Time Inconsistency with Heterogeneous Agents*

Eduardo Dávila† Andreas Schaab‡

September 2025

Abstract

This paper studies the time inconsistency of policy in heterogeneous agent economies, where the need to aggregate individual welfare gains creates a distinct source of time inconsistency. We establish that welfare and efficiency assessments are time inconsistency for different reasons, even when agents have identical preferences and when no inconsistency would emerge if individual welfare gains were considered in isolation. Only utilitarian social welfare functions ensure that welfare assessments are time-consistent. In turn, it is impossible to always have time-consistent efficiency assessments that satisfy the compensation principle and are useful to define Pareto frontiers when markets are incomplete. As a result, even when optimal policy is time-consistent, its justification on grounds of efficiency and redistribution changes as time passes and uncertainty is realized. We characterize these forces in a general dynamic stochastic environment and illustrate their policy relevance through applications to anticipated relief policies, insurance, aggregate investment, and labor taxation.

JEL Codes: D52, D60, D70, E60, H21

Keywords: time inconsistency, heterogeneous agents, incomplete markets, welfare, efficiency, redistribution, risk-sharing

*We would like to thank Marco Bassetto, Joel Flynn, Joao Guerreiro, Marina Halac, Ricardo Reis, Karthik Sastry, Jón Steinsson, Ludwig Straub, Alireza Tabaz-Salehi, and Pierre Yared for helpful conversations, as well as the audiences at various conferences and seminars for their comments. This material is based upon work supported by the National Science Foundation under Grant DGE-2237790.

†Yale University and NBER. eduardo.davila@yale.edu

‡University of California, Berkeley, and NBER. schaab@berkeley.edu

1 Introduction

Concerns about time consistency are central to dynamic policy analysis. A policy desirable from today's perspective is called *time-consistent* if it remains desirable when revisited in the future. Because only time-consistent policies can be implemented credibly, understanding why time inconsistency emerges is indispensable for normative economics. Building on Strotz (1956) and Kydland and Prescott (1977), a large literature has examined the origins and implications of time inconsistency, typically in representative-agent environments.[1] This paper advances the study of the time consistency of policies in environments with heterogeneous agents, where the need to aggregate individual welfare gains emerges as a new source of time inconsistency.

Our central contribution characterizes the distinct forms of time inconsistency that welfare and efficiency assessments exhibit in heterogeneous agent economies with incomplete markets, even when individuals have identical preferences and when no inconsistency would emerge if individuals were in isolation. We establish our results in a general dynamic stochastic environment with heterogeneous agents, making minimal assumptions about the structure of the economy.

We set the stage for our analysis of heterogeneity by first providing a new characterization of the forms of inconsistency that emerge in representative/single agent scenarios. This serves as a useful benchmark, as by shutting down these forces we can isolate the novel sources of inconsistency caused by heterogeneity, which is our ultimate goal. We then show that the welfare assessments of welfarist planners exhibit a distinct form of time inconsistency that arises from the need to make interpersonal welfare comparisons (in utils) across heterogeneous agents. This form of inconsistency is governed by a particular cross-sectional covariance determined by the ratio of ex ante and ex post Pareto weights. A corollary of this result is that only (equal- or unequal-weighted) utilitarian social welfare functions — with linear coefficients on individual utilities — ensure time consistency due to interpersonal welfare comparisons. Intuitively, when the planner's marginal rate of substitution between individuals in utils changes over time, the desirability of a policy can also shift. Only utilitarian social welfare functions are immune to this form of time inconsistency because Pareto weights are fixed across dates and histories.

Because welfarist welfare assessments combine efficiency and redistribution assessments, it

[1] Strotz (1956) shows that non-exponential discounting implies inconsistent intertemporal preferences. When the planner inherits such preferences, policy becomes time inconsistent. Kydland and Prescott (1977) show that when the planner can influence individual behavior today by making promises about policy in the future, she may have an incentive to renege on these promises ex post once individual decisions are sunk.

is natural to also investigate whether these components are themselves time-consistent. Our main result shows they are not: it is impossible to always have time-consistent efficiency (and redistribution) assessments with forward-looking numeraires in heterogeneous agent economies with incomplete markets, even when the overall welfare assessment is time-consistent. This impossibility results implies that even when a welfare-maximizing planner assesses policies consistently, the rationale that justifies a particular policy — whether welfare gains are due to efficiency or redistribution considerations — changes over time. This is perhaps the most important practical takeaway of our results, as it identifies an unavoidable tension between ensuring that policies are time-consistent and efficiency-maximizing.

Efficiency assessments become time inconsistent when the welfare numeraire — the unit used to evaluate individual welfare changes — varies over time. We refer to this as *numeraire inconsistency* and again associate the resulting time consistency problem with a particular cross-sectional covariance. The inconsistency arises because, when markets are incomplete, the rate at which individuals (and the planner) convert utility gains into consumption-equivalent gains — needed to define efficiency gains — changes over time. As we explain in detail, using forward-looking welfare numeraires — that is, choosing units based on contemporaneous and future consumption, as it is routinely done in practice — is necessary to ensure that efficiency assessments satisfy the compensation principle and are useful to define Pareto frontiers at all times in which assessments are made. When markets are complete, numeraire inconsistency cannot arise, because individuals' relative valuations are equalized at all times. Backward-looking numeraires would preserve time-consistency, but only at the cost of rendering efficiency assessments meaningless for policy analysis.

Finally, we examine the inconsistency of the different sources of efficiency gains. We show that all sources of efficiency gains — aggregate-efficiency, risk-sharing, and intertemporal-sharing — are vulnerable to time-inconsistency when markets are incomplete, except when a perturbation solely impacts aggregate net consumption at a single point in time. The inconsistency of the different sources arises because individuals' valuations of consumption across dates and histories vary with the perspective from which the assessment is made.

Motivating Example. Before presenting the general results, we illustrate the main ideas of the paper in a stylized dynamic model of taxation in which a distortionary labor income tax finances a lump-sum transfer. We consider a two-period economy with ex-ante identical households who face labor productivity risk, and contrast the forces that determine the optimal labor income tax ex ante and ex post.

We illustrate four results from this stylized example. First, we show that the attribution of welfare gains to efficiency and redistribution changes with the perspective of the assessments and is therefore time-inconsistent. Ex ante, the planner perceives that a tax increase yields risk-sharing gains, while there are no redistribution gains. Ex post, however, risk-sharing no longer plays any role, and all gains are solely attributed to redistribution. Second, the production efficiency loss due to distortionary taxation is the same ex ante and ex post. This result highlights that specific sources of efficiency gains may be time-consistent in particular situations, as explained in Section 4.3. Third, the overall welfare assessment of the policy is time-consistent under the utilitarian SWF. Fourth, the optimal policy is time-inconsistent for all other SWFs, since the utilitarian SWF is the only one that values ex post redistribution and ex ante risk-sharing equivalently.

Applications. We present three applications that demonstrate how the time inconsistency forces identified in this paper provides a useful new perspective on applied policy questions. Each application is designed with a dual objective: first, capture a realistic scenario in which the time inconsistency forces due to heterogeneity that we have studied play an important role, and second, illustrate how a particular source of efficiency gains — intertemporal-sharing, risk-sharing, and aggregate-efficiency — can lead to the time inconsistency of policies.

Anticipated Relief Policies: Our first application identifies a particular form of time inconsistency in the design of anticipated relief policies — policies intended to support individuals who experience a temporary spell of low consumption followed by a gradual, anticipated recovery. Relief policies of this sort are common, including the COVID-19 assistance programs, childcare support and parental leave, unemployment insurance, retraining programs, and affirmative action. In each case, the efficiency-maximizing policy involves supporting low-consumption individuals early in their period of hardship, phasing out as consumption recovers and approaches long-run levels. However, ex post, consumption is still below its long-run level at times in which the originally transfer policy phased out, opening the door to a new round of transfers. The efficiency-maximizing transfer is therefore extended ex post beyond the ex ante optimal time of expiry.

Risk-Sharing Policies: Our second application reveals a time inconsistency that arises when designing risk-sharing policies. In economies where agents face uncertain outcomes, the ex ante efficiency-maximizing policy calls for contingent transfers from lucky to unlucky agents at a terminal date. However, as uncertainty partially resolves, terminal states that appeared "bad" for one of the agents from an ex ante perspective may seem relatively favorable given new information about systematic differences between agents. The direction of transfers is then reversed, undoing earlier risk-sharing arrangements. This dynamic inconsistency mirrors real-world patterns, such as the EU's shift in support for peripheral countries after the 2008–2011 crisis. While ex ante support was justified by symmetric risks, the revelation that peripheral economies are persistently weaker led to demands for austerity and reduced assistance.

Aggregate Investment Policies: Our third application shows that efficiency-maximizing planners become present biased when assessing investment policies in the present of consumption trends, generating a type of time inconsistency that increases the hurdle rate for future aggregate investments as time goes by. Individuals with declining consumption are always more willing to undertake investment policies but less willing to pay for them over time, while individuals with growing consumption are always less willing to undertake investment policies, but disproportionally more willing to pay for them as time goes by. This creates a form of present bias where aggregate investment policies become less attractive over time from an efficiency perspective, as their benefits increasingly reflect the preferences of more impatient (growing-consumption) individuals. This application identifies a new source of present bias in efficiency assessments that is distinct from discount factor heterogeneity, complementing the results of Jackson and Yariv (2014, 2015).

Related Literature. This paper contributes to the study of time inconsistency and the analysis of heterogeneous agent economies. Building on the classic insights of Strotz (1956) and Kydland and Prescott (1977), a vast literature has examined the time consistency of government policies, especially fiscal and monetary policy.[2] Much of this work relies on representative agent models. A smaller set of contributions, including Werning (2007), Chang (2022), Bilbiie (2024) and Dávila and Schaab (2023a), study time consistency when individuals are heterogeneous but restrict attention to linear social welfare functions.[3] Besides our reformulation — in the dual — of the forms of time inconsistency in single-agent scenarios, we contribute to this body of work in three ways. First, we characterize the conditions under which an overall welfare assessment

[2] See Barro and Gordon (1983), Lucas and Stokey (1983), Rogoff (1985), Chari and Kehoe (1990), Aiyagari, Marcet, Sargent, and Seppälä (2002), Athey, Atkeson, and Kehoe (2005), Amador, Werning, and Angeletos (2006), Klein, Krusell, and Rios-Rull (2008), Halac and Yared (2014), Marcet and Marimon (2019), Dovis and Kirpalani (2020, 2021), Sublet (2023), and Clayton and Schaab (2025), among many others.

[3] In Werning (2007), the usual inconsistency of capital taxation in a representative agent environment is present because the planner lacks lump-sum taxation and capital expropriation once accumulation decisions are sunk mimics a non-distortionary tax. With heterogeneity, optimal capital taxation may remain time inconsistent even in the presence of a lump-sum tax because it helps the planner redistribute across agents, which is otherwise only possible with distortionary labor taxes. This mechanism is in the spirit of Kydland and Prescott (1977) since interpersonal comparisons are not source of time inconsistency because the SWF is linear.

is time-consistent with heterogeneous agents, showing that time inconsistency materializes when the SWF is non-linear. Second, we prove that efficiency and redistribution assessments are invariably time inconsistent with incomplete markets, even when the SWF is linear. Third, we provide novel characterizations of the forms of inconsistency of the different sources of efficiency.

Our results also connect to the social choice literature on aggregation of individual preferences, that follows Arrow (1950)'s impossibility result. More closely, Zuber (2011) shows that time consistency generically requires a linear social welfare function. Our approach based on welfare assessments is complementary to his, as it allows us to associate time inconsistency in interpersonal welfare comparisons with a specific cross-sectional covariance. A strand of this literature studies the aggregation of heterogeneous discount factors and shows that time inconsistency in social choice is an unavoidable consequence of individual heterogeneity in time preferences (Marglin, 1963; Feldstein, 1964; Jackson and Yariv, 2014, 2015; Adams, Cherchye, De Rock, and Verriest, 2014; Chambers and Echenique, 2018). Halevy (2015) and Millner and Heal (2018) distinguish between time invariance and time consistency.[4] They clarify that while time invariant and stationary social welfare functions necessarily imply time-consistent welfare assessments, time invariance is not required for time consistency. Relatedly, Mongin (1995, 1998) shows that belief heterogeneity also implies time inconsistent aggregate welfare assessments. In this paper, we consider time invariant SWFs but allow for non-stationarity. To highlight that none our findings are not driven by preference heterogeneity, we assume in the body of the paper that all individuals have identical preferences.

The tension between efficiency or risk-sharing and redistribution plays an important role in our paper and has long been recognized in economics. Marshall (1920), Pigou (1920) and Knight (1921) already noted that schemes providing risk-sharing ex ante imply redistributive transfers ex post once uncertainty has realized. Arrow (1963) explicitly states that:

> *[T]he preference for redistribution expressed in government taxation (...) can be reinterpreted as desire for insurance (...). Thus, optimality, in a context which*

[4]The social choice literature often discusses two properties of SWF's: time invariance and stationarity (Koopmans, 1960; Halevy, 2015; Millner and Heal, 2018). A SWF is *time invariant* if it does not directly depend on the perspective of the assessment s^k. A SWF is *stationary* if it implies constant relative valuations between individuals over time. The welfarist criteria that we consider are time invariant but not stationary. The functional form of $\mathcal{W}(\cdot)$ does not vary with the perspective of the assessment s^k. But the relative valuation of individuals may not remain constant across dates and histories. This is precisely encoded in the dependence of the Pareto weights $\alpha^i_{s^k}$ on the assessment perspective s^k. This violation of stationarity is what opens the door to time inconsistency. Indeed, Halevy (2015) showed that time invariance and stationarity would necessarily imply time consistency.

includes risk-bearing, includes much that appears to be motivated by distributional value judgments.

Our results provide a novel and practical formalization of these ideas, leveraging the formal definitions of redistribution, efficiency, and its sources introduced in Dávila and Schaab (2024).

Our paper also shares this theme with subsequent work in the public finance literature. The motivating example in Section 2 builds on Varian (1980) and Piketty and Saez (2013b). Much of the early work on optimal taxation (Mirrlees, 1971; Diamond, 1998; Saez, 2001) focuses on static environments and appeals to an equivalence between ex post redistribution and ex ante risk-sharing (behind the veil of ignorance). Our results highlight an important qualification: this equivalence holds only under a utilitarian SWF. Non-utilitarian SWFs — widely used in applied work (Atkinson, 1970; Atkinson and Stiglitz, 1976; Stern, 1976; Boskin and Sheshinski, 1978; Diamond, 1998; Saez, 2001; Piketty and Saez, 2012, 2013a; Saez and Stantcheva, 2016; Heathcote, Storesletten, and Violante, 2017) lead to a time consistency problem. The interplay between ex ante risk-sharing and ex post redistribution has been explored in the new dynamic public finance, for instance, by Farhi and Werning (2013) and Golosov, Troshkin, and Tsyvinski (2016). In particular, Farhi and Werning (2013) observe that a utilitarian SWF implies an equivalence between ex ante risk-sharing and ex post redistribution. We study and rationalize the time inconsistency not only of overall welfare assessments but also of efficiency and redistribution assessments.

Lastly, a literature on political economy studies time consistency focusing on the tension between efficiency and redistribution (Alesina and Rodrik, 1994; Persson and Tabellini, 1994; Krusell, Quadrini, and Ríos-Rull, 1997; Farhi, Sleet, Werning, and Yeltekin, 2012; Bisin, Lizzeri, and Yariv, 2015). Papers in this literature typically consider voting equilibria or coalition shifts rather than the welfarist perspective we focus on.

2 A Motivating Example

We illustrate the main ideas of the paper in a stylized dynamic model of taxation in which a distortionary labor income tax finances a lump-sum transfer.

2.1 Environment

We consider an economy with two periods, indexed by $t \in \{0, 1\}$, and a unit measure of households, indexed by $i \in [0, 1]$. Households are ex ante identical but face uninsurable idiosyncratic earnings risk in period 1.

Households. Household i's lifetime utility from a period-0 perspective is

$$V_0^i = u\left(c_0^i, \ell_0^i\right) + \beta \mathbb{E}\left[u\left(c_1^i, \ell_1^i\right)\right], \tag{1}$$

where $u(\cdot)$ is the instantaneous utility from consumption and hours worked and β is the discount factor. We denote i's lifetime utility from a period-1 perspective, after uncertainty is realized, by $V_1^i = u\left(c_1^i, \ell_1^i\right)$. Households face the budget constraints

$$c_0^i = w_0 \ell_0^i \qquad \text{and} \qquad c_1^i = (1-\tau)\, w_1 z_1^i \ell_1^i + T.$$

In period 0, household i supplies ℓ_0^i hours of labor at a wage w_0. In period 1, households draw idiosyncratic labor productivity z_1^i from a density $g(z)$, normalized such that $\int_0^1 z_1^i di = \int z g(z) dz = 1$. Finally, τ denotes a labor income tax and T is a lump-sum transfer.

Firms. A profit-maximizing representative firm produces the consumption good with labor each period according to the production function

$$Y_t = L_t, \tag{2}$$

where L_t denotes aggregate effective labor input.

Equilibrium. For a given tax policy τ, a competitive equilibrium comprises an allocation $\{Y_t, L_t, c_t^i, \ell_t^i\}_{i,t}$ and a transfer T such that (i) households optimize, (ii) firms maximize profits, (iii) the government's budget is balanced, so

$$T = \tau \int_0^1 z_1^i \ell_1^i \, di,$$

and (iv) the markets for goods and effective labor clear in each period:

$$Y_t = \int_0^1 c_t^i \, di \qquad \text{and} \qquad L_t = \int_0^1 z_t^i \ell_t^i \, di.$$

2.2 Optimal Policy

Social welfare from a period-t perspective is

$$W_t = \mathcal{W}\left(\left\{V_t^i\right\}_i\right) = \left(\int_0^1 \left(V_t^i\right)^{1-\phi} di\right)^{\frac{1}{1-\phi}}. \tag{3}$$

This social welfare function (SWF), often referred to as isoelastic (Atkinson, 1970), is widely used to parameterize the planner's concern for redistribution (Benabou, 2002; Heathcote, Storesletten, and Violante, 2017). In particular, (3) nests the utilitarian, Nash, and Rawlsian SWFs for the parameter values $\phi \in \{0, 1, \infty\}$, respectively. We contrast the ex ante optimal policy, set in period 0 to maximize W_0, with the ex post optimal policy, set in period 1 to maximize W_1.

Ex Ante Policy. From a period-0 perspective, the optimal tax trades off *risk-sharing* gains against production efficiency losses. Risk-averse households would like to smooth consumption across idiosyncratic earnings shocks, but are unable to do so because of the assumed market incompleteness. A labor tax with uniform transfers in period 1 substitutes for the missing insurance, achieving welfare gains from risk-sharing. At the same time, the tax distorts labor supply, creating production efficiency losses. As shown in the Appendix, the ex ante optimal tax balances both forces:

$$0 = \underbrace{-\tau \frac{dL_1}{d\tau}}_{\substack{\text{Production} \\ \text{Efficiency}}} + \underbrace{\mathbb{C}ov_i \left[\frac{u'\left(c_1^i\right)}{\mathbb{E}\left[u'\left(c_1^i\right)\right]}, -z_1^i \ell_1^i \right]}_{\text{Risk-Sharing}}. \tag{4}$$

Crucially, all social welfare functions agree on the ex ante optimal tax and the risk-sharing and production efficient assessments, irrespective of their concern for redistribution. In particular, (4) does not depend on ϕ.

Ex Post Policy. Once period 1 materializes, all uncertainty is resolved, leaving no further scope to insure idiosyncratic earnings risk. Would the planner choose the same policy if allowed to re-optimize at this point? Not necessarily. From a period-1 perspective, the optimal tax trades off *redistribution* gains against production efficiency losses. On the one hand, the planner has a desire to reallocate consumption to less well off individuals. On the other, the tax still distorts labor supply, creating production efficiency losses. As shown in the Appendix, the ex post optimal tax balances both forces:

$$0 = \underbrace{-\tau \frac{dL_1}{d\tau}}_{\substack{\text{Production} \\ \text{Efficiency}}} + \underbrace{\mathbb{C}ov_i \left[\frac{\left(V_1^i\right)^{-\phi} u'\left(c_1^i\right)}{\int_0^1 \left(V_1^i\right)^{-\phi} u'\left(c_1^i\right) di}, -z_1^i \ell_1^i \right]}_{\text{Redistribution}}. \tag{5}$$

The redistribution gains now critically depend on the underlying social welfare function through the inequality-aversion parameter ϕ. However, the production efficiency loss remains invariant to the choice of SWF.

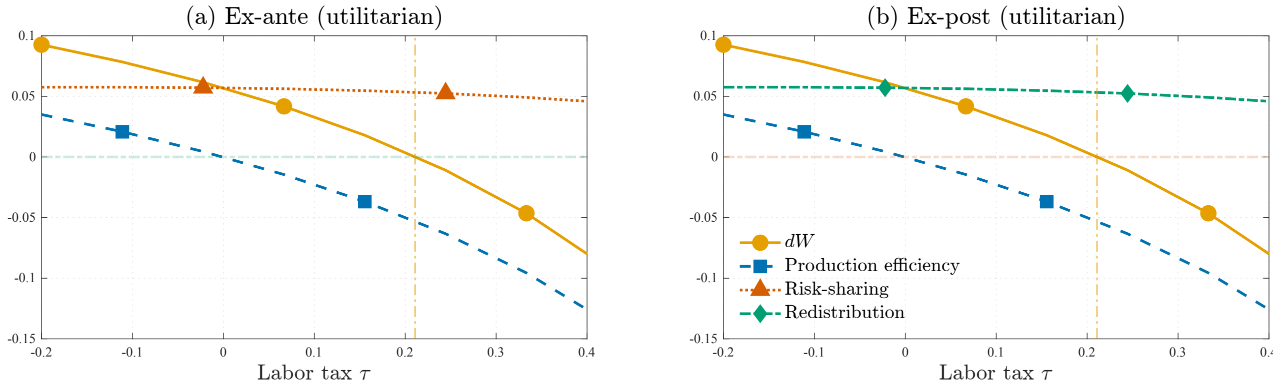

Figure 1: Utilitarian Welfare Assessments

Note. This figure displays ex ante (left panel) and ex post (right panel) welfare assessments of a marginal tax increase $d\tau$ under a utilitarian SWF ($\phi = 0$), as a function of the tax rate τ. The colored lines correspond to the overall welfare assessments (yellow) and their decompositions into gains and losses due to production efficiency (blue), risk-sharing (red), and redistribution (green). The ex ante and ex post optimal tax rates correspond to the τ at which the yellow lines cross 0, that is $\frac{dW_0}{d\tau} = 0$ in the left panel and $\frac{dW_1}{d\tau} = 0$ in right panel — illustrated by the dashed vertical lines. Whenever a line lies above (below) 0 at a given τ, the planner associates gains (losses) with a marginal tax increase due to that particular motive. We set $\beta = 0.96$ and assume isoelastic preferences, $u(c, \ell) = \frac{1}{1-\gamma}c^{1-\gamma} - \frac{1}{1+\eta}\ell^{1+\eta} + \bar{u}$, with $\gamma = \eta = 2$, and $\bar{u} = 3$ to ensure that $V_1^i > 0$ for all households. The distribution of idiosyncratic shocks is log-normal with mean 1 and standard deviation 0.4.

2.3 Lessons for Time Inconsistency

Comparing the ex ante and ex post welfare assessments in 4 and 5 allows us to illustrate four key ideas that we develop in the rest of the paper. Figure 1 illustrates the first three.

1. *(Efficiency and redistribution assessments are time inconsistent)* Assessments of efficiency and redistribution gains and losses shift over time and are therefore time inconsistent. Ex ante, the planner perceives that a tax increase yields risk-sharing gains (red line) for tax rates up to and exceeding 40%, while the there are no redistribution gains (green line), as shown in the left panel of Figure 1. Ex post, however, risk-sharing no longer plays any role and all gains are solely attributed to redistribution, as shown in the right panel of Figure 1. Ex ante, the planner therefore attributes all gains and losses to efficiency; whereas ex post, the planner trades off an efficiency loss against a redistribution gain. The main result of this paper shows that this time inconsistency of efficiency and redistribution assessments is a general feature of heterogeneous agent incomplete markets economies. What the planner perceives as the sources of welfare gains and losses invariably changes over time in these environments.

2. *(Production efficiency assessment are time consistent)* The planner perceives the same

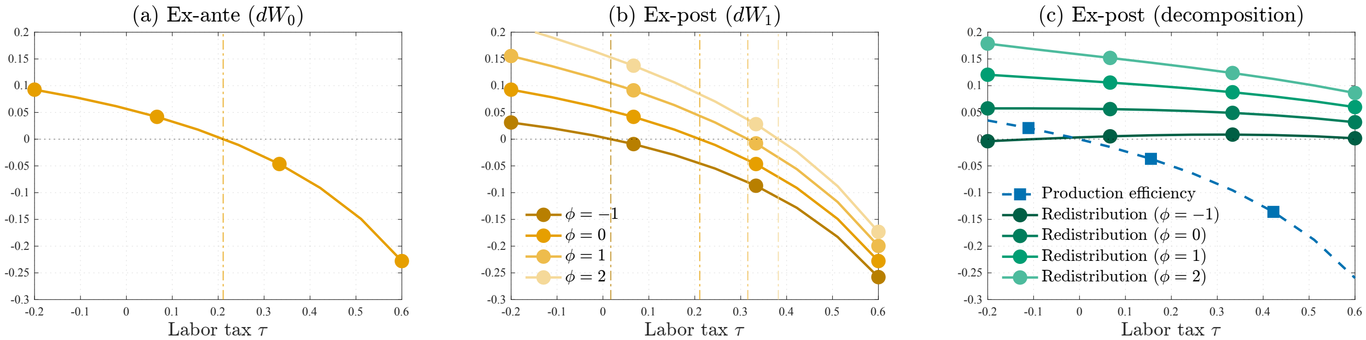

Figure 2: Time Inconsistency of Welfare Assessments

Note. This figure displays ex ante (left panel) and ex post (middle panel) welfare assessments of a marginal tax increase $d\tau$ under alternative SWFs. It also displays (right panel) a decomposition of the ex post assessments. Each line corresponds to a particular SWF: inequality-loving ($\phi = -1$), utilitarian ($\phi = 0$), Nash ($\phi = 1$), and inequality-averse ($\phi = 2$). All SWFs agree on the ex ante welfare assessment $\frac{dW_0}{d\tau}$ as well as on the ex post production efficiency assessment: There is a single yellow line in the left panel and a single blue line in the right panel. Different SWFs disagree on the ex post welfare assessment in middle panel because they disagree on the ex post redistribution in the right panel. Figure 2 uses the same parametrization as Figure 1.

production efficiency loss $\tau \frac{dL_1}{d\tau}$ from taxation ex ante and ex post. From both perspectives, production efficiency considerations suggest an optimal tax of 0. The blue lines in both panels of Figure 1 are identical and cross 0 at $\tau = 0$. We show in Section 4.2 that this insight generalizes: While risk-sharing and redistribution assessments change over time, assessments of static production efficiency gains and losses are time-consistent. This result, which is not special to the utilitarian SWF but holds for all ϕ, underscores that the source of the inconsistency is interpersonal.

3. *(Utilitarian welfare assessments are time consistent)* The overall welfare assessment $\frac{dW_0}{d\tau}$ is time-consistent under the utilitarian SWF ($\phi = 0$), despite the inconsistency of the risk-sharing and redistribution assessments. The yellow lines in both panels of Figure 1 are identical and cross 0 at $\tau = 0.22$. Hence, even though the justification for taxation varies starkly between the ex ante and ex post perspectives, the optimal policy remains time-consistent under the utilitarian SWF. Formally, notice that the covariance term in (5) becomes identical to that in (4) for $\phi = 0$. This consistency is special to the utilitarian SWF and does not hold for any $\phi \neq 0$ as we show next.

Figure 2 contrasts welfare assessments across different social welfare functions.This implies our fourth and final observation:

4. *(Non-utilitarian welfare assessments are time inconsistent).* Only the utilitarian SWF ($\phi = 0$) values ex post redistribution the same as ex ante risk-sharing. Non-utilitarian

SWFs perceive redistribution gains ex-post higher or lower than the ex ante risk-sharing gains. As ϕ increases, the planner perceives the redistribution gains from the social transfer ex post as more valuable than the risk-sharing gains ex ante (Panel c). The ex post optimal tax thus increases with ϕ and only coincides with the ex ante optimal policy for $\phi = 0$. This leads to time inconsistency.

3 Environment

Our notation closely follows Chapter 8 of Ljungqvist and Sargent (2018). We consider an economy populated by $I \geq 1$ individuals, indexed by $i \in \mathcal{I} = \{1, \ldots, I\}$. At each date $t \in \{0, \ldots, T\}$, where $T \leq \infty$, there is a realization of a stochastic event s_t, where s_0 is predetermined. We denote the history of events up to date t by $s^t = (s_0, s_1, \ldots, s_t)$, and the conditional probability of observing history s^t starting from history s^k by $\pi_t\left(s^t | s^k\right)$. We use the term *perspective* to refer to the history at which a planner conducts a welfare assessment.

Preferences. To rule out time consistency problems due to inconsistent individual preferences in the spirit of Strotz (1956), we assume that all individuals have expected utility preferences with exponential discounting. Formally, individual i's continuation lifetime utility from the perspective of history s^k, $V^i_{s^k}$, is given by

(Preferences) $$V^i_{s^k} = \sum_{t \geq k} \beta^{t-k} \sum_{s^t \geq s^k} \pi_t\left(s^t | s^k\right) u\left(c^i_t\left(s^t\right)\right), \tag{6}$$

where $c^i_t(s^t)$ denotes individual i's consumption at history s^t, and where $\sum_{s^t \geq s^k}$ indicates summation over all histories that follow from s^k. We assume that $u(\cdot)$ is well-behaved and Inada conditions apply.

To simplify the exposition, we assume in the main text that all individuals have identical preferences. We allow for heterogeneity in flow utilities, discount factors, and beliefs in the Online Appendix.

Social Welfare Function. We study welfare assessments for welfarist planners, those who evaluate outcomes using a social welfare function based on individual utilities.[5] Formally, social

[5]The welfarist approach is widely used because it is Paretian, that is, it can be used to conclude that every Pareto-improving perturbation is desirable, and because nonwelfarist approaches violate the Pareto principle (Kaplow and Shavell, 2001).

welfare from the perspective of history s^k, W_{s^k}, is given by

$$\text{(Social Welfare Function)} \qquad W_{s^k} = \mathcal{W}\left(V^1_{s^k}, \ldots, V^i_{s^k}, \ldots, V^I_{s^k}\right), \tag{7}$$

where individual lifetime utilities $V^i_{s^k}$ are defined in (6). The critical restriction encoded in (7) is that $\mathcal{W}(\cdot)$ is a time- and history-invariant function of individual lifetime utilities as perceived at the time of the assessment, but of nothing else. We define individual i's Pareto weight from the perspective of history s^k as

$$\text{(Pareto Weight)} \qquad \alpha^i_{s^k} = \frac{\partial \mathcal{W}(\cdot)}{\partial V^i_{s^k}} > 0. \tag{8}$$

Note that Pareto weights depend on the perspective because continuation lifetime utilities vary across histories, not because the SWF $\mathcal{W}(\cdot)$ varies over time or across histories.

It is useful to distinguish between linear and non-linear social welfare functions. We refer to — equal- or unequal-weighted — linear SWFs as *utilitarian*. When $\mathcal{W}(\cdot)$ is linear, $W_{s^k} = \sum_i \alpha^i V^i_{s^k}$, and the Pareto weights $\alpha^i_{s^k} = \alpha^i$ are time- and history-invariant. This is no longer the case when $\mathcal{W}(\cdot)$ is non-linear.[6]

Consumption Functions from Particular Perspectives. We assume that individual consumption allocations depend on perturbation parameters $\boldsymbol{\theta} = \{\theta_t(s^t)\}_{t,s^t}$, where each $\theta_t(s^t)$ is a scalar. Concretely, we assume that individual i's consumption at history s^t from the perspective of some other history s^k can be expressed as a smooth consumption function

$$\text{(Consumption Function)} \qquad \mathcal{C}^i_{t|s^k}\left(s^t, \boldsymbol{\theta}\right). \tag{9}$$

The mapping between individual consumption and $\boldsymbol{\theta}$ may arise from a competitive equilibrium or other equilibrium notions. Since our results do not depend on the specific micro-foundation of $\mathcal{C}^i_{t|s^k}(\cdot)$, we take the dependence of consumption on the perturbation parameters $\boldsymbol{\theta}$ as given. It is natural to think of $\boldsymbol{\theta}$ as policy instruments, such as taxes/transfers (fiscal policy) or nominal interest rates (monetary policy), but $\boldsymbol{\theta}$ may also represent exogenous primitives such as productivity or endowments, or simply index allocations directly chosen by a planner.[7] To simplify the exposition, we typically suppress the dependence of $\mathcal{C}^i_{t|s^k}(\cdot)$, $V^i_{s^k}$, and W_{s^k} on $\boldsymbol{\theta}$.

[6]Important examples of non-linear SWFs include isoelastic (Atkinson, 1970), Nash (Nash, 1950; Kaneko and Nakamura, 1979), and Rawlsian (Rawls, 1971). Non-linear social welfare functions are often used in applied work in macroeconomics and public finance (Atkinson, 1970; Atkinson and Stiglitz, 1976; Stern, 1976; Boskin and Sheshinski, 1978; Diamond, 1998; Saez, 2001; Benabou, 2002; Piketty and Saez, 2012, 2013a; Saez and Stantcheva, 2016; Heathcote, Storesletten, and Violante, 2017).

[7]Optimizing over $\boldsymbol{\theta}$ taking the consumption function as given corresponds to a *dual* approach in optimal policy (Dávila and Schaab, 2023a; Auclert, Cai, Rognlie, and Straub, 2024).

Critically, the consumption function (9) depends on the perspective s^k, because the response of consumption to a perturbation at a particular history may differ depending on when the perturbation is assessed; formally,

$$\frac{\partial \mathcal{C}^i_{t|s^k}\left(s^t, \boldsymbol{\theta}\right)}{\partial \boldsymbol{\theta}} \neq \frac{\partial \mathcal{C}^i_{t|s^\ell}\left(s^t, \boldsymbol{\theta}\right)}{\partial \boldsymbol{\theta}}.$$

For instance, the consumption response at history s^k to a contemporaneous tax change at that history s^k differs across perspectives: from the perspective s^k, agents take the tax as given, while from an earlier s^0 perspective, they can adjust behavior in anticipation, altering equilibrium allocations between date 0 and s^k. We formalize these ideas in Appendix C, which presents a general, microfounded derivation of consumption functions (9) using a sequence-space representation (Auclert, Bardóczy, Rognlie, and Straub, 2021).

4 Time Inconsistency

We proceed in three steps. First, we examine the consistency of overall welfare assessments. Next, we do the same for efficiency assessments. Finally, we analyze the consistency of the sources of efficiency.

4.1 Inconsistency of Welfare Assessments

We begin by characterizing the conditions under which an overall welfare assessment is time-consistent. We define time consistency as follows.

Definition. *(Time Consistency of Welfare Assessments) The welfare assessment of a perturbation* $d\theta_t\left(s^t\right)$*, from the perspective of* s^0, $\frac{dW_{s^0}}{d\theta_t(s^t)}$*, is time-consistent if later assessments* $\frac{dW_{s^k}}{d\theta_t(s^t)}$ *share the same sign for all histories* $s^0 < s^k \leq s^t$*. The welfare assessment is time inconsistent if there is a history* s^k *for which this does not hold.*

A welfare assessment is therefore time-consistent if all later assessments of the same perturbation have the same sign and hence agree on whether the perturbation is desirable or not. If $\boldsymbol{\theta}$ is a policy, then consistency in the sign of the welfare assessment ensures the consistency of the optimal policy.[8]

[8]For a policy $\boldsymbol{\theta}^*$ to be optimal from the perspective of s^0, no feasible policy can yield welfare gains, that is, $dW_{s^0}/d\theta_t(s^t) = 0$ for all $\theta_t(s^t)$ at $\boldsymbol{\theta} = \boldsymbol{\theta}^*$. Therefore, if ex post assessments agree on the sign with the ex ante assessment for all perturbations, then $\boldsymbol{\theta}^*$ is a time-consistent optimal policy. This explains why our definition of consistency is based on the sign of the assessments and why we define time consistency in marginal form.

Remark. (Perturbation to $\theta_1(\bar{s}^1)$: θ as placeholder) For ease of exposition, in the main text we always consider a perturbation $d\theta_1(\bar{s}^1)$ at a particular history $\bar{s}^1$ at date 1, which we abbreviate as "$d\theta$". We call $\frac{dW_{s^0}}{d\theta} = \frac{dW_{s^0}}{d\theta_1(\bar{s}^1)}$ the *ex ante* assessment and $\frac{dW_{\bar{s}^1}}{d\theta} = \frac{dW_{\bar{s}^1}}{d\theta_1(\bar{s}^1)}$ the *ex post* assessment at history $\bar{s}^1$. Time consistency then requires that $\frac{dW_{s^0}}{d\theta}$ and $\frac{dW_{\bar{s}^1}}{d\theta}$ have the same sign since $\bar{s}^1$ is the only relevant history for reassessing $d\theta_1(\bar{s}^1)$.

4.1.1 Benchmark: Single Agent

As a benchmark, we first study the consistency of welfare assessments with a single agent, when $I = 1$. Even when individual preferences are consistent in the sense of Strotz (1956), three distinct forms of time inconsistency can still emerge in a single agent environment.

Proposition 1. *(Inconsistency of Welfare Assessments: Single Agent) In a single individual economy ($I = 1$), the ex ante welfare assessment of perturbation $d\theta$ is time inconsistent when one of the following holds: The perturbation affects*

(i) (Time goes by) *the consumption allocation at date* 0*,*

(ii) (Road not taken) *the consumption allocation at a history $s^1 \neq \bar{s}^1$ at date* 1*,*

(iii) (Consumption function shifts) *the consumption function at history $\bar{s}^1$ (and subsequent histories) from perspective s^0 differently than from perspective $\bar{s}^1$, so $\frac{\partial \mathcal{C}^i_{1|s^0}(\bar{s}^1)}{\partial\theta} \neq \frac{\partial \mathcal{C}^i_{1|\bar{s}^1}(\bar{s}^1)}{\partial\theta}$.*

To interpret Proposition 1, we can write the ex ante welfare assessment as

$$\frac{dW_{s^0}}{d\theta} = \sum_i \alpha^i_{s^0} \frac{dV^i_{s^0}}{d\theta} = \alpha^i_{s^0} \sum_{t\geq 0} \beta^t \sum_{s^t \geq s^0} \pi_t\left(s^t|s^0\right) u'\left(c^i_t\left(s^t\right)\right) \frac{\partial \mathcal{C}^i_{t|s^0}\left(s^t\right)}{\partial\theta},$$

where the second equality follows because $I = 1$, and the ex post assessment at history $\bar{s}^1$ as

$$\frac{dW_{\bar{s}^1}}{d\theta} = \sum_i \alpha^i_{\bar{s}^1} \frac{dV^i_{\bar{s}^1}}{d\theta} = \alpha^i_{\bar{s}^1} \sum_{t\geq 1} \beta^{t-1} \sum_{s^t \geq \bar{s}^1} \pi_t\left(s^t|\bar{s}^1\right) u'\left(c^i_t\left(s^t\right)\right) \frac{\partial \mathcal{C}^i_{t|\bar{s}^1}\left(s^t\right)}{\partial\theta}.$$

Comparing these two expressions reveals three forms of time inconsistency. First, the ex ante assessment includes a date-0 term, whereas the ex post assessment sums over periods from date 1 onwards. If the perturbation affects date-0 consumption, this affects the ex ante but not the ex post assessment, and we refer to this form of time inconsistency as "*time goes by*". Second, the ex ante assessment includes terms for histories $s^1 \neq \bar{s}^1$. Ex post, these unrealized histories are no longer considered, and we refer to this source of form inconsistency as the "*road not taken*". Third, the perturbation may affect the consumption allocation at $\bar{s}^1$ (and subsequent histories) differently depending on the perspective, by changing the consumption function and possibly influencing state variables like capital. Note that the individual weights $\alpha^i_{s^0}$ and $\alpha^i_{\bar{s}^1}$ are strictly positive ex-ante and ex-post and can therefore not be a source of inconsistency. Although formalized differently, the three forms of inconsistency identified in Proposition 1 are already present in Kydland and Prescott (1977).

4.1.2 Heterogeneous Agents

After reviewing the single agent benchmark, we now study the time consistency of welfare assessments with heterogeneous agents. Proposition 2 shows that, in the presence of heterogeneity, interpersonal welfare comparisons emerge as a new form of time inconsistency.

Proposition 2. *(Inconsistency of Welfare Assessments: Heterogeneous Agents) In a heterogeneous agent economy ($I > 1$), the ex ante welfare assessment of perturbation $d\theta$ is time inconsistent when either one of (i) – (iii) in Proposition 1 holds, or*

(iv) (Interpersonal welfare comparisons) *the ratio of ex-ante to ex-post Pareto weights differs across individuals, so $\frac{\alpha^i_{\bar{s}^1}}{\alpha^i_{s^0}} \neq \frac{\alpha^j_{\bar{s}^1}}{\alpha^j_{s^0}}$ for at least two individuals i and j.*

Proposition 2 establishes that a distinct form of time inconsistency — arising from interpersonal welfare comparisons — emerges with heterogeneous agents. To isolate this new mechanism, we abstract from conditions (i) – (iii) going forward and focus on condition (iv), under a "No-Kydland–Prescott" assumption.

Assumption. *(No-Kydland-Prescott) We assume that (i) $\frac{\partial \mathcal{C}^i_{0|s^0}(s^0)}{\partial \theta} = 0$; (ii) $\frac{\partial \mathcal{C}^i_{1|s^0}(s^1)}{\partial \theta} = 0$ for $s^1 \neq \bar{s}^1$; and (iii) $\frac{\partial \mathcal{C}^i_{t|s^0}(s^t)}{\partial \theta} = \frac{\partial \mathcal{C}^i_{t|\bar{s}^1}(s^t)}{\partial \theta}$, for $s^t \geq \bar{s}^1$, $\forall i$.*

In No-Kydland-Prescott scenarios, we can write the ex ante welfare assessment of the perturbation as

$$\frac{dW_{s^0}}{d\theta} = \underbrace{\beta\pi\left(\bar{s}^1\right)\left(\frac{1}{I}\sum_i \frac{\alpha^i_{s^0}}{\alpha^i_{\bar{s}^1}}\right)}_{>0} \frac{dW_{\bar{s}^1}}{d\theta} + \beta\pi\left(\bar{s}^1\right)\mathbb{C}ov^{\Sigma}_i\left[\frac{\alpha^i_{s^0}}{\alpha^i_{\bar{s}^1}}, \alpha^i_{\bar{s}^1}\frac{dV^i_{\bar{s}^1}}{d\theta}\right], \tag{10}$$

where $\mathbb{C}ov^{\Sigma}_i[\cdot,\cdot] = I \cdot \mathbb{C}ov_i[\cdot,\cdot]$ denotes a cross-sectional covariance-sum. This expression decomposes the ex ante welfare assessment $\frac{dW_{s^0}}{d\theta}$ into a term proportional to the ex post assessment $\frac{dW_{\bar{s}^1}}{d\theta}$ and a cross-sectional covariance term. Because the coefficient on $\frac{dW_{\bar{s}^1}}{d\theta}$ is strictly positive, ex ante and ex post assessments always agree on sign when the covariance term is zero. Hence, time inconsistency arises only when the covariance term is non-zero, and large enough to overturn the sign of $\frac{dW_{\bar{s}1}}{d\theta}$. This can only occur when the ratio of ex ante and ex post Pareto weights $\frac{\alpha^i_{s0}}{\alpha^i_{\bar{s}1}}$ differs across individuals, precisely condition (iv) of Proposition 2. Intuitively, the ratio $\frac{\alpha^i_{s0}}{\alpha^i_{\bar{s}1}}$ differs across individuals when the planner's valuation of individual gains in utils shifts over time.

Are there social welfare functions for which the ratio of Pareto weights remains constant across individuals, ensuring time consistency of welfare assessments? Only utilitarian social welfare functions ensure time consistency.

Corollary. *(Time Consistency under Utilitarian SWFs) Only utilitarian social welfare functions ensure time consistency due to interpersonal welfare comparisons.*

When the social welfare function $\mathcal{W}(\cdot)$ is utilitarian and therefore linear, individual Pareto weights are constant across histories for each individual, so $\alpha^i = \alpha^i_{s^t}$ and $\frac{\alpha^i_{s0}}{\alpha^i_{\bar{s}1}} = 1$, $\forall i$. The covariance term in (10) therefore vanishes and the welfare assessment $\frac{dW_{s0}}{d\theta}$ is time-consistent in No-Kydland-Prescott scenarios. The (equal- or unequal-weighted) utilitarian SWF is the only one for which this is the case. This result goes back to Zuber (2011) who showed, using a different approach, that welfare assessments based on additively separable social welfare functions are time-consistent as long as individuals have a common discount factor. For more general SWFs, the weight the planner assigns to a particular individual changes endogenously with the allocation because in that case the Pareto weight $\alpha^i_{s^t}$ is a function of lifetime utilities $V^i_{s^t}$ evaluated from the perspective of s^t.[9]

4.2 Inconsistency of Efficiency Assessments

It is well-understood that a welfare assessment combines efficiency and redistribution assessments. It is thus natural to also investigate whether these are time-consistent. Our main result shows that efficiency and redistribution assessments are — under natural circumstances — time inconsistent in heterogeneous agent economies with incomplete markets, even when

[9] An interesting question is whether the analysis of Kydland and Prescott (1977) already nests Proposition 2 for the heterogeneous agent case after an appropriate relabeling. Suppose we relabel their "x_t" as the vector of all individuals' consumption. The question then becomes whether their key term "$\partial X_1/\partial \pi_2$" already captures condition (iv) of Proposition 2. The answer is no: with heterogeneous agents, their term "$\partial S/\partial x_2$" differs across ex ante and ex post assessments, which is precisely the content of Proposition 2. Kydland and Prescott (1977) implicitly rule out heterogeneity by imposing this term to be constant across perspectives. Utilitarian social welfare functions are the special case where "$\partial S/\partial x_2$" has the same sign ex ante and ex post. A useful restatement of the Corollary would therefore be: Only under the utilitarian social welfare function does the analysis of Kydland and Prescott (1977) extend to heterogeneous agent environments by relabeling their "x_t" as the vector of all individuals' consumption.

the overall welfare assessment is time-consistent. Therefore, even when a welfare-maximizing planner assesses policies consistently, the rationale that justifies a particular policy — whether welfare gains are due to efficiency or redistribution considerations — changes over time.

4.2.1 Normalized Welfare Gains

A welfare assessment

$$\frac{dW_{s^t}}{d\theta} = \sum_i \alpha^i_{s^t} \frac{dV^i_{s^t}}{d\theta}$$

cannot directly reveal how a planner makes tradeoffs in meaningful units across individuals, because individual utilities are ordinal (measured in utils) and not inherently comparable. To make comparisons meaningful, individual gains must be expressed in a common unit — a *welfare numeraire*. Formally, individual i's normalized welfare gains or willingness-to-pay for a perturbation, expressed in units of the welfare numeraire from the perspective of history s^t, are

$$\frac{\frac{dV^i_{s^t}}{d\theta}}{\lambda^i_{s^t}},$$

where $\lambda^i_{s^t}$ is an individual normalizing factor that converts utils into units of the chosen numeraire.[10] The choice of welfare numeraire — and hence of $\lambda^i_{s^t}$ — typically depends on the perspective from which the assessment is made, a dependence central to our results.

We can also express the overall welfare assessment in units of the welfare numeraire, dividing by $\frac{1}{I}\sum_i \alpha^i_{s^t}\lambda^i_{s^t}$. Hence, a *normalized* welfare assessment expressed in units of the welfare numeraire from the perspective of history s^t is

$$\frac{dW^\lambda_{s^t}}{d\theta} = \frac{\frac{dW_{s^t}}{d\theta}}{\frac{1}{I}\sum_i \alpha^i_{s^t}\lambda^i_{s^t}} = \sum_i \omega^i_{s^t} \frac{\frac{dV^i_{s^t}}{d\theta}}{\lambda^i_{s^t}}, \tag{11}$$

where $\omega^i_{s^t} = \frac{\alpha^i_{s^t}\lambda^i_{s^t}}{\frac{1}{I}\sum_i \alpha^i_{s^t}\lambda^i_{s^t}}$ is a (normalized) individual weight, where $\frac{1}{I}\sum_i \omega^i_{s^t} = 1$. The individual weight $\omega^i_{s^t}$ encodes how a planner trades off welfare gains across individuals. For example, if $\omega^i_{s^t} = 1.3$, the planner perceives providing individual i with a 1% increase in units of the welfare numeraire as equivalent to providing everyone a 1.3% increase. If $\frac{dW^\lambda_{s^t}}{d\theta} = 1.1$, the planner perceives the perturbation from the perspective of s^t equal to distributing 1.1 units of welfare numeraire equally across individuals.

[10]The dimension of $\lambda^i_{s^t}$ is $\dim\left(\lambda^i_{s^t}\right) = \frac{\text{history-}s^t \text{ individual } i \text{ utils}}{\text{history-}s^t \text{ welfare numeraire}}$, so $\dim\left(\frac{\frac{dV^i_{s^t}}{d\theta}}{\lambda^i_{s^t}}\right) = \frac{\text{history-}s^t \text{ welfare numeraire}}{\text{units of } \theta}$, $\forall i$. The only restriction when choosing the welfare numeraire is that $\lambda^i_{s^t}$ must be strictly positive for all individuals.

4.2.2 Efficiency vs. Redistribution

Given a welfare numeraire, a normalized welfare assessment admits a unique decomposition into Kaldor-Hicks efficiency and redistribution (Dávila and Schaab, 2024), given by

$$\frac{dW_{s^t}^{\lambda}}{d\theta} = \sum_i \omega_{s^t}^i \frac{\frac{dV_{s^t}^i}{d\theta}}{\lambda_{s^t}^i} = \underbrace{\sum_i \frac{\frac{dV_{s^t}^i}{d\theta}}{\lambda_{s^t}^i}}_{\Xi_{s^t}^E \text{(Efficiency)}} + \underbrace{\mathbb{C}ov_i^{\Sigma}\left[\omega_{s^t}^i, \frac{\frac{dV_{s^t}^i}{d\theta}}{\lambda_{s^t}^i}\right]}_{\Xi_{s^t}^{RD} \text{(Redistribution)}}. \tag{12}$$

The efficiency component $\Xi_{s^t}^E$ corresponds to Kaldor-Hicks efficiency gains — the sum of individual willingness to pay in units of the welfare numeraire — from the perspective of s^t. This is a useful notion for at least three reasons. First, it is invariant to preference-preserving transformations of utilities, unlike the overall welfare assessment. Second, it is independent of the social welfare function, which only affects the redistribution component through individual weights. Third, it satisfies the compensation principle; perturbations with $\Xi_{s^t}^E > 0$ can be turned into Pareto improvements if compensating transfers in units of the welfare numeraire are feasible. For this reason, Kaldor-Hicks efficiency is useful to define Pareto frontiers.

The redistribution component $\Xi_{s^t}^{RD}$ captures the equity concerns embedded in a particular social welfare function: $\Xi_{s^t}^{RD}$ is positive when the individuals relatively favored in a perturbation are those relatively preferred by the planner, that is, have a higher individual weight $\omega_{s^t}^i$.

4.2.3 Inconsistency of Efficiency Assessments

We now characterize the time consistency of efficiency assessments, which we define as follows.[11]

Definition. *(Time Consistency of Efficiency Assessments) The efficiency assessment of a perturbation $d\theta_t(s^t)$, from the perspective of s^0, $\Xi_{s^0}^E$, is time-consistent if later assessments $\Xi_{s^k}^E$ share the same sign for all histories $s^0 < s^k \leq s^t$. The efficiency assessment is time inconsistent if there is a history s^k for which this does not hold.*

As before, we consider a perturbation $d\theta_1(\bar{s}^1)$, denoted by $d\theta$, and compare ex ante and ex post assessments:

$$\Xi_{s^0}^E = \sum_i \frac{\frac{dV_{s^0}^i}{d\theta}}{\lambda_{s^0}^i} \qquad \text{and} \qquad \Xi_{\bar{s}^1}^E = \sum_i \frac{\frac{dV_{\bar{s}^1}^i}{d\theta}}{\lambda_{\bar{s}^1}^i}.$$

[11] We use identical definitions of time inconsistency for redistribution assessments, and for the assessments of the different sources of efficiency in Section 4.3.

While Pareto weights play no role in the consistency of efficiency assessments, the choice of welfare numeraires — the units in which individual gains are expressed in each assessment — does. Proposition 3 shows that changes in individuals' valuations for these numeraires can give rise to time inconsistency in efficiency assessments.

Proposition 3. *(Inconsistency of Efficiency Assessments) In a heterogeneous agent economy ($I > 1$), the ex ante efficiency assessment of perturbation $d\theta$ is time inconsistent when either one of (i) – (iii) in Proposition 1 holds or:*

(v) (Numeraire inconsistency) *the ratio of individual normalizing factors differs across individuals, so $\frac{\lambda^i_{\bar{s}^1}}{\lambda^i_{s^0}} \neq \frac{\lambda^j_{\bar{s}^1}}{\lambda^j_{s^0}}$ for at least two individuals i and j.*

This result shows that efficiency assessments with heterogeneous agents can be time inconsistent even when Pareto weights $\alpha^i_{s^t}$ remain fixed over time. Hence, even when the overall welfare assessment is time-consistent — as in the utilitarian case — the attribution of welfare gains and losses to efficiency and redistribution changes over time. This new form of inconsistency in efficiency assessments stems not from a change in the social valuation of individual utilities, but from a change in the units used to aggregate individual consumption-equivalents gains from the perturbation.

To interpret this proposition, note that in No-Kydland-Prescott scenarios, we can write the ex ante efficiency assessments of the perturbation as

$$\Xi^E_{s^0} = \underbrace{\beta\pi\left(\bar{s}^1\right)\left(\frac{1}{I}\sum_i \frac{\lambda^i_{\bar{s}^1}}{\lambda^i_{s^0}}\right)}_{>0} \Xi^E_{\bar{s}^1} + \beta\pi\left(\bar{s}^1\right)\mathbb{C}ov^{\Sigma}_i\left[\frac{\lambda^i_{\bar{s}^1}}{\lambda^i_{s^0}}, \frac{\frac{dV^i_{\bar{s}^1}}{d\theta}}{\lambda^i_{\bar{s}^1}}\right]. \tag{13}$$

This expression — which resembles (10) — decomposes the ex ante efficiency assessment $\Xi^E_{s^0}$ into a term proportional to the ex post assessment $\Xi^E_{\bar{s}^1}$ and a cross sectional covariance term. Because the coefficient on $\Xi^E_{\bar{s}^1}$ is strictly positive, time inconsistency arises only when the covariance term is non-zero, and large enough to overturn the sign of $\Xi^E_{\bar{s}^1}$. This can only occur when the ratio of ex ante and ex post normalizing factors differs across individuals, precisely condition (v) of Proposition 3. Intuitively, ratios $\frac{\lambda^i_{\bar{s}^1}}{\lambda^i_{s^0}}$ that differ across individuals imply that individuals value gains utility in units of the welfare numeraire differently over time.

The inconsistency of welfare assessments — which we studied in Section 4.1 — stemmed from cross-sectional dispersion in ratios of Pareto weights, $\frac{\alpha^i_{s^0}}{\alpha^i_{\bar{s}^1}}$, which captures shifts in the valuation of individual's utility gains across perspectives. By contrast, the inconsistency of efficiency assessments depends on the cross-sectional dispersion of normalizing factors, $\lambda^i_{\bar{s}^1}/\lambda^i_{s^0}$, which captures shifts in the valuation that an individual attaches to the ex ante and ex post welfare numeraires. The ratio $\lambda^i_{\bar{s}^1}/\lambda^i_{s^0}$ captures the relative valuation of individual i for 1 unit of ex ante welfare numeraire in units of ex post welfare numeraire. Intuitively, when this ratio is large, the individual's willingness to pay for the same utility gain is higher when expressed in units of ex ante welfare numeraire. The covariance term in equation (13) is therefore positive when those individuals with a large willingness to pay ex post — so $\frac{1}{\lambda^i_{\bar{s}^1}}\frac{dV^i_{\bar{s}^1}}{d\theta}$ is large — also have a higher relative valuation for welfare gains ex post — so $\lambda^i_{\bar{s}^1}/\lambda^i_{s^0}$ is large. Hence, the corollary to Proposition 2 no longer applies to efficiency assessments: Even though a utilitarian planner puts the same weight on individual gains and losses in utils ex ante and ex post, the individual's own valuation of the perturbation in units of the welfare numeraires may differ regardless.

4.2.4 Impossibility Result

We now present what is arguably the central result of this paper: an impossibility theorem for the time consistency of efficiency assessments in heterogeneous agent environments. In particular, we show that it is generally not possible to make time-consistent efficiency assessments that satisfy the compensation principle in heterogeneous agent incomplete markets environments.

Forward- and Backward-looking Numeraires. To formally state this result, we first must distinguish between forward- and backward-looking welfare numeraires.

Definition. *(Forward-and Backward-Looking Numeraires) A welfare numeraire is forward-looking from the perspective of history s^k if it exclusively a function of consumption at s^k or at continuation histories, that is, individual i's normalizing factor $\lambda^i_{s^k}$ is a function only of $\{c^i_t(s^t)\}_{t\geq k, s^t \geq s^k}$. A welfare numeraire is backward-looking from the perspective of history s^k if it is also a function of consumption at past or date-k unrealized histories*

Forward-looking welfare numeraires are the ones routinely used to assess policies in practice. They nest, for instance, the practice of using net present values or perpetual consumption (Lucas, 1987) to express welfare gains.[12] Expressing individuals' willingness to pay in these forward-looking units is useful for two reasons.

[12] The most common welfare numeraires in practice — all forward-looking — are, once expressed from the $\bar{s}^1$-perspective: i) contemporaneous consumption, with $\lambda^i_{\bar{s}^1} = u'\left(c^i_t\left(\bar{s}^1\right)\right)$ ii) perpetual unit consumption, with $\lambda^i_{\bar{s}^1} = \sum_{t\geq 1}\sum_{s^t\geq \bar{s}^1}\beta^{t-1}\pi\left(s^t|\bar{s}^1\right)u'\left(c^i_t\left(s^t\right)\right)$; and iii) perpetual aggregate consumption, with $\lambda^i_{\bar{s}^1} = \sum_{t\geq 1}\sum_{s^t\geq \bar{s}^1}\beta^{t-1}\pi\left(s^t|\bar{s}^1\right)u'\left(c^i_t\left(s^t\right)\right)c_t\left(s^t\right)$, where $c_t\left(s^t\right) = \sum_i c^i_t\left(s^t\right)$.

First, with forward-looking numeraires, if an efficiency assessment from perspective s^t is positive ($\Xi^E_{s^t} > 0$), the compensation principle applies: hypothetical transfers of that numeraire can turn the perturbation into a Pareto improvement, because such transfers can occur contemporaneously or in the future. Second, no such positive efficiency assessments exist at Pareto efficient allocations for any forward-looking numeraire, ensuring that efficiency assessments are indeed useful to define Pareto frontiers.

By contrast, backward-looking numeraires lack these properties. A positive efficiency assessment from perspective s^t does not mean that the compensation principle holds, in the sense that compensating transfers would have had to occur in the past, something infeasible once s^t is reached. More importantly, it may be possible to find positive efficiency assessments even at Pareto efficient allocations, because nothing ensures that agents can freely trade the backward-looking numeraire once the history s^t has been reached.

Impossibility of Time-Consistent Efficiency Assessments. The considerations just discussed suggest that for efficiency assessments to be useful, they must be based on a forward-looking welfare numeraire. But forward-looking numeraires invite time consistency problems, as formalized in the following Theorem.

Theorem. *(Impossibility of Time-Consistent Efficiency Assessments) In heterogeneous agent incomplete markets economies, it is impossible to always have time-consistent efficiency assessments with forward-looking numeraires, even in No-Kydland-Prescott scenarios and when using utilitarian social welfare functions.*

This result establishes that in the presence of heterogeneity and market incompleteness it is typically not possible to make efficiency assessments that satisfy the compensation principle and are useful to define Pareto frontiers. This occurs because the rate at which individuals (and the planner) exchange utility gains for consumption-equivalent gains changes over time. For instance, with contemporaneous consumption as welfare numeraire ex ante and ex post, $\lambda^i_{s^0} = u'\left(c^i_0\left(s^0\right)\right)$ and $\lambda^i_{\bar{s}^1} = u'(c^j_1\left(\bar{s}^1\right))$, so unless individuals consume identical amounts in both histories, the covariance term in (13) will be non-zero.[13] More generally, whenever markets are incomplete, individuals' valuation for forward-looking consumption will change as time unfolds, which yields the result.

[13]While one might consider consumption at history $\bar{s}^1$ as the numeraire for the ex ante assessment at date 0, so that trivially $\lambda^i_{s^0} = \lambda^i_{\bar{s}^1}$, this logic works only in simple settings. If a policy also affects allocations in another history $\tilde{s}^1$, time consistency requires that all ex post assessments at date 1 share the sign of the ex ante assessment. But if the ex ante assessment uses history $\bar{s}^1$ consumption as numeraire and history $\bar{s}^1$ realizes, the inconsistency reappears.

We highlight that this result is not about the forms of inconsistency identified by Kydland and Prescott (1977), which also apply to single individual environments — see conditions (i) – (iii) of Proposition 1 — or the use of non-utilitarian social welfare functions, which already make overall welfare assessments time-inconstant — see condition (iv) of Proposition 2. The key forces driving this impossibility result are unique to heterogeneous agent environments with $I > 1$, regardless of the choice of social welfare function. It is also important to highlight that this Theorem is a statement of impossibility for all perturbations in the spirit of Arrow (1950). While specific perturbations my be time-consistent, it is always possible to construct perturbations $d\theta$ whose ex ante and ex post efficiency assessments differ in sign.

Remark. (Practical Implication: Sources of Welfare Gains Change over Time) The main practical implication of this result is that while a utilitarian planner can consistently assess policies, her justification on grounds of efficiency and redistribution changes as time passes and uncertainty is realized. For instance, a policy that ex ante is justified on efficiency grounds, perhaps by improving risk-sharing, ex post can only be desirable purely for a redistributive motive. We illustrate this phenomenon in our applications.

4.2.5 Can Efficiency Assessments be Time-Consistent?

Efficiency assessments avoid time inconsistency in two scenarios: i) with backward-looking numeraires, or ii) when markets are complete.

Backward-looking Welfare Numeraires. Proposition 3 suggests a natural approach to ensure the time consistency of efficiency assessments: choosing ex ante and ex post welfare numeraires such that $\lambda^i_{s^0} = \lambda^i_{\bar{s}^1}$, $\forall i$, so that individual valuations of the numeraire are constant across time and histories and the covariance term in (13) drops out. One natural candidate to ensure $\lambda^i_{s^0} = \lambda^i_{\bar{s}^1}$, $\forall i$, is to use the backward-looking numeraire associated with $\lambda^i_{s^0}$ when making the ex post assessment in history $\bar{s}^1$. The resulting efficiency assessment would then compare gains across individuals in date-0 consumption units, which would mean using a backward-looking welfare numeraire ex post.[14] However, as explained below, efficiency assessments based on backward-looking numeraires no longer satisfy the compensation principle or can be used to define a Pareto frontier, which makes the notion of efficiency in this case not useful.

[14]The idea of using a backward-looking number resembles the use of promise-keeping constraints and recursive multipliers (Marcet and Marimon, 2019; Ljungqvist and Sargent, 2018), as their use requires keeping track of allocations in the past and in histories that did not realize. Using a backward-looking numeraire is akin to keeping track of exchange rate at which a planner traded off utility gains into a common unit at the time of the ex-ante assessment.

Complete Markets. Under complete markets, marginal rates of substitution between histories are equalized across all individuals, so

$$\frac{u'\left(c_1^i\left(\bar{s}^1\right)\right)}{u'\left(c_0^i\left(s^0\right)\right)}=\frac{u'(c_1^j\left(\bar{s}^1\right))}{u'(c_0^j\left(s^0\right))} \tag{14}$$

for all individuals i and j. But this property immediately implies that

$$\frac{\lambda_{\bar{s}^1}^i}{\lambda_{s^0}^i}=\frac{\lambda_{\bar{s}^1}^j}{\lambda_{s^0}^j}, \quad \forall i,j \tag{15}$$

which guarantees the time consistency of efficiency assessments. Intuitively, when markets are complete, individuals' relative valuation of consumption at any two histories is equalized (14). Any valid welfare numeraire must correspond to a bundle of consumption at different histories and $\lambda_{s^t}^i$ represents individual i's valuation of that bundle. Therefore, if (14) holds for all pairs of histories, then valuations for any combination of histories must also be equalized, which directly implies (15). Complete markets therefore guarantee that individuals' valuations of the ex ante and ex post welfare numeraires evolve proportionally. Condition (v) of Proposition 3 therefore never applies in complete markets environments. And under the no-KP assumption, which rules out conditions (i) – (iii) of Proposition 1, efficiency assessments are therefore always time-consistent when markets are complete.

4.3 Inconsistency of Sources of Efficiency

Finally, we explore the inconsistency of the different sources of efficiency gains. In dynamic stochastic economies with heterogeneous agents, the efficiency gains defined in (12) can be due to i) *aggregate efficiency*, which captures the gains from changes in discounted aggregate consumption, ii) *risk-sharing*, which captures the gains from reallocating consumption toward individuals who value it more at specific histories, and iii) *intertemporal-sharing*, which captures the gains from reallocating consumption toward individuals who value it more at specific dates. The following Lemma derives Dávila and Schaab (2024) decomposition in our environment, using perpetual unit consumption as welfare numeraire.[15]

Lemma. *(Sources of Efficiency) The efficiency assessment $\Xi_{s^k}^E$ of a perturbation from the perspective of history s^k can be decomposed into aggregate-efficiency, risk-sharing, and intertemporal-*

[15]This is the unique decomposition in which efficiency gains can be expressed as the discounted sum — using an aggregate discount factor — of aggregate consumption gains, Ξ^{AE}, and two terms that capture differences in valuations across dates and histories.

sharing components, as follows:

$$\Xi^{E}_{s^k} = \Xi^{AE}_{s^k} + \Xi^{RS}_{s^k} + \Xi^{IS}_{s^k}, \qquad \textit{where}$$

$$\Xi^{AE}_{s^k} = \sum_{t \geq k} \omega_{t|s^k} \sum_{s^t \geq s^k} \omega_{t|s^k}\left(s^t\right) \Xi^{AE}_{t|s^k}\left(s^t\right) \quad \textit{where} \quad \Xi^{AE}_{t|s^k}\left(s^t\right) = \sum_i \frac{\partial \mathcal{C}^i_{t|s^k}\left(s^t\right)}{\partial \theta}$$

$$\Xi^{RS}_{s^k} = \sum_{t \geq k} \omega_{t|s^k} \sum_{s^t \geq s^k} \omega_{t|s^k}\left(s^t\right) \Xi^{RS}_{t|s^k}\left(s^t\right) \quad \textit{where} \quad \Xi^{RS}_{t|s^k}\left(s^t\right) = \mathbb{C}ov^{\Sigma}_i \left[\frac{\omega^i_{t|s^k}\left(s^t\right)}{\omega_{t|s^k}\left(s^t\right)}, \frac{\partial \mathcal{C}^i_{t|s^k}\left(s^t\right)}{\partial \theta}\right]$$

$$\Xi^{IS}_{s^k} = \sum_{t \geq k} \omega_{t|s^k} \sum_{s^t \geq s^k} \omega_{t|s^k}\left(s^t\right) \Xi^{IS}_{t|s^k}\left(s^t\right) \quad \textit{where} \quad \Xi^{IS}_{t|s^k}\left(s^t\right) = \mathbb{C}ov^{\Sigma}_i \left[\frac{\omega^i_{t|s^k}}{\omega_{t|s^k}}, \frac{\omega^i_{t|s^k}\left(s^t\right)}{\omega_{t|s^k}\left(s^t\right)} \frac{\partial \mathcal{C}^i_{t|s^k}\left(s^t\right)}{\partial \theta}\right],$$

where $\omega_{t|s^k} = \frac{1}{I}\sum_i \omega^i_{t|s^k}$ *and* $\omega_{t|s^k}\left(s^t\right) = \frac{1}{I}\sum_i \omega^i_{t|s^k}\left(s^t\right)$ *define cross-sectional averages of (normalized) dynamic and stochastic weights, in turn defined as*

$$\omega^i_{t|s^k} = \frac{\beta^{t-k}\sum_{s^t \geq s^k} \pi\left(s^t|s^k\right) u'\left(\mathcal{C}^i_{t|s^k}\left(s^t\right)\right)}{\sum_{t \geq k} \beta^{t-k} \sum_{s^t \geq s^k} \pi\left(s^t|s^k\right) u'\left(\mathcal{C}^i_{t|s^k}\left(s^t\right)\right)} \quad \textit{and} \quad \omega^i_{t|s^k}\left(s^t\right) = \frac{\pi\left(s^t|s^k\right) u'\left(\mathcal{C}^i_{t|s^k}\left(s^t\right)\right)}{\sum_{s^t \geq s^k} \pi\left(s^t|s^k\right) u'\left(\mathcal{C}^i_{t|s^k}\left(s^t\right)\right)}.$$

The dynamic weight $\omega^i_{t|s^k}$ measures how individual i trades off date-t consumption against perpetual consumption from perspective s^k. For example, if $\omega^i_{t|s^k} = 0.1$, the planner perceives a unit increase in consumption assigned to individual i at date t as equivalent to a 0.1% increase in perpetual consumption assigned to that individual. The stochastic weight $\omega^i_{t|s^k}\left(s^t\right)$ measures how individual i trades off history-s^t consumption against unconditional consumption at date-t from perspective s^k. For example, if $\omega^i_{t|s^k} = 0.1$, the planner perceives a unit increase in consumption assigned to individual i at history s^t as equivalent to a 0.1% increase in consumption at all histories at date t.

The following proposition shows that the inconsistency of efficiency assessments extends to their underlying sources, although through different mechanisms. The ultimate cause of time inconsistency in these cases is the fact that the valuations that individuals attribute to consumption at different dates and histories changes with the perspective of the assessment.

Proposition 4. *(Inconsistency of Sources of Efficiency) In No-Kydland-Prescott scenarios, for any social welfare function, and with forward-looking welfare numeraires:*

a) If markets are complete, risk- and intertemporal-sharing assessments are zero, $\Xi^{RS}_{s^k} = \Xi^{IS}_{s^k} = 0$*, and (static and intertemporal) aggregate efficiency assessments* $\Xi^{AE}_{s^k}$ *are time-consistent.*

b) If markets are incomplete, static aggregate-efficiency assessments $\Xi^{AE}_{s^k}\left(s^t\right)$ *are time-*

consistent, but intertemporal aggregate-efficiency assessments, as well as risk- and intertemporal-sharing assessments are time inconsistent.

Consistent with the impossibility theorem shown above, complete markets ensures that efficiency assessments are time inconsistent. This is achieved because risk- and intertemporal-sharing gains are zero at all times, and the fact that all individual valuations move in lockstep ensure that aggregate-efficiency considerations are time-consistent, similar to a single-individual scenario.

However, part b) of Proposition 4 underscores that assessments of static aggregate efficiency gains are special.[16] In general, an overall aggregate-efficiency assessment $\Xi_{s^k}^{AE}$ is time inconsistent not because its components at specific histories $\Xi_{s^k}^{AE}\left(s^t\right)$ differ across perspectives (in the Non-Kydland-Prescott case), but because the aggregate weights used to discount them, $\omega_{t|s^k}$ and $\omega_{t|s^k}\left(s^t\right)$, vary with the perspective of the assessment. Hence, static assessments of aggregate efficiency are time-consistent, but once multiple periods or histories are involved, the change in the relevant aggregate weights generates inconsistency.

By contrast, both risk-sharing and intertemporal-sharing. $\Xi_{s^k}^{RS}$ and $\Xi_{s^k}^{IS}$, are time inconsistent not only because of shifting aggregate weights, $\omega_{t|s^k}$ and $\omega_{t|s^k}\left(s^t\right)$, but also because its components at specific histories, $\Xi_{s^k}^{RS}\left(s^t\right)$ and $\Xi_{s^k}^{IS}\left(s^t\right)$, themselves depend on the perspective of the assessment. For risk-sharing, the relevant covariance terms reflect conditional probabilities that change between ex ante and ex post evaluations, so $\Xi_{t|s^0}^{RS}\left(s^t\right) \neq \Xi_{t|s^k}^{RS}\left(s^t\right)$. A similar logic applies to intertemporal-sharing, where the interaction between dynamic and stochastic weights leads to assessments that differ across perspectives, so $\Xi_{t|s^0}^{IS}\left(s^t\right) \neq \Xi_{t|s^k}^{IS}\left(s^t\right)$.

Together, these results underscore that the evolution of the different sources of efficiency is sophisticated and non-trivial. While no general results are available at the level of generality we have considered so far, in realistic scenarios clear patterns emerge in the evolution of the different sources of efficiency, as our applications illustrate

5 Applications

This section presents three applications that demonstrate how the time inconsistency forces identified in this paper arise in practical settings. Each application is designed with a dual

[16] With factors of production, aggregate-efficiency gains include gains from production efficiency, as in Section 2. With multiple goods, aggregate efficiency gains also include exchange efficiency gains, as in an Edgeworth box. See Dávila and Schaab (2023b).

objective: first, capture a realistic scenario in which the time inconsistency forces due to heterogeneity that we have studied play an important role, and second, illustrate how a particular source of efficiency gains — intertemporal-sharing, risk-sharing, and aggregate-efficiency — can lead to the time inconsistency of policies, as summarized in Table 1.

Table 1: Summary of Applications

#	Application	Ξ^{AE}	Ξ^{RS}	Ξ^{IS}
1	Anticipated Relief Policies	$= 0$	$= 0$	✓
2	Risk-Sharing Policies	$= 0$	✓	$= 0$
3	Aggregate Investment Policies	✓	$= 0$	$= 0$

Note: This table summarizes the determinants of efficiency gains present in each application.

5.1 Application 1: Anticipated Relief Policies

This application identifies a particular form of time inconsistency that arises in the context of anticipated relief policies — transfer schemes intended to support individuals experiencing transitory spells of low consumption followed by a gradual, anticipated recovery. When markets are incomplete, a government relief policy can therefore generate efficiency (intertemporal-sharing) gains at a particular date by transferring resources towards individuals whose consumption is below their long-run consumption level, replicating the consumption-smoothing individuals would choose if borrowing were feasible.

Ex ante, the efficiency-maximizing transfer policy that uses permanent consumption as welfare numeraire supports low-consumption individuals early on, phasing out once their consumption approaches its long-run level. Over time, a gradual recovery takes place. However, ex post, consumption is still below its long-run level at times in which the originally transfer policy has phased out, opening the door to a new round of transfers. The efficiency-maximizing transfer is therefore extended ex post beyond the ex ante optimal time of expiry.

The result is a dynamic inconsistency in the design of anticipated relief policies, where an efficiency maximizing planner finds it optimal ex post to extend the relief longer than originally intended. Relief policies of this sort are common. A salient example are the COVID-19 assistance programs, designed to temporarily support households disproportionately affected by the pandemic.[17] But the time consistency problem identified here also emerges in other

[17]These programs were initially enacted through the Coronavirus Aid, Relief, and Economic Security (CARES) Act, in March 2020. They were subsequently expanded through the Consolidated Appropriations Act, in

policy domains, including i) childcare support and parental leave, ii) unemployment insurance, iii) retraining programs, and iv) affirmative action policies. In each case, policies designed ex ante to offset a foreseeable period of hardship and intended to phase out once recovery is underway are extended ex post.

5.1.1 Environment

We consider a deterministic finite-horizon economy with dates $t \in \{0, 1, \ldots, T\}$. There are two individuals with identical preferences, indexed by $i \in \{A, B\}$, and there is a single good that appears as an endowment.

Preferences and Endowments. Individual i's lifetime utility from the perspective of date k is given by

$$V^i_{|k} = \sum_{t \geq k} \beta^{t-k} u\left(c^i_t\right).$$

Individual i's endowment at date t is denoted by y^i_t. Individual A has a constant endowment of $y^A_t = y$. Individual B starts with the same endowment as individual A, but experiences a transitory spell of low but gradually recovering endowments that starts at date $\underline{T}$ and lasts until $\overline{T}$. Hence, individual B's endowment, y^B_t, is given by

$$y^B_t = \begin{cases} y, & t < \underline{T} \\ \underline{y} + \frac{t-\underline{T}}{\overline{T}-\underline{T}} \left(y - \underline{y}\right), & \underline{T} \leq t \leq \overline{T} \\ y, & t > \overline{T}. \end{cases}$$

We adopt this formulation so that there is a terminal date for the individual A's low endowment spell. Identical results arise if individual A's endowment recovered exponentially, as in Application 2.

Relief Policy. We focus on financial autarky for illustration, so individual i's consumption at date t is

$$c^i_t = y^i_t + \theta M^i_t,$$

December 2020, and, more notably, the American Rescue Plan Act (ARPA), in March 2021. The timing of ARPA — which took place once the recovery was well underway — is consistent with the time inconsistency forces presented in this application.

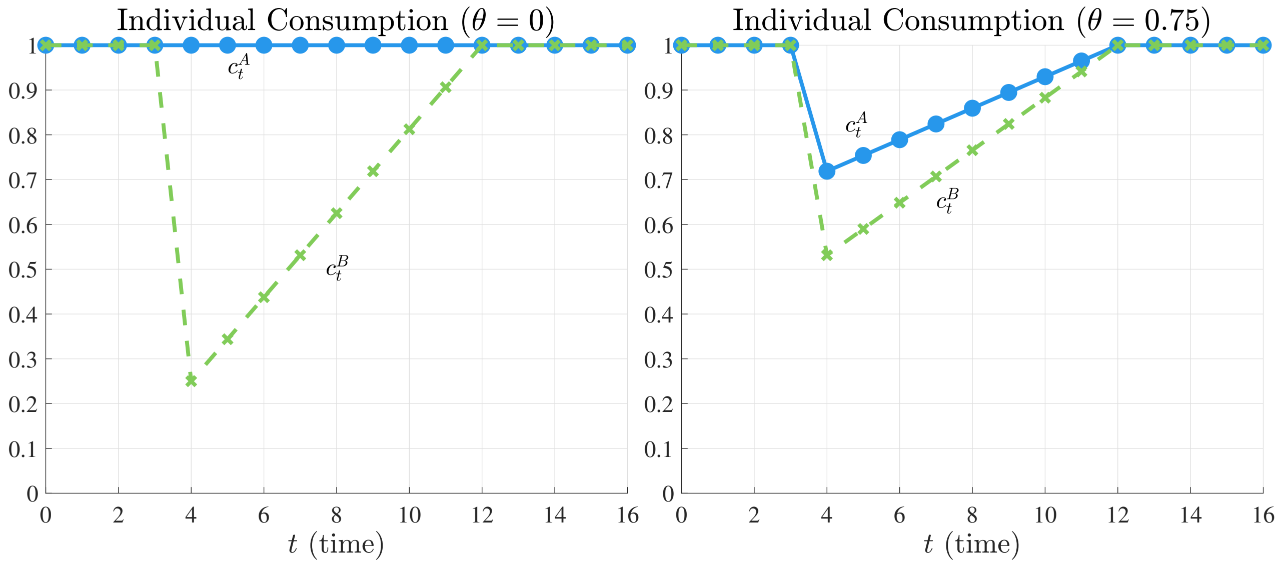

Figure 3: Consumption Profiles (Application 1)

Note. This figure shows consumption for individual A (green dashed line) and individual B (solid blue line) before (left panel, $\theta = 0$) and after (right panel, $\theta = 0.75$) the relief policy is implemented.

where M_t^i represents a transfer scheme that provides relief to individual B during the low-endowment spell, with

$$M_t^A = \frac{1}{2}\left(y_t^B - y_t^A\right) \qquad \text{and} \qquad M_t^B = -M_t^A.$$

As $\theta \to 1$, both individuals consume identical amounts. Figure 3 shows the consumption profiles for both individuals without the relief policy ($\theta = 0$) in the left panel, and after a sizable transfer ($\theta = 0.75$) in right panel.

Parametrization. We assume log utility $u(c) = \log(c)$ and interpret a date in the model as a quarter, setting a discount factor $\beta = (0.95)^{1/4}$. We normalize individual A's endowment level to $y = 1$ and consider a 75% initial consumption drop for individual B, so $\underline{y} = 0.25$. Finally, we assume that the low-endowment spell starts at date $\underline{T} = 4$ and lasts until $\overline{T} = 12$, where $T = 16$. For some applications (e.g. COVID relief policies), it may be natural to set $\underline{T} = 0$. The results are qualitatively unchanged in that case.

5.1.2 Weights, Assessments, and Time Inconsistency

We study welfare assessments of the relief policy M_t^i under an equal-weighted utilitarian social welfare function. Normalized welfare gains from the perspective of date k are thus given by

$$\frac{dW_{|k}^{\lambda}}{d\theta} = \sum_i \omega_{|k}^i \sum_{t \geq k} \omega_{t|k}^i \frac{dc_t^i}{d\theta},$$

where i's normalized individual weight and date-t's normalized dynamic weight — both from the perspective of date k and using perpetual consumption as welfare numeraire — are respectively given by

$$\omega_{|k}^i = \frac{\sum_{t \geq k} \beta^{t-k} u'\left(c_t^i\right)}{\sum_i \sum_{t \geq k} \beta^{t-k} u'\left(c_t^i\right)} \quad \text{and} \quad \omega_{t|k}^i = \frac{\beta^{t-k} u'\left(c_t^i\right)}{\sum_{t \geq k} \beta^{t-k} u'\left(c_t^i\right)}. \tag{16}$$

Normalized welfare gains can be decomposed into efficiency and redistribution components, where $\frac{dW_{|k}^{\lambda}}{d\theta} = \Xi_{|k}^E + \Xi_{|k}^{RD}$. Efficiency gains are thus given by

$$\Xi_{|k}^E = \sum_i \sum_{t \geq k} \omega_{t|k}^i \frac{dc_t^i}{d\theta} = \sum_{t \geq k} \mathbb{C}ov_i^{\Sigma} \left[\omega_{t|k}^i, \frac{dc_t^i}{d\theta}\right],$$

where the second equality follows because this is an endowment economy, with $\sum_i \frac{dc_t^i}{d\theta} = 0$, and where $\mathbb{C}ov_i^{\Sigma}[\cdot,\cdot] = I \cdot \mathbb{C}ov_i[\cdot,\cdot]$ denotes a covariance-sum. Since this is a single good deterministic economy, all efficiency gains are due to reallocating resources towards individuals with relatively higher valuations for consumption at particular dates, that is, they are due to intertemporal-sharing, so $\Xi_{|k}^E = \Xi_{|k}^{IS}$. Redistribution gains, which capture whether individuals relative preferred by the planner are relatively favored by the policy, are given by

$$\Xi_{|k}^{RD} = \mathbb{C}ov_i^{\Sigma} \left[\omega_{|k}^i, \sum_{t \geq k} \omega_{t|k}^i \frac{dc_t^i}{d\theta}\right],$$

where individual weights $\omega_{|k}^i$ are defined in (16).

Normalized Individual Weights. The individual weight $\omega_{|k}^i$ captures the value that a planner assigns to the gains of a individual i in units of the welfare numeraire from a particular perspective k, shaping the welfare gains due to redistribution. The left panel of Figure 4 shows the evolution of these weights over time . The ratio $\omega_{|k}^B / \omega_{|k}^A$ represents the planner's relative valuation of the two individuals in units of welfare numeraire. This ratio peaks at about $\frac{1.25}{0.75} \approx 1.67$ at date $\underline{T} = 4$, the date at which the difference in terms of permanent consumption is largest between individuals, before converging back to 1 by date $\overline{T} = 12$. At date $\underline{T} = 4$, the planner values giving a unit of numeraire to individual B by 1.67 more than to individual

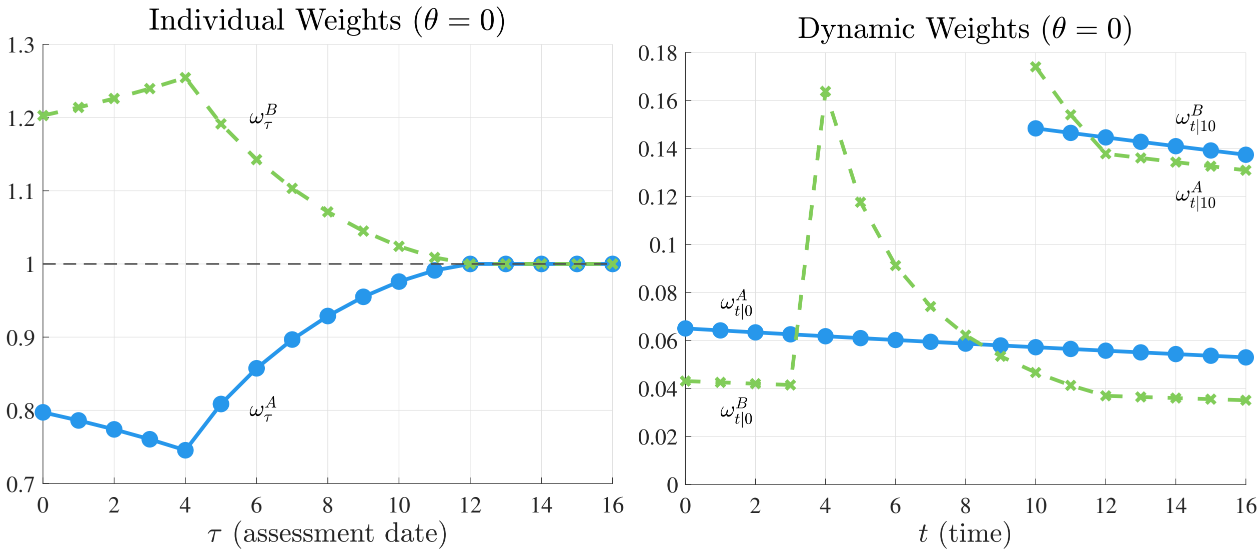

Figure 4: Individual and Dynamic Weights (Application 1)

Note. The left panel of this figure shows the evolution of the normalized individual weights for both individuals as a function of the assessment date for an equal-weighted utilitarian planner when $\theta = 0$. The right panel shows the normalized dynamic weights as a function of time for two particular assessments: $k = 0$ and $k = 10$. The solid blue line corresponds to individual A's weights, and the green dashed line corresponds to individual B's.

A. Consistent with this largest difference, the welfare gains attributed to redistribution are maximal at the onset of the low-endowment spell ($\underline{T} = 4$), as shown in Figure 5.

Normalized Dynamic Weights. The dynamic weights $\omega_{t|k}^i$ define marginal rates of substitution for each individual between date-t and perpetual aggregate consumption from a particular perspective k, and capture whose consumption the planner values more at each date. By construction, these dynamic weights sum to 1 for each individual because giving a one unit of consumption at a particular date for all dates is equivalent to providing one unit of perpetual consumption.[18] The right panel of Figure 4 shows the dynamic weights for each individual for two different perspectives $\omega_{t|k}^i$ against calendar time t: the ex ante perspective at date $k = 0$ and the ex post perspective at date $k = 10$. Cross-sectional dispersion in normalized dynamic weights $\omega_{t|k}^i$ at a particular date indicates that there is scope for efficiency gains from reallocating consumption. Intuitively, consumption-smoothing efficiency gains can be achieved by transferring consumption from the individual with a smaller dynamic weight $\omega_{t|k}^i$ at date t to the individual with the larger one.

When the low-endowment spell materializes at date $\underline{T} = 4$, $\frac{\omega_{4|0}^B}{\omega_{4|0}^A} \approx \frac{0.16}{0.06} = 2.67$, so the

[18]While there are good arguments to use perpetual consumption as welfare numeraire, there are other reasonable choices of welfare numeraires in this environment, for instance, date-0 consumption.

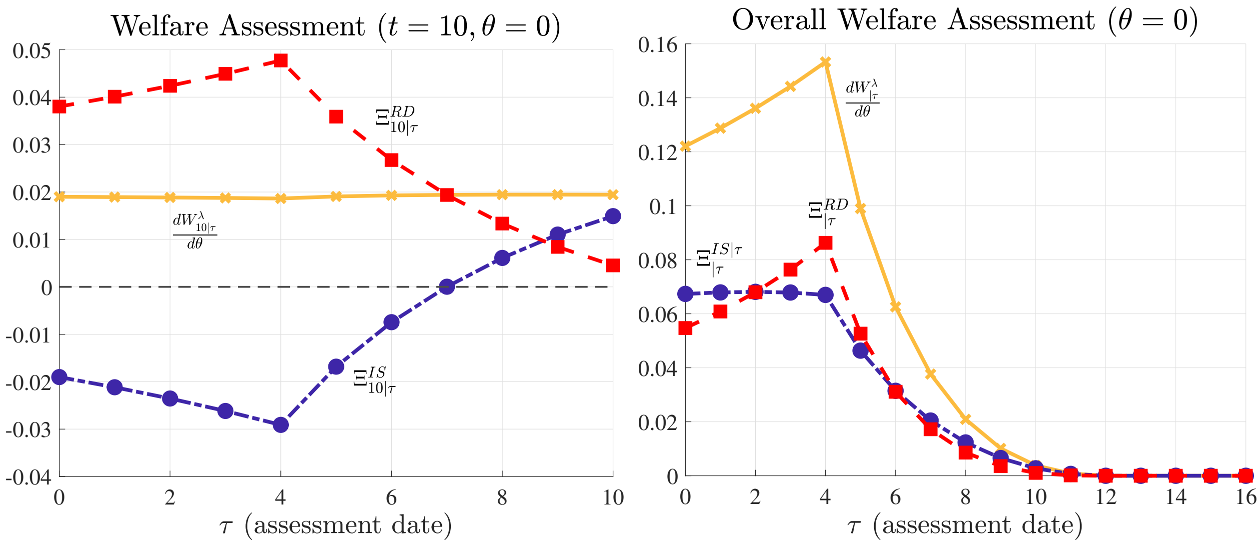

Figure 5: Welfare Assessment of Relief Policy (Application 1)

Note. This figure shows how the planner makes welfare, efficiency (intertemporal-sharing), and redistribution assessments as a function of the assessment date when $\theta = 0$. The left panel exclusively shows the welfare, efficiency, and redistribution assessments that pertain date-10. That is, $\Xi^{IS}_{10|\tau} = \mathbb{C}ov^{\Sigma}_{i}\left[\frac{\omega^{i}_{10|\tau}}{\omega_{10|\tau}}, \frac{dc^{i}_{10}}{d\theta}\right]$, $\Xi^{RD}_{10|\tau} = \mathbb{C}ov^{\Sigma}_{i}\left[\omega^{i}_{|\tau}, \frac{\omega^{i}_{10|\tau}}{\omega_{10|\tau}}\frac{dc^{i}_{10}}{d\theta}\right]$ and $\frac{dW^{\lambda}_{10|\tau}}{d\theta} = \Xi^{IS}_{10|\tau} + \Xi^{RD}_{10|\tau}$. The right panel accounts for all future times from a particular assessment date.

planner regards consumption reallocation towards B to be efficiency maximizing. And since we have $\omega^{B}_{t|0} > \omega^{A}_{t|0}$ for $t \in [4, 8]$, the planner finds a relief policy that transfers from A to B and that starts at date 4 and expires at date 8 desirable. After $\overline{T} = 8$, individual B's consumption is closer to its long-run/permanent level from a date 0 perspective, so the normalized dynamic weight of individual B falls below that of individual A. This has the interpretation that it is at those times when individual B would pay back the early transfer to individual A if possible.

As time goes by, Figure 4 illustrates how the dynamic weights look from a date $k = 10$ perspective. From this ex post perspective, the worst of the low-endowment spell for individual B now lies in the past. So while consumption at date 10 from a date 0 perspective was perceived to be below B's long-run level, consumption at date 10 from a date 10 perspective now appears to still be below its long-run level from date 10 onwards. As a result, individual B's normalized dynamic weight $\omega^{B}_{t|10}$ is still larger than individual A's at dates 10 and 11, unlike from the ex ante perspective. Therefore, the planner finds it desirable to extend the relief policy ex post and continue transferring from individual A to individual B until date 12.

Time Inconsistency. Figure 5 illustrates the time inconsistency of efficiency assessments that emerges in this setting. In particular, the left panel shows how the planner assesses the relief policy that materializes at date 10 from different assessment perspectives (horizontal axis). Consistent with our explanation of dynamic weights, this figure shows that the efficiency gains attributed to a relief policy at date 10 are perceived to be initially negative, as $\Xi^{E}_{10|0} = \Xi^{IS}_{10|0} < 0$ ex ante, eventually turning positive as time goes by, so $\Xi^{E}_{10|8} = \Xi^{IS}_{10|8} > 0$. By date 8, the planner finds it valuable to extend the relief policy on efficiency grounds. This illustrates the time consistency forces that lead the planner to extend the relief policy beyond the originally intended expiration date.

Consistent with our explanation of individual weights, the left panel of Figure 5 also illustrates that i) the redistribution assessment of the date-10 component of the relief policy is highest when permanent consumption differences are higher, at $\underline{T} = 4$, and ii) the overall welfare assessment of the date-10 component of the relief policy largely remains unchanged as the assessment date changes, masking substantial variation in its efficiency and redistribution components. The right panel of Figure 5 shows the assessment of the relief policy as a whole — that is, not only the impact of the date-10 components, but the impact of the policy over all future dates — from different assessments perspectives. It shows that the relief policy seems desirable on both efficiency and redistribution grounds when considered as a whole from all perspectives. However, our previous analysis shows that this apparent uniformity in the desire to implement the policy masks subtle economic forces.

5.2 Application 2: Risk-Sharing Policies

This application identifies a different form of time inconsistency that arises when designing policies to share risks that will materialize in the future. When markets are incomplete, a policy that transfers resources towards ex-post unlucky agents is perceived as increasing efficiency from an ex ante perspective, replicating the consumption-smoothing that the agents would choose if insurance markets were open.

In particular, in a symmetric economy with two ex-ante identical agents that consume at a terminal date, the ex-ante efficiency-maximizing policy would transfer resources from one agent to the other in half of the terminal histories. Over time, uncertainty is partially realized, becoming clear than one agent will do better than the other at the terminal date. Hence, at interim dates, the new efficiency-maximizing policy calls for reversing the direction of the transfer in intermediate histories, which were bad in relative terms for one of the agents ex-ante, but that now are perceived as relatively good scenarios from the new perspective.[19] While

[19]Note that in this application we focus on the inconsistency of risk-sharing (efficiency) assessments *before*

ex-ante symmetry helps to present the results, similar forces apply when agents are ex-ante different.

The result is a dynamic inconsistency in the design of risk-sharing policies: transfers that were initially desirable on efficiency grounds are later dismantled ex post. A salient example of the time inconsistency forces studied here is the Stability and Growth Pact (SGP) and related mechanisms aimed to insure EU member states against shocks through common rules, access to credit, and ECB support. Ex-ante, all EU countries were perceived to face similar risks. However, between 2008 and 2011, Greece and other peripheral countries experience worse-than-expected macroeconomic shocks. This led to a shift in beliefs: future states that had been considered "bad" for peripheral countries ex ante were now seen as relatively favorable, given the realization that these countries would systematically experience worse outcomes overall. As a result, EU support policies began requiring countries like Greece to contribute more — for example, through austerity measures — in these relatively good states, contrary to the original plan.

But the time inconsistency problem identified here also emerges in other situations in which agent share risks, including i) relief policies against natural disasters (e.g. FEMA), ii) student loan forgiveness plans, iii) sectoral labor market support policies. In each case, policies designed ex ante to better share risk among similar agents in some states are perceived as counterproductive for risk-sharing purposes once time has gone by and it is clear than one of the agents will systematically be worse off.

5.2.1 Environment

We consider a finite-horizon economy with dates $t \in \{0, 1, \ldots, T\}$. There are two individuals with identical preferences, indexed by $i \in \{A, B\}$, and there is a single good that appears as an endowment.

Preferences and Endowments. Both individuals exclusively consume at the terminal date T. Individual's i utility at date T is given by

$$V_T^i(n) = u\left(c_T^i(n)\right),$$

a terminal date. Once all risks are resolved, efficiency gains are zero, and welfare assessments can only only driven by redistribution.

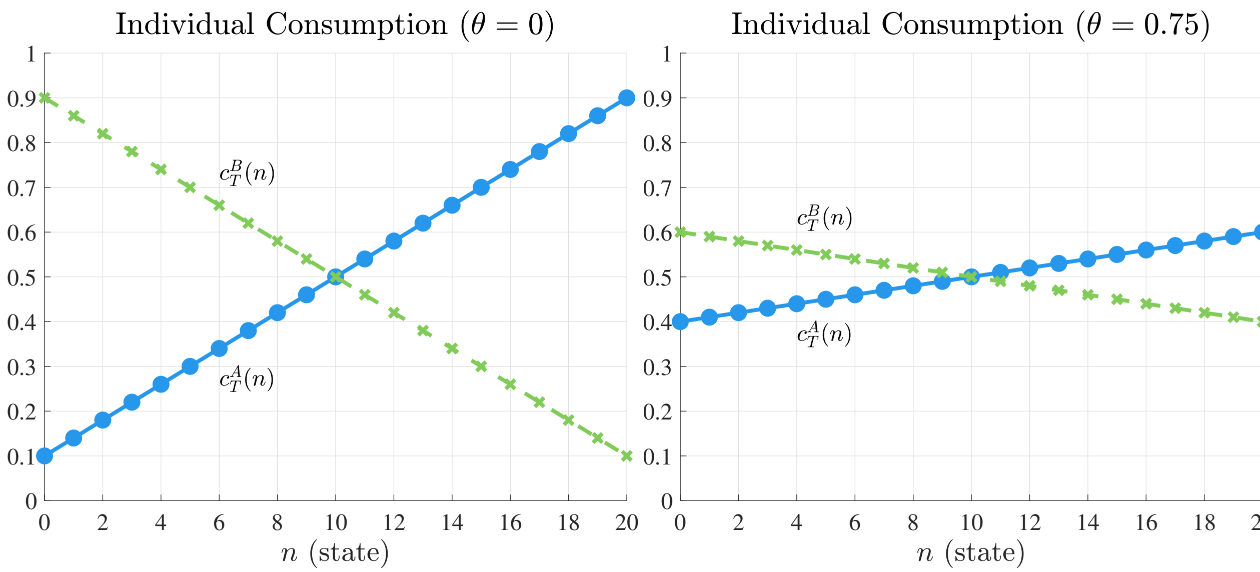

Figure 6: Consumption Profiles (Application 2)

Note. This figure shows consumption for individual A (green dashed line) and individual B (solid blue line) before (left panel, $\theta = 0$) and after (right panel, $\theta = 0.75$) a relief policy is implemented.

where n indexes the state of the economy. Individual i's endowment at the terminal date T in state n, denoted by $y_T^i(n)$, is given by

$$y_T^i(n) = \chi_T^i(n) y_T,$$

where $\chi_T^i(n)$ denotes i's share of the aggregate endowment y_T, normalized to $y_T = 1$ regardless of the state. We assume that A's endowment share is given by

$$\chi_T^A(n) = \chi_0 + \overline{g} \cdot n - \underline{g} \cdot (T - n),$$

where $\chi_T^B(n) = 1 - \chi_T^A(n)$, where χ_0, $\overline{g}$, and $\underline{g}$ are positive parameters, and where the state evolves as a binomial process. That is, at each date t, the state either increases by one with probability $q \in [0, 1]$, or remains the same with probability $1 - q$, starting from an initial state $n = 0$. Therefore, there are $n = T$ possible states at the terminal date.[20] Hence, from the perspective of an earlier date $t \leq T$, individual i's utility satisfies the recursion

$$V_t^i(n) = q V_{t+1}^i(n+1) + (1-q) V_{t+1}^i(n),$$

[20] The probability of reaching state n at date t starting from state $\underline{n}$ at date $\underline{t}$ follows a binomial distribution, with

$$\pi_{t|\underline{t}}(n|\underline{n}) = \binom{\Delta t}{\Delta n} q^{\Delta n} (1-q)^{\Delta t - \Delta n},$$

where $\Delta n = n - \underline{n}$ and $\Delta t = t - \underline{t}$.

reflecting the binomial evolution of n. That is, while consumption exclusively takes place at date T, the realizations of the binomial process before then convey information about individual endowments, shaping normative assessments.

Risk-Sharing Policy. Considering a financial autarky scenario, as in Application 1, individual i's consumption at date t in state n is given by

$$c_T^i(n) = y_T^i(n) + \theta M_T^i(n),$$

where $M_t^i(n)$ represents a state contingent policy that transfers resources from the lower endowment to the higher endowment individual. Formally,

$$M_T^A(n) = \frac{1}{2}\left(y_T^B(n) - y_T^A(n)\right), \quad \text{where} \quad M_T^B(n) = -M_T^A(n).$$

As $\theta \to 1$, both individuals consume identical amounts. Figure 6 shows the consumption profiles for both individuals without the risk-sharing policy ($\theta = 0$) in the left panel, and once partially implemented ($\theta = 0.75$) in right panel.

Parametrization. We assume isoelastic preferences $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, with $\gamma = 2$, and interpret a date in the model as one year, setting a discount factor $\beta = 0.95$. We assume that $T = 20$, $q = 0.5$, making up and down moves equally likely, and set $\overline{g} = \underline{g} = 0.02$, so that the highest and lowest possible terminal consumption shares for individual A are 0.9 and 0.1, respectively.

5.2.2 Weights, Assessments, and Time Inconsistency

We study welfare assessments of the risk-sharing policy M_t^i under an equal-weighted utilitarian social welfare function. Normalized welfare gains from the perspective of state m at date k can be expressed as

$$\frac{dW_{|k.m}^{\lambda}}{d\theta} = \sum_i \omega_{|k,m}^i \sum_n \omega_{T|k,m}^i(n) \frac{dc_T^i(n)}{d\theta},$$

where i's normalized individual weight and state n's normalized dynamic weight — both from the perspective of date k and using perpetual consumption as welfare numeraire— are respectively given by

$$\omega_{|k,m}^i = \frac{\sum_n \pi_{T|t}(n|m)\, u'(c_T^i(n))}{\sum_i \sum_n \pi_{T|t}(n|m)\, u'(c_T^i(n))} \quad \text{and} \quad \omega_{T|k,m}^i(n) = \frac{\pi_{T|k}(n|m)\, u'(c_T^i(n'))}{\sum_{\tilde{n}} \pi_{T|k}(\tilde{n}|m)\, u'(c_T^i(\tilde{n}))}. \tag{17}$$

Normalized welfare gains from the perspective of state m at date k can be decomposed into efficiency and redistribution components, where $\frac{dW_{|k,m}^{\lambda}}{d\theta} = \Xi_{|k,m}^{RS} + \Xi_{|k,m}^{RD}$. Efficiency and

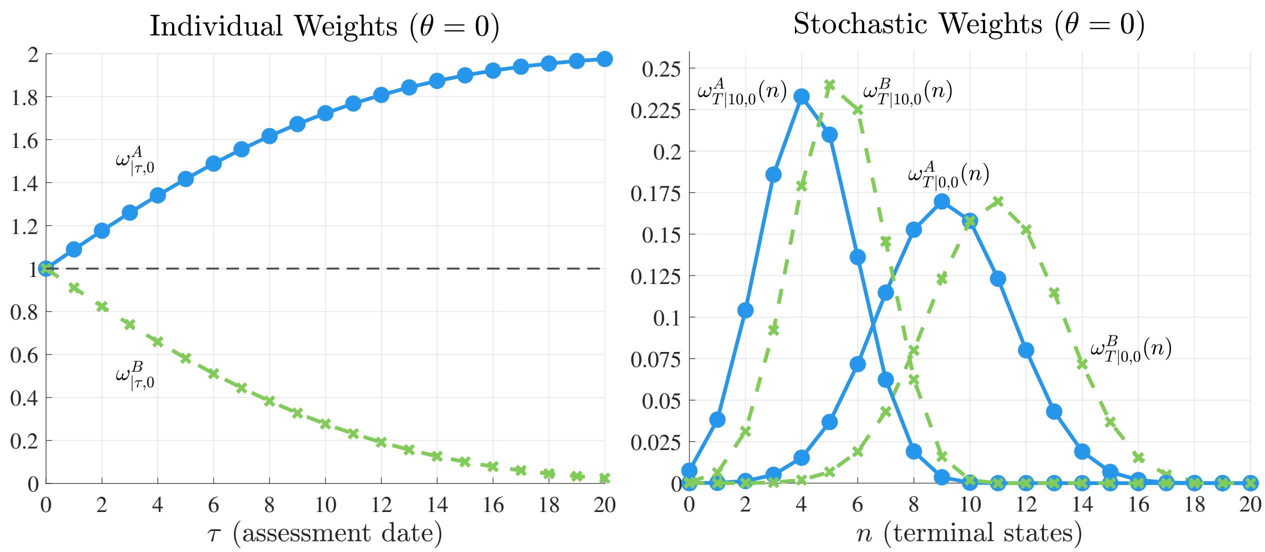

Figure 7: Individual and Stochastic Weights (Application 2)

Note. The left panel of this figure shows the evolution of the normalized individual weights assigned by the equal-weighted utilitarian planner to individual A (solid blue) and B (dashed green) over time when the state is $n = 0$ and when $\theta = 0$. As times goes by and the state remains at 0, the planner assigns a higher weight to individual A, who is more likely to be worse off at the terminal date. The right panel of this figure shows the stochastic weights associated with the terminal date when the state is $n = 0$ at two particular assessments times: $t = 0$ and $t = 10$. The solid blue lines correspond to individual A's weights and the green dashed lines to individual B's.

redistribution gains — using unconditional consumption across all states at date T as welfare numeraire — are thus given by

$$\Xi^{RS}_{|k,m} = \sum_n \mathbb{C}ov^{\Sigma}_i \left[\omega^i_{T|k,m}(n), \frac{dc^i_T(n)}{d\theta}\right] \quad \text{and} \quad \Xi^{RD}_{|k,m} = \mathbb{C}ov^{\Sigma}_i \left[\omega^i_{|k,m}, \sum_n \omega^i_{T|k,m}(n) \frac{dc^i_T(n)}{d\theta}\right].$$

Since this is a single good deterministic economy, all efficiency gains are due to reallocating resources towards individuals with relatively higher valuations for consumption at particular states, that is, they are due to risk-sharing, so $\Xi^E_{|k} = \Xi^{RS}_{|k,m}$.

Normalized Individual Weights. The individual weight $\omega^i_{|k,m}$ captures the value that a planner assigns to the gains of individual i in units of the welfare numeraire, from the perspective of date t and state m. The left panel of Figure 7 shows the evolution of these weights over time when the state remains as $n = 0$, a situation in which A continuously experiences negative shocks. Ex ante, when agents are symmetric, individual weights begin equal, but diverge as uncertainty is partially resolved against individual A. As early realizations of the state process imply that A is doing increasingly worse, then $\omega^A_{|\tau,0} > \omega^B_{|\tau,0}$ reflecting the planner's higher valuation of consumption for agent B. This valuation captures the increased concern for the individual who is becoming worse off.

Normalized Stochastic Weights. The stochastic weights $\omega^{i}_{T|k,m}(n)$ define marginal rates of substitution between consumption at terminal state n and unconditional consumption across all terminal states, conditional on being at state n at date k. By construction, these weights sum to 1 for each individual. The right panel of Figure 7 shows the stochastic weights for each individual from two different perspectives: an ex ante perspective at date 0 and state 0 and an ex post perspective at date 10 and state 0. Cross-sectional dispersion in normalized stochastic weights indicates that there is scope for efficiency gains from reallocating consumption at particular terminal states. Intuitively, consumption-smoothing efficiency gains can be achieved by transferring consumption from the individual with a smaller stochastic weight at a particular terminal state to the individual with the larger one.

Initially, a policy that reallocates consumption from individual B to individual A in the terminal states $[0, 10)$ and from A to B in the terminal states $(10, 20]$ is desirable from an efficiency perspective. This is precisely a risk-sharing policy that supports the individual with the worse ex-post realizations. As time goes by, in a situation such as the one shown in the right panel of Figure 7, individual A has only received bad news between date 0 and date 10. From that point onwards, the terminal states 5 through 9, which were relatively bad from an ex ante perspective, are now relatively good states. This is reflected on the fact that the stochastic weights in those states are higher now for B relative to A. Therefore, if allowed to reassess the policy at date 10 and state 0, the planner would find policies that reallocate consumption from individual B to A in the terminal states $[0, 10]$ undesirable.

Time Inconsistency. Figure 8 illustrates the time inconsistency in the planner's efficiency assessments. The left panel plots the efficiency, redistribution, and welfare gains associated with the impact of the policy on the terminal state 6, as perceived from different assessment dates when the state remains $n = 0$ at all times. Initially, the policy of reallocating resources from B to A in the terminal state 6 appears desirable on efficiency grounds, as it generates risk-sharing gains from an ex ante perspective. However, as time progresses and the likely endowment distribution across terminal states shifts, the planner's assessment changes. In particular, if the state remains at $n = 0$ as times goes by, the planner understands that it is more likely that individual A will have a lower endowment compared to B at the terminal date. As a result, transferring resources from B to A in state 6 ceases to be efficiency improving, as those are states in which A should compensating B for the desirable risk-sharing transfers in the really bad terminal scenarios. This is a form of time inconsistency: transfers that were

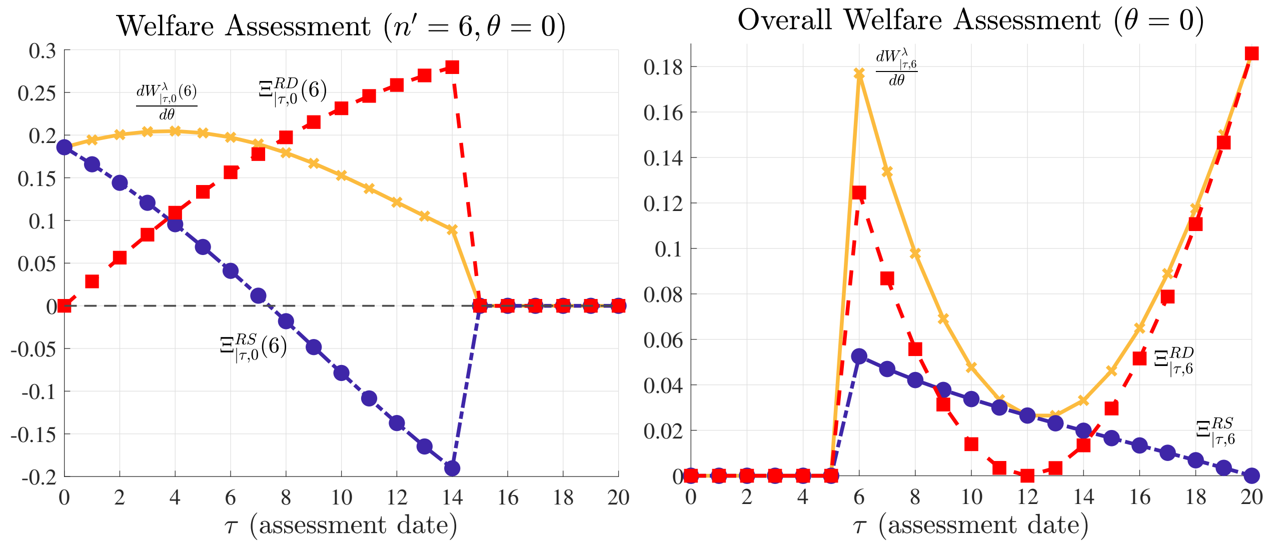

Figure 8: Welfare Assessment of Risk-Sharing Policy (Application 2)

Note. This figure shows how the planner's welfare assessment of the risk-sharing policy changes with the assessment date when $\theta = 0$. The left panel isolates the welfare gain of the policy in a given terminal state $n = 6$, decomposed into efficiency/risk-sharing (blue), redistribution (red), and the total welfare gain (yellow). While initially justified by risk-sharing, transferring resources from B to A at the terminal state 6 eventually becomes undesirable. The right panel reports the assessment of the entire policy, showing that although total welfare gains remain non-negative, its components change over time.

optimal ex ante are dismantled ex post. A real-world example is the shift in EU support policies after the 2008 financial crisis. Ex ante, support was motivated by symmetric risks. But once peripheral countries like Greece experienced systematically bad outcomes, future "bad" states were perceived as relatively good — leading to less generous transfers and conditionality requirements.

The right panel of Figure 8 shows the assessment of the risk-sharing policy as a whole—not just to a specific terminal state. It illustrates the evolution of the assessment from the perspective of a fixed state $n = 6$. While the total welfare gain from the policy is consistently positive (due to the use of a utilitarian welfare function), its composition and overall values change significantly over time, in particular due to redistribution considerations. Starting from date 6, if the economy is state $n = 6$ this means that individual A is likely to be better off at the terminal date, which makes the policy good for redistribution purposes. As time goes by and the economy remains at the state $n = 6$, this is bad news for A, so redistribution gains for the policy become lower until $t = 12$. As time evolves, if the state remains at $n = 6$, then it is more and more likely that individual B would do well at the terminal, increasing again the redistribution gains from the policy. When accounting for all terminal states, the right panel

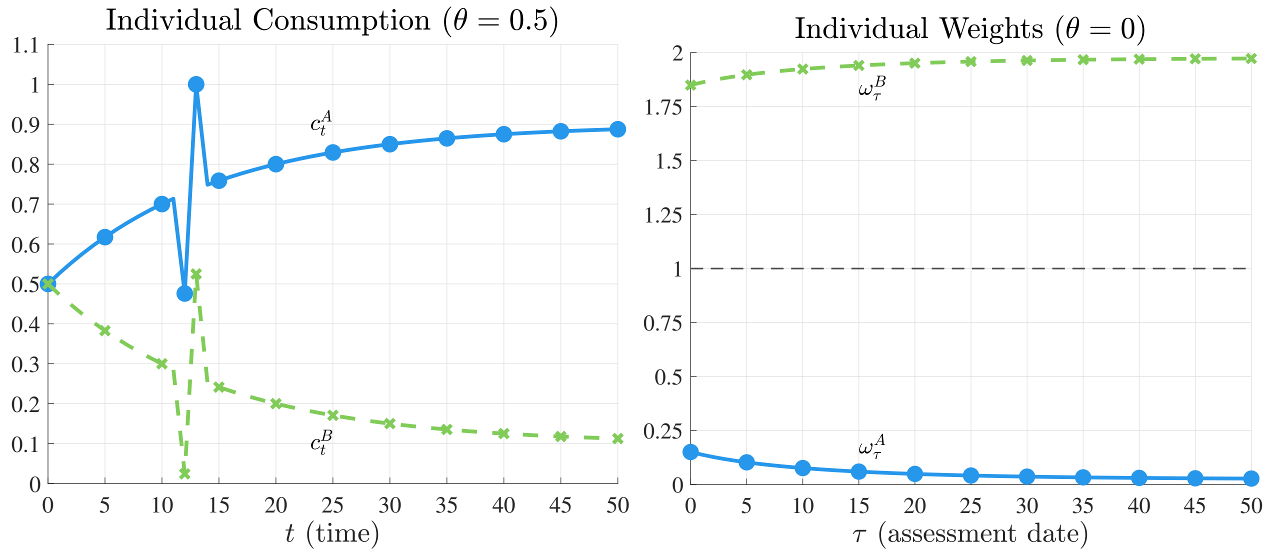

Figure 9: Consumption Profiles and Individual Weights (Application 3)

Note. The left panel shows consumption for individual A (green dashed line) and individual B (solid blue line) after (right panel, $\theta = 0.5$) the aggregate investment policy is implemented. The right panel shows the evolution of the normalized individual weights for individual A (green dashed line) and individual B (solid blued line) as a function of the assessment date for an equal-weighted utilitarian planner when $\theta = 0$.

of Figure 8 shows how the overall risk-sharing gains slowly decrease in magnitude as time goes by, although our previously analysis shows that this effect masks changes in the risk-sharing gains at particular states relative to others.

5.3 Application 3: Aggregate Investment Polices

This application identifies a third form of time inconsistency that arises in heterogeneous agent economies. It shows that persistent trends in consumption inequality — either increasing or decreasing inequality — make efficiency-maximizing planners present biased, reducing their willingness to undertake future *aggregate* investment policies as time unfolds.

Two conditions make an efficiency-maximizing planner present biased. First, consumption trends are needed. Consumption trends endogenously determine whether an individual is more or less patient, which in turn determines the willingness-to-pay for an investment policies, with increasing-consumption individuals being more impatient than shrinking-consumption individuals. Second, individual permanent consumption must change over time, which pins down whether an individual is willing to pay more or less for the policy in relative terms as time unfolds.

With increasing inequality, the shrinking-consumption individual is always more willing to undertake investment policies, but disproportionally less willing to pay for them as time unfolds, as he becomes relatively poorer in permanent terms over time. The growing-consumption individual is instead always less willing to undertake investment policies, but disproportionally more willing to pay for them as time unfolds, as he becomes relatively richer in permanent terms over time. Given these patterns, efficiency gains from aggregate investment policies disproportionally represent the valuation of more impatient individuals as time unfolds, making these policies less attractive over time.

This reduction in aggregate-efficiency gains over time gets compensated by an increase in the redistribution gains: the aggregate investment policy disproportionally benefits the shrinking-consumption individual, and this individual becomes poorer on a permanent basis as time unfolds, increasing the redistribution gains.

5.3.1 Environment

We consider a deterministic infinite-horizon economy with dates $t \in \{0, 1, \ldots\}$. There are two individuals with identical preferences, indexed by $i \in \{A, B\}$.

Preferences and Endowments. Individual i's lifetime utility from the perspective of date k is defined as

$$V_k^i = \sum_{t \geq k} \beta^{t-k} u\left(c_t^i\right).$$

Individual i's endowment at date t is given by

$$y_t^i = \chi_t^i y_t,$$

where $y_t = 1$ denotes the aggregate endowment, and where χ_t^i denotes individual i's endowment share at date t. Individual A's endowment share starts at $\chi_t^A = \underline{\chi}^A$ and exponentially converges to $\chi_t^A = \overline{\chi}^A$, where a parameter $\psi \geq 0$ controls the speed of convergence. Formally,

$$\chi_t^A = e^{-\psi t} \underline{\chi}^A + \left(1 - e^{-\psi t}\right) \overline{\chi}^A.$$

Individual B's endowment share is given by $\chi_t^A = 1 - \chi_t^B$.

Aggregate Investment Policy. Individual i's consumption at date t is given by

$$c_t^i = y_t^i + \theta M_t^i,$$

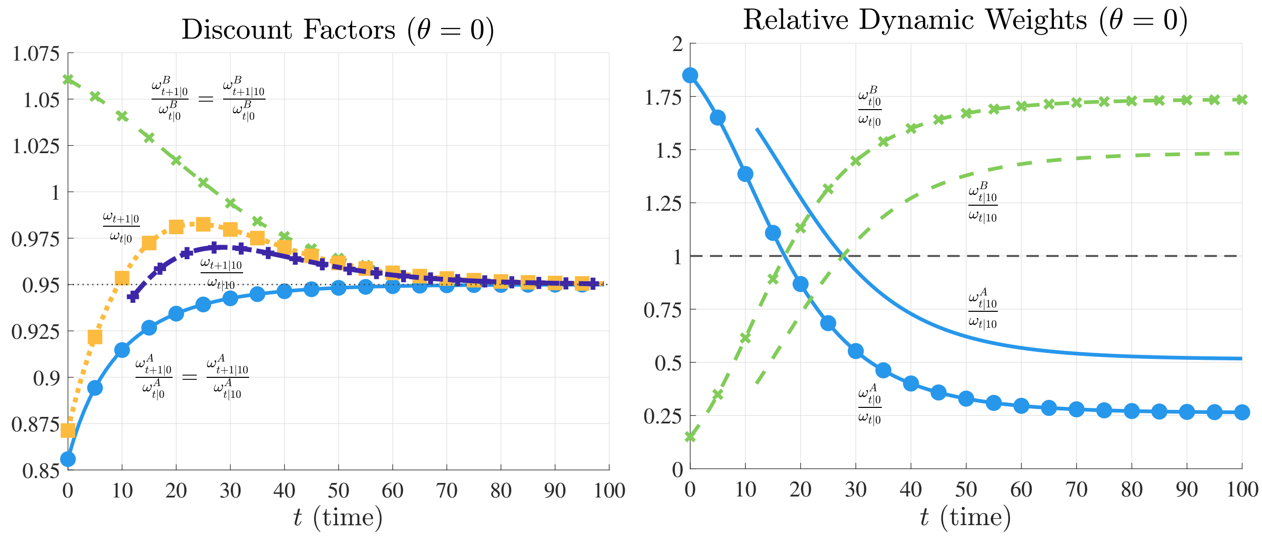

Figure 10: Discount Factors (Application 3)

Note. The left panel shows individual and aggregate discount factors between dates t and $t+1$ as a function of time from two different perspectives: $k=0$ and $k=10$. Individual discount factors, solid blue for A and dashed green for B, are identical regardless of the perspective. Aggregate discount factors between any two dates are lower as time goes by. The right panel shows the relative dynamic weights for both individuals from two different perspectives: $k=0$ and $k=10$. Relative dynamic weights are critical to determine the evolution of aggregate discount factors, as shown in Equation (19).

where M_t^i captures the consumption impact of an aggregate investment policy. There exists the possibility to invest 1 unit of consumption at date $\bar{\tau}$, receiving a return of z units at $\bar{\tau}+1$, where the cost and benefits of investment accrue symmetrically to each individual, so

$$M_{\bar{\tau}}^i = \frac{1}{2} M_{\bar{\tau}} = -1 \quad \text{and} \quad M_{\bar{\tau}+1}^i = \frac{1}{2} M_{\bar{\tau}+1} = z.$$

Parametrization. We assume isoelastic preferences, with $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ and $\gamma = 2$, and interpret a date in the model as one year, setting a discount factor $\beta = 0.95$. We calibrate the convergence parameter to $\psi = (\ln 2)/4$ so that the half-life — the time it takes for the gap between χ_t^A and $\overline{\chi}^A$ to close — is 4 years. The investment policy is made at date $\bar{\tau} = 12$, and has a gross return $z = 1.05$. Given this calibration, we consider policies such that $\theta \in [0, 0.5]$, to ensure that individual consumption is strictly positive.

5.3.2 Weights, Assessments, and Time Inconsistency

We study welfare assessments of the aggregate investment policy under an equal-weighted utilitarian social welfare function. Normalized welfare gains, efficiency and redistribution gains, normalized individual weights, and normalized dynamic weights are defined as in Application

1. Because we now consider a policy that impacts aggregate consumption and symmetrically impacts both individuals, efficiency gains are now solely due to aggregate efficiency. Formally, we can express efficiency gains as

$$\Xi_{|k}^{E} = \Xi_{|k}^{AE} = \sum_{i} \sum_{t \geq k} \omega_{t|k}^{i} \frac{dc_{t}^{i}}{d\theta} = \omega_{\bar{\tau}+1|k} z - \omega_{\bar{\tau}|k}, \quad \text{where} \quad \omega_{\bar{\tau}|k} = \frac{1}{I} \sum_{i} \omega_{\bar{\tau}|k}^{i}, \tag{18}$$

with dynamic weights $\omega_{t|k}^{i}$ defined as in (16), and where the second equality follows from the form of the particular policy considered. It follows directly from (18), that it is desirable to carry the aggregate investment policy from a date k perspective as long as

$$\frac{\omega_{\bar{\tau}+1|k}}{\omega_{\bar{\tau}|k}} > \frac{1}{z}, \quad \text{where} \quad \frac{\omega_{\bar{\tau}+1|k}}{\omega_{\bar{\tau}|k}} = \frac{1}{I} \sum_{i} \frac{\omega_{\bar{\tau}+1}^{i}}{\omega_{\bar{\tau}}^{i}} + \mathbb{C}ov_{i} \left[\frac{\omega_{\bar{\tau}|k}^{i}}{\omega_{\bar{\tau}|k}}, \frac{\omega_{\bar{\tau}+1}^{i}}{\omega_{\bar{\tau}}^{i}} \right]. \tag{19}$$

Since individual preferences are time-consistent in this case, the term $\frac{\omega_{\bar{\tau}+1}^{i}}{\omega_{\bar{\tau}}^{i}}$ does not depend on the perspective k, highlighting that the aggregate discount rate varies depending on the perspective k depending on whether individuals with higher relative dynamic weights for date $\bar{\tau}$ are those with higher discount factors between $\bar{\tau}+1$ and $\bar{\tau}$, as we further illustrate below.

Normalized Individual Weights. The individual weight $\omega_{|k}^{i}$ captures how much the planner values the welfare gains of individual i in units of the welfare numeraire from the perspective of date k. The right panel of Figure 9 shows the evolution of the individual weights over time. Because individual B's consumption share is declining over time, his normalized individual weight increases: the planner assigns more value to the gains of the poorer individual. In contrast, individual A's growing permanent consumption makes his weight decline. This increase in the dispersion of individual weights over time contributes to explaining why redistribution gains in importance as the assessment date advances.

Discount Factors. Because in this application we consider an aggregate investment policy, we show the behavior of the discount factors, which are in turn a function of normalized dynamic weights. The left panel of Figure 10 shows individual and aggregate discount factors between dates $\bar{\tau}$ and $\bar{\tau}+1$ as a function of time from two different perspectives: $k=0$ and $k=10$. Because individual preferences are time-consistent, individual discount factors, $\frac{\omega_{\bar{\tau}+1}^{i}}{\omega_{\bar{\tau}}^{i}}$, are invariant to the time of the assessment. Intuitively, the growing-consumption individual A discounts the future more aggressively than individual B, who has a discount rate even higher than unity.

However, aggregate discount factors between any two dates are lower as time goes by, which generates a form of present bias and is the source of time inconsistency in this application.In order to understand why aggregate discount factors are lower over time, it is useful to show the relative dynamic weights — which are critical to determine the evolution of aggregate discount factors, as shown in Equation (19) — from both perspectives: $k = 0$ and $k = 10$. Relative discount factors capture the relative valuation of an individual between date τ and permanent consumption from different perspectives. Intuitively, as time goes by, individual A values more consumption at a given date relative to permanent consumption, since he is becoming richer. The opposite occurs for individual B. This explains why the aggregate discount rate approaches the individual discount rate of A as time goes by.

Time Inconsistency. Figure 11 illustrates the resulting time inconsistency in the welfare assessments of aggregate investment policies. The left panel show shows the efficiency, redistribution, and total welfare gains from the policy. Ex ante, the efficiency gains from the policy are strictly positive, as the returns from the investment policy are applied a higher discount factor (lower discount rate). As time goes by, the shrinking-consumption individual B — who is is always more willing to undertake investment policies — becomes less willing to pay for the policy, as he becomes relatively poorer in permanent terms over time. Hence, efficiency gains from aggregate investment policies disproportionally represent the valuation of more impatient individuals as time unfolds, making these policies less attractive over time. Figure 11 shows that the investment policy, carried out at date $\bar{\tau} = 12$, is perceived to be efficiency-increasing between dates 0 and 4, but efficiency-decreasing after date 4.

Figure 11 also shows that this reduction in aggregate-efficiency gains over time gets compensated by an increase in the redistribution gains: the aggregate investment policy disproportionally benefits the shrinking-consumption individual (with a higher individual weight), and this individual becomes poorer on a permanent basis as time unfolds, increasing the redistribution gains.

The right panel of Figure 11 reports the implied hurdle rate for an aggregate investment policy when assessed at different times. This is another way of illustrating the time-consistency of efficiency assessments. Initially, at date 0, investments made at date 12 with an (annual) return of roughly 4% or higher are found to be desirable from an efficiency perspective. By date 12, only investment with an (annual) return of roughly 6% or higher are found to be desirable from an efficiency perspective.

We conclude this application with two remarks.

Remark. (Time Inconsistency with Shrinking Inequality) The same form of present bias time inconsistency applies in economies in which individual consumption profiles converge. As in

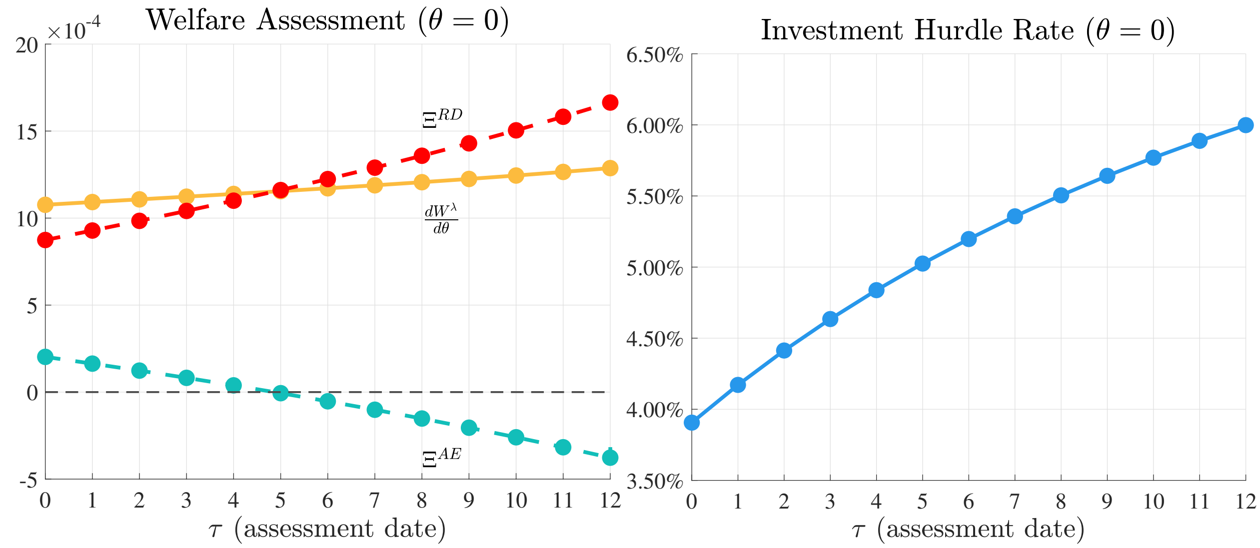

Figure 11: Welfare Assessment and Investment Hurdle Rate (Application 3)

Note. The left panel shows the welfare (yellow), efficiency/aggregate-efficiency (light green), and redistribution (red) assessments of the aggregate investment policy as a function of the assessment date when $\theta = 0$. Over time, the efficiency gains from the policy shrink and turn negative, illustrating the present bias and time inconsistency of the planner's efficiency assessment. The right panel shows the corresponding efficiency-maximizing investment hurdle rate as a function of the assessment date, that is, the rate of return of that makes the investment desirable from an efficiency perspective. As time goes by, the hurdle rate increases, illustrating the present bias and time inconsistency of the planner's efficiency assessment.

the growing inequality case, the shrinking-consumption individual is always more willing to undertake investment policies, but disproportionally less willing to pay for them as time unfolds, as he becomes relatively poorer in permanent terms over time. The growing-consumption individual is instead always less willing to undertake investment policies, but disproportionally more willing to pay for them as time unfolds, as he becomes relatively richer in permanent terms over time. Efficiency gains from aggregate investment policies disproportionally represent the valuation of more impatient individuals as time unfolds, making these policies less attractive over time.

Remark. (Relation to Jackson and Yariv (2014, 2015)) Jackson and Yariv (2014, 2015) show that when individuals have heterogeneous discount factors, utilitarian welfare assessments are time inconsistent. In this application, individuals have identical preferences so utilitarian welfare assessments are time-consistent. However, the inconsistency of efficiency and redistribution components in this applications has a resemblance to the Jackson and Yariv (2014, 2015) result. Intuitively, when individuals have identical preferences but experience different consumption profiles, they endogenously have different discount factors when expressed in consumption units, not utils. This difference in valuations is the source of the time inconsistency identified here.

6 Conclusion

This paper systematically studies the time inconsistency of welfare assessments in economies with heterogeneous agents, establishing three central lessons for applied work.

First, even when individual preferences are identical and time-consistent (Strotz, 1956), and forward-looking behavior of the form identified in Kydland and Prescott (1977) is absent, only utilitarian social welfare functions ensure the time consistency of welfare assessment due to interpersonal welfare comparisons. Non-linear social welfare functions lead to time inconsistency because the weights assigned by a planner to different individuals' utility gains vary over time.

Second, it is generally not possible to make time-consistent efficiency assessments that satisfy the compensation principle in heterogeneous agent incomplete markets environments, even when all other requirements for consistency are met. Therefore, while a utilitarian planner consistently assesses policies, the justification on grounds of efficiency and redistribution varies as time passes and uncertainty is realized. For instance, a policy justified ex ante on efficiency grounds, perhaps by improving risk-sharing, may be desirable ex post solely for redistributive reasons. This impossibility result requires the incompleteness of markets and the use of forward-looking units to express efficiency assessments, conditions that lead to changes over time in the rates at which individuals translate utility gains into consumption equivalents.

Finally, all sources of efficiency gains are prone to time-inconsistency problems when markets are incomplete — unless a perturbation solely impacts aggregate net consumption at a single point in time. Applications to i) anticipated relief policies, ii) risk-sharing policies, and iii) aggregate investment policies illustrate how efficiency-based interventions are systematically challenged by the evolving nature of welfare assessments.

References

Adams, A., L. Cherchye, B. De Rock, and E. Verriest (2014): "Consume now or later? Time inconsistency, collective choice, and revealed preference," *American Economic Review*, 104(12), 4147–4183.

Aiyagari, S. R., A. Marcet, T. J. Sargent, and J. Seppälä (2002): "Optimal taxation without state-contingent debt," *Journal of Political Economy*, 110(6), 1220–1254.

Alesina, A., and D. Rodrik (1994): "Distributive politics and economic growth," *The quarterly journal of economics*, 109(2), 465–490.

Amador, M., I. Werning, and G. M. Angeletos (2006): "Commitment vs. flexibility," *Econometrica*, 74(2), 365–396.

Arrow, K. J. (1950): "A difficulty in the concept of social welfare," *Journal of political economy*, 58(4), 328–346.

——— (1963): "Uncertainty and the Welfare Economics of Medical Care," *American Economic Review*, 53(5), 941–973.

Athey, S., A. Atkeson, and P. J. Kehoe (2005): "The optimal degree of discretion in monetary policy," *Econometrica*, 73(5), 1431–1475.

Atkinson, A. B. (1970): "On the measurement of inequality," *Journal of economic theory*, 2(3), 244–263.

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

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021): "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 89(5), 2375–2408.

Auclert, A., M. Cai, M. Rognlie, and L. Straub (2024): "Optimal long-run fiscal policy with heterogeneous agents," .

Barro, R. J., and D. B. Gordon (1983): "Rules, discretion and reputation in a model of monetary policy," *Journal of Monetary Economics*, 12(1), 101–121.

Benabou, R. (2002): "Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?," *Econometrica*, 70(2), 481–517.

Bilbiie, F. O. (2024): "Monetary policy and heterogeneity: An analytical framework," *Review of Economic Studies*, p. rdae066.

Bisin, A., A. Lizzeri, and L. Yariv (2015): "Government policy with time inconsistent voters," *American Economic Review*, 105(6), 1711–1737.

Boskin, M. J., and E. Sheshinski (1978): "Optimal redistributive taxation when individual welfare depends upon relative income," *The Quarterly Journal of Economics*, 92(4), 589–601.

Chambers, C. P., and F. Echenique (2018): "On multiple discount rates," *Econometrica*, 86(4), 1325–1346.

Chang, R. (2022): "Should Central Banks Have an Inequality Objective?," Discussion paper, National Bureau of Economic Research.

Chari, V. V., and P. J. Kehoe (1990): "Sustainable plans," *Journal of political economy*, 98(4), 783–802.

Clayton, C., and A. Schaab (2025): "A theory of dynamic inflation targets," *American Economic Review*, 115(2), 448–490.

Dávila, E., and A. Schaab (2023a): "Optimal monetary policy with heterogeneous agents: discretion, commitment, and timeless policy," .

——— (2023b): "Welfare Accounting," .

——— (2024): "Welfare Assessments with Heterogeneous Individuals," .

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

Dovis, A., and R. Kirpalani (2020): "Fiscal rules, bailouts, and reputation in federal governments," *American Economic Review*, 110(3), 860–888.

——— (2021): "Rules without commitment: Reputation and incentives," *The Review of Economic Studies*, 88(6), 2833–2856.

Farhi, E., C. Sleet, I. Werning, and S. Yeltekin (2012): "Non-linear capital taxation without commitment," *Review of Economic Studies*, 79(4), 1469–1493.

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

Feldstein, M. S. (1964): "The social time preference discount rate in cost benefit analysis," *The economic journal*, 74(294), 360–379.

Golosov, M., M. Troshkin, and A. Tsyvinski (2016): "Redistribution and social insurance," *American Economic Review*, 106(2), 359–386.

Halac, M., and P. Yared (2014): "Fiscal Rules and Discretion Under Persistent Shocks," *Econometrica*, 82(5), 1557–1614.

Halevy, Y. (2015): "Time consistency: Stationarity and time invariance," *Econometrica*, 83(1), 335–352.

Heathcote, J., K. Storesletten, and G. L. Violante (2017): "Optimal tax progressivity: An analytical framework," *The Quarterly Journal of Economics*, 132(4), 1693–1754.

Jackson, M. O., and L. Yariv (2014): "Present bias and collective dynamic choice in the lab," *American Economic Review*, 104(12), 4184–4204.

——— (2015): "Collective dynamic choice: the necessity of time inconsistency," *American*

Economic Journal: Microeconomics, 7(4), 150–178.

Kaneko, M., and K. Nakamura (1979): "The Nash social welfare function," *Econometrica: Journal of the Econometric Society*, pp. 423–435.

Kaplow, L., and S. Shavell (2001): "Any non-welfarist method of policy assessment violates the Pareto principle," *Journal of Political Economy*, 109(2), 281–286.

Klein, P., P. Krusell, and J.-V. Rios-Rull (2008): "Time-consistent public policy," *The Review of Economic Studies*, 75(3), 789–808.

Knight, F. H. (1921): *Risk, uncertainty and profit*. Houghton Mifflin.

Koopmans, T. C. (1960): "Stationary ordinal utility and impatience," *Econometrica: Journal of the Econometric Society*, pp. 287–309.

Krusell, P., V. Quadrini, and J.-V. Ríos-Rull (1997): "Politico-economic equilibrium and economic growth," *Journal of Economic Dynamics and Control*, 21(1), 243–272.

Kydland, F. E., and E. C. Prescott (1977): "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85(3), 473–491.

Ljungqvist, L., and T. J. Sargent (2018): *Recursive macroeconomic theory*. MIT press.

Lucas, R. E. (1987): *Models of business cycles*, vol. 26. Basil Blackwell Oxford.

Lucas, R. E., and N. L. Stokey (1983): "Optimal fiscal and monetary policy in an economy without capital," *Journal of monetary Economics*, 12(1), 55–93.

Marcet, A., and R. Marimon (2019): "Recursive contracts," *Econometrica*, 87(5), 1589–1631.

Marglin, S. A. (1963): "The social rate of discount and the optimal rate of investment," *The Quarterly Journal of Economics*, 77(1), 95–111.

Marshall, A. (1920): *Principles of Economics, 8th edn (1920).*

Millner, A., and G. Heal (2018): "Time consistency and time invariance in collective intertemporal choice," *Journal of Economic Theory*, 176, 158–169.

Mirrlees, J. A. (1971): "An exploration in the theory of optimum income taxation," *The review of economic studies*, 38(2), 175–208.

Mongin, P. (1995): "Consistent bayesian aggregation," *Journal of Economic Theory*, 66(2), 313–351.

——— (1998): "The paradox of the Bayesian experts and state-dependent utility theory," *Journal of Mathematical Economics*, 29(3), 331–361.

Nash, J. F. (1950): "The bargaining problem," *Econometrica*, 18(2), 155–162.

Persson, T., and G. E. Tabellini (1994): "Is inequality harmful for growth?," *The*

American Economic Review, 84(3), 600–621.

Pigou, A. (1920): *The economics of welfare*. Macmillan.

Piketty, T., and E. Saez (2012): "A theory of optimal capital taxation," Discussion paper, National Bureau of Economic Research.

——— (2013a): "A theory of optimal inheritance taxation," *Econometrica*, 81(5), 1851–1886.

——— (2013b): "Optimal labor income taxation," in *Handbook of public economics*, vol. 5, pp. 391–474. Elsevier.

Rawls, J. (1971): "A theory of justice," .

Rogoff, K. (1985): "The Optimal Degree of Commitment to an Intermediate Monetary Target," *The Quarterly Journal of Economics*, 100(4), 1169–1189.

Saez, E. (2001): "Using elasticities to derive optimal income tax rates," *The review of economic studies*, 68(1), 205–229.

Saez, E., and S. Stantcheva (2016): "Generalized social marginal welfare weights for optimal tax theory," *American Economic Review*, 106(01), 24–45.

Stern, N. H. (1976): "On the specification of models of optimum income taxation," *Journal of Public economics*, 6(1-2), 123–162.

Strotz, R. H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 23(3), 165–180.

Sublet, G. (2023): "The Optimal Degree of Discretion in Fiscal Policy," .

Varian, H. R. (1980): "Redistributive taxation as social insurance," *Journal of public Economics*, 14(1), 49–68.

Werning, I. (2007): "Optimal fiscal policy with redistribution," *The Quarterly Journal of Economics*, 122(3), 925–967.

Zuber, S. (2011): "The aggregation of preferences: can we ignore the past?," *Theory and decision*, 70, 367–384.

Online Appendix

A Proofs and Derivations: Section 2

Equilibrium Characterization at $t = 0$. Individual optimization implies that

$$u^i_{\ell,0} = -u^i_{c,0} w_0$$
$$u^i_{\ell,1} = -u^i_{c,1}(1-\tau) w_1 z^i_1.$$

Profit maximization implies that $w_t = 1$. Note that the choice of the optimal tax has no impact on the allocation in period 0.

Ex-Ante Optimal Policy. From a date-0 perspective, the Ramsey problem is to choose a tax rate τ to maximize social welfare subject to the competitive equilibrium conditions. At an optimum, the necessary optimality condition is $\frac{dW_0}{d\tau} = 0$, which can be equivalently expressed in units of welfare numeraire as

$$\frac{1}{\int_0^1 \lambda^i_0 \alpha^i_0 \, di} \frac{dW_0}{d\tau} = \int_0^1 \frac{1}{\lambda^i_0} \frac{dV^i_0}{d\tau} \, di + \mathbb{C}ov_i \left[\omega^i_0, \frac{dV^i_0}{d\tau} \right] = 0,$$

where normalized individual weights are given by $\omega^i_0 = \frac{\alpha^i_0 \lambda^i_0}{\int_0^1 \alpha^i_0 \lambda^i_0 \, di}$. This expression decomposes the welfare gain $\frac{dW_0}{d\tau}$ into an efficiency and a redistribution gain. For any choice of social welfare function $\mathcal{W}(\cdot)$ and for any choice of welfare ex ante numeraire λ^i_0, the RHS of the above covariance is constant across i because individuals are identical from a date-0 perspective. Therefore, the impact of anticipated tax policy is symmetric across individuals. The covariance term thus vanishes.

Next, we plug in for lifetime utility and rewrite the Ramsey FOC as

$$\int_0^1 \frac{1}{\lambda^i_0} \frac{dV^i_0}{d\tau} \, di = \int_0^1 \frac{1}{\lambda^i_0} \beta \, \mathbb{E}_0 \left[u^i_{c,1} \frac{dc^i_1}{d\tau} + u^i_{\ell,1} \frac{d\ell^i_1}{d\tau} \right] di = \int_0^1 \frac{1}{\lambda^i_0} \beta \, \mathbb{E}_0 \, u^i_{c,1} \left[\frac{dc^i_1}{d\tau} - (1-\tau) z^i_1 \frac{d\ell^i_1}{d\tau} \right] di,$$

where we use the fact that tax policy has no impact on the date 0 allocation, and the third line uses household i's date 1 FOC. Now we decompose the valuation as follows

$$
\begin{aligned}
0 &= \int_0^1 \mathbb{E}_0 \frac{\beta u_{c,1}^i}{\lambda_0^i} \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] di \\
&= \int_0^1 \mathbb{E}_0 \frac{\beta u_{c,1}^i}{\lambda_0^i} \frac{(\mathbb{E}_0 \beta u_{c,1}^i)}{(\mathbb{E}_0 \beta u_{c,1}^i)} \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] di \\
&= \int_0^1 \mathbb{E}_0 \left\{ \frac{\beta (\mathbb{E}_0 u_{c,1}^i)}{\lambda_0^i} \frac{u_{c,1}^i}{(\mathbb{E}_0 u_{c,1}^i)} \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] \right\} di,
\end{aligned}
$$

where sometimes for convenience we use the shorthand notation

$$
\omega_1^i = \frac{\beta (\mathbb{E}_0 u_{c,1}^i)}{\lambda_0^i} \qquad \text{and} \qquad \tilde{\omega}_1^i = \frac{u_{c,1}^i}{(\mathbb{E}_0 u_{c,1}^i)}
$$

where $\tilde{\omega}_1^i$ is a random variable from a date-0 perspective, whereas ω_1^i is not.

Next, we swap the expectation operator and perform a cross-sectional covariance decomposition

$$
\begin{aligned}
0 &= \mathbb{E}_0 \int_0^1 \frac{\beta (\mathbb{E}_0 u_{c,1}^i)}{\lambda_0^i} \frac{u_{c,1}^i}{(\mathbb{E}_0 u_{c,1}^i)} \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] di \\
&= \left(\int_0^1 \frac{\beta (\mathbb{E}_0 u_{c,1}^i)}{\lambda_0^i} di \right) \left(\int_0^1 \frac{u_{c,1}^i}{(\mathbb{E}_0 u_{c,1}^i)} \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] di \right) \\
&\quad + \mathbb{C}ov_i \left(\frac{\beta (\mathbb{E}_0 u_{c,1}^i)}{\lambda_0^i}, \frac{u_{c,1}^i}{(\mathbb{E}_0 u_{c,1}^i)} \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] \right)
\end{aligned}
$$

Notice that $\frac{\beta (\mathbb{E}_0 u_{c,1}^i)}{\lambda_0^i}$ is constant across all i. That's because households are symmetric from the perspective of period 0: they behave exactly the same in period 0, and they are still the same in expectation in period 1. Therefore, the covariance term here vanishes. And we are left with

$$
0 = \bar{\omega}_1 \int_0^1 \frac{u_{c,1}^i}{\mathbb{E}_0 u_{c,1}^i} \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] di
$$

where we set

$$
\bar{\omega}_1 = \frac{\beta (\mathbb{E}_0 u_{c,1}^i)}{\lambda_0^i},
$$

which is the same for all i.

One final cross-sectional covariance decomposition yields

$$
0 = \bar{\omega}_1 \left(\int_0^1 \frac{u_{c,1}^i}{(\mathbb{E}_0 u_{c,1}^i)} di \right) \int_0^1 \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] di + \bar{\omega}_1 \mathbb{C}ov_i \left(\frac{u_{c,1}^i}{(\mathbb{E}_0 u_{c,1}^i)}, \frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right).
$$

Notice that

$$\int_0^1 \frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}\, di = 1$$

due to a law of large numbers.

Now notice that the individual budget constraint implies

$$\frac{dc^i_1}{d\tau} = (1-\tau) z^i_1 \frac{d\ell^i_1}{d\tau} - z^i_1 \ell^i_1 + \frac{dT}{d\tau},$$

where $\frac{dT}{d\tau} = L_1 + \tau \frac{dL_1}{d\tau}$. Therefore, we have

$$\begin{aligned} 0 &= \bar{\omega}_1 \int_0^1 \left[- z^i_1 \ell^i_1 + \frac{dT}{d\tau} \right] di + \bar{\omega}_1 \mathbb{C}ov_i \left(\frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}, -z^i_1 \ell^i_1 \right) \\ &= \bar{\omega}_1 \tau \frac{dL_1}{d\tau} + \bar{\omega}_1 \mathbb{C}ov_i \left(\frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}, -z^i_1 \ell^i_1 \right), \end{aligned}$$

where the second term in the first equality follows because dT is constant across individuals and therefore drops out from the covariance. We can now write this as

$$0 = \tau \frac{dL_1}{d\tau} - \mathbb{C}ov_i \left(\frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}, z^i_1 \ell^i_1 \right),$$

which corresponds to equation (4) in the main text.

Ex Post Optimal Policy. The relevant optimality condition at $t = 1$ is $u^i_{\ell,1} = -u^i_{c,1}(1-\tau) z^i_1$. At an optimum, the necessary optimality condition in $\frac{dW_1}{d\tau} = 0$, which can be written as

$$\frac{1}{\int_0^1 \alpha^i_1 \lambda^i_1\, di} \frac{dW_1}{d\tau} = \int_0^1 \frac{1}{\lambda^i_1} \frac{dV^i_1}{d\tau}\, di + \mathbb{C}ov_i \left[\omega^i_1, \frac{1}{\lambda^i_1} \frac{dV^i_1}{d\tau} \right],$$

where we define $\omega^i_1 = \frac{\alpha^i_1 \lambda^i_1}{\int_0^1 \alpha^i_1 \lambda^i_1\, di}$ from the perspective of date 1. The only reasonable numeraire in this case is period 1 consumption, so we have

$$\lambda^i_1 = u^i_{c,1}.$$

We can write efficiency as

$$\begin{aligned} \int_0^1 \frac{1}{\lambda^i_1} \frac{dV^i_1}{d\tau}\, di &= \int_0^1 \frac{1}{\lambda^i_1} \left[u^i_{c,1} \frac{dc^i_1}{d\tau} + u^i_{\ell,1} \frac{d\ell^i_1}{d\tau} \right] di \\ &= \int_0^1 \frac{1}{\lambda^i_1} u^i_{c,1} \left[\frac{dc^i_1}{d\tau} - (1-\tau) z^i_1 \frac{d\ell^i_1}{d\tau} \right] di \\ &= \int_0^1 \left[\frac{dc^i_1}{d\tau} - (1-\tau) z^i_1 \frac{d\ell^i_1}{d\tau} \right] di, \end{aligned}$$

where the last line follows due to $\lambda_1^i = u_{c,1}^i$. We now use the individual budget constraint, then the government budget constraint, and aggregate, yielding

$$\begin{aligned}\int_0^1 \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau}\right] di &= \int_0^1 \left[- z_1^i \ell_1^i + \frac{dT}{d\tau}\right] di \\ &= \int_0^1 \left[- z_1^i \ell_1^i + L_1 + \tau \frac{dL_1}{d\tau}\right] di \\ &= \tau \frac{dL_1}{d\tau}.\end{aligned}$$

Therefore, we have the ex post optimality condition

$$\begin{aligned}0 &= \tau \frac{dL_1}{d\tau} + \mathbb{C}ov_i \left[\frac{\alpha_1^i u_{c,1}^i}{\int_0^1 \alpha_1^i u_{c,1}^i \, di}, \frac{1}{\lambda_1^i} \frac{dV_1^i}{d\tau}\right] \\ &= \tau \frac{dL_1}{d\tau} + \mathbb{C}ov_i \left[\frac{\alpha_1^i u_{c,1}^i}{\int_0^1 \alpha_1^i u_{c,1}^i \, di}, \frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau}\right] \\ &= \tau \frac{dL_1}{d\tau} + \mathbb{C}ov_i \left[\frac{\alpha_1^i u_{c,1}^i}{\int_0^1 \alpha_1^i u_{c,1}^i \, di}, -z_1^i \ell_1^i\right].\end{aligned}$$

And finally, we have

$$\alpha_1^i = \frac{\partial \mathcal{W}(\{V_1^i\}_i)}{\partial V_1^i} = W_1^{\phi} (V_1^i)^{-\phi}.$$

This leaves us with

$$0 = \tau \frac{dL_1}{d\tau} + \mathbb{C}ov_i \left[\frac{(V_1^i)^{-\phi} u_{c,1}^i}{\int_0^1 (V_1^i)^{-\phi} u_{c,1}^i \, di}, -z_1^i \ell_1^i\right],$$

which corresponds to equation (5) in the main text.

B Proofs and Derivations: Section 4

Proof of Proposition 1. (Inconsistency of Welfare Assessments: Single Agent)

Proof. The result is a special case of Proposition 2. With a representative agent and $I = 1$, we have $\alpha^i_{s^t} = 1$ without loss for all s^t. Therefore, we can decompose the ex ante assessment as

$$\begin{aligned}\frac{dW_{s^0}}{d\theta} &= \beta^k \pi(s^k) \frac{dW_{s^k}}{d\theta} \\ &+ \beta^k \pi(s^k) \left(\frac{dV_{s^0}(s^k)}{d\theta} - \frac{dV_{s^k}}{d\theta} \right) \\ &+ \sum_{t \geq k} \sum_{s^t \ngeq s^k} \beta^t \pi(s^t) u'(c_t(s^t)) \frac{\partial \mathcal{C}_t(s^t \mid s^0)}{\partial \theta} \\ &+ \sum_{0 \leq t < k} \sum_{s^t} \beta^t \pi(s^t) u'(c_t(s^t)) \frac{\partial \mathcal{C}_t(s^t \mid s^0)}{\partial \theta}.\end{aligned}$$

Therefore, the only way in which $\frac{dW_{s^0}}{d\theta} \neq \beta^t \pi(s^t) \frac{dW_{s^t}}{d\theta}$ is for the second, third or fourth lines to be non-zero. This requires that one of the conditions (i) – (iii) in Proposition 1 is satisfied. □

Proof of Proposition 2. (Inconsistency of Welfare Assessments: Heterogeneous Agents)

Proof. We prove the result for welfare assessments from a date-0 perspective (without loss of generality) to make the notation easier. Consider three dates $0 < k \leq T$. We compare welfare assessments of the perturbation $d\theta_T(s^T)$ for a particular history s^T. We simply refer to this perturbation as "$d\theta$", again to make the notation easier. From a date-0 perspective and history s^0, we have

$$\frac{dW_{s^0}}{d\theta} = \sum_i \alpha^i_{s^0} \frac{dV^i_{s^0}}{d\theta} = \sum_i \alpha^i_{s^0} \sum_t \sum_{s^t} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t, \boldsymbol{\theta} \mid s^0)}{\partial \theta}.$$

From the perspective of some history s^k at the later date k, we have

$$\frac{dW_{s^k}}{d\theta} = \sum_i \alpha^i_{s^k} \frac{dV^i_{s^k}}{d\theta} = \sum_i \alpha^i_{s^k} \sum_{t \geq k} \sum_{s^t \geq s^k} (\beta^i)^{t-k} \pi(s^t \mid s^k) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t, \boldsymbol{\theta} \mid s^k)}{\partial \theta}.$$

Notice that we simply use $c^i_t(s^t)$ when writing marginal utility rather than the consumption function. When evaluated at the status quo, the consumption allocation (in levels rather than changes) does not depend on the perspective of the assessment, which lets us write $\mathcal{C}^i_t(s^t, \bar{\boldsymbol{\theta}} \mid s^0) = \mathcal{C}^i_t(s^t, \bar{\boldsymbol{\theta}} \mid s^k) = c^i_t(s^t)$, when evaluated at a particular $\bar{\boldsymbol{\theta}}$.

We now derive an expression that constructively isolates each potential source of time inconsistency. □

First, we look at the *interpersonal welfare comparison* source. Notice that we can write the ex ante assessment as

$$\frac{dW_{s^0}}{d\theta} = \sum_i \frac{\alpha^i_{s^0}}{\alpha^i_{s^k}} \alpha^i_{s^k} \sum_t \sum_{s^t} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial \theta}$$
$$= \left(\frac{1}{I} \sum_i \frac{\alpha^i_{s^0}}{\alpha^i_{s^k}} \right) \sum_i \alpha^i_{s^k} \sum_t \sum_{s^t} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial \theta} + \mathbb{C}ov^{\Sigma}_i \left(\frac{\alpha^i_{s^0}}{\alpha^i_{s^k}},\ \alpha^i_{s^k} \frac{dV^i_{s^0}}{d\theta} \right).$$

Next, we look at the *Kydland-Prescott* source .We now split the sum $\sum_t \sum_{s^t}$ of the first term. We can write

$$\frac{dW_{s^0}}{d\theta} = \left(\frac{1}{I} \sum_i \frac{\alpha^i_{s^0}}{\alpha^i_{s^k}} \right) \Bigg\{ \sum_i \alpha^i_{s^k} \sum_{t \geq k} \sum_{s^t \geq s^k} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial \theta}$$
$$+ \sum_i \alpha^i_{s^k} \sum_{t \geq k} \sum_{s^t \ngeq s^k} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial \theta}$$
$$+ \sum_i \alpha^i_{s^k} \sum_{0 \leq t < k} \sum_{s^t} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial \theta} \Bigg\}$$
$$+ \mathbb{C}ov^{\Sigma}_i \left(\frac{\alpha^i_{s^0}}{\alpha^i_{s^k}},\ \alpha^i_{s^k} \frac{dV^i_{s^0}}{d\theta} \right).$$

Notice that the second and third lines correspond precisely to conditions (i) and (ii) of Proposition 1.

Next, we define the continuation lifetime utility from s^k onwards but from the perspective of s^0 as

$$V^i_{s^0}(s^k) = \sum_{t \geq k} \sum_{s^t \geq s^k} (\beta^i)^{t-k} \pi(s^t \mid s^k) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial \theta}.$$

In other words, the only difference between this and the actual lifetime utility from the perspective of s^k, $V^i_{s^k}$, is that the consumption function is evaluated from different perspectives.

Proof. Also notice that for each individual i we have

$$(\beta^i)^t \pi(s^t) = (\beta^i)^k \pi(s^k) \cdot (\beta^i)^{t-k} \pi(s^t \mid s^k).$$

Using this, we can write

$$\frac{dW_{s^0}}{d\theta} = \left(\frac{1}{I}\sum_i \frac{\alpha^i_{s^0}}{\alpha^i_{s^k}}\right)\Bigg\{ \sum_i \alpha^i_{s^k} \sum_{t\geq k}\sum_{s^t\geq s^k} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^k)}{\partial\theta}$$
$$+ \sum_i \alpha^i_{s^k} \sum_{t\geq k}\sum_{s^t\geq s^k} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \left(\frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial\theta} - \frac{\partial \mathcal{C}^i_t(s^t \mid s^k)}{\partial\theta}\right)$$
$$+ \sum_i \alpha^i_{s^k} \sum_{t\geq k}\sum_{s^t\ngeq s^k} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial\theta}$$
$$+ \sum_i \alpha^i_{s^k} \sum_{0\leq t<k}\sum_{s^t} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial\theta}\Bigg\}$$
$$+ \mathbb{C}ov^{\Sigma}_i\left(\frac{\alpha^i_{s^0}}{\alpha^i_{s^k}},\ \alpha^i_{s^k}\frac{dV^i_{s^0}}{d\theta}\right),$$

or simply

$$\frac{dW_{s^0}}{d\theta} = \left(\frac{1}{I}\sum_i \frac{\alpha^i_{s^0}}{\alpha^i_{s^k}}\right)\Bigg\{ \sum_i \alpha^i_{s^k} (\beta^i)^k \pi(s^k) \frac{dV^i_{s^k}}{d\theta}$$
$$+ \sum_i \alpha^i_{s^k} (\beta^i)^k \pi(s^k) \left(\frac{dV^i_{s^0}(s^k)}{d\theta} - \frac{dV^i_{s^k}}{d\theta}\right)$$
$$+ \sum_i \alpha^i_{s^k} \sum_{t\geq k}\sum_{s^t\ngeq s^k} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial\theta}$$
$$+ \sum_i \alpha^i_{s^k} \sum_{0\leq t<k}\sum_{s^t} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial\theta}\Bigg\}$$
$$+ \mathbb{C}ov^{\Sigma}_i\left(\frac{\alpha^i_{s^0}}{\alpha^i_{s^k}},\ \alpha^i_{s^k}\frac{dV^i_{s^0}}{d\theta}\right),$$

where the second line corresponds to condition (iii) of Proposition 1 and the last line corresponds to condition (iv) of Proposition 2. □

Common discount factor. We conclude the proof by assuming a common discount factor $\beta^i = \beta$ as in the main text. In that case, we have

$$\sum_i \alpha^i_{s^k}(\beta^i)^k\pi(s^k)\frac{dV^i_{s^k}}{d\theta} = \beta^k\pi(s^k)\sum_i \alpha^i_{s^k}\frac{dV^i_{s^k}}{d\theta} = \beta^k\pi(s^k)\frac{dW_{s^k}}{d\theta}.$$

Therefore, the first line above becomes

$$\underbrace{\left(\frac{1}{I}\sum_i \frac{\alpha^i_{s^0}}{\alpha^i_{s^k}}\right)\beta^k\pi(s^k)}_{>0}\frac{dW_{s^k}}{d\theta}.$$

This tells us that we can express the ex ante welfare assessment $\frac{dW_{s^0}}{d\theta}$ as a function of the ex post assessment $\frac{dW_{s^k}}{d\theta}$. The scalar coefficient here is necessarily strictly positive. So this term will never be a source of time inconsistency. This concludes the proof. Heterogeneity in β^i implies that aggregation of discount factors can lead to a time consistency problem as discussed in Jackson and Yariv (2014, 2015).

Proof of Proposition 3. (Inconsistency of Efficiency Assessments)

Proof. We consider the Kaldor-Hicks efficiency assessment of a perturbation $d\theta_T(s^T)$, which we refer to simply as "$d\theta$" to make the notation easier. As in previous proofs, we compare the ex ante assessment at s^0 with ex post assessments at arbitrary histories s^k.

The ex ante efficiency assessment can be written as

$$\begin{aligned}
\Xi^E_{s^0} &= \sum_i \frac{1}{\lambda^i_{s^0}} \frac{dV^i_{s^0}}{d\theta} \\
&= \sum_i \frac{\lambda^i_{s^k}}{\lambda^i_{s^0}} \frac{1}{\lambda^i_{s^k}} \frac{dV^i_{s^0}}{d\theta} \\
&= \left(\frac{1}{I} \sum_i \frac{\lambda^i_{s^k}}{\lambda^i_{s^0}} \right) \sum_i \frac{1}{\lambda^i_{s^k}} \frac{dV^i_{s^0}}{d\theta} + \mathbb{C}ov^{\Sigma}_i \left(\frac{\lambda^i_{s^k}}{\lambda^i_{s^0}}, \frac{1}{\lambda^i_{s^k}} \frac{dV^i_{s^0}}{d\theta} \right).
\end{aligned}$$

The covariance term corresponds precisely to condition (v) of Proposition 3.

We can now rewrite the first term, splitting the sum, as

$$\begin{aligned}
\Xi^E_{s^0} = \left(\frac{1}{I} \sum_i \frac{\lambda^i_{s^k}}{\lambda^i_{s^0}} \right) \Bigg\{ & \sum_i \frac{1}{\lambda^i_{s^k}} \beta^k \pi(s^k) \frac{dV^i_{s^k}}{d\theta} \\
& + \sum_i \frac{1}{\lambda^i_{s^k}} \beta^k \pi(s^k) \left(\frac{dV^i_{s^0}(s^k)}{d\theta} - \frac{dV^i_{s^k}}{d\theta} \right) \\
& + \sum_i \frac{1}{\lambda^i_{s^k}} \sum_{t \geq k} \sum_{s^t \not\succeq s^k} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial \theta} \\
& + \sum_i \frac{1}{\lambda^i_{s^k}} \sum_{0 \leq t < k} \sum_{s^t} (\beta^i)^t \pi(s^t) u'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid s^0)}{\partial \theta} \Bigg\} \\
& + \mathbb{C}ov^{\Sigma}_i \left(\frac{\lambda^i_{s^k}}{\lambda^i_{s^0}}, \frac{1}{\lambda^i_{s^k}} \frac{dV^i_{s^0}}{d\theta} \right).
\end{aligned}$$

Lines two, three and four correspond precisely to conditions (i) – (iii) of Proposition 1. And the last line corresponds to the new numeraire inconsistency condition (v) in Proposition 3.

Finally, notice that Assumption 4.1.2, which rules out time inconsistency á la Kydland-

Prescott, implies

$$\frac{dV^i_{s^0}}{d\theta} = \beta^k \pi(s^k) \frac{dV^i_{s^k}}{d\theta}.$$

Under 4.1.2, therefore, we have

$$\Xi^E_{s^0} = \left(\frac{1}{I} \sum_i \frac{\lambda^i_{s^k}}{\lambda^i_{s^0}} \right) \beta^k \pi(s^k) \Xi^E_{s^k} + \beta^k \pi(s^k) \mathbb{C}ov^{\Sigma}_i \left(\frac{\lambda^i_{s^k}}{\lambda^i_{s^0}}, \frac{1}{\lambda^i_{s^k}} \frac{dV^i_{s^k}}{d\theta} \right)$$

corresponding to equation (13) in the main text. □

Proof of Theorem. (Impossibility of Time-Consistent Efficiency Assessments)

Proof. It is sufficient to prove this result under Assumption 4.1.2, which rules out time consistency problems á la Kydland and Prescott (1977). Under 4.1.2, time inconsistency is therefore governed by the value of

$$\mathbb{C}ov_i \left[\frac{\lambda^i_{s^1}}{\lambda^i_{s^0}}, \frac{\frac{dV^i_{s^1}(s^1)}{d\theta}}{\lambda^i_{s^1}} \right].$$

Restricting to forward-looking numeraires, which means that $\lambda^i_{s^t}$ must be a function of $\{c^i_k(s^k)\}_{k \geq t, s^k \geq s^t}$. Suppose that the forward-looking numeraire at date 0 includes date 0 consumption. See footnote 4.2.3 for a discussion of why this assumption is without loss since the Theorem is "for all $d\theta$". Under incomplete markets, we will therefore be able to find a pair of histories at which individual MRS are not equalized, which then also implies that

$$\frac{\lambda^i_{s^t}}{\lambda^i_{s^0}}$$

are not equalized in the cross section. And since different perturbations are associated with different values of $\frac{1}{\lambda^i_{s^t}} \frac{dV^i_{s^t}}{d\theta}$, we can always find a perturbations such that the second term in equation (13) is sufficiently large as to change the sign of $\Xi^E_{s^t}$ relative to $\Xi^E_{s^0}$. Hence, it is possible to find a perturbation for which the efficiency assessment based on a forward-looking numeraire is time inconsistent. □

Properties of Welfare Numeraires. Here we formally establish the following two properties of welfare numeraires.

(i) The perturbation $d\theta$ is Paretian with transfers in units of a forward-looking numeraire if and only if $\Xi^E_{s^t} > 0$ in units of that numeraire.

(ii) It is impossible to find feasible perturbations from Pareto efficient allocations that satisfy $\Xi^E_{s^t} > 0$ for any forward-looking numeraire.

To prove (i), note that if a perturbation $d\theta$ is a Pareto improvement, then it must be that $\frac{dV^i_{s^t}}{d\theta} \geq 0$ for all i, with at least one inequality. Therefore, the sum of individual welfare gains expressed in any valid (with $\lambda^i_{s^t} > 0$) welfare numeraire (forward or backward-looking) must be strictly positive. So we must also have $\Xi^E_{s^t} = \sum_i \frac{1}{\lambda^i_{s^t}} \frac{dV^i_{s^t}}{d\theta} > 0$. If instead there is a feasible perturbation for which $\Xi^E_{s^t} = \sum_i \frac{1}{\lambda^i_{s^t}} \frac{dV^i_{s^t}}{d\theta} > 0$ where $\lambda^i_{s^t}$ corresponds to a forward-looking numeraire, then it must be possible to transfer resources from winners to losers in units of that welfare numeraire so that $\frac{1}{\lambda^i_{s^t}} \frac{dV^i_{s^t}}{d\theta} \geq 0$ for all i, with at least one strict inequality, while ensuring that $\Xi^E_{s^t}$ does not change. This direction is only valid for forward-looking numeraires: If we chose a backward-looking numeraire, the transfers needed to construct the Pareto-improving perturbations would not be feasible, as it is impossible to transfer resources in the past.

To prove (ii), we proceed by contradiction. Suppose that $\Xi^E_{s^t} > 0$ for some feasible perturbation of a given allocation. Then by virtue of (i), we can reallocate resources among individuals to find a Pareto improvement with transfers. But this means that the original allocation was not Pareto efficient, leading to a contradiction. It is important to note that this result only applies to the set of Pareto efficient allocations that solve the Pareto Problem (first-best allocations). It does not apply to constrained efficient allocations. If we chose a backward-looking numeraire, the transfers needed to construct the Pareto-improving perturbations would not be feasible, as it is impossible to transfer resources in the past.

Proof of Proposition 4. (Inconsistency of Sources of Efficiency)

Proof. a) If markets are complete, risk-sharing and intertemporal-sharing are zero by Proposition 4 in Dávila and Schaab (2024). As valuations are proportional, $\Xi^{AE}_{s^k}$ is necessarily time-consistent.

b) Recall that the history s^k aggregate efficiency assessment is defined as

$$\Xi^{AE}_{s^k} = \sum_{t \geq k} \omega_{t|s^k} \sum_{s^t \geq s^k} \omega_{t|s^k}(s^t) \sum_i \frac{\partial \mathcal{C}^i_t(s^t \mid s^k)}{\partial \theta}$$

When the perturbation $d\theta$ generates a *static* aggregate efficiency gain or loss, this means that

$$\frac{\partial \mathcal{C}^i_t(s^t \mid s^k)}{\partial \theta} = 0$$

for all $t \neq T$ and $s^t \neq s^T$. Therefore, we can write the static aggregate efficiency assessment for such a perturbation from the perspective of history s^k as

$$\Xi_{s^k}^{AE} = \omega_{T|s^k}\,\omega_{T|s^k}(s^T)\sum_i \frac{\partial \mathcal{C}_T^i(s^T \mid s^k)}{\partial \theta}.$$

When Assumption 4.1.2 is satisfied, then

$$\frac{\partial \mathcal{C}_T^i(s^T \mid s^k)}{\partial \theta} = \frac{\partial \mathcal{C}_T^i(s^T \mid s^t)}{\partial \theta} = \frac{dc_T^i(s^T)}{d\theta},$$

for any s^t and s^k. Therefore, under 4.1.2, the perceived effect of perturbation $d\theta$ on aggregate consumption at history s^T does not change with the perspective of the assessment. The only sources of time inconsistency can therefore be the aggregate valuations $\omega_{T|s^k}$ and $\omega_{T|s^k}(s^T)$. But since $\lambda_{s^k}^i > 0$ and $u'(c_t^i(s^t)) > 0$ for all t, s^t and i, these two weights must also be strictly positive. Therefore, we have

$$\Xi_{s^0}^{AE} = \underbrace{\frac{\bar{\omega}_{T|s^0}\,\bar{\omega}_{T|s^0}(s^T)}{\bar{\omega}_{T|s^k}\,\bar{\omega}_{T|s^k}(s^T)}}_{>0}\,\Xi_{s^k}^{AE}.$$

This tells us that "static" aggregate efficiency assessments are always time-consistent under Assumption 4.1.2. □

C Microfoundation of the Consumption Function

There are I individuals indexed by i. Time is discrete and we capture uncertainty using the usual history notation, allowing for both aggregate and idiosyncratic uncertainty. We assume that each individual i makes a single decision $c_t^i(s^t) \in \mathbb{R}$ at date t and conditional on the realization of history s^t.

Behavior. We focus on environments in which behavior is characterized by a *policy function*

$$C\left(\phi, x, X, \vec{p}, \vec{\theta}\right).$$

Our notation is as follows:

- x denotes idiosyncratic time-varying state variables (ex-post heterogeneity)
- ϕ denotes idiosyncratic permanent state variables or types (ex-ante heterogeneity)
- X denotes aggregate state variables
- $\vec{p}$ denotes the "continuation price process", i.e., a stochastic process with a particular initialization (see below)
- $\vec{\theta}$ denotes the "continuation policy process" (see below)

Notice that we refer as "prices" to every (aggregate) macroeconomic variable that affects (the decision problems of) agents directly. We also introduce the aggregate state X at this point; while it is not a direct argument of the policy function $C(\cdot)$, it will affect the determination of prices below. Formally, we have

$$C : \mathbb{R}^{N_\phi} \times \mathbb{R}^{N_x} \times \mathcal{L}^{N_p} \times \mathcal{L} \to \mathbb{R}.$$

We denote the space of stochastic processes under consideration (initialized in a certain way) by $\mathcal{L}$. There is a single policy process θ.

Using this consumption function, we can now express individual i's consumption decision at date t in history s^t as

$$c_t^i(s^t) = C\left(\phi^i, x_t^i(s^t), \left\{p_\ell(s^\ell), \theta_\ell(s^\ell)\right\}_{\ell \geq t, s^\ell \geq s^t}\right).$$

In other words, given agent i's type and state, and given the pair of stochastic processes $\vec{p}$ and $\vec{\theta}$ initialized at $p_t(s^t)$ and $\theta_t(s^t)$, this policy function determines behavior $c_t^i(s^t)$.

Environment. The environment is defined as a description of how state variables and prices are determined. Notice that the permanent types ϕ^i never change. First, we have idiosyncratic state variables, which evolve according to the law of motion

$$x^i_{t+1}(s^{t+1}) = h\Big(\phi^i, x^i_t(s^t), c^i_t(s^t), p_t(s^t), \theta_t(s^t), s_{t+1}\Big)$$

Second, we have the aggregate state variables. We assume they evolve according to the law of motion

$$X_{t+1}(s^{t+1}) = k\Big(X_t(s^t), p_t(s^t), \theta_t(s^t), \Big\{c^i_t(s^t)\Big\}_i, s_{t+1}\Big).$$

Third and finally, we have macroeconomic prices. We assume that they solve "market clearing conditions" given by

$$p_t(s^t) = m\Big(X_t(s^t), \theta_t(s^t), \Big\{c^i_t(s^t)\Big\}_i\Big).$$

We do not let $m(\cdot)$ explicitly depend on s_t. This is without loss because, as explained above, we already allow for the aggregate state $X_t(s^t)$ to depend on s_t directly, and so aggregate "shocks" can always be written as one of the aggregate states.

Dual representation. We now derive what we will refer to as a "dual representation" of behavior. Notice first that we can iterate on the law of motion for $x^i_t(s^t)$ and arrive at

$$x^i_t(s^t) = H_t\Big(x^i_0(s^0), \phi^i, \Big\{c^i_k(s^k), p_k(s^k), \theta_k(s^k)\Big\}_{0\leq k<t, s^0\leq s^k<s^t}, s^t\Big)$$

So what matters for the determination of individual i's state in history s^t is the initial condition $x^i_0(s^0)$, her type, her behavior since the world started, all the shocks she drew, and then the prices and policies she faced along the way. Crucially, the individual state is *backward-looking.* Finally, we index $H_t(\cdot)$ by t because the size of the input arguments changes with time. We have

$$H_t : \mathbb{R}^{N_x} \times \mathbb{R}^{N_\phi} \times \mathbb{R}^t \times \mathbb{R}^{t\times N_p} \times \mathbb{R}^t \times \mathcal{S}^t \to \mathbb{R}^{N_x}$$

Second, for the aggregate state variables, we can also iterate backwards and arrive at

$$X_t(s^t) = K_t\Big(X_0(s^0), \Big\{p_k(s^k), \theta_k(s^k)\Big\}_{0\leq k<t, s^0\leq s^k<s^t}, \Big\{c^i_k(s^k)\Big\}_{i,0\leq k<t, s^0\leq s^k<s^t}, s^t\Big)$$

Notice that in both cases the equalities are strict, i.e., $k < t$ and $s^k < s^t$. Notice also that the aggregate state depends on all agents' behavior.

We can now put together the previous equation with the market clearing condition, yielding

$$p_t(s^t) = M_t\Big(X_0(s^0), \Big\{\theta_k(s^k)\Big\}_{0\le k\le t,\ s^0\le s^k\le s^t}, \Big\{c_k^i(s^k)\Big\}_{i,\ 0\le k\le t,\ s^0\le s^k\le s^t}\Big).$$

Several observations are in order: First, the inequalities are now weak. That's because $p_t(s^t)$ depends both on contemporaneous policy $\theta_t(s^t)$ directly and on prior policy indirectly through the aggregate state. The same is the case for consumption. Second, one could in principle subsume the initial condition $X_0(s^0)$ directly into the function $M(\cdot)$ since it is policy-invariant. But since initial conditions will be critical when assessing policy implications from different perspectives, we leave the dependence explicit here. Finally, we must have the function $M_t(\cdot)$ depend on calendar time because the size of the input arguments changes over time. Notice, for example, that the second argument of $M_t(\cdot)$ is a $t \times 1$ vector. In other words,

$$M_t : \mathbb{R}^{N_X} \times \mathbb{R}^t \times \mathbb{R}^{t\times I} \to \mathbb{R}^{N_p}.$$

It is a function of vectors (rather than matrices to account for histories) because the only objects that matter are the realizations of policy and consumption *along the realized history.*

We can now use the previous expressions and rewrite the law of motion for individual state variables via $H_t(\cdot)$ as

$$\begin{aligned} x_t^i(s^t) &= H_t\Big(x_0^i(s^0), \phi^i, \Big\{c_k^i(s^k), p_k(s^k), \theta_k(s^k)\Big\}_{0\le k<t, s^0\le s^k<s^t}, s^t\Big) \\ &= \mathcal{H}_t\Big(\phi^i, x_0^i(s^0), X_0(s^0), \Big\{c_k^i(s^k)\Big\}_{0\le k<t}^{s^0\le s^k<s^t}, \Big\{\theta_k(s^k)\Big\}_{0\le k\le t}^{s^0\le s^k\le s^t}, \Big\{c_k^j(s^k)\Big\}_{j,\ 0\le k\le t}^{s^0\le s^k\le s^t}, s^t\Big), \end{aligned}$$

where we index brackets $\{\cdot\}$ by sub- and superscripts only to make the notation more compact. This notation makes explicit the dependence of $x_t^i(s^t)$ on consumption decisions through different channels. At the cost of conflating these channels, we can simplify and arrive at

$$x_t^i(s^t) = \mathcal{H}_t\Big(\phi^i, x_0^i(s^0), X_0(s^0), \Big\{\theta_k(s^k)\Big\}_{0\le k\le t}^{s^0\le s^k\le s^t}, \Big\{c_k^j(s^k)\Big\}_{j,\ 0\le k\le t}^{s^0\le s^k\le s^t}, s^t\Big).$$

Finally, we can go back to the policy function that determines behavior. Reproducing from above for convenience,

$$c_t^i(s^t) = C\Big(\phi^i, x_t^i(s^t), X_t(s^t), \Big\{p_\ell(s^\ell), \theta_\ell(s^\ell)\Big\}_{\ell\ge t, s^\ell\ge s^t}\Big).$$

Notice that using the price equation, we can rewrite the aggregate state as

$$X_t(s^t) = \mathcal{K}_t\Big(X_0(s^0), \Big\{\theta_k(s^k)\Big\}_{0\le k<t}^{s^0\le s^k<s^t}, \Big\{c_k^i(s^k)\Big\}_{i,\,0\le k<t}^{s^0\le s^k<s^t}, s^t\Big)$$

Finally, we can plug into the consumption function, which gives us

$$c_t^i(s^t) = \tilde{\mathcal{C}}_t\Bigg(\phi^i, x_0^i(s^0), X_0(s^0), \Big\{\theta_k(s^k)\Big\}_{0\le k\le t}^{s^0\le s^k\le s^t}, \Big\{c_k^j(s^k)\Big\}_{j,\, 0\le k\le t}^{s^0\le s^k\le s^t}, \Big\{p_\ell(s^\ell), \theta_\ell(s^\ell)\Big\}_{\ell\ge t}^{s^\ell\ge s^t}, s^t\Bigg),$$

or simply

$$c_t^i(s^t) = \mathcal{C}_t\Bigg(s^0, \Big\{\theta_\ell(s^\ell)\Big\}_{\ell\ge 0}^{s^\ell\ge s^0}\Bigg),$$

after expressing the forward-looking price process as a function of the policy process. This derivation illustrates constructively the various channels through which a perturbation in a policy parameter $\theta_k(s^k)$ affects consumption of different individuals at different horizons.