Risk Sharing and Risk Reduction with Moral Hazard*

Adrien Wicht† Ramon Marimon‡ Luca Zavalloni§

September, 2026

Abstract

Official lending programs aim to stabilize and strengthen financially stressed economies, while having incentives for the country to act with the same aim. However, the capacity of a country to improve its risk profile depends on whether the risk distribution is exogenous and, in this case, if the risk can be improved with effort, or endogenous, and risk can be reduced by the choice of distribution. The program designer knows the country's feasible distributions but cannot contract on effort, in the exogenous case, or on a specific choice of distribution, in the endogenous case; it must design appropriate incentives. The *canonical moral hazard theory* addresses the former and incentives rely on performance: rewards and punishments, which disrupt risk sharing. The *flexible moral hazard theory* addresses the latter and incentives are provided by ex ante cost compensations, which do not undermine risk-sharing. In a general framework of dynamic contracts: canonical moral hazard incentives lead, in general, to *welfare immiseration*; instead, dynamic flexible moral hazard incentives lead to *welfare bliss*, but also no-risk (*risk bliss*). We introduce a novel entropy-based cost where no-risk it too costly. We also show that back-loading incentives through subcontracts reconciles canonical moral hazard with improved risk sharing. We calibrate these different programs to stressed Euro Area countries and compare them as dynamic constraint-efficient contracts with enforcement and different incentive constraints, unveiling sizeable distinct effects on business-cycle dynamics and welfare.

Keywords: official lending, limited enforcement, moral hazard, risk sharing, recursive contract
JEL classification: E43, E44, E47, E62, F34, F36, F37

*The views expressed are those of the authors and do not necessarily reflect those of the European Stability Mechanism, which support is acknowledged. We would like to thank the participants at UPF, CREi, the University of Basel, the 2024 ADEMU Workshop at the BSE Summer Forum, the 2024 EEA meeting in Rotterdam, the 2024 SED annual meeting in Barcelona, and the 2025 World Congress of the Econometric Society in Seoul. Wicht gratefully acknowledges the financial support from the Swiss National Science Foundation (grant number 215537).

†University of Basel

‡European University Institute, BSE, CREi, CEPR and NBER

§European Stability Mechanism

1 Introduction

A central challenge in the design of international official lending programs is to reconcile two objectives that are inherently in tension: providing insurance against adverse macroeconomic shocks while preserving incentives for policy effort. Countries are exposed to both exogenous risks, whose realization lies outside their control, and endogenous risks, whose likelihood depends on policy choices and institutional decisions. Even when shocks themselves are exogenous their consequences are often shaped by preventive policies and broader socio-economic structures. The scope for effective risk sharing therefore depends not only on the nature of risks but also on the capacity of countries to improve their risk profile and on the contractual environment through which incentives for risk reduction are provided. Understanding how different contractual environments resolve this trade-off is essential for evaluating long-term lending arrangements, especially those designed to stabilize financially stressed countries.

This paper develops a more general theory of optimal dynamic contracts with moral hazard between a sovereign borrower and a Financial Stability Fund – hereafter, the Fund – that is explicitly motivated by the institutional features of official lending programs. The contract must deliver insurance against macroeconomic shocks, remain incentive compatible (IC) when choices, actions or policies of countries are not contractable (e.g. unobservable), and respect limited enforcement (LE) on both sides. The borrower acts as an agent, while the Fund acts as a principal in a principal-agent contracting relationship. We show that how moral hazard (MH) is modeled is central for risk sharing, the design of incentives, and long-run welfare dynamics, and that the standard formulation misses key features of official lending arrangements.

To start, it is useful to distinguish three types of risk: *i*) purely exogenous, for which the country can at most have an insurance or a risk-sharing contract; *ii*) semi-endogenous, when the risk distribution is exogenous but the country can act – in particular, with effort – to improve the risk profile, for example preventive measures can reduce the damage of a natural catastrophe, and *iii*) purely endogenous when the country can choose the risk distribution, for example, defense policy is, to a large extend, to choose risk distributions. In contrast with (*i*), there is a moral hazard problem with (*ii*) and (*iii*), as already said.

Since the pioneer work of Holmström (1979) on the moral hazard problem for the semi-endogenous case (*ii*), it has become the *canonical moral hazard* model: risk-distributions, with a common support, are known by both parties and with non-contractable effort the agent can improve the risk profile, where improvement means first-order stochastic dominance across the known distributions. This implies, first, that the effort has a global effect on the risk distribution and, second, that realized outcomes convey information about effort, thus incentives, based on performance, take the form of rewards and punishments.

Recently, Georgiadis et al. (2024) have focused in the purely endogenous case (*iii*) with a novel *model of flexible moral hazard*: without a constraint of a common support, the agent chooses a distribution, where first-order stochastic dominance requires a more costly distribution. This implies, first, that there may be *local flexibility* in how distributions are perturbed and, second, that performance may not be revealing, so incentives cannot rely on ex post performance, indeed there are only rewards increasing in distribution costs, which requires that also the principal knows the cost of the distributions.

In sum, this distinction between (*ii*) and (*iii*), between *canonical* and *flexible* moral hazard, substantially changes the set of contractual environments that can be analyzed and captures specific policy decisions faced by sovereign borrowers, such as the design of social protection systems or regulatory frameworks, whose consequences materialize through changes in the underlying risk profile rather than through observable effort. We build on the flexible MH framework but depart from it in two ways that prove essential. First, we embed it in a fully dynamic contracting problem. Second, we generalize the cost of choosing a distribution: Georgiadis et al. (2024) tie it to the mean alone, whereas we make risk reduction itself costly through an entropy-based cost. As we show below, this second change is what gives risk sharing any value at all under flexible MH.

The analysis delivers three main insights. First, the scope for risk sharing depends crucially on the specification of MH. The canonical approach applies the general contract-enforcement principle of "the carrot or the stick" with dynamic contracts: to further increase the value of the contract to the agent when the outcome is good and to further diminish it if it is bad. This means the *ex post* value of the contract varies more with than without incentives. For instance, the value decreases for bad outcomes that may occur due to bad luck rather than a lack of effort. When the contract is a risk-sharing agreement between a risk-averse agent and a risk-neutral principal, this provision of incentives disrupts the full risk-sharing that could be achieved with observable effort. Under flexible MH, by contrast, incentives are not tied to realized outcomes: the optimal contract instead compensates *ex ante* the agent for the marginal cost of the distribution it. As a result, incentive provision does does not disrupt risk-sharing.

Our second insight is that dynamic incentives need not imply immiseration. In the canonical framework, the literature has shown that, absent full information,[1] constrained efficient allocations are achieved only at the cost of immiseration. Atkeson and Lucas (1992) find that the cross-sectional consumption path spreads out indefinitely, so that immiseration arises at

[1] And unbounded utility.

the individual level. Under flexible MH, this mechanism is reversed. Incentives are provided through transfers that compensate effort costs *ex ante*, rather than through *ex post* realizations, so that low realizations are not punished and risk sharing is preserved. Formally, the agent's relative Pareto weight follows a submartingale, and expected welfare drifts upward over time rather than collapsing toward the lower bound. We label this *bliss*, the mirror image of immiseration. With two-sided limited enforcement constraints, the borrower's LE stops the downward utility drift preventing immiseration, while the lender's LE may stops the upward utility drift preventing bliss.

The third insight is that back-loading incentives reconciles canonical MH with risk sharing. We show that a contract structured as a sequence of subcontracts, in the spirit of Marcet and Marimon (1992), improves risk sharing for the borrower relative to the canonical benchmark. Within each subcontract, incentives are latent and accumulate over time. Rewards and punishments are realized only at reset points defined by a binding LE constraint. This structure preserves incentives while reducing the reliance on outcome-contingent distortions during normal times. Interestingly, it also brings the contract closer, on one hand, to flexible MH since, if it is the borrower's LE that binds punishment is not feasible, therefore the "punish and reward" principle becomes "stick to the bound and reward even more", and, on the other hand, to existing official lending programs, where it is common that a relatively short-term program is followed by another, with a new risk assessment based on past performance.

To highlight the *flexible* vs the *canonical* mechanisms, we begin by analyzing a static version of the contracting problem. The static environment allows us to isolate the role of MH in shaping risk sharing, abstracting from intertemporal considerations. Under flexible MH, we show that when the cost of effort is purely linked to the mean of the outcome distribution, as in Georgiadis et al. (2024), the borrower can fully self-insure by choosing degenerate (i.e. Dirac) distributions, rendering contracts with the Fund useless in terms of risk sharing. This highlights an important limitation of the original flexible framework. If reducing risk is not itself sufficiently costly, the borrower optimally eliminates uncertainty altogether, leaving no role for risk sharing. Introducing an entropy-based component to the cost function fundamentally alters this conclusion: optimal distributions become interior, and risk sharing emerges as a mutually beneficial arrangement. More precisely, the optimal distribution is a logit-like distribution that assigns positive probability mass to every element of the support. Hence, the Fund risk-sharing contract has value under flexible MH only when risk reduction is itself costly.

We then turn to the dynamic contracting problem, where canonical and flexible MH have opposite long run dynamics. Using the recursive Lagrangian approach of Marcet and Marimon (2019), we characterize the contract through the law of motion of the relative

Pareto weight between the borrower and the Fund, which cumulates the multipliers on the LE and IC constraints period by period. As before, the optimal distribution under flexible MH remains a Dirac under mean-based cost and a logit-like distribution under entropy-based cost. This second outcome can therefore rationalize Gumbel-distributed stochastic choice structures in quantitative work. We show that the multiplier on the IC constraint is strictly increasing in the endowment, since a higher realization raises both. That is, the cost the borrower must be compensated for and the surplus the Fund has to share. Crucially, incentives operate on this multiplier rather than on the endowment realization itself, so that, whenever the borrower's LE constraint is slack and the borrower is not too impatient relative to the Fund, the relative Pareto weight is a left-bounded submartingale: expected continuation utility drifts upward rather than collapsing. We call this *bliss*, the mirror image of immiseration. Under canonical MH, the same law of motion behaves very differently: an inverse-Euler-equation argument shows that the IC constraint leaves the *expected* Pareto weight unaffected, even though it distorts consumption state by state: rewards for good realizations and punishments for bad ones which exactly cancel out in expectation, so that, as in Atkeson and Lucas (1992), immiseration occurs along individual paths.

We also bring closer the canonical MH to the flexible MH by restoring risk sharing without weakening incentives in the canonical. For this, we back-load incentives based on the following logic. If the LE constraints are not binding in the flexible MH, incentive provisions never distort risk-sharing. In opposition, in the canonical MH, incentive provisions always distort risk sharing. We therefore consider a long-term canonical MH contract as a sequence of subcontracts. Within subcontracts consumption follows a purely deterministic path, decaying at the rate of relative impatience exactly as if MH were absent. The multiplier on the IC constraint becomes a *latent* variable that accumulates rewards and punishments without paying them out. When one of the LE constraints binds the subcontract terminates and a new one starts with an updated Pareto weight accounting for the accumulated incentives. This reset rule is asymmetric in a precise sense: since the borrower's outside option is increasing in the endowment, it is a *high* realization that triggers the borrower's LE constraint, so consumption jumps up with no punishment attached. In opposition, a *low* realization can trigger the Fund's LE constraint, in which case consumption falls below its deterministic path. Because it is the infrequent binding of enforcement, rather than the realized endowment as such, that governs when consumption departs from its deterministic path, full risk sharing is preserved throughout each subcontract.

To assess the quantitative relevance of these theoretical mechanisms, we embed the Fund contract into a broader environment with production and an outside option given by temporary autarky with incomplete markets and defaultable debt (IMD). The model is calibrated to stressed Euro Area economies (Portugal, Italy, Greece and Spain), where official lending programs have played a central role. Importantly, we consider that the risk profile of a country has two components: a purely exogenous common productivity component (*i*) and an endogenous idiosyncratic component that reflects domestic policy choices, either semi-endogenous (*ii*) or purely endogenous (*iii*). This decomposition allows us to study environments in which part of the country's risk is insurable while another part depends on incentives, thereby providing a quantitative framework for comparing alternative MH specifications.

The quantitative analysis serves two distinct purposes. First, as a positive exercise, each MH specification is calibrated on its own terms to match key business-cycle and debt moments. This allows us to assess whether alternative MH specifications have empirical content and can reproduce observed macroeconomic dynamics. Second, as a normative exercise, we fix a common parameterization across contracts and compare welfare and risk-sharing outcomes holding fundamentals constant. This isolates the role of contract design per se, rather than differences in preferences or technology.

We find that the different MH specifications emdebed in an IMD economy can match the data well. On targeted moments, the entropy-cost flexible MH records the lowest mean absolute deviation, followed by the mean-cost flexible MH. Matching the data in the flexible MH requires more impatience, a lower output cost of default and a lower post-default reentry parobability compared to the canonical MH. On untargeted moments, canonical MH has the best overall fit, but the ranking flips across variables: canonical MH matches the volatility of the spread and of the primary surplus well, yet generates a pro-cyclical, rather than counter-cyclical, spread – a feature that both flexible specifications capture, at the cost of missing the primary surplus moments.

Regarding the different Fund contracts, entropy is the sharpest dividing line. It is the lowest under mean-cost flexible MH, but is the highest of all contracts and genuinely time-varying under entropy-cost flexible MH. Both flexible contracts sustain substantially higher effort and output than canonical MH, yet only the entropy-cost contract converts this into higher consumption, and only it raises mean consumption relative to its own IMD benchmark. Entropy-cost MH delivers the lowest volatility of every variable and lowers each of them relative to IMD, whereas mean-cost MH delivers the highest. Canonical and back-loaded MH, whose IC multiplier is small, behave almost identically and work through co-movement rather than volatility, making labor and the primary surplus markedly more pro-cyclical. Entropy-cost flexible MH combines lower volatility with a stronger consumption-output correlation, while mean-cost flexible MH enlarges volatility and loosens the link between consumption and output.

Welfare gains, computed as consumption-equivalent changes, are immense. The entropy-cost flexible MH Pareto dominate all the others from the borrower's side, with gains above 100% relative to canonical MH, followed by the mean-cost flexible MH, with gains of roughly 80% on average. As expected, the Fund contract under back-loaded MH is welfare improving relative to canonical MH for the borrower, with the gain coming from the improved insurance.

The paper is organized as follows. Section 1.1 reviews the literature. Section 2 exposes the environment. Sections 3 and 4 expose the static and dynamic contracts, respectively. Section 5 contains the quantitative analysis. Section 6 concludes. The Appendix contains the proofs and the details on the data used for the calibration.

1.1 Literature Review

We derive optimal contracts combining LE and MH constraints. Our analysis is close to Atkeson (1991) and Thomas and Worrall (1994) who study lending contracts in international contexts. However, they model MH with respect to consuming or investing the borrowed funds, while we focus on risk management policies. Quadrini (2004) also combines MH and LE to study when and how contracts are renegotiation-proof. Similarly, Dovis (2019) shows that the combination of MH and LE can generate a region of *ex post* inefficiency. This is not the focus of our analysis as our contract is both *ex ante* and *ex post* efficient. In addition, Müller et al. (2019) study dynamic sovereign lending contracts with MH, with respect to reform policy efforts, and LE. Our approach is more general as it assumes a richer distribution support and distinguishes between local and global flexibility. We solve the optimal contract by means of the Lagrangian approach of Marcet and Marimon (2019) which has been widely used to account for LE (e.g. Kehoe and Perri (2002), Ferrari et al. (2024) and Wicht (2025)) and its combination with MH (e.g. Simpson-Bell (2020) and Ábrahám et al. (2025)).

Our research contributes to the literature on MH within dynamic macroeconomic models. Building upon the seminal work of Prescott and Townsend (1984), which demonstrated a constrained efficient allocation can be the allocation of competitive equilibrium if the space of contracts satisfy the corresponding incentive compatibility constraints, we extend the flexible MH approach introduced by Georgiadis et al. (2024) to a dynamic framework with a richer cost specification. We then compare our model's incentive structures with those in the canonical dynamic MH model in the spirit of Atkeson and Lucas (1992). These authors find that the cross-sectional consumption path spreads out indefinitely but without a downward trend. Hence immiseration arises at the individual and not at the societal level. In contrast, the flexible MH approach leads to what we term *bliss*, the antithesis of immiseration, and does not disrupt risk sharing, as incentives are not contingent on realized outcomes.[2] While

[2]Generally, studies on repeated MH highlight stark long-run implications for consumption. In unobserved

Georgiadis et al. (2024) focus on mean-based cost, we show that conclusions about optimal distribution and risk sharing completely reverse once one adopts an entropy-based cost. Finally, we bring the two MH approaches closer by back-loading incentives in the canonical one.

Our study offers an alternative interpretation of utility shocks used in quantitative models with discrete choices (Rust, 1988; Iskhakov et al., 2017).[3] We show that the flexible MH with entropy-based cost can rationalize the use of such shocks, as it endogenously generates logit-like distributions. Our analysis thus reframes the role of these shocks and provides guidance on calibrating the variance of the utility shock.

Our work more closely contributes to the recent literature on the design of an optimal stability Fund. Roch and Uhlig (2018), Liu et al. (2020) and Callegari et al. (2023) focus on the lender's side of the contract and therefore disregard MH issues. In opposition, Dovis and Kirpalani (2023) account for MH and show that the provision of effort is back-loaded. We build on Ábrahám et al. (2025), where defaultable sovereign debt is transformed into a safe Fund contract, which accounts for MH. They assume that the Fund has an exclusivity contract unlike Liu et al. (2020) and Callegari et al. (2023) who model a Fund which absorbs a minimal amount of debt. We exploit the fact that incentive compatibility constraints are disruptions to perfect risk-sharing. Our contribution is twofold. First, we provide a more comprehensive analysis of MH in the Fund design, describing and contrasting different provisions of incentives and their interactions with LE constraints. Importantly, we show the conditions under which Fund contracts can improve risk sharing or cannot. Second, we offer a quantitative exploration using a benchmark calibration for the Euro Area stressed countries.[4]

endowment economies, Green (1987) and Thomas and Worrall (1990) show that agents' consumption immiserates with almost all agents experiencing utility tending to negative infinity. Phelan (1998) argues that the critical determinant behind this result is the shape of the utility function: whether marginal utility remains bounded away from zero as consumption increases governs whether agents' utilities diverge negatively, positively, or spread without drift. In light of this, Phelan (1995) shows that limited commitment acts as an endogenous stopper for immiseration.

[3]For the use of utility shocks in sovereign debt models, see Mihalache (2020), Dvorkin et al. (2021), Mateos-Planas et al. (2022) and Mateos-Planas et al. (2023).

[4]As in Ábrahám et al. (2025) we have a theorem on existence (and uniqueness) of the constrained–efficient contract, but, in contrast with them, and as in Liu et al. (2020) and Callegari et al. (2023), we do not provide a market decentralization of the contract to obtain the constrained-efficient version of the two main welfare theorems.

2 Environment

Consider a small open economy with a single homogeneous consumption good in discrete time $t \in \{1, \ldots, T\}$. A benevolent government acts as a representative agent and takes decisions on behalf of the small open economy.

In each period, the government receives a stochastic endowment $y \in Y \equiv [y_1, \ldots, y_n]$ which is drawn from a probability distribution π. It is able to generate any *distribution* over Y. The choice of distribution is not contractible. Accordingly, we denote by $\mathcal{M}$ the set of set of Borel probability measures on Y and by δ_y the Dirac measures that assigns probability mass 1 to outcome $y \in Y$. Additionally, we define the borrower's capacity in the choice of distributions as *local flexibility*.

Definition 1 (Local Flexibility). *We say that the borrower enjoys local flexibility, when the target distribution $\ddot{\pi} \in \mathcal{N} \subset \mathcal{M}$ is such that for every $\pi \in \mathcal{M}$, there is some $\epsilon > 0$ for which $\ddot{\pi} + \epsilon(\pi - \ddot{\pi}) \in \mathcal{N}$ is feasible.*

Flexible MH corresponds to a distribution choice satisfying Definition 1. In opposition, canonical MH does not allow *local* but only *global* flexibility in the choice of distribution. That is the borrower loses its capacity to manipulate in an arbitrary way the relative likelihood of any collection of y.

The government discounts the future at the rate β, satisfying $\beta \leq 1/(1+r)$, where r is the risk-free world interest rate. The fact that the government is less patient than the lenders implies that it would like to front-load consumption. The government's utility can be defined by $U : \mathbb{R}^+ \times \mathcal{M} \to \mathbb{R}$ and is additively separable. So, if the government chooses a distribution π and consumes c then its instantaneous payoff is $U(c, \pi) \equiv u(c) - v(\pi)$. We make standard assumptions on preferences of consumption. For the distribution choice, we assume that the cost of effort is continuous, convex and differentiable. We denote by $v_\pi : Y \to \mathbb{R}$ and $\omega_\pi : Y \to \mathbb{R}$ the first and the second derivative of the cost function, respectively.

Assumption 1 (Monotonicity, Differentiability and Convexity). *The utility function from consumption, $u : \mathbb{R}^+ \to \mathbb{R}$, is continuous, strictly increasing and strictly concave. The cost of distribution, $v : \mathcal{M} \to \mathbb{R}$, is continuous, convex and differentiable.*

We additionally consider a flexible functional form governed by two key parameters $\alpha \geq 0$ and $\gamma \geq 0$. The latter implies the cost increases in first-order stochastic dominance, while the former implies the cost decreases in the entropy.

Assumption 2 (Monotonicity in Entropy and in First-Order Stochastic Dominance). *Let $\alpha \geq 0$, $\gamma \geq 0$ and $l : Y \to \mathbb{R}$ be a linearly increasing and continuous function. The* *distribution cost function is given by*

$$v(\pi) = -\alpha H(\pi) + \gamma \sum_{y \in Y} \pi(y) l(y).$$

where $H(\pi) = -\sum_{y \in Y} \pi(y) \log(\pi(y))$ *denotes the* entropy *of the distribution* π*.*

Note that an equivalent definition of first-order stochastic dominance is that "for any non-decreasing $u : \mathbb{R} \to \mathbb{R}$ there is mean dominance". Therefore, with $\alpha = 0 \wedge \gamma > 0$ and Assumption 2, $v(\pi)$ perfectly satisfies Georgiadis et al. (2024) assumption that the distribution cost increases with first-order stochastic dominance,[5] as well as their implicit assumption that there is no cost for *risk-reduction*. While with $\alpha > 0 \wedge \gamma = 0$, only the *risk-reduction* cost, in terms of entropy, is taken into account. In fact, the distribution with the highest entropy is the uniform distribution $1/|Y|$, whereas the Dirac measure δ_y has the lowest entropy (i.e. $H(\delta_y) = 0$). In sum, with $(\alpha, \gamma) > 0$, $v(\pi)$ accounts for both: increasing mean and risk reduction.

3 The Static Contract

To gain insight into the provision of incentives, we first consider that $T = 1$ meaning that the small open economy lives only one period. We first consider the case of flexible MH before linking it to the canonical MH where Definition 1 does not hold anymore.

3.1 Flexible Moral Hazard

Before deriving the principal-agent contracting problem between the Fund and the borrower, we determine the borrower's maximization problem in autarky as

$$\max_{\pi \in \mathcal{M}} \left\{ z(\pi) \equiv \sum_{y \in Y} \pi(y) u(y) - v(\pi) \right\}$$

Note now that an equivalent definition of second-order stochastic dominance is that "for any non-decreasing (not necessarily strictly) concave $u : \mathbb{R} \to \mathbb{R}$ there is mean dominance" and first-order stochastic dominance implies second-order stochastic dominance. Therefore,

[5]Note that the mapping $\pi \mapsto \sum_y \pi(y) u(y)$ is linear in π, which is also convenient. For the optimal distribution under non-linear cost, see Georgiadis et al. (2024, Corollary 3).

following the last comment, $z(\pi)$ with $\alpha > 0, \gamma \geq 0$ can be seen as generalized form of *mean-variance analysis* in which, as we next show, setting the variance to zero is not an optimal choice either.

Attaching the multiplier λ to the constraint, the first-order condition is given by

$$u(y) - v_\pi(y) = \lambda.$$

This condition holds for all y in the support of π. The borrower directly selects a distribution over outcomes, not merely an action that stochastically determines outcomes. This implies that the borrower can arbitrarily distort the relative likelihood of any subset of endowment realizations. Since these realizations are fully manipulable as of Definition 1, they carry no informational content. As a result, the provision of incentives cannot rely on realized endowments. Instead, incentives must be provided entirely through compensating the *marginal cost* of assigning probability mass to each outcome in the support of Y.

To see this, we can define the contracting problem between the borrower and the Fund. The contract specifies the implementation of a state contigent consumption $c(y)$ together with a distribution π resulting in a transfer to the Fund of $y - c(y)$. Formally,

$$\begin{aligned} \max_{c(y),\pi\in\mathcal{M}} \; & \mu_b \left[\sum_{y\in Y} \pi(y)u(c(y)) - v(\pi)\right] + \mu_l \left[\sum_{y\in Y} \pi(y)(y - c(y))\right] \\ \text{s.t.} \quad & \max_{\pi\in\mathcal{M}} \left\{\sum_{y\in Y} \pi(y)\, u(c(y)) - v(\pi)\right\} \\ & \sum_y \pi(y) = 1. \end{aligned} \tag{1}$$

We denote by χ the multiplier attached to (1). Given this, the maximization problem can be rewritten as

$$\max_{c(y),\pi\in\mathcal{M}} \; \mu_b \left[\sum_{y\in Y} \pi(y)u(c(y)) - v(\pi)\right] + \mu_l \left[\sum_{y\in Y} \pi(y)(y - c(y))\right] \tag{2}$$

$$\text{s.t.} \quad u(c(y)) - v_\pi(y) \geq \chi \text{ for all } y \in \operatorname{supp}(\pi), \tag{3}$$

where $\mu_b > 0$ and $\mu_l > 0$ are the Pareto weights assigned to the agent and the principal, respectively. By defining the IC constraint as (3), we use the first-order approach of Georgiadis et al. (2024). That is we replace the agent's full optimization problem with respect to π by its necessary and sufficient first-order condition.[6] Attaching the multiplier $\varepsilon(y)$ to the IC constraint (3), the first-order condition with respect to $c(y)$ is simply

$$u_c(c(y)) = \frac{\mu_l}{\mu_b + \frac{\varepsilon(y)}{\pi(y)}}.$$

[6] Convexity of v is not needed for necessity, only for sufficiency.

The first-order condition with respect to π is then

$$\mu_b \left[u(c(y)) - v_\pi(y)\right] + \mu_l \left[y - c(y)\right] - \varepsilon(y)\omega_\pi(y) = 0,$$

With contractible π, the contracting problem (2) would not have the IC constraint (3) and the first-order condition would be given by the first two terms, with the second one accounting for the social value of choosing the distribution. In contrast, with non-contractible π, as we assume, constraint (3) is present and the first term is equal to zero. The above expression therefore reduces to $\varepsilon(y)\omega_\pi(y) = \mu_b\chi + \mu_l \left[y - c(y)\right]$. Hence, the contract accounts for the external effect of effort on the Fund's value through its effect on the IC constraint.

These two first-order conditions are key in understanding the difference in allocation when the cost function depends on the entropy (i.e. $\alpha > 0$) or not (i.e. $\alpha = 0$). The following proposition characterize these two cases.

Proposition 1 (Optimal Allocation in Static Contract)**.** *Under Assumptions 1 and 2:*

1. *If $\alpha = 0$ and $\gamma > 0$, the optimal distribution π^* is such that $\pi^* = \delta_y$ for some $y \in Y$. Unless $y - c(y) < 0$ for $y \in supp(\pi^*)$, the agent is better off in autarky.*

2. *If $\alpha > 0$ and $\gamma \geq 0$, the optimal distribution π^* is such that $H(\pi^*) > 0$ and for all $y \in Y$*

$$\pi^*(y) = \frac{e^{(u(c(y)) - \gamma l(y))/\alpha}}{\sum_{y' \in Y} e^{(u(c(y')) - \gamma l(y'))/\alpha}} > 0. \tag{4}$$

with $c(y) > c(y')$ for all $y > y'$ and the agent is strictly better off than in autarky.

Under mean-based cost ($\alpha = 0$ and $\gamma > 0$), the optimal distribution is a Dirac meaning the borrower achieves perfect consumption smoothing. This is true with or without contract with the Fund. However, unless the Fund makes losses (i.e. $y - c(y) \leq 0$), the borrower is worse off contracting with the Fund as it must share part of the endowment y. We can therefore conclude that under Assumption 2 with $\alpha = 0$ and $\gamma > 0$ the contract is about rent-sharing – as opposed to risk-sharing – between the borrower and the Fund as the borrower can completely eliminate risk by itself. In other words, risk-sharing contracts are ineffective.

Under entropy-based cost ($\alpha > 0$ and $\gamma \geq 0$), risk reduction becomes costly. The conclusion we reach is reversed as the optimal distribution is interior and the Fund can make the agent strictly better off relative to autarky. In this case, the borrower does not find it optimal to perfectly smooth consumption as Dirac measures are too costly. Full consumption smoothing is also suboptimal for the Fund.[7] The Fund however offers partial insurance

[7] Under a uniform distribution $\pi = 1/|Y|$, the entropy cost is the lowest possible. Moreover, it holds that $\varepsilon(y)$ is constant for all y implying that $c(y)$ is constant from the first-order condition. This is however too costly to sustain for the Fund.

which improves welfare upon the autarkic allocation. The welfare improvement comes from the Fund providing insurance that reduces consumption variability, which allows the agent to economize on the entropy cost of distorting the distribution. Hence, risk-sharing contracts becomes effective when risk reduction is costly under flexible MH.

Note that the optimal distribution (4) takes the familiar logit form which is usually obtained through Gumbel-distributed utility shocks. Hence, the flexible MH friction with entropy-based offers an alternative interpretation of such shocks. In that logic, this distribution converges to a Dirac measure δ_y as $\alpha \to 0$.

3.2 Canonical Moral Hazard

So far, the borrower enjoyed local flexibility as per Definition 1. This is not anymore the case in the canonical MH as the agent's choice of distribution is restricted to a subset of $\mathcal{Q} \subset \mathcal{M}$. This means that the borrower can only affect the distribution *globally*. We model this restriction as follows. Define Q as a mixture distribution $Q(e) = \varpi(e)Q_L + (1-\varpi(e))Q_H$ for $Q_L, Q_H \in \mathcal{Q}$ with a weighting function $\varpi : [0,1] \to [0,1]$. The agent can manipulate the weights by choosing the (non-contractible) effort $e \in [0,1]$ to change Q. We assume $\varpi(e)$ is continuous, strictly decreasing and concave.

In this context, for the first-order approach of Rogerson (1985) to be valid, the cumulative distribution function of y' should be differentiable, convex and satisfy the monotone likelihood-ratio condition. This is the purpose of Assumption 3.

Assumption 3 (*Differentiability, Monotonicity and Convexity*)**.** *If $e \geq \tilde{e} > 0$ the ratio $\frac{Q(y'|\tilde{e})}{Q(y'|e)}$ is nonincreasing in y', and, for every e, $Q_p(e) = \sum_{p \subseteq Y} Q(y'|e)$ is differentiable in e, with $\partial_e Q_p(e) \leq 0$ and $\partial_e^2 Q_p(e) \geq 0$.*

The cost function still satisfies Assumption 2 with the implicit dependency on e. More precisely, it is given by $v(Q(e)) = -\alpha H(Q(e)) + \gamma \sum_{y \in Y} Q(y|e) l(y)$. The cost of effort therefore remains a mapping from $\mathcal{M}$ to $\mathbb{R}$. In autarky, the agent then solves

$$\max_{e \in [0,1]} \left\{ \sum_{y \in Y} Q(y|e)\, u(y) - v(Q(e)) \right\}. \tag{5}$$

Taking the first-order condition, we get

$$v_{Q(e)}(e) = \sum_{y \in Y} \frac{\partial Q(y|e)}{\partial e} u(y). \tag{6}$$

The marginal cost is now weighted by the expected marginal density. One could have considered an alternative function $\hat{v} : [0,1] \to \mathbb{R}$ which is a mapping from $[0,1]$ instead of

$\mathcal{M}$ and have obtained $\tilde{v}_e(e)$ as the marginal cost. This however only affects the quantitative property of the optimal effort e.

More importantly, the IC constraint in (6) relies on the informativeness of the realization of y'. This is because the information content of a specific realization can be directly measured by the relative likelihood given that Q is contractible, whereas e is not. In the case of flexible MH, the probability distribution is a choice variable which is not contractible precluding any informativeness of the shock realization. Hence, unlike flexible MH, canonical MH supposes the target distribution is known as the borrower can only affect the distribution *globally*.

Given this, the contracting problem between the borrower and the Fund under canonical MH reads

$$\max_{c(y),e\in[0,1]} \mu_b \left[\sum_{y\in Y} Q(y|e)u(c(y)) - v(Q(e))\right] + \mu_l \left[\sum_{y\in Y} Q(y|e)(y - c(y))\right]$$
$$\text{s.t.} \quad v_{Q(e)}(e) = \sum_{y\in Y} \frac{\partial Q(y|e)}{\partial e} u(c(y)).$$

Attaching the multiplier ι to the IC constraint, the first-order condition with respect to $c(y)$ is simply

$$u_c(c(y)) = \frac{\mu_l}{\mu_b + \frac{\partial_e Q(y'|e)\iota}{\pi(y)}}.$$

One directly notices the appearance of the first-derivative of $Q(e)$ in the determination of consumption. Moreover, the multiplier attached to (6) is non contingent. This reflects the main difference with the flexible MH approach we discussed previously. As the choice of distribution is restricted, the realization of y' is informative about the relative likelihood of exerting the optimal effort. In particular, $\partial_e Q(y'|e)$ measures how strongly the Fund can infer from y' that the borrower did not exert the assumed effort. Accordingly, the Fund will punish the borrower when a bad outcome realizes (i.e. $\partial_e Q(y'|e) < 0$) and will reward when a good outcome realizes (i.e. $\partial_e Q(y'|e) > 0$). Note that bad outcomes may occur here due to bad luck rather than a lack of effort. Thus, in terms of risk sharing, this provision of incentives is very costly for the borrower since it bears part of the risk of the endowment realization.

The first-order condition with respect to e is then

$$\mu_b \left[\sum_{y\in Y} \frac{\partial Q(y|e)}{\partial e} u(c(y)) - v_{Q(e)}(e)\right] +$$
$$\mu_l \sum_{y\in Y} \frac{\partial Q(y|e)}{\partial e} [y - c(y)] - \iota \left[\omega_{Q(e)}(e) - \sum_{y\in Y} \frac{\partial^2 Q(y|e)}{\partial e^2} u(c(y))\right] = 0,$$

This condition has a similar structure as for the flexible MH since it balances the marginal cost of effort with the benefits. The first term is the life-time utility benefit of effort to the borrower (which equates zero given (6)); the second term is the marginal benefit of effort to the Fund; the third term accounts for the marginal relaxation/tightening effect of the MH constraint (6) when there is a change in effort. Again, the contract accounts for the external effect of effort on the Fund's value through its effect on the IC constraint.

4 The Dynamic Contract

We turn to the infinite-horizon small open economy (i.e. $T \to \infty$) and analyze the dynamic contract which establishes a long-term relationship between the borrower and the Fund by defining a state-contingent sequence of consumption and shock distribution that maximises the life-time utility of both contracting parties given some initial conditions.

Besides MH, we introduce two-sided LE constraints. The dynamic contract is therefore self-enforcing and seeks to provide risk-sharing between the borrower and the Fund to the extent possible. Both LE and MH frictions preclude perfect risk-sharing in general.

4.1 Flexible Moral Hazard

As before we characterize the contract under flexible MH before turning to the contract under canonical MH.

4.1.1 The Constraints

Given the LE and MH frictions, the Fund has to account for three different constraints. The first one is the LE constraint of the borrower. For any $y^t, t \geq 0$, it should be that

$$\mathbb{E}_t \left[\sum_{j=t}^{\infty} \beta^{j-t} U(c(y^j), \pi_{j+1}) \right] \geq V^D(y_t). \tag{7}$$

The notation is implicit about the fact that expectations are conditional on the implemented distributions of $\{y^j\}_{j=t}^{\infty}$. At this stage, we do not specify the exact form of the borrower's outside option and simply assume it is strictly increasing in y_t.

The second constraint is the LE constraint of the Fund. For any $y^t, t \geq 0$, it should hold that

$$\mathbb{E}_t\left[\sum_{j=t}^{\infty}\left(\frac{1}{1+r}\right)^{j-t}\left(y_j - c(y^j)\right)\right] \geq Z(y_t). \tag{8}$$

The finite outside option of the Fund $Z(y_t) \leq 0$ is increasing in y_t and measures the extent of *ex post* redistribution the Fund is willing to tolerate. That is, if $Z(y_t) < 0$ the Fund is allowed to make a permanent loss in terms of lifetime expected net present value – i.e. the Fund can find better investment opportunities in the international financial market and if it does not renege it is because it has committed to sustaining $Z(y_t) < 0$. Clearly, the level of $Z(y_t)$ has an important impact on the amount of risk sharing in our environment and it can thus be interpreted as the extent of solidarity the Fund is willing to accept in state y_t, as in Tirole (2015).

Finally, the last constraint is the IC constraint. Define $V^b(y^t) = \mathbb{E}_t[\sum_{j=0}^{\infty}\beta^j U(c(y^{t+j}), \pi_{t+j+1})]$ as the value of the borrower at time t. For any $y^t, t \geq 0$ and a given consumption schedules $\{c(y^t)\}_{t=0}^{\infty}$, the optimal vector of distributions from the borrower's perspective is the outcome of

$$\max_{\pi_{t+1}\in\mathcal{M}} \sum_{t=0}^{\infty}\beta^t U(c(y^t), \pi_{t+1})$$
$$\sum_{y_{t+1}\in Y}\pi_{t+1}(y_{t+1}) = 1 \ \forall t \geq 0.$$

Attaching the multiplier χ_t to the constraint at time t, the first-order condition for all y_{t+1} in the support of π_{t+1} is given by

$$\beta V^b(y^{t+1}) - v_{\pi_{t+1}}(y_{t+1}) = \chi_t. \tag{9}$$

This condition is very similar to its static equivalent (3). The only difference is that the marginal benefit corresponds to the borrower's continuation value as opposed to its instantaneous utility.

4.1.2 The Contracting Problem

In its extensive form, the dynamic contract specifies that in state $y^t = (y_0, \ldots, y_t)$, the borrower consumes $c(y^t)$ and chooses the distribution π_{t+1}, resulting in a transfer to the Fund of $y_t - c(y^t)$. With two-sided LE and MH constraints, an optimal dynamic contract is a solution to the following problem

$$\max_{\{c(y^t),\pi_{t+1}\}} \mathbb{E}_0\left[\mu_{b,0}\sum_{t=0}^{\infty}\beta^t U(c(y^t), \pi_{t+1}) + \mu_{l,0}\sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^t\left[y_t - c(y^t)\right]\right] \tag{10}$$
$$\text{s.t.} \quad (7),\ (8), (9),\ \forall t \geq 0.$$

Note that $(\mu_{b,0}, \mu_{l,0})$ are the initial Pareto weights, which are key for our interpretation of the Fund contract as a risk-sharing contract. Given (7), (8) and (9), we take the following interiority assumption to ensure the uniform boundedness of the Lagrange multipliers.

Assumption 4 (Interiority). *There is an $\epsilon > 0$, such that, for all y_0 there is an allocation $\{\tilde{c}(y^t), \tilde{\pi}_t\}_{t=0}^{\infty}$ satisfying constraints (7) and (8) when, on the right-hand side, $V^D(y_t)$ and $Z(y_t)$ are replaced by $V^D(y_t) + \epsilon$ and $Z(y_t) + \epsilon$, respectively, and similarly, when in (9) $v_{\pi_{t+1}}(y^{t+1})$ is replaced by $v_{\pi_{t+1}}(y^{t+1}) + \epsilon$ and $=$ is replaced by $\geq$.*

For constraints (7) and (8), this assumption requires that, in spite of the LE constraints, there are strictly positive rents to be shared since otherwise there may not be a constrained-efficient risk-sharing contract. The last part of this assumption is satisfied if a distribution exists that generate a marginal benefit above the marginal cost.[8]

Following Marcet and Marimon (2019) and Mele (2011), we can rewrite the Fund contract problem as a saddle-point Lagrangian problem:

$$
\begin{aligned}
\text{SP} \min_{\{\gamma_b(y^t),\gamma_l(y^t),\xi(y^{t+1})\}} \max_{\{c(y^t),\pi_{t+1}\}} & \Bigg\{ \mathbb{E}_0 \Bigg[\sum_{t=0}^{\infty} \beta^t \Big(\mu_{b,t}(y^t) U\left(c(y^t), \pi_{t+1}\right) \\
& - \xi(y^{t+1}) \left(\chi_t + v_{\pi_{t+1}}(y^{t+1})\right) + \gamma_b(y^t) \big[U(c(y^t), \pi_{t+1}) - V^D(y_t)\big] \Big) \\
& + \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t \Big(\mu_{l,t+1}(y^t)[y_t - c(y^t)] - \gamma_l(y^t)\big[y_t - c(y^t) - Z(y^t)\big] \Big) \Bigg] \Bigg\} \\
\text{s.t.} \qquad & \mu_{b,t+1}(y^{t+1}) = \mu_{b,t}(y^t) + \gamma_b(y^t) + \xi(y^{t+1}), \\
& \mu_{l,t+1}(y^t) = \mu_{l,t}(y^t) + \gamma_l(y^t), \\
& \mu_{b,0}(y^0) \equiv \mu_{b,0}, \mu_{l,0}(y^0) \equiv \mu_{l,0} \text{ given},
\end{aligned}
$$

where $\gamma_b(y^t)$, $\gamma_l(y^t)$ and $\xi(y^{t+1})$ are the Lagrange multipliers of the LE constraints in (7) and in (8), and the IC constraint in (9), respectively.

The above formulation of the problem defines two new co-state variables $\mu_b(y^t)$ and $\mu_l(y^t)$, which represent the temporary Pareto weights of the borrower and the Fund respectively. These variables are initialized at the original Pareto weights $(\mu_{b,0}, \mu_{l,0})$ and become time-variant because of the LE and MH frictions. In particular, a binding LE constraint of the borrower (Fund) will imply a higher co-state variable of the borrower (Fund) so that it does

[8]The first part of the assumption can easily be satisfied since there are gains from risk-sharing in a contract between a risk-averse borrower and a risk-neutral Fund as long as autarky is not the unique constrained efficient allocation. The second part of the assumption is also easily met under various cost functions $v(\pi)$ if full risk sharing is not the only feasible allocation.

not leave the contract. In addition, the MH friction implies that the borrower's co-state variable will increase as $\xi(y^{t+1}) \geq 0$ under Assumption 4.

Given the homogeneity of degree one of the maximization problem in $(\mu_{b,t}, \mu_{l,t})$, only relative Pareto weights, defined as $x_t(y^t) \equiv \mu_{l,t}(y^t)/\mu_{b,t}(y^t)$, matter for the allocations, and this allows us to reduce the dimensionality of the co-state vector and write the problem recursively by using a convenient normalization. Let $\eta \equiv \beta(1+r) \leq 1$ and normalize the multipliers as follows

$$\nu_b(y^t) = \frac{\gamma_b(y^t)}{\mu_{b,t}(y^t)},\ \nu_l(y^t) = \frac{\gamma_l(y^t)}{\mu_{l,t}(y^t)} \text{ and } \varrho(y^t, y^{t+1}) = \frac{\xi(y^{t+1})}{\mu_{b,t}(y^t)}.$$

The Saddle-Point Functional Equation (SPFE) – i.e. the saddle-point version of Bellman's equation – is given by

$$\begin{aligned} FV(y,x) = \mathrm{SP} \min_{\{\nu_b,\nu_l,\varrho\}} \max_{\{c,\pi\}} \Big\{ & x\Big[(1+\nu_b)U(c,\pi) - \nu_b V^D(y)\Big] + \Big[(1+\nu_l)[y-c] - \nu_l Z(y)\Big] \\ & + \sum_{y' \in Y} \pi(y') \left[\frac{1+\nu_l}{1+r} FV(y', x'(y')) - x\varrho(y')(\upsilon_\pi(y') + \chi)\right] \Big\} \end{aligned} \tag{11}$$

$$\text{s.t.} \quad x'(y') \equiv \overline{x}'(y) + \hat{x}'(y') = \left[\frac{1+\nu_b}{1+\nu_l} + \frac{\varrho(y')}{1+\nu_l}\right]\eta x. \tag{12}$$

Equation (12) gives the law of motion of the relative Pareto weight in recursive form. The prospective weight $x'(y')$ can be separated into two parts: the update due to the borrower's LE constraint $\overline{x}'(y)$ and the update due to the IC constraint $\hat{x}'(y')$ (both relative to the Fund's updated weight). The value functions can be decomposed as follows

$$FV(y,x) = xV^b(y,x) + V^l(y,x) \text{ with} \tag{13}$$

$$V^l(y,x) = y - c + \frac{1}{1+r}\sum_{y' \in Y} \pi(y') V^l(y', x'(y')), \tag{14}$$

$$V^b(y,x) = U(c,\pi) + \beta \sum_{y' \in Y} \pi(y') V^b(y', x'(y')). \tag{15}$$

Taking the first-order condition with respect to c, one gets

$$u_c(c) = \frac{1+\nu_l}{1+\nu_b}\frac{1}{x}. \tag{16}$$

With respect to π, one gets

$$\frac{1}{1+r}V^l(y', x'(y')) + x'(y')\left[\beta V^b(y', x'(y')) - (\upsilon_\pi(y') + \chi)\right] + x(1+\nu_b)\chi = \pi(y')x\varrho(y')\omega_\pi(y').$$

Since (9) eventually holds with equality, this expression reduces to

$$\frac{1}{1+r}V^l(y', x'(y')) + x(1+\nu_b)\chi = \pi(y')x\varrho(y')\omega_\pi(y'). \tag{17}$$

These two first-order conditions are key in understanding the difference in allocation when $\alpha = 0$ or $\alpha > 0$. We reach similar conclusions about the optimal distribution in the dynamic contract as in the static contract. In that logic, Proposition 2 is the dynamic equivalent of Proposition 1.

Proposition 2 (Optimal Allocation in Dynamic Contract)**.** *Under Assumptions 1 and 2:*

1. *If $\alpha = 0$ and $\gamma > 0$, the optimal distribution π^* is such that $\pi^* = \delta_y$ for some $y \in Y$. Unless $Z(y) < 0$ for $y \in supp(\pi^*)$, the agent is strictly better off in autarky.*

2. *If $\alpha > 0$ and $\gamma \geq 0$, the optimal distribution π^* is such that $H(\pi^*) > 0$ and for all $y \in Y$*

$$\pi^*(y) = \frac{e^{(\beta V^b(y,x(y)) - \gamma l(y))/\alpha}}{\sum_{y' \in Y} e^{(\beta V^b(y',x(y')) - \gamma l(y'))/\alpha}} > 0. \tag{18}$$

with $c(y) > c(y')$ for all $y > y'$ and the agent is strictly better off than in autarky.

Under mean-based cost ($\alpha = 0$ and $\gamma > 0$), the optimal distribution is a Dirac meaning the borrower achieves perfect consumption smoothing. Similar to the static contract, unless the Fund makes losses, the agent is worse off contracting with the Fund as it must share part of the endowment y.

Under entropy-based cost ($\alpha > 0$ and $\gamma \geq 0$), the optimal distribution takes the familiar logit form where the borrower puts more probability mass where future continuation values are higher. Moreover, the dynamic contract helps the borrower smooth consumption.

The following lemma provides a characterization of the Lagrange multiplier attached to the IC constraint.

Lemma 1 (Bliss)**.** *Under Assumptions 1, 2 and 4, for $y \geq \tilde{y}$ it holds that $\varrho(y) > \varrho(\tilde{y}) \geq 0$. When $\eta = 1$ and $\nu_l(s) = 0$ in all states, $\mathbb{E}x'(y') \geq x(y)$.*[9]

The lemma is made of two parts. The first part states that the multiplier attached to the IC constraint is strictly increasing in y. This comes from two elements. On the one hand, the borrower needs to be compensated for incurring a larger cost. This is the case when $\gamma > 0$ as the cost is larger for larger y. On the other hand, the value of the Fund increases in y as there is more surplus to share. The multiplier therefore needs to increase to satisfy (17).

[9] Note that a sufficient condition for $\nu_l(y) = 0$ is that $Z(y)$ is negative enough for all y.

The second part of the lemma states that the law of motion of the relative Pareto weight is a left bounded positive submartingale. To see this, take the expectations of the law of motion

$$\mathbb{E}x'(y') \equiv \mathbb{E}_t\left[\bar{x}'(y) + \hat{x}'(y')\right] = \mathbb{E}\left[\frac{1+\nu_b(y)}{1+\nu_l(y)}x(y) + \frac{\varrho(y')}{1+\nu_l(y)}x(y)\right]\eta,$$

where $\hat{x}'(y')$ accounts for the dynamic effect of the IC constraint. There are two forces working against the multiplier on the IC constraint $\varrho(y')$: impatience $\eta \leq 1$ and the Fund's LE constraint $\nu_l(y) \geq 0$. Hence, with neither the Fund's LE constraint nor impatience, we get that $\mathbb{E}x'(y') \geq x(y)$. This means that without either one of these two elements, the relative Pareto weight increases as a submartingale. This is the reverse of the standard *immiseration* result in Green (1987), Thomas and Worrall (1990), Atkeson and Lucas (1992) and Phelan (1998). We call it *bliss*.

This last result is important as it uncovers a different provision of incentives than the canonical MH. The Fund does not rely on the realizations of y' to provide incentives but compensates the borrower for the cost it incurred. Such mechanism is more risk sharing friendly as there is no punishment for low realizations of y'. In the next section we analyze such problem and contrast it with the outcome of Lemma 1.

For completeness of the argument, we show existence and uniqueness extending the proof of Marcet and Marimon (2019) and Ábrahám et al. (2025) to our environment.

Proposition 3 (Existence and Uniqueness)**.** *Given Assumptions 1, 2 and 4, for any y_0, $x(y_0)$, and outside options $\{V^D(y), Z(y)\}$, there is a unique recursive constrained-efficient dynamic contract.*

4.2 Canonical Moral Hazard

We switch to the canonical MH meaning that the borrower loses its capacity to manipulate in an arbitrary way the relative likelihood of any collection of y' as given by Definition 1. We analyze which implication this has for the dynamic contract.

4.2.1 The Contracting Problem

Following the previous section, we can formulate the IC constraint as

$$v_{Q(e)}(e) = \beta \sum_{y' \in Y} \frac{\partial Q(y'|e)}{\partial e} V^b(y', x'(y')). \tag{19}$$

Similarly, we find that the SPFE is given by

$$\begin{aligned}
FV(y,x) = \text{SP} \min_{\{\nu_b,\nu_l,\varrho\}} \max_{\{c,e\}} \Big\{ & x\Big[(1+\nu_b)U(c,Q(e)) - \nu_b V^D(y) - \varrho v_{Q(e)}(e)\Big] \\
& + \big[(1+\nu_l)(y-c) - \nu_l Z(y)\big] + \frac{1+\nu_l}{1+r}\sum_{y'\in Y} Q(y'|e)FV(y',x'(y'))\Big\} \\
\text{s.t. } & x'(y') \equiv \bar{x}'(y) + \hat{x}'(y') = \left[\frac{1+\nu_b}{1+\nu_l} + \frac{\varphi(y')}{1+\nu_l}\right]\eta x, \\
& \varphi(x') = \varrho\frac{\partial_e Q(y'|e)}{Q(y'|e)}.
\end{aligned} \tag{20}$$

The Fund's value functions can be decomposed as in the case with flexible MH. Similarly, the policy functions for consumption is the solution to (16). This is because of additive separability in the utility function. Hence, the formulation of the MH does not directly affect the formulation of this first-order condition.

With the need of risk-sharing, it is reasonable to model contracts with LE constraints that satisfy the following property.

Definition 2. *The LE constraints (7) and (8) satisfy the 'no-free-lunch condition' if, given (y,x), whenever $\nu_b > 0$, then $\partial_e Q(y'|e) > 0$ and whenever $\nu_l > 0$, then $\partial_e Q(y'|e) < 0$, respectively.*

Conversely, if $\partial_e Q(y'|e) = 0$ (or the inequality signs were reversed) exercising more effort would not have any effect on the LE constraints (or a perverse effect) and, on those grounds, MH would not be an issue. The following lemma provides a characterization of the interaction between LE and MH constraints.

Lemma 2 (IC Constraint Multiplier in Canonical MH)**.** *Under Assumptions 1 and 3,*

1. *LE constraints have an effect on the expected law of motion of the Pareto weights, when they are binding; in contrast, the IC constraints does not, even if they bind; i.e.* $\mathbb{E}x' = \bar{x}'(y)$.

2. *If LE constraints satisfy Definition 2, the IC constraint make the borrower's LE constraint (7) more likely to bind and the Fund's LE constraint (8) less likely to bind and, in both cases,* $\mathbb{E}_t\frac{1}{u'(c')}$ *increases.*

To see the first point, note that, given (20), $\sum_{y'\in Y} Q(y'|e)\hat{x}'(y') = 0$, since independently of effort we have $\sum_{y'\in Y} Q(y'|e) = 1$ implying that $\sum_{y'\in Y} \partial_e Q(y'|e) = 0$. Therefore $\mathbb{E}\hat{x}' = 0$ and $\mathbb{E}x' = \bar{x}'(s)$. Alternatively, the expected law of motion of x can also be expressed as

$$\mathbb{E}x' = \mathbb{E}\left[\frac{1}{u'(c')}\frac{1+\nu_l'(y',x'(y'))}{1+\nu_b'(y',x'(y'))}\right] = \frac{1}{u'(c)}\eta,$$

where the last equality is the *inverse Euler equation* of the recursive contract (Ábrahám et al. (2025), Lemma 4). This result has to be properly understood. At the individual level, the consumption path spreads out depending on whether the borrower draws a good endowment (i.e. $\partial_e Q(y'|e) > 0$) or a bad endowment (i.e. $\partial_e Q(y'|e) < 0$). However, over the long-run, these effects cancel out each other. Hence, similar to Atkeson and Lucas (1992), immiseration arises at the individual level and not the aggregate.

To see the second point, note that, since the LE constraint multipliers are either zero or at most one of the two is positive, we have that

$$\mathbb{E}_t \frac{1}{u'(c')} = \mathbb{E}_t \left[x' \frac{1 + \nu_b'(y', x'(y'))}{1 + \nu_l'(y', x'(y'))} \right].$$

If LE constraints satisfy Definition 2, the borrower's LE constraint is more likely to bind, while the Fund's LE constraint is less likely to bind and, as a result, in both cases expected consumption increases.

Under Assumption 3 and suitable adaptation of Assumption 4, Ábrahám et al. (2025) show existence and uniqueness of the dynamic contract under canonical MH. The interested reader can directly refer to this study.

4.2.2 Back Loading Incentives

In the canonical MH problem, the provision of incentives is generated by a system of rewards and punishments associated with the moral-hazard constraint (19). Such system hinders consumption smoothing even across the states in which the LE constraint is not binding. In opposition, in the flexible MH, the dynamic contract does not hinder consumption smoothing. We aim at improving the risk sharing property of the canonical MH with a Fund contract consisting of an infinite sequence of subcontracts, whereby within each subcontract rewards and punishments are *back-loaded* to the end. The main idea behind back-loading is that within the subcontract the borrower enjoys perfect consumption smoothing adjusted for the relative impatience of the borrower.

The length of each subcontract is directly determined by the binding LE constraints. The reason is that whenever a subcontract would violate one of the LE constraints, one of the contracting parties would find it optimal to terminate the contract. Hence, the binding LE constraints endogenously determine the subcontract's end. When this happens, we say that the subcontract resets.

The Fund contract can be expressed as the solution to a sequence of subcontracts. As a subcontract resets when one of the LE constraint binds, the start of a subcontract is such that

$$
\begin{aligned}
\hat{FV}(y,x) = \min_{\{\nu_b,\nu_l,\varrho\}} \max_{\{c,e\}} \Bigg\{ & x\big[(1+\nu_b)(u(c)-v(Q(e))) - \nu_b V^D(y) - \varrho v_{Q(e)}(e)\big] \\
& + [(1+\nu_l)(y-c) - \nu_l Z(y)] \\
& + \frac{1+\nu_l}{1+r}\mathbb{E}\left[\mathbb{I}_{\{(y',x'(y'))\}}\hat{FV}(y',x'(y')) + (1-\mathbb{I}_{\{(y',x'(y'))\}})\overline{FV}(y',x'(y'),\bar{x}') \mid e\right] \Bigg\} \\
\text{s.t.} \quad & x'(y') = \eta x \frac{1+\nu_b+\varphi(y')}{1+\nu_l}, \\
& \bar{x}' = \eta x \frac{1+\nu_b}{1+\nu_l},
\end{aligned}
$$

where $\varphi(y') = \varrho\frac{\partial_e Q(y'|e)}{Q(y'|e)}$ as before and $\mathbb{I}_{\{(y',x'(y'))\}}$ is an indicator function which takes value one when one of the LE constraints is binding – i.e. $\nu_b(y',x'(y')) + \nu_l(s',x'(g')) > 0$. In other words, $\mathbb{I}_{\{(y',x'(y'))\}} = 1$ indicates when the subcontract resets. Then within the subcontract

$$
\begin{aligned}
\overline{FV}(y,x,\bar{x}) = \min_{\{\varrho\}} \max_{\{c,e\}} \Bigg\{ & \bar{x}u(c) - x\big[v(Q(e)) + \varrho v_{Q(e)}(e)\big] + (y-c) \\
& + \frac{1}{1+r}\mathbb{E}\left[\mathbb{I}_{\{(y',x'(y'))\}}\hat{FV}(y',x'(y')) + (1-\mathbb{I}_{\{(y',x'(y'))\}})\overline{FV}(y',x'(y'),\bar{x}') \mid e\right] \Bigg\} \\
\text{s.t.} \quad & x'(y') = \eta x\left[1+\varphi(y')\right], \\
& \bar{x}' = \eta \bar{x}.
\end{aligned}
$$

As it can be seen, within the subcontract, $x'(y')$ is the *latent multiplier* which cumulates past incentives. The consumption policy is given by the same first-order condition, (16), resulting in $c(y,\overline{x})$, while the optimal effort $e(y,x)$ requires a reformulation of the original first-order condition. For this, it is useful to recall that, as in the benchmark Fund contracts, $\hat{FV}(y,x) = x\hat{V}^b(y,x) + V^l(y,x)$. However, $\overline{FV}(y,x,\overline{x})$ depends on $\overline{x}$ and x; therefore, we first decompose

$$
\begin{aligned}
\overline{V}^b_1(y,\overline{x}) &= u(c(y,\overline{x})) + \beta\mathbb{E}\left[\overline{V}^b_1(y',x'(y'))\right], \\
\overline{V}^b_2(y,x) &= -v(e(y,x)) + \beta\mathbb{E}\left[\overline{V}^b_2(y',x'(y'))\right],
\end{aligned}
$$

then the value of the borrower is simply $\overline{V}^b(y,x,\overline{x}) = \overline{V}^b_1(y,\overline{x}) + \overline{V}^b_2(y,x)$ implying that

$$
\overline{FV}^b(y,x,\overline{x}) = \overline{x}\,\overline{V}^b_1(y,\overline{x}) + x\overline{V}^b_2(y,x) + \overline{V}^l(y,x,\overline{x}).
$$

Note that, except for the distinction between x and $\overline{x}$, $\overline{FV}(y,x,\overline{x})$ is the same as $\hat{FV}(y,x)$ since we can always incorporate in the minimization $\{\nu_b,\nu_l\}$ which will satisfy $\nu_b = \nu_l = 0$, by construction, within the subcontract. This allows us to express a unique first-order condition for the effort policy.

The presence of subcontracts enhances risk sharing compared to the canonical MH. The reason behind this is that, within a subcontract, there is perfect consumption smoothing adjusted for the relative impatience of the borrower. As one can see from the law of motion of $\bar{x}$ together with (16), consumption is entirely deterministic and decays at rate η as long as the subcontract runs. Conversely, when the subcontract resets, consumption is adjusted up if the borrower's LE constraint binds and is adjusted down otherwise. More precisely, when the endowment increases, the borrower's LE constraint binds due to the strict monotonicity of $V^D(y)$. This triggers the end of the subcontract and a larger borrower's value since $\nu_b(y, x) > \nu_b(\tilde{y}, x)$. Hence, there is no (net) punishment when the borrower's LE constraint is binding. This design makes the dynamic contract more risk-sharing friendly because it is the infrequent binding of enforcement, rather than the realized endowment, that governs when consumption departs from its deterministic path.

It is important to note that the back-loaded structure mirrors the existing framework of sovereign lending programs implemented by international multilateral lenders, such as the International Monetary Fund. These lenders offer relatively short-term lending programs which are often followed by subsequent arrangements, each contingent on a new risk assessment that takes into account the borrowing country's previous performance. This iterative process ensures that the lending is aligned with the evolving economic conditions and reform progress of the recipient country, thereby aiming to enhance the effectiveness of the financial support provided.

5 Quantitative Analysis

We extend the dynamic contract to endogenous labor supply, production, and an autarky outside option with incomplete markets and defaultable debt (IMD), and calibrate it to the Euro Area stressed countries (Portugal, Italy, Greece and Spain) to compare the different contracts in terms of business cycle properties and welfare.

The calibration serves two distinct purposes, and we keep them separate throughout. The IMD calibration is *positive*: each MH specification is calibrated on its own terms, letting its parameters adjust freely, to check whether it has empirical content and can match the data. The Fund calibration is *normative*: we fix a common set of parameters across contracts – taken from the canonical MH calibration – so that differences across Fund contracts reflect the contract design alone, and welfare comparisons are meaningful.

5.1 The Quantitative Dynamic Contract

For the quantitative model, we expand the Fund contract in several dimensions. We expose these changes in the Fund under flexible MH. The other Fund contracts are derived in the Appendix. First, the borrower can produce goods using a decreasing-returns labour technology $y = \theta f(n)$, where $f'(n) > 0$, $f''(n) < 0$, $n \in [0,1]$ denotes labor and θ is a productivity shock. The shock is composed of two parts

$$\theta = \varkappa\zeta + (1-\varkappa)g,$$

where $g \in G$, $\zeta \in L$ and $s \equiv (g, \zeta)$. The borrower can generate any distribution π^g over G, while the distribution π^ζ over L is exogenous. In other words, only part of the shock (i.e. g) is endogenous or semi-endogenous, while the rest (i.e. ζ) is purely exogenous.

Labor supply is endogenous. Accordingly, the instantaneous utility is given by $U : \mathbb{R}^+ \times [0,1] \times \mathcal{M} \to \mathbb{R}$. So, if the borrower chooses a distribution π^g, supplies labor n and consumes c then its payoff is $U(c, n, \pi^g) \equiv u(c) + h(1-n) - v(\pi^g)$ where $1-n$ denotes leisure.

Second, we assume the borrower's outside option corresponds to the autarky value in a IMD economy,[10]

$$V^D(s) = \max_{n,\pi^g} \Big\{ U(\theta^d f(n), n, \pi^g) + \beta \sum_{g' \in G} \pi^g(g') \sum_{\zeta' \in L} \pi^\zeta(\zeta'|\zeta) \big[(1-\varsigma)V^D(s') + \varsigma J(s', 0)\big] \Big\},$$

where $\theta^d \leq \theta$ contains the penalty for defaulting and $\varsigma \geq 0$ is the probability to re-access the private bond market. Furthermore, $J(\cdot)$ corresponds to the value of reintegrating the private bond market without the Fund. More precisely, $J(s,b) = \max_{D \in \{0,1\}} \{(1-D)V^P(s,b) + DV^D(s)\}$, with

$$V^P(s,b) = \max_{\{c,n,\pi^g,b'\}} \left\{ U(c,n,\pi^g) + \beta \sum_{g' \in G} \pi^g(g') \sum_{\zeta' \in L} \pi^\zeta(\zeta'|\zeta) \big[J(s',b')\big] \right\} \tag{21}$$
$$\text{s.t. } c + q(s,b')(b' - \delta b) \leq \theta f(n) + (1 - \delta + \delta\kappa)b.$$

In the private bond market, the government can borrow long-term defaultable bonds, b', at a unit price of $q(s,b')$. A fraction $1-\delta$ of each bond matures today and the remaining fraction δ is rolled-over and pays a coupon κ. Private lenders are competitive and the price of one unit of private bond is given by $q(s,b') = \frac{1}{1+r} \sum_{g' \in G} \pi^g(g') \sum_{\zeta' \in L} \pi^\zeta(\zeta'|\zeta)(1 - D(s',b'))[1 - \delta + \delta\kappa + \delta q(s',b'')]$ where $D(\cdot)$ is the default policy taking value one in case of default and zero otherwise.

[10]See for instance Aguiar and Gopinath (2006), Arellano (2008), Chatterjee and Eyigungor (2012) and Aguiar and Amador (2021).

In this extended environment with endogenous labor supply and IMD autarky value, the Fund's contract in sequential form is given by

$$\max_{\{c(s^t),n(s^t),\pi^g_{t+1}\}} \mathbb{E}_0\left[\alpha_{b,0}\sum_{t=0}^{\infty}\beta^t U(c(s^t),n(s^t),\pi^g_{t+1}) + \alpha_{l,0}\sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^t\left[\theta_t f(n(s^t)) - c(s^t)\right]\right]$$

$$\text{s.t.}\quad \mathbb{E}_t\left[\sum_{j=t}^{\infty}\beta^{j-t}U(c(s^j),n(s^j),\pi^g_{j+1})\right] \geq V^D(s_t)$$

$$\mathbb{E}_t\left[\sum_{j=t}^{\infty}\left(\frac{1}{1+r}\right)^{j-t}\left(\theta_j f(n(s^j)) - c(s^j)\right)\right] \geq Z(s_t)$$

$$\beta\sum_{\zeta^{t+1}\in Z}\pi^\zeta(\zeta^{t+1}|\zeta_t)V^b(s^{t+1}) - v_{\pi^g_{t+1}}(g^{t+1}) = \chi_t.$$

As one can see the different constraints are easily adapted to the extensions we consider. The Fund's contract in recursive form is then

$$FV(s,x) = \text{SP}\min_{\{\nu_b,\nu_l,\varrho\}}\max_{\{c,n,\pi^g\}}\Bigg\{x\left[(1+\nu_b)U(c,n,\pi^g) - \nu_b V^D(s)\right] + \left[(1+\nu_l)[\theta f(n) - c] - \nu_l Z(s)\right]$$
$$+\sum_{g'\in G}\pi^g(g')\left[\sum_{\zeta'\in Z}\pi^\zeta(\zeta'|z)\frac{1+\nu_l}{1+r}FV(s',x'(g')) - x\varrho(g')\left(v_{\pi^g}(g')+\chi\right)\right]\Bigg\}$$
$$\text{s.t.}\qquad x'(g') \equiv \overline{x}'(s) + \hat{x}'(g') = \left[\frac{1+\nu_b}{1+\nu_l} + \frac{\varrho(g')}{1+\nu_l}\right]\eta x.$$

5.2 Calibration

Parameters are chosen to match moments of the Euro Area stressed countries over 1980–2019, with the model period set to one year. Throughout, CMH stands for Canonical MH, BMH for back-loaded MH, FeMH for entropy-cost flexible MH and FmMH for mean-cost flexible MH; we denote the primary surplus by $\tau = \theta f(n) - c$ and output by $y = \theta f(n)$. As explained above, the IMD calibration lets each specification's parameters adjust on its own terms (a positive exercise), while the Fund calibration imposes a common parameterization across contracts (a normative exercise). The Appendix contains more information about the data.

Table 1 lists all parameters, split into a first group common to all MH specifications (drawn from the literature or the data) and a second group specific to each specification, each targeting specific moments of the IMD economy.

The borrower's utility is additively separable in consumption, leisure and effort, with $u(c) = \frac{c^{1-\sigma_c}-1}{1-\sigma_c}$ and $h(1-n) = \upsilon\frac{(1-n)^{1-\sigma_l}-1}{1-\sigma_l}$. We set σ_c to the standard literature value and choose υ and σ_l to match the average n and the correlation of n with y. The discount factor

Table 1: Parameters

Parameter	Value			Description	Targeted Moment
A. Literature					
σ_c	1			Risk aversion	
Z	0			Fund mutualization	
B. Data					
r	0.0198			Risk-free rate	Average German short-term rate
δ	0.814			Bond maturity	Average bond maturity
κ	0.0760			Bond coupon rate	Average bond coupon rate
ℓ	0.5696			Labor share	Average labor share
$\varkappa$	0.3769			Share of ζ shock	
ρ_ζ	0.9784			Autocorrelation ζ shock	
ρ_g	0.9947			Autocorrelation g shock	Labor productivity
std(v)	0.0264			Standard deviation ζ shock	
std(u)	0.021			Standard deviation g shock	
C. Model	CMH	FmMH	FeMH		
β	0.943	0.926	0.938	Discount factor	Average $-b/y$
σ_l	0.1	0.1	0.1	Labor elasticity	Correlation n and y
υ	1.525	1.528	1.524	Leisure utility weight	Average n
γ	-	0.22	-	Effort disutility (mean)	Correlation c and y
α	0.041	-	0.040	Effort disutility (entropy)	
ψ	0.281	0.795	0.7995	Output default cost	Average spread
ς	0.33	0.025	0.02	Market re-access probability	Relative volatility c
ϵ	0.00001	0.00001	0.00001	Utility shock variance	Convergence

Note: CMH stands for Canonical MH, BMH for back-loaded MH, FeMH for entropy-cost flexible MH and FmMH for mean-cost flexible MH. See the Appendix for more information about the data.

β matches the average debt ratio, the risk-free rate r the average German short-term real rate (2000-2019), and the bond parameters (δ, κ) the average maturity and coupon rate of debt. The probability of market re-access ς matches the relative volatility of consumption, and the (asymmetric) default penalty $\theta^d = \min\{\theta, (1-\psi)\underline{\theta} + \psi\overline{\theta}\}$, as in Arellano (2008), has ψ chosen to match the average spread.

Following Ábrahám et al. (2025), Liu et al. (2020) and Callegari et al. (2023), the Fund's participation constraint is set to $Z = 0$: the Fund is not built on an assumption of solidarity, and so does not require expected permanent transfers to the borrower.

The production technology is $f(n) = n^\ell$ with the labor share ℓ equating the average labor share in the panel of countries for the period considered. For the labor productivity, $\theta = \varkappa\zeta + (1-\varkappa)g$, we estimate two AR(1) processes treating them as latent state variables in a state-space model. More precisely, $\zeta_t = \rho_\zeta \zeta_{t-1} + v_t$ is a common shock and $g_{i,t} = \rho_g g_{i,t-1} + u_{i,t}$ is an idiosyncratic shock at time t for country i. Since our model does not have any capital accumulation, we use the time series for the labor productivity $\theta_{i,t}$ for the four Euro Area stressed countries between 1980 and 2019. The estimated parameters are reported in Table 1. We further discretize the shock process using the method of Rouwenhorst (1995) with 7 grid points for each process leading to a total of 49 labor productivity states θ.

This delivers the transition matrix $Q = \varpi(e)Q_L + (1 - \varpi(e))Q_H$, $Q_L, Q_H \in \mathcal{Q}$, but the decomposition into Q_H and Q_L is not unique. We pin it down by requiring Q_L to be maximally entropy-dominant over Q_H. To achieve this, we fix $\varpi(\bar{e}) = 0.3$ for the average effort $\bar{e}$ and determine α accordingly to match the correlation between c and y. We then apply a mass transfer algorithm detailed in to recover Q_H and $Q_L = \frac{Q-(1-\varpi(e))Q_H}{\varpi(e)}$. We use $\varpi(e) = 1 - e$, which gives $\frac{\partial Q(g'|g,e)}{\partial e} = Q_L(g'|g) - Q_H(g'|g)$ and $\frac{\partial^2 Q(g'|g,e)}{\partial e^2} = 0$. For the cost function, $v(\pi) = \gamma \sum_{y \in Y} \pi(y)l(y)$ under mean-cost flexible MH and $v(\pi) = -\alpha H(\pi)$ under the other specifications.

As noted by Chatterjee and Eyigungor (2012), the computation of the IMD economy with long-term debt requires some form of randomization. Building on the most recent literature, we introduce a utility shock to the default choice which follows a Type 1 Extreme Value (i.e. Gumbel) distribution with a scale parameter ϵ. We pick the smallest possible value of ϵ leading to convergence.

5.3 IMD Outcome

We now ask whether each MH specification, calibrated on its own terms, has empirical content. Table 2 reports the calibration outcome for the three MH specifications.

All three specifications match the targeted means well, with the entropy-cost flexible MH giving the lowest mean absolute deviation, followed by the mean-cost flexible MH. Matching these means requires more impatience and a lower output cost of default under flexible MH. The targeted second moments are harder to match uniformly: despite a low probability of regaining market access, both flexible MH specifications have difficulty to generate consumption volatility relative to output as high as under canonical MH, though they match the labor-output correlation that canonical MH misses.[11]

[11] Notice that we could not reduce σ_l below 0.1 as the bisection algorithm becomes unstable.

Table 2: IMD Outcome

Variables	Targeted	Data	IMD Economy		
			CMH	FmMH	FeMH
A. Means					
b'/y (%)	×	85.70	85.96	85.80	86.38
n (%)	×	36.09	36.10	36.09	36.12
Spread (%)	×	2.20	2.22	1.99	2.20
B. Volatilities					
std(spread)		1.20	1.49	1.14	0.95
std(c)/std(y)	×	0.99	0.96	0.84	0.78
std(n)/std(y)		0.99	0.28	0.53	0.47
std(τ/y)/std(y)		1.03	0.93	1.85	1.56
C. Correlations					
corr(c, y)	×	0.80	0.97	0.62	0.70
corr(n, y)	×	0.66	0.29	0.60	0.64
corr($\tau/y, y$)		0.17	0.24	0.55	0.59
corr(spread, y)		-0.18	0.08	-0.14	-0.11
D. Mean absolute deviations					
Targeted moments			0.14	0.09	0.06
All moments			0.34	0.40	0.41

Note: CMH stands for Canonical MH, BMH for back-loaded MH, FeMH for entropy-cost flexible MH and FmMH for mean-cost flexible MH. $\tau \equiv \theta f(n) - c$ corresponds to the primary surplus and $y = \theta f(n)$ to the output. Mean absolute deviation is computed as $\sum_i |\frac{\hat{m}_i - m_i}{m_i}|$ where m_i denotes the ith data moment and $\hat{m}_i$ the model counterpart. See the Appendix contains more information about the data.

Untargeted moments are more mixed, and no specification dominates: canonical MH has the best overall fit, closely followed by the other two, and all three miss the volatility of labor relative to output. Canonical MH matches the spread's volatility and the primary surplus' second moments well but produces a pro-cyclical rather than counter-cyclical spread. The two flexible MH specifications get the spread's volatility and correlation right but miss the primary surplus moments.

Figure D.1 in the Appendix depicts a simulated path of debt and g in steady state for the three IMD economies.

5.4 Fund Outcome

We now turn to the Fund economy and switch to a normative approach: all Fund contracts share the same parameters, taken from the canonical MH calibration, so that differences across contracts reflect design rather than parameterization. The only adjustment is $\gamma = 0.5154$, set so that the ex post cost of effort equates across the two flexible MH specifications. We first discuss the key policy functions, then the steady-state allocation of the different Fund contracts, and finally welfare.

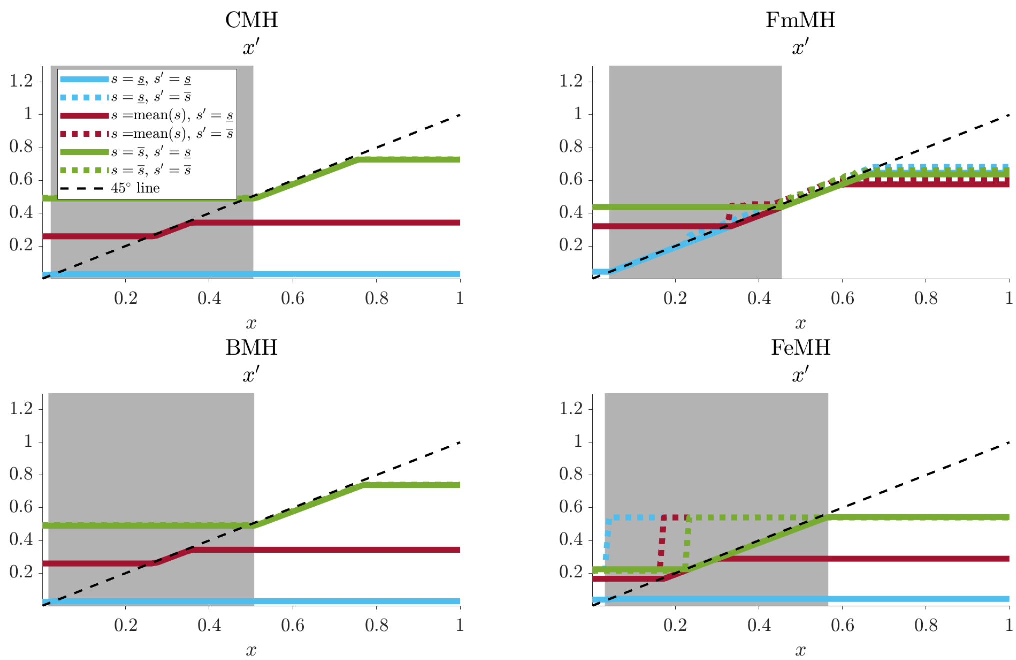

Note: The figure depicts the policy of relative Pareto weights. The gray region corresponds to the ergodic set.

Figure 1: Relative Pareto Weights and Ergodic Set

Figure 1 depicts the dynamic of the relative Pareto weight. The red lines relate to the highest, the green lines to the mean and the blue lines to the lowest values of s. The dotted lines relate to the highest value of s' and the solid lines to the lowest value of s'. The figure also represents the ergodic set of the relative Pareto weights for the different Fund contracts in gray. The ergodic set gives the steady state of the Fund contract in which we conduct simulations. In each specification, the horizontal line on the left hand side is determined by the borrower's binding LE constraint, while the horizontal line on the right hand side is determined by the Fund's binding LE constraint. The line rejoining both horizontal lines is determined by the allocation when none of the LE constraints binds. Two differences stand out between canonical and flexible MH. First, the ergodic set is the widest under entropy-cost flexible MH, followed by canonical MH. Second, the multiplier on the IC constraint is larger under both flexible MH specifications: the dotted and solid lines are almost indistinguishable under canonical MH but clearly separated under flexible MH, especially entropy-cost. This means that the IC multiplier is generally small under canonical MH, which is also why canonical and back-loaded MH generate similar business cycle moments as we see next.

Table 3: Fund outcome

Variables	Fund Economy				Fund relative to IMD (Δ)			
	CMH	FmMH	FeMH	BMH	CMH	FmMH	FeMH	BMH
A. Means								
mean(y)	0.217	0.418	0.425	0.216	0.004	-0.027	-0.018	0.003
mean(g)	0.367	0.712	0.977	0.367	0.000	-0.292	-0.010	0.000
mean(c)	0.210	0.316	0.424	0.211	-0.001	-0.116	0.007	0.000
mean(n)	0.366	0.443	0.358	0.364	0.005	0.077	-0.021	0.003
mean(τ)	0.008	0.102	0.001	0.006	0.005	0.089	-0.025	0.003
mean(H)	0.091	0.000	0.644	0.091	0.000	0.000	-0.363	0.000
B. Volatilities								
std(y)	0.062	0.262	0.048	0.061	0.002	0.172	-0.020	0.002
std(g)	0.146	0.424	0.126	0.146	0.000	0.233	-0.022	0.000
std(c)	0.053	0.047	0.044	0.054	-0.004	0.026	-0.007	-0.003
std(n)	0.016	0.280	0.004	0.018	-0.001	0.208	-0.036	0.001
std(τ)	0.011	0.255	0.004	0.013	-0.003	0.171	-0.054	-0.001
std(H)	0.011	0.000	0.121	0.011	0.000	0.000	-0.018	0.000
C. Correlations								
corr(g, y)	0.865	0.975	0.862	0.861	0.014	0.094	-0.021	0.011
corr(c, y)	0.988	0.230	0.999	0.977	0.022	-0.192	0.415	0.011
corr(n, y)	0.658	0.980	0.928	0.582	0.371	0.020	0.273	0.295
corr(τ, y)	0.746	0.983	0.938	0.648	0.446	0.023	0.264	0.349
corr(H, y)	-0.387	0.000	0.354	-0.373	-0.010	0.000	0.755	0.004

Note: FeMH stands for entropy-cost Flexible Moral Hazard and FmMH for mean-cost Flexible Moral Hazard. The category "Fund relative to IMD (Δ)" corresponds to the moment in the Fund economy minus the same moment in the respective IMD economy.

Table 3 reports the business cycle properties of the four Fund contracts together with their change relative to the corresponding IMD economy. The entropy of the chosen distribution, $H(\pi^g)$, is the most direct summary of what separates the two flexible specifications. Under mean-cost flexible MH, entropy is exactly zero, with zero volatility, in both the Fund and the IMD economy, since the borrower always selects a Dirac measure on g. Under entropy-cost flexible MH, it is the highest of all contracts and genuinely time-varying, so that the optimal distribution stays interior along the whole simulated path. Canonical and back-loaded MH inherit a small and nearly constant entropy from the fixed distributions of the canonical framework. Relative to IMD, the contract under entropy-cost flexible MH lowers mean entropy and turns it from counter to pro-cyclical: once partially insured by the Fund, the borrower is willing to pay the entropy cost of concentrating its distribution, and does so most in high endowment states.

Regarding the means, both flexible contracts sustain much higher g and output than canonical MH, but only the entropy-cost contract converts this into higher consumption, and it is also the only contract that raises mean consumption relative to its own IMD benchmark. Under mean-cost flexible MH, by contrast, the Fund contract is costly in levels since sustaining a high g is expensive when the cost depends on the mean alone.

The reported volatilities separate the two flexible contracts even more sharply. Entropy-cost MH delivers the lowest volatility of every variable in the table and lowers each of them relative to IMD, for output, consumption, labor and the primary surplus. In opposition, mean-cost MH delivers the highest volatilities and raises every one of them. Canonical and back-loaded MH are again nearly indistinguishable, leaving output volatility essentially unchanged and reducing consumption volatility only slightly, consistent with the small IC multiplier documented in Figure 1.

Finally, the correlation panel shows that canonical and back-loaded MH work almost entirely through co-movement rather than through volatility, making labor and the primary surplus markedly more pro-cyclical while barely moving the already high consumption-output correlation. Entropy-cost flexible MH combines lower volatility with a stronger co-movement of consumption with output, whereas mean-cost flexible MH is the only contract that lowers it. Thus, under mean cost, risk sharing shows up as a weaker link between consumption and output rather than as reduced volatility.

We end with a welfare analysis in steady state, computed as consumption-equivalent changes. Denoting the borrower's value in the benchmark and alternative cases by $V^b(\theta)$ and $\ddot{V}^b(\theta)$, the borrower's gain is $(\exp\left[(\ddot{V}^b(\theta) - V^b(\theta))(1-\beta)\right] - 1) \times 100$. For lenders, we compute $\ddot{V}^l(\theta)/V^l(\theta) \times 100$. Table 4 reports these gains: the upper panel compares the different IMD economies to the IMD economy under canonical MH (the benchmark), the lower panel does the same for the Fund contracts.

When comparing the different IMD economies, the Pareto ranking is clear. The mean-cost flexible MH dominates followed by the entropy-cost flexible MH. The reverse order holds for private lenders. The magnitude of the gains is immense, especially for the borrower with average gains far above 100%.

Looking at the Fund contracts, the Fund contract under entropy-cost flexible MH largely dominates all the other contracts on the borrower's side. Again, the magnitude of these gains is immense. On the Fund's side, the mean-cost flexible MH dominates. Regarding the

Table 4: Welfare Gains

	IMD vs. CMH IMD					
	Borrower			Private Lenders		
	FmMH	FeMH		FmMH	FeMH	
Average	137.40	122.70		3.55	4.81	
	Fund vs. CMH Fund					
	Borrower			Fund		
	FmMH	FeMH	BMH	FmMH	FeMH	BMH
Average	73.00	137.50	0.61	5.54	-3.67	-0.55

Note: CMH stands for Canonical Moral Hazard, BMH for Back-Loaded Moral Hazard, FeMH for entopy-cost Flexible Moral Hazard and FmMH for mean-cost Flexible Moral Hazard. The borrower's welfare gains for a specific θ correspond to $(\exp\left[(\ddot{V}^b(\theta) - V^b(\theta))(1-\beta)\right] - 1) \times 100$ where $\ddot{V}^b$ and V^b are the values of the borrower in the benchmark and the alternative case, respectively. For the lenders, welfare gains are given by $\ddot{V}^l(\theta)/V^l(\theta) \times 100$.

back-loaded MH, there are positive welfare gains on average for the borrower as one would expect. The Fund records welfare losses on average, though.

6 Conclusion

This paper studies how the design of incentives in dynamic contracts shapes long-run risk sharing when actions are not contractable (e.g. unobservable) and commitment is limited. We contribute to the theory and, hopefully, to the official lending practices. To start, we emphasize the importance of discriminating, not only between exogenous and endogenous risks, but also among different types of endogenous risk according to the capacity that borrowers have to reduce their risks. By comparing alternative moral hazard (MH) environments within a unified contracting framework, we show that the way incentives are provided is central to understanding whether risk sharing is preserved or undermined over time. The analysis highlights that MH does not, by itself, imply poor risk sharing; rather, it is the informational content of realized outcomes and the associated incentive mechanisms that drive long-run distortions.

Our theoretical results identify a sharp contrast between canonical and flexible MH. When – as in the canonical MH – incentives rely on outcome-contingent rewards and punishments, risk sharing is necessarily disrupted and consumption paths spread out over time. In contrast, when – as in the flexible MH – agents can locally distort probability distributions at a cost, incentives can be provided by compensating marginal effort costs without conditioning on realized outcomes.

In this almost fifty years since Holmström (1979), the canonical MH has been fully developed as theory, less as an application to official lending programs; in contrast, the flexible MH theory of Georgiadis et al. (2024) is still novel and we have contributed to it in three ways. First, by developing its dynamic contract formulation; second, by showing that it reverts the *immiseration* result of the dynamic canonical MH (Atkeson and Lucas (1992)) to *bliss* in the flexible MH, and third and importantly, by introducing a cost for risk reduction, in particular, entropy-based costs resulting in constrained-efficient interior distributions and effective risk sharing, in contrast to the original mean-based costs.

We also propose a back-loaded incentive scheme that bridges the two MH approaches. By postponing rewards and punishments until enforcement constraints bind, which may, in turn, constraint punishments or rewards. This contractual design improves risk sharing while preserving incentives and aligns closely with the structure of official lending programs. Our quantitative analysis confirms that these differences in incentive provision translate into substantial variation in macroeconomic dynamics and welfare in a calibrated model of sovereign lending for Euro Area stressed countries.

The counterpart to the borrowers capacity to reduce risks is the lenders capacity to design and implement incentive compatibility constraints. For instance, in the canonical MH the distribution is assumed to be known and the effort is non-contractable (e.g. non-observable), while in the canonical MH the feasible distributions are assumed to be known – in particular, their cost – but the specific choice is not contractable, in fact, incentives are *ex-ante* the choice is made. Nevertheless, for some risks, the lender may be able to target neither, but risk-profile improvements may be detected by contractable – possibly, state-contingent – conditions. This is a third type of risk contract (not MH): a state-contingent conditionality contract. Non state-contingent conditionality has been, and is, widely used by official lenders (e.g. conditions that need to be fulfilled for disbursements to take place), without a theory, which is now developed in Marimon et al. (2026).

References

Ábrahám, Á., E. Carceles-Poveda, Y. Liu, and R. Marimon (2025): "On the Optimal Design of a Financial Stability Fund," *The Review of Economic Studies*, Forthcoming.

Aguiar, M. and M. Amador (2021): *The Economics of Sovereign Debt and Default*, Princeton University Press.

Aguiar, M. and G. Gopinath (2006): "Defaultable Debt, Interest Rates and the Current Account," *Journal of International Economics*, 69, 64–83.

Arellano, C. (2008): "Default Risk and Income Fluctuations in Emerging Economies," *American Economic Review*, 98, 690–712.

Atkeson, A. (1991): "International Lending with Moral Hazard and Risk of Repudiation," *Econometrica*, 59, 1069–1089.

Atkeson, A. and R. E. Lucas (1992): "On Efficient Distribution with Private Information," *The Review of Economic Studies*, 59, 427–453.

Bocola, L., G. Bornstein, and A. Dovis (2019): "Quantitative Sovereign Default Models and the European Debt Crisis," *Journal of International Economics*, 118, 20–30.

Callegari, G., R. Marimon, A. Wicht, and L. Zavalloni (2023): "On a Lender of Last Resort with a Central Bank and a Stability Fund," *Review of Economic Dynamics*, 50, 106–130.

Chatterjee, S. and B. Eyigungor (2012): "Maturity, Indebtedness, and Default Risk," *American Economic Review*, 102, 2674–2699.

Dovis, A. (2019): "Efficient Sovereign Default," *Review of Economic Studies*, 86, 282–312.

Dovis, A. and R. Kirpalani (2023): "On the Design of a Robust Lender of Last Resort," *Unpublished Manuscript*.

Dvorkin, M., J. Sánchez, H. Sapriza, and E. Yurdagul (2021): "Sovereign Debt Restructurings," *American Economic Journal: Macroeconomics*, 13, 26–77.

Ferrari, A., Y. Liu, R. Marimon, and C. Simpson-Bell (2024): "Fiscal and Currency Union with Default and Exit," .

Georgiadis, G., D. Ravid, and B. Szentes (2024): "Flexible Moral Hazard Problems," *Econometrica*, 92, 387–409.

Green, E. J. (1987): "Lending and the Smoothing of Uninsurable Income," in *Contractual Arrangements for Intertemporal Trade*, ed. by E. C. Prescott and N. Wallace, Minneapolis: University of Minnesota Press.

Holmström, B. (1979): "Moral Hazard and Observability," *The Bell Journal of Economics*.

Iskhakov, F., T. H. Jørgensen, J. Rust, and B. Schjerning (2017): "The Endogenous Grid Method for Discrete-Continuous Dynamic Choice Models with (or without) Taste Shocks," *Quantitative Economics*, 8, 317–365.

Kehoe, P. and F. Perri (2002): "International Business Cycles with Endogenous Incomplete Markets," *Econometrica*, 70, 907–928.

Liu, Y., R. Marimon, and A. Wicht (2020): "Making Sovereign Debt Safe with a Financial Stability Fund," .

Marcet, A. and R. Marimon (1992): "Communication, commitment, and growth," *Journal of Economic Theory*, 58, 219–249.

——— (2019): "Recursive Contracts," *Econometrica*, 87, 1589–1631.

Marimon, R., A. Lenarcic, A. Erce, and D. Clancy (2026): "A Theory of Conditionality: Evidence from the European debt crisis," .

Mateos-Planas, X., S. McCrary, J.-V. Rios-Rull, and A. Wicht (2022): "The Generalized Euler Equation and the Bankruptcy-Sovereign Default Problem," .

——— (2023): "Commitment in the Canonical Sovereign Default Model," *University of Pennsylvania.*

Mele, A. (2011): "Repeated moral hazard and recursive Lagrangeans," Working Paper.

Mihalache, G. (2020): "Sovereign Default Resolution through Maturity Extension," *Journal of International Economics*, 125, 103–126.

Müller, A., K. Storesletten, and F. Zilibotti (2019): "Sovereign Debt and Structural Reforms," *American Economic Review*, 109, 4220–4259.

Phelan, C. (1995): "Repeated Moral Hazard and One-Sided Commitment," *Journal of Economic Theory*, 66, 488–506.

——— (1998): "On the Long Run Implications of Repeated Moral Hazard," *Journal of Economic Theory*, 79, 174–191.

Prescott, E. C. and R. M. Townsend (1984): "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," *Econometrica*, 52, 21–46.

Quadrini, V. (2004): "Investment and Liquidation in Renegotiation-Proof Contracts with Moral Hazard," *Journal of Monetary Economics*, 51, 713–751.

Roch, F. and H. Uhlig (2018): "The Dynamics of Sovereign Debt Crises and Bailouts," *Journal of International Economics*, 114, 1–13.

Rogerson, W. P. (1985): "The First-Order Approach to Principal-Agent Problems," *Econometrica*, 53, 1357–1367.

Rouwenhorst, K. G. (1995): "Asset Pricing Implications of Equilibrium Business Cycle Models," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, Princeton: Princeton University Press, 294–330.

Rust, J. (1988): "Statistical Models of Discrete Choice Processes," *Transportation Research Part B: Methodological*, 22, 125–158.

Simpson-Bell, C. (2020): "Risk Sharing and Policy Convergence in Economic Unions," *PhD thesis, European University Institute.*

Thomas, J. and T. Worrall (1990): "Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem," *Journal of Economic Theory*, 51, 367–390.

——— (1994): "Foreign Direct Investment and the Risk of Expropriation," *Review of Economic Studies*, 61, 81–108.

Tirole, J. (2015): "Country Solidarity in Sovereign Crises," *American Economic Review*, 105, 2333–2363.

Wicht, A. (2025): "Efficient Sovereign Debt Buybacks," *University of Basel.*

Online Appendix (Not For Publication)

A Proofs

A.1 Preliminary lemmas

We first prove a few preliminary lemmas before proving the propositions and lemmas located in the main text. We start with the characteristics of the borrower's and the Fund's value.

Lemma A.1. *Under Assumption 1, in the Fund contract under flexible MH:*

1. *When none of the LE constraints binds, $x'(y,x,y')$ and $c(y,x)$ are strictly increasing in x, $\pi(y,x)$ is strictly decreasing in entropy in x if $\alpha > 0$ and strictly increasing in first-order stochastic dominance in x if $\gamma > 0$, $V^b(y,x)$ is strictly increasing and strictly concave in x and $V^l(y,x)$ is strictly decreasing and concave in x.*

2. *When one of the LE constraints binds, $x'(y,x,y')$, $c(y,x)$, $\pi(y,x)$, $V^b(y,x)$ and $V^l(y,x)$ are all constant in x.*

3. *$V^b(s', x'(y,x,y'))$ is strictly increasing and concave in y' and $V^l(s', x'(y,x,y'))$ is strictly increasing and concave in y'.*

Proof. For the flexible MH, the values are given by (13)-(15). We first show that

$$\partial_x FV(y,x) = V^b(y,x) + x\partial_x V^b(y,x) + \partial_x V^l(y,x) = V^b(y,x),$$

which implies the efficient risk-sharing property: $x\partial_x V^b(y,x) = -\partial_x V^l(y,x)$. This comes from the envelope condition stating that $\partial_x FV = V^b(y,x)$. At the same time, the decomposition in FV leads to $\partial_x FV(y,x) = V^b(y,x) + x\partial_x V^b(y,x) + \partial_x V^l(y,x)$. Combining these two equations delivers the desired result.

We second show that $x'(y,x,g')$ and $c(y,x)$ are strictly increasing in x and $\pi(y,x)$ is monotone (in entropy or first-order stochastic dominance) in x when none of the LE constraints binds. For $x'(y')$, the statement directly follows from the law of motion of the relative Pareto weight in (12). For consumption, the statement follows from the first-order conditions (16). For $\pi(y,x)$, the statement follows from (9) and $V^b(y,x)$ being strictly increasing in x as we show next.

We third show the properties of $V^b(y,x)$ and $V^l(y,x)$ when none of the LE constraints binds. By definition, when x increases the Fund gives more weight to the borrower. As a result, $V^b(y,x)$ is strictly monotone in x. Then, using our first result, $\partial_x V^l(y,x) < 0$ so that $V^l(y,x)$ is strictly decreasing in x. We show concavity below.

We fourth show that all policies and value functions are constant when none of the LE constraints binds. This follows from the fact that the policies, value functions and multipliers are evaluated when the constraints are binding as solutions to the saddle-point problem.

We fifth show that $x'(y')$, $V^b(y', x'(y'))$ and $V^l(y', x'(y'))$ increasing in y'. For $x'(y')$, the statement directly follows from the fact that $\varrho(y') \geq \varrho(\tilde{y}') \geq 0$ for $y' \geq \tilde{y}'$ as shown in Lemma 1. Regarding the value function, we have seen that $FV(y,x)$ is increasing in x. Given x, a higher y means a higher surplus and therefore a higher $FV(y,x)$ and, through risk-sharing, a higher $V^b(y,x)$ and $V^l(y,x)$.

Finally we can show concavity of $V^l(y,x)$ and strict concavity of $V^b(y,x)$ in (y,x), when the LE constraints do not bind. When these constraints bind, the two values are simply constant. From the law of motion of the relative Pareto weight, when no LE constraint binds, $x'(y')$ linearly depends on x. Additionally, from (16), $c(y,x)$ linearly depends on x and is constant in y. Given this and the fact that the Fund is risk neutral, $V^l(y,x)$ is concave in (y,x). Similarly, given the risk aversion of the borrower, $V^b(y,x)$ is strictly concave in x and concave in y. □

Ábrahám et al. (2025, Lemmas 1-3) provide a proof of the same lemma in the case of canonical MH. To extend the proof to the back-loaded MH, we need to distinguish between the main and the latent relative Pareto weight. Other than that, the same type of argument ought to apply.

A.2 Proof of Proposition 1

We first consider the case in which $\alpha = 0$ and $\gamma > 0$. When the borrower is in autarky, the first-order condition is both necessary and sufficient. Sufficiency is ensured by the convexity of the cost function under Assumption 1. Necessity comes from the differentiability of the cost function and the fact that we assume standard probability simplex with linear equality constraint.

Given this, under Assumption 2 with $\alpha = 0$ and $\gamma > 0$, the marginal cost of a specific distribution π is given by

$$v_\pi(y) = \gamma l(y),$$

which is linear in y. From (3), this gives

$$u(y) = \lambda + \gamma l(y).$$

The right-hand side is linear in y, while the left-hand side is strictly concave in y. Hence, there can be tangency at at most one point. Therefore, the support of π^* contains at most one point. Alternatively, define $\Lambda_\pi(y) = u(y) - \lambda - v_\pi(y)$. As a result, it holds that

$$\text{supp } \pi \subseteq \underset{y \in Y}{\text{argmax}} \ \Lambda_\pi(y).$$

Since $u(y)$ is strictly concave and $v_\pi(y)$ is linear in y, $\Lambda_\pi(y)$ is strictly concave in y. Hence, there is only one y that maximizes $\Lambda_\pi(y)$. The problem of choosing an optimal distribution is therefore the same as the one of choosing y directly. The borrower's expected utility is therefore $u(y)$ for $y \in \text{supp}(\delta_y)$.

The principal-agent problem is a weighted sum of the borrower's problem which we already show is convex and the Fund's problem which is linear. Sufficiency and necessity comes from the previous argument and the fact that we assume interiority of the solution. Given this, the maximization problem can be rewritten as

$$\max_{\pi} \mu_b \left[\sum_{y \in Y} \pi(y) \left(v_\pi(y) + \chi \right) - v(\pi) \right] + \mu_l \left[\sum_{y \in Y} \pi(y) (y - u^{-1} \left(v_\pi(y) + \chi \right)) \right].$$

where we simply plugged (3) directly inside the objective function. The first-order condition is then given by

$$\mu_b \left[\chi + \pi(y) \omega_\pi(y) \right] + \mu_l \left[y - c(y) - \frac{\pi(y)\omega_\pi(y)}{u_c(c(y))} \right] = \chi,$$

where we use the fact that $(u^{-1}(x))' = \frac{1}{u_c(u^{-1}(x))}$. We can rewrite this expression as

$$\Lambda_\pi(y) \equiv \left(\mu_b - \frac{\mu_l}{u_c(c(y))} \right) \pi(y)\omega_\pi(y) + \mu_l \left[y - c(y) \right] - \chi(1 - \mu_b).$$

Under Assumption 2 with $\alpha = 0$ and $\gamma > 0$, $\omega_\pi(y) = 0$ for all y. We have as above

$$\text{supp } \pi \subseteq \underset{y \in Y}{\text{argmax}} \ \Lambda_\pi(y).$$

From the property of the utility function $c(y) = u^{-1}(v_\pi(y) + \chi)$ is strictly increasing and strictly convex in y. As a result, $\Lambda_\pi(y)$ is strictly concave in y and has a single maximizer.

Now observe that whenever $y - c(y) > 0$ with optimal distribution δ_y, the agent's expected utility is lower than in autarky since $u(c(y)) < u(y)$ for $y \in \text{supp}(\delta_y)$.

We then consider the case in which $\alpha > 0$ and $\gamma \geq 0$. Using the same argument as before, the first-order condition is both necessary and sufficient. With $\alpha > 0$, the marginal cost is $v_\pi(y) = \alpha(1 + \log \pi(y)) + \gamma l(y)$ which gives

$$\pi(y) = \exp\left(\frac{u(y) - \gamma l(y) - \lambda}{\alpha} - 1\right) = \exp\left(\frac{u(y) - \gamma l(y)}{\alpha}\right) \exp\left(\frac{-\lambda}{\alpha} - 1\right).$$

Define $A \equiv \exp\left(\frac{-\lambda}{\alpha} - 1\right)$, it must then hold that

$$\sum_{y \in Y} A \exp\left((u(y) - \gamma l(y))/\alpha\right) = 1,$$

which gives $A = \frac{1}{\sum_{y' \in Y} \exp((u(y') - \gamma l(y'))/\alpha)}$. Thus, the optimal distribution takes the familiar logit form given which converges to a Dirac measure δ_y as $\alpha \downarrow 0$.

Let's pass to the principal-agent problem. Under Assumption 2 with $\alpha > 0$ and $\gamma \geq 0$, for a given $\{c(y)\}_{y \in Y}$, the optimal distribution is the outcome of

$$\begin{aligned} &\max_{\pi \in \mathcal{M}} \left\{ \sum_{y \in Y} \pi(y)\, u(c(y)) - v(\pi) \right\} \\ &\text{s.t.} \quad \sum_{y \in Y} \pi(y) = 1. \end{aligned}$$

Following the same steps as above, the optimal distribution is given by (18). This means that the maximization problem can be rewritten as

$$\begin{aligned} &\max_{c(y)} \ \mu_b \left[\sum_{y \in Y} \pi(y) u(c(y)) - v(\pi)\right] + \mu_l \left[\sum_{y \in Y} \pi(y)(y - c(y))\right] \\ &\text{s.t.} \quad \pi(y) = \frac{e^{(u(c(y)) - \gamma l(y))/\alpha}}{\sum_{y' \in Y} e^{(u(c(y')) - \gamma l(y'))/\alpha}}. \end{aligned}$$

Plugging the constraint directly into the objective function, the agent's value becomes

$$\begin{aligned} \sum_{y \in Y} \pi(y) u(c(y)) - v(\pi) =& \sum_{y \in Y} \frac{e^{(u(c(y)) - \gamma l(y))/\alpha}}{\sum_{y' \in Y} e^{(u(c(y')) - \gamma l(y'))/\alpha}} u(c(y)) \\ &- \sum_{y \in Y} \frac{e^{(u(c(y)) - \gamma l(y))/\alpha}}{\sum_{y' \in Y} e^{(u(c(y')) - \gamma l(y'))/\alpha}} \left[\alpha \log\left(\frac{e^{(u(c(y)) - \gamma l(y))/\alpha}}{\sum_{y' \in Y} e^{(u(c(y')) - \gamma l(y'))/\alpha}}\right) + \gamma l(y)\right] \\ =& \alpha \log\left(\sum_{y' \in Y} e^{(u(c(y')) - \gamma l(y'))/\alpha}\right). \end{aligned}$$

Given this, the first-order condition of the principal-agent problem with respect to $c(y)$ is

$$\mu_b \pi(y) u_c(c(y)) + \mu_l \left[\pi(y)(1 - \pi(y)) \frac{u_c(c(y))}{\alpha}[y - c(y)] - \pi(y)\right] = 0.$$

Dividing by $\pi(y)$ and rearranging

$$u_c(c(y))\left[\mu_b + \mu_l(1-\pi(y))\frac{y-c(y)}{\alpha}\right] = \mu_l. \tag{A.1}$$

Given this, we obtain for any two states $y_1 < y_2$

$$\frac{u_c(c(y_1))}{u_c(c(y_2))} = \frac{\mu_b + \mu_l(1-\pi(y_2))\frac{y_2-c(y_2)}{\alpha}}{\mu_b + \mu_l(1-\pi(y_1))\frac{y_1-c(y_1)}{\alpha}}$$

Assume by contradiction that $c(y_1) > c(y_2)$. This implies that $\pi(y_1) > \pi(y_2)$ by (18) and thus the numerator on the right-hand side is larger than the denominator. This means that $u_c(c(y_1)) > u_c(c(y_2))$ and by the strict concavity of u that $c(y_1) < c(y_2)$, a contradiction. So consumption increases with endowment, but less than one-for-one. We have in other words partial insurance.

Full consumption smoothing is not optimal for $\alpha > 0$. One directly sees that setting $c(y) = \bar{c}$ for all y violates the first-order condition (A.1) as y is stochastic. However, as $\alpha \to 0$, the optimal distribution becomes a Dirac as shown above and $c(y) = y$.

Let's finally show that the principal can improve upon the autarkic allocation. Observe that the autarkic allocation $c(y) = y$ for all y is not optimal as (A.1) would imply $u_c(y) = \frac{\mu_l}{\mu_b}$. This cannot be true as (μ_b, μ_l) is fixed while y is stochastic. To see how the principal improves upon autarky, define by $\overline{y}$ and by $\underline{y}$ the highest and lowest endowment, respectively. The agent's value is $\alpha \log W$ with $W \equiv \sum_{y' \in Y} e^{(u(c(y'))-\gamma l(y'))/\alpha}$. Denote by $W_A = \sum_{y' \in Y} e^{(u(y')-\gamma l(y'))/\alpha}$ the value of W in autarky, and by $\pi_A(y)$ the optimal distribution in autarky given above. The principal can set $c(\overline{y}) = \overline{y} - \epsilon$ for some $\epsilon > 0$ and $c(\underline{y}) = \underline{y} + \frac{\pi_A(\overline{y})}{\pi_A(\underline{y})}\epsilon$. We then have

$$W^\epsilon = e^{(u(\overline{y}-\epsilon)-\gamma l(\overline{y}))/\alpha} + e^{(u(\underline{y}+\frac{\pi_A(\overline{y})}{\pi_A(\underline{y})}\epsilon)-\gamma l(\underline{y}))/\alpha} + \sum_{y' \neq \overline{y}, \underline{y}} e^{(u(c(y'))-\gamma l(y'))/\alpha}.$$

Taking the derivative with respect to ϵ evaluated at $\epsilon = 0$, we obtain

$$\begin{aligned}\left.\frac{\partial W^\epsilon}{\partial \epsilon}\right|_{\epsilon=0} &= -\, e^{(u(\overline{y})-\gamma l(\overline{y}))/\alpha}\frac{u_c(\overline{y})}{\alpha} + e^{(u(\underline{y}-\gamma l(\underline{y}))/\alpha}\frac{u_c(\underline{y})}{\alpha}\frac{\pi_A(\overline{y})}{\pi_A(\underline{y})} \\ &= \frac{W_A \pi_A(\overline{y})}{\alpha}\left[u_c(\underline{y}) - u_c(\overline{y})\right] > 0,\end{aligned}$$

where we use the fact that $\pi_A(y) = e^{(u(y)-\gamma l(y))/\alpha}/W_A$. The inequality comes from the fact that $\overline{y} > \underline{y}$ and u is strictly concave. Hence, smoothing consumption (i.e. reducing the variance of c) increases W when we reduce the spread of c relative to y while maintaining the budget constraint. Since $\alpha > 0$ and the logarithm is a strictly increasing function, the agent is strictly better off than in autarky.

Given that the borrower can be made strictly better off than autarky under a budget-neutral consumption reallocation, the Fund cannot be worse off than in autarky.

A.3 Proof of Proposition 2

We first consider the case in which $\alpha = 0$ and $\gamma > 0$. In autarky, the agent solves

$$V_A^b(y) = \max_{\pi \in \mathcal{M}} \left\{ u(y) - v(\pi) + \beta \sum_{y' \in Y} \pi(y') V_A^b(y') \right\} \tag{A.2}$$

$$\text{s.t.} \quad \sum_{y \in Y} \pi(y) = 1. \tag{A.3}$$

Attaching the multiplier λ to the constraint, the first-order condition for all y in the support of π is given by

$$\beta V_A^b(y) - v_\pi(y) = \lambda. \tag{A.4}$$

Under Assumption 2 with $\alpha = 0$ and $\gamma > 0$, the marginal cost of a specific distribution π is given by

$$v_\pi(y) = \gamma l(y),$$

which is linear in y. From (A.4), this gives

$$\beta V_A^b(y) = \lambda + \gamma l(y).$$

The right-hand side is linear in y, while the left-hand side is strictly concave in y since only $u(y)$ depends on y. Hence, there can be tangency at at most one point. Therefore, the support of π^* contains at most one point. Alternatively, define $\Lambda_\pi(y) = \beta V_A^b(y) - \lambda - v_\pi(y)$. Therefore

$$\text{supp } \pi \subseteq \operatorname*{argmax}_{y \in Y} \ \Lambda_\pi(y).$$

Since $\beta V_A^b(y)$ is strictly concave and $v_\pi(y)$ is linear in y, $\Lambda_\pi(y)$ is strictly concave in y. Hence, there is only one y that maximizes $\Lambda_\pi(y)$. The problem of choosing an optimal distribution is therefore the same as the one of choosing y directly. The borrower's expected utility is simply $u(y)/(1-\beta)$ for $y \in \text{supp}(\delta_y)$.

Let's pass to the principal-agent problem. Under Assumption 2 with $\alpha = 0$ and $\gamma > 0$, $\omega_\pi(y) = 0$ for all y. This means that (17) becomes

$$\frac{1}{1+r} V^l(y', x'(y')) = \chi(1 - x(1 + \nu_b)).$$

Since $V^D(y')$ is strictly increasing in y' so does $V^b(y', x(y'))$ and thus $V^l(y', x(y'))$ is strictly decreasing in y'. There is therefore only one y' that satisfies (17). In other words, the optimal distribution has at most one y' in its support.

Now observe that whenever $y - c(y) > 0$ with optimal distribution δ_y, the agent's expected lifetime utility is simply $u(c(y))/(1-\beta)$ which is lower than in autarky since $u(c(y)) < u(y)$ for $y \in \text{supp}(\delta_y)$.

We then consider the case in which $\alpha > 0$ and $\gamma \geq 0$. For a given $\{c(y, x(y))\}_{y \in Y}$, the optimal distribution is the outcome of

$$V^b(y, x) = \max_{\pi \in \mathcal{M}} \left\{ u(c(y, x)) - v(\pi) + \beta \sum_{y' \in Y} \pi(y') V^b(y', x(x')) \right\}$$

$$\text{s.t.} \quad \sum_{y \in Y} \pi(y) = 1.$$

Following the same steps as in the proof of Proposition 1, the optimal distribution is given by (18). Plugging the constraint directly into the objective function, the agent's value becomes

$$\begin{aligned} V^b(y, x) =& u(c(y, x)) - \sum_{y' \in Y} \frac{e^{(\beta V^b(y', x') - \gamma l(y'))/\alpha}}{\sum_{y'' \in Y} e^{(V^b(y'', x') - \gamma l(y''))/\alpha}} \left[\alpha \log \left(\frac{e^{(V^b(y', x') - \gamma l(y'))/\alpha}}{\sum_{y'' \in Y} e^{(V^b(y'', x') - \gamma l(y''))/\alpha}} \right) + \gamma l(y') \right] \\ &+ \beta \sum_{y' \in Y} \frac{e^{(\beta V^b(y', x') - \gamma l(y'))/\alpha}}{\sum_{y'' \in Y} e^{(V^b(y'', x') - \gamma l(y''))/\alpha}} V^b(y', x') \\ =& u(c(y, x)) + \alpha \log \left(\sum_{y' \in Y} e^{(\beta V^b(y', x') - \gamma l(y'))/\alpha} \right) \end{aligned}$$

Given (16) and (12), the first-order condition with respect to c becomes

$$u_c(c) = \frac{\eta}{x'(y)},$$

Fix x and consider any two states $y_1 < y_2$. Assume by contradiction that $c(y_1) > c(y_2)$. This implies that $x'(y_1) > x'(y_2)$ and therefore $V^D(y_1) = V^b(y_1, x) > V^b(y_2, x) \geq V^D(y_2)$. This is a contradiction as $V^D(y_1) < V^D(y_2)$ by assumption. So consumption increases with endowment.

Full consumption smoothing is not optimal for $\alpha > 0$. One directly sees that setting $c(y) = \bar{c}$ for all y violates the first-order condition (16) as $V^D(y)$ is strictly increasing in y. However, as $\alpha \to 0$, the optimal distribution becomes a Dirac as shown before and $c(y) = y$.

Let's finally show that the principal can improve upon the autarkic allocation. Define by $\overline{y}$ and by $\underline{y}$ the highest and lowest endowment, respectively. The agent's value is $u(c) + W$

where $W \equiv \alpha \log \sum_{y' \in Y} e^{(u(\beta V^b(y',x')-\gamma l(y'))/\alpha}$. Denote by W_A the value of W in autarky, and by $\pi_A(y)$ the optimal distribution in autarky where $c(y) = y$. Consider a one-shot deviation in which in the next period the principal sets $c(\overline{y}) = \overline{y} - \epsilon$ for some $\epsilon > 0$ and $c(\underline{y}) = \underline{y} + \frac{\pi_A(\overline{y})}{\pi_A(\underline{y})}\epsilon$. In all the following periods, the principal then set $c(y) = y$ for all $y \in Y$. We then have

$$W^\epsilon = e^{(\beta u(\overline{y}-\epsilon)+\beta^2 \mathbb{E} V^b(y')-\gamma l(\overline{y}))/\alpha} + e^{(\beta u(\underline{y}+\frac{\pi_A(\overline{y})}{\pi_A(\underline{y})}\epsilon)+\beta^2 \mathbb{E} V^b(y')-\gamma l(\underline{y}))/\alpha} + \sum_{y' \neq \overline{y}, \underline{y}} e^{(\beta V^b(y')-\gamma l(y'))/\alpha}.$$

Taking the derivative with respect to ϵ evaluated at $\epsilon = 0$, we obtain

$$\begin{aligned} \left.\frac{\partial W^\epsilon}{\partial \epsilon}\right|_{\epsilon=0} &= - e^{(V^b(\overline{y})-\gamma l(\overline{y}))/\alpha} \frac{\beta u_c(\overline{y})}{\alpha} + e^{(V^b(\underline{y})-\gamma l(\underline{y}))/\alpha} \frac{\beta u_c(\underline{y})}{\alpha} \frac{\pi_A(\overline{y})}{\pi_A(\underline{y})} \\ &= \frac{W_A \pi_A(\overline{y})}{\alpha} \beta \left[u_c(\underline{y}) - u_c(\overline{y}) \right] > 0, \end{aligned}$$

where we use the fact that $\pi_A(y) = e^{(V^b(y)-\gamma l(y))/\alpha}/W_A$. The inequality comes from the fact that $\overline{y} > \underline{y}$ and u is strictly concave. Hence, smoothing consumption (i.e. reducing the variance of c) increases W when we reduce the spread of c relative to y while maintaining the budget constraint. Since $\alpha > 0$ and the logarithm is a strictly increasing function, the agent is strictly better off than in autarky.

Moreover, given that the borrower can be made strictly better off than autarky under a budget-neutral consumption reallocation, the Fund cannot be worse off than in autarky.

A.4 Proof of Proposition 3

To show existence we use the argument in the proof of Theorem 3 of Marcet and Marimon (2019) and the proof of Proposition 1 of Ábrahám et al. (2025).

Marcet and Marimon (2019) make the following necessary assumptions: A1 a well defined Markov chain process for y, A2 continuity in $\{c, \pi\}$ and measurability in y, A3 non-empty feasible sets, A4 uniform boundedness, A5 convex technologies, A6 concavity for the lenders and strict concavity for the borrower, and a strict interiority condition. Assumption A1, A2, A5 and A6 are trivially met given Assumption 1. Since feasible c and π are bounded, payoffs functions are bounded as well. $V^D(y)$ is (strictly) monotone in y which ensures that A4 is met. Whether A3 is satisfied depends on the initial condition $(y_0, x_0(y_0))$. Assumption 4 ensures feasibility and that the strict interiority condition is satisfied.

Similar to Ábrahám et al. (2025), we consider a relaxed contracting problem which is the same as the original contracting problem except that (9) is replaced by a weak inequality version. More precisely,

$$\beta V^b(y^{t+1}) - v_{\pi_{t+1}}(y^{t+1}) \geq 0. \tag{A.5}$$

Taking the derivative of this expression leads to

$$-w_{\pi_{t+1}}(y^{t+1}) \leq 0,$$

where the inequality follows from Assumption 2 stating that the second derivative is non negative. As a result, (A.5) defines a convex set of feasible distribution choices.

It should be noted that Theorem 3 in Marcet and Marimon (2019) is the recursive, saddle-point, representation corresponding to the original contract problem. To obtain the recursive formulation of the contract, we have normalized the co-state variable. We relied on the the homogeneity of degree one in (μ_b, μ_l) to redefine the contracting problem using x – i.e. effectively $(x, 1)$ – as a co-state variable. Given this and the fact that multipliers are uniformly bounded, the theorem applies. That is, if we define the set of of feasible Lagrange multipliers by $L = \{(\mu_b, \mu_l) \in \mathbb{R}^2_+\}$ and the set of feasible allocations by $A = \{c \in \mathbb{R}^+, \pi \in \mathcal{M}\}$, the correspondence $SP : A \times L \to A \times L$ mapping non-empty, convex, and compact sets to themselves, is non-empty, convex-valued, and upper hemicontinuous. I can therefore apply Kakutani's fixed point theorem and existence immediately follows.

Given this, we need to show that the relaxed contracting problem has the same solution as the original contracting problem. For this it suffices to show that $\varrho(y, x, y') > 0$. Assume by contradiction that $\varrho(y, x, y') = 0$. Define $\widetilde{FV}(y', x'(y')) \equiv x'(y')V^b(y', x'(y')) + \left[V^l(y', x'(y')) - V^l(\underline{y}', x'(\underline{y}))\right]$. Since the added terms do not depend on π, the first-order condition can be rewritten as

$$0 = x\,(1+\nu_b)\left[\beta V^b(y', x'(y')) - v_\pi(y')\right] + \frac{1+\nu_l}{1+r}\left[V^l(y', x'(y')) - V^l(\underline{y}, x'(\underline{y}))\right],$$

where $V^l(\underline{y}, x'(\underline{y}))$ was added as it is independent of π. Given the monotonicity of V^l, the second line is (weakly) positive meaning that the first line needs to be (weakly) negative. This contradicts (A.5).

Finally, as FV is monotone in x, constant when either of the LE constraints are binding and concave when both are slack, we can directly use the argument of Marcet and Marimon (2019) who show that the saddle point functional equation of the SPFE is a contraction mapping. The strict concavity/convexity assumptions on u and v imply that the allocation is unique.

A.5 Proof of Lemma 1

Define $\widetilde{FV}(y', x'(y')) \equiv x'(y')V^b(y', x'(y')) + \left[V^l(y', x'(y')) - V^l(\underline{y}', x'(\underline{y}))\right]$. Since the added terms do not depend on π, the chosen distribution maximizes the Fund's objective function if it maximizes

$$\sum_{y'\in Y}\pi(y')\left[\frac{1+\nu_l}{1+r}\widetilde{FV}(y',x'(y'))-x\varrho(y')(v_\pi(y')+\chi)\right]$$

The first-order condition with respect to π is then

$$\begin{aligned}&\frac{1}{1+r}\left[V^l(y',x'(y'))-V^l(\underline{y}',x'(\underline{y}))\right]+x(1+\nu_b+\varrho(y'))\left[\beta V^b(y',x'(y'))-(v_\pi(y')+\chi)\right]\\&=\pi(y')x\varrho(y')\omega_\pi(y')-x(1+\nu_b)\chi.\end{aligned}$$

Given (9), the second term equates zero. This means that for $y'=\underline{y}$, $\varrho(y')\geq 0$. Moreover, since $V^l(y',x'(y'))>V^l(\underline{y}',x'(\underline{y}))$ for $y'>\underline{y}$ from Lemma A.1, it holds that $\varrho(y')>\varrho(\underline{y})\geq 0$.

When $\eta=1$ and $\nu_l=0$ in all states, the law of motion of the relative Pareto weight simplifies to the following expression

$$\mathbb{E}x'(y')\equiv\mathbb{E}\left[\bar{x}'(y)+\hat{x}'(y')\right]=\mathbb{E}\left[(1+\nu_b x+\varrho(y')x\right].$$

Since $(\nu_b,\varrho(y'))\geq 0$, we get that $\mathbb{E}x'(y')\geq x$.

A.6 Proof of Lemma 2

Regarding the first part of the lemma, note that, given (20),

$$\sum_{y'\in Y}Q(y'|y,e(y))\hat{x}'(y')=0,$$

since independently of effort $\sum_{y'\in Y}Q(y'|y,e(y))=1$, hence $\sum_{y'\in Y}\partial_e Q(y'|y,e(y))=0$. Therefore $\mathbb{E}\hat{x}'=0$ and $\mathbb{E}x'=\bar{x}'(y)$. Alternatively, the expected law of motion of x can also be expressed as

$$\mathbb{E}x'=\mathbb{E}\left[\frac{1}{u'(c')}\frac{1+\nu_l(y')}{1+\nu_b(y')}\right]=\frac{1}{u'(c)}\eta,$$

where the last equality is the *inverse Euler equation* of the recursive contract (Ábrahám et al. (2025), Lemma 4).

Regarding the second part of the lemma, note that, since the LE multipliers are either zero or at most one of the two is positive, we can have the following decomposition

$$\mathbb{E}\frac{1}{u'(c')}=\mathbb{E}\left[x'\frac{1+\nu_b(y')}{1+\nu_l(y')}\right]=\mathbb{E}x'+\mathbb{E}x'\nu_b(y)-\mathbb{E}x'\frac{\nu_l(y)}{1+\nu_l(y)},$$

where $\mathbb{E}x'=\eta x$ and, without incentive constraints, the last two terms simply denote the change in the relative Pareto weight when either the no-default or the sustainability constraints binds. However, if LE constraints satisfy the 'no-free-lunch condition', (7) is more likely to bind, while the (8) is less likely to bind and, as a result, in both cases expected consumption increases.

B Quantitative Fund Contracts

In this section, we derive the quantitative version of the Fund contracts with endogenous labor and IMD autarky under canonical, back-loaded and restricted flexible MH.

B.1 Canonical moral hazard

We assume independence between ζ' and g'. This implies that a single shock variable, g', depends on effort. Given this, the Fund contract under canonical MH in recursive form is given by

$$\begin{aligned} FV(s,x) = \mathrm{SP} \min_{\{\nu_b,\nu_l,\varrho\}} \max_{\{c,n,e\}} \Big\{ & x\Big[(1+\nu_b)U(c,n,Q(e)) - \nu_b V^D(s) - \varrho v_{Q(e)}(e)\Big] \\ & + \big[(1+\nu_l)(\theta f(n) - c) - \nu_l Z(s)\big] + \frac{1+\nu_l}{1+r}\sum_{g'\in G} Q(g'|e) \sum_{\zeta'\in L} \pi^\zeta(\zeta'|\zeta) FV(s',x') \Big\} \\ \text{s.t. } x'(g') \equiv \bar{x}'(s) + \hat{x}'(g') = & \left[\frac{1+\nu_b}{1+\nu_l} + \frac{\varphi(g')}{1+\nu_l}\right]\eta x, \\ \varphi(g') = \varrho & \frac{\partial_e Q(g'|e)}{Q(g'|e)}. \end{aligned}$$

The effort policy $e(s,x)$ is determined by the first order condition of the SPFE with respect to e,

$$\begin{aligned} v_{Q(e)}(e) = & \beta \sum_{g'\in G} \partial_e Q(g'|e) \sum_{\zeta'\in L} \pi^\zeta(\zeta'|\zeta) V^b(s',x'(g')) \\ & + \frac{1+\nu_l(s,x)}{1+\nu_b(s,x)} \frac{1}{x} \frac{1}{1+r} \sum_{g'\in G} \partial_e Q(g'|e) \sum_{\zeta'\in L} \pi^\zeta(\zeta'|\zeta) V^l(s',x'(g')) \\ & - \frac{\varrho(g,x)}{1+\nu_b(s,x)} \left[\omega_{Q(e)}(e) - \beta \sum_{g'\in G} \partial_e^2 Q(g'|e) \sum_{\zeta'\in L} \pi^\zeta(\zeta'|\zeta) \hat{V}^b(s',x'(g'))\right]. \end{aligned}$$

Since the IC constraint is given by

$$v_{Q(e)}(e) = \beta \sum_{g'\in G} \partial_e Q(g'|e) \sum_{\zeta'\in L} \pi^\zeta(\zeta'|\zeta) V^b(s',x'(g')),$$

the first-order condition simplifies to

$$\begin{aligned} & \frac{1}{1+r} \sum_{g'\in G} \partial_e Q(g'|e) \sum_{\zeta'\in L} \pi^\zeta(\zeta'|\zeta) V^l(s',x'(g')) \\ & \quad = \vartheta(s,x) \Bigg[\omega_{Q(e)}(e) - \beta \sum_{g'\in G} \partial_e^2 Q(g'|e) \sum_{\zeta'\in L} \pi^\zeta(\zeta'|\zeta) V^b(s',x'(g')), \end{aligned}$$

where $\vartheta(s,x) \equiv \frac{x\varrho(g,x)}{1+\nu_l(s,x)}$.

B.2 Back-loaded moral hazard

As a subcontract resets when one of the LE constraint binds, the start of a subcontract is such that

$$
\begin{aligned}
\hat{FV}(s,x) = \min_{\{\nu_b,\nu_l,\varrho\}} \max_{\{c,n,e\}} \Big\{ & x\big[(1+\nu_b)(u(c)+h(1-n)-v(Q(e))) - \nu_b V^D(s) - \varrho v'(e)\big] \\
& + [(1+\nu_l)(\theta f(n)-c) - \nu_l Z(s)] \\
& + \frac{1+\nu_l}{1+r}\mathbb{E}\left[\mathbb{I}_{\{(x'(g'),s')\}}\hat{FV}(x'(g'),s') + (1-\mathbb{I}_{\{(x'(g'),s')\}})\overline{FV}(x'(g'),s',\bar{x}') \mid s,e\right] \Big\} \\
\text{s.t.} \quad & x'(g') = \eta x \frac{1+\nu_b+\varphi(g')}{1+\nu_l}, \\
& \bar{x}' = \eta x \frac{1+\nu_b}{1+\nu_l},
\end{aligned}
$$

Then within the subcontract

$$
\begin{aligned}
\overline{FV}(s,x,\bar{x}) = \min_{\{\varrho\}} \max_{\{c,n,e\}} \Big\{ & \bar{x}\,[u(c)+h(1-n)] - x\big[v(e)+\varrho v_{Q(e)}(e)\big] + (\theta f(n)-c) \\
& + \frac{1}{1+r}\mathbb{E}\left[\mathbb{I}_{\{(x'(g'),s')\}}\hat{FV}(x'(g'),s') + (1-\mathbb{I}_{\{(x'(g'),s')\}})\overline{FV}(x'(g'),s',\bar{x}') \mid s,e\right] \Big\} \\
\text{s.t.} \quad & x'(g') = \eta x\,[1+\varphi(g')]\,, \\
& \bar{x}' = \eta \bar{x}.
\end{aligned}
$$

We decompose the value as

$$
\begin{aligned}
\overline{V}^b_1(\overline{x},s) &= u(c(\overline{x},s)) + h(1-n(\overline{x},s)) + \beta\mathbb{E}\overline{V}^b_1(x'(g'),s'), \\
\overline{V}^b_2(s,x) &= -v(Q(e(x,s))) + \beta\mathbb{E}\overline{V}^b_2(x'(g'),s'),
\end{aligned}
$$

then the value of the borrower is simply $\overline{V}^b(s,x,\overline{x}) = \overline{V}^b_1(\overline{x},s) + \overline{V}^b_2(s,x)$ implying that

$$
\overline{FV}^b(s,x,\overline{x}) = \overline{x}\overline{V}^b_1(\overline{x},s) + x\overline{V}^b_2(s,x) + \overline{V}^l(s,x,\overline{x}).
$$

Given this, the first-order condition for the effort policy is

$$
\begin{aligned}
v_{Q(e)}(e) = & \beta\sum_{g'\in G}\partial_e Q(g'|e)\sum_{\zeta'\in L}\pi^{\zeta}(\zeta'|\zeta)W^b(x'(g'),s',\bar{x}') \\
& + \frac{1+\nu_l}{1+\nu_b}\frac{1}{x}\frac{1}{1+r}\sum_{g'\in G}\partial_e Q(g'|e)\sum_{\zeta'\in L}\pi^{\zeta}(\zeta'|\zeta)W^l(x'(g'),s',\bar{x}') \\
& - \frac{\varrho(g,x)}{1+\nu_b}\Bigg[-\beta\sum_{g'\in G}\partial^2_e Q(g'|e)\sum_{\zeta'\in L}\pi^{\zeta}(\zeta'|\zeta)\left[\mathbb{I}_{\{(x'(g'),s')\}}\overline{V}^b_2(x'(g'),s') + (1-\mathbb{I}_{\{(x'(g'),s')\}})\hat{V}^b(x'(g'),s')\right] \\
& + \omega_{Q(e)}(e)\Bigg].
\end{aligned}
$$

where $W^b(x'(g'), s', \bar{x}') = \mathbb{I}_{\{(x'(g'),s')\}}\hat{V}^b(x'(g'), s') + (1-\mathbb{I}_{\{(x'(g'),s')\}})\overline{V}^b(x'(g'), s', \bar{x}')$ and $W^l(x'(g'), s', \bar{x}') = \mathbb{I}_{\{(x'(g'),s')\}}\hat{V}^l(x'(g'), s') + (1-\mathbb{I}_{\{(x'(g'),s'\bar{x}')\}})\overline{V}^l(x'(g'), s', \bar{x}')$. As before this expression can be decomposed into the IC constraint

$$v_{Q(e)}(e) = \beta \sum_{g' \in G} \partial_e Q(g'|e) \sum_{\zeta' \in L} \pi^{\zeta}(\zeta'|\zeta) W^b(x'(g'), s', \bar{x}'),$$

which determines $e(s, x)$, and

$$\begin{aligned}
&\frac{1}{1+r} \sum_{g' \in G} \partial_e Q(g'|e) \sum_{\zeta' \in L} \pi^{\zeta}(\zeta'|\zeta) \left[W^l(x'(g'), s', \bar{x}')\right] \pi^{\zeta}(\mathrm{d}\zeta') \\
&= \vartheta(s,x)\Bigg[-\beta \sum_{g' \in G} \partial_e^2 Q(g'|e) \sum_{\zeta' \in L} \pi^{\zeta}(\zeta'|\zeta) \left[\mathbb{I}_{\{(x'(g'),s')\}}\overline{V}^b_2(x'(g'), s') + (1-\mathbb{I}_{\{(x'(g'),s')\}})\hat{V}^b(x'(g'), s')\right] \\
&\qquad + \omega_{Q(e)}(e)\Bigg],
\end{aligned}$$

which determines $\vartheta(s, x) \equiv \frac{x\varrho(g,x)}{1+\nu_l}$ defined as in Section 4.2.

C Data

C.1 Data Sources

Table C.1 reports the source of data used for the calibration of the model. We follow the same methodology as Ábrahám et al. (2025).

Labor Input. For the aggregate labor input $n_{i,t}$, we use two series from AMECO, the aggregate working hours $H_{i,t}$ and the total employment $E_{i,t}$ of each country over the period 1980-2019. We calculate the normalized labor input as $n_{i,t} = H_{i,t}/(E_{i,t} \times 5200)$, assuming 100 hours of allocatable time per worker per week. However, for most of the data moment computations, we use $H_{i,t}$ directly.

Consumption. We fit the observed fiscal behavior across the selected countries, so that we use directly the data measures of household and government consumption and trade balance to calibrate the model.

Government. We use the general government consolidated gross debt. As noted by Bocola et al. (2019), matching the overall public debt allows a quantitative sovereign default model to better fit crisis dynamics. Regarding the risk-free rate, we take the average real short-term yield of Germany after the introduction of the euro from 2000 to 2019. Similarly, the interest rate spread corresponds to the difference with the nominal long-term yield of

Table C.1: Data Sources and Definitions

Series	Times	Sources	Unit
Output	1980–2019	AMECO (OVGD)[a]	1 billion 2015 constant euro
Consumption	1980–2019	AMECO (OCNT)[b]	1 billion 2015 constant euro
Working hours	1980–2019	AMECO (NLHT,NLHA)[c]	1 million hours
Employment	1980–2019	AMECO (NETD)	1000 persons
Government debt	1980–2019	AMECO (EDP)	end-of-year percentage of GDP
Debt service	1980–2019	AMECO (UYIGE)[d]	end-of-year percentage of GDP
Trade balance	1980–2019	AMECO (OXGS, OMGS)[e]	end-of-year percentage of GDP
Bond yields	1980–2019	AMECO (ILN,ISN,ISRV)[f]	percentage
Debt maturity	1980–2019	OECD, EuroStat, ESM[g]	years
Labor share	1980–2019	AMECO[h]	percentage

[a] Strings in parentheses indicate AMECO labels of data series.

[b] PWT 8.1 values for Greece in 1980–1982.

[c] Total and average working hours.

[d] AMECO for 1995–2019; European Commission General Government Data (GDD 2002) for 1980–1995.

[e] Exports minus imports of goods and services.

[f] Nominal long-term yield, nominal and real short-term yield. A few missing values for Greece and Portugal replaced by Eurostat long-term government bond yields.

[g] Average across different data sources, identical to Ábrahám et al. (2025).

[h] Compensation of employees (UWCD) plus gross operating surplus (UOGD) minus gross operating surplus adjusted for imputed compensation of self-employed (UQGD), then divided by nominal GDP (UVGD).

Germany between 2000 and 2019. We compute the coupon rate as the ratio of debt service over debt. Finally, as noted by Ábrahám et al. (2025), the information on the maturity structure of the government debt is not comprehensive for the country considered. The overall time coverage is unequal across countries: 1998–2015 for Greece, 1991–2015 for Spain, 1990–2015 for Italy, and 1995–2015 for Portugal.

C.2 Productivity Shock Estimation

The decomposition algorithm takes a transition matrix P and splits it into two transition matrices P^H and P^L such that $P = \varkappa P^H + (1-\varkappa)P^L$ for a fixed weight $\varkappa \in (0, 1)$. Intuitively, P^H is a "high-persistence" matrix in which future states are relatively predictable given the current one, while P^L is a "high-noise" matrix in which transitions are more dispersed across states. We formalize this distinction through the entropy of each row: a row with all its mass concentrated on a single state has zero entropy (perfectly predictable), while a row that spreads mass evenly across all states has maximum entropy (completely unpredictable). We require that every row of P^H has weakly lower entropy than the corresponding row of P^L, so that P^H always represents the more predictable regime. The weight $\varkappa$ can then be interpreted as the probability that a given transition is drawn from the high-persistence regime rather than the high-noise regime.

Because many pairs (P^H, P^L) can satisfy these conditions, we need an additional restriction to pin down a unique decomposition. We do so by choosing P^H to be as similar as possible to a target matrix that blends P itself with a fully deterministic matrix — one that places all probability on the most likely successor state. Specifically, for each row i, the target is

$$t_i(\lambda) = (1-\lambda)\, p_i + \lambda\, e_{k^*}, \qquad k^* = \underset{j}{\operatorname{argmax}}\; p_{ij}, \tag{C.1}$$

where e_{k^*} is a vector that places all mass on state k^* and $\lambda \in [0,1]$ is a tuning parameter. When $\lambda = 0$ the target is simply the original row p_i, producing a decomposition where P^H and P^L are as similar to each other as possible while still respecting the entropy ordering. When $\lambda = 1$ the target is fully deterministic, pushing P^H toward a Dirac and P^L toward the uniform distribution. Intermediate values of λ produce decompositions that lie between these two extremes. To save on parameters, we set $\lambda = \varkappa$.

In practice, the decomposition is solved separately for each row of P via a quadratic program that finds the p_i^H closest to the target $t_i(\lambda)$, subject to both p_i^H and the implied $p_i^L = (p_i - \varkappa\, p_i^H)/(1-\varkappa)$ being valid probability vectors. Since entropy is a nonlinear function, the entropy ordering $H(p_i^H) \leq H(p_i^L)$ cannot be imposed as a standard linear constraint. We therefore enforce it through an iterative procedure: after each solve, if the entropy ordering is violated we add a linearized constraint that pushes p_i^H to be more concentrated than p_i^L, and re-solve. This loop converges quickly in practice. The output is a decomposition that satisfies $P = \varkappa P^H + (1-\varkappa)P^L$ exactly, with P^H always being the lower-entropy (more predictable) matrix.

D Additional Quantitative Results

This section presents additional quantitative results for the IMD economies and the Fund contracts

Figure D.1 depicts a simulated path of debt and g in steady state for the three IMD economies. The black line depicts the exogenous ζ shock which is the same for all three specifications. The gray dotted lines represent the occurrence of defaults. Debt dynamics

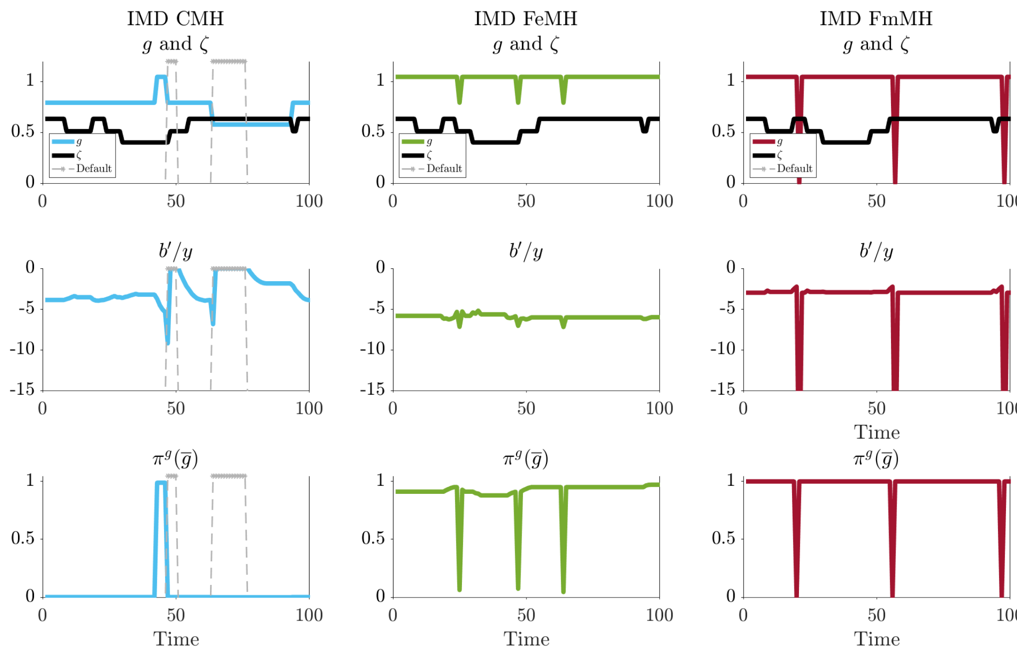

Note: The figure depicts a simulated path of debt and effort in steady state in the different IMD economies.

Figure D.1: Simulation – IMD economies

differ sharply across specifications: swings in the debt ratio are largest under mean-cost flexible MH and smallest under entropy-cost flexible MH. The two flexible MH specifications also share a common pattern absent from canonical MH: g, and the probability of $\overline{g}$, rise together with ζ.

Figure D.2 compares the probability distribution of g across the IMD and Fund economies. Among Fund contracts, the two flexible MH specifications move in opposite directions as x rises: mean-cost MH puts less weight on $\overline{g}$, while entropy-cost MH (and, to a lesser extent, canonical MH) puts more. Among IMD economies, both flexible MH specifications put more mass on $\overline{g}$ at specific debt levels forming an inverse U shape, while under canonical MH the probability is unchanged over the depicted states.

Figure D.3 depicts a simulated path of the different Fund contracts. Canonical and back-loaded MH follow a similar dynamic for the relative Pareto weight, reflecting the small IC multiplier noted above (Figure 1): the latent Pareto weight x' (solid yellow) closely tracks $\bar{x}'$ (dotted orange). Under flexible MH, movements in g are more pronounced under mean-based cost. However, the relative Pareto weight seems equally oscillatory under the two flexible MH specifications. This is in line with the large volatility of output and the low volatility of consumption under mean-cost flexible MH.

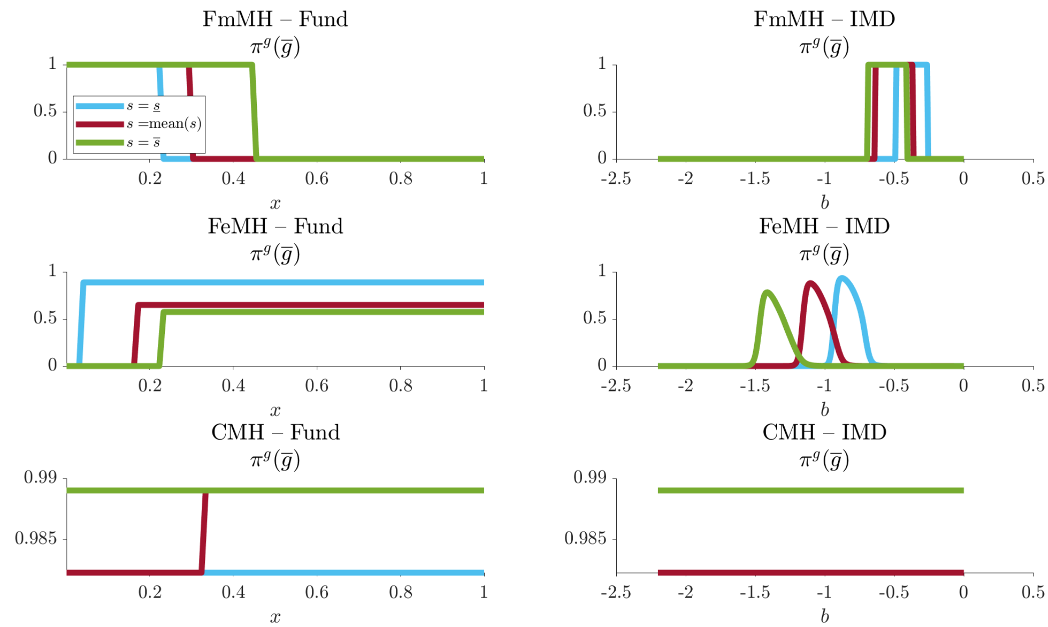

Note: The figure depicts the main policy functions as a function of the relative Pareto weight x for the Fund economies and as a function of b for the IMD economies.

Figure D.2: Probability distribution in the IMD and Fund economies

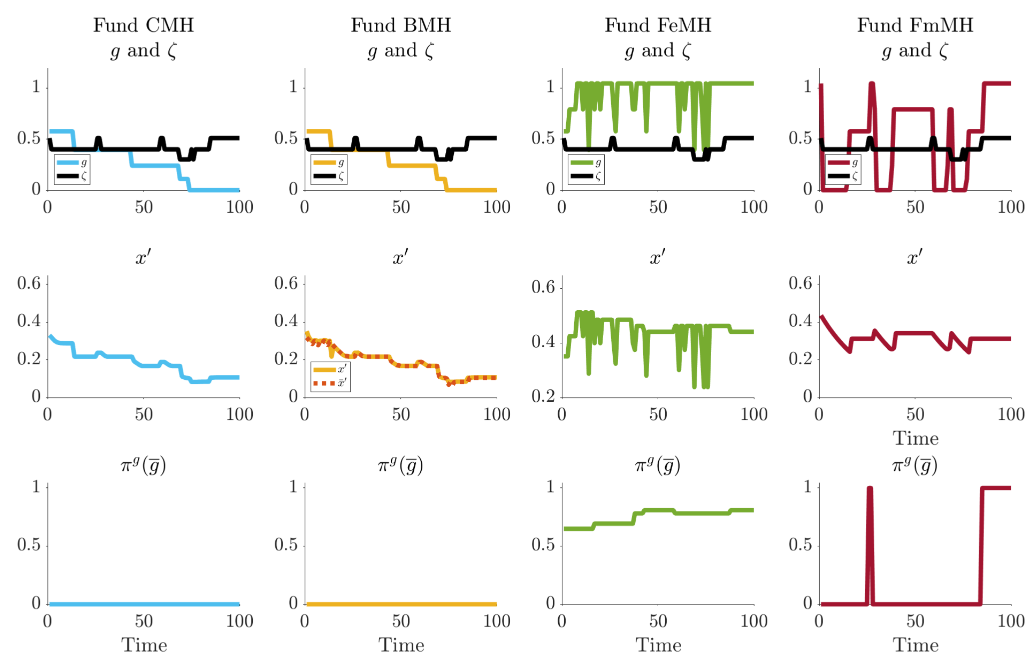

Note: The figure depicts a simulated path of debt and effort in steady state in the different Fund economies.

Figure D.3: Simulation – Fund economies