_t^{m,\hat{c}} + b_t + s_t - \hat{c}_t\big)\mathrm{d}t$$

and satisfies the no Ponzi condition [NP-A] at the market rate $r = \rho$. The agent's joint strategy $(m, \hat{c})$ is *optimal given contract* s if it maximizes his lifetime utility from the consumption process $\hat{c}$ among all strategies satisfying the above constraints.

A contract is *F-HS-IC* if it makes truthful reporting—i.e., some joint strategy $(m^* \equiv 0, \hat{c})$—optimal for the agent.[11] It is *F-NS-IC* it makes truthful reporting and no savings—i.e., the joint strategy $(m^* \equiv 0, \hat{c} = s + b)$—optimal for the agent. It is *NS-FO-IC* if (i) it is FO-IC and (ii) conditional on truthful reporting, the consumption process $\hat{c} = s + y$ satisfies the agent's Euler equation [5.7] at rate $r = \rho$. Intuitively, properties (i)–(ii) defining NS-FO-IC contracts are the infinitesimal optimality conditions implied by F-NS-IC.[12] Finally, *Contract PPI is implementable as an F-HS-IC contract* if there

[11] In an analogous discrete-time setting, Doepke and Townsend (2006) show that it is without loss of generality (in terms of implementable consumption processes and payoffs) to focus on F-HS-IC contracts.

[12] Cf. Footnotes 56 and 57 for possible technical caveats to this intuition.

is an F-HS-IC contract under which the agent's optimal joint strategy $(m, \hat{c})$ satisfies $m = m^* \equiv 0$ (truthful reporting) and $\hat{c} = c$ from [3.4] with $y = y^* = b$ (consumption is the same as that under Contract PPI and truthful reporting). We define implementability as an F-NS-IC contract in the obvious analogous manner.

Theorem 4. *Given any $\lambda \geqslant 0$, Contract PPI satisfies the following properties:*

(i) *It is the unique (hence, optimal) NS-FO-IC contract.*

(ii) *It is implementable as an F-HS-IC contract for any $F \subseteq \{m : m$ is b-adapted$\}$.*

Proof. *Point (i):* Under any NS-FO-IC contract, the agent's flow utility $u_t \equiv u(s_t + y_t)$ and marginal flow utility $u'(s_t + y_t) \equiv -\theta u_t$ are P^*-martingales. Plugging this into [3.1]–[3.2] and using Tonelli's Theorem to interchange the order of integration yields $q_t \equiv u_t/\rho$ and $p_t \equiv \theta u_t/(\rho + \lambda)$. The only FO-IC contract with these properties is Contract PPI.
Point (ii): As noted in Section 5.3, Contract PPI can be implemented as an SI Contract with zero taxes ($r = \rho$); moreover, this can be done with deterministic flow transfers rather than a lump-sum transfer at $t = 0$, per Footnote 36. This implies the present result because (a) those transfers are independent of the agents' reports by construction and (b) from the agent's perspective, saving via the ambient market is a perfect substitute for saving via the principal at rate ρ. □

The intuition for Theorem 4(i) is familiar from Allen's (1985) two-period model. With hidden savings, the agent only cares about the net present value (NPV) of the contract's transfers, so HS-IC requires that the agent receive the same NPV along every path of reports. Thus, it is *as if* all transfers were made in lump-sum at time $t = 0$, as in the SI Contract with no taxes ($r = \rho$), which implements Contract PPI. Two aspects of Theorem 4(ii) warrant elaboration:

- Theorem 4(ii) states that Contract PPI is implementable even when the agent's misreporting strategies are permitted to violate the no Ponzi constraint [**NP-*m***]. This might seem to contradict Observation 2 and Theorem 2, but it does not. The above proof shows that, with hidden savings, Contract PPI can be implemented *without communication* via determinstic transfers by *having the agent save and consume outside of the contract*. In such implementations, the no Ponzi constraint on assets [**NP-A**]—which is necessary for the agent's self-insurance problem to be well-posed—effectively imposes the same restrictions on the agent's consumption that [**NP-*m***] does in PPI's model without hidden savings. This suggests that [**NP-*m***] is needed for PPI's model to be well-behaved: without it, we would reach the unreasonable conclusion that Contract PPI is not IC in the original model (without hidden savings) but is HS-IC in the hidden savings model, wherein the agent has access to more deviations.
- Theorem 4(ii) does *not* state that Contract PPI is always implementable as an F-NS-IC contract. In fact, the obvious adaptation of Observation 2 implies that it is

not $\left[\mathscr{M}^{\mathrm{LAC}}_{\leqslant} \cap \mathscr{M}^r\right]$-NS-IC for any $r > \rho$. Again, this suggests that [**NP-*m***] (with an appropriately chosen rate r) is needed for PPI's model to be well-behaved: without it, restricting attention to NS-IC contracts would be *with* loss of generality, undercutting a key simplification on which much of the hidden savings literature is based (e.g., Cole and Kocherlakota 2001; DeMarzo and Sannikov 2006; He et al. 2017).

I. Incentive Compatibility of DR-SICs without Jump Reports

In this appendix, we consider the following corollary of Theorem 2:

Theorem 5. *For any given $r > 0$, every DR-SIC (b_0, q_0, r) is $\mathscr{M}^r$-IC. Per Fact 1, such contracts are also F-IC for any smaller strategy space $F \subseteq \mathscr{M}^r$.*

Our goal is to prove Theorem 5 from first principles, *without reference to the "extended" reporting problem from Section 6*, and to highlight some subtleties that arise in such an analysis. We first introduce the requisite definitions, then discuss the relevant subtleties, and finally present the proof (sketch) of Theorem 5.

Preliminaries. Fix a DR-SIC with rate $r > 0$. When the agent's strategy space is $\mathscr{M}^r$, she is constrained to misreports with absolutely continuous sample paths, viz., $m_t \equiv \int_0^t \Delta_\tau \mathrm{d}\tau$. Thus, as in PPI, her reporting problem can be viewed as one of stochastic control with states $(A_t^v, y_t, m_t) \in \mathbb{R}^3$ and control $\Delta_t \in \mathbb{R}$. As in Appendix B, $V^{\mathrm{NJ}}(A_t^v, y_t, m_t)$ denotes the agent's value function in this problem.

For a smooth function $\psi \in \mathbf{C}^2(\mathbb{R}^3)$ and $\Delta_t \in \mathbb{R}$, define the *infinitesimal generator*[13]

$$\begin{aligned}\mathscr{L}^{\Delta_t}\psi(A_t^v, y_t, m_t) := &\left[\mu - \lambda\cdot(y_t - m_t) + \Delta_t\right]\psi_y(A_t^v, y_t, m_t) + \Delta_t\cdot\psi_m(A_t^v, y_t, m_t)\\ &+ \left[\bar{A}(r;\lambda) + \frac{\lambda}{r+\lambda}y_t\right]\psi_A(A_t^v, y_t, m_t) + \frac{\sigma^2}{2}\cdot\psi_{yy}(A_t^v, y_t, m_t),\end{aligned}$$

where $\bar{A}(r;\lambda)$ is the constant from [**5.5**]; define the *Hamiltonian*

$$[\mathscr{H}] \qquad \mathscr{H}(A_t^v, y_t, m_t \mid \psi) := u(\hat{C}(A_t^v, y_t) - m_t) + \sup_{\Delta_t\in\mathbb{R}}\left[\mathscr{L}^{\Delta_t}\psi(A_t^v, y_t, m_t)\right],$$

where $\hat{C}(A_t^v, y_t)$ is from [**5.4**]; and define the *("no jump") HJB equation*

$$[\textbf{HJB-NJ}] \qquad \rho\psi(A_t^v, y_t, m_t) = \mathscr{H}(A_t^v, y_t, m_t \mid \psi).$$

The relevant notion of a "solution" to [**HJB-NJ**] will be that of a *supersolution*.

Definition I.1. A (continuous) locally bounded function $F : \mathbb{R}^3 \to \mathbb{R}$ is:

[13] We have $\mathrm{d}m_t = \Delta_t\mathrm{d}t$, $\mathrm{d}y_t = \mathrm{d}b_t + \mathrm{d}m_t$ (where $\mathrm{d}b_t$ satisfies [**2.1**] and $b_t \equiv y_t - m_t$), and $\mathrm{d}A_t^v = [\bar{A}(r;\lambda) + \frac{\lambda}{r+\lambda}y_t]\mathrm{d}t$ by [**6.1**].

(i) A *viscosity supersolution* of [**HJB-NJ**] if

$$\rho\psi(A_t^v, y_t, m_t) \geqslant \mathscr{H}(A_t^v, y_t, m_t \mid \psi) \tag{I.1}$$

for all $(A_t^v, y_t, m_t) \in \mathbb{R}^3$ and $\psi \in \mathbf{C}^2(\mathbb{R}^3)$ such that (A_t^v, y_t, m_t) is a minimizer of $F - \psi$.

(ii) A *classical supersolution* of [**HJB-NJ**] if it is a viscosity supersolution and in $\mathbf{C}^2(\mathbb{R}^3)$.

(iii) A *classical solution* to [**HJB-NJ**] if it is a classical supersolution and satisfies [**I.1**] (for $\psi = F$) with equality everywhere.

Definition I.1 is standard in the stochastic control literature (e.g., Pham 2009, Ch. 3–4; Touzi 2018, Ch. 2 and 6). For essentially any (maximization) control problem, if the value function is locally bounded (but possibly non-smooth), it is necessarily a viscosity supersolution of the relevant HJB equation (Touzi 2018, Proposition 6.2). If the value function is also smooth, it is necessarily a classical supersolution of the HJB equation (Touzi 2018, Proposition 2.4).

Discussion. Generally, even if a value function is smooth, additional regularity conditions on the Hamiltonian are needed to conclude that it is a classical solution, rather than just a supersolution (Touzi 2018, Proposition 2.5).[14] It is known that *the requisite regularity conditions typically fail in settings where the control variable is unbounded and enters linearly in the infinitesimal generator, as in the present formulation of the agent's reporting problem* (e.g., Pham 2009, Sec. 3.4.2 and 4.5). In such cases, the value function only satisfies the HJB equation in the weaker sense of being a supersolution.

To illustrate, recall from Lemma B.1 that, given the agent's value function $V^{\mathrm{DR}} \in \mathbf{C}(\mathbb{R}^2)$ from Section 6.4 in the extended reporting problem (with feasible set $\mathscr{M}^r_{\text{ext}}$), we can deduce that $V^{\mathrm{NJ}}(A_t^v, y_t, m_t) = V^{\mathrm{DR}}(A_t^v, y_t - m_t)$, and hence $V^{\mathrm{NJ}} \in \mathbf{C}(\mathbb{R}^3)$. This fact motivates the following observation:

Lemma I.2. The function $F(A_t^v, y_t, m_t) := V^{\mathrm{DR}}(A_t^v, y_t - m_t)$ is a classical supersolution of [**HJB-NJ**]. However, it is *not* a classical solution: [**I.1**] (with $\psi = F$) holds with equality at (A_t^v, y_t, m_t) iff $m_t = 0$.

Proof. By [**6.9**], we have $V^{\mathrm{DR}}(A_t^v, y_t - m_t) = \hat{V}^{\mathrm{DR}} \exp\left[-\theta r\left(A_t^v + \frac{b_t}{r+\lambda}\right)\right]$ for $\hat{V}^{\mathrm{DR}} = -\frac{1}{r}\exp\left[\theta \bar{A}(r;\lambda)\right]$ (where $\bar{A}(r;\lambda)$ is from [**5.5**]). Thus, for F defined as above, we have $F_A = -\theta r F$, $F_y = F_A/(r+\lambda)$, $F_m + F_y = 0$, and $F_{yy} = \theta^2 r^2 F/(r+\lambda)^2$. Furthermore, for each $(A_t^v, y_t, m_t) \in \mathbb{R}^3$, $u(\hat{C}(A_t^v, y_t) - m_t) = \exp\left[\frac{\theta\lambda}{r+\lambda} m_t\right] \cdot rF(A_t^v, y_t - m_t)$. These

[14] These conditions concern the Hamiltonian's continuity, when viewed as a function of the partial derivatives of ψ. For all other HJB equations stated in this paper (viz., [**6.9**] in Section 6.4, [**E.2**] in Appendix E, and [**J.6**] and [**J.9**] in Appendix J.3.1) it can be shown that the regularity conditions in Touzi (2018, Proposition 2.5) are satisfied because the control variables enter the strictly concave/convex "flow return" functions, yielding interior optima and allowing one to solve for the relevant Hamiltonians in closed-form. This justifies our focus on classical solutions elsewhere in the paper.

properties imply that

$$\rho F(A_t^v, y_t, m_t) - \mathscr{H}(A_t^v, y_t, m_t \mid F) = \underbrace{-rF(A_t^v, y_t, m_t)}_{>0} \cdot \underbrace{\left(\exp\left[\frac{\theta\lambda}{r+\lambda} m_t\right] - 1 - \frac{\theta\lambda}{r+\lambda} m_t\right)}_{\geqslant 0,\text{ with equality iff } m_t = 0}$$

where (strict) positivity of the last term follows from strict convexity of $x \mapsto e^x$. □

One way to "restore equality" in the agent's HJB equation is to reformulate the agent's problem by expanding his strategy space and treating $m_t \in \mathbb{R}$ as a control (rather than state) variable, as in Section 6. Another way is to keep $\Delta_t \in \mathbb{R}$ as the control, but reformulate [**HJB-NJ**]. For instance, noting that $\mathscr{H}(A_t^v, y_t, m_t \mid \psi) < \infty$ iff $\psi_y(A_t^v, y_t, m_t) + \psi_m(A_t^v, y_t, m_t) = 0$, [**HJB-NJ**] is equivalent to the variational inequality

$$\min\left\{\rho\psi(\cdot) - \mathscr{H}(\cdot \mid \psi), \left|\psi_y(\cdot) + \psi_m(\cdot)\right|\right\} = 0$$

described in Pham's (2009) treatment of *singular control* problems. Alternatively, one can reformulate [**HJB-NJ**] as the variational inequality in the associated *impulse control* problem (Oksendal and Sulem 2019) where setting $\Delta_t = \pm\infty$ is viewed as inducing a jump in m_t, as in Strulovici (2022).

Proof of Theorem 5. We work directly with [**I.1**], the inequality version of [**HJB-NJ**]. We describe the main steps, only sketching some details for brevity.

Step 1: Shape of Value Function. It is not *a priori* clear that the value function V^{NJ} coincides with V^{DR} (cf. Lemma B.1). But by applying the same controls at different states, it can be shown that

$$V^{\text{NJ}}(A_t^v, y_t, m_t) = \hat{V}^{\text{NJ}} \cdot h(m_t) \cdot \exp\left[-\theta r \left(A_t^v + \frac{y_t}{r+\lambda}\right)\right] \quad \text{[I.2]}$$

for some constant $\hat{V}^{\text{NJ}} < 0$ and convex function $h : \mathbb{R} \to \mathbb{R}_{++}$. Without loss of generality, we can normalize $h(0) := 1$.

Step 2: Determining the h Function. It is not *a priori* clear that h is smooth. But Step 1 implies that V^{NJ} is locally bounded, so Touzi (2018, Proposition 6.2) implies that V^{NJ} is a viscosity supersolution of [**HJB-NJ**]. By standard smooth approximation results,[15] there exists a dense subset $D \subseteq \mathbb{R}$ such that, at every point $m_t \in D$, there exists a function $\phi_{(m_t)} \in \mathrm{C}^2(\mathbb{R})$ satisfying $\phi_{(m_t)}(m_t) = h(m_t)$ and $\phi_{(m_t)}(\cdot) \geqslant h(\cdot)$; hence, h is differentiable at m_t and $h'(m_t) = \phi'_{(m_t)}(m_t)$. Thus, for any (A_t^v, y_t, m_t) with $m_t \in D$, the

[15]See, for instance, Lemma 8(g), Theorem 9, and associated discussion in Katzourakis (2015, Ch. 2).

function

$$\text{[I.3]} \qquad \psi_{(m_t)}(m_t)(A_t^v, y_t, m_t) := \hat{V}^{\mathrm{NJ}} \cdot \phi_{(m_t)}(m_t) \cdot \exp\left[-\theta r\left(A_t^v + \frac{y_t}{r+\lambda}\right)\right]$$

satisfies the conditions of Definition I.1(i). Plugging [I.3] into [$\mathscr{H}$], we see that $\mathscr{H}(A_t^v, y_t, m_t \mid \psi_{(m_t)}) < \infty$—a necessary condition for [I.1]—iff $\phi'_{(m_t)}(m_t) = f(r;\lambda)\phi_{(m_t)}(m_t)$, which is equivalent to $h'(m_t) = f(r;\lambda)h(m_t)$. As $h'(\cdot) = f(r;\lambda)h(\cdot)$ on the dense set $D \subseteq \mathbb{R}$ and h is convex (hence continuous and directionally differentiable on $\mathbb{R}$), it can be shown that $h'(\cdot) = f(r;\lambda)h(\cdot)$ on $\mathbb{R}$. The unique solution to this ODE satisfying $h(0) = 1$ is

$$\text{[I.4]} \qquad h(m_t) = \exp\left[f(r;\lambda)m_t\right].$$

Step 3: Determining the $\hat{V}^{NJ}$ Constant. Combining [I.1], [I.2], and [I.4] yields $\hat{V}^{\mathrm{NJ}} \geqslant -\frac{1}{r}\exp\left[\theta\bar{A}(r;\lambda)\right]$. We conclude that this inequality holds with equality because $V^{\mathrm{DR}}(A_t^v, y_t - m_t) \geqslant V^{\mathrm{NJ}}(A_t^v, y_t, m_t)$ by construction (cf. Section 6.4).[16]

Step 4: Optimal Strategy. Together, [I.2] and [I.4] imply that, for all states (A_t^v, y_t, m_t), the supremum in [$\mathscr{H}$] is attained at $\Delta_t = 0$ (in fact, at any $\Delta_t \in \mathbb{R}$). Step 3 and the same calculations underlying Lemma I.2 imply that V^{NJ} satisfies [I.1] with equality when $m_t = 0$. Thus, integrating [I.1] forward from an truthful initial state $(A_0^v, y_0, m_0 = 0)$ implies that truthful reporting attains $V^{\mathrm{NJ}}(A_0^v, y_0, m_0 = 0)$, i.e., [IC] holds. □

J. Optimal FO-IC Contracts

This appendix presents supporting details for the discussion of fully optimal contracts in Section 7.1. In Appendix J.3.1 (see Remark J.10), we also independently verify Steps 1–3 and 5 of PPI's derivation of Contract PPI, as described in Appendix A.

J.1. Domain of Implementable (Marginal) Promised Utilities

Lemma J.1. For each $q_0 < 0$, the following hold:

(i) If $\lambda > 0$, then under any FO-IC contract $k_t \in (0, \theta)$ for all $t \geqslant 0$. Furthermore, for each $k_0 \in (0, \theta)$ there exists an FO-IC contract with $k_t \equiv k_0$.

(ii) If $\lambda = 0$, then under any FO-IC contract $k_t = \theta$ for all $t \geqslant 0$. Furthermore, an FO-IC contract exists.

[16] This is the only place in the proof that we reference the extended reporting problem from Section 6; alternatively, we could reach the same conclusion by appealing to the (equivalent) self-insurance problem from Section 5.1. We do not know if this step can be avoided: [I.1] only places a lower bound on $\hat{V}^{\mathrm{NJ}}$ and *in principle* could hold as a strict inequality everywhere, so it seems necessary to appeal elsewhere for an upper bound on $\hat{V}^{\mathrm{NJ}}$. (One could instead conjecture that $\hat{V}^{\mathrm{NJ}} = -\frac{1}{r}\exp\left[\theta\bar{A}(r;\lambda)\right]$ and then appeal to a "verification theorem" as desribed in Section 4.2, but this would entail additional technical restrictions on the agent's strategy space beyond those embodied in $\mathscr{M}^r$. See Pham (2009, Theorem 3.5.3) for details.)

Proof. Point (i): Let $\lambda > 0$ and fix an FO-IC contract. Since $\lambda > 0$ and $u(\cdot) < 0$, we have $u_t < e^{-\lambda t}u_t < 0$ for all $t > 0$, so [3.1]–[3.2] imply that $p_t > \theta q_t$. Dividing through by $q_t < 0$ yields $k_t = p_t/q_t < \theta$, while $q_t, p_t < 0$ implies $k_t > 0$. DR-SICs demonstrate the existence claim (see [5.15] and recall that $f(\cdot;\lambda)$ has range $(0,\theta)$).

Point (ii): Let $\lambda = 0$ and fix an FO-IC contract. Since $\lambda = 0$, we have $u_t \equiv e^{-\lambda t}u_t$. Thus, [3.1] and [3.2] imply that $p_t = \theta q_t$, and dividing through by $q_t < 0$ yields $k_t = p_t/q_t = \theta$. Contract PPI demonstrates the existence claim. □

Lemma J.1 allows us to define the *domain* D of implementable (q,p) pairs by

[J.1]
$$D := \begin{cases} \{(q,p) \in \mathbb{R}^2_{--} : p/q \in (0,\theta)\} & \text{if } \lambda > 0 \\ \{(q,p) \in \mathbb{R}^2_{--} : p = \theta q\} & \text{if } \lambda = 0. \end{cases}$$

J.2. Permanent Shocks

Theorem 6. *If $\lambda = 0$, then Contract PPI satisfies the following properties:*

(i) *It is the unique optimal FO-IC contract.*

(ii) *It is F-IC for any feasible set $F \subseteq \{m : m$ is b-adapted$\}$.*

Proof. Point (i). For any FO-IC contract, Lemma J.1 implies that $p_t \equiv \theta q_t$. Display [A.1] then implies that, under truthful reporting, promised utility satisfies

[J.2]
$$q_t = q_0 \exp\left[\int_0^t \left(\rho - \beta_\tau - \frac{\sigma^2\theta^2}{2}\right) \mathrm{d}\tau - \sigma\theta W_t\right],$$

where $\beta_t \equiv u_t/q_t$. The agent's recommended consumption process is then $c_t \equiv \bar{c}(q_t,\beta_t) := -\log(-q_t\beta_t)/\theta$. Thus, substituting the transfer process $s_t \equiv \bar{c}(q_t,\beta_t) - y_t$ into [2.2] and ignoring terms that do not involve β, the principal minimizes

$$\mathsf{E}_0^*\left[\int_0^\infty e^{-\rho t}\left(-\log(\beta_t) + \int_0^t \beta_\tau \mathrm{d}\tau\right)\mathrm{d}t\right]$$

over all b-adapted, strictly positive β processes. This objective is strictly convex, so the optimal β process is unique, deterministic, and satisfies the pointwise FOC[17]

$$\frac{\mathrm{d}}{\mathrm{d}\beta_t}\left[-\log(\beta_t) + \frac{\beta_t}{\rho}\right] = 0 \quad \text{for all } t \geqslant 0.$$

[17] Focusing on deterministic β and applying integration by parts, the objective in [J.3] becomes $\int_0^\infty e^{-\rho t}\left(-\log(\beta_t) + \beta_t/\rho\right)\mathrm{d}t$. Note that the flow cost in this transformed objective satisfies an Inada condition at $\beta_t = 0$, implying that the optimal process must be strictly positive. Thus, the pointwise first-order condition from this transformed objective is [J.3].

Thus, the optimal β is $\beta_t \equiv \rho$, which upon substitution into [J.2] yields Contract PPI.
Point (ii). Under Contract PPI, $c_t = \bar{c}(q_0;\rho) + \bar{A}(\rho;0)t + y_t$ when $\lambda = 0$. Thus, the transfer process $s_t \equiv c_t - y_t$ is deterministic, i.e., report-independent. □

Remark J.2. Theorem 6(i) can be strengthened: *Contract PPI is the unique optimal contract satisfying* $\gamma_t + p_t \geqslant 0$ *for all* $t \geqslant 0$, the one-sided variant of [**FO-IC**] that is appropriate under NHB or IML (cf. Footnote 57). We adapt the above proof as follows. Defining $\hat{k}_t := -\gamma_t/q_t$, the constraint becomes $\hat{k}_t \geqslant \theta$. Promised utility satisfies $q_t = q_0 \exp\left[\int_0^t \left(\rho - \beta_\tau - \frac{\sigma^2 \hat{k}_\tau^2}{2}\right) \mathrm{d}\tau - \int_0^t \sigma \hat{k}_\tau \mathrm{d}W_\tau\right]$. The principal's problem is then additively separable in β and $\hat{k}$. The optimization over β is unchanged. The portion of the principal's objective involving $\hat{k}$ is $\mathbf{E}_0^*\left[\int_0^\infty e^{-\rho t}\left(\frac{\sigma^2}{2}\int_0^t \hat{k}_\tau^2 \mathrm{d}\tau + \sigma \int_0^t \hat{k}_\tau \mathrm{d}W_\tau\right)\right]$. The stochastic integral vanishes in expectation, and (constrained) minimization of the first term yields $\hat{k}_t \equiv \theta$.

J.3. Transient Shocks

Let $J(y_0, q_0, p_0)$ denote the principal's value function over FO-IC contracts given the initial condition (y_0, q_0, p_0) (see [**J.4**] below for details). We call $J : \mathbb{R} \times D \to \mathbb{R}$ the principal's *FO value function*.

Definition J.3. The environment is *regular* if:

(i) An optimal FO-IC contract (as defined in Appendix A) exists.

(ii) The principal's FO value function J is twice continuously differentiable.

Regularity is a technical assumption, which is implicitly adopted in PPI. It is needed to analyze the principal's FO problem with standard stochastic control techniques.

Theorem 7. *Suppose that* $\lambda > 0$*. If the environment is regular, then the optimal DR-SIC is not an optimal FO-IC contract.*

We prove Theorem 7 in Appendix J.3.1 below. Recall that DR-SICs are equivalent to Stationary contracts, i.e., those with constant $k_t = p_t/q_t$ processes (Lemma G.1). Thus, Theorem 7 equivalently states that the optimal FO-IC contract is not Stationary, consistent with Implication 2 at the end of Appendix A. Of course, if the first-order approach is valid—viz., every IC contract is FO-IC, and the optimal FO-IC contract is IC—then Theorem 7 also implies that the fully optimal contract outperforms the optimal DR-SIC/Stationary contract. It is an open question whether the first-order approach is valid in PPI's model with $\lambda > 0$.

J.3.1. *Proof of Theorem 7*

By the Martingale Representation Theorem, the agent's promised utility process q (defined in [**3.1**]) under any FO-IC contract satisfies [**A.1**] and FO-IC, and his marginal promised utility process p (defined in [3.2]) satisfies

$$\text{[J.3]} \qquad \mathrm{d}p_t = [(\rho+\lambda)p_t - \theta u_t]\,\mathrm{d}t + Q_t\sigma \mathrm{d}W_t^y,$$

where $\sigma \mathrm{d}W_t^y \equiv \mathrm{d}y_t - (\mu - \lambda y_t)\mathrm{d}t$ and Q is a y-adapted process.

Definition J.4. The principal's *auxiliary first-order (FO) problem* is

$$\text{[J.4]} \qquad J(y_0, q_0, p_0) := \inf_{(c,Q)\in \mathscr{A}_P(y_0,q_0,p_0)} \mathbf{E}_0^* \left[\int_0^\infty e^{-\rho t}\,(c_t - b_t)\,\mathrm{d}t\right]$$

where $\mathscr{A}_P(y_0, q_0, p_0)$ consists of the y-adapted processes (c, Q) such that [A.1] and [J.3] have unique solutions satisfying [FO-IC] and the transversality conditions: for all $t \geqslant 0$, $\lim_{T\to\infty} \mathbf{E}_t^*\left[e^{-\rho(T-t)}q_T\right] = \lim_{T\to\infty} \mathbf{E}_t^*\left[e^{-\rho(T-t)}p_T\right] = 0.$[18] The *first-stage FO problem* is

$$\text{[J.5]} \qquad \inf_{p_0<0 \text{ s.t. } (q_0,p_0)\,\in\, D} J(y_0, q_0, p_0),$$

and the *FO problem* is the joint optimization [J.4]–[J.5].

Definition J.3(ii) requires that $J \in \mathbf{C}^2(\mathbb{R}\times D)$. Standard arguments (Yong and Zhou 1999, Theorem 3.3; Touzi 2018, Propositions 2.4–2.5) then imply that J is a classical solution to the HJB equation

$$\begin{aligned}
\rho J(y_t, q_t, p_t) = \min_{(c_t,Q_t)\in\mathbb{R}^2} \Big[& c_t - y_t + J_y(y_t, q_t)\cdot(\mu - \lambda y_t) + J_q(y_t, q_t, p_t)\cdot(\rho q_t - u(c_t)) \\
\text{[J.6]} \qquad & + J_p(y_t, q_t, p_t)\cdot((\rho+\lambda)p_t + \theta u(c_t)) \\
& + \frac{\sigma^2}{2}J_{yy}(y_t, q_t, p_t) + \frac{\sigma^2 p_t^2}{2}J_{qq}(y_t, q_t, p_t) + \frac{\sigma^2 Q_t^2}{2}J_{pp}(y_t, q_t, p_t) \\
& - \sigma^2 p_t J_{yq}(y_t, q_t, p_t) + \sigma^2 Q_t J_{yp}(y_t, q_t, p_t) - \sigma^2 p_t Q_t J_{qp}(y_t, q_t, p_t)\Big].
\end{aligned}$$

(Cf. display (19) on p. 1249 of PPI.) We wish to rewrite [J.6] in terms of (y_t, q_t, k_t). This requires two lemmas, the latter of which appears in PPI as a conjecture.

Lemma J.5. Under any FO-IC contract and truthful reporting, the $k_t \equiv p_t/q_t$ process satisfies

$$\text{[J.7]} \qquad \mathrm{d}k_t = \left[(\beta_t+\lambda)\,k_t - \theta\beta_t + \sigma^2 k_t\left(k_t^2 - \hat{Q}_t\right)\right]\mathrm{d}t + \sigma\left[k_t^2 - \hat{Q}_t\right]\mathrm{d}W_t,$$

where the process $\hat{Q} = (\hat{Q}_t)_{t\geqslant 0}$ is defined as $\hat{Q}_t := -Q_t/q_t$.

Proof. Apply Ito's lemma to [A.1] and [J.3] under truthtelling ($W_t^y \equiv W_t$). □

[18] As in Section 2, we also implicitly restrict attention to (c, Q) processes such that the double-integral defining $J(y_0, q_0, p_0)$ is well-defined.

Lemma J.6. Let $\lambda > 0$. If the environment is regular, then J satisfies $J(y_t, q_t, p_t) \equiv \hat{J}(y_0, q_0, p_0/q_0)$, where

[**J.8**] $$\hat{J}(y_0, q_0, k_0) := -\frac{y_0}{\rho + \lambda} - \frac{\log(-q_0)}{\rho\theta} + h(k_0)$$

for some function $h \in \mathbf{C}^2\left((0, \theta)\right)$.

Proof. Regularity implies that J is well-defined and finite-valued and, given [**J.8**], also that $h \in \mathbf{C}^2\left((0, \theta)\right)$. Let $(q_0, p_0) \in D$, $y_0 \in \mathbb{R}$, and $\alpha > 0$ be given for Steps 1-2 below.
Step 1: We assert that $J(y_0, q_0, p_0) = J(0, q_0, p_0) - y_0/(\rho + \lambda)$. Let $(c, Q) \in \mathcal{A}_P(y_0, q_0, p_0)$ be given. Define $g_t := y_0 e^{-\lambda t}$, $\tilde{c}_t(y) := c_t(y + g)$, and $\tilde{Q}_t(y) := Q_t(y + g)$.[19] We have $(\tilde{c}, \tilde{Q}) \in \mathcal{A}_P(0, q_0, p_0)$. Let $\mathbf{P}^{*,(y_0)}$ denote the distribution over report paths starting from y_0 and $\mathbf{P}^{*,(0)}$ denote the distribution starting from $\tilde{y}_0 = 0$, assuming truthful reporting. The law of (c, Q) under $\mathbf{P}^{*,(y_0)}$ equals the law of $(\tilde{c}, \tilde{Q})$ under $\mathbf{P}^{*,(0)}$. Thus,

$$\mathbf{E}_0^{*,(y_0)}\left[\int_0^\infty e^{-\rho t}\left(c_t - y_t\right)\,\mathrm{d}t\right] = \mathbf{E}_0^{*,(0)}\left[\int_0^\infty e^{-\rho t}\left(\tilde{c}_t - y_t\right)\,\mathrm{d}t\right] - \underbrace{\int_0^\infty e^{-\rho t} g_t\,\mathrm{d}t}_{= y_0/(\rho+\lambda)}.$$

Step 2: We assert that $J(y_0, \alpha q_0, \alpha p_0) = J(y_0, q_0, p_0) - \log(\alpha)/(\rho\theta)$. Let $(c, Q) \in \mathcal{A}_P(y_0, q_0, p_0)$ be given. Define $\tilde{c}_t := c_t - \log(\alpha)/\theta$ and $\tilde{Q}_t := \alpha Q_t$. Note that $u(\tilde{c}_t) \equiv \alpha u(c_t)$. Display [**A.1**], [**J.3**], and [**FO-IC**] then imply that $(\tilde{c}, \tilde{Q}) \in \mathcal{A}_P(y_0, \alpha q_0, \alpha p_0)$. The distribution over endowment paths is the same at both initial states, so the principal's cost of $\tilde{c}$ at $(y_0, \alpha q_0, \alpha p_0)$ equals her cost of c at (y_0, q_0, p_0) plus $\int_0^t e^{-\rho t}\left(-\log(\alpha)/\theta\right)\,\mathrm{d}t = -\log(\alpha)/(\theta\rho)$.
Step 3: Fix $(q_0, p_0) \in D$ and $y_0 \in \mathbb{R}$. Combining Steps 1 and 2 with $\alpha = -p_0/q_0$ yields

$$J(y_0, q_0, p_0) = -\frac{y_0}{\rho + \lambda} - \frac{\log(-q_0)}{\rho\theta} + J\left(0, -1, -p_0/q_0\right).$$

Defining $h(k_0) := J\left(0, -1, -k_0\right)$ and $\hat{J}$ as in [**J.8**] completes the proof. □

Using Lemmas J.5 and J.6, we can write the HJB equation [**J.6**] in terms of (y_t, q_t, k_t) as

[**J.9**] $$\begin{aligned}\rho\hat{J}(y_t, q_t, k_t) = \min_{\beta_t > 0, \hat{Q}_t \in \mathbb{R}} &\Big\{\overline{c}(q_t, \beta_t) - y_t - \frac{1}{\rho + \lambda}\left[\mu - \lambda y_t\right] - \frac{1}{\rho\theta}\left[\rho - \beta_t\right] \\ &+ h'(k_t)\left[\left(\beta_t + \lambda\right)k_t - \theta\beta_t + \sigma^2 k_t\big(k_t^2 - \hat{Q}_t\big)\right] \\ &+ \frac{\sigma^2 k_t^2}{2\rho\theta} + \frac{\sigma^2(k_t^2 - \hat{Q}_t)^2}{2} h''(k_t)\Big\},\end{aligned}$$

[19] That is, the path of $(\tilde{c}, \tilde{Q})$ when the agent reports the endowment path $\hat{y} \in \mathbf{C}\left([0, \infty)\right)$ is the same as the path of (c, Q) when he reports the endowment path $\hat{y} + g$.

where $\beta_t = u(c_t)/q_t$ and $\overline{c}(q_t, \beta_t) = -\log(-\beta_t q_t)/\theta$.

Lemma J.7. Let $\lambda > 0$. If the environment is regular, then any optimal $\beta^\dagger$ satisfies $\beta_t^\dagger \equiv \hat{\beta}(k_t)$, where $\hat{\beta} : (0, \theta) \to \mathbb{R}$ is defined by

$$\hat{\beta}(k_t) := \frac{1}{1/\rho + \theta(k_t - \theta)h'(k_t)}. \tag{J.10}$$

Proof. Eliminating terms on the RHS of [**J.9**] yields the following minimization over β_t, which any optimal $\beta^\dagger$ process must a.e. satisfy

$$\min_{\beta_t > 0} \left[-\log(\beta_t) + \frac{\beta_t}{\rho} + \beta_t \cdot h'(k_t)(k_t - \theta) \right]. \tag{J.11}$$

The unique solution to [**J.11**] is interior and characterized by the FOC [**J.10**]. □

Lemma J.8. Let $\lambda > 0$. If the environment is regular, then any $k_0^\dagger \in \arg\min_{k_0 \in (0,\theta)} h(k_0)$ satisfies $\hat{\beta}(k_0^\dagger) = \rho$.

Proof. Because $h \in \mathbf{C}^2((0, \theta))$ by Definition J.3(i) and Lemma J.6, any such $k_0^\dagger$ must satisfy the FOC $h'(k_0^\dagger) = 0$. Plugging this into [**J.10**] completes the proof. □

Lemma J.9. Let $\lambda > 0$. If the environment is regular and an optimal FO-IC contract is Stationary (i.e., induces a constant k process), then that contract is Contract PPI.

Proof. Consider any optimal FO-IC contract that is Stationary. Lemma J.5 and the unique decomposition property for Ito processes imply that $k_t^2 \equiv \hat{Q}_t^\dagger$ (so k has zero volatility) and thus that $\beta_t^\dagger(\lambda + k_t) - \theta\beta_t^\dagger \equiv 0$ (so k has zero drift). Furthermore, [**J.5**] and [**J.8**] imply that $k_t \equiv k_0^\dagger \in \arg\min_{k_0 \in (0,\theta)} h(k_0)$. Then Lemma J.8 implies that $\hat{\beta}_t^\dagger \equiv \rho$ and Lemma G.1 implies that $k_t \equiv f(\rho; \lambda) = k_0^*$. This yields Contract PPI. □

Proof of Theorem 7. Suppose the environment is regular. Theorem 1(i) implies that the optimal SI Contract / DR-SIC strictly dominates Contract PPI. Every DR-SIC is FO-IC (by construction) and Stationary (by Lemma G.1). If an optimal FO-IC were Stationary, then Lemma J.9 would imply that it is Contract PPI, a contradiction. □

Remark J.10. The above work confirms Steps 1–3 and 5 of PPI's derivation of Contract PPI, as described in Appendix A. The conjecture in Step 1 is established by Lemma J.6. Step 2 follows from Lemma J.6 and [**J.9**]. Step 3 follows from plugging the optimal $\hat{Q}_t$ from [**J.9**] into [**J.7**]. Step 5 is Lemma J.9.

K. Properties of Brownian Motion and OU Process

This appendix collects auxiliary facts about Brownian motion (BM) and OU processes.

Lemma K.1 (SLLN for BM)**.** Let $W = (W_t)_{t\geqslant 0}$ be a standard Brownian motion. Then $\lim_{t\to\infty} W_t/g(t) = 0$ for any function $g : \mathbb{R}_+ \to \mathbb{R}_+$ such that $\lim_{t\to\infty} t/g(t) < \infty$.

Proof. The $g(t) = t$ case is standard (e.g., Problem 9.4 in Karatzas and Shreve (1998, p. 104)). Thus, more generally, $\lim_{t\to\infty} W_t/g(t) = \lim_{t\to\infty} (W_t/t) \cdot \lim_{t\to\infty} (t/g(t)) = 0$. □

Let b be an OU process as defined by the equation [2.1], the solution to which is

$$b_t = b_0 e^{-\lambda t} + \mu \underbrace{\left(\frac{1-e^{-\lambda t}}{\lambda}\right)}_{= t \text{ when } \lambda = 0} + \underbrace{e^{-\lambda t}\int_0^t \sigma e^{\lambda\tau}\,\mathrm{d}W_\tau}_{=:\,X_t}. \tag{K.1}$$

We invoke $X = (X_t)_{t\geqslant 0}$ from [K.1] below. Note that $X_t \equiv b_t$ when $b_0 = \mu = 0$.

Lemma K.2. Given X_t from [K.1], the following holds for each $t \geqslant 0$:

(i) When $\lambda = 0$,

$$\int_0^t b_\tau\,\mathrm{d}\tau = b_0 t + \frac{1}{2}\mu t^2 + \sigma\left(tW_t - \int_0^t \tau\,\mathrm{d}W_\tau\right) \tag{K.2}$$

(ii) When $\lambda > 0$,

$$\int_0^t b_\tau\,\mathrm{d}\tau = b_0\left(\frac{1-e^{-\lambda t}}{\lambda}\right) + \frac{\mu}{\lambda}\left(t - \frac{1-e^{-\lambda t}}{\lambda}\right) + \frac{\sigma W_t - X_t}{\lambda} \tag{K.3}$$

Proof. In both points (i) and (ii), the deterministic terms follow from straightforward integration of the first two terms in [K.1]. The stochastic terms follow from stochastic integration by parts calculations:

Point (i): When $\lambda = 0$, we have $X_t = \sigma W_t$. Itô's lemma yields applied to tW_t yields $tW_t = \int_0^t W_\tau\,\mathrm{d}\tau + \int_0^t \tau\,\mathrm{d}W_\tau$. Thus, $\int_0^t X_\tau\,\mathrm{d}\tau = \sigma\left(tW_t - \int_0^t \tau\,\mathrm{d}W_\tau\right)$, as desired.

Point (ii): When $\lambda > 0$, we have $X_t = e^{-\lambda t}\int_0^t \sigma e^{\lambda\tau}\,\mathrm{d}W_\tau =: e^{-\lambda t}Y_t$. Itô's lemma applied to X_t yields $\mathrm{d}X_t = -\lambda X_t\,\mathrm{d}t + e^{-\lambda t}\,\mathrm{d}Y_t = -\lambda X_t\,\mathrm{d}t + \sigma\,\mathrm{d}W_t$. Putting this in integral form and rearrangning yields $\int_0^t X_\tau\,\mathrm{d}\tau = (\sigma W_t)/\lambda - X_t/\lambda$, as desired. □

Lemma K.3. The following hold (almost surely):

(i) If $\lambda > 0$, then $\lim_{t\to\infty} b_t/t = 0$. If $\lambda = 0$, then $\lim_{t\to\infty} b_t/t = \mu$,

(ii) For all $\lambda \geqslant 0$ and $\alpha > 0$, $\lim_{t\to\infty} e^{-\alpha t} b_t = 0$, and

(iii) For all $\lambda \geqslant 0$ and $\alpha > 0$, $\lim_{t\to\infty} e^{-\alpha t}\int_0^t b_\tau\,\mathrm{d}\tau = 0$.

Proof. We consider each point of the lemma in turn.

Point (i): When $\lambda = 0$, the result is immediate from [K.1] and Lemma K.1. When $\lambda > 0$, we have

$$\lim_{t\to\infty}\frac{1}{t}\left[b_0 e^{-\lambda t} + \mu\left(\frac{1-e^{-\lambda t}}{\lambda}\right)\right] = 0.$$

Thus, it suffices to show that $\lim_{t\to\infty} X_t/t = 0$ for X_t in [K.1]. Defining the *time-change* $v(t) := e^{2\lambda t} - 1$ (with inverse $t(v) := \log(v+1)/(2\lambda)$), the process $B_v := \frac{\sqrt{2\lambda}}{\sigma} e^{\lambda t} X_{t(v)}$ is a standard BM (Karatzas and Shreve 1998, p. 174). By construction,

$$\frac{X_t}{t} = \frac{\sigma\sqrt{2\lambda} B_{v(t)}}{\sqrt{v(t)}\log\left(v(t)\right)}. \tag{K.4}$$

It suffices to show that the RHS of [K.4] goes to zero as $v \to \infty$. To this end, the Law of the Iterated Logarithm (Mörters and Peres 2010, p. 119) implies that

$$\limsup_{v\to\infty} \frac{|B_v|}{\sqrt{2v\log\left(\log(v)\right)}} = 1 \tag{K.5}$$

and L'Hôpital's rule implies that

$$\lim_{v\to\infty} \frac{\sqrt{2v\log\left(\log(v)\right)}}{\sqrt{v}\log\left(v\right)} = \lim_{v\to\infty} \frac{1}{\sqrt{2}\log(v)\sqrt{\log\left(\log(v)\right)}} = 0. \tag{K.6}$$

Combining [K.5] with [K.6] yields the desired conclusion that

$$\limsup_{v\to\infty} \frac{|B_v|}{\sqrt{v}\log\left(v\right)} = \limsup_{v\to\infty} \frac{|B_v|}{\sqrt{2v\log\left(\log(v)\right)}} \frac{\sqrt{2v\log\left(\log(v)\right)}}{\sqrt{v}\log\left(v\right)} = 0.$$

Point (ii): Let $\alpha > 0$ be given. Then

$$\limsup_{t\to\infty} e^{-\alpha t}|b_t| = \limsup_{t\to\infty} \frac{|b_t|}{t}\cdot\frac{t}{e^{\alpha t}} = \mathbf{1}(\lambda = 0)|\mu| \cdot \lim_{t\to\infty}\frac{t}{e^{\alpha t}} = 0$$

where the second equality follows from point (i).
Point (iii): First, suppose that $\lambda > 0$. Lemma K.2(ii) yields

$$e^{-\alpha t}\int_0^t b_\tau \,\mathrm{d}\tau = e^{-\alpha t}\left[b_0\left(\frac{1-e^{-\lambda t}}{\lambda}\right) + \frac{\mu}{\lambda}\left(t - \frac{1-e^{-\lambda t}}{\lambda}\right)\right] + \frac{\sigma}{\lambda}e^{-\alpha t}W_t - \frac{e^{-\alpha t}}{\lambda}X_t$$

The first term clearly goes to zero as $t \to \infty$. The second term also goes to zero by Lemma K.1. The third term goes to zero by point (ii) of the present lemma. Next, suppose that $\lambda = 0$. Lemma K.2(i) yields

$$e^{-\alpha t}\int_0^t b_\tau \,\mathrm{d}\tau = e^{-\alpha t}\left[b_0 t + \frac{1}{2}\mu t^2\right] + \sigma e^{-\alpha t} t W_t - \sigma e^{-\alpha t}\int_0^t \tau \,\mathrm{d}W_\tau \tag{K.7}$$

The first term clearly goes to zero as $t \to \infty$, as does the second term because

$$\lim_{t\to\infty} e^{-\alpha t} t W_t = \lim_{t\to\infty} \frac{W_t}{e^{\alpha t/2}} \cdot \underbrace{\lim_{t\to\infty} \frac{t}{e^{\alpha t/2}}}_{=0}$$

and $\lim_{t\to\infty} W_t/e^{\alpha t/2} = 0$ by Lemma K.1. For the final term in [K.7], define $Z_t := \int_0^t \tau \, \mathrm{d}W_\tau$. Using the time-change $\phi(t) := t^3/3$ (with inverse $t(\phi) := (3\phi)^{1/3}$), the process $\hat{B}_{\phi(t)} := Z_t$ is a standard BM. Therefore,

$$\lim_{t\to\infty} e^{-\alpha t} Z_t = \lim_{\phi\to\infty} e^{-\alpha t(\phi)} \hat{B}_\phi = \lim_{\phi\to\infty} \frac{\hat{B}_\phi}{\phi} \cdot \frac{\phi}{\exp(\alpha(3\phi)^{1/3})} = 0,$$

where the last equality is by Lemma K.1 and L'Hôpital's rule. □

Lemma K.4. For any $\alpha > 0$, the following holds:

$$\int_0^\infty e^{-\alpha t} b_t \, \mathrm{d}t = \frac{b_0}{\alpha+\lambda} + \frac{\mu}{\alpha(\alpha+\lambda)} + \frac{\sigma}{\alpha+\lambda} \int_0^\infty e^{-\alpha t} \, \mathrm{d}W_t \tag{K.8}$$

Proof. We integrate [K.1], discounted by $e^{-\alpha t}$. The first two terms yield

$$\int_0^\infty e^{-\alpha t} \left[b_0 e^{-\lambda t} + \mu \underbrace{\left(\frac{1 - e^{-\lambda t}}{\lambda} \right)}_{= t \text{ when } \lambda = 0} \right] \mathrm{d}t = \frac{b_0}{\alpha+\lambda} + \frac{\mu}{\alpha(\alpha+\lambda)} \tag{K.9}$$

For the final $\int_0^\infty e^{-\alpha t} X_t \, \mathrm{d}t$ term, applying Itô's lemma twice yields

$$e^{-\alpha T} X_T = -(\alpha+\lambda) \int_0^T e^{-\alpha t} X_t \, \mathrm{d}t + \int_0^T e^{-\alpha t} \sigma \, \mathrm{d}W_t.$$

Rearranging and letting $T \to \infty$, we obtain

$$\int_0^\infty e^{-\alpha t} X_t \, \mathrm{d}t = \frac{\sigma}{\alpha+\lambda} \int_0^\infty e^{-\alpha t} \, \mathrm{d}W_t - \underbrace{\frac{1}{\alpha+\lambda} \lim_{T\to\infty} e^{-\alpha T} X_T}_{= 0 \text{ by Lemma K.3(ii)}}. \tag{K.10}$$

Combining [K.1], [K.9], and [K.10] completes the proof. □