EIEF WORKING PAPER SERIES

Einaudi Institute for Economics and Finance

EIEF Working Paper 25/12

November 2025

Hyperbolic Discounting with Random Gratification

By

Facundo Piguillem
(EIEF and CEPR)

Liyan Shi
(CMU and CEPR)

Hyperbolic Discounting with Random Gratification*

Facundo Piguillem[1] and Liyan Shi[2]

[1]EIEF and CEPR
[2]CMU and CEPR

September 2025

Abstract

We analyze the dynamic problem of decision makers with quasi-hyperbolic discounting and random shocks to temptation. We show that this problem is equivalent to that of a standard consumer-saver who assigns biased weights to future shocks. This equivalence provides a straightforward methodology for characterizing the Markov equilibrium with hyperbolic agents. Through this equivalent problem, we provide conditions for the existence and uniqueness of the equilibrium. If the weights constitute a probability measure, the decision maker can be interpreted as optimistically biased, ensuring a unique equilibrium with continuous decision rules and implying no value for commitment devices. Otherwise, there is intertemporal "conflict" between present and future selves: if the conflict is limited, uniqueness is guaranteed.

JEL classification: E2, E7, H1.

Keywords: Quasi-hyperbolic discounting, optimism, political economy, time inconsistency.

*We thank Manuel Amador, Roland Bénabou, Ali Shourideh, and seminar participants at CMU Tepper, Minneapolis Fed, EIEF, SED Winter Meetings 2024, Cornell, Princeton, Padova, LMU Munich, Ca' Foscari Venice, UNC, St. Louis Fed, Boston University, Philly Fed, Rice-LEMMA Monetary Conference, SAET 2025 Meetings, and SITE 2025 Conference for their insightful discussions and comments. Email: facundo.piguillem@gmail.com, liyans@andrew.cmu.edu

1 Introduction

Self-control problems, usually represented by preferences with hyperbolic discounting, have long been recognized in numerous economic and social choices (see, e.g., Strotz, 1955; Phelps and Pollak, 1968). They affect many aspects of life, from seemingly mundane everyday struggles such as "the diet starts tomorrow" to significant economic decisions like retirement saving. Beyond individual behavior, the outcomes of many political decisions resemble those of individuals with time-inconsistent preferences, even when all parties participating in decision-making are endowed with standard time-consistent preferences. These issues are of paramount importance as a rationalization for many corrective policies, forming the basis for minimum savings requirements (Amador, Werning, and Angeletos, 2006) and giving rise to the need for fiscal rules (Halac and Yared, 2014, 2018).

Although widely recognized as a relevant feature of economic choices, models capturing such behavior are difficult to work with. Textbook dynamic-programming tools typically break down, and these settings are plagued with multiplicity of equilibria and discontinuous decision rules (see Krusell and Smith, 2003; Chatterjee and Eyigungor, 2016; Cao and Werning, 2018). The difficulties arise from strategic incentives inherent in these intrapersonal games: if my future self is going to squander the money, I may prefer to spend it now to prevent that outcome. Many efforts have been devoted to tackle the intractability. One approach is to assume *naïveté* over *sophistication*: if agents are unaware of their future self-control problems, the strategic responses that generate technical difficulties are eliminated. Another approach is the elegant *instantaneous gratification* developed by Harris and Laibson (2013), where present self is in control only for a split second. However, these workarounds are not well-suited in many important settings, where successive decision makers are strategic and in power for some duration.[1]

We propose a new methodology to analyze a broad class of hyperbolic models with taste shocks, drawing on tools from the mechanism-design literature. We show that the hyperbolic agent's problem is equivalent to that of a consumer-saver who discounts exponentially but forms optimistically biased "expectations" of future shocks. This equivalence yields three main insights. First, it establishes a connection between present bias and optimism bias, two literatures that have largely developed independently. We show that these behavioral phenomena are equivalent under certain conditions. Second, while the original problem is

[1]In political-economy applications that map to hyperbolic discounting, decision makers typically remain in power for nontrivial duration. The instantaneous-gratification model assumes instant transition and is at odds with observed tenures. It can also lead to extreme predictions. For example, default risk is central to both household and sovereign decisions. Felli, Piguillem, and Shi (2025) show that with short-term debt the instantaneous-gratification model implies a de facto borrowing limit, so default risk never arises in equilibrium.

elusive, the transformed problem facilitates a sharp characterization of equilibrium properties, such as existence, uniqueness, and continuity. Third, the equivalence offers insights for the demand for commitment devices.

To fix ideas, we consider a present–future selves model of consumption–saving and study the Markov equilibrium of the intrapersonal game. As in Harris and Laibson (2013), the present self loses control to the future self stochastically, but with two differences: time can be either discrete or continuous, and when the temptation arrives, a taste shock is also realized. These shocks embody the random gratification we emphasize, a feature common in such models. In a nutshell, this is essentially the β-δ framework with stochastic taste shocks.

As in the mechanism-design literature, we apply the generalized envelope theorem of Milgrom and Segal (2002) and Sinander (2022) to any candidate equilibrium. This allows us to express the continuation value function when the agent is out of control in terms of value function of the agent in control. We show that the mapping between the two is linear and involves only the underlying shock distribution and the present bias parameter. Using this mapping, we can eliminate the equation keeping track of the continuation value and collapse the usual dual Bellman equations into a single equation resembling a standard recursive consumption-savings problem.

In the transformed problem, although the decision maker discounts future geometrically, present bias is not lost: it is instead captured by a distorted distribution of future shocks, which is dominated by the true one, implying an optimistic bias. Intuitively, the optimistically biased agent forms rosy expectations of the future, believing that rainy days are unlikely tomorrow, and thus convinces herself to spend more today. ***Sophistication***, the present self's awareness of future lack of self-control, in the original problem is not lost either: it is reflected by the lack of expectations updating in the transformed problem. This failure to learn precisely captures the present self's awareness that her future selves will remain optimistically biased and continue to take biased actions.

We have thus far discussed expectations in a loose way, because the distorted distribution of shocks may not be a well-defined probability measure, in the sense that it may imply negative weights for some taste shocks. Technically, our optimistic agent computes expectations using a signed measure rather than a standard probability measure. While this adds some technical difficulties, it also provides a deeper and sharper understanding of the underlying mechanisms at play.

The distorted distribution introduces a notion of ***conflict***, which we distinguish from ***disagreement***. The present self always disagrees with her future selves about desirable actions, but they may not necessarily be in conflict with each other—actions that improve a future agent's value may also improve the present self's value. However, when a negative weight is assigned to a future state, any value-improving action by and for this future type reduces the present self's value. As a result, the agent in control does not want the future self to maximize. It is in this sense that we say that conflict arises. The negative weights also reflect strategic complementarities in these games that can generate equilibrium multiplicity: if I am going to act irresponsibly tomorrow, I should behave badly today.

When the distorted distribution features negative weights, characterizing the solution is more complex. Even though the equilibrium outcome can be written as the solution to a recursive optimization problem, the associated Bellman operator between current and future value functions is not guaranteed to be monotone, so the standard tools no longer apply. Nevertheless, we provide sufficient conditions under which, although monotonicity may be lost, the mapping remains contractive, thereby yielding a unique equilibrium. This holds as long as the cumulative negative weights are not too large, that is, conflict is limited.

In the complete absence of conflicts, the transformed problem is effectively a standard consumption-savings problem. In this case, the monotonicity of the Bellman operator is restored, and textbook dynamic-programming tools are back in business, delivering the existence, uniqueness, continuity of the equilibrium. We also establish necessary and sufficient conditions under which all weights are positive: the true shock distribution is sufficiently fat-tailed relative to the degree of present bias. While random shocks generally act as a smoothing force, this no-conflict condition formalizes how much dispersion is required to smooth equilibrium outcomes.

We also show that intertemporal conflict, rather than mere disagreement, drives the demand for commitment. To assess the value of commitment, consider a delegation problem in which the agent losing control delegates decisions to her future self of unknown type. Recall that when there is no conflict, although the present self disagrees with her future self about the precise desirable action, any value-improving action for the future self still improves the present self's value. The present self therefore would not wish to constrain future decisions and prefers full flexibility to respond to shocks. That is, they agree to disagree. Hence, the conventional wisdom that hyperbolic discounting necessarily entails a demand for commitment need not hold, and the use of commitment devices may not provide as clear an empirical identification for hyperbolic preference.

The transformation used to study dynamic intrapersonal games can have broader applicability. For instance, the equivalence extends to situations of future bias, mapping to a pessimistic agent. While our analysis is largely in discrete time, the same transformation and equivalence hold in continuous time. In the same way the lack of monotonicity precludes the use of Blackwell's sufficient conditions in discrete time, the comparison principle is also lost in continuous time. Nevertheless, we present a modified comparison principle that is applicable to these settings. The conditions for the existence of a unique (viscosity) solution are analogous to those in discrete time. In both settings, the transformation lends itself to robust efficient numerical tools for quantitative applications.

1.1 Literature

Since the seminal work of Strotz (1955), a large literature has sought to understand the subject. Early important contributions date back to quasi-hyperbolic setting of Phelps and Pollak (1968) and the highly influential work of Laibson (1997). Yet, the technical difficulties to analyze these types of settings have slowed down the progress.

Many complications related to the subject were pointed out early on, as in Krusell and Smith (2003), and others surfaced later, as in Chatterjee and Eyigungor (2016) and Cao and Werning (2018). Although some have been resolved in particular settings, they remain unresolved in more complicated and empirically relevant contexts. Nowadays, aside from the simple, yet restrictive, approach by Harris and Laibson (2013), there is no clear systematic way to approach these problems. In some cases, this simple approach has offered very interesting results, as in Maxted, Laibson, and Moll (2024) and Maxted (2025). Another approach is to change the game-theoretical framework altogether, as in Bassetto, Huo, and Ríos-Rull (2024), but it is applicable under very specific conditions. As a result, other than the time-consistent preference for commitment of Gul and Pesendorfer (2001, 2004), even in the most recent studies, the workaround is to assume that agents are unsophisticated, as in Choukhmane and Palmer (2024) and Choukhmane (forthcoming), thereby abstracting from strategic behavior in individual decision making.

At the same time, other literatures, such as dynamic mechanism design, have been progressing and generating new tools to overcome technical difficulties, as seen in Amador et al. (2006), Amador and Bagwell (2013), Halac and Yared (2014, 2018, 2022), etc. In this regard, new theoretical results that are crucial as inputs in mechanism design, like those by Milgrom and Segal (2002) and Sinander (2022), lay the groundwork for bridging the gap for the hyperbolic-discounting toolkit. Our transformation, which recasts the problem as choice under a distorted distribution, resembles the robust-control approach of Hansen and Sargent (2001, 2008), both enabling sharp characterizations of equilibrium.

Models with hyperbolic discounting are especially relevant in political economy, where it is possible to show that the usual dynamic "political friction" is isomorphic to quasi-hyperbolic discounting (see, for instance, Piguillem and Riboni, 2020). There is a large body of literature building on the seminal works by Alesina and Tabellini (1990) and Persson and Svensson (1989), and more recently Battaglini and Coate (2008). This body of work mostly focuses on debt patterns, but the same arguments apply to every dynamic decision, including the highly relevant climate policies. Initially, the findings by Barro (1979) seemed to indicate that such friction was not present, since the U.S. economy appeared to follow an optimal path. However, the work by Roubini and Sachs (1989) presented suggestive evidence that it is indeed the case. Apart from some attempts to disentangle whether the friction was driven by ideological reasons or the "tragedy of the commons," Pettersson-Lidbom (2001), the quantification of the friction itself has been elusive. For this reason, much of the work has moved to theoretical grounds, designing rules to prevent inefficient outcomes, e.g., Azzimonti, Battaglini, and Coate (2016) and Cunha and Ornelas (2017), or developing models that generate patterns consistent with the data, e.g., Müller, Storesletten, and Zilibotti (2016). Yet, a precise, agreed-upon measure does not exist.

Finally, the explicit link that we uncover between hyperbolic discounting and biased expectations relates to the behavioral literature addressing non-bayesian updating, as in Epstein (2006) and Epstein, Noor, and Sandroni (2008), projection bias, as in Loewenstein, O'Donoghue, and Rabin (2003), and strategic ignorance and self-confidence, as in Carrillo and Mariotti (2000) and Bénabou and Tirole (2002).

2 A random gratification model

2.1 Environment

Time is discrete and infinite, $t \in \{0, 1, 2, \cdots\}$. In each period, there is an agent in control of decisions, who receives an exogenous source of income y and faces a consumption choice c. The agent can save and borrow through a risk-free bond b at an interest rate r, subject to an exogenous ad hoc borrowing limit $\underline{b} > -y/r$.

The agent attributes a value to her consumption needs, which we denote as her "taste" type θ. Specifically, the agent derives a utility from consumption:

$$\theta u(c),$$

where $u(\cdot)$ is strictly increasing and strictly concave, with first-order and second-order derivatives satisfying $u'(\cdot) > 0$ and $u''(\cdot) < 0$. A high taste state is thus associated with a higher marginal utility from consumption, representing a "rainy day" when the agent needs to incur a high level of consumption.

Present-future selves. The agent loses control with probability $\lambda \in (0, 1]$ each period. A new agent then takes control and draws a new taste, which is *i.i.d.* across agents. The

Figure 1: Present-future selves

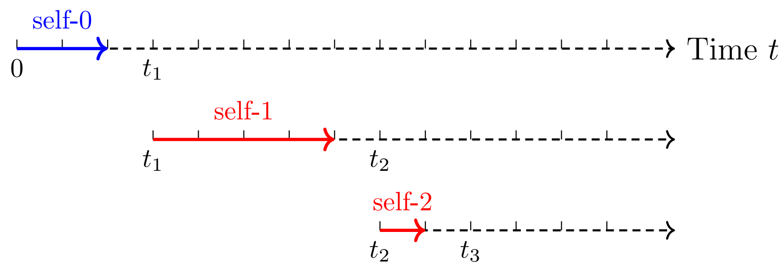

Notes: The figure shows an example timeline, where self-0 is in control for 2 periods, self-1 in control for 4 periods, and self 2 for 1 period.

taste is drawn from a set $\Theta \equiv [\underline{\theta}, \bar{\theta}]$ according to a cumulative distribution function $F(\cdot)$, with an expected value $\mathbb{E}[\theta] = 1$. The distribution function $F(\theta)$ admits a differentiable and bounded density $f(\theta)$.[2]

All agents, whether in control or not, discount the future exponentially with a discount factor $\delta < 1$. However, the present self discounts the utilities of future selves by an additional factor $\beta \in (0, 1]$. We refer to β as the present-bias parameter.

The environment encompasses a *deterministic present-biased* model as a limiting case: the transition from present to future selves occurs every period, and the taste for consumption is constant, i.e., $\lambda = 1$ and $\underline{\theta} = \bar{\theta} = 1$. While we interpret the present-future selves as an intrapersonal self-control issue and focus on household consumption-savings problem, the interpretation extends to present-future governments in dynamic political economy.

Hyperbolic agent's problem. We consider a Markov equilibrium of a game between present self and future selves. That is, agents' decisions depend only on the current state, $\{\theta, b\}$, and do not depend on the history of past realizations or on time. The Markov equilibrium follows the standard approach commonly used in the literature. Alternative concepts include subgame perfect equilibrium and "organizational" equilibrium as discussed by Bassetto et al. (2024). The problem of the agent in control is:

$$w(\theta, b) = \sup_{c, b'} \Big\{ \theta u(c) + \delta \Big[(1 - \lambda) w(\theta, b') + \lambda \beta \mathbb{E}[v(\theta', b')] \Big] \Big\} \tag{P}$$

[2] In the context of government debt choice, shifts in preferences may reflect changes in the constituency's views on the social value of spending. Such evolving opinions could, in turn, influence the country's fiscal policy stance. Another interpretation is that preferences vary in response to business cycles. For instance, Amador et al. (2006) show that under exponential utility, income shocks are equivalent to taste shocks.

subject to

$$c + b' \leq y + (1 + r)b \tag{1}$$

$$b' \geq \underline{b} \tag{2}$$

$$v(\theta, b) = \theta u(c(\theta, b)) + \delta\Big[(1 - \lambda)v(\theta, b'(\theta, b)) + \lambda\mathbb{E}[v(\theta', b'(\theta, b))]\Big]. \tag{3}$$

To understand problem (P), we first note that, as long as $\lambda = 0$, it represents a standard consumption-savings problem, subject to the budget constraint (1) and the borrowing constraint (2). However, when $\lambda > 0$, the present bias starts to play a role. The last term in (P) captures the effect of random control: with probability λ the present self loses control and is replaced by her future self, whose choices lead to a continuation value assessed as $\beta\mathbb{E}[v(\theta', b')]$. Here, two features are noteworthy. First, the future self's taste is not yet known, requiring an expectation over future types $\theta' \in \Theta$. Second, and more importantly, the present self discounts the unbiased continuation value by the additional factor β. Nevertheless, the present self has the discretion to decide her consumption and savings while in control. We denote these decisions by the functions $\{c(\theta, b), b'(\theta, b)\}$.

For a given state $\{\theta, b\}$, the unbiased continuation value $v(\theta, b)$ follows the companion "accounting" equation (3). One key difference between (P) and (3) is the bias factor β in the former equation: once the present self is out of control, all future allocations are uniformly discounted by β. Having accounted for this bias, all future allocations are geometrically discounted by a discount factor δ. Equation (3) also captures that the present self takes as given that her future selves will follow their own optimal choices $\{c(\theta, b), b'(\theta, b)\}$. Sophistication is properly captured here. The present self is aware of future self-control problems: when she is no longer in control, her future selves will exhibit the same biases as she did.

We focus our analysis on cases where $0 < \beta \leq 1$, leaving aside the case with extreme myopia, $\beta = 0$. When $\beta = 0$, the present self behaves as a time-consistent agent who assigns a geometric discount factor $\delta(1 - \lambda)$ to the future. In this case, the problem coincides with the consumption-savings problem of a time-consistent agent. Thus, the bulk of knowledge from that literature can be applied immediately. Furthermore, in Section 4.2 we discuss the consequences of future bias, i.e. $\beta > 1$.

Definition 1 (Markov Equilibrium)**.** A Markov equilibrium is a collection of value functions and decision functions

$$\{w(\theta, b), v(\theta, b), c(\theta, b), b'(\theta, b)\}_{\theta\in\Theta, b\geq\underline{b}}$$

that solve problem (P).

2.2 Maintained assumptions

We introduce the following assumptions, which will be maintained for the remainder of this paper. The first assumption concerns the asset position.

Assumption 1 (Boundedness)**.** *The borrowing limit satisfies* $\underline{b} > -\frac{y}{r}$ *and there is a maximum allowed level of assets* $\bar{b} < \infty$*, so that* $b \in [\underline{b}, \bar{b}]$.

This assumption guarantees the boundedness of the problem. While the assumption may appear restrictive, it is imposed only for technical convenience. It is not essential for the main results to hold and can therefore be relaxed. Doing so, however, would introduce additional complexity in notation and definitions, with little gain in understanding our main conceptual contributions.[3] As an alternative, we can assume that $u(\cdot)$ is bounded or that consumption is bounded above and below from zero.[4] Moreover, it is possible to provide conditions under which the assumption is without loss of generality: as long as the interest rate is sufficiently low relative to the discount factor and the lowest type $\underline{\theta}$, the upper bound on the asset position $\bar{b}$ never binds in equilibrium.

Assumption 2 (Distribution)**.** *The density function* $f(\cdot)$ *is differentiable and satisfies* $|f(\theta)| < \bar{f}$ *and* $|f'(\theta)| < \bar{f}$ *for all* $\theta \in \Theta$*, where* $\bar{f} < \infty$*.*

The assumption requires the distribution of taste types to be smooth, which in turn guarantees that the adjusted weights that we introduce in the next section are bounded. In essence, the assumption only requires some dispersion in the taste shocks and thus rules out deterministic hyperbolic models. This is a mild condition and is satisfied by all commonly used distribution functions.

3 Transformation

In this section, we show that the Markov equilibrium can be equivalently represented as the recursive problem of an agent who assigns biased weights to future taste shocks.

3.1 Mapping to a recursive problem

We start by establishing a useful relation in the following lemma.

[3]See Alvarez and Stokey (1998) and Rincón-Zapatero and Rodríguez-Palmero (2003) for existence and uniqueness results of dynamic programming in environments with unbounded returns.

[4]We rule out pathological equilibria in which the present self consumes zero, anticipating that future selves will consume zero. When the utility of zero consumption is negative infinity—for example, under a log utility function—such pathological cases can indeed be the basis of an equilibrium. We disregard these cases as in Harris and Laibson (2013) and Cao and Werning (2016).

Lemma 1 (Value functions)**.** *In any Markov Equilibrium, for all $b \in [\underline{b}, \bar{b}]$, the value functions are absolutely continuous in θ and, for almost all $\theta \in \Theta$, satisfy*

$$\beta v(\theta, b) = w(\theta, b) - (1-\beta)\theta w_\theta(\theta, b). \tag{4}$$

Proof. See Appendix A.1. □

The relation in Lemma 1 follows from the linearity of the value functions in the (exogenous) state θ and the parameter β. The key step is to appeal to the Envelope Theorem in the sequential representation of problem (P). It may appear surprising to many readers that only Assumption 1 is required since the classical Envelope Theorem relies on the differentiability of the optimized indirect value function. However, the modern version of the Envelope Theorem does not rely on it, as shown by Milgrom and Segal (2002) and Sinander (2022), and holds for arbitrary choice sets. For this reason, the relationship in Lemma 1 holds *almost* everywhere. When it fails, it does so only in a set of measure zero. What is needed is that the partial derivative of the utility function with respect to θ exists, and that the family $\{\theta u(c_t)\}_{\forall c_t}$ is *absolutely equicontinuous* in θ. The first condition is obvious given the linearity, while Assumption 1 ensures that the second condition is also met.[5] As pointed out by the latter paper, this assumption can be relaxed, so it is by no means necessary for Lemma 1. We have chosen to work with bounded payoffs to emphasize our conceptual contribution and to avoid dealing with technical complications, which, despite generalizing the results, would add little to the substance of the insights.

Another key feature of problem (P) is that the constraint set does not directly depend on θ. Departures from this independence would not invalidate the lemma per se, but they could introduce additional terms to the relationship. Such type-dependent constraint sets arise, for instance, when either the interest rate, the borrowing constraint, or both depend directly on θ. Nevertheless, in applications to government debt, Felli et al. (2025) extend this result to an environment with sovereign default, where both the interest rate and the borrowing limit endogenously depend on θ.

Lemma 1 allows us to reduce the complexity of problem (P). Specifically, we can replace the hard-to-handle continuation value function $v(\theta, b)$ using the relation in equation (4) and subsequently drop the companion equation (3). This substitution introduces the term $w_\theta(\theta, b)$, which can be further eliminated through integration by parts. This step is possible because Lemma 1 also implies that $w(\theta, b)$ is absolutely continuous in θ. It is in this process

[5] Our assumption, i.e., bounded payoffs, is strong and can be significantly relaxed. It makes sure that the environment also satisfies the conditions for Theorem 2 of Milgrom and Segal (2002) and, more specifically, corresponds in essence to Example 1 in Sinander (2022). For a definition of *absolute equicontinuity* see Definition 1 in the latter.

that a distorted "distribution" appears, as we present in the following proposition.

Proposition 1 (Equivalence)**.** *The collection $\{w(\theta,b), c(\theta,b), b'(\theta,b)\}$ is part of a Markov Equilibrium if and only if it solves the following recursive problem:*

$$w(\theta,b) = \sup_{c,b'} \left\{ \theta u(c) + \delta \left[(1-\lambda) w(\theta,b') + \lambda \int_{\underline{\theta}}^{\bar{\theta}} w(\theta',b') dG(\theta') \right] \right\}, \qquad (\hat{P})$$

subject to the budget constraint (1) and the borrowing limit (2), where the recursive agent assigns the following cumulative weights to the taste shocks:

$$G(\theta) = \begin{cases} F(\theta) + (1-\beta)\,\theta f(\theta), & \text{for } \theta \in [\underline{\theta}, \bar{\theta}) \\ 1, & \text{for } \theta = \bar{\theta}. \end{cases} \qquad (5)$$

Proof. See Appendix A.2. □

Behind the equivalence result in Proposition 1 is the mapping in the continuation value functions. Specifically, when the agent loses control, her expected continuation value can be expressed as follows:

$$\beta \int_{\underline{\theta}}^{\bar{\theta}} v(\theta,b)\, dF(\theta) = \int_{\underline{\theta}}^{\bar{\theta}} w(\theta,b)\, dG(\theta). \qquad (6)$$

The equivalence suggests that the Markov equilibrium can instead be characterized by problem ($\hat{P}$), which involves a single Bellman equation and a single value function as in standard consumption-savings problems. However, this equivalence does not necessarily imply that the equilibrium features are the same as those in the standard consumption-savings model. The key distinction lies in the adjusted function $G(\theta)$, which may not constitute a well-defined probability measure. For this reason, we will refer to it as a weighting function rather than a distribution function.

3.2 Properties of the biased weights

We now study the properties of the weighting function $G(\theta)$ and compare it to the true distribution of the taste shocks.

A story of biased expectations. One interpretation of Proposition 1 is that the Markov Equilibrium of an agent with hyperbolic discounting is equivalent to that of an agent with exponential discounting but biased expectations about future shocks. This bias is persistent: even after observing shock realizations, she does not update or correct her expectation accordingly. Although the agent appears to discount the future in a time-consistent manner, she updates her expectations inconsistently.

To see the biases in expectations more clearly, a quick inspection of equation (5) reveals that the transformation tends to shift weights downward, and some straightforward calculations show that:

Corollary 1 (Optimism)**.** *The equivalent agent is optimistically biased about the future:*

(i) $F(\cdot)$ *dominates* $G(\cdot)$: $F(\theta) \leq G(\theta)$, $\forall\theta \in \Theta$, *with strict inequality for some* $\theta \in \Theta$.

(ii) The adjusted weights imply a weighted average $\int_{\underline{\theta}}^{\bar{\theta}} \theta dG(\theta) = \beta$.

Proof. See Appendix A.3. □

Corollary 1 states that the equivalent agent assigns more weights to shocks with lower future spending needs, resulting in an expected value of $\int_{\underline{\theta}}^{\bar{\theta}} \theta dG(\theta) = \beta$, which is below the true expected value $\mathbb{E}[\theta] = 1$. We say that the equivalent agent is **optimistic**: she acts as if her future consumption needs would be lower than those implied by the true distribution $F(\theta)$. Intuitively, as the agent anticipates lower consumption needs in the future, believing that the odds of "rainy days" when she would need to spend a lot of money are lower, she justifiably spends more and saves less today. From this perspective, present-biased agents are behaviorally equivalent to agents who display an optimistic bias about the future. In other words, myopia can serve as a microfoundation for optimism, and, conversely, optimism can provide a foundation for myopia.

Sophistication, the present self's awareness of future lack of self-control, in the original problem is reflected by no updating of expectations in the transformed problem. The optimistic agent forms a biased expectation and does not learn about the true distribution of shocks, even after disproportionately many occurrences of high taste events. This persistent bias in expectations precisely captures the present self's awareness that her future selves will remain optimistically biased and continue to take biased actions.

The transformation of the hyperbolic agent's problem bears resemblance to the robust control approach introduced by Hansen and Sargent (2001, 2008). Their framework reformulates the problem of a decision maker concerned about model misspecification to one in which the decision maker behaves as if she distorts probabilities to guard against potential misspecification. This probability distortion captures a form of pessimism or caution, by evaluating actions as if adverse scenarios were more likely to occur than suggested by the benchmark model. In a similar vein, our transformation reinterprets the hyperbolic agent's intertemporal inconsistencies as optimistic biases in the probability assessment of future events. Thus, just as robust control captures concerns about misspecification through distorted beliefs, our approach recasts present bias in terms of biased expectations.

Our interpretation of "expectation" thus far has been in a general sense, not strictly confined to the conventional definition in probability theory. The weights assigned to the taste states are given by:

$$dG(\theta) = \begin{cases} (1-\beta)\underline{\theta} f(\underline{\theta}), & \text{for } \theta = \underline{\theta} \\ \left[(2-\beta) f(\theta) + (1-\beta)\theta f'(\theta)\right] d\theta, & \text{for } \theta \in \left(\underline{\theta}, \bar{\theta}\right) \\ -(1-\beta)\bar{\theta} f(\bar{\theta}), & \text{for } \theta = \bar{\theta}. \end{cases} \tag{7}$$

The weights in equation (7) capture how the transformation in equation (5) tends to shift the weights downward. At the bottom of the taste distribution, since there is no type below $\underline{\theta}$, a positive mass of size $(1-\beta)\underline{\theta} f(\underline{\theta})$ can accumulate here. At the top, a mass of size $(1-\beta)\bar{\theta} f\left(\bar{\theta}\right)$ might be taken away. In the interior, the adjusted weights depend on how the density changes as taste increases. Thus, the weights may include negative values in the interior or at the top of the taste distribution. In this sense, $G(\theta)$ is not a probability measure but rather a signed measure.

We characterize conditions under which $G(\theta)$ indeed constitutes a well-defined probability measure, i.e., the weights in equation (7) are positive for all possible taste states. We establish a necessary and sufficient condition for these scenarios below.

Condition 1 (No conflict)**.** *The taste distribution satisfies:*

(i) The speed at which the density decreases is bounded below:

$$\frac{\theta f'(\theta)}{f(\theta)} \geq -\frac{2-\beta}{1-\beta}, \ \forall \theta \in \Theta. \tag{8}$$

(ii) The taste distribution is unbounded at the top:

$$\bar{\theta} = \infty \ \textit{and} \ \lim_{\theta \to \infty} \theta f(\theta) = 0.$$

Condition 1 characterizes all cases in which $G(\theta)$ constitutes a well-defined probability measure. It consists of two items. Item (i) guarantees that the weight in the interior of the taste distribution is positive: as long as the density does not decrease too fast, as in (8), the adjusted weights are positive. Item (ii) guarantees that the weight at the top is positive. At the top, a mass of size $(1-\beta)\bar{\theta} f\left(\bar{\theta}\right)$ might be taken away. To guarantee this mass is zero, if $\bar{\theta}$ is bounded, it requires that the density $f\left(\bar{\theta}\right) = 0$, which would be a problem since

Figure 2: A Pareto distribution example

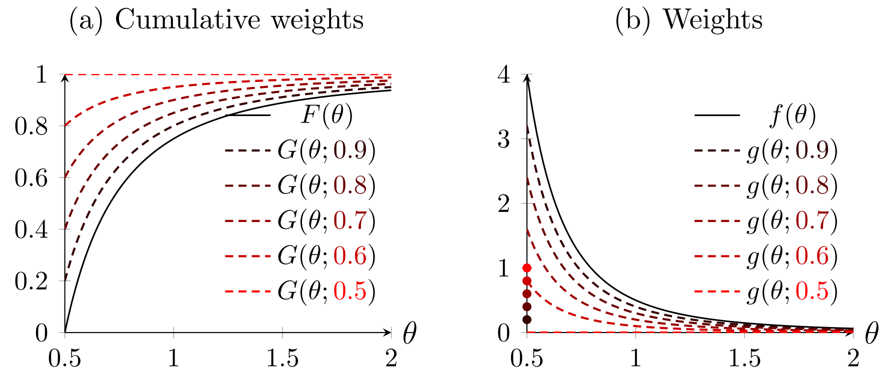

Notes: $F(\theta)$ is a Pareto distribution with parameter $\zeta = 2$ for $\theta \in [0.5, \infty)$. For a given value of β, panel (a) plots the adjusted cumulative weights $G(\theta; \beta)$, while panel (b) plots the weights $dG(\theta; \beta)$ at mass points and the density $g(\theta; \beta) = G'(\theta; \beta)$ at non-mass points.

(8) will be violated at $\bar{\theta}$. Hence, only when the taste distribution is unbounded, $\bar{\theta} = \infty$, a negative mass at the top is avoided.[6] Revisiting Corollary 1, we can now say that the true distribution $F(\cdot)$ first-order stochastic dominates the biased distribution $G(\cdot)$.

Condition 1 essentially requires that the distribution be sufficiently *fat-tailed* relative to the degree of present bias. A broad class of fat-tailed distributions would satisfy this condition. We provide such an example below.

Example 1 (Pareto distribution)**.** *Suppose θ follows a Pareto distribution with a tail parameter $\zeta > 1$ for $\theta \in \Theta = [\underline{\theta}, \infty)$. The lower bound $\underline{\theta} = \frac{\zeta - 1}{\zeta}$ such that the mean is normalized to 1. Condition 1 holds if the tail parameter satisfies $\zeta \leq \frac{1}{1-\beta}$.*

The Pareto distribution exhibits a constant speed of decline in the density, $\theta f'(\theta) / f(\theta) = -\zeta - 1$, $\forall \theta \in \Theta$. The bound in (8) is satisfied when $\zeta \leq \frac{1}{1-\beta}$. To illustrate how present bias translates into a biased expectation, Figure 2 is constructed with $\zeta = 2$, considering varying degrees of bias $\beta \geq 1/2$. Panel (a) shows that the adjusted cumulative weight $G(\theta; \beta)$ always lies above the true one $F(\theta)$, and hence is first-order stochastic dominated by the latter. As the bias becomes more severe, that is, as β decreases, $G(\theta; \beta)$ shifts further upward. Panel (b) depicts the weights, which move downward everywhere relative to the true density, except at the lower bound where a mass accumulates. When $\beta = 0.5$, the optimistic bias is at the extreme of Condition 1: the agent behaves as if she assigns the entire weight to the lowest possible future spending need and expects that event to occur with certainty.

[6] There is no distribution with finite $\bar{\theta}$ that simultaneously satisfies *(i)* and $f(\bar{\theta}) = 0$. We thank Manuel Amador for this observation, who also provided us with a proof, which is available upon request.

Figure 3: A uniform distribution example

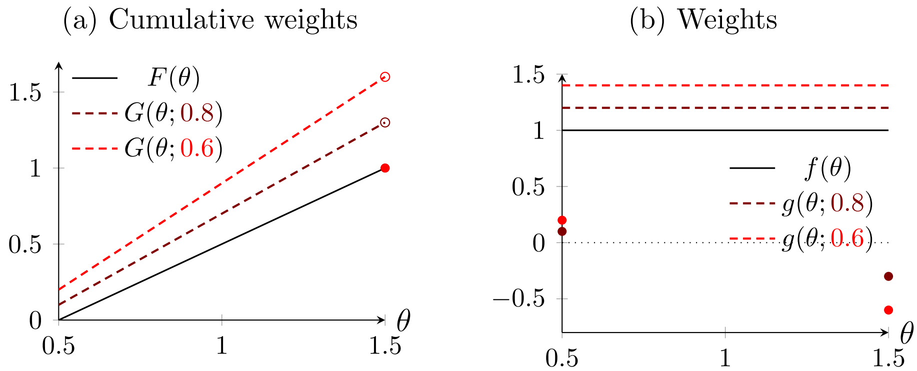

Notes: $F(\theta)$ is uniform distribution over the domain $[0.5, 1.5]$. For a given value of β, panel (a) plots the adjusted cumulative weights $G(\theta; \beta)$, while panel (b) plots the weights $dG(\theta; \beta)$ at mass points and the density $g(\theta; \beta) = G'(\theta; \beta)$ at non-mass points.

A story of disagreement and conflict. The discrepancies between the adjusted weights and the true distribution have two additional implications. Bear in mind that for the agent in control, the expected continuation value upon losing control can be transformed according to equation (6).

First, there is **disagreement** between present and future selves. Specifically, the present self evaluates the continuation value $\beta v(\theta, b)$ in a different way than the future type θ, who values its decisions according to $w(\theta, b)$. Through the lens of Lemma 1, the fact that $\beta v(\theta, b) \neq w(\theta, b)$ translates to the gap in the weights: $F(\theta) \neq G(\theta)$ almost everywhere, even when the present self is still assigning a positive density $dG(\theta) \geq 0$ to that event.

Second, there can be **conflict** between present and future selves. A conflict arises whenever the adjusted weight is negative. What does a negative weight mean? A negative weight $dG(\theta) < 0$ implies that an increase in the value obtained by a future θ-self, $w(\theta, b)$, *reduces* the present self's continuation value. It further implies that the desired actions of the future θ-self make a negative contribution to the expected continuation value of the present self. As a result, the present self does not want the future self to maximize its value and, if possible, would prevent that type from doing it. This manipulation of future actions could occur, for instance, through the use of commitment devices.

Disagreement and conflict have a key difference. Present bias always leads to disagreement, as the present and future selves disagree on the desired actions to be taken. However, if there is no conflict, $dG(\theta) \geq 0$, the present self would agree to let future selves optimize at their discretion: *they agree to disagree*. In contrast, when there is conflict, the present self will actively seek to control the actions of future selves. We show later that this distinction plays an important role in determining the efficiency properties of the equilibrium.

We conclude the discussion with an example where Condition 1 fails, leading to conflicts.

Example 2 (Uniform distribution)**.** *Suppose* $\theta \sim U\left(\underline{\theta}, \bar{\theta}\right)$, *where* $\underline{\theta} + \bar{\theta} = 2$.

In Example 2, the uniform distribution features a constant density, $\theta f'(\theta)/f(\theta) = 0$, $\forall \theta \in \Theta$, thus satisfying *(i)* of Condition 1. However, given that the type distribution is bounded at the top, item *(ii)* of Condition 1 fails, leading to a negative adjusted weight at the top. That is, there is conflict at the top. In Figure 3, we provide an illustration by setting $\underline{\theta} = 0.5$ for $\beta \in \{0.8, 0.6\}$. In panel (a), the cumulative weights are above the true cumulative density function before reaching the top of the taste distribution and jump down to 1 at the top, implying a negative weight there. In panel (b), as present bias becomes more severe, that is, as β decreases, the negative mass increases, leading to a bigger conflict.

4 Equilibrium characterization

We now proceed to analyze the properties of the equilibrium. Working directly with the original problem (P) can be cumbersome and may require imposing numerous assumptions. Instead, with problem ($\hat{P}$), one can appeal to a wide range of tools readily available for studying recursive problems.

We first establish the existence of equilibrium in general. In problem ($\hat{P}$), since the adjusted weights always sum up to one, $\int_{\Theta} dG(\theta) = 1$, and the discount factor $\delta < 1$, a quick observation of the Bellman operator makes clear that it exhibits discounting. One must then be tempted to appeal to Blackwell's sufficient conditions and argue that it is a contraction. However, whenever there are negative weights, the operator is no longer guaranteed to be monotone, and thus one of Blackwell's sufficient conditions is not satisfied.[7] Nonetheless, the failure of monotonicity does not necessarily preclude the existence or uniqueness of equilibrium. In the subsequent analysis, we show that problem ($\hat{P}$) admits a solution, thereby guaranteeing the existence of a Markov equilibrium.

Proposition 2 (Existence)**.** *There exists at least one solution* $\{w(\theta, b), c(\theta, b), b'(\theta, b)\}_{\theta \in \Theta, b \geq \underline{b}}$ *to problem* ($\hat{P}$).

Proof. See Appendix A.5. □

Proposition 2 establishes that the existence of a Markov equilibrium is not a concern. Its existence is guaranteed under fairly weak conditions. Here, we appeal to the Schauder Fixed-Point Theorem adapted to bounded and continuous functions. By contrast, the potential

[7] It is well established that numerical solutions of quasi-hyperbolic discounting problems can often encounter convergence problems (see Chatterjee and Eyigungor, 2016). These issues could originate from the lack of monotonicity, which is straightforward to see in problem ($\hat{P}$), but difficult to grasp in problem (P).

multiplicity of equilibria presents a more substantive challenge. To address this, we next analyze economies with limited conflict and those without conflict.

4.1 Limited conflict

Although problem ($\hat{P}$) has a standard structure, some complications may arise. Indeed, in general, the Bellman operator of a deterministic present-biased consumer is not a contraction mapping, and only in the narrow limit when $\beta \to 1$ it is guaranteed to be contractive (see Harris and Laibson, 2001, 2003). In our random environment, the conditions for contraction and uniqueness can be relaxed. We provide sufficient conditions—the conflict between present and future selves is limited—for the equilibrium to be unique. Loosely speaking, limited conflict requires that the total measure of negative weights is not too large relative to the degree of present bias.

One obvious starting point is to consider economies when there is no conflict between present and future selves at all. The problem ($\hat{P}$) becomes a standard time-consistent consumption-savings problem and shares all of its standard properties. We can apply textbook dynamic programming methods, specifically Blackwell's sufficient conditions for contraction mapping, which not only delivers the existence of its solution but also the uniqueness. As discussed earlier, when conflict is present, the operator can be non-monotone. To understand the uniqueness properties, we need to resort to other tools. One could conjecture that by limiting the amount of conflict the uniqueness could be restored. This conjecture turns out to be true, and we establish sufficient conditions under which it holds.

For convenience, we separate the type set into two subsets: a set of positive measures $\Theta^+ = \{\theta \in \Theta : dG(\theta) \geq 0\}$ and a set of negative measures $\Theta^- = \{\theta \in \Theta : dG(\theta) < 0\}$. By the Jordan Decomposition Theorem these two sets exist and the decomposition is unique. We then let $G^+ = \int_{\Theta^+} dG(\theta)$ denote the measure of the positive set and $G^- = -\int_{\Theta^-} dG(\theta)$ the measure of the negative set such that $G^+ - G^- = 1$.

Condition 2 (Limited conflict)**.** *The cumulative negative weights are limited:*

$$G^- < \frac{1-\delta}{2\lambda\delta}. \tag{9}$$

Naturally, Condition 2 is a weaker condition than Condition 1, as the latter would imply that $G^- = 0$. Moreover, in the condition specified in (9), the dependency on β is implicit in the ranges where negative weights can arise. For instance, in Examples 1 and 2, this condition would translate into a lower bound on β. In the first example, $G^- = 0$ as long as $\beta \geq (\zeta - 1)/\zeta$. In the second example, $G^- = (1-\beta)\bar{\theta} f(\bar{\theta}) > 0$. Hence, the normalization

$\mathbb{E}(\theta) = 1$ implies that the inequality (9) is satisfied as long as

$$(1-\beta)\frac{\bar{\theta}}{2(\bar{\theta}-1)} < \frac{1-\delta}{2\lambda\delta},$$

which is satisfied for a wide range of parameters, including the possibility of the whole range $\beta \in (0, 1]$. With a more severe present bias, it requires a heavier time discounting, a larger shock dispersion, or a lower arrival probability of temptation. Then, we have

Proposition 3 (Uniqueness: limited conflict)**.** *If Condition 2 is satisfied, then there exists a unique solution* $\{w(\theta,b), c(\theta,b), b'(\theta,b)\}_{\theta\in\Theta, b\geq \underline{b}}$ *to problem* ($\hat{P}$).

Proof. See Appendix A.6. □

The proof of Proposition 3 relies on the Contraction Mapping Theorem, which is standard and can be found in many places, e.g. Stokey, Lucas, and Prescott (1989).[8] There are two main complications for its direct application. First, the appealing properties of probability measures can no longer be taken for granted with our weighting function, which is a signed measure, but Lemma A1 in the appendix makes sure that the continuity and boundedness of the operator is preserved, as long as the signed measure is finite. Second, given the operator may not be monotone, additional conditions need to be imposed on negative measures to guarantee it is contractive.

To see the role of Condition 2 more clearly, we rewrite inequality (9) as:

$$\delta\left[1-\lambda+\lambda\int_{\theta\in\Theta}|g(\theta)|d\theta\right] < 1.$$

Here, the total variation of the signed measure $G(\theta)$ is $\int_{\theta\in\Theta}|g(\theta)|d\theta = G^{+} + G^{-} = 1 + 2G^{-}$. The cumulative negative weights drive how much monotonicity may be lost. They also reflect strategic complementarities in this type of games: if I am going to act irresponsibly tomorrow, I should behave badly today. In this sense, the good value tomorrow has a negative impact today. If $G^{-} = 0$, then as any other probability measure its total variation is 1, and problem ($\hat{P}$) is always contractive. When $G^{-} > 0$, the total variation is larger than 1, generating an expansive force. Hence, for the mapping to remain contractive, this force must be counterbalanced by either a lower discount factor or a lower arrival probability of temptation.

We conclude this discussion by noting that, beyond its theoretical advantages, the transformed problem offers numerical benefits. The original problem is difficult to work with in

[8] In particular, see Theorems 9.6, 9.7 and 9.8 in Chapter 9. Since $G(\cdot)$ is not necessarily a probability measure, Lemma 9.5 must be substituted by our Lemma A1 in Appendix A.4.

quantitative analysis; as Chatterjee and Eyigungor (2016) show, numerical algorithms with the inevitable discretization often fail to find an equilibrium, even when an equilibrium is known analytically. By contrast, our transformed problem is robust in finding an equilibrium with the same efficiency required to find a solution in standard consumption-savings models. Hence, it provides a powerful tool for quantitative analysis.

4.2 Future bias

So far we have focused on the case with $\beta < 1$, leading to a present bias. It is straightforward to extend Lemma 1 and Proposition 1 to the case in which $\beta > 1$. This would capture the situations where agents have a future bias, potentially leading to excessive savings or starting an activity too early. Although present bias preferences are more prevalent in the literature than future bias, experimental evidence shows that the latter are also present (see for example Benhabib, Bisin, and Schotter, 2010; Takeuchi, 2011).

In this case, our expectation interpretation is reversed to that of a pessimistic agent. Intuitively, as the agent anticipates higher consumption needs in the future, believing that the odds of "rainy days" when she would need to incur high spending are higher, she spends less and saves more today. That is, future-biased agents are behaviorally equivalent to agents who display a pessimistic bias about the future.

Although most results on equilibrium characterization continue to hold, additional technical assumptions may be required to ensure the existence of a well-defined solution. For instance, whenever $\delta\beta > 1$, even when the utility function is bounded and the state space is compact, the value function may be unbounded. Nevertheless, aside from these additional required assumptions, it is straightforward to show that the limited-conflict Condition 2 for Proposition 3 would remain unaffected. A quick observation of equation (7) reveals that when $\beta > 1$, any negative measure on the boundaries must occur at the bottom rather than the top. For example, returning to the uniform distribution example, this condition is satisfied whenever $(\beta - 1)\frac{\underline{\theta}}{2(1-\underline{\theta})} < \frac{1-\delta}{2\lambda\delta}$. Therefore, for any combination of parameters $\lambda, \delta, \underline{\theta}$, as long as β is not too large, this condition holds, and the equilibrium is unique.

To describe situations without conflict, the inequality (8) in Condition 1 must be reversed to rule out interior negative weights, provided that $\beta < 2$. This limits the speed at which the density increases. Accordingly, it is satisfied by any decreasing density. For example, under the exponential distribution, since $\underline{\theta} = 0$, the adjusted weights would constitute a well-defined distribution function. We emphasize that this brief discussion is not intended as a thorough analysis, but rather as illustrative of potential outcomes. We leave a detailed analysis for future research.

4.3 No conflict and value of commitment

We now turn to the analysis of economies without conflict. Recall that Condition 1, which ensures the absence of conflict, is stronger than Condition 2. Naturally, a unique equilibrium exists. Furthermore, the problem ($\hat{P}$) resembles a standard time-consistent consumption-savings problem and shares all of its standard properties. We can apply textbook dynamic programming methods, specifically Blackwell's sufficient conditions for contraction mapping, which delivers not only the existence and uniqueness of a solution, but also additional regularity properties.

We first show that the equilibrium is continuous. That is, the consumption and saving decisions are continuously increasing in the asset position.

Proposition 4 (Continuity: no conflict)**.** *If Condition 1 is satisfied, then the equilibrium value function $w(\theta, b)$ is strictly increasing and strictly concave for all $\theta \in \Theta$; the decision functions $\{c(\theta, b), b'(\theta, b)\}$ are continuous and increasing in b.*

Proof. The proof is standard and therefore omitted. □

The continuity of the decision functions in Proposition 4 contrasts with jumps observed in deterministic taste models (see, for example Krusell and Smith, 2003; Chatterjee and Eyigungor, 2016). It is intuitive that such jumps, if they occur, result from agents' desire to manipulate future outcomes. In the absence of conflict—and thus no incentive for future manipulation—agents would smooth their consumption. It is also natural that taste shocks play a role in smoothing agents' decisions. The no-conflict Condition 1 guarantees that these shocks are large enough to yield smooth behavior.

Value of commitment. Another implication of no conflict is that agents place no value on commitment devices. This is by no means obvious. On the one hand, since economies without conflict are well-behaved, one might conclude that no intervention is necessary. On the other hand, although conflict is absent, disagreement between present and future selves still persists. The former may therefore have incentives to manipulate the latter. This also contrasts with the common wisdom that the inherent time inconsistency leads to a willingness to pay for assets serving as commitment devices, such as illiquid assets (mortgages, 401k plans, and IRAs) that lock resources away from future selves.

To formalize this insight, we appeal to the *optimal delegation* approach: a principal delegates decisions to an informed but biased agent (see Amador et al., 2006; Amador and Bagwell, 2013). In our setting, the agent in control is privately informed about her type θ, which is unobservable to or unverifiable by the principal. Otherwise, first-best allocations that fully correct for the bias would be attainable. Hence, the proposed allocations must be incentive-compatible. Moreover, since contingent transfers are unavailable to provide incentives, the principal instead restricts the agent's action set. This leads to a trade-off between *flexibility and commitment*: whether to grant the agent the flexibility to respond to shocks or to impose commitment by limiting the action set.

We implement this approach by considering an "Overlapping Principals" problem, where each agent plays the dual role of agent and principal. The agent in control must abide by the rules set by the previous principal and, upon losing control, designs a mechanism that binds the next agent. As in the Markov equilibrium, the principal can only choose mechanisms, $A = \{c(\theta, b), b'(\theta, b)\}_{\theta \in \Theta, b \geq \underline{b}}$, that depend on the payoff-relevant state $\{\theta, b\}$. That is, the exiting agent, leaving asset b to the next agent and acting as the principal, chooses a static incentive-compatible mechanism A to maximize the ex-ante expected continuation value $\beta \int_{\underline{\theta}}^{\bar{\theta}} v(\theta, b; A)\, dF(\theta)$. Further details on the setup are provided in Appendix A.7.

As shown by Amador et al. (2006), the restriction to static mechanisms is without loss of generality when $\lambda = 1$ so that the taste type is *i.i.d.* over time. However, when $\lambda < 1$ and tastes are persistent, introducing history dependence could improve outcomes (see Halac and Yared, 2014). To implement such contract, one would need to resort to equilibrium concepts weaker than Markov. For this reason, we see our welfare criterion as choosing the best Markov Equilibrium.

The equivalence transformation is also instrumental in analyzing the solution to the delegation problem. Specifically, the relationship established in Lemma 1 for equilibrium allocations extends to any incentive-compatible allocation, as we formalize in Lemma A2 in the appendix. This is because the envelope condition used to establish this relationship also holds under incentive-compatible mechanisms. Consequently, we obtain the same mapping between the continuation values as given in equation (6):

$$\beta \int_{\underline{\theta}}^{\bar{\theta}} v(\theta, b; A)\, dF(\theta) = \int_{\underline{\theta}}^{\bar{\theta}} w(\theta, b; A)\, dG(\theta). \tag{10}$$

When there is no conflict, all weights $dG(\theta)$ are positive. Any optimizing future action makes a positive contribution to the continuation value. In this case, the agent does not wish to put any restriction on her future choice set. Hence, flexibility dominates commitment: full flexibility and no commitment are optimal. We formalize these insights in the following proposition.

Proposition 5 (Efficiency: no conflict)**.** *If Condition 1 is satisfied, the equilibrium is constrained efficient. In other words, the agents do not value commitment devices, nor would they want to blind themselves to any information about their tastes.*

Figure 4: Commitment vs. flexibility

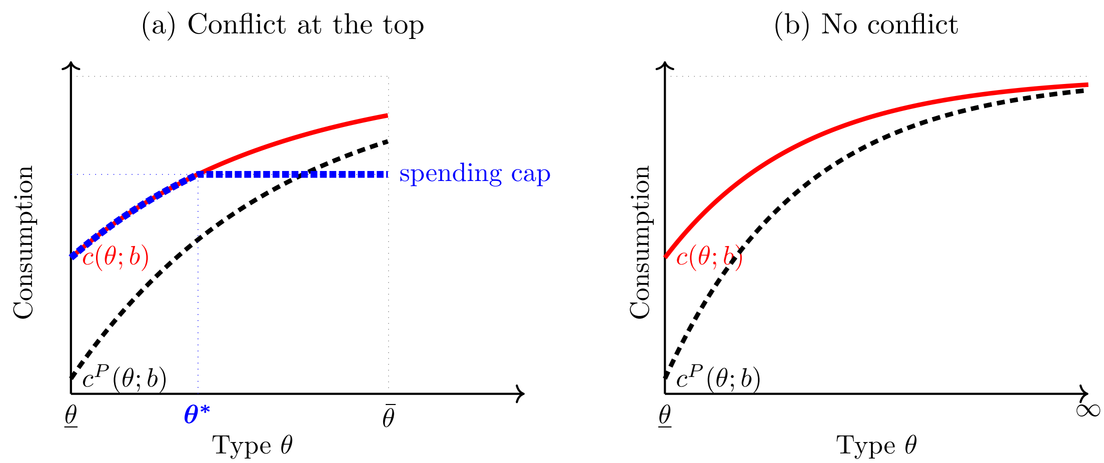

Notes: In each panel, the red solid line represents the equilibrium consumption $c(\theta, b)$ absent any rules, while the black line depicts the first-best choices $c^P(\theta, b)$. In panel (a), the blue densely dotted line represents the optimal incentive-compatible allocation: it features a threshold type θ^* above which all types are bunched.

Proof. See Appendix A.7. □

A common approach to identifying present bias and sophistication is to look at the willingness to pay for commitment devices, such as illiquid assets (Kocherlakota, 2001). The no-commitment result in Proposition 5 suggests that the use of commitment devices may not provide as clear an identification. It also suggests that even if we do not find clear evidence of willingness to pay for commitment or the use of commitment devices, it does not necessarily invalidate myopia, nor does it necessarily imply naïveté over sophistication.

To prevent conflict from arising, Condition 1 requires that the density does not decrease too fast relative to the degree of present bias and that the taste type is unbounded. When either condition fails, conflict will be present. For example, if the density does not decrease too fast but the taste type is bounded, as in Example 2, the adjusted weight becomes negative only at the upper bound. This "conflict at the top" scenario relates our results to the minimum savings rules prescribed by Amador et al. (2006) and the conditions under which such rules are optimal. The transformation in (10) helps to understand why. Given that conflict is present, commitment holds value relative to full flexibility. Specifically, because the conflict arises at the very top and a negative weight is assigned there, any action that improves the value obtained by this type makes a negative contribution to the expected continuation value. It is therefore beneficial to restrict the choices of this type, either by capping their spending or, equivalently, by imposing a minimum savings requirement, as illustrated in panel (a) of Figure 4.

Value of information. To understand the value of information, consider that, instead of constraining her future choices, the agent now can blind certain information from her future selves. When there is no conflict, the agent does not value commitment and prefers her future self to respond optimally to the information available about future states. It follows that the agent values information, as she would not choose to withhold any of it. This implies that full information is optimal.

Another immediate implication is that, from the perspective of the equivalent optimistically biased agent, she would not wish to alter her future belief, which in this case is a well-defined expectation with positive weights everywhere. In other words, it is optimal for the agent to stay optimistic and not learn. This result aligns with situations in which strategic ignorance and self-confidence are not desirable, as discussed by Carrillo and Mariotti (2000) and Bénabou and Tirole (2002).[9]

5 Continuous-time approach

In this section, we analyze the continuous-time limit of the discrete-time model presented in Section 2. The continuous-time approach often aligns closely with its discrete-time counterpart, with no substantive differences in many instances. In this case, we show that while the transformation and the conditions for existence and uniqueness are analogous to those in discrete time, we can also show that decision functions are uniquely determined at each debt level. Moreover, continuous time allows for the possibility of applying alternative numerical methodologies that often work better in applied settings.

The environment is identical to the one laid out in Section 2, with the following two mappings of parameters. First, the discount rate is redefined as $\rho = -\log(\delta) \in (0, \infty)$. Second, the transition of the present self to a future self occurs at a Poisson arrival rate $\tilde{\lambda} = -\log(1-\lambda) \in (0, \infty)$.[10] For simplicity, the maintained assumptions, Assumptions 1 and 2, remain in force here.

5.1 Transformation

Markov equilibrium in continuous time. The agent's value function when in control, $w(\theta, b)$, together with the companion value function $v(\theta, b)$, satisfy the following Hamilton-

[9] While our setup and information structure differ from those in Bénabou and Tirole (2002), Condition 1(i) coincides with one of the two cases in Proposition 4 of Bénabou and Tirole (2002).

[10] As $\tilde{\lambda} \to \infty$, with $\underline{\theta} \to 1$ and $\bar{\theta} \to 1$, the model converges to the instantaneous gratification limit by Harris and Laibson (2013).

Jacobi-Bellman (HJB) equations: $\forall \theta \in \Theta, b \in [\underline{b}, \bar{b}]$,

$$\rho w(\theta, b) = \sup_{c \geq 0} \Big\{ \theta u(c) + (y + rb - c) w_b(\theta, b) + \tilde{\lambda} \left(\beta \mathbb{E}[v(\theta', b)] - w(\theta, b) \right) \Big\} \tag{11}$$

$$\rho v(\theta, b) = \theta u(c(\theta, b)) + (y + rb - c(\theta, b)) v_b(\theta, b) + \tilde{\lambda} \left(\mathbb{E}[v(\theta', b)] - v(\theta, b) \right). \tag{12}$$

Parallel to problem (P), the current agent's decision function is denoted by $c(\theta, b)$, and she evaluates the continuation value given future agents' decisions.

The following proposition presents the continuous-time analogue of the equivalence result in Proposition 1.

Proposition 6 (Equivalence: continuous time)**.** *The value function $w(\theta, b)$ and decision function $c(\theta, b)$ are part of a Markov Equilibrium if and only if they solve the recursive problem: $\forall \theta \in \Theta, b \in [\underline{b}, \bar{b}]$,*

$$\rho w(\theta, b) = \sup_{c \geq 0} \left\{ \theta u(c) + (y + rb - c) w_b(\theta, b) + \tilde{\lambda} \left(\int_{\underline{\theta}}^{\bar{\theta}} w(\theta', b) dG(\theta') - w(\theta, b) \right) \right\}. \tag{13}$$

Proof. See Appendix A.8. □

Given the similarities between the continuous- and discrete-time versions of the equilibrium, it is clear that the economics behind it is the same as previously discussed. One may wonder, though, if the continuous-time setting is helpful in providing a sharper characterization. To answer this question, we study the viscosity solution to the HJB equation (13), which is a partial differential equation. These types of solution coincide almost everywhere with the classical solution and are instrumental in handling potential kinks. For this reason, they have become widely used in settings with idiosyncratic income risk. In addition, when they exist, they are mostly unique. Last but not least, the first-order necessary condition in equation (13) generates:

$$\theta u'(c(\theta, b)) = w_b(\theta, b), \quad a.e.$$

Since the marginal utility $u'(\cdot)$ is continuous, the properties of w directly translate into those of $c(\theta, b)$. To begin with, it is immediate that there is a unique solution for each combination of states, something that we were not able to establish in the discrete-time setting whenever conflict arises $G^- > 0$. In addition, the continuity of w_b is passed on directly to $c(\theta, b)$.

5.2 Characterization

Before proceeding, we introduce the following definition.

Definition 2 (Viscosity solution)**.** For any set $\Omega \in \mathbb{R}^2$, let $C(\Omega)$ be the set of continuous functions and $C^1(\Omega)$ the set of continuously differentiable functions. Then,

(i) A function $w \in C(\Omega)$ is a viscosity subsolution if for any $\phi \in C^1(\Omega)$ and any $(\theta, b) \in \Omega$ where $(w - \phi)$ attains a local maximum, then:

$$(\rho + \tilde{\lambda})w(\theta, b) \leq \max_{c \geq 0} \left\{ \theta u(c) + (y + rb - c)\phi_b(\theta, b) + \tilde{\lambda} \int w(\theta', b) dG(\theta') \right\}.$$

(ii) A function $w \in C(\Omega)$ is a viscosity supersolution if for any $\phi \in C^1(\Omega)$ and any $(\theta, b) \in \Omega$ where $(w - \phi)$ attains a local minimum, then:

$$(\rho + \tilde{\lambda})w(\theta, b) \geq \max_{c \geq 0} \left\{ \theta u(c) + (y + rb - c)\phi_b(\theta, b) + \tilde{\lambda} \int w(\theta', b) dG(\theta') \right\}.$$

In the standard definition, w is a **viscosity solution**, if it is both a viscosity subsolution and a viscosity supersolution.[11] In general to single out the appropriate solution to a problem a border or boundary condition is necessary. However, in problems with state constraints, e.g., a borrowing limit, the border values can be difficult to identify. Fortunately, in problems with such constraints, a slight modification to the definition is enough to ensure that one is pinning down the right solution and also to obtain uniqueness. See Soner (1986) and Capuzzo-Dolcetta and Lions (1990) for the earlier treatments of viscosity solutions with state constraints.

To be precise, define $\bar{\Omega} = [\underline{\theta}, \bar{\theta}] \times [\underline{b}, \bar{b}]$, and let Ω be the interior of $\bar{\Omega}$. Since our problem has the state constraints $b \in [\underline{b}, \bar{b}]$, $w \in C(\bar{\Omega})$ is a **Constrained Viscosity Solution** to problem (13) if it is a viscosity subsolution on $\bar{\Omega}$ and a viscosity supersolution on Ω (only in the interior).

We also introduce a condition to limit the extent of conflicts.

Condition 3 (Limited conflict: continuous time)**.** *The cumulative negative weights are limited:*

$$G^- < \frac{\rho}{2\tilde{\lambda}}. \tag{14}$$

Condition 3 is analogous to Condition 2 for generating uniqueness in discrete time. In both cases, we face the same problem and departure from the standard one: lack of monotonicity of the operator. Monotonicity is a key element in Blackwell's sufficient conditions, and it is also extremely useful for establishing uniqueness of viscosity solutions, by appealing

[11]See Fleming and Soner (2006), Chapter 2, for an extensive discussion and alternative (equivalent) definitions.

to the Comparison Theorem, as it implies an ordering of solutions. Quasi-hyperbolic discounting renders most operators non-monotone. The total variation, $G^+ + G^-$, determines how much monotonicity is lost. As long as the agent is not too forward looking relative to the extent of conflict, that is ρ is large relative to G^-, the sophisticated strategic behavior does not generate major complications.

Analogously to the discrete-time setting, where we lost Blackwell's sufficient conditions, in the continuous-time environment we lose the comparison principle due to the potential negative weights. As a result, we provide conditions for uniqueness relying on a modified version of the comparison principle in the following proposition.

Proposition 7 (Uniqueness: continuous time)**.** *If Condition 3 is satisfied, there exists at most one Constrained Viscosity Solution to problem* (13).

Proof. See Appendix A.9. □

We have purposely avoided many technical complications by restricting the analysis to bounded assets and consumption. This is without loss of generality. We can easily extend the proofs to the case in which both assets and the utility function are unbounded. This would require bounding the speed at which the utility can grow.[12] These additional technicalities are common in the standard problem. By incorporating them in our analysis we would have added clutter with little gains in economic insights.

6 Conclusion

In this paper, we offer a new method to study Markov equilibria with quasi-hyperbolic discounting. These settings are of great interest in the consumption-savings literature, where hyperbolic preferences are considered a leading candidate for rationalizing many puzzling behaviors. Beyond that, such settings also arise naturally in dynamic games with heterogeneous decision makers and social interactions. Nevertheless, whether behavioral or microfounded, these environments are technically challenging and have generated a large literature approaching the problem from different directions.

Our approach embeds several interesting features. First, it reduces the problem to a standard recursive maximization problem. This opens the door to using a wide range of existing tools that were previously not applicable in such settings. Second, it does not require specific knowledge of either discrete- or continuous-time settings. It works equally well in both cases. The applied economist can directly use their favorite toolbox. Finally,

[12]See, for instance, Theorem 9.1 in Fleming and Soner (2006), where a proof with bounded and unbounded sets is provided.

the transformed problem is informative about the efficiency and potential corrective policies that could be implemented. Usually, hyperbolic discounting is viewed as synonymous with corrective policies, either by governments or through self-generated commitment devices. We show that this is not necessarily the case.

Our method can have broader applicability. Our proofs rely on readily available tools with small modifications. The conditions we provide for existence and uniqueness are sufficient, but by no means necessary. It would be interesting to explore, focusing on the implications of signed measures for economics, whether these conditions can be related. Future work can also extend the method to similar dynamic games beyond the quasi-hyperbolic setting here.

References

Alesina, A. and G. Tabellini (1990): "A Positive Theory of Fiscal Deficits and Government Debt," *Review of Economic Studies*, 57, 403–414.

Alvarez, F. and N. L. Stokey (1998): "Dynamic Programming with Homogeneous Functions," *Journal of Economic Theory*, 82, 167–189.

Amador, M. and K. Bagwell (2013): "The Theory of Optimal Delegation With an Application to Tariff Caps," *Econometrica*, 81, 1541–1599.

Amador, M., I. Werning, and G.-M. Angeletos (2006): "Commitment vs. Flexibility," *Econometrica*, 74, 365–396.

Azzimonti, M., M. Battaglini, and S. Coate (2016): "The Costs and Benefits of Balanced Budget Rules: Lessons from a Political Economy Model of Fiscal Policy," *Journal of Public Economics*, 136, 45–61.

Barro, R. J. (1979): "On the Determination of the Public Debt," *Journal of Political Economy*, 87, 940–971.

Bassetto, M., Z. Huo, and J.-V. Ríos-Rull (2024): "Institution Building without Commitment," *American Economic Review*, 114, 3427–68.

Battaglini, M. and S. Coate (2008): "A Dynamic Theory of Public Spending, Taxation, and Debt," *American Economic Review*, 98, 201–36.

Bénabou, R. and J. Tirole (2002): "Self-Confidence and Personal Motivation," *The Quarterly Journal of Economics*, 117, 871–915.

Benhabib, J., A. Bisin, and A. Schotter (2010): “Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs,” *Games and Economic Behavior*, 69, 205–223.

Cao, D. and I. Werning (2016): “Dynamic Savings Choices with Disagreements,” Working Paper 22007, National Bureau of Economic Research.

——— (2018): “Saving and Dissaving With Hyperbolic Discounting,” *Econometrica*, 86, 805–857.

Capuzzo-Dolcetta, I. and P.-L. Lions (1990): “Hamilton-Jacobi Equations with State Constraints,” *Transactions of the American Mathematical Society*, 318, 643–683.

Carrillo, J. D. and T. Mariotti (2000): “Strategic Ignorance as a Self-Disciplining Device,” *The Review of Economic Studies*, 67, 529–544.

Chatterjee, S. and B. Eyigungor (2016): “Continuous Markov Equilibria with Quasi-Geometric Discounting,” *Journal of Economic Theory*, 163, 467 – 494.

Choukhmane, T. (forthcoming): “Default Options and Retirement Saving Dynamics,” Tech. rep., American Economic Review.

Choukhmane, T. and C. Palmer (2024): “How Do Consumers Finance Increased Retirement Savings,” Working paper, MIT Sloan.

Cunha, A. B. and E. Ornelas (2017): “The Limits of Political Compromise: Debt Ceilings and Political Turnover,” *Journal of the European Economic Association*, 16, 781–824.

Epstein, L. G. (2006): “An Axiomatic Model of Non-Bayesian Updating,” *The Review of Economic Studies*, 73, 413–436.

Epstein, L. G., J. Noor, and A. Sandroni (2008): “Non-Bayesian Updating: A Theoretical Framework,” *Theoretical Economics*, 3, 193–229.

Felli, C., F. Piguillem, and L. Shi (2025): “Fiscal Rules and Discretion with Risk of Default,” CEPR Discussion Papers 16211, C.E.P.R. Discussion Papers.

Fleming, W. H. and H. M. Soner (2006): *Controlled Markov Processes and Viscosity Solutions*, vol. 25 of *Stochastic Modelling and Applied Probability*, New York: Springer, 2 ed.

Gul, F. and W. Pesendorfer (2001): “Temptation and Self-Control,” *Econometrica*, 69, 1403–1435.

——— (2004): "Self-Control and the Theory of Consumption," *Econometrica*, 72, 119–158.

Halac, M. and P. Yared (2014): "Fiscal Rules and Discretion Under Persistent Shocks," *Econometrica*, 82, 1557–1614.

——— (2018): "Fiscal Rules and Discretion in a World Economy," *American Economic Review*, 108, 2305–34.

——— (2022): "Fiscal Rules and Discretion Under Limited Enforcement," *Econometrica*, 90, 2093–2127.

Hansen, L. P. and T. J. Sargent (2001): "Robust Control and Model Uncertainty," *American Economic Review*, 91, 60–66.

——— (2008): *Robustness*, Princeton University Press.

Harris, C. and D. Laibson (2001): "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 69, 935–957.

——— (2003): *Hyperbolic Discounting and Consumption*, Cambridge University Press, 258–297, Econometric Society Monographs.

——— (2013): "Instantaneous Gratification," *The Quarterly Journal of Economics*, 128, 205–248.

Kocherlakota, N. (2001): "Looking for Evidence of Time-Inconsistent Preferences in Asset Market Data," *Federal Reserve Bank of Minneapolis Quarterly Review*, 25, 13–24.

Krusell, P. and A. J. Smith (2003): "Consumption-Savings Decisions with Quasi-Geometric Discounting," *Econometrica*, 71, 365–375.

Laibson, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–477.

Loewenstein, G., T. O'Donoghue, and M. Rabin (2003): "Projection Bias in Predicting Future Utility," *The Quarterly Journal of Economics*, 118, 1209–1248.

Maxted, P. (2025): "Present Bias Unconstrained: Consumption, Welfare, and the Present-Bias Dilemma," *The Quarterly Journal of Economics*, qjaf030.

Maxted, P., D. Laibson, and B. Moll (2024): "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy," *The Quarterly Journal of Economics*, 140, 691–743.

Milgrom, P. and I. Segal (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70, 583–601.

Müller, A., K. Storesletten, and F. Zilibotti (2016): "The Political Color of Fiscal Responsibility," *Journal of the European Economic Association*, 14, 252–302.

Persson, T. and L. E. O. Svensson (1989): "Why a Stubborn Conservative would Run a Deficit: Policy with Time-Inconsistent Preferences," *The Quarterly Journal of Economics*, 104, 325–345.

Pettersson-Lidbom, P. (2001): "An Empirical Investigation of the Strategic Use of Debt," *Journal of Political Economy*, 109, 570–583.

Phelps, E. and R. Pollak (1968): "On Second-Best National Saving and Game-Equilibrium Growth," *The Review of Economic Studies*, 35, 185–199.

Piguillem, F. and A. Riboni (2020): "Fiscal Rules as Bargaining Chips," *The Review of Economic Studies*.

Rincón-Zapatero, J. P. and C. Rodríguez-Palmero (2003): "Existence and Uniqueness of Solutions to the Bellman Equation in the Unbounded Case," *Econometrica*, 71, 1519–1555.

Roubini, N. and J. D. Sachs (1989): "Political and Economic Determinants of Budget Deficits in the Industrial Democracies," *European Economic Review*, 33, 903–933.

Sinander, L. (2022): "The Converse Envelope Theorem," *Econometrica*, 90, 2795–2819.

Soner, H. M. (1986): "Optimal Control with State-Space Constraint I," *SIAM Journal on Control and Optimization*, 24, 552–561.

Stokey, N. L., R. E. Lucas, and E. C. Prescott (1989): *Recursive Methods in Economic Dynamics*, Harvard University Press.

Strotz, R. H. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165–180.

Takeuchi, K. (2011): "Non-Parametric Test of Time Consistency: Present Bias and Future bias," *Games and Economic Behavior*, 71, 456–478.

A Proofs

A.1 Proof of Lemma 1

We introduce and work with the sequential representation of (P). Consider an agent that comes into control with an initial asset $b_0 \in [\underline{b}, \bar{b}]$. Call the sequence $\underset{\sim}{b} = \{b_t\}_{t=0}^{\infty}$ a plan. Let

$$\Pi(b_0) = \left\{ \{b_t\}_{t=0}^{\infty} : b_t \leq y + (1+r)b_{t-1} \text{ and } \underline{b} \leq b_t \leq \bar{b}, \forall t = 1, 2, \cdots \right\}$$

be the set of *feasible* plans from b_0. Let $U(b_t, b_{t+1}) \equiv u\left(y + (1+r)b_t - b_{t+1}\right)$. The agent's problem can then be written in sequential form as

$$w\left(\theta, b_0\right) = \sup_{\underset{\sim}{b} \in \Pi(b_0)} \left\{ \sum_{t=0}^{\infty} \left(\delta(1-\lambda)\right)^t \left(\theta U(b_t, b_{t+1}) + \beta\delta\lambda \mathbb{E}\left[v\left(\theta', b_{t+1}\right)\right]\right) \right\}.$$

Or, in a more compact fashion:

$$w(\theta, b_0) = \sup_{\underset{\sim}{b} \in \Pi(b_0)} \left\{ \theta\psi_0(\underset{\sim}{b}) + \psi_1(\underset{\sim}{b}) \right\},$$

where $\psi_0(\underset{\sim}{b}) \equiv \sum_{t=0}^{\infty} \left(\delta(1-\lambda)\right)^t U(b_t, b_{t+1})$ and $\psi_1(\underset{\sim}{b}) \equiv \sum_{t=0}^{\infty} \left(\delta(1-\lambda)\right)^t \beta\delta\lambda\mathbb{E}\left[v\left(\theta', b_{t+1}\right)\right]$.[13]

This problem is akin to a standard optimization problem with payoff $R(\theta, \underset{\sim}{b}) = \theta\psi_0(\underset{\sim}{b}) + \psi_1(\underset{\sim}{b})$. The partial $\partial R(\theta, \underset{\sim}{b})/\partial\theta = \psi_0(\underset{\sim}{b})$ always exists. Since $u(\cdot)$ is continuous and, by Assumption 1, evaluated in a compact set, it is bounded. Thus, for $\delta < 1$, $\psi_0(\underset{\sim}{b})$ is bounded too. Hence, the problem satisfies the *basic assumptions* in Sinander (2022): the Envelope Theorem and its converse hold. Thus, the value functions are absolutely continuous and yield

$$w_\theta\left(\theta, b_0\right) = \psi_0(\underset{\sim}{b}^*) = \sum_{t=0}^{\infty} \left(\delta(1-\lambda)\right)^t U(b_t^*, b_{t+1}^*) = \sum_{t=0}^{\infty} (\delta(1-\lambda))^t u\left(c_t^*\right), \tag{15}$$

where the superscript $*$ indicates that the variables are evaluated at the optimal choices.

In any equilibrium, the following equations also hold:

$$w\left(\theta, b_0\right) = \sum_{t=0}^{\infty} (\delta(1-\lambda))^t \left(\theta u\left(c_t^*\right) + \beta\delta\lambda\mathbb{E}\left[v\left(\theta', b_{t+1}^*\right)\right]\right) \tag{16}$$

$$v\left(\theta, b_0\right) = \sum_{t=0}^{\infty} (\delta(1-\lambda))^t \left(\theta u\left(c_t^*\right) + \delta\lambda\mathbb{E}\left[v\left(\theta', b_{t+1}^*\right)\right]\right). \tag{17}$$

[13] The Principle of Optimality delivers the equivalence between the sequential and recursive problems (see Stokey et al., 1989, Chapter 4). The required conditions are satisfied: (i) $\Pi(b_0)$ is nonempty, for all $b_0 \in [\underline{b}, \bar{b}]$; (ii) $\lim_{T\to\infty} \sum_{t=0}^{T} \left(\delta(1-\lambda)\right)^t \left(\theta U(b_t, b_{t+1}) + \beta\delta\lambda\mathbb{E}\left[v\left(\theta', b_{t+1}\right)\right]\right)$ exists, for all $b_0 \in [\underline{b}, \bar{b}]$ and $\underset{\sim}{b} \in \Pi(b_0)$.

Combining equations (16)-(17), we obtain that the equilibrium value functions satisfy

$$w\left(\theta,b_0\right)-\beta v\left(\theta,b_0\right)=(1-\beta)\theta\sum_{t=0}^{\infty}(\delta(1-\lambda))^t u\left(c_t^*\right).$$

Since the summation in the last expression above is equal to $w_\theta\left(\theta,b_0\right)$, according to equation (15), we obtain the relation in (4). ☐

A.2 Proof of Proposition 1

Only if. Suppose first a Markov Equilibrium $\{v(\theta,b),w(\theta,b),c(\theta,b),b'(\theta,b)\}$ exists. By Lemma 1 we can use the relation in equation (4) to write:

$$\beta\mathbb{E}\left[v\left(\theta,b\right)\right]=\int_{\underline{\theta}}^{\bar{\theta}}\left[w\left(\theta,b\right)-(1-\beta)\,\theta w_\theta\left(\theta,b\right)\right]dF\left(\theta\right).$$

Since $F(\theta)$ is differentiable and $w(\theta,b)$ absolutely continuous, the integral above exists. Integrating by parts:

$$\begin{aligned}\int_{\underline{\theta}}^{\bar{\theta}}\theta w_\theta\left(\theta,b\right)dF\left(\theta\right)&=\int_{\underline{\theta}}^{\bar{\theta}}\theta f\left(\theta\right)dw\left(\theta,b\right)\\&=\bar{\theta}f\left(\bar{\theta}\right)w\left(\bar{\theta},b\right)-\underline{\theta}f\left(\underline{\theta}\right)w\left(\underline{\theta},b\right)-\int_{\underline{\theta}}^{\bar{\theta}}w\left(\theta,b\right)d\left[\theta f\left(\theta\right)\right].\end{aligned}$$

Using the last equation in the first:

$$\begin{aligned}\beta\mathbb{E}\left[v\left(\theta,b\right)\right]&=\int_{\underline{\theta}}^{\bar{\theta}}w\left(\theta,b\right)d\left[F\left(\theta\right)+(1-\beta)\theta f\left(\theta\right)\right]+(1-\beta)\left[\underline{\theta}f\left(\underline{\theta}\right)w\left(\underline{\theta},b\right)-\bar{\theta}f\left(\bar{\theta}\right)w\left(\bar{\theta},b\right)\right]\\&=\int_{\underline{\theta}}^{\bar{\theta}}w\left(\theta,b\right)dG\left(\theta\right).\end{aligned}\tag{18}$$

In the last step we have collected all the terms multiplying $w(\theta,b)$ into $dG(\theta)$. It is straightforward then that problems (P) and ($\hat{P}$) share the same objective function. Since the constraint sets are identical and $c(\theta,b)$ and $b'(\theta,b)$ are a solution to problem (P), it must that $c(\theta,b)$ and $b'(\theta,b)$ are also a solution to problem ($\hat{P}$).

If. Now suppose $\{w(\theta,b),c(\theta,b),b'(\theta,b)\}$ is a solution to problem ($\hat{P}$). In the solution, the value function $w(\theta,b)$ must be differentiable almost everywhere with respect to θ, since the assumptions of Theorem 2 in Milgrom and Segal (2002) are satisfied. We define $v(\theta,b)$ according to equation (4), which together with the solution constitutes an equilibrium. ☐

A.3 Proof of Corollary 1

Since $F(\theta)$ is continuously differentiable, $G(\theta)$ is also continuously differentiable for $\theta \in (\underline{\theta}, \bar{\theta})$. However, there are potentially mass points at the ends of the distribution. In particular, the weight assigned to $\bar{\theta}$ is $-(1-\beta)\bar{\theta} f(\bar{\theta})$, which is strictly negative if $f(\bar{\theta}) > 0$. Taking into account the potential accumulation points in the borders, the weighted value can be computed as:

$$\int_{\underline{\theta}}^{\bar{\theta}} \theta dG(\theta) = \underline{\theta} \times (1-\beta)\underline{\theta} f(\underline{\theta}) + \int_{\underline{\theta}_+}^{\bar{\theta}_-} \theta dG(\theta) - \bar{\theta} \times (1-\beta)\bar{\theta} f(\bar{\theta}).$$

Using integration by parts:

$$\begin{aligned}\int_{\underline{\theta}_+}^{\bar{\theta}_-} \theta dG(\theta) &= 1 + (1-\beta)\int_{\underline{\theta}_+}^{\bar{\theta}_-} \theta d[\theta f(\theta)] \\ &= 1 + (1-\beta)\left[\theta^2 f(\theta)\,|_{\underline{\theta}_+}^{\bar{\theta}_-} - \int_{\underline{\theta}_+}^{\bar{\theta}_-} \theta f(\theta)\, d\theta\right] = \beta + \theta \times (1-\beta)\,\theta f(\theta)\,|_{\underline{\theta}_+}^{\bar{\theta}_-}.\end{aligned}$$

Replacing this result in the first equation makes $\int_{\underline{\theta}}^{\bar{\theta}} \theta dG(\theta)$ equal to β.

A.4 Some stepping-stone results

We write down the Bellman operator for problem ($\hat{P}$). To keep the notation compact, let

$$\Gamma(b) = \{b' : b' \leq y + (1+r)b \text{ and } b' \in [\underline{b}, \bar{b}]\}$$

and define the operator T as:

$$(Tw)(\theta, b) = \sup_{b' \in \Gamma(b)} \left\{\theta U(b, b') + \delta\left[(1-\lambda)w(\theta, b') + \lambda \int_{\underline{\theta}}^{\bar{\theta}} w(\theta', b') dG(\theta')\right]\right\}.$$

Let X be a bounded subset of R^2, and let $C(X)$ be the space of bounded continuous functions on X. Then we have,

Lemma A1 (Self-mapping)**.** *The operator T maps $C(X)$ into itself.*

Proof. We first show that the integral involving $G(\cdot)$ does not affect the basic properties of functions in $C(X)$. Define the operator Mw as:

$$Mw(b) = \int_{\Theta} w(\theta, b) dG(\theta)$$

We show that Mw is bounded if w is bounded. Suppose w is bounded by some constant $\bar{w} < \infty$. To show that we can find a constant $B < \infty$ such that $\|Mw\| \leq B$ for all $b \in [\underline{b}, \bar{b}]$:

$$\begin{aligned}
\|Mw\| &= \left\| \int_{\Theta} w(\theta, b) dG(\theta) \right\| \\
&= \left\| w(\underline{\theta}, b)(1-\beta)\underline{\theta} f\left(\underline{\theta}\right) - w(\bar{\theta}, b)(1-\beta)\bar{\theta} f\left(\bar{\theta}\right) + \int_{\underline{\theta}_+}^{\bar{\theta}_-} w(\theta, b) dG\left(\theta\right) \right\| \\
&\leq \left\| w(\underline{\theta}, b)(1-\beta)\underline{\theta} f\left(\underline{\theta}\right) \right\| + \left\| w(\bar{\theta}, b)(1-\beta)\bar{\theta} f\left(\bar{\theta}\right) \right\| + \left\| \int_{\underline{\theta}_+}^{\bar{\theta}_-} w(\theta, b) dG\left(\theta\right) \right\| \\
&\leq (1-\beta)\left[\underline{\theta} f\left(\underline{\theta}\right) + \bar{\theta} f\left(\bar{\theta}\right)\right] \bar{w} + \left\| \int_{\underline{\theta}_+}^{\bar{\theta}_-} w(\theta, b)\left[(2-\beta) f\left(\theta\right) + (1-\beta)\theta f'\left(\theta\right)\right] d\theta \right\| \\
&\leq (1-\beta)\left[\underline{\theta} f\left(\underline{\theta}\right) + \bar{\theta} f\left(\bar{\theta}\right)\right] \bar{w} + \int_{\underline{\theta}_+}^{\bar{\theta}_-} \|w(\theta, b)\| \left[(2-\beta) f\left(\theta\right) + (1-\beta)\theta \left\|f'\left(\theta\right)\right\|\right] d\theta \\
&\leq (1-\beta)\left[\underline{\theta} f\left(\underline{\theta}\right) + \bar{\theta} f\left(\bar{\theta}\right)\right] \bar{w} + \left[2-\beta+(1-\beta)\bar{f}\right] \bar{w} \\
&= \left[2-\beta+(1-\beta)\left(\underline{\theta} f\left(\underline{\theta}\right) + \bar{\theta} f\left(\bar{\theta}\right) + \bar{f}\right)\right] \bar{w}.
\end{aligned}$$

The first line follows from the definition of the operator. The second line uses the expression for $dG(\theta)$ in (7); the third by the distance properties; and the fourth due to the bound of w. The fifth and sixth lines are due to the bounds of w and f' and true density integration to one. Hence, $B = \left[2-\beta+(1-\beta)\left(\underline{\theta} f\left(\underline{\theta}\right) + \bar{\theta} f\left(\bar{\theta}\right) + \bar{f}\right)\right] \bar{w}$ is a bound for Mw.

For the continuity of the operator, choose a convergent sequence $b_n \to b$, let $d_n(\theta) = w(\theta, b) - w(\theta, b_n)$ and recall that w is bounded by W. Then

$$\begin{aligned}
|Mw(b) - Mw(b_n)| &= \left| \int_{\Theta} d_n(\theta) dG(\theta) \right| \\
&= \left| (1-\beta)\left[d_n(\underline{\theta})\underline{\theta} f\left(\underline{\theta}\right) - d_n(\bar{\theta})\bar{\theta} f\left(\bar{\theta}\right)\right] + \int_{\underline{\theta}_+}^{\bar{\theta}_-} d_n(\theta) dG\left(\theta\right) \right| \\
&\leq (1-\beta)\left[\left|d_n(\underline{\theta})\right| \underline{\theta} f\left(\underline{\theta}\right) + \left|d_n(\bar{\theta})\right| \bar{\theta} f\left(\bar{\theta}\right)\right] + \left| \int_{\underline{\theta}_+}^{\bar{\theta}_-} d_n(\theta) dG\left(\theta\right) \right|.
\end{aligned}$$

Because w is continuous, both $|d(\underline{\theta})|$ and $|d(\underline{\theta})|$ vanish as $n \to \infty$. Similarly each term inside the integral converges pointwise to the zero function. Since each term in the sequence is bounded by the constant $2W$, the Lebesgue Dominated Convergence Theorem implies:

$$\lim_{n\to\infty} \int_{\underline{\theta}_+}^{\bar{\theta}_-} d_n(\theta) dG\left(\theta\right) = \int_{\underline{\theta}_+}^{\bar{\theta}_-} \lim_{n\to\infty} d_n(\theta) dG\left(\theta\right) = 0$$

Thus, the integral term also vanishes. Hence, Mw is continuous.

Now consider the operator T. Note that the constraint set $\Gamma(b)$ is a continuous compact value correspondence for each b. Hence, for any $w \in C(X)$, by the Theorem of Maximum a solution exists, so we can replace the sup with a max. Furthermore, $(Tw)(\theta, b)$ is continuous in both θ and b. For further reference, bear in mind that this also implies that the operator T is continuous. □

Lemma A1 states that even though our weighting function may not constitute a probability measure, it still retains the essential properties to apply fixed point theorems and potentially also the Contraction Mapping Theorem. In particular, it preserves continuity and boundedness. The fact that the density is bounded and differentiable in Assumption 2 ensures that nothing pathological happens.

Before proceeding, it is useful to observe that the range of admissible solutions is limited by the resources. Thus, the set in which the value function lies can be narrowed. However, determining this set would depend on the cardinality of $u(\cdot)$. Without loss of generality, suppose that $u : [\underline{b}, \bar{b}] \to \mathbb{R}_+$. Define $\underline{u} = u(y + r\underline{b})$ and $\bar{u} = u(y + (1 + r)\bar{b})$. Then it is straightforward to show that

$$\underline{w} = \underline{\theta}\frac{\underline{u}}{1 - \beta\delta} \leq w(\theta, b) \leq \bar{\theta}\frac{\bar{u}}{1 - \delta} = \bar{w}, \quad \forall \theta, b.$$

Let $W \subset C(X)$ be the space of continuous functions $w : [\underline{\theta}, \bar{\theta}] \times [\underline{b}, \bar{b}] \to [\underline{w}, \bar{w}]$.

A.5 Proof of Proposition 2

To prove existence, we first introduce the definition of equicontinuity (see Stokey et al., 1989, page 520) and the Schauder Fixed-Point Theorem adapted to bounded and continuous functions (see Stokey et al., 1989, Theorem 17.4).

Definition A1 (Equicontinuity)**.** A subset W of $C(X)$ is equicontinuous if for every $\varepsilon > 0$ there exists $\delta > 0$ such that

$$|x - y| < \delta \text{ implies } |w(x) - w(y)| < \varepsilon, \text{ for all } w \in W.$$

Theorem A1 (Schauder Fixed-Point Theorem)**.** *Let X be a bounded subset of R^n, and let $C(X)$ be the space of bounded continuous functions on X. Let $W \subset C(X)$ be nonempty, closed, bounded and convex. If the mapping $T : W \to W$ is continuous and the family $T(W)$ is equicontinuous, then T has a fixed point in W.*

We proceed in the following steps. We first show that W is a closed, convex subset of $C(X)$. Second, the operator T is continuous by Lemma A1. Third, we show that the family of functions generated by the operator $T(W)$ is equicontinuous. To this end, we use the Heine–Cantor Theorem which states that the continuous function evaluated in a compact set is also uniformly continuous. By the self-mapping of Lemma A1, Tw is continuous, and Assumption 1 ensures that (θ, b) lie in a compact set, therefore Tw is uniformly continuous. Thus, the family $T(W)$ is equicontinuous. We then use Schauder Fixed-Point Theorem. $\square$

The existence result relies on the Schauder Fixed-point Theorem. Assumption 1 plays an important role in ensuring boundedness and simplifying the proof. The cardinality assumption, on the other hand, is inconsequential. It just make sure that $T : W \to W$. If for instance $u : [\underline{b}, \bar{b}] \to \mathbb{R}_-$, we could redefine $\underline{w} = \bar{\theta}\frac{u}{1-\delta}$ and $\bar{w} = \underline{\theta}\frac{\bar{u}}{1-\beta\delta}$, and everything goes through without changes. If instead the u changes signs, since it is bounded below by $\underline{b}$, it is always possible to add a positive constant without altering the properties of the equilibrium.

A.6 Proof of Proposition 3

Consider any two functions $h, j \in C(X)$. Define the metric:

$$d(h, j) = \sup_{\theta, b} |h(\theta, b) - j(\theta, b)|.$$

For any θ, b, we obtain the following:

$$h(\theta, b) - j(\theta, b) \leq \sup_{\theta, b} |h(\theta, b) - j(\theta, b)| = d(h, j) \tag{19}$$

$$h(\theta, b) - j(\theta, b) \geq -\sup_{\theta, b} |h(\theta, b) - j(\theta, b)| = -d(h, j). \tag{20}$$

Use inequality (19) and integrate over Θ^+:

$$\int_{\Theta^+} h(\theta, b) dG(\theta) \leq \int_{\Theta^+} j(\theta, b) dG(\theta) + d(h, j) \int_{\Theta^+} dG(\theta'), \quad \forall b.$$

Similarly, use inequality (20) and integrate over Θ^-:

$$\int_{\Theta^-} h(\theta, b) dG(\theta) \leq \int_{\Theta^-} j(\theta, b) dG(\theta) - d(h, j) \int_{\Theta^-} dG(\theta), \quad \forall b.$$

Add up the inequalities above and recall that $G^+ = \int_{\Theta^+} dG(\theta)$ and $G^- = -\int_{\Theta^-} dG(\theta)$:

$$\int h(\theta, b) dG(\theta) \leq \int j(\theta, b) dG(\theta) + d(h, j) \left(G^+ + G^-\right), \quad \forall b. \tag{21}$$

Since $G^+ - G^- = 1$, the last term $G^+ + G^- = 1 + 2G^-$.

Assess inequality (21) at b' and multiply it by $\lambda\delta > 0$, combine with $h(\theta, b') \leq j(\theta, b') + d(h, j)$, and add $\theta U(b, b')$:

$$\begin{aligned}
&\theta U(b,b') + \delta\left((1-\lambda)\, h(\theta,b') + \lambda \int h(\theta',b')dG(\theta')\right)\\
&\leq \theta U(b,b') + \delta\left((1-\lambda)\, j(\theta,b') + \lambda \int j(\theta',b')dG(\theta')\right) + \gamma d(h,j), \quad \forall \theta, b, b',
\end{aligned}$$

where

$$\gamma = \delta\left(1 - \lambda + \lambda(1 + 2G^-)\right) = \delta\left(1 + 2\lambda G^-\right).$$

Since the inequality above holds for all b', we obtain

$$\begin{aligned}
&\sup_{b'}\left\{\theta U(b,b') + \delta\left((1-\lambda)\, h(\theta,b') + \lambda \int h(\theta',b')dG(\theta')\right)\right\}\\
&\leq \sup_{b'}\left\{\theta U(b,b') + \delta\left((1-\lambda)\, j(\theta,b') + \lambda \int j(\theta',b')dG(\theta')\right)\right\} + \gamma d(h,j).
\end{aligned}$$

That is,

$$(Th)(\theta, b) \leq (Tj)(\theta, b) + \gamma d(h, j).$$

Reversing the order of h and j, we get

$$|(Th)(\theta, b) - (Tj)(\theta, b)| \leq \gamma d(h, j), \quad \forall \theta, b.$$

Since it is for all θ and b, we can take the supremum:

$$d(Th, Tj) \leq \gamma d(h, j).$$

Condition 2 guarantees that $\gamma = \delta\,(1 + 2\lambda G^-) < 1$, making sure that it is a contraction. Lemma A1 shows that the operator T maps bounded and continuous function into itself. Since it is a contraction, by the Contraction Mapping Theorem there exists a unique solution.

A.7 Proof of Proposition 5

We propose a welfare criterion that ranks Markov perfect equilibria. We show that it is equivalent to the solution of a planner who maximizes ex-ante welfare before observing the realization of θ. We consider an "Overlapping Principals" problem, in which each agent is not only an agent but also acts as a principal designing rules governing the behavior of future agents. That is, when in control, every agent is committed to respecting the rules imposed by the previous one, but she sets rules that would restrict the actions of future agents. Furthermore, to retain the recursive Markovian structure we assume that the principal can choose rules depending only on the payoff relevant states and, thus, there is no record keeping.

To be precise, we consider incentive-compatible direct mechanisms. A mechanism $A = \{c(\theta, b), b'(\theta, b)\}_{\theta\in\Theta, b\geq \underline{b}}$ specifies a *feasible* consumption and savings choice depending on the type θ and the current stock of savings. This rule starts to operate when the agent comes into control and stays in place while she retains control. The agent reports her type every period and proposes a mechanism A' that binds the next agent, if she were to lose control.[14] Let $\tilde{w}(\tilde{\theta}, \theta, b; A)$ be the payoff of a type-θ agent who reports type $\tilde{\theta}$, given debt level b and mechanism A. Let $w(\theta, b; A) = \tilde{w}(\theta, \theta, b; A)$. Then,

$$\tilde{w}(\tilde{\theta}, \theta, b; A) = \max_{A'} \left\{ \theta u(c(\tilde{\theta}, b)) + \delta \Big[(1-\lambda) w(\theta, b'(\tilde{\theta}, b); A) + \lambda\beta\mathbb{E}[v(\theta', b'(\tilde{\theta}, b); A')] \Big] \right\} \quad (22)$$

with the companion value:

$$v(\theta, b; A) = \theta u(c(\theta, b)) + \delta \Big[(1-\lambda) v(\theta, b'(\theta, b); A) + \lambda\mathbb{E}[v(\theta', b'(\theta, b); A'^{*})] \Big],$$

where A'^{*}, the proposed mechanism, solves the problem in (22).

Truth telling implies $w(\theta, b; A) = \max_{\tilde{\theta}} \tilde{w}\left(\tilde{\theta}, \theta, b; A\right)$. In equation (22), the agent anticipates that if she were to stay in control, she would report again her true type, hence the continuation value in this case $w(\theta, b'(\tilde{\theta}, b); A)$ already imposes truth telling. If she were to lose control, because the next type θ' is $i.i.d.$ from θ, the optimal continuation mechanism is independent of the current taste θ. Thus when the current self loses control, her incentives are aligned with the previous self, given any state she is in. As a result, the maximization problem reduces to:

$$\max_{A} \beta\mathbb{E}[v(\theta', b; A)].$$

Discussion. This proposed welfare criterion is de facto choosing the best Markov equilibrium. One way to see this problem is as choosing the best static mechanism. If $\lambda = 1$, the restriction to history independence would be irrelevant, since as shown by Amador et al. (2006), the best unrestricted mechanism is static. If $\lambda < 1$, the restriction to history independence would imply that the best Markov equilibrium could be improve in a more general setting where contracts with history dependency are allowed, as in Halac and Yared (2014).

Under any incentive-compatible allocation A, the envelope condition (15) holds too.

[14] Because events are deterministic while the present agent is in control, it is irrelevant whether she sets the rules for the next agent before or upon losing control, so long as this occurs before the realization of the new type and the new rule only binds the future agent.

Hence the relation in Lemma 1 under any equilibrium choices extends to incentive-compatible direct mechanisms:

Lemma A2 (Value functions: IC). *Under any incentive-compatible allocation A, the relation in equation (4) holds too: at any $b \in [\underline{b}, \bar{b}]$, for almost all $\theta \in \Theta$,*

$$\beta v\left(\theta, b; A\right) = w\left(\theta, b; A\right) - (1-\beta)\theta w_\theta\left(\theta, b; A\right). \tag{23}$$

Proof. The proof follows the same steps as Lemma 1 and is omitted for brevity. □

Lemma A2 suggests that it is equivalent to look at the transformed continuation value:

$$\beta \int_{\underline{\theta}}^{\bar{\theta}} v\left(\theta, b; A\right) dF(\theta) = \int_{\underline{\theta}}^{\bar{\theta}} w\left(\theta, b; A\right) dG(\theta), \ \forall b$$

To see that leaving the agent unconstrained is optimal, guess that the incentive compatibility constraints are not binding. Then for each b, we choose allocation for each type:

$$\begin{aligned}
&\max_{\{c(\theta), b'(\theta)\}} \int_{\underline{\theta}}^{\bar{\theta}} w\left(\theta, b; A\right) dG(\theta) \\
&= \max_{\{c(\theta), b'(\theta)\}} \int_{\underline{\theta}}^{\bar{\theta}} \left[\theta u(c(\theta)) + \delta\left((1-\lambda) w(\theta, b'(\theta); A) + \lambda \int_{\underline{\theta}}^{\bar{\theta}} w(\theta', b'(\theta); A) dG(\theta')\right)\right] dG(\theta) \\
&= \int_{\underline{\theta}}^{\bar{\theta}} \left[\max_{\{c, b'\}} \left\{\theta u(c) + \delta\left((1-\lambda) w(\theta, b'; A) + \lambda \int_{\underline{\theta}}^{\bar{\theta}} w(\theta', b'; A) dG(\theta')\right)\right\}\right] dG(\theta) \\
&= \int_{\underline{\theta}}^{\bar{\theta}} \left[\max_{\{c, b'\}} \left\{\theta u(c) + \delta\left((1-\lambda) w(\theta, b'; A) + \lambda\beta \int_{\underline{\theta}}^{\bar{\theta}} v(\theta', b'; A) dF(\theta')\right)\right\}\right] dG(\theta).
\end{aligned}$$

The first equality follows from the definition of w, the second from the fact that all weights $dG(\theta)$ are positive, and the last equality is due again to Lemma A2. It follows then that the optimal allocations coincide with the unconstrained optimal choice of each agent, confirming the conjecture that the IC constraints are not binding and the proposition's statement.

A.8 Proof of Proposition 6

We omit most of the details because the proof follows the same steps as Lemma 1 and Proposition 1. Replacing summations over time with integrals, and following the same steps that all the results also hold in continuous time, and thus implies equation (6). Substituting it into the HJB equation (11), we obtain the transformed HJB equation (13).

A.9 Proof of Proposition 7

We provide two proofs. The first proof assumes that the solution is differentiable everywhere and sidesteps potential issues of boundary of the state space. This highly simplifies the proof and highlights the role played by Condition 3. The second proof is lengthy, since it must deal with points of non-differentiability and boundary of the state space, but it closely follows the standard approach in the literature.

Proof with differentiability. Suppose by contradiction that there exist two viscosity solutions. Let w_1 and w_2 denote the two solutions. By definition, each is both a subsolution and supersolution. Then

$$M = \max_{\bar{\Omega}} |w_1 - w_2| = \max\{\max_{\bar{\Omega}}(w_1 - w_2), \max_{\bar{\Omega}}(w_2 - w_1)\} > 0.$$

There are two possibilities: $M = \max_{\bar{\Omega}}(w_1 - w_2)$ or $M = \max_{\bar{\Omega}}(w_2 - w_1)$. We show that each case leads to a contradiction.

Case 1: $M = \max_{\bar{\Omega}}(w_1 - w_2)$**.** Suppose it is attained at, say, $(\theta^*, b^*) \in \Omega$ in the interior. Since $w_1 - w_2$ attains local maximum at interior (θ^*, b^*) and since w_1 and w_2 are differentiable, then $w_{1b}(\theta^*, b^*) = w_{2b}(\theta^*, b^*)$.

We use the fact tha w_1 is a subsolution and w_2 is a supersolution. Since w_1 is a subsolution, set $\phi = w_1$. Similarly, because w_2 is a supersolution, set $\phi = w_2$. Thus,

$$(\rho + \tilde{\lambda}) w_1(\theta^*, b^*) \le \max_{c \ge 0} \left\{ \theta^* u(c) + (y + rb^* - c) w_{1b}(\theta^*, b^*) + \tilde{\lambda} \int_{\underline{\theta}}^{\bar{\theta}} w_1(\theta', b^*) dG(\theta') \right\}$$
$$(\rho + \tilde{\lambda}) w_2(\theta^*, b^*) \ge \max_{c \ge 0} \left\{ \theta^* u(c) + (y + rb^* - c) w_{2b}(\theta^*, b^*) + \tilde{\lambda} \int_{\underline{\theta}}^{\bar{\theta}} w_2(\theta', b^*) dG(\theta') \right\}.$$

Subtracting the second from the first, and since $w_{1b}(\theta^*, b^*) = w_{2b}(\theta^*, b^*)$:

$$(\rho + \tilde{\lambda})(w_1(\theta^*, b^*) - w_2(\theta^*, b^*)) \le \tilde{\lambda} \int_{\underline{\theta}}^{\bar{\theta}} [w_1(\theta', b^*) - w_2(\theta', b^*)] dG(\theta').$$

Notice that on the right-hand side of the last equation, due to the possibility that $dG(\theta) < 0$, we lose the usual monotonicity and in turn the standard Comparison Theorem. Nevertheless, with this modified approach we have $|w_1 - w_2| \le M$, so we can still obtain an upper bound:

$$\int_{\underline{\theta}}^{\bar{\theta}} [w_1(\theta', b^*) - w_2(\theta', b^*)] dG(\theta') \le \int_{\underline{\theta}}^{\bar{\theta}} |w_1(\theta', b^*) - w_2(\theta', b^*)| \cdot |g(\theta')| d\theta' \le M(G^+ + G^-).$$

Therefore,

$$(\rho + \tilde{\lambda})M \leq \tilde{\lambda} M(G^{+} + G^{-}).$$

Since $G^{+} - G^{-} = 1$, this gives:

$$\rho M \leq \tilde{\lambda} M \cdot 2G^{-}.$$

Since $M > 0$, then $\rho \leq 2\tilde{\lambda}G^{-}$. This contradicts (14) in Condition 3.

Case 2: $M = \max_{\bar{\Omega}}(w_2 - w_1)$**.** We now use the fact that w_2 is subsolution and w_1 is supersolution. We repeat the steps in the previous case and again arrive to a contradiction.

It must be that $\max_{\bar{\Omega}}(w_2 - w_1) = 0$ and uniqueness follows.

Proof without differentiability. Because the solution may have countably many points of non-differentiability, and because at the boundary of $\bar{\Omega}$ the solution does not need to be a supersolution, the previous intuitive proof may fail. Hence, we prove this proposition following closely the steps in Capuzzo-Dolcetta and Lions (1990). This allows for a straightforward comparison with the literature and highlights the main difference. We start by proving a result analogous to their Theorem III.1. For any function ϕ, define $\phi^{+} = \max\{\phi, 0\}$, then:

Suppose there are two Constrained Viscosity Solutions $U, V \in C(\bar{\Omega})$. As before, there are two possibilities: $\max_{\bar{\Omega}}(U - V) \geq \max_{\bar{\Omega}}(V - U)$, ensuring that $\max_{\bar{\Omega}} |U - V| = \max_{\bar{\Omega}}(U - V)$ or $\max_{\bar{\Omega}}(V - U) \geq \max_{\bar{\Omega}}(U - V)$, ensuring $\max_{\bar{\Omega}} |U - V| = \max_{\bar{\Omega}}(V - U)$. In the following theorem, we consider the first case and it extends to the second case.

Theorem A2 (Modified Comparison)**.** *Let $\bar{\Omega} = [\underline{\theta}, \bar{\theta}] \times [\underline{b}, \bar{b}]$, define:*

$$H(\theta, b, p) = \max_{c \geq 0} \{\theta u(c) + (y + rb - c)p\}.$$

And assume:

(H1) $|H(\theta, b, p) - H(\theta', b', p)| \leq \omega(|(\theta, b) - (\theta', b')|(1 + |p|))$ *for some modulus ω.*

(H2) $|H(\theta, b, p) - H(\theta, b, q)| \leq \mu(|p - q|)$ *for some modulus μ.*

Let $U \in C(\bar{\Omega})$ be a viscosity subsolution of

$$(\rho + \tilde{\lambda})U = H(\theta, b, U_b) + \chi(\theta, b) + \tilde{\lambda} \int_{\underline{\theta}}^{\bar{\theta}} U(\theta', b) dG(\theta') \tag{24}$$

in $\bar{\Omega}$, where $\chi \in C(\bar{\Omega})$, and $V \in C(\bar{\Omega})$ be a viscosity supersolution of (13) in Ω. Suppose $\max_{\bar{\Omega}}(U - V) \geq \max_{\bar{\Omega}}(V - U)$*. Then:*

$$\max_{\bar{\Omega}}(U - V)^{+} \leq \frac{1}{\rho - 2\tilde{\lambda}G^{-}} \max_{\bar{\Omega}} \chi^{+}.$$

Remark. This theorem immediately delivers the result in Proposition 7. First, since we are restricting the state space to be compact, the assumptions on the utility function directly imply *(H1)* and *(H2)*. In addition, the result is true for any continuous χ. Thus, by taking $\chi \equiv 0$, it implies $U \leq V$ everywhere and further $U = V$.

Proof of Theorem A2. Suppose, by contradiction, that

$$\max_{\bar{\Omega}}(U-V)^{+} > \frac{1}{\rho - 2\tilde{\lambda}G^{-}} \max_{\bar{\Omega}} \chi^{+}.$$

Define $M := \max_{\bar{\Omega}}(U-V)^{+} > 0$. If the maximum is attained in the interior of $\bar{\Omega}$, i.e., in Ω, standard viscosity solution techniques apply (see, for instance, Theorem 9.1 in Chapter II of Fleming and Soner (2006)). Therefore, in this proof we consider

$$M = (U-V)^{+}(\theta^{*}, b^{*}) = (U-V)(\theta^{*}, b^{*}) \text{ for some } (\theta^{*}, b^{*}) \in \partial\Omega.$$

To reproduce the arguments for the case in which the maximum is attained in Ω rather than on $\partial\Omega$, it suffices to set $\zeta \equiv 0$ in what follows.

Step 1: Doubling of variables sequence. Consider the auxiliary function:

$$\Phi_{\varepsilon}(\theta, b, \theta', b') = U(\theta, b) - V(\theta', b') - \frac{1}{\varepsilon^{2}}|(\theta, b) - (\theta', b') - \varepsilon T(\theta', b')|^{2},$$

where $T(\theta', b') = \zeta(\theta', b')n(\theta', b')$ is defined as follows.[15] For ε_0 small enough, define a closed neighborhood of $\partial\Omega$: $\Gamma_0 = \{(\theta', b') : \text{dist}((\theta', b'), \partial\Omega) \leq \varepsilon_0\}$.

(i) $\zeta \equiv 1$ if $(\theta', b') \in \Gamma_0$, $\zeta \in C^1(\bar{\Omega})$; $\zeta \equiv 0$ if $(\theta', b') \in \bar{\Omega} - \Gamma_0$.

(ii) $n(\theta', b') = -\nabla d(\theta', b')$ is the inward unit normal, where $d(\theta', b') = \text{dist}((\theta', b'), \partial\Omega)$. We take ε_0 small enough so that d is differentiable on Γ_0.

Step 2: Sequence analysis. Start by setting a lower bound on the maximum. Since $\chi^{+} \geq 0$ and $\rho \geq 2\tilde{\lambda}G^{-}$, it must be that $M > 0$, and by continuity of U:

$$\max_{\bar{\Omega}\times\bar{\Omega}} \Phi_{\varepsilon} \geq \Phi_{\varepsilon}((\theta^{*}, b^{*}), (\theta^{*}, b^{*}) - \varepsilon n(\theta^{*}, b^{*})) \geq M - \omega(\varepsilon),$$

where ω is the modulus of continuity of V.

[15] For example, T in our rectangular domain would be: At a boundary point like $(\bar{\theta}, b')$, we have $n(\bar{\theta}, b') = (-1, 0)$ (pointing inward). So near this boundary: $T(\bar{\theta}, b') = \zeta(\bar{\theta}, b')(-1, 0) = (-1, 0)$. This means $(\bar{\theta}, b') + \varepsilon T(\bar{\theta}, b') = (\bar{\theta} - \varepsilon, b')$, which is pushed inside Ω.

Let $((\theta_\varepsilon, b_\varepsilon), (\theta'_\varepsilon, b'_\varepsilon)) \in \bar{\Omega} \times \bar{\Omega}$ be a maximum point of Φ_ε. Note that

$$\max_{\bar{\Omega}\times\bar{\Omega}} \Phi_\varepsilon \leq M - \frac{1}{\varepsilon^2}|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon) - \varepsilon T(\theta'_\varepsilon, b'_\varepsilon)|^2 + \omega(|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon)|).$$

Combining the last two inequalities we obtain:

$$\frac{1}{\varepsilon^2}|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon) - \varepsilon T(\theta'_\varepsilon, b'_\varepsilon)|^2 \leq \omega(|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon)|) + \omega(\varepsilon). \tag{25}$$

Claim 1. $|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon) - \varepsilon T(\theta'_\varepsilon, b'_\varepsilon)| \leq \varepsilon\delta(\varepsilon)$, *where* $\delta(\varepsilon) \to 0$ *as* $\varepsilon \to 0$.

Proof. We first show that

$$|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon)| \leq C\varepsilon,$$

for some constant C. Equation (25) immediately implies $|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon) - \varepsilon T(\theta'_\varepsilon, b'_\varepsilon)| \leq C_1\varepsilon$, for some C_1 when ε is small enough such that $\sqrt{\omega(|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon)|) + \omega(\varepsilon)} \leq C_1$. Therefore, $|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon)| \leq |(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon) - \varepsilon T(\theta'_\varepsilon, b'_\varepsilon)| + \varepsilon|T(\theta'_\varepsilon, b'_\varepsilon)| \leq C_1\varepsilon + C_2\varepsilon = C\varepsilon$. Combining $|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon)| \leq C\varepsilon$ and (25), we obtain the result. □

Claim 2. $(\theta'_\varepsilon, b'_\varepsilon) \in \Omega$ *(interior) for* ε *small enough.*

Proof. Suppose by contradiction that $(\theta'_\varepsilon, b'_\varepsilon) \in \partial\Omega$. Then, $T(\theta'_\varepsilon, b'_\varepsilon) = n(\theta'_\varepsilon, b'_\varepsilon)$ since $\zeta = 1$ near $\partial\Omega$. It follows that $(\theta'_\varepsilon, b'_\varepsilon) + \varepsilon n(\theta'_\varepsilon, b'_\varepsilon)$ is at distance ε inside Ω. Since n is the unit inward, it must be that $|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon) - \varepsilon T(\theta'_\varepsilon, b'_\varepsilon)| \geq \varepsilon$. However, for ε small enough such that $\delta(\varepsilon) < 1$, this is in contradiction with Claim 1. □

Step 3: Applying viscosity solution properties. Recall that U and V are not required to be differentiable. To this end, as is standard in the literature, we replace U_b and V_b with alternative functions p_ε and q_ε, respectively. These functions are such that if U_b and V_b exist they would equal p_ε and q_ε, but the latter ones always exist even when the former are not well defined. Since $(\theta'_\varepsilon, b'_\varepsilon) \in \Omega$ and $((\theta_\varepsilon, b_\varepsilon), (\theta'_\varepsilon, b'_\varepsilon))$ maximizes Φ_ε, we can compute the gradients. With respect to (θ', b') so that $\nabla_{(\theta',b')}\Phi_\varepsilon = 0$ at $(\theta'_\varepsilon, b'_\varepsilon)$ generates:

$$\nabla V(\theta'_\varepsilon, b'_\varepsilon) = \frac{2}{\varepsilon^2}(I - \varepsilon\nabla T(\theta'_\varepsilon, b'_\varepsilon))[(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon) - \varepsilon T(\theta'_\varepsilon, b'_\varepsilon)],$$

where I is the identity matrix. In particular, the b-component:[16]

$$q_\varepsilon := V_b(\theta'_\varepsilon, b'_\varepsilon) = \frac{2}{\varepsilon^2}(1 - \varepsilon\nabla_b T_b(\theta'_\varepsilon, b'_\varepsilon))[b_\varepsilon - b'_\varepsilon - \varepsilon T_b(\theta'_\varepsilon, b'_\varepsilon)] - \frac{2}{\varepsilon^2}\varepsilon\nabla_b T_\theta(\theta'_\varepsilon, b'_\varepsilon)[\theta_\varepsilon - \theta'_\varepsilon - \varepsilon T_\theta(\theta'_\varepsilon, b'_\varepsilon)].$$

[16] The T_θ derivatives appear because $T(\theta, b)$ is a vector-valued function of both θ and b, so when differentiating the squared norm with respect to b, we include how both components of T change with b.

Similarly,

$$p_\varepsilon := U_b(\theta_\varepsilon, b_\varepsilon) = \frac{2}{\varepsilon^2}[b_\varepsilon - b'_\varepsilon - \varepsilon T_b(\theta'_\varepsilon, b'_\varepsilon)].$$

Thus, in what follows we replace U_b and V_b (which may not exist) with p_ε and q_ε (that are always well defined). Since U is a viscosity subsolution at $(\theta_\varepsilon, b_\varepsilon)$:

$$(\rho + \tilde{\lambda})U(\theta_\varepsilon, b_\varepsilon) \leq H(\theta_\varepsilon, b_\varepsilon, p_\varepsilon) + \chi(\theta_\varepsilon, b_\varepsilon) + \tilde{\lambda} \int_{\underline{\theta}}^{\bar{\theta}} U(\theta', b_\varepsilon) dG(\theta'). \tag{26}$$

Since V is a viscosity supersolution at $(\theta'_\varepsilon, b'_\varepsilon) \in \Omega$:

$$(\rho + \tilde{\lambda})V(\theta'_\varepsilon, b'_\varepsilon) \geq H(\theta'_\varepsilon, b'_\varepsilon, q_\varepsilon) + \tilde{\lambda} \int_{\underline{\theta}}^{\bar{\theta}} V(\theta', b'_\varepsilon) dG(\theta'). \tag{27}$$

Step 4: Combining the inequalities. Subtracting (27) from (26):

$$\begin{aligned}
&(\rho + \tilde{\lambda})[U(\theta_\varepsilon, b_\varepsilon) - V(\theta'_\varepsilon, b'_\varepsilon)] \\
\leq\ & \tilde{\lambda} \int_{\underline{\theta}}^{\bar{\theta}} [U(\theta', b_\varepsilon) - V(\theta', b'_\varepsilon)]\, dG(\theta') + H(\theta_\varepsilon, b_\varepsilon, p_\varepsilon) - H(\theta'_\varepsilon, b'_\varepsilon, q_\varepsilon) + \chi(\theta'_\varepsilon, b'_\varepsilon).
\end{aligned}$$

So far, we have followed the standard procedure. The last term is particular to our setting. The boundedness assumption highly simplifies the proof. Since f is bounded, and so is g, and U, V are continuous functions on the compact set $\bar{\Omega}$, they are also Lipschitz continuous, then:

$$\begin{aligned}
\int_{\underline{\theta}}^{\bar{\theta}} [U(\theta', b_\varepsilon) - V(\theta', b'_\varepsilon)]\, dG(\theta') &\leq \int_{\underline{\theta}}^{\bar{\theta}} |U(\theta', b_\varepsilon) - V(\theta', b'_\varepsilon)||g(\theta')| d\theta' \\
&\leq \int_{\underline{\theta}}^{\bar{\theta}} [|U(\theta', b_\varepsilon) - U(\theta', b'_\varepsilon)| + |U(\theta', b'_\varepsilon) - V(\theta', b'_\varepsilon)|]\, |g(\theta')| d\theta' \\
&\leq M(G^+ + G^-) + \int_{\underline{\theta}}^{\bar{\theta}} |U(\theta', b_\varepsilon) - U(\theta', b'_\varepsilon)||g(\theta')| d\theta' \\
&\leq M(G^+ + G^-) + \int_{\underline{\theta}}^{\bar{\theta}} L_2(\theta')|b_\varepsilon - b'_\varepsilon||g(\theta')| d\theta'.
\end{aligned}$$

The second inequality follows from the triangle inequality, the third from the definition of $|U(\theta', b'_\varepsilon) - V(\theta', b'_\varepsilon)| \leq M$, the fourth from the Lipschitz continuity of U, with $L_2(\theta)$ being the local Lipschitz constant.

Step 5: Bound on Hamiltonian difference. Bounding $H(\theta_\varepsilon, b_\varepsilon, p_\varepsilon) - H(\theta'_\varepsilon, b'_\varepsilon, q_\varepsilon)$.

Case 1: If $(\theta'_\varepsilon, b'_\varepsilon) \notin \Gamma_0$ (away from boundary), then $T(\theta'_\varepsilon, b'_\varepsilon) = 0$, so

$$p_\varepsilon = q_\varepsilon = \frac{2}{\varepsilon^2}(b_\varepsilon - b'_\varepsilon).$$

Using *(H1)*:

$$|H(\theta_\varepsilon, b_\varepsilon, p_\varepsilon) - H(\theta'_\varepsilon, b'_\varepsilon, p_\varepsilon)| \leq \omega(C\varepsilon(1 + |p_\varepsilon|)).$$

From inequality (25) and the definition of p_ε, it is clear that $|p_\varepsilon| \leq \eta(\varepsilon) = \omega(|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon)|) + \omega(\varepsilon)$ as the sequence converges to the boundary. Thus, we get a bound that goes to zero as $\varepsilon \to 0$

Case 2: If $(\theta'_\varepsilon, b'_\varepsilon) \in \Gamma_0$ (near boundary), we have

$$|p_\varepsilon - q_\varepsilon| \leq \left| \frac{2}{\varepsilon^2} \varepsilon \nabla T(\theta'_\varepsilon, b'_\varepsilon)[(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon) - \varepsilon T(\theta'_\varepsilon, b'_\varepsilon)] \right| \leq C\delta(\varepsilon).$$

The second inequality follows from Claim 1 and $\delta(\varepsilon) \to 0$.

Using *(H1)* and *(H2)*:

$$\begin{aligned}|H(\theta_\varepsilon, b_\varepsilon, p_\varepsilon) - H(\theta'_\varepsilon, b'_\varepsilon, q_\varepsilon)| &\leq |H(\theta_\varepsilon, b_\varepsilon, p_\varepsilon) - H(\theta'_\varepsilon, b'_\varepsilon, p_\varepsilon)| + |H(\theta'_\varepsilon, b'_\varepsilon, p_\varepsilon) - H(\theta'_\varepsilon, b'_\varepsilon, q_\varepsilon)| \\ &\leq \omega(|(\theta_\varepsilon, b_\varepsilon) - (\theta'_\varepsilon, b'_\varepsilon)|(1 + |p_\varepsilon|)) + \mu(|p_\varepsilon - q_\varepsilon|) \\ &\leq \omega(C\varepsilon(1 + |p_\varepsilon|)) + \mu(C\delta(\varepsilon)).\end{aligned}$$

Combining all bounds:

$$(\rho + \tilde{\lambda})[U(\theta_\varepsilon, b_\varepsilon) - V(\theta'_\varepsilon, b'_\varepsilon)] \leq \tilde{\lambda} M(G^+ + G^-) + \kappa(\varepsilon) + \chi(\theta'_\varepsilon, b'_\varepsilon),$$

where $\kappa(\varepsilon) = \int_{\underline{\theta}}^{\bar{\theta}} L_2(\theta')|b_\varepsilon - b'_\varepsilon||g(\theta')|d\theta' + \omega(C\varepsilon(1 + |p_\varepsilon|)) + \mu(C\delta(\varepsilon)) \to 0$ as $\varepsilon \to 0$. As a result, taking $\varepsilon \to 0$, we have $(\theta_\varepsilon, b_\varepsilon), (\theta'_\varepsilon, b'_\varepsilon) \to (\theta^*, b^*)$, which implies:

$$(\rho+\tilde{\lambda})M = (\rho+\tilde{\lambda})[U(\theta^*, b^*) - V(\theta^*, b^*)] \leq \tilde{\lambda} M(G^+ + G^-) + \chi(\theta^*, b^*) \leq \tilde{\lambda} M(G^+ + G^-) + \max_{\bar{\Omega}} \chi^+.$$

Then, using $G^+ - G^- = 1$ we obtain:

$$\rho M \leq 2\tilde{\lambda} M G^- + \max_{\bar{\Omega}} \chi^+ \quad \Rightarrow \quad M(\rho - 2\tilde{\lambda} G^-) \leq \max_{\bar{\Omega}} \chi^+.$$

As long as Condition 3 is satisfied, this is in contradiction with the initial assumption. $\square$