Time Inconsistency with Heterogeneous Agents*

Eduardo Dávila† Andreas Schaab‡

August 2025

Abstract

This paper studies the time inconsistency of policies in heterogeneous agent (HA) environments. HA models feature three sources of time inconsistency: inconsistent individual preferences (Strotz, 1956), forward-looking behavior (Kydland and Prescott, 1977), and interpersonal welfare comparisons. When individual preferences are symmetric, only the utilitarian social welfare function implies time consistent aggregate welfare assessments. The attribution of welfare gains to efficiency and redistribution, on the other hand, invariably changes over time when markets are incomplete and is therefore time inconsistent, irrespective of the social welfare function and the extent of heterogeneity in individual preferences. Therefore, even when optimal policy is time consistent, its justification on grounds of efficiency and redistribution changes as time passes and uncertainty is realized. These results offer a useful new perspective on time consistency problems in the design of anticipated relief policies, social insurance, labor taxation, investment policies, and social contracts.

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

Time consistency is a central concern in 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 the sources of time inconsistency is indispensable for normative economics. This paper studies the time consistency of policies in environments with heterogeneous agents and incomplete markets.

In representative agent (RA) environments, two well-understood forces lead to the time inconsistency of policies: inconsistent individual preferences (Strotz, 1956) and forward-looking behavior (Kydland and Prescott, 1977).[1] In heterogeneous agent (HA) environments, there is a third source of time inconsistency: aggregate welfare assessments. When using a social welfare function (SWF) to judge the desirability of a policy, the planner aggregates welfare gains and losses across heterogeneous individuals. These interpersonal welfare comparisons can emerge as a new source of time inconsistency, in addition to those identified by Strotz (1956) and Kydland and Prescott (1977). Intuitively, when the planner's marginal rate of substitution between two individuals in utils changes over time, so too can the perceived desirability of the policy. Only the utilitarian SWF is immune to this form of time inconsistency because individuals' Pareto weights are constant across dates and histories.[2] This observation was first formalized by Zuber (2011) but its intuitive importance had already been acknowledged — albeit only informally — by earlier work dating to at least Harsanyi (1955).

In this paper, we revisit the question of time consistency in heterogeneous agent environments and study the role of incomplete markets. Our main contribution is to unpack the aggregate welfare assessment and characterize the time consistency of the distinct normative considerations that make up the overall welfare assessment. Our main result shows that, in HA incomplete markets economies, assessments of efficiency and redistribution gains and losses invariably change over time and are therefore time inconsistent. Even when the overall welfare assessment is consistent in the special case of a linear SWF and with homogeneous discount factors and beliefs, the attribution of welfare gains and losses to efficiency and redistribution invariably

[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) 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.

[2] Heterogeneity in discount factors and beliefs can lead to time inconsistent aggregate welfare assessments even when the SWF is utilitarian, as shown by Jackson and Yariv (2014, 2015) and Mongin (1995, 1998), respectively.

changes over time.

We begin in Section 2 with an illustrative example. The world is inhabited by a continuum of households and there are two periods, today and tomorrow. Households work and consume in both periods, but there are no financial markets so they cannot save. Today, households are identical. Tomorrow, they face symmetric labor productivity risk. We consider a tax-and-transfer policy that raises revenue from a distortionary labor income tax to provide a uniform, lump-sum transfer to households. And we compare the assessments of this policy from the perspectives of today and tomorrow under different SWFs. This example allows us to highlight four simple observations.

First, the attribution of gains and losses from the optimal policy to efficiency and redistribution changes with the perspective of the assessment and is therefore time inconsistent. Ex ante, households are identical. They would prefer smoothing their consumption across states of the world tomorrow but are unable to do so. The social transfer helps insure households against idiosyncratic risk because the burden of taxation falls disproportionately on the more productive while the transfer is uniform. From the perspective of today, the promise of labor income taxation tomorrow achieves welfare gains due to risk-sharing. On the other hand, the tax is distortionary because it discourages households from supplying labor, generating a welfare loss from production efficiency. Suppose time passes and tomorrow's state of the world is realized. From this ex post perspective, the planner no longer perceives any scope to insure households against idiosyncratic risk. What the planner perceived as risk-sharing gains ex ante now appears as pure redistribution. Therefore, the planner initially attributes both gains and losses from labor taxation to efficiency; whereas ex post, the planner trades off redistribution gains against production efficiency losses. This is a concrete example of the time inconsistency of efficiency and redistribution assessments.

Second, the perceived production efficiency loss due to distortionary taxation is the same ex ante and ex post. While risk-sharing and redistribution assessments change over time, assessments of production efficiency gains and losses are time consistent. This highlights that some components of the overall efficiency and welfare assessment are time consistent, while others are not.

Third, the overall welfare assessment of tax policy is time consistent under the utilitarian SWF, highlighting the result of Zuber (2011). Even so, the underlying assessments of efficiency and redistribution gains and losses are not. Even though the optimal policy is time consistent, therefore, its justification varies with the perspective of the assessment.

Fourth, optimal policy is time inconsistent for all other SWFs. Only the utilitarian SWF values ex post redistribution and ex ante risk-sharing equivalently. Non-utilitarian SWFs over- or under-weight redistribution relative to the utilitarian benchmark. Non-utilitarian planners therefore perceive the redistribution gains from the social transfer ex post as more or less valuable than the risk-sharing gains ex ante. This leads to a time consistency problem.

We generalize these four observations to a large class of dynamic stochastic environments with heterogeneous agents and incomplete markets. The framework we introduce in Section 3 nests most HA models currently found in the applied literature, leveraging a general sequence-space representation in the spirit of Auclert, Bardóczy, Rognlie, and Straub (2021). Our only restrictive assumption is that preferences take the separable expected utility form with homogeneous discount factors and beliefs. This allows us to clarify our contribution relative to the literature that studies time inconsistency resulting from heterogeneity in individual preferences and beliefs.

We set the stage in Section 4 by presenting a systematic characterization of the sources of time inconsistency for RA and HA welfare assessments, revisiting the results of Kydland and Prescott (1977) and Zuber (2011). Relative to Zuber (2011) and the subsequent social choice literature, we analytically characterize the relationship between ex ante and ex post assessments. We show that the time consistency problem from interpersonal welfare comparisons is governed by a particular cross-sectional covariance. when this covariance takes the opposite sign from the welfare assessment and is sufficiently large can time inconsistency arise. Our result also clarifies the role of incomplete markets. Indeed, when markets are complete, even non-utilitarian social welfare functions lead to time consistent assessments as long as individual preferences are symmetric.

We then present our main result in Section 4.2: The attribution of welfare gains and losses to efficiency and redistribution invariably changes over time and is therefore time inconsistent when markets are incomplete. Even in the special case of a linear SWF and homogeneous discount factors and beliefs, the justification of optimal policy on grounds of efficiency or redistribution varies with the assessment perspective. The source of time inconsistency in efficiency assessments is a change in the reference unit, or welfare numeraire, used to express and compare individual welfare gains and losses. We refer to this as *numeraire inconsistency* and again associate the resulting time consistency problem with a particular cross-sectional covariance term. When markets are complete, numeraire inconsistency can never arise because individuals' relative valuations of bundles of goods are equalized. The time inconsistency in

efficiency and redistribution assessments we identify is therefore a direct result of incomplete markets.

Applications. Our goal in this paper is to show that the time inconsistency of efficiency and redistribution assessments provides a useful new perspective on applied policy questions. To do so, we present three applications, each highlighting a different source of efficiency gains — intertemporal-sharing, risk-sharing, or aggregate-efficiency — that can lead to the time inconsistency of policies in realistic scenarios.

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 provides transfers to low-consumption individuals early in their recovery, phasing out as consumption approaches long-run levels. 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.

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 optimal 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 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. those on time inconsistency and on heterogeneous agent incomplete markets economies. Following the foundational papers of Strotz (1956) and Kydland and Prescott (1977), a vast body of work has analyzed the time consistency of government policies, especially fiscal and monetary policy.[3] Much of this tradition has worked with RA models. Notable exceptions include Werning (2007), Chang (2022), Bilbiie (2024) and Dávila and Schaab (2023) who study time consistency in the presence of heterogeneity.[4] All of these papers use linear SWFs to make interpersonal welfare comparisons. Our contribution to this body of work is three-fold. First, we show that interpersonal welfare comparisons emerge as a new source of time inconsistency when the SWF is non-linear and we characterize a specific cross-sectional covariance that governs the resulting time consistency problem. Second, we establish that efficiency and redistribution assessments are invariably time inconsistent when markets are incomplete, even for the special case of a linear SWF. Third, we present a unified characterization of the sources of time inconsistency in heterogeneous agent environments — distinguishing between time consistency problems á la Strotz (1956), Kydland and Prescott (1977) and due to interpersonal welfare comparisons — using a general sequence-space representation that nests most models that are currently employed in applied work. This allows us in particular to clarify the role of incomplete markets for time consistency as distinct from heterogeneity in individual preferences.

[3] See Kydland and Prescott (1980), 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), 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.

[4] Werning (2007) discusses the time consistency of capital taxation. The usual time consistency problem in the RA environment arises because the planner lacks lump-sum taxation and expropriating capital ex post mimics a non-distortionary tax since accumulation decisions are already sunk. Werning (2007) shows that in the presence of heterogeneity, optimal capital taxation may be 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 time consistency problem is therefore in the spirit of Kydland and Prescott (1977). Indeed, Werning (2007) exclusively focuses on welfare assessments under a linear SWF with common discount factor and beliefs; aggregate welfare assessments consequently do not emerge as a distinct source of time inconsistency in his environment.

Our paper is therefore also part of the rapidly growing literature in macroeconomics that uses sequence-space methods to study heterogeneous agent models with incomplete markets. See Auclert, Rognlie, and Straub (2025) for a recent survey. We derive our results for a general class of HA models leveraging a representation of individual consumption functions that builds on the approach developed in Auclert, Bardóczy, Rognlie, and Straub (2021). While most recent work in this literature has been positive, our focus is instead on the normative implications of heterogeneity.

The main result of this paper is a characterization of the time consistency of efficiency and redistribution assessments. Distinguishing notions like efficiency, risk-sharing and redistribution in simple environments like our example of Section 2 is straightforward. But this becomes non-trivial in general incomplete markets economies. Our results therefore build on prior work that has developed formal characterizations of efficiency, risk-sharing and redistribution in heterogeneous agent environments: Kaldor (1939), Hicks (1939), Benabou (2002), Bhandari, Evans, Golosov, and Sargent (2024), and in particular Dávila and Schaab (2024).

Outside of macroeconomics, there is a long tradition in social choice theory of studying the possibility of aggregating individual preferences. Arrow (1950)'s impossibility result presents an early benchmark, and subsequent work has explored alternative formulations and relaxations of Arrow's criteria. Zuber (2011) shows that time consistency generically requires a linear social welfare function. 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. 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 focus on time invariant SWFs but allow for non-stationarity. We assume homogeneous discount factors and beliefs from the outset to clearly distinguish our results from this body of work. Our contribution to this literature is three-fold. First, our marginal approach based on welfare assessments is complementary to Zuber (2011)'s result as it allows us to associate time inconsistency in interpersonal welfare comparisons with a specific cross-sectional covariance term. Second, we clarify the role incomplete markets play for the time inconsistency of welfare assessments,

as distinct from heterogeneity in individual preferences. Third, we show that efficiency and redistribution assessments are invariably inconsistent even when the SWF is stationary and discount factors and beliefs are homogeneous.

The tension between efficiency or risk-sharing and redistribution plays an important role in our paper and has been an enduring theme of economic thought. The early work of Marshall (1920), Pigou (1920) and Knight (1921) already recognizes that schemes designed to provide risk-sharing ex ante necessarily imply redistributive transfers ex post once uncertainty has realized. Arrow (1963) is often credited as providing an early formal statement that the pooling of risk can raise all individuals' expected utility ex ante but requires redistribution ex post:

> *[T]he preference for redistribution expressed in government taxation... can be reinterpreted as desire for insurance... Thus, optimality, in a context which includes risk-bearing, includes much that appears to be motivated by distributional value judgments.*

In a sense, our results provide a novel and practical formalization of these ideas.

Our paper also shares this theme with subsequent work in the public finance literature. The simple example we develop in Section 2 builds on Varian (1980) and Piketty and Saez (2013b). Much of the early work on optimal taxation focuses on static environments and appeals to an equivalence between ex post redistribution and ex ante risk-sharing (behind the veil of ignorance).[5] Our results emphasize an important qualification to this reasoning: Only the utilitarian SWF values risk-sharing and redistribution equivalently, whereas non-utilitarian SWFs lead to a time consistency problem. Much applied work in public finance has employed non-utilitarian SWFs (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). The implications of dynamics and uncertainty have been the focus of a new dynamic public finance literature. And the interplay between ex ante risk-sharing and ex post redistribution has been explored, for instance, by Farhi and Werning (2013) and Golosov, Troshkin, and Tsyvinski (2016). Farhi and Werning (2013) observe that a utilitarian SWF implies an equivalence between ex ante risk-sharing and ex post redistribution, which breaks down for non-utilitarian SWFs. In contrast to this literature, our focus is on characterizing time inconsistency in not only welfare but also efficiency and redistribution assessments.

Lastly, there is long tradition in political economy that studies time consistency and shares our focus on the tension between efficiency and redistribution. Important examples include

[5] See for example Mirrlees (1971), Diamond (1998) and Saez (2001).

Alesina and Rodrik (1994), Persson and Tabellini (1994), Krusell, Quadrini, and R\'\ios-Rull (1997), and Bisin, Lizzeri, and Yariv (2015). Farhi, Sleet, Werning, and Yeltekin (2012) study a dynamic Mirrleesian model where policy is chosen sequentially without commitment. Time inconsistency arises in their setting due to the interplay between dynamic incentive provision and redistribution; but this time consistency problem is due to forward-looking behavior á la Kydland and Prescott (1977) rather than interpersonal welfare comparisons: Capital accumulation decisions are sunk ex post, and so a utilitarian planner with redistribution concern but without commitment is tempted to expropriate capital ex post. Papers in this literature often study 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 transfer policy.

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

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

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 $V_1^i = u(c_1^i, \ell_1^i)$. 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

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

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:

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

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:

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

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.

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

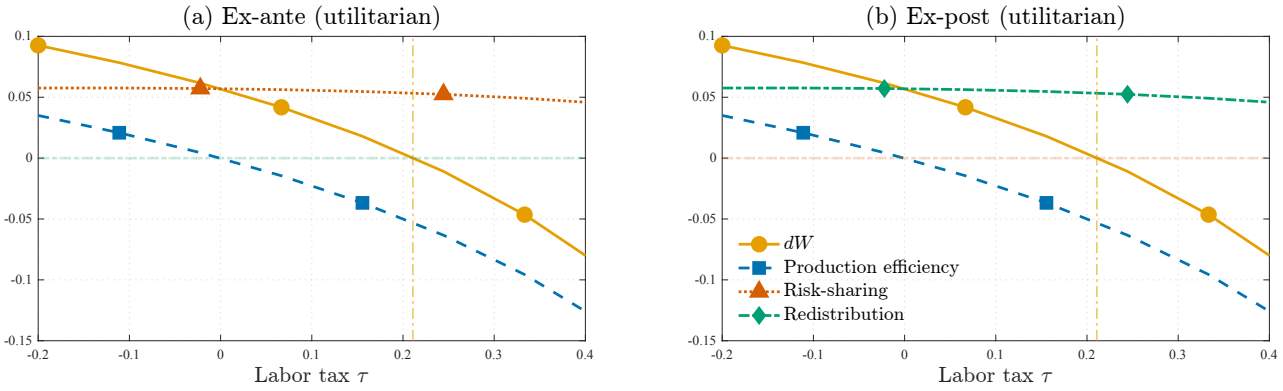


Figure 1: Utilitarian Welfare Assessments

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.

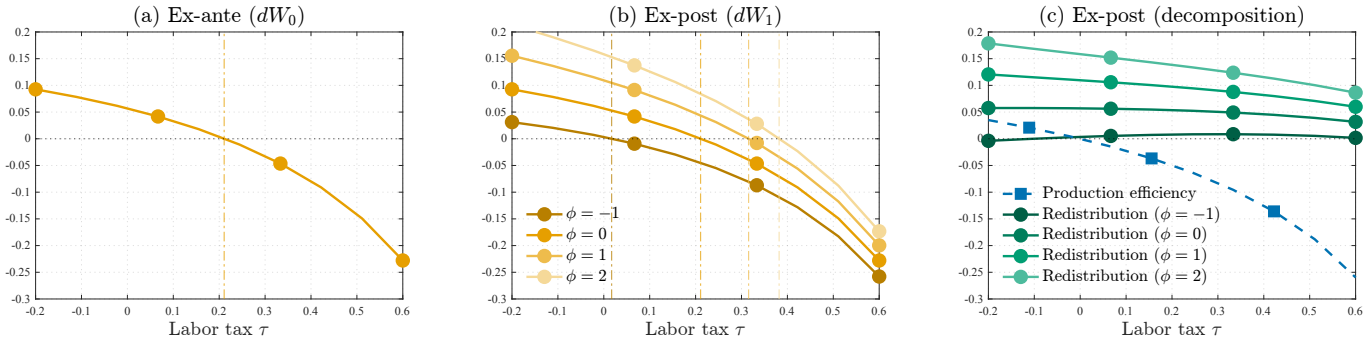


Figure 2: Time Inconsistency of Welfare Assessments

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.

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 remains unchanged)* The planner perceives the same 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 a time consistency problem.

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 evaluates 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

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

where $c^i_t\left(s^t\right)$ 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 $u\left(\cdot\right)$ is well-behaved and Inada conditions apply. To simplify the exposition, we assume in the main text that all individuals have identical preferences, while allowing 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.[6] Formally, social 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 $\mathcal{W}\left(\cdot\right)$ 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

[6] The welfarist approach is widely used because it is Paretian, that is, it concludes that every Pareto-improving perturbation is desirable, and because nonwelfarist approaches violate the Pareto principle (XX).

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 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.[7]

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 nominal interest rates (monetary policy) or taxes/transfers (fiscal policy), but $\boldsymbol{\theta}$ may also represent exogenous primitives such as productivity or endowments, or simply index allocations directly chosen by a planner.[8]

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,

$$\partial \mathcal{C}^i_t(s^t, \boldsymbol{\theta} \mid s^k)/\partial\theta \neq \partial \mathcal{C}^i_t(s^t, \boldsymbol{\theta} \mid s^\ell)/\partial\theta.$$

For instance, the response of consumption at history s^k to a tax change at that history from the perspective of that history s^k is different from the response from an earlier s^0 perspective, as individuals have the opportunity in the latter case to react to the tax in anticipation by

[7] 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).

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

changing their behavior, and thus equilibrium allocations, between date 0 and history 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 study the time inconsistency of overall welfare assessments. Next, we do the same with efficiency and redistribution assessments. Finally, we analyze the time inconsistency of the sources of efficiency gains, in particular, aggregate-efficiency, risk-sharing, and intertemporal-sharing.

4.1 Inconsistency of Welfare Assessments

We being by asking under what conditions an overall welfare assessment is time consistent. We define time consistency as follows.

Definition 1. *(Time Consistency of Welfare Assessments) The welfare assessment of a perturbation $d\theta_t(s^t)$, 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. In that case, all later assessments agree on whether the perturbation is desirable or not. If $\boldsymbol{\theta}$ is a policy set by a planner, then consistency in the sign of the welfare assessment ensures the time consistency of the optimal policy.[9]

For ease of exposition, in the main text we focus on the specific case of a perturbation $d\theta_1(\bar{s}^1)$ at a particular history $\bar{s}^1$ at date 1, which we refer to simply as "$d\theta$". We refer to $\frac{dW_{s^0}}{d\theta}$ as the *ex ante* assessment and to $\frac{dW_{\bar{s}^1}}{d\theta}$ as 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 from the perspective of which the perturbation $d\theta_1(\bar{s}^1)$ might be reassessed. Our proofs in Appendix B are for arbitrary perturbations.

[9]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 explain why we focus on sign consistency in this paper, and why Definition 1 is the appropriate definition of consistency in marginal form.

Representative agent benchmark. To set the stage, we review the representative agent (RA) benchmark with $I = 1$. When individual preferences are consistent (Strotz, 1956), there are three distinct time consistency problems than can emerge in environments with a representative agent.

Proposition 1 (Kydland and Prescott, 1977)**.** *Suppose $I = 1$. For the date 0 welfare assessment of perturbation $d\theta$ to be time inconsistent, one of the following must be true:*

(i) (Forward-looking behavior) *θ affects the consumption allocation at date 0*

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

(iii) (State variables) $\frac{\partial \mathcal{C}_1^i(\bar{s}^1 \mid s^0)}{\partial \theta} \neq \frac{\partial \mathcal{C}_1^i(\bar{s}^1 \mid \bar{s}^1)}{\partial \theta}$

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

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

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

$$\frac{dW_{\bar{s}^1}}{d\theta} = \alpha^i_{\bar{s}^1} \sum_{t \geq 1} \beta^{t-1} \sum_{s^t \geq \bar{s}^1} \pi(s^t \mid \bar{s}^1) \, (u^i)'(c^i_t(s^t)) \frac{\partial \mathcal{C}^i_t(s^t \mid \bar{s}^1)}{\partial \theta}.$$

Comparing these two expressions highlights three distinct source of time inconsistency. First, the ex ante assessment features terms corresponding to date 0, whereas the ex post assessment sums over periods $\sum_{t \geq 1}$ from date 1 onwards. If the perturbation $d\theta$ affects consumption at date 0 due to *forward-looking behavior*, then this will affect the ex ante but not the ex post assessment. Second, the ex ante assessment features terms corresponding to the perturbation's impact on the allocation at histories $s^1 \neq \bar{s}^1$. Ex post, on the other hand, these alternative histories did not realize and so the planner no longer considers these terms. We refer to this source of time inconsistency as *the road not taken*. Third, the perturbation's effect on the consumption allocation at history $\bar{s}^1$ may differ depending on the perspective or time of the announcement. In that case, $\frac{\partial \mathcal{C}_1^i(\bar{s}^1 \mid s^0)}{\partial \theta} \neq \frac{\partial \mathcal{C}_1^i(\bar{s}^1 \mid \bar{s}^1)}{\partial \theta}$, which corresponds to condition (iii) of Proposition 1.[10] Finally, notice that individual weights $\alpha^i_{s^0}$ and $\alpha^i_{\bar{s}^1}$ must be strictly positive

[10]Notice that (i) is a prerequisite of (iii). Condition (iii) requires that the initial conditions (state variables) we take as given from the perspective of history $\bar{s}^1$ are affected by an announcement of the perturbation at date 0. This can only happen when the allocation at date 0 is affected, which in turn changes the state variables taken as given from the perspective of $\bar{s}^1$. Nonetheless, we find it useful to explicitly distinguish between conditions (i) and (iii) because they describe different economic channels through which time inconsistency can emerge.

and can therefore not be a source of inconsistency.[11]

Welfare assessments with heterogeneous agents. Having set the stage by reviewing the RA benchmark, we now study the time consistency of welfare assessments in heterogeneous agent (HA) environments with $I > 1$. Our next result shows that, in the presence of heterogeneity, interpersonal welfare comparisons emerge as a new source of time inconsistency.

Proposition 2 (Inconsistency of Welfare Assessments)**.** *Suppose $I > 1$. For the date 0 welfare assessment of perturbation $d\theta$ to be time inconsistent, either one of (i) – (iii) in Proposition 1 must be true or:*

(iv) (Interpersonal welfare comparisons) $\alpha^i_{\bar{s}^1}/\alpha^i_{s^0} \neq \alpha^j_{\bar{s}^1}/\alpha^j_{s^0}$ *for at least two individuals i and j .*

Proposition 2 demonstrates that a new source of time inconsistency emerges in heterogeneous agent environments, in addition to those identified by Strotz (1956) and Kydland and Prescott (1977). With heterogeneous agents, interpersonal welfare comparisons can lead to time consistency problems.

To explain the intuition behind Proposition 2, it is useful to consider the case where conditions (i) – (iii) do not apply, allowing us to focus on the new condition (iv). We make the following "no-KP" assumption, ruling out time consistency problems in the spirit of Kydland and Prescott (1977).

Assumption 1. *ass:A1 (No-Kydland-Prescott) We assume that (i) $\frac{\partial \mathcal{C}^i_0(s^0 \mid s^0)}{\partial \theta} = 0$; (ii) $\frac{\partial \mathcal{C}^i_1(s^1 \mid s^0)}{\partial \theta} = 0$ for $s^1 \neq \bar{s}^1$; and (iii) $\frac{\partial \mathcal{C}^i_1(\bar{s}^1 \mid s^0)}{\partial \theta} = \frac{\partial \mathcal{C}^i_1(\bar{s}^1 \mid \bar{s}^1)}{\partial \theta}$ for all i.*

In environments that satisfy Assumption 1, we can write the ex ante welfare assessment of the perturbation $d\theta$ as

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

[11] Kydland and Prescott (1977) focus on environments without uncertainty. Their main analysis only features conditions (i) and (iii) of Proposition 1. Specifically, the main condition in their text can be written using our notation as $\partial \mathcal{C}^i_0(s^0 \mid s^0)/\partial\theta \neq 0$. That is, the allocation at date 0 is affected by the date 0 announcement of a perturbation $d\theta$. This can lead to (i) or (iii) because, in order for state variables at the beginning of date 1 to be affected, necessarily behavior at date 0 must be affected. They note that the case of uncertainty can be treated by relabeling periods as histories in the spirit of Arrow-Debreu. While we find it useful to treat (ii) as a distinct source of time inconsistency, their paper is to be credited with identifying time consistency problems (i) – (iii). Proposition 1 should therefore be interpreted as a restatement of Kydland and Prescott (1977).

which follows directly from our proof in Appendix B.2. Equation (10) decomposes the ex ante welfare assessment $\frac{dW_{s^0}}{d\theta}$ into a term that is proportional to the ex post assessment $\frac{dW_{\bar{s}^1}}{d\theta}$ and a cross-sectional covariance term. When the covariance term is 0, the ex ante and ex post assessments necessarily share the same sign because the coefficient on the later is strictly positive.[12] Time inconsistency therefore requires that the covariance term is non-zero — in fact, it must have the opposite sign from $\frac{dW_{\bar{s}^1}}{d\theta}$ — and is sufficiently large.

The covariance term can only be non-zero when the ratio of Pareto weights $\alpha^i_{s^0}/\alpha^i_{\bar{s}^1}$ differs across individuals. This precisely corresponds to condition (iv) of Proposition 2. $\alpha^i_{s^t}$ represents the planner's valuation of individual i's gains and losses in utility units from the perspective of history s^t. Whenever the ratios $\frac{\alpha^i_{s^0}}{\alpha^i_{\bar{s}^1}}$ are not equalized in the cross section, the planner's relative valuation of individuals in utils changes over time. The covariance term in equation (10) is positive whenever the planner puts a smaller weight ex post on those individuals who have a relatively large welfare gain from the perturbation.

According to Proposition 2, interpersonal welfare comparisons become a new source of time inconsistency in HA environments. We now show that utilitarian is the unique social welfare function that ensures time consistency in interpersonal welfare comparisons.

Corollary 1 (Time Consistency under Utilitarian SWF)**.** *The utilitarian social welfare function is the only welfarist criterion for which interpersonal welfare comparisons do not lead to time consistency problems because the Pareto weights $\alpha^i_{s^t} = \alpha^i$ are constant across histories for all individuals i.*

When the social welfare function $\mathcal{W}(\cdot)$ is utilitarian and therefore linear, the individual Pareto weights $\alpha^i_{s^t}$ are constant across histories for each i. The covariance term in (10) therefore vanishes, so that the welfare assessment $\frac{dW_{s^0}}{d\theta}$ is time consistent under assumption 1. The utilitarian SWF is the only one for which this is the case. This result goes back to Zuber (2011) who showed 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 $\alpha^i_{s^t}$ is a function of lifetime utilities $V^i_{s^t}$ evaluated from the perspective of s^t.[13]

[12] To see this, note that the individual Pareto weights $\alpha^i_{s^0}$ and $\alpha^i_{\bar{s}^1}$ must be strictly positive. As are the discount factor β and the probability $\pi(\bar{s}^1)$ since we drop all histories from our analysis that realize with probability 0 without loss of generality.

[13] An interesting question is whether the analysis of Kydland and Prescott (1977) already nests Proposition 2

Corollary 2 (Role of Incomplete Markets)**.** *Assume that the SWF $\mathcal{W}(\cdot)$ is symmetric and individual preferences are symmetric. Then welfare assessments are time consistent even for non-linear SWFs as long as markets are complete so that $u'(c_1^i(\bar{s}^1))/u'(c_0^i(s^0)) = u'(c_1^j(\bar{s}^1))/u'(c_0^j(s^0))$.*

Corollary 2 clarifies the role incomplete markets play for the time consistency of welfare assessments, as distinct from individual preference heterogeneity. In particular, when markets are complete and individual marginal rates of substitution between consumption at different histories are equalized, then welfare assessments are time consistent even when the social welfare function is non-linear.

4.2 Inconsistency of Efficiency Assessments

We present our main result in this section: Efficiency and redistribution assessments are invariably time inconsistent in heterogeneous agent incomplete markets economies. Even when the overall welfare assessment $\frac{dW_{s^0}}{d\theta}$ is time consistent as in the special case of Corollary 1, whether welfare gains and losses are attributed to efficiency or redistribution therefore changes over time.

4.2.1 Efficiency

The notion of efficiency we work with in this paper is that of Kaldor-Hicks efficiency (Kaldor, 1939; Hicks, 1939). It defines the efficiency gain from a policy perturbation as the sum of individual willingness to pay expressed in a common unit.

Welfare numeraire and individual willingness to pay. The welfare assessment $\frac{dW_{s^t}}{d\theta} = \sum_i \alpha^i_{s^t} \frac{dV^i_{s^t}}{d\theta}$ compares gains and losses across different individuals in utils. The units of $\frac{dV^i_{s^t}}{d\theta}$ are history s^t individual utils. Those of $\frac{dW_{s^t}}{d\theta}$ are social utils. And the weights $\alpha^i_{s^t}$ represent the planner's valuation of individual i's gains and losses in utils.

for the heterogeneous agent case after an appropriate relabeling. Suppose we relabel their "x_t" as the vector of all individuals' consumption. Then the question 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 "$\partial S/\partial x_2$" would differ across the ex ante and ex post assessments, which is precisely the content of Proposition 2. Kydland and Prescott (1977) notationally force this term to be constant across assessment perspectives, which in general rules out heterogeneous agent environments. However, the utilitarian social welfare function is precisely the special case where their term "$\partial S/\partial x_2$" would have the same sign from the ex ante and ex post perspective. A useful restatement of Corollary 1 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.

Making interpersonal welfare comparisons in utility units is not meaningful since utility is *ordinal* — its units are meaningless as any monotonic transformation implies the same preference ordering. We start by expressing individual gains and losses in a comparable unit (such as a consumption-equivalent) and we refer to this unit as the *welfare numeraire*. Individual i's welfare gain or loss from the perturbation can then be expressed in units of the welfare numeraire as

$$\frac{\frac{dV^i_{s^t}}{d\theta}}{\lambda^i_{s^t}}, \quad \text{where} \quad \lambda^i_{s^t} > 0 \quad \text{and} \quad \dim(\lambda^i_{s^t}) = \frac{\text{history } s^t \text{ individual } i \text{ utils}}{\text{history } s^t \text{ welfare numeraire}}.$$

We use the *individual normalizing factor* $\lambda^i_{s^t}$ to express individual gains and losses in a common unit — the welfare numeraire. This welfare numeraire is chosen at the time of the assessment, history s^t. The units of the normalizing factor are $\dim(\lambda^i_{s^t}) = \frac{\text{history } s^t \text{ individual } i \text{ utils}}{\text{history } s^t \text{ welfare numeraire}}$. We have therefore expressed individual gains and losses as a *willingness to pay* in units of welfare numeraire, i.e.,

$$\dim\left(\frac{\frac{dV^i_{s^t}}{d\theta}}{\lambda^i_{s^t}}\right) = \frac{\frac{\text{history } s^t \text{ individual } i \text{ utils}}{\text{units of } \theta}}{\frac{\text{history } s^t \text{ individual } i \text{ utils}}{\text{history } s^t \text{ welfare numeraire}}} = \frac{\text{history } s^t \text{ welfare numeraire}}{\text{units of } \theta}.$$

Dividing $\frac{dV^i_{s^t}}{d\theta}$ by the normalizing factor $\lambda^i_{s^t}$ expresses individual gains and losses as a willingness to pay in units of welfare numeraire from the perspective of the point in time at which the welfare assessment takes place.[14]

The normalizing factor $\lambda^i_{s^t}$ and therefore the welfare numeraire itself may depend on the time of the welfare assessment s^t. Why is this the case? Suppose we choose date 0 consumption as the unit in which we express the welfare assessment and the interpersonal welfare comparisons which they aggregate at that point in time. If we reconsider such a welfare assessment at a later date, for example from the perspective of history s^1, expressing individual willingness to pay in units of date 0 consumption is no longer feasible and a different welfare numeraire must be chosen, such as history s^1 consumption for example. As a consequence, the units in which we express welfare assessments depends on the perspective of the assessment via $\lambda^i_{s^t}$ and may change over time.

Kaldor-Hicks efficiency. Having expressed individual welfare gains and losses in a common and meaningfully comparable unit, we now revisit the overall welfare assessment that compares

[14] To express individual willingness to pay in units of date t history s^t consumption, for example, we set $\lambda^i_{s^t} = u'(c^i_t(s^t))$. Its units $\dim(u'(c^i_t(s^t)))$ are individual i utils per unit of consumption. And therefore the units of $\frac{dV^i_{s^t}}{d\theta}/\lambda^i_{s^t}$ would be units of date t history s^t consumption per unit of θ.

and aggregates them. To express the assessment $\frac{dW_{s^t}}{d\theta}$ in units of the welfare numeraire as well, we divide by $\frac{1}{I}\sum_i \alpha^i_{s^t}\lambda^i_{s^t}$. A *normalized* welfare assessment from the perspective of history s^t can then be expressed in units of the welfare numeraire as

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

where $\omega^i_{s^t} = \frac{1}{\frac{1}{I}\sum_i \alpha^i_{s^t}\lambda^i_{s^t}}\alpha^i_{s^t}\lambda^i_{s^t}$ is a normalized individual welfare weight. We use the superscript $^\lambda$ notation to emphasize that the units of $\frac{dW^\lambda_{s^t}}{d\theta}$ are in welfare numeraire. If $\frac{dW^\lambda_{s^t}}{d\theta} = 1.1$, for example, then the planner values the policy change $d\theta$ equal to distributing 1.1 units of welfare numeraire equally across individuals at history s^t.

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

$$\frac{dW^\lambda_{s^t}}{d\theta} = \sum_i \omega^i_{s^t}\frac{\frac{dV^i_{s^0}}{d\theta}}{\lambda^i_{s^t}} = \underbrace{\sum_i \frac{\frac{dV^i_{s^t}}{d\theta}}{\lambda^i_{s^t}}}_{\Xi^E_{s^t}\text{(Efficiency)}} + \underbrace{\mathbb{C}ov^\Sigma_i\left(\omega^i_{s^t}, \frac{\frac{dV^i_{s^t}}{d\theta}}{\lambda^i_{s^t}}\right)}_{\Xi^{RD}_{s^t}\text{(Redistribution)}}. \tag{12}$$

The efficiency component $\Xi^E_{s^0}$ corresponds to Kaldor-Hicks efficiency gains — the unweighted sum of individual willingness to pay — from the perspective of s^0. The normalized individual welfare weight $\omega^i_{s^0}$ encodes how the SWF $\mathcal{W}$ trades off welfare gains and losses across individuals in units of welfare numeraire. For instance if $\omega^i_{s^0} = 1.3$, then the planner finds the welfare gain from giving 1 unit of numeraire to individual i equivalent to distributing 1.3 units equally across all individuals. This welfare weight may depend on the time of the assessment s^0 through both the unnormalized weight $\alpha^i_{s^0}$ and the choice of numeraire via $\lambda^i_{s^0}$.

The sum of willingness to pay $\Xi^E_{s^t}$ in units of welfare numeraire is a useful measure of efficiency gains and losses for at least three reasons. First, it is invariant to preference-preserving transformations of utilities $u^i(\cdot)$, whereas the overall welfare assessment dW_{s_t} is not. Second, it is also invariant to the choice of SWF $\mathcal{W}(\cdot)$. Notice from the definition of efficiency in (12) that $\Xi^E_{s^t}$ is invariant to $\alpha^i_{s^t}$, which only affect the normalized welfare weights $\omega^i_{s^t}$ and therefore the redistribution component. Finally, Kaldor-Hicks efficiency satisfies the compensation principle. When the Kaldor-Hicks efficiency assessment of a perturbation $d\theta$ is positive, then the perturbation can be turned into a Pareto improvement by implementing a set of compensating transfers across individuals. In fact, it is the only welfare criterion that satisfies this property, which we will refer to as "Paretian with transfers".

4.2.2 The Inconsistency of Efficiency Assessments

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

Definition 2. Consider a perturbation $d\theta_t(s^t)$. The efficiency assessment $\Xi^E_{s^0}$ from the perspective of s^0 is *time consistent* if later assessments $\Xi^E_{s^k}$ share the same sign for all histories $s^0 < s^k \leq s^t$. The efficiency assessment is time *in*consistent if there is a history s^k for which this does not hold.

As before, we focus on a perturbation $d\theta_1(\bar{s}^1)$ in the main text, which we simply refer to as $d\theta$. And we compare the ex ante assessment from the perspective of s^0 with the ex post assessment at $\bar{s}^1$,

$$\Xi^E_{s^0} = \sum_i \frac{\frac{dV^i_{s^0}}{d\theta}}{\lambda^i_{s^0}} \qquad \text{and} \qquad \Xi^E_{\bar{s}^1} = \sum_i \frac{\frac{dV^i_{\bar{s}^1}}{d\theta}}{\lambda^i_{\bar{s}^1}}.$$

Our proofs in Appendix B are stated for the general case. Two observations follow from comparing the ex ante and ex post efficiency assessments. First, neither $\Xi^E_{s^0}$ nor $\Xi^E_{\bar{s}^1}$ depends on the Pareto weights $\alpha^i_{s^0}$ or $\alpha^i_{\bar{s}^1}$. As discussed in Section 4.2.1, Kaldor-Hicks efficiency is invariant to the underlying social welfare function. It follows that condition (iv) of Proposition 2, which concerns the evolution of relative Pareto weights over time, does not apply when evaluating efficiency.

Second, comparing ex ante and ex post efficiency assessments does require taking a stance on the welfare numeraires — that is, the units in which individual willingness to pay is expressed in each assessment. The next result shows that inconsistency in individuals' valuations for these numeraires can give rise to time inconsistency in efficiency assessments.

Proposition 3 (Inconsistency of Efficiency Assessments)**.** *Suppose $I > 1$. For the date 0 efficiency assessment of perturbation $d\theta$ to be time inconsistent, either one of (i) – (iii) in Proposition 1 must be true or:*

(v) (Numeraire inconsistency) $\lambda^i_{\bar{s}^1}/\lambda^i_{s^0} \neq \lambda^j_{\bar{s}^1}/\lambda^j_{s^0}$ *for at least two individuals i and j.*

Proposition 3 establishes that efficiency assessments in heterogeneous agent environments can be time inconsistent even when Pareto weights $\alpha^i_{s^t}$ remain fixed over time. This form of inconsistency arises independently of the underlying social welfare function. In particular, even

when the overall welfare assessment is time consistent — as in the special case of Corollary 1 — the attribution of welfare gains and losses to efficiency and redistribution invariably changes over time.

Proposition 3 therefore identifies a novel source of time inconsistency that emerges in heterogeneous agent environments due to the choice of welfare numeraires. While time inconsistency in the overall welfare assessment results from the evolution of Pareto weights $\alpha^i_{s^t}$ under a non-linear social welfare function, efficiency assessments remain invariant to such changes. Instead, a new condition — numeraire inconsistency — emerges as the relevant criterion. The inconsistency stems not from a change in the social valuation of individual utilities, but from a change in the reference units used to aggregate individual consumption-equivalent valuations.[15]

What then explains the inconsistency of efficiency assessments? To build intuition for the economics underlying Proposition 3, we derive a useful representation of the relationship between the ex ante and ex post efficiency assessments. Under assumption 1, we can write

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

In other words, the ex ante and ex post assessments are proportional to each other (the scalar coefficient on $\Xi^i_{\bar{s}^1}$ is always positive), except for the covariance term on the RHS.[16] If the covariance term were 0, the assessment would be time consistent since the sign would always be preserved. However, whenever the covariance term has the opposite sign from $\Xi^E_{\bar{s}^0}$ and is sufficiently large, the efficiency assessment becomes inconsistent. And the covariance term is positive when those individuals who have a large willingness to pay for the perturbation in units of ex post numeraire also have a high relative valuation for ex post numeraire relative to ex ante numeraire.

In Section 4, we studied welfare assessments that compared gains and losses across individuals in utility units. Time inconsistency emerged whenever the planner's relative valuation for individuals in utils changed over time — the ratio $\frac{\alpha^i_{\bar{s}^1}}{\alpha^i_{s^0}}$ in the covariance term of equation (10).

[15] In any RA environment, welfare and efficiency coincide, so our discussion of time consistency in Proposition 1 also applies to efficiency assessments. In HA environments, however, we show that the conditions for time consistency of the efficiency assessment are different from those for welfare assessments.

[16] The scalar multiplying $\Xi^E_{\bar{s}^1}$ is always positive because welfare numeraires $\lambda^i_{s^t}$ and the discount factor β are strictly positive by assumption. The probability of history $\bar{s}^1$ is weakly positive. But whenever $\pi(\bar{s}^1)$ is 0, we also have $\Xi^E_{s^0} = 0$, in which case the sign is again preserved. We assume $\pi(\bar{s}^1) > 0$ in the main text for simplicity.

Utilitarian welfare assessments remain time consistent precisely because the planner's ex ante and ex post valuation of gains and losses for an individual i remain constant in utils. When we make efficiency assessments, however, we compare gains and losses across individuals in a common welfare numeraire. The normalizing factor $\lambda^i_{s^t}$ captures individual i's valuation of this numeraire from the perspective of history s^t. That is, $\lambda^i_{s^t}$ captures how many utils individual i derives from one unit of the numeraire. And 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 util 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.

Individuals' valuations for the ex ante and ex post numeraires may differ independently from the underlying social welfare function. Corollary 1 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 numeraires may differ regardless.

Forward- and backward-looking numeraires. Proposition 3 suggests a natural approach to ensure the time consistency of efficiency assessments: choosing ex ante and ex post numeraires such that $\lambda^i_{s^0} = \lambda^i_{\bar{s}^1}$ for all i. If each individual's valuation of the numeraire is constant across time and histories, then the covariance term in (13) drops out. It will be useful to distinguish between two classes of possible welfare numeraires: forward- and backward-looking numeraires.

Definition 3 (Forward-Looking Welfare Numeraire)**.** A welfare numeraire is *forward-looking* from the perspective of history s^k if individual i's valuation $\lambda^i_{s^k}$ is a function only of $\{c^i_t(s^t)\}_{t\geq k,\, s^t \geq s^k}$.

If $\lambda^i_{s^k}$ is also a function of $\{c^i_t(s^t)\}_{t<k,\, s^t}$ or $\{c^i_t(s^t)\}_{t,\, s^t \not\geq s^k}$, on the other hand, we say that the associated numeraire is *backward-looking*. In particular, the numeraire $\lambda^i_{\bar{s}^1}$ we use for the ex post efficiency assessment is forward-looking if it is a function only of consumption at history $\bar{s}^1$ or at future histories s^k that can be reached from $\bar{s}^1$. When we use a forward-looking numeraire, we express individuals' willingness to pay in units of contemporaneous or future consumption. Differences in marginal utilities of past consumption do not influence the valuation. The two most common examples of forward-looking numeraires

are contemporaneous consumption, with $\lambda^i_{\bar{s}^1} = (u^i)'(c^i_1(\bar{s}^1))$, and perpetual consumption, with $\lambda^i_{\bar{s}^1} = \sum_{t\geq 1}\sum_{s^t\geq\bar{s}^1}\beta^{t-1}\pi(s^t \mid \bar{s}^1)(u^i)'(c^i_t(s^t))$.

One natural candidate to ensure $\lambda^i_{s^0} = \lambda^i_{\bar{s}^1}$ for all 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 and losses across individuals in units of past (last period's) consumption. Efficiency assessments based on backward-looking numeraires no longer satisfy the compensation principle, however, as we show in Proposition 4 below. In other words, even if the implied sum of willingness to pay is positive in units of the backward-looking numeraire, it would not be possible to construct (hypothetical) compensating transfers to turn the perturbation into a Pareto improvement. That is because individuals have no value for history s^0 consumption from the perspective of history $\bar{s}^1$. So we cannot use this unit for compensation. In summary, while it appears straightforward to ensure the time consistency of efficiency assessments by making the ex post assessment in units of a backward-looking numeraire, such assessments no longer satisfy the compensation principle. We formalize these observations in the next Proposition.

Proposition 4 (Properties of Welfare Numeraires)**.** *Forward-looking welfare numeraires satisfy the following two conditions:*

1. *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.*

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

Efficiency assessments must therefore use forward-looking numeraires in order to be Paretian with transfers. But forward-looking numeraires naturally invite time consistency problems. Consider the case where we use contemporaneous consumption for both the ex ante and ex post assessments, so $\lambda^i_{s^0} = (u^i)'(c_0(s^0))$ and $\lambda^i_{\bar{s}^1} = (u^i)'(c^i_1(\bar{s}^1))$. Whenever individuals' valuations for consumption at these two histories are not equalized, the covariance term in equation (13) will be non-zero.[17]

[17] One might consider using consumption at history $\bar{s}^1$ as the numeraire for the ex ante assessment at date 0. Both assessments would then be based on forward-looking numeraires and trivially satisfy $\lambda^i_{s^0} = \lambda^i_{\bar{s}^1}$. This will not be possible in all but the simplest settings, however. Consider a policy that affects the allocation not only in history $\bar{s}^1$ but also in a different history $\tilde{s}^1$ at date 1. Time consistency then requires that ex post assessments in all relevant histories at date 1 share the same sign as the ex ante assessment at date 0. If we

Complete markets. There is a special class of environments in which numeraire inconsistency never arises and efficiency assessments are therefore time consistent under assumption 1. This is true whenever marginal rates of substitution between histories are equalized across all individuals — a condition satisfied in complete markets environments. Under complete markets,

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

for all individuals i and j. The next result shows that this condition guarantees numeraire consistency.

Proposition 5 (Consistency under Complete Markets)**.** *In environments in which all marginal rates of substitution are equalized, so equation* (14) *is satisfied, we have*

$$\frac{\lambda^i_{\bar{s}^1}}{\lambda^i_{s^0}} = \frac{\lambda^j_{\bar{s}^1}}{\lambda^j_{s^0}} \qquad \textit{for all } i \textit{ and } j. \tag{15}$$

Assumption 1 then guarantees the time consistency of efficiency assessments.

When markets are complete, individuals' relative valuation of consumption at any two histories is equalized (equation 14). Any valid welfare numeraire must correspond to a bundle of consumption at different histories and $\lambda^i_{s^t}$ 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 are equalized. Condition (v) of Proposition 3 therefore never applies in complete markets environments. And under the no-KP assumption 1, which rules out conditions (i) – (iii) of Proposition 1, efficiency assessments are therefore always time consistent.

4.2.3 Impossibility Result

Drawing on our discussion thus far, we now present the central result of this paper: We show that it is generally not possible to make efficiency assessments in heterogeneous agent incomplete markets environments that both (i) satisfy the compensation principle and (ii) are

make the ex ante assessment using history $\bar{s}^1$ consumption as the numeraire and history $\bar{s}^1$ realizes ex post, no numeraire inconsistency will arise and the assessment will be consistent. If history $\tilde{s}^1$ realizes, however, numeraire inconsistency will again be a problem as long as individuals' valuations of consumption across these two histories are not equalized.

time consistent. Both qualifications are critical. This result only applies to heterogeneous agent economies with $I > 1$. It also requires that marginal rates of substitution across consumption in different histories are not equalized for all individuals in the cross section. We summarize these conditions in the following definition.

Definition 4. We refer as *heterogeneous agent incomplete markets* economies to environments with $I > 1$ where there are at least two individuals i and j whose marginal rates of substitution between histories s^0 and $\bar{s}^1$ are not equalized,

$$\frac{(u^i)'(c_1^i(\bar{s}^1))}{(u^i)'(c_0^i(s^0))} \neq \frac{(u^j)'(c_1^j(\bar{s}^1))}{(u^j)'(c_0^j(s^0))}. \tag{16}$$

We are now ready to state our main result.

Theorem 1 (Impossibility Result)**.** *In heterogeneous agent incomplete markets economies, it is not possible to make efficiency assessments that satisfy the following three properties for all $d\theta$:*

P1 *Forward-looking numeraires*

P2 *Time consistency*

In the presence of heterogeneity and market incompleteness, it is not possible to make efficiency assessments that simultaneously satisfy the desirable properties P1 – P3.[18] Theorem 1 shows that there is an intrinsic tension between making efficiency assessments that satisfy the compensation principle (are Paretian with transfers) and are time consistent.
This result is not about one of the sources of time inconsistency identified by Kydland and Prescott (1977), i.e., conditions (i) – (iii) of Proposition 1. Theorem 1 applies even to environments that additionally satisfy assumption 1, shutting down time consistency problems á la Kydland and Prescott (1977), which also apply to representative agent environments. Instead, the key forces driving Theorem 1 are unique to heterogeneous agent environments with $I > 1$.

A key step in the proof of Theorem 1 shows that property P2 implies that P3 cannot be true for all $d\theta$ whenever markets are incomplete. The argument builds on our discussion in Section

[18]Theorem 1 is a statement of generic impossibility in the spirit of Arrow (1950). What this means is that, taking as given a particular environment that satisfies $I > 1$ and (16), it is impossible to make an efficiency assessment that satisfies P1 – P3 *for all* $d\theta$. In other words, we can always construct a perturbation $d\theta$ whose efficiency assessment violates one of P1 – P3. The "for all $d\theta$" condition is therefore parallel in spirit to Arrow (1950)'s condition of universal domain.

4.2.2. Kaldor-Hicks efficiency uniquely satisfies the compensation principle — and is therefore Paretian with transfers — when it is based on forward-looking numeraires. And we show in our proof in Appendix B.5 that making both the ex ante and ex post efficiency assessments using forward-looking numeraires (generically) leads to time inconsistency. That is, it is possible to construct perturbations $d\theta$ whose ex ante and ex post efficiency assessments will differ in sign.

4.3 Inconsistency of The Origins of Efficiency Gains

We have so far characterized the time consistency of the overall welfare assessment and of the attribution of welfare gains and losses to efficiency and redistribution. We now unpack the efficiency assessment and discuss the time consistency of aggregate efficiency, risk-sharing, and intertemporal-sharing assessments.

To that end, we start by defining normalized dynamic and stochastic welfare weights. Like the normalized individual welfare weights $\omega^i_{s^t}$, these weights capture the key marginal rates of substitution that govern the planner's relative valuation of changes in consumption bundles. Given a welfare numeraire $\lambda^i_{s^t}$, we define

$$\omega^i_{t|s^k} = \frac{1}{\lambda^i_{s^k}} \beta^{t-k} \sum_{s^t \geq s^k} \pi(s^t \mid s^k) u'(c^i_t(s^t))$$

as the normalized dynamic welfare weight for consumption of individual i in period t from the perspective of history s^k. Intuitively, $\omega^i_{t|s^k}$ captures individual i's marginal rate of substitution between a unit of consumption at date t (in all histories) and one unit of welfare numeraire. The former is valued at $\beta^{t-k} \sum_{s^t \geq s^k} \pi(s^t \mid s^k) u'(c^i_t(s^t))$ and the latter at $\lambda^i_{s^k}$. Since the social welfare function $\mathcal{W}(\cdot)$ is non-paternalistic by construction, the planner's MRS between these two bundles coincides with individual i's private MRS. If $\omega^i_{t|s^k} = 0.2$, for example, then individual i (and therefore the planner) is indifferent between 1 additional unit of consumption at date t and 0.2 units of welfare numeraire.

We also define the normalized stochastic welfare weight

$$\omega^i_{t|s^k}(s^t) = \frac{\pi(s^t \mid s^k) u'(c^i_t(s^t))}{\sum_{s^t \geq s^k} \pi(s^t \mid s^k) u'(c^i_t(s^t))}$$

for consumption of individual i in history s^t from the perspective of history s^k. $\omega^i_{t|s^k}(s^t)$ captures individual i's MRS between a unit of consumption in history s^t and a unit of consumption at date t (in all histories). If $\omega^i_{t|s^k}(s^t) = 0.4$, then individual i (and therefore the planner) is indifferent between 1 additional unit of consumption in history s^t and 0.4 units of consumption

in all histories at date t.

Finally, we define the cross-sectional averages of these normalized welfare weights by

$$\bar{\omega}_{t|s^k} = \frac{1}{I}\sum_i \omega^i_{t|s^k} \qquad \text{and} \qquad \bar{\omega}_{t|s^k}(s^t) = \frac{1}{I}\sum_i \omega^i_{t|s^k}(s^t).$$

We are now ready to characterize the sources of efficiency gains and losses from the perspective of history s^k.

Lemma 1 (Origins of Efficiency)**.** *The efficiency assessment $\Xi^E_{s^k}$ of a perturbation $d\theta$ from the perspective of history s^k can be decomposed into aggregate efficiency, risk-sharing, and intertemporal-sharing components,*

$$\Xi^E_{s^k} = \Xi^{AE}_{s^k} + \Xi^{RS}_{s^k} + \Xi^{IS}_{s^k},$$

where

$$\begin{aligned}
\Xi^{AE}_{s^k} &= \sum_{t\geq k} \bar{\omega}_{t|s^k} \sum_{s^t \geq s^k} \bar{\omega}_{t|s^k}(s^t) \sum_i \frac{\partial \mathcal{C}^i_t(s^t \mid s^k)}{\partial \theta} \\
\Xi^{RS}_{s^k} &= \sum_{t\geq k} \bar{\omega}_{t|s^k} \sum_{s^t \geq s^k} \mathbb{C}ov^{\Sigma}_i\left(\omega^i_{t|s^k}(s^t),\ \frac{\partial \mathcal{C}^i_t(s^t \mid s^k)}{\partial \theta}\right) \\
\Xi^{IS}_{s^k} &= \sum_{t\geq k} \mathbb{C}ov^{\Sigma}_i\left(\omega^i_{t|s^k},\ \sum_{s^t \geq s^k} \omega^i_{t|s^k}(s^t) \sum_i \frac{\partial \mathcal{C}^i_t(s^t \mid s^k)}{\partial \theta}\right)
\end{aligned}$$

Lemma 1 characterizes the sources of efficiency gains and losses from the perspective of history s^k, building on results in Dávila and Schaab (2024).

Proposition 6 (Aggregate Efficiency)**.** *Under assumption 1, assessments of static aggregate efficiency gains and losses are time consistent.*

This result underscores that assessments of production efficiency gains and losses are special. While risk-sharing and redistribution assessments change over time, assessments of production efficiency gains and losses are time consistent. This highlights that some components of the overall efficiency and welfare assessment are time consistent, while others are not.

5 Applications

This section presents three applications that demonstrate how the time inconsistency forces identified in this paper arise in practical settings. Each application illustrates 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.[19] But the time consistency problem identified here also emerges in other policy

[19]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 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.

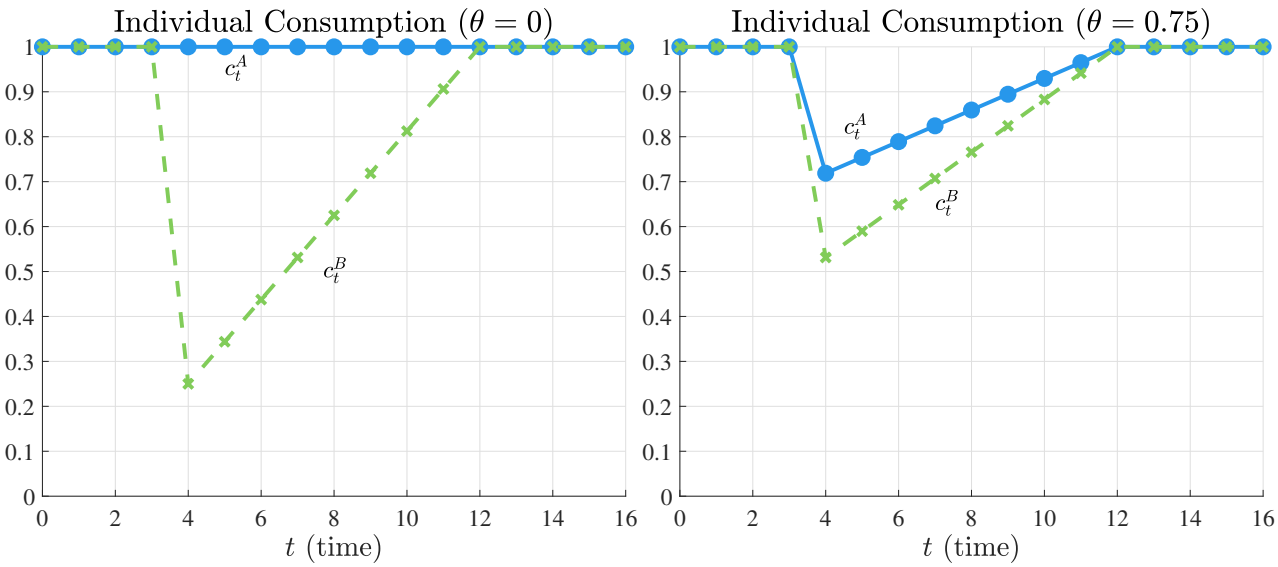


Figure 3: Consumption Profiles (Application 1)

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.

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

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

Individual i's endowment at date t is denoted by y_t^i. Individual A has a constant endowment of $y_t^A = y$. 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_t^B, is given by

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

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

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

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

$$\omega^i_{|k} = \frac{\sum_{t\geq k} \beta^{t-k} u'\left(c^i_t\right)}{\sum_i \sum_{t\geq k} \beta^{t-k} u'\left(c^i_t\right)} \quad \text{and} \quad \omega^i_{t|k} = \frac{\beta^{t-k} u'\left(c^i_t\right)}{\sum_{t\geq k} \beta^{t-k} u'\left(c^i_t\right)}. \tag{17}$$

Normalized welfare gains can be decomposed into efficiency and redistribution components, where $\frac{dW^\lambda_{|k}}{d\theta} = \Xi^E_{|k} + \Xi^{RD}_{|k}$. Efficiency gains are thus given by

$$\Xi^E_{|k} = \sum_i \sum_{t\geq k} \omega^i_{t|k} \frac{dc^i_t}{d\theta} = \sum_{t\geq k} \mathbb{C}ov^\Sigma_i \left[\omega^i_{t|k}, \frac{dc^i_t}{d\theta}\right],$$

where the second equality follows because this is an endowment economy, with $\sum_i \frac{dc^i_t}{d\theta} = 0$, and where $\mathbb{C}ov^\Sigma_i[\cdot,\cdot] = I \cdot \mathbb{C}ov_i[\cdot,\cdot]$ 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 $\Xi^E_{|k} = \Xi^{IS}_{|k}$. Redistribution gains, which capture whether individuals relative preferred by the planner are relatively favored by the policy, are given by

$$\Xi^{RD}_{|k} = \mathbb{C}ov^\Sigma_i \left[\omega^i_{|k}, \sum_{t\geq k} \omega^i_{t|k} \frac{dc^i_t}{d\theta}\right],$$

where individual weights $\omega^i_{|k}$ are defined in (17).

Normalized Individual Weights. The individual weight $\omega^i_{|k}$ 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 $\omega^B_{|k}/\omega^A_{|k}$ 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 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 $\omega^i_{t|k}$ 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

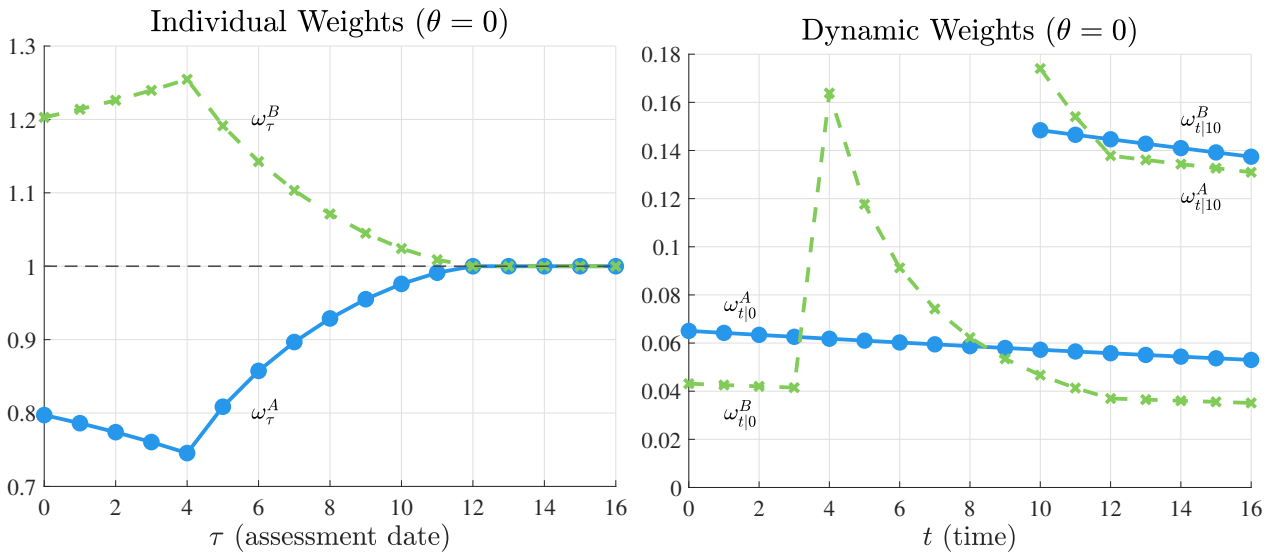


Figure 4: Individual and Dynamic Weights (Application 1)

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.

of perpetual consumption.[20] The right panel of Figure 4 shows the dynamic weights for each individual for two different perspectives $\omega^i_{t|k}$ 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 $\omega^i_{t|k}$ 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 $\omega^i_{t|k}$ at date t to the individual with the larger one.

When the low-endowment spell materializes at date $\underline{T} = 4$, $\frac{\omega^B_{4|0}}{\omega^A_{4|0}} \approx \frac{0.16}{0.06} = 2.67$, so the 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 $k = 10$

[20] 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.

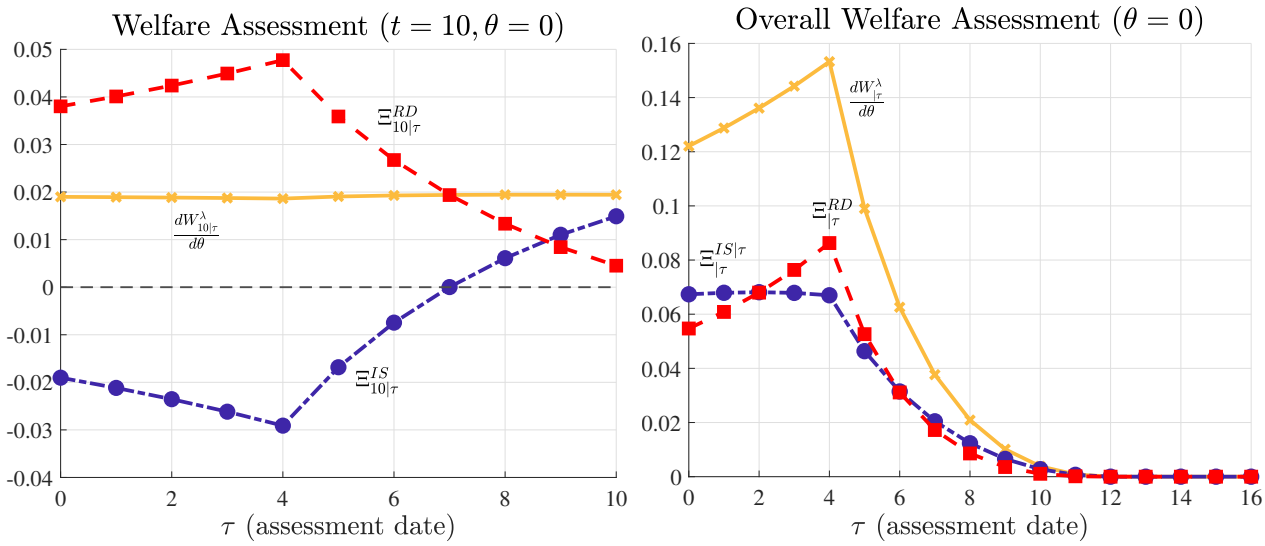


Figure 5: Welfare Assessment of Relief Policy (Application 1)

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.

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.[21] While 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

[21] Note that in this application we focus on the inconsistency of risk-sharing (efficiency) assessments *before* a terminal date. Once all risks are resolved, efficiency gains are zero, and welfare assessments can only only driven by redistribution.

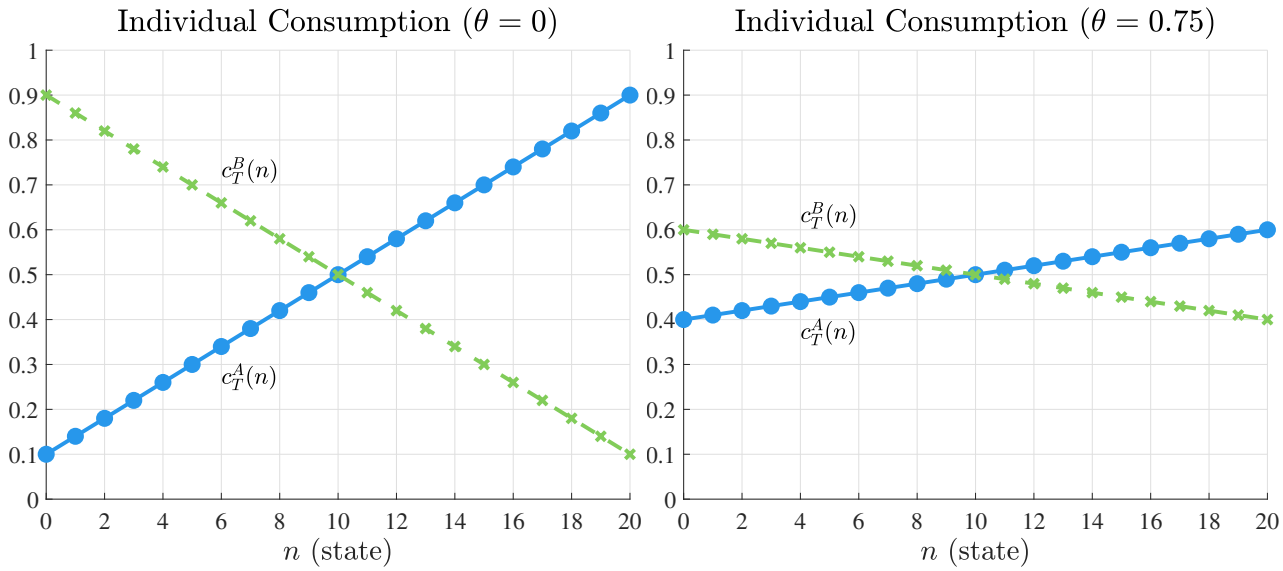


Figure 6: Consumption Profiles (Application 2)

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.

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

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

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.[22] Hence, from the perspective of an earlier date $t \leq T$, individual i's utility satisfies the recursion

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

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.

[22] 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 $\Delta t = t - \underline{t}$.

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

$$\omega_{|k,m}^i = \frac{\sum_n \pi_{T|t}(n|m)\, u'\left(c_T^i(n)\right)}{\sum_i \sum_n \pi_{T|t}(n|m)\, u'\left(c_T^i(n)\right)} \quad \text{and} \quad \omega_{T|k,m}^i(n) = \frac{\pi_{T|k}(n|m)\, u'\left(c_T^i(n')\right)}{\sum_{\tilde{n}} \pi_{T|k}(\tilde{n}|m)\, u'\left(c_T^i(\tilde{n})\right)}. \tag{18}$$

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 redistribution gains — using unconditional consumption across all states at date T as welfare

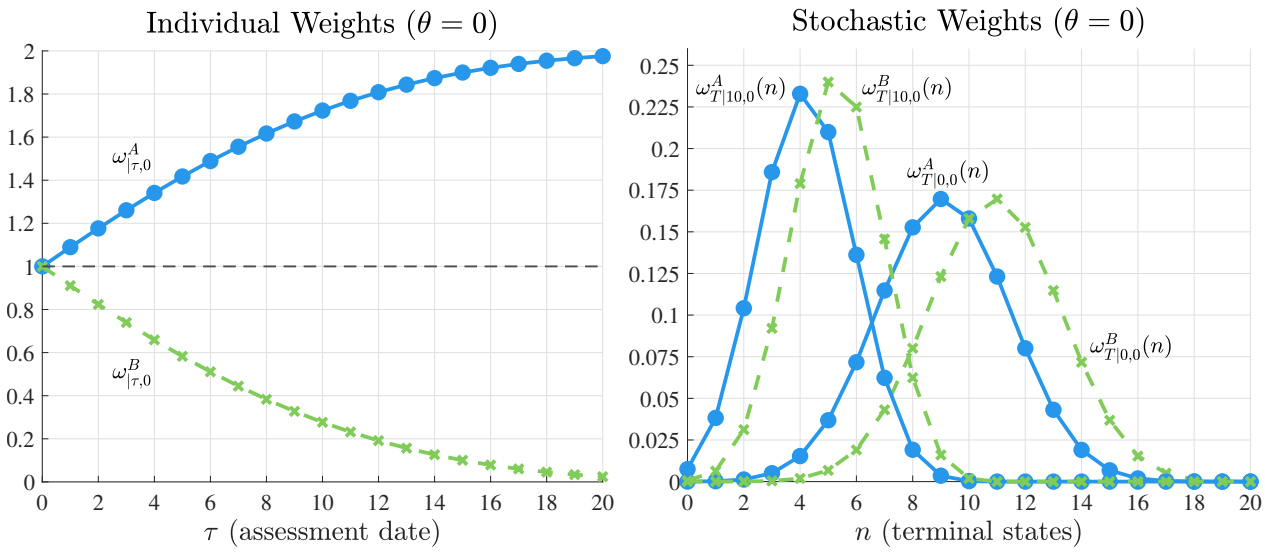


Figure 7: Individual and Stochastic Weights (Application 2)

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.

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 optimal ex ante are dismantled ex post. A real-world example is the shift in EU support

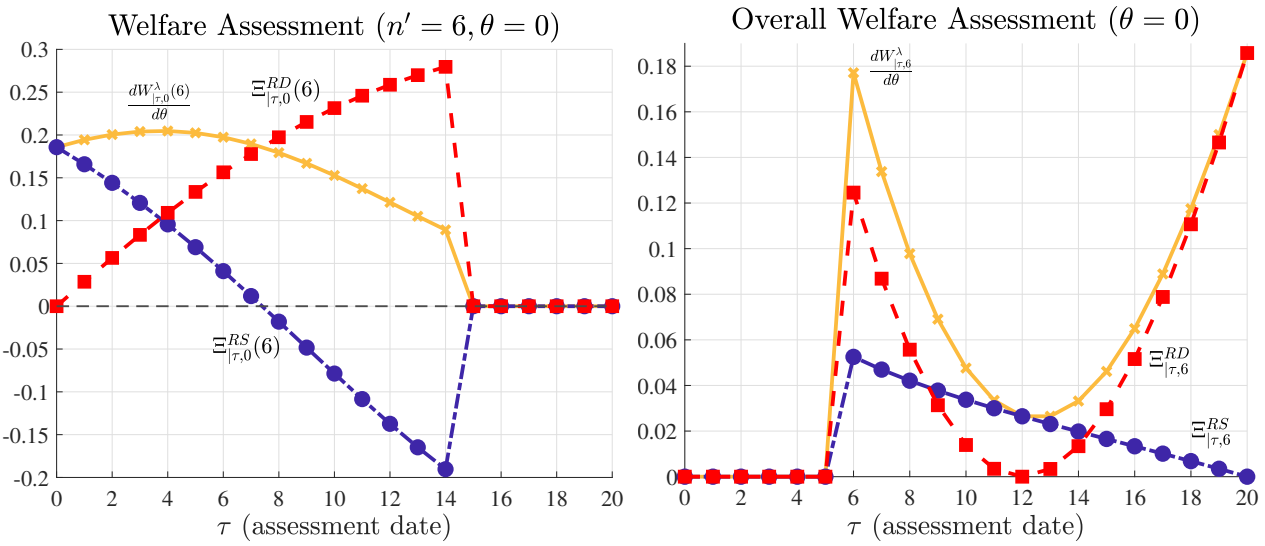


Figure 8: Welfare Assessment of Risk-Sharing Policy (Application 2)

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.

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 of Figure 8 shows how the overall risk-sharing gains slowly decrease in magnitude as time goes by,

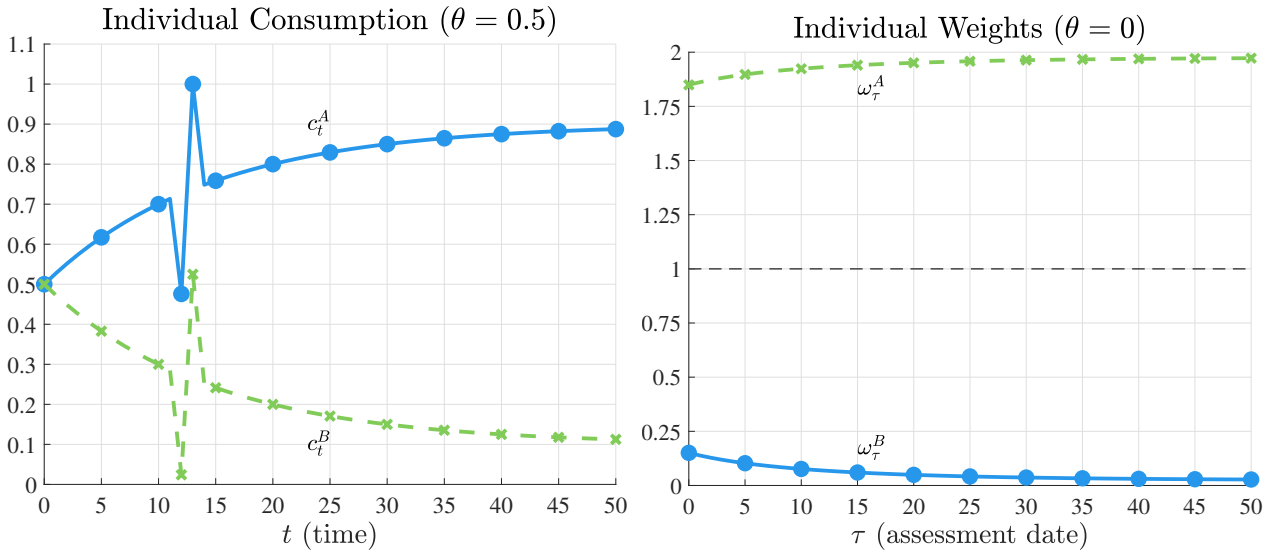


Figure 9: Consumption Profiles and Individual Weights (Application 3)

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 $\theta = 0$.

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

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

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 $\bar{\tau} + 1$,

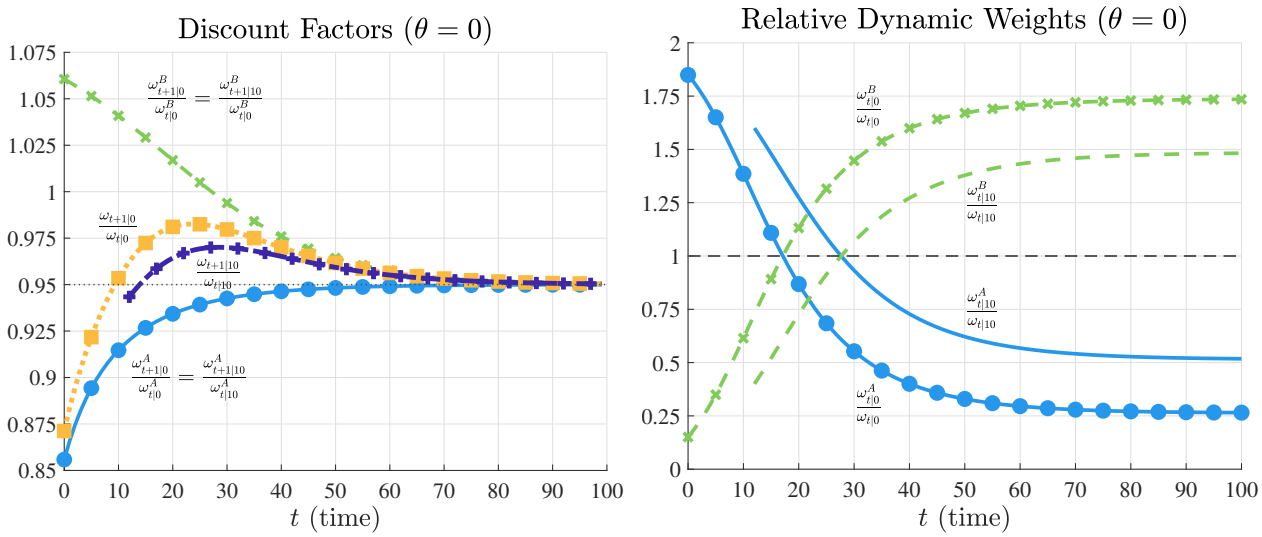


Figure 10: Discount Factors (Application 3)

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 (20).

where the cost and benefits of investment accrue symmetrically to each individual, so

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

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^A_t 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

$$\Xi^{E}_{|k} = \Xi^{AE}_{|k} = \sum_i \sum_{t \geq k} \omega^i_{t|k} \frac{dc^i_t}{d\theta} = \omega_{\bar{\tau}+1|k} z - \omega_{\bar{\tau}|k}, \quad \text{where} \quad \omega_{\bar{\tau}|k} = \frac{1}{I} \sum_i \omega^i_{\bar{\tau}|k}, \tag{19}$$

with dynamic weights $\omega^i_{t|k}$ defined as in (17), and where the second equality follows from the form of the particular policy considered. It follows directly from (19), that it is desirable to carry the aggregate investment policy from a date k perspective as long as

$$\frac{\omega_{\bar{\tau}+1|k}}{\omega_{\bar{\tau}|k}} > \frac{1}{z}, \quad \text{where} \quad \frac{\omega_{\bar{\tau}+1|k}}{\omega_{\bar{\tau}|k}} = \frac{1}{I} \sum_i \frac{\omega^i_{\bar{\tau}+1}}{\omega^i_{\bar{\tau}}} + \mathbb{C}ov_i \left[\frac{\omega^i_{\bar{\tau}|k}}{\omega_{\bar{\tau}|k}}, \frac{\omega^i_{\bar{\tau}+1}}{\omega^i_{\bar{\tau}}} \right]. \tag{20}$$

Since individual preferences are time consistent in this case, the term $\frac{\omega^i_{\tau+1}}{\omega^i_{\tau}}$ 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 $\omega^i_{|k}$ 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, $\frac{\omega^i_{\bar{\tau}+1}}{\omega^i_{\bar{\tau}}}$, 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 (20) — 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 1. *(Time Inconsistency with Shrinking Inequality)* The same form of present bias time inconsistency applies in economies in which individual consumption profiles converge. As in 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,

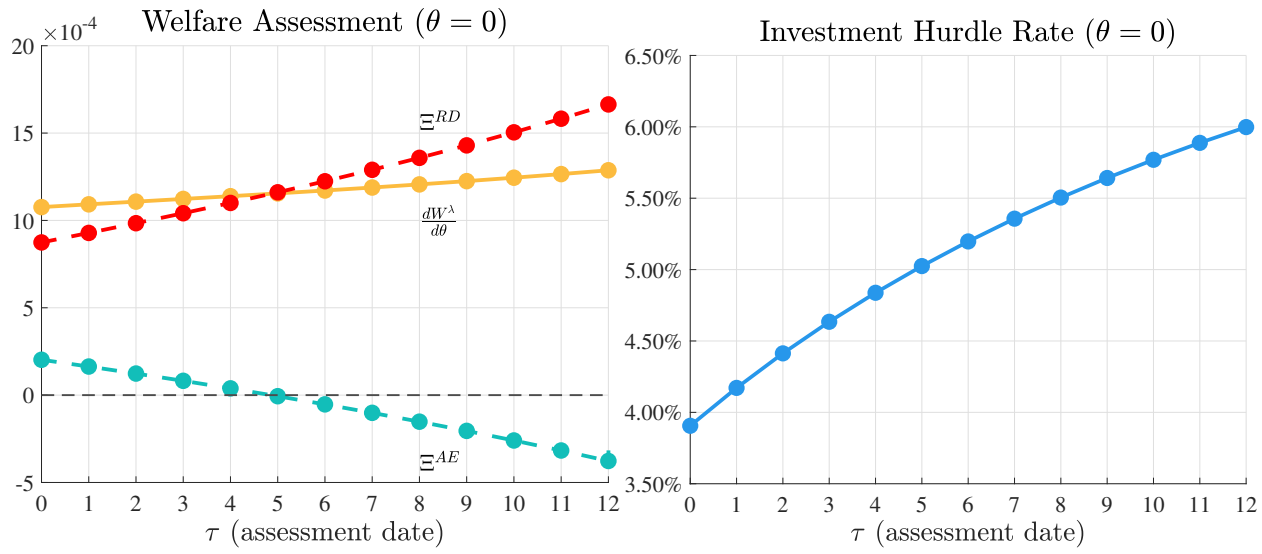


Figure 11: Welfare Assessment and Investment Hurdle Rate (Application 3)

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.

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 2. *(Relation to Jackson and Yariv (2014, 2015))* JYCITE 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. The time inconsistency that we identify in this application, which applies to the efficiency and redistribution components, has a resemblance to the JYCITE 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 studies the time inconsistency of policies in environments with heterogeneous agent and incomplete markets. Specifically, three broad conclusions emerge to guide future applied work on optimal policy design with heterogeneous agents.

First, interpersonal welfare comparisons emerge as a new source of time inconsistency when markets are incomplete or individual preferences are not symmetric. In the presence of discount factor or belief heterogeneity, no time-invariant social welfare function allows for time consistent assessments. When discount factors and beliefs are homogeneous, there is a sharp distinction between linear and non-linear social welfare functions. As long as markets are incomplete, only the utilitarian criterion implies time consistent assessments. When markets are complete and individual preferences symmetric, then welfare assessments under both linear and non-linear SWFs become time consistent. In summary, heterogeneous agent models may thus feature three distinct sources of time inconsistency: dynamically inconsistent individual preferences (Strotz, 1956), forward-looking behavior (Kydland and Prescott, 1977), and interpersonal welfare comparisons.

Second, the attribution of welfare gains to efficiency and redistribution is invariably time inconsistent when markets are incomplete, irrespective of the social welfare function and the extent of heterogeneity in individual preferences. In other words, even when the overall welfare assessment is time consistent in one of the special cases discussed above, whether those gains and losses are perceived as efficiency or redistribution invariably changes over time.

Finally, assessments of production efficiency gains stand out as a special case. Unlike risk-sharing and redistribution, production efficiency assessments can be time consistent for arbitrary social welfare functions and preference heterogeneity. This happens whenever the perturbation under consideration generates a production efficiency gain or loss at a single date and history. We show that this applies to a broad class of applied policy questions and therefore presents a relevant exception to the broader prevalence of time inconsistency.

Our results suggest the following rules-of-thumb for applied work on optimal policy design with heterogeneous agents: When using a utilitarian SWF in an environment without discount factor or belief heterogeneity, optimal policy is time consistent. However, the justification of optimal policy on efficiency or redistribution grounds changes as time passes and uncertainty is resolved. When either studying environments with discount factor and belief heterogeneity or using non-utilitarian SWFs, optimal policy will generally be time inconsistent.

It may seem a natural takeaway from this discussion to restrict the study of optimal policy design to environments with symmetric individual preferences and utilitarian planning objectives, but that restriction is untenably limiting. Mounting empirical evidence documents sizable dispersion in individual rates of time preference and beliefs. And many policy questions naturally arise in the context of non-utilitarian objectives — for example, when independent institutions or agencies set policies according to a narrowly defined mandate. Whenever such a mandate is not utilitarian, the implied optimal policy may be time inconsistent as a direct consequence of the interpersonal welfare comparisons implied by the mandate. Confronting these time consistency problems in heterogeneous agent models seems a valuable avenue for future work.

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. 1

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. 3

Alesina, A., and D. Rodrik (1994): "Distributive politics and economic growth," *The quarterly journal of economics*, 109(2), 465–490. 1

Arrow, K. J. (1950): "A difficulty in the concept of social welfare," *Journal of political economy*, 58(4), 328–346. 1, 18

——— (1963): "Uncertainty and the Welfare Economics of Medical Care," *American Economic Review*, 53(5), 941–973. 1

Athey, S., A. Atkeson, and P. J. Kehoe (2005): "The optimal degree of discretion in monetary policy," *Econometrica*, 73(5), 1431–1475. 3

Atkinson, A. B. (1970): "On the measurement of inequality," *Journal of economic theory*, 2(3), 244–263. 1, 2.2, 7

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. 1, 7

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. 1, 3

Auclert, A., M. Cai, M. Rognlie, and L. Straub (2024): "Optimal long-run fiscal policy with heterogeneous agents," . 8

Auclert, A., M. Rognlie, and L. Straub (2025): "Fiscal and monetary policy with heterogeneous agents," *Annual Review of Economics*, 17. 1

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. 3

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. 1, 2.2, 7

Bhandari, A., D. Evans, M. Golosov, and T. Sargent (2024): "Efficiency, insurance, and redistribution effects of government policies," Discussion paper, Working paper. 1

Bilbiie, F. O. (2024): "Monetary policy and heterogeneity: An analytical framework," *Review of Economic Studies*, p. rdae066. 1

Bisin, A., A. Lizzeri, and L. Yariv (2015): "Government policy with time inconsistent voters," *American Economic Review*, 105(6), 1711–1737. 1

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. 1, 7

Chambers, C. P., and F. Echenique (2018): "On multiple discount rates," *Econometrica*, 86(4), 1325–1346. 1

Chang, R. (2022): "Should Central Banks Have an Inequality Objective?," Discussion paper, National Bureau of Economic Research. 1

Chari, V. V., and P. J. Kehoe (1990): "Sustainable plans," *Journal of political economy*, 98(4), 783–802. 3

Clayton, C., and A. Schaab (2025): "A theory of dynamic inflation targets," *American Economic Review*, 115(2), 448–490. 3

Dávila, E., and A. Schaab (2023): "Optimal monetary policy with heterogeneous agents: discretion, commitment, and timeless policy," . 1, 8

——— (2024): "Welfare Assessments with Heterogeneous Individuals," . 1, 4.2.1, 4.3

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. 1, 5, 7

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

——— (2021): "Rules without commitment: Reputation and incentives," *The Review of Economic Studies*, 88(6), 2833–2856. 3

Farhi, E., C. Sleet, I. Werning, and S. Yeltekin (2012): "Non-linear capital taxation without commitment," *Review of Economic Studies*, 79(4), 1469–1493. 1

Farhi, E., and I. Werning (2013): "Insurance and taxation over the life cycle," *Review of Economic Studies*, 80(2), 596–635. 1

Feldstein, M. S. (1964): "The social time preference discount rate in cost benefit analysis," *The economic journal*, 74(294), 360–379. 1

Golosov, M., M. Troshkin, and A. Tsyvinski (2016): "Redistribution and social insurance," *American Economic Review*, 106(2), 359–386. 1

Halac, M., and P. Yared (2014): "Fiscal Rules and Discretion Under Persistent Shocks," *Econometrica*, 82(5), 1557–1614. 3

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

Harsanyi, J. C. (1955): "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility," *Journal of political economy*, 63(4), 309–321. 1

Heathcote, J., K. Storesletten, and G. L. Violante (2017): "Optimal tax progressivity: An analytical framework," *The Quarterly Journal of Economics*, 132(4), 1693–1754. 2.2, 7

Hicks, J. R. (1939): "The foundations of welfare economics," *The economic journal*, 49(196), 696–712. 1, 4.2.1

Jackson, M. O., and L. Yariv (2014): "Present bias and collective dynamic choice in the lab," *American Economic Review*, 104(12), 4184–4204. 2, 1, 2, B.2

——— (2015): "Collective dynamic choice: the necessity of time inconsistency," *American Economic Journal: Microeconomics*, 7(4), 150–178. 2, 1, 2, B.2

Kaldor, N. (1939): "Welfare propositions of economics and interpersonal comparisons of utility," *The economic journal*, 49(195), 549–552. 1, 4.2.1

Kaneko, M., and K. Nakamura (1979): "The Nash social welfare function," *Econometrica: Journal of the Econometric Society*, pp. 423–435. 7

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

Knight, F. H. (1921): *Risk, uncertainty and profit*. Houghton Mifflin. 1

Krusell, P., V. Quadrini, and J.-V. R\'\ios-Rull (1997): "Politico-economic equilibrium

and economic growth," *Journal of Economic Dynamics and Control*, 21(1), 243–272. 1

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. (document), 1, 1, 4, 1, 4.1, 11, 13, 4.2.3, 6, B.5

——— (1980): "Dynamic optimal taxation, rational expectations and optimal control," *Journal of Economic Dynamics and control*, 2, 79–91. 3

Ljungqvist, L., and T. J. Sargent (2018): *Recursive macroeconomic theory*. MIT press. 3

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. 3

Marcet, A., and R. Marimon (2019): "Recursive contracts," *Econometrica*, 87(5), 1589–1631. 3

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

Marshall, A. (1920): *Principles of Economics, 8th edn (1920)*. 1

Millner, A., and G. Heal (2018): "Time consistency and time invariance in collective intertemporal choice," *Journal of Economic Theory*, 176, 158–169. 1

Mirrlees, J. A. (1971): "An exploration in the theory of optimum income taxation," *The review of economic studies*, 38(2), 175–208. 5

Mongin, P. (1995): "Consistent bayesian aggregation," *Journal of Economic Theory*, 66(2), 313–351. 2, 1

——— (1998): "The paradox of the Bayesian experts and state-dependent utility theory," *Journal of Mathematical Economics*, 29(3), 331–361. 2, 1

Nash, J. F. (1950): "The bargaining problem," *Econometrica*, 18(2), 155–162. 7

Persson, T., and G. E. Tabellini (1994): "Is inequality harmful for growth?," *The American Economic Review*, 84(3), 600–621. 1

Pigou, A. (1920): *The economics of welfare*. Macmillan. 1

Piketty, T., and E. Saez (2012): "A theory of optimal capital taxation," Discussion paper, National Bureau of Economic Research. 1, 7

——— (2013a): "A theory of optimal inheritance taxation," *Econometrica*, 81(5), 1851–1886. 1, 7

——— (2013b): "Optimal labor income taxation," in *Handbook of public economics*, vol. 5, pp. 391–474. Elsevier. 1

Rawls, J. (1971): "A theory of justice," . 7

Rogoff, K. (1985): "The Optimal Degree of Commitment to an Intermediate Monetary Target," *The Quarterly Journal of Economics*, 100(4), 1169–1189. 3

Saez, E. (2001): "Using elasticities to derive optimal income tax rates," *The review of economic studies*, 68(1), 205–229. 1, 5, 7

Saez, E., and S. Stantcheva (2016): "Generalized social marginal welfare weights for optimal tax theory," *American Economic Review*, 106(01), 24–45. 1, 7

Stern, N. H. (1976): "On the specification of models of optimum income taxation," *Journal of Public economics*, 6(1-2), 123–162. 1, 7

Strotz, R. H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 23(3), 165–180. (document), 1, 1, 3, 4.1, 4.1, 6

Sublet, G. (2023): "The Optimal Degree of Discretion in Fiscal Policy," . 3

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

Werning, I. (2007): "Optimal fiscal policy with redistribution," *The Quarterly Journal of Economics*, 122(3), 925–967. 1, 4

Zuber, S. (2011): "The aggregation of preferences: can we ignore the past?," *Theory and decision*, 70, 367–384. 1, 4.1

Online Appendix

A Proofs and Derivations: Section

In this Appendix, we provide additional derivations and details for Section 2. Appendix A.1 characterizes the implementability conditions for the ex ante and ex post planning problems. Appendices A.2 and A.3 present derivations of equations (4) and (5) in the main text. Finally, Appendix A.4 presents details of our numerical implementation.

A.1 Implementability Conditions

Recall the household budget constraints from the main text, $c_0^i = w_0 \ell_0^i$ and $c_1^i = (1-\tau) w_1 z_1^i \ell_1^i + T$. Maximizing preference simplies the FOCs

$$\begin{aligned} u_{\ell,0}^i &= -u_{c,0}^i w_0 \\ u_{\ell,1}^i &= -u_{c,1}^i (1-\tau) w_1 z_1^i. \end{aligned}$$

The firm's profit maximization optimality condition implies

$$w_t = 1.$$

We use this firm optimality condition to solve out for the real wage. The remaining equilibrium conditions can therefore be written as

$$\begin{aligned} u_{\ell,0}^i &= -u_{c,0}^i \\ u_{\ell,1}^i &= -u_{c,1}^i (1-\tau) z_1^i \\ c_0^i &= \ell_0^i \\ c_1^i &= (1-\tau) z_1^i \ell_1^i + T \\ Y_t &= L_t \\ T &= \tau \int_0^1 z_1^i \ell_1^i \, di \\ Y_t &= \int_0^1 c_t^i \, di \\ L_t &= \int_0^1 z_t^i \ell_t^i \, di. \end{aligned}$$

Notice that fiscal policy has no impact on the date 0 allocation.

A.2 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 conditions of competitive equilibrium listed above. At an optimum, the necessary optimality condition is

$$\begin{aligned} 0 &= \frac{dW_0}{d\tau} \\ 0 &= \frac{1}{\int_0^1 \lambda_0^i \alpha_0^i \, di} \frac{dW_0}{d\tau} \\ 0 &= \int_0^1 \frac{1}{\lambda_0^i} \frac{dV_0^i}{d\tau} \, di + \mathbb{C}ov_i \left(\omega_0^i, \, \frac{dV_0^i}{d\tau} \right), \end{aligned}$$

where we define the normalized individual welfare weight

$$\omega_0^i = \frac{\alpha_0^i \lambda_0^i}{\int_0^1 \alpha_0^i \lambda_0^i \, di}.$$

The second line represents a normalization that expresses the first-order condition in units of welfare numeraire. And the third line decomposes the welfare gain $\frac{dW_0}{d\tau}$ into an efficiency and a redistribution gain.

Next, notice that for any choice of social welfare function $\mathcal{W}(\cdot)$ and for any choice of welfare ex ante numeraire λ_0^i, 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

$$\begin{aligned} 0 &= \int_0^1 \frac{1}{\lambda_0^i} \frac{dV_0^i}{d\tau} \, di \\ 0 &= \int_0^1 \frac{1}{\lambda_0^i} \beta \, \mathbb{E}_0 \left[u_{c,1}^i \frac{dc_1^i}{d\tau} + u_{\ell,1}^i \frac{d\ell_1^i}{d\tau} \right] di \\ 0 &= \int_0^1 \frac{1}{\lambda_0^i} \beta \, \mathbb{E}_0 \, u_{c,1}^i \left[\frac{dc_1^i}{d\tau} - (1-\tau) z_1^i \frac{d\ell_1^i}{d\tau} \right] di \end{aligned}$$

where the second line uses 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

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

where sometimes for convenience we use the shorthand notation

$$\omega^i_1 = \frac{\beta(\mathbb{E}_0 u^i_{c,1})}{\lambda^i_0} \qquad \text{and} \qquad \tilde{\omega}^i_1 = \frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}$$

where $\tilde{\omega}^i_1$ is a random variable from a date-0 perspective, whereas ω^i_1 is not.

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

$$\begin{aligned}
0 &= \mathbb{E}_0 \int_0^1 \frac{\beta(\mathbb{E}_0 u^i_{c,1})}{\lambda^i_0}\frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}\left[\frac{dc^i_1}{d\tau} - (1-\tau)z^i_1\frac{d\ell^i_1}{d\tau}\right] di \\
&= \left(\int_0^1 \frac{\beta(\mathbb{E}_0 u^i_{c,1})}{\lambda^i_0}\, di\right)\left(\int_0^1 \frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}\left[\frac{dc^i_1}{d\tau} - (1-\tau)z^i_1\frac{d\ell^i_1}{d\tau}\right] di\right) \\
&\quad + \mathbb{C}ov_i\left(\frac{\beta(\mathbb{E}_0 u^i_{c,1})}{\lambda^i_0}, \frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}\left[\frac{dc^i_1}{d\tau} - (1-\tau)z^i_1\frac{d\ell^i_1}{d\tau}\right]\right)
\end{aligned}$$

Notice that $\frac{\beta(\mathbb{E}_0 u^i_{c,1})}{\lambda^i_0}$ 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

$$0 = \bar{\omega}_1 \int_0^1 \frac{u^i_{c,1}}{(\mathbb{E}_0 u^i_{c,1})}\left[\frac{dc^i_1}{d\tau} - (1-\tau)z^i_1\frac{d\ell^i_1}{d\tau}\right] di$$

where we set

$$\bar{\omega}_1 = \frac{\beta(\mathbb{E}_0 u^i_{c,1})}{\lambda^i_0},$$

which is the same for all i.

One final cross-sectional covariance decomposition yields

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

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.

A.3 Ex Post Optimal Policy

We now characterize the optimality condition for optimal tax policy from the ex post perspective of period 1. The relevant equilibrium conditions are still given by

$$\begin{aligned} u^i_{\ell,1} &= -u^i_{c,1}(1-\tau) z^i_1 \\ c^i_1 &= (1-\tau) z^i_1 \ell^i_1 + T \\ Y_1 &= L_1 \\ T &= \tau \int_0^1 z^i_1 \ell^i_1\, di \\ Y_1 &= \int_0^1 c^i_1\, di \\ L_1 &= \int_0^1 z^i_1 \ell^i_1\, di. \end{aligned}$$

But now the welfare decomposition works differently. We have

$$\begin{aligned}
0 &= \frac{dW_1}{d\tau} \\
0 &= \frac{1}{\int_0^1 \alpha_1^i \lambda_1^i \, di} \frac{dW_1}{d\tau} \\
0 &= \int_0^1 \frac{1}{\lambda_1^i} \frac{dV_1^i}{d\tau} \, di + \mathbb{C}ov_i \left(\omega_1^i, \, \frac{1}{\lambda_1^i} \frac{dV_1^i}{d\tau} \right),
\end{aligned}$$

where we define

$$\omega_1^i = \frac{\alpha_1^i \lambda_1^i}{\int_0^1 \alpha_1^i \lambda_1^i \, di}$$

from the perspective of date 1. Also notice that the only sensible numeraire in this case is period 1 consumption, so we have

$$\lambda_1^i = u_{c,1}^i.$$

Let's now unpack the efficiency term. We have

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

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

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

This leaves us with

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

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

A.4 Numerics

We start with the competitive equilibrium conditions

$$\begin{aligned}
u_{\ell,1}^i &= -u_{c,1}^i (1-\tau) z_1^i \\
c_1^i &= (1-\tau) z_1^i \ell_1^i + T \\
Y_1 &= L_1 \\
T &= \tau \int_0^1 z_1^i \ell_1^i \, di \\
Y_1 &= \int_0^1 c_1^i \, di \\
L_1 &= \int_0^1 z_1^i \ell_1^i \, di.
\end{aligned}$$

Notice that the only thing that changes between the ex ante and ex post assessments are the valuations, not the allocation. So we only have to work this out once, and then compute valuations from the two perspectives.

We start by solving out for Y_1 and T, yielding

$$\begin{aligned} u^i_{\ell,1} &= -u^i_{c,1}(1-\tau)z^i_1 \\ c^i_1 &= (1-\tau)z^i_1\ell^i_1 + \tau L_1 \\ L_1 &= \int_0^1 c^i_1\, di \\ L_1 &= \int_0^1 z^i_1\ell^i_1\, di. \end{aligned}$$

Now we drop one of the market clearing conditions by Walras' law. Say the goods market clearing condition, leaving us with

$$\begin{aligned} u^i_{\ell,1} &= -u^i_{c,1}(1-\tau)z^i_1 \\ c^i_1 &= (1-\tau)z^i_1\ell^i_1 + \tau L_1 \\ L_1 &= \int_0^1 z^i_1\ell^i_1\, di. \end{aligned}$$

So taking as given τ, we have 3 equations in the 3 unknowns L_1 and $\{c^i_1, \ell^i_1\}$.

Algorithm to solve in levels. Fix τ. Guess L_1. Then use a non-linear solver to solve the first two conditions for c^i_1 and ℓ^i_1, then use the last equation as the gap in a Newton. Also, we now rewrite things in terms of the realized z, which yields

$$\begin{aligned} u_{\ell,1}(z) &= -u_{c,1}(z)(1-\tau)z \\ c_1(z) &= (1-\tau)z\ell_1(z) + \tau L_1 \\ L_1 &= \int z\ell_1(z)g(z)\, dz \end{aligned}$$

Now we use isoelastic preferences with

$$u(c,\ell) = \frac{1}{1-\gamma}c^{1-\gamma} - \frac{1}{1+\eta}\ell^{1+\eta}.$$

This lets us rewrite the first condition as

$$-\ell_1(z)^\eta = -c_1(z)^{-\gamma}(1-\tau)z$$

So if $\gamma = \eta = 2$, we have

$$\ell_1(z) = \frac{1}{c_1(z)}\sqrt{(1-\tau)z}$$

And using the budget constraint,

$$(1-\tau)z\ell_1(z)^2 + \tau L_1 \ell_1(z) = \sqrt{(1-\tau)z}$$

Algorithm to solve derivatives. Next, we fully differentiate to solve for $\frac{dL_1}{d\tau}$. We have

$$\begin{aligned} u_{\ell\ell,1}(z)\frac{d\ell_1(z)}{d\tau} &= u_{c,1}(z)z - (1-\tau)zu_{cc,1}(z)\frac{dc_1(z)}{d\tau} \\ \frac{dc_1(z)}{d\tau} &= -z\ell_1(z) + (1-\tau)z\frac{d\ell_1(z)}{d\tau} + \tau\frac{dL_1}{d\tau} + L_1 \\ \frac{dL_1}{d\tau} &= \int z\frac{d\ell_1(z)}{d\tau}g(z)\,dz \end{aligned}$$

We can rewrite the first equation as

$$\begin{aligned} \frac{u_{\ell\ell,1}(z)}{u_{\ell,1}(z)}\frac{u_{\ell,1}(z)}{u_{c,1}(z)}\frac{d\ell_1(z)}{d\tau} &= z - (1-\tau)z\frac{u_{cc,1}(z)}{u_{c,1}(z)}\frac{dc_1(z)}{d\tau} \\ -\eta\frac{1}{\ell_1(z)}(1-\tau)z\frac{d\ell_1(z)}{d\tau} &= z + (1-\tau)z\gamma\frac{1}{c_1(z)}\frac{dc_1(z)}{d\tau} \\ -\eta\frac{1}{\ell_1(z)}\frac{d\ell_1(z)}{d\tau} &= \frac{1}{1-\tau} + \gamma\frac{1}{c_1(z)}\frac{dc_1(z)}{d\tau} \end{aligned}$$

So since this is linear, we can actually solve for it explicitly. We get

$$\frac{d\ell_1(z)}{d\tau} = -\frac{1}{\eta}\ell_1(z)\left[\frac{1}{1-\tau} + \gamma\frac{1}{c_1(z)}\frac{dc_1(z)}{d\tau}\right]$$

and therefore

$$\frac{dc_1(z)}{d\tau} = -z\ell_1(z) - \frac{1}{\eta}\ell_1(z)\left[\frac{1}{1-\tau} + \gamma\frac{1}{c_1(z)}\frac{dc_1(z)}{d\tau}\right](1-\tau)z + \tau\frac{dL_1}{d\tau} + L_1$$

or simply

$$\begin{aligned} \frac{dc_1(z)}{d\tau} &= -z\ell_1(z) - \frac{1}{\eta}\ell_1(z)z - \frac{\gamma}{\eta}\ell_1(z)\frac{1}{c_1(z)}\frac{dc_1(z)}{d\tau}(1-\tau)z + \tau\frac{dL_1}{d\tau} + L_1 \\ \frac{dc_1(z)}{d\tau} &= -\frac{1+\eta}{\eta}z\ell_1(z) - \frac{\gamma}{\eta}(1-\tau)z\frac{\ell_1(z)}{c_1(z)}\frac{dc_1(z)}{d\tau} + \tau\frac{dL_1}{d\tau} + L_1 \end{aligned}$$

and therefore

$$\frac{dc_1(z)}{d\tau} = \frac{1}{1+\frac{\gamma}{\eta}(1-\tau)z\frac{\ell_1(z)}{c_1(z)}}\left[-\frac{1+\eta}{\eta}z\ell_1(z) + \tau\frac{dL_1}{d\tau} + L_1\right]$$

This gives us

$$\frac{d\ell_1(z)}{d\tau} = -\frac{1}{\eta}\ell_1(z)\frac{1}{1-\tau} - \frac{\gamma}{\eta}\frac{\ell_1(z)}{c_1(z)}\frac{1}{1+\frac{\gamma}{\eta}(1-\tau)z\frac{\ell_1(z)}{c_1(z)}}\left[-\frac{1+\eta}{\eta}z\ell_1(z) + \tau\frac{dL_1}{d\tau} + L_1\right]$$

or simply

$$\frac{d\ell_1(z)}{d\tau} = -\frac{1}{\eta}\ell_1(z)\frac{1}{1-\tau} - \frac{1}{(1-\tau)z + \frac{\eta}{\gamma}\frac{c_1(z)}{\ell_1(z)}}\left[-\frac{1+\eta}{\eta}z\ell_1(z) + \tau\frac{dL_1}{d\tau} + L_1\right]$$

And so finally we arrive at

$$\begin{aligned}
\frac{dL_1}{d\tau} &= \int z\frac{d\ell_1(z)}{d\tau}g(z)\,dz \\
&= \int z\left\{ -\frac{1}{\eta}\ell_1(z)\frac{1}{1-\tau} - \frac{1}{(1-\tau)z + \frac{\eta}{\gamma}\frac{c_1(z)}{\ell_1(z)}}\left[-\frac{1+\eta}{\eta}z\ell_1(z) + \tau\frac{dL_1}{d\tau} + L_1\right]\right\}g(z)\,dz \\
&= \int z\left\{ -\frac{1}{\eta}\ell_1(z)\frac{1}{1-\tau}\right\}g(z)\,dz - \int z\left\{ \frac{1}{(1-\tau)z + \frac{\eta}{\gamma}\frac{c_1(z)}{\ell_1(z)}}\left[-\frac{1+\eta}{\eta}z\ell_1(z) + \tau\frac{dL_1}{d\tau} + L_1\right]\right\}g(z)\,dz
\end{aligned}$$

where the first term becomes

$$\int z\left\{ -\frac{1}{\eta}\ell_1(z)\frac{1}{1-\tau}\right\}g(z)\,dz = -\frac{1}{\eta}\frac{1}{1-\tau}L_1$$

And the second term becomes

$$+\frac{1+\eta}{\eta}\int\left\{\frac{1}{(1-\tau) + \frac{\eta}{\gamma}\frac{c_1(z)}{z\ell_1(z)}}z\ell_1(z)\right\}g(z)\,dz - \left[\tau\frac{dL_1}{d\tau} + L_1\right]\int\frac{1}{(1-\tau) + \frac{\eta}{\gamma}\frac{c_1(z)}{z\ell_1(z)}}g(z)\,dz$$

Now we denote

$$x(z) = \frac{1}{(1-\tau) + \frac{\eta}{\gamma}\frac{c_1(z)}{z\ell_1(z)}}$$

and can therefore write

$$\frac{dL_1}{d\tau} = -\frac{1}{\eta}\frac{1}{1-\tau}L_1 + \frac{1+\eta}{\eta}\int x(z)z\ell_1(z)g(z)\,dz - \tau\frac{dL_1}{d\tau}\int x(z)g(z)\,dz - L_1\int x(z)g(z)\,dz$$

And so therefore

$$\frac{dL_1}{d\tau} = \frac{1}{1+\tau X}\left(-\frac{1}{\eta}\frac{1}{1-\tau}L_1 + \frac{1+\eta}{\eta}\int x(z)z\ell_1(z)g(z)\,dz - L_1\int x(z)g(z)\,dz\right)$$

where $X = \int x(z)g(z)\,dz$.

B Proofs

B.1 Proof of Proposition 1

Proof. The result is a special case of Proposition 2. See Appendix B.2 below for the proof. 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

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

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 Propposition 1 is satisfied. □

B.2 Proof of Proposition 2

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.

Interpersonal welfare comparisons. Notice that we can write the ex ante assessment as

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

Kydland-Prescott. We now split the sum $\sum_t \sum_{s^t}$ of the first term. We can write

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

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.

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

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

or simply

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

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). □

B.3 Proof of Proposition 3

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 1, 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 1, 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. □

B.4 Proof of Proposition 4

Proof. We start with a proof of condition (i). 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.

Now we prove condition (ii) 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. □

B.5 Proof of Theorem 1

Proof. Our proof of Proposition 3 presents a systematic characterization of the sources of time inconsistency. It is sufficient to prove Theorem 1 under Assumption 1, which rules out time consistency problems á la Kydland and Prescott (1977). Under 1, time inconsistency is therefore governed by the value of

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

Property P1 of Theorem 1 restricts us 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 22 for a discussion of why this assumption is without loss since Theorem 1 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. □

B.6 Proof of Proposition 6

Proof. Recall that the history s^k aggregate efficiency assessment is defined as

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

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

$$\Xi^{AE}_{s^k} = \bar{\omega}_{T|s^k}\, \bar{\omega}_{T|s^k}(s^T) \sum_i \frac{\partial \mathcal{C}^i_T(s^T \mid s^k)}{\partial \theta}$$

When Assumption 1 is satisfied, then

$$\frac{\partial \mathcal{C}^i_T(s^T \mid s^k)}{\partial \theta} = \frac{\partial \mathcal{C}^i_T(s^T \mid s^t)}{\partial \theta} = \frac{dc^i_T(s^T)}{d\theta}$$

for any s^t and s^k. Therefore, under 1, 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 $\bar{\omega}_{T|s^k}$ and $\bar{\omega}_{T|s^k}(s^T)$. But since $\lambda^i_{s^k} > 0$ and $u'(c^i_t(s^t)) > 0$ for all t, s^t and i, these two weights must also be strictly positive. Therefore, we have

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

This tells us that "static" aggregate efficiency assessments are always time consistent under Assumption 1, concluding our proof. □

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*

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

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

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

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.

This concludes our description of the environment.

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

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

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

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

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

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

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

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

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

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

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

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

or simply

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

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.

Contents

1 Introduction **2**

2 A Motivating Example **9**

2.1 Environment 9

2.2 Optimal Policy 10

2.3 Lessons for Time Inconsistency 11

3 Environment **14**

4 Time Inconsistency **16**

4.1 Inconsistency of Welfare Assessments 16

4.2 Inconsistency of Efficiency Assessments 20

4.2.1 Efficiency 20

4.2.2 The Inconsistency of Efficiency Assessments 23

4.2.3 Impossibility Result 27

4.3 Inconsistency of The Origins of Efficiency Gains 29

5 Applications **30**

5.1 Application 1: Anticipated Relief Policies 31

5.1.1 Environment 32

5.1.2 Weights, Assessments, and Time Inconsistency 33

5.2 Application 2: Risk-Sharing Policies 37

5.2.1 Environment 39

5.2.2 Weights, Assessments, and Time Inconsistency 40

5.3 Application 3: Aggregate Investment Polices 44

5.3.1 Environment 45

5.3.2 Weights, Assessments, and Time Inconsistency 46

6 Conclusion **50**

A Proofs and Derivations: Section **OA-1**

A.1 Implementability Conditions OA-1

A.2 Ex Ante Optimal Policy OA-2

A.3 Ex Post Optimal Policy OA-4
A.4 Numerics OA-6

B Proofs **OA-10**
B.1 Proof of Proposition 1 . . . . . OA-10
B.2 Proof of Proposition 2 . . . . . OA-10
B.3 Proof of Proposition 3 . . . . . OA-13
B.4 Proof of Proposition 4 . . . . . OA-14
B.5 Proof of Theorem 1 . . . . . OA-15
B.6 Proof of Proposition 6 . . . . . OA-15

C Microfoundation of the Consumption Function **OA-17**