Time-Consistent Individuals, Time-Inconsistent Households

Andrew Hertzberg
Federal Reserve Bank of Philadelphia

WP 26-20

PUBLISHED
April 2026

ISSN: 1962-5361

DOI: https://doi.org/10.21799/frbp.wp.2026.20

Time-Consistent Individuals, Time-Inconsistent Households

ANDREW HERTZBERG*

ABSTRACT

I present a model of consumption and savings for a multi-person household in which members are imperfectly altruistic, derive utility from both private and shared public goods, and share wealth. I show that, despite having standard exponential time preferences, the household is time-inconsistent: members save too little and overspend on private consumption goods. The household remains time-inconsistent even when members save separately, because the possibility of voluntary transfers or joint contribution to the public good preserves the dynamic commons problem. The household will choose to share wealth when the risk sharing benefits outweigh the utility cost of overconsumption. JEL Codes: D13, H31, D91

*The Federal Reserve Bank of Philadelphia; E-mail: andrew.hertzberg@phil.frb.org. I thank Patrick Bolton, Pierre-Andre Chiappori, Daniel Gottlieb, Ulrike Malmendier, Daniel Paravisini, Bryan Routledge, Tano Santos, Daniel Wolfenzon, and seminar participants at Anderson School of Management, Columbia Business School, Carlson School of Management, Copenhagen Business School, Dartmouth College, Fuqua School of Business, Rady School of Management, and Stockholm School of Economics, as well as conference participants at the American Economic Association Meetings, Corporate Finance NBER Fall Session, Household Finance NBER Summer Institute, Western Finance Association Conference, and World Finance Conference. I thank Anna Benoit, Cassandra Cole, Marc Folch, and Tara Philip for outstanding research assistance. All remaining errors are my own. I have read the *Journal of Finance's* disclosure policy and have no conflicts of interest to disclose.

Standard theory of household financial decision-making is based upon the behavior of a single maximizing agent. However, a considerable body of empirical work strongly rejects the unitary model as an adequate description of household decision-making (see, for example, Lundberg et al. (1997), Browning and Chiappori (1998), Phipps and Burton (1998), Ashraf (2009), and Cesarini et al. (2017)). The key lesson from this literature is that household members have distinct preferences and are self-interested, caring more about the utility from their own consumption than their partners' consumption. These empirical findings have led to considerable theoretical work that reconsiders the static theory of household demand for goods and labor supply (see, for example, Chiappori (1988), Browning et al. (1994), Browning and Chiappori (1998), and Chiappori et al. (2002)). Little attention has been directed to the question of how financial decision-making will be affected by the presence of distinct and imperfectly aligned preferences within a multi-person household.[1] This paper provides a new framework for addressing this question. In particular, I explore whether multi-person households are time-consistent or whether they are inherently unable to carry out optimal consumption and savings plans.

I propose a model of the household that comprises two members who are connected in three ways. First, wealth is shared. Second, members derive utility both from private consumption and shared non-rival public consumption goods such as children and housing. Third, following the evidence cited above, household members are altruistic but remain self-interested. That is, member A cares more about the utility from his consumption than B cares about A's consumption and vice versa. I further assume that both members have the same exponential time

[1]Recent exceptions include Mazzocco (2004), Mazzocco (2007), and Schaner (2015).

preferences and therefore agree on the ex ante optimal path of total household consumption and savings. This ensures that any time-inconsistency is due solely to the strategic interaction between household members.

I characterize the household's equilibrium consumption path as a subgame perfect Nash equilibrium in consumption choices. This is the equilibrium that obtains when household members are unable to commit ex ante to their future consumption choices.[2] I show that the household is time-inconsitent: members systematically overconsume relative to the agreed optimal ex ante consumption and savings plan. Time-inconsistency arises because both members unilaterally deviate from the optimal plan and increase their own private consumption at the expense of shared future savings. The central insight is that in a multi-person household, shared savings is subject to a dynamic commons problem and is therefore under-provided in equilibrium.[3] To measure the cost of the household's time-inconsistency problem, I solve for the fraction of total wealth that the household would be willing to pay ex ante to achieve the full commitment consumption path. I show this is monotonically increasing in the degree of self-interest.

Having shown that the household overconsumes when savings are shared in a joint account, I next study household behavior when members save in separate accounts that the other cannot access. If household members consume and save in complete isolation, they will do so optimally. However, even when unable to directly access the other's savings, household members remain linked by their altruism and consumption of a shared public good. As a result, if the wealth difference

[2] Technically this requires that members are unable to enforce contracts between themselves written contingent on their consumption choices.

[3] Dynamic commons problems have been used to study national underinvestment (Lancaster (1973), Tornell and Velasco (1992)), overexploitation of natural resources (Levhari and Mirman (1980)), and sovereign debt (Amador (2008)).

between members is sufficiently large, the member with relatively more wealth will voluntarily transfer savings to the other member, while, if the wealth of both members is sufficiently close, both will contribute to the public good. Anticipating either possibility occurring in the future creates the same time-inconsistency that occurs when wealth is shared. This occurs because both members recognize that an additional dollar of personal saving today may offset a future contribution from the other member to the purchase of the public good or lower the amount the other member may voluntarily transfer to them in the future. This result suggests that the phenomenon of time-inconsistency can be applied to extended families (for example, the interaction between adult siblings) in settings where wealth is held separately but where the possibility of future transfers or joint contribution to a public good may arise. I also show that with separate savings, the household's marginal propensity to consume depends on both the distribution of wealth within the household and which household member receives the additional dollar.

Finally, I examine the conditions under which household members will choose to save in separate individual accounts rather than share wealth. I show that the household will adopt joint accounts when its exposure to relative wealth shocks is sufficiently large and the initial wealth of each member is sufficiently close. The intuition is as follows: a larger standard deviation of intra-household relative wealth shocks increases the risk pooling benefit of a shared account and, by increasing the likelihood of future financial dependence within the household, reduces the disciplining effect that separate accounts has on the undersavings problem.

This paper makes several contributions to the literature on household savings behavior. The first is to show that the intra-household pattern of ownership and control of assets can impact consumption and savings decisions. This contrasts with the unitary model of the household in which only the combined household balance sheet matters, and which is hence unable to rationalize the choice between either arrangement. Survey evidence suggests that both separate and shared savings are popular in developed countries.[4] Consistent with the results in this paper, Ashraf (2009) and Schaner (2015) show that household savings decisions vary on whether assets are shared or held separately by household members.

By showing that time-inconsistency arises naturally in a multi-person household, the framework I propose rationalizes the use of commitment technologies that limit the ability of individual members to unilaterally deviate from jointly agreed consumption and savings plans. As an example, in the US, the 1984 Retirement Equity Act mandates that all retirement plans covered by the 1974 Employee Retirement Income Security Act (which include all defined benefit plans, Individual Retirement Accounts, and all 401(k) plans) require joint approval by both spouses before funds can be withdrawn or loans can be taken against such savings. Aura (2005) shows that the introduction of this law increased household saving. In the context of developing economies, savings commitment technologies such as rotating savings and credit associations are motivated by the ability of one spouse to limit the ability of a partner to overconsume out of shared wealth (Anderson and Baland (2002), Collins et al. (2009)).

This paper contributes to a large theoretical literature that studies household decision-making when members have misaligned preferences (see Lundberg and

[4]The 2002 General Social Survey (Smith et al. (2011)) finds that 53% of all married households in the United States share all financial wealth. Similar survey evidence for the United States and the United Kingdom is presented by Treas (1993) and Vogler et al. (2006). A 2006 survey of Japanese wives found that 50% held secret savings (referred to in Japan as "hesokuri") (see Alexy (2007)). Evidence from developing countries indicates that household risk sharing is limited and hence implies that some assets are effectively owned separately (Robinson (2012), Duflo and Udry (2004)).

Pollak (2007), Browning et al. (2006), and Donni and Chiappori (2011) for comprehensive surveys). In many of these papers, static household decision-making is often modeled as the outcome of an efficient bargaining process, and the focus is on studying the determinants of threat points and bargaining weights of each household member.[5] Evidence on whether households are able to achieve Pareto-efficient allocations is mixed. Several papers find intra-household consumption and labor-supply patterns that are consistent with static efficiency (see, for example, Browning and Chiappori (1998), Chiappori et al. (2002), and Bobonis (2009)) while others document choices that are clear breaches of efficiency (Castilla (2015), Almås et al. (2018), Conlon et al. (2021), Afzal et al. (2022), and Choukhmane et al. (2022)).[6] Tests for dynamic efficiency are generally rejected (see, for example, Udry (1996), Duflo and Udry (2004), Mazzocco (2007), De Mel et al. (2009), Robinson (2012), and Lise and Yamada (2019)). Several recent papers argue that information asymmetries between members may be the underlying cause of inefficiencies at the household level (see, for, example Bloch and Rao (2002), Ashraf et al. (2014), Chen (2013), Ambler (2015), Jakiela and Ozier (2016), Apedo-Amah et al. (2020), Ashraf et al. (2020), and Baseler (2023)). Motivated by this evidence, I study the equilibrium that obtains when commitment is not possible, as the focus of this paper is intertemporal household decision-making.[7]

Finally, this paper is related to the literature that studies individuals with

[5]As an exception non-cooperative decision-making within the household has been considered by Lundberg and Pollak (1993), Chen and Woolley (2001), and Lundberg and Pollak (2003).

[6]Dauphin et al. (2018) argues that many tests of household efficiency are under powered and uses a meta analysis to show that additional implications of efficiency are often rejected in the same settings.

[7]Lundberg and Pollak (2003) argue that Pareto-efficient bargaining may break down in a dynamic context because current decisions may affect future bargaining power.

dynamically inconsistent time preferences.[8] The results of this paper highlight that self-interest within multi-person households can produce time-inconsistency, and rationalize commitment technologies and behavior that cannot be explained by the literature focusing on individual time preferences or self-control (see, for example, Thaler and Benartzi (2004), Ashraf et al. (2006), Beshears et al. (2011)).

The paper proceeds as follows. Section I outlines the base model of household consumption. Section II characterizes the equilibrium consumption choices of the household and shows that they are time-inconsistent. Section III studies household decision-making when members are able to save in separate accounts. Section IV studies the household choice to save in separate or joint accounts. Section V discusses empirical implications of the model, including strategies that the household might adopt to mitigate time-inconsistency. All derivations are provided in the Internet Appendix.[9]

I. Model of Household Consumption

The household has two members, indexed by i, labeled A and B. Time is discrete and indexed by t. Both household members consume from $t = 1$ until the end of their lives at T. I assume that the household remains together for their entire lives with certainty and abstract from endogenous household formation.

[8]For theory see Thaler and Shefrin (1981), Laibson (1997), Harris and Laibson (2001), and Laibson et al. (2003), and for evidence see Ainslie (2001), Frederick et al. (2002), and Shapiro (2005).

[9]The Internet Appendix is available in the online version of the article on the *Journal of Finance* website.

A. *Preferences*

In each period, $t \geq 1$ member i derives utility from two goods. The first is a private consumption good denoted by $C_{i,t}$. The second is a nonrival public consumption good that both members share, denoted by H_t. The public consumption good captures one defining characteristic of being in a household: members share consumption such as housing, appliances, children, and care for extended family members. The public and private consumption goods are aggregated into a composite consumption good using a constant elasticity of substitution aggregator so that the composite consumption of member i in period t is

$$g_{i,t} = \left[(1 - \mu_H) \, C_{i,t}^{\frac{\varepsilon - 1}{\varepsilon}} + \mu_H H_t^{\frac{\varepsilon - 1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon - 1}}, \tag{1}$$

where $\mu_H \in [0, 1]$ is the relative utility weight that members place on public consumption. The elasticity of substitution between the two goods is measured by $\varepsilon \geq 0$. The total level of household public consumption is the sum of the amount purchased by both members in each period: $H_t = H_{A,t} + H_{B,t}$.

The second defining characteristic of the household is that its members are altruistic. I capture this by supposing that the utility of each member is determined by the amount of the composite good that both members consume. Each member places weight $\frac{1+\Delta}{2}$ on their own consumption and weight $\frac{1-\Delta}{2}$ on the consumption of the other member. I focus on the case in which the altruism between household members is imperfect in the sense that each member cares more about their own consumption than that of their partner: $\Delta \in [0, 1]$. In words, Δ measures the degree of self-interest within the household.[10] The utility derived by each member

[10] The framework can also be used to study the case in which members care more about each

in a given period is determined by the amount of the composite good that each member consumes according to

$$q_{i,t} = \left[\left(\frac{1+\Delta}{2} \right) g_{i,t}^{\frac{\xi-1}{\xi}} + \left(\frac{1-\Delta}{2} \right) g_{j,t}^{\frac{\xi-1}{\xi}} \right]^{\frac{\xi}{\xi-1}} \tag{2}$$

where $\xi \geq 0$ is the elasticity of substitution between each member's consumption.

Members have Epstein-Zin (Epstein and Zin (1989, 1991)) recursive preferences over the expected stream of future household consumption. The objective of member i at t is

$$V_{i,t} = \left[(q_{i,t})^{1-\frac{1}{z}} + \delta \left(E_t V_{i,t+1}^{1-\gamma} \right)^{\frac{1-\frac{1}{z}}{1-\gamma}} \right]^{\frac{1}{1-\frac{1}{z}}} \tag{3}$$

where $z \geq 0$ captures the elasticity of intertemporal substitution and $\gamma \geq 0$ captures the degree of risk aversion. When $\gamma = z$, these preferences reduce to the standard case of state-separable constant relative risk aversion utility. Both household members discount utility from future consumption using the exponential discount factor δ.[11] Note that these are standard time preferences, so if the members were to act in isolation, the optimal consumption plan for each household member would be time-consistent.

---

other than themselves ($\Delta < 0$). Since the evidence on household consumption decisions suggests that this is generally not the case, I do not focus on this scenario.

[11] Adams et al. (2014) and Jackson and Yariv (2015) show that heterogeneous discount rates within the household can also produce time inconsistency. For related evidence, see Ringdal and Sjursen (2021).

B. *Household Budget Constraint*

The household starts life at $t = 1$ with wealth W_1, which is taken as given, and has no income.[12] The third defining characteristic of the household is that all wealth is shared so that both household members have full access to the remaining combined wealth in each period. This assumption is made to align the framework with the way the budget constraint is treated in any standard unitary model of intertemporal decision-making. Moreover, the 2002 General Social Survey (Smith et al. (2011)) finds that 53% of all married households in the United States share all financial wealth, suggesting that this is the most empirically relevant characterization of the household budget constraint. I study household behavior when household members save separately in Section III. The private consumption good serves as the numeraire, and the relative price of the public good is p_H. Any wealth not consumed by the household is saved between periods at a gross risk-free interest rate of R. Household financial wealth evolves according to

$$W_{t+1} = R\left(W_t - X_t\right), \tag{4}$$

where

$$X_t = C_{A,t} + C_{B,t} + p_H\left(H_{A,t} + H_{B,t}\right) \tag{5}$$

is total household expenditure in period t.

[12] This is identical to assuming that household labor supply decisions are fixed and the household is free to borrow and lend at R each period. In this case, W_t is the present value of all future income plus (minus) any savings (debt) that the household has in period t.

C. Decision-Making

Household members cannot commit to a consumption path. As a result, household members are unable to enforce mutually agreed levels of consumption, either in the present or the future. Household members noncooperatively and simultaneously decide how much of the household wealth W_t to spend on their own private consumption $C_{i,t} \geq 0$ and on their contribution to public consumption $H_{i,t} \geq 0$ in each period. The dynamic equilibrium consumption path will be the Nash subgame-perfect solution to the consumption game between these two members. Let a single "*" denote the noncooperative equilibrium quantities $C^*_{i,t}$, $D^*_{i,t}$, and $H^*_{i,t}$.

Since both members make consumption decisions simultaneously, it is possible that both members could attempt to spend more than total household wealth. To avoid this problem I assume that both members are able to consume, at most, half the total household wealth in any single period[13]:

$$C_{i,t} + p_H H_{i,t} \leq \frac{W_t}{2}. \tag{6}$$

By imposing (6) I ensure $X_t \leq W_t$, and hence have a well-defined budget constraint for each household member's consumption problem each period. In the Internet Appendix, I show that (6) does not bind in any period $t < T$.[14]

[13]If the model were extended to allow household income, then this constraint would impose the condition that credit markets will not allow the household to raise debt in excess of the present value of all future income.

[14]In Section II of the Internet Appendix, I solve an otherwise-identical model in which household members make consecutive consumption decisions. In that setting (6) is replaced with the standard budget constraint $C_{i,t} + p_D D_{i,t} + p_H H_{i,t} \leq W_t$. I show that the equilibrium studied here is arbitrarily close to the unique equilibrium from the consecutive-move model as the model approaches continuous time, which demonstrates that (6) does not drive the results below.

D. *Full Commitment Problem and the Value of Commitment*

To evaluate the optimality of the non-cooperative equilibrium consumption path, I compare it to the consumption path that would obtain if the household were able to fully commit to consumption choices. Consider the problem that the household would face in setting a full-commitment path. Whenever $\Delta > 0$, household members disagree over the optimal allocation. However, any allocation that they would choose must be Pareto-optimal, and hence I characterize the solution to the following full-commitment Pareto problem:

$$\max_{\left\{\{C_{i,t},D_{i,t},H_t\}_{t=1}^{t=T}\right\}_{i\in\{A,B\}}} \Pi = \eta V_{A,1} + (1-\eta) V_{B,1} \tag{7}$$

$$\text{subject to } W_1 - \sum_{x=1}^{T} R^{-(x-1)} \left[C_{A,x} + C_{B,x} + p_H H_x\right] \geq 0 \text{ and} \tag{8}$$

$$\{C_{A,t}, C_{B,t}, H_t\}_{t=1}^{t=T} \geq 0, \tag{9}$$

where $\eta \in [0,1]$ is the Pareto weight placed on the objective of member A. Let a double "**" denote the full-commitment Pareto-optimal consumption quantities $C_{i,t}^{**}(\eta)$, and $H_t^{**}(\eta)$ that solve this problem.

To quantify the welfare loss incurred by the household's time-inconsistency, I calculate how much the household would be willing to pay at the start of $t = 1$ for a technology that allowed them to commit to an optimal consumption path. Let $V_{i,1}^{*}(W_1)$ be the discounted lifetime utility that household member i achieves absent commitment as a function of initial household wealth. Let $V_{i,1}^{**}(W_1(1-\phi), \eta)$ be the counterpart for the case in which the household spends a fraction ϕ of its initial wealth W_1 to achieve the full commitment plan that places weight η on the preferences of member A. The value of commitment ϕ^{**} is the most the household will pay while ensuring that there exists a weight η such that the purchase is a Pareto improvement for both members. Formally, ϕ^{**} solves

$$\phi^{**} = \max_{\phi,\eta} \phi \tag{10}$$

$$\text{subject to } V_{i,1}^{**}\left(W_1\left(1-\phi\right),\eta\right) \geq V_{i,0}^{*}\left(W_1\right) \text{ for } i \in \left\{A,B\right\}, \text{ and } \eta \in \left[0,1\right]. \tag{11}$$

The solution to this problem defines the minimum wealth $W_1\left(1-\phi^{**}\right)$ a household with commitment would need to achieve an allocation that weakly Pareto-dominates the allocation without commitment. As such, it can be thought of as the multi-person household counterpart to the money metric utility function proposed by Samuelson (1974) and Varian (1990).

II. Consumption Choices and Time-Inconsistency

I start by studying a special case of the model that has a tractable analytical solution. In particular, I focus on the case in which the aggregation of goods is Cobb-Douglas ($\varepsilon = 1$) and the elasticity of intertemporal substitution is unity ($z = 1$).[15] In Section D, I examine the results of the model outside this special case. All derivations and proofs are in Section I of the Internet Appendix.

[15]When household members share wealth the equilibrium allocation will be symmetric and riskless and thus will be unaffected by ξ and γ. These parameters will play a role in Section III.

A. Example: Two Periods

To illustrate the source of time-inconsistency within the household, I sketch a solution to the two period version of the model ($T = 2$). If the household were able to commit to an optimal consumption plan at $t = 0$, it would satisfy the standard Euler equation

$$\frac{1}{X_t^{**}} = R\delta \frac{1}{X_{t+1}^{**}}. \tag{12}$$

Under this plan, total consumption expenditure would be allocated between the public and private goods each period according to their preference weights,

$$C_{i,t}^{**} = \frac{1-\mu_H}{2} X_t^{**} \text{ and } p_H H_t^{**} = \mu_H X_t^{**}, \tag{13}$$

where I have assumed without loss of generality a symmetric Pareto weight in the planning problem ($\eta = 0.5$). I show that this plan is time-inconsistent by illustrating that both members have an incentive to unilaterally deviate from this consumption allocation at $t = 1$. To do so, note that when z= 1 (log utility), the value function that each member will face in the second period can be summarized as

$$V_{i,2} = \ln X_2 + \text{Constant}, \tag{14}$$

where the second term is a constant that is invariant to X_2 and hence irrelevant for the problem at $t = 1$. Using (14), consider a deviation by member i from the optimal consumption path at $t = 1$ whereby she increases her private consumption above $C_{i,1}^{**}$ by a small amount $\epsilon_C > 0$, recognizing that this will reduce combined household savings and hence consumption at $t = 2$ by $R\epsilon_C$. The marginal benefit from such a deviation is[16]

$$\frac{\partial}{\partial \epsilon_C}\left\{\frac{1+\Delta}{2}(1-\mu_H)\ln\left(C_{i,1}^{**}+\epsilon_C\right)+\delta\ln\left(X_2^{**}-R\epsilon_C\right)\right\}\Big|_{\epsilon_C=0}=\underbrace{\frac{1}{X_1^{**}}}_{\text{[MU at }t=1]}-\underbrace{\left(\frac{R}{1+\Delta}\right)}_{\text{[PRTSS]}}\delta\underbrace{\frac{1}{X_2^{**}}}_{\text{[MU at }t=2]}\geq 0 \tag{15}$$

where the first equality follows using (13). Comparing (15) to the Euler equation that governs the optimal allocation (12) yields that member i strictly prefers to make such a deviation from the optimal consumption plan whenever altruism is imperfect: $\Delta > 0$. The incentive to deviate stems from the difference between the social return to saving, R, and each member's private return to shared saving, $\frac{R}{1+\Delta}$ (the term denoted $PRTSS$). When considering the trade-off between consumption and saving, each member recognizes that every dollar saved will be shared and, since altruism is imperfect, they will not internalize the full combined household benefit of those shared savings. Thus, a trade-off between the benefit of *private* consumption and the return to *shared* saving produces an incentive for each member to deviate from the ex-ante optimal plan. As the degree of self-interest within the household rises, as measured by Δ, the wedge between the social and private return to saving increases.

To highlight the importance of private consumption in generating time-inconsistency, consider a second possible deviation from the optimal household allocation given by (12). Suppose that member i contemplates a deviation whereby she increases public consumption above H_1^{**} by a small amount $\frac{\epsilon_H}{p_H} > 0$, recognizing that this will reduce combined household savings and hence consumption at $t = 2$ by $R\epsilon_H$.

[16] I omit all terms unaffected by the deviation from i's objective: $\frac{1-\Delta}{2}\mu\ln C_{j,1}^{**}+(1-\mu)\ln H_1^{**}$.

The marginal benefit from such a deviation is[17]

$$\frac{\partial}{\partial \epsilon_H} \left\{ \mu_H \ln \left(H_1^{**} + \frac{\epsilon_H}{p_H} \right) + \delta \ln \left(X_2^{**} - R\epsilon_H \right) \right\} \bigg|_{\epsilon_H = 0} = \frac{1}{X_1^{**}} - R\delta \frac{1}{X_2^{**}} = 0, \quad (16)$$

where, as before, the first equality follows using (13). In this case, there is no intertemporal distortion. Each member trades off the *shared* benefit of consuming the public good today and the *shared* benefit of saving. Since the same concern for the combined household is present in both terms, there is no distortion to the relative intertemporal trade-off.

B. *Multi-Period Solution*

I now characterize household time-inconsistency by characterizing the household consumption path for any T. With full commitment, the optimal consumption path that the household would choose at $t = 0$ would specify the following level of total expenditure each period

$$X_t^{**} = \frac{1}{1 + \sum_{x=1}^{T-t} \delta^x} W_t. \quad (17)$$

The allocation of X_t^{**} remains according to (13).[18]

Without commitment, the equilibrium levels of private and public consumption expenditures by member i in period $t < T$ is

$$C_{i,t}^{*} = \frac{\frac{1+\Delta}{2}\left(1 - \mu_H\right)}{1 + \left(1 - \mu_H\right)\Delta + \sum_{x=1}^{T-t} \delta^x} W_t \text{ and } p_H H_t^{*} = \frac{\mu_H}{1 + \left(1 - \mu_H\right)\Delta + \sum_{x=1}^{T-t} \delta^x} W_t. \quad (18)$$

[17] I omit all terms unaffected by the deviation from i's objective: $\frac{1+\Delta}{2}\mu \ln C_{i,1}^{**} + \frac{1-\Delta}{2}\mu \ln C_{j,1}^{**}$.

[18] For any pareto weight η, $C_{A,t}^{**}(\eta) = \mu\eta X_t^{**}$ and $C_{B,t}^{**}(\eta) = \mu(1-\eta) X_t^{**}$.

Total equilibrium household expenditure in period t is

$$X_t^* = \frac{1}{1 + \frac{1}{1+(1-\mu_H)\Delta} \sum_{x=1}^{T-t} \delta^x} W_t. \tag{19}$$

Comparing (17) and (19) yields the following proposition.

PROPOSITION 1: *When $(1 - \mu_H)\Delta > 0$, the household is time-inconsistent. Each period, the fraction of remaining wealth that the household consumes is strictly higher in the noncooperative equilibrium than under the full-commitment Pareto optimum.*

Proposition 1 highlights that the household will be time-inconsistent whenever altruism is imperfect and members derive utility from private consumption. When this is true, household savings is subject to a dynamic commons problem. As a result, the household will exceed the optimal level of consumption early in life and, through the intertemporal budget constraint, consume below the optimal level later in life. For example, if $T = 50$, $\delta = 0.95$, $R = \frac{1}{0.95}$, and $\Delta = \mu_H = 0.4$, the household will spend 22% above the optimal level in the first year and consume less than 51% of the level that the household would like to commit to in each of the last five years of its life.

C. *Quantifying the Inefficiency*

Having shown that the consumption allocation achieved in the noncooperative solution is inefficient, I now quantify this inefficiency. A household in which $\Delta = \mu_H = 0.4$ incurs a utility cost of time-inconsistency that is equivalent to giving up 3.54% of household wealth at $t = 1$. Panel A of Figure 1 shows that the value of commitment increases monotonically with the degree of self-interest within the family (Δ). To understand this comparative static, start by considering the scenario in which the household does not consume any public goods ($\mu_H = 0$). In this case, the suboptimality of equilibrium household consumption comes purely from an intertemporal inefficiency whereby, for a given level of household wealth at any $t < T$, the total level of household expenditure exceeds the full-commitment optimum. Comparing (19) and (17) shows that the degree of the intertemporal inefficiency, captured by $\frac{X_i^*}{X_i^{**}}$, increases monotonically with the degree of selfishness in the household (Δ). This is shown in Panel B of Figure 1. Intuitively, the dynamic commons problem is more severe when household members are less altruistic.

Holding Δ constant, increasing the role of the public good (μ_H) partially realigns intertemporal preferences and thereby mitigates the intertemporal inefficiency (this is shown in Panel B of Figure 1 and corroborated for the full range of μ_H in Panel A of Figure IA1). When the household consumes a public good, $\mu_H > 0$, an allocative inefficiency is also present whereby, for a given level of wealth at any t, the household spends a smaller fraction of their expenditures on the public good than under the full-commitment optimum: $\frac{H_t^{**}}{X_t^{**}} - \frac{H_t^*}{X_t^*} > 0$. Panels B and C of Figure IA1 show that the underprovision of the public good is monotonically increasing in selfishness (Δ) and nonmonotonic in the utility weight on the public good (μ_H).[19] Panel A of Figure 1 shows that effect of μ_H on the intertemporal inefficiency always dominates any effect on allocative inefficiency and thus the value

[19]The nonmonotonic effect of μ_H on allocative inefficiency comes from two countervailing effects: an increase in μ_H creates a larger role for public consumption (which mechanically increases the magnitude of the allocative inefficiency) while also aligning the objective of each member (which decreases the inefficiency).

Figure 1. The Value of Commitment and Intertemporal Inefficiency

This plot shows how the value of commitment (ϕ^{**}) in Panel A and the degree of intertemporal inefficiency ($\frac{X_1^*}{X_1^{**}}$) in Panel B vary with the degree of self-interest within the household (Δ). The plot is drawn using the following parameters: household members have an exponential discount factor of $\delta = 0.95$, the gross interest rate is $R = 1/0.95$, the household exists for $T = 50$ years, the elasticity of substitution between goods is unity $\varepsilon = 1$, the intertemporal elasticity of substitution is unity $z = 1$, and the relative price of the public good is one $p_H = 1$. Since the utility function is homothetic, both outcomes are invariant to the level of W_1. Both figures are drawn for the case where household members place low, intermediate and high weight on the value of public consumption ($\mu_H = 0, 0.4, 0.8$). These comparative statics are generated assuming household wealth is shared between both members.

Panel A. Value of Commitment

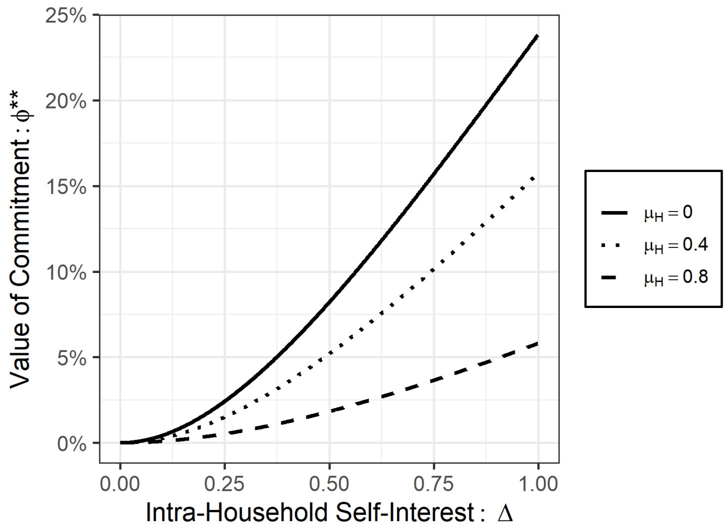

Panel B. Intertemporal Inefficiency

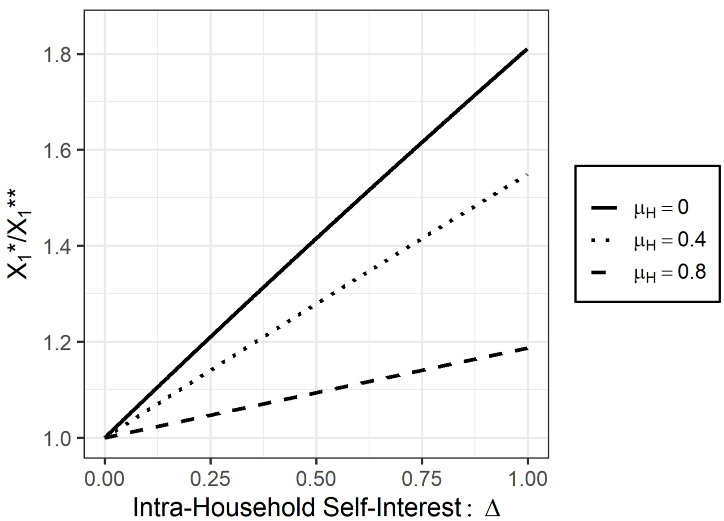

of commitment is unambiguously decreasing in the concern for public consumption (μ_H).

The comparative statics imply that two otherwise-identical households will differ in their tendency to overconsume if the households have different internal levels of altruism or place different utility weight on public consumption goods such as children. While the model does not explicitly model a change in preferences over time, these comparative statics also imply that a change in altruism, for example, prior to divorce, or increased utility weight on public consumption, for example, due to the birth of a child, will alter the tendency of the household to overconsume.

D. General Case

I now discuss the equilibrium household allocation beyond the special case of log Cobb-Douglas ($\varepsilon = z = 1$) considered so far. I demonstrate the robustness of the results described above and produce additional comparative statics.

D.1. Robustness

The central result, that household savings decisions are time-inconsistent, comes from the basic principal that, in a multi-person household, savings are effectively a public good and as such subject to a dynamic commons problem whenever members are self-interested. This principal does not rely on the assumption of log Cobb-Douglas preferences. To demonstrate this generality, in Figure IA2 I show that the value of commitment (Panels A and B) and the degree of intertemporal inefficiency (Panels C and D) are monotonically increasing in the degree of self-interest within the household, regardless of the elasticity of intertemporal substitution (z) or the elasticity of substitution between public and private consumption

(ε).

D.2. Comparative Statics in μ_H

The effect of a change in utility weight on public consumption depends on the elasticity of substitution between goods. In most cases, increasing the weight on public consumption aligns the incentives of each member and monotonically reduces the value of commitment (see Figure IA3, Panel A). However, when the two goods are near-perfect substitutes, a pronounced allocative inefficiency can exist for intermediate values of the weight on public consumption. In this case, a nonmonotonic relationship between this weight and the value of commitment arises.[20]

D.3. Comparative Statics in z

The value of commitment has an inverted U-shaped relationship with the elasticity of intertemporal substitution. Figure IA4 shows that this holds regardless of the level of household self-interest (Panel A), the elasticity of substitution between goods (Panel B), or the utility weight on public consumption (Panel C). This relationship is produced by the interplay of two opposing forces. First, when the elasticity of intertemporal substitution is zero, there is no incentive for household members to overconsume early in life, and as a result, the time-inconsistency problem does not exist. By this logic, increasing the elasticity of intertemporal substitution unambiguously increases the degree of the time-inconsistency, as is shown in Panel D of Figure IA4. In contrast, the utility cost of any given intertemporal inefficiency is reduced by lowering the elasticity of substitution. This

[20]Additional discussion of this nonmonotonicity is provided in the Internet Appendix.

countervailing force dominates only when the elasticity is sufficiently high to produce the inverted U-shaped relationship.

III. Separate Accounts

Thus far I show that imperfectly altruistic household members who save in a joint account undersave. In this section, I study household behavior when members are able to break their direct financial interdependence by saving in separate accounts that the other cannot access. For example, this would correspond to each member of a married couple holding assets such as bank accounts and retirement accounts solely in their own names. Moreover, to study extended family settings (for example, adult children interacting with each other and their parents and grandparents), it is natural to assume that each member has their own privately-held wealth that others cannot directly access. This framework also more closely resembles transnational households that remain linked by voluntary remittances (see Ambler (2015) and Ratha et al. (2016)). If there were no other connection between family members, having separate accounts would trivially eliminate the inefficiency highlighted thus far, by making each member fully autonomous. However, even with separate accounts, household members remain indirectly linked by altruism and their concern for shared public consumption. Altruism among household members may lead each member to anticipate making or receiving transfers from another. Shared concern for public consumption may lead each member to anticipate jointly contributing to the purchase of that good. In both cases, the future consumption of one member will depend on the wealth of the other. In this section, I show that these indirect links are sufficient to generate time-inconsistent overconsumption. I defer any analysis of the endogenous choice to save in shared versus separate accounts to Section IV.

A. *Setup with Separate Accounts*

To incorporate individual ownership of financial wealth, I introduce a measure of the wealth that each member has at the beginning of a period: $W_{i,t}$. Within each period, the timing of events is as follows. First, household members simultaneously choose a nonnegative amount $\Psi_{i,t} \in [0, W_{i,t}]$ to transfer to the other member. Transfers are chosen as noncooperative Nash best responses. Once transfers have been made, each member takes as given their wealth, net of transfers, $\widetilde{W}_{i,t} = W_{i,t} - \Psi_{i,t} + \Psi_{j,t}$, and selects a level of private consumption $C_{i,t} \geq 0$ and a nonnegative contribution to the purchase of the public consumption good $H_{i,t} \geq 0$. I assume that the consumption choices and wealth of each member are common knowledge.[21] I allow for the possibility that $C_{i,t} + p_H H_{i,t} > \widetilde{W}_{i,t}$, so that members can borrow, at gross interest rate R, against transfers they will receive in the future. Implicit in this formulation is that member i must repay any debts carried from the previous period (captured by $W_{i,t} < 0$) before consuming out of remaining wealth $\widetilde{W}_{i,t}$. Note that neither member is compelled to repay the debt of her partner, although she may choose to do so through a voluntary transfer. No new borrowing is possible at $t = T$, and in equilibrium all loans will be repaid with certainty.[22] The intertemporal budget constraint of each member is analogous to the scenario with

[21] Several papers consider the role of asymmetric information within the household, including Ambler (2015), Jakiela and Ozier (2016), and Castilla (2019).

[22] This is assured because marginal utility is infinite at zero consumption.

shared wealth

$$W_{A,t} = R\left(\widetilde{W}_{A,t-1} - C_{A,t-1} - p_H H_{A,t-1} + \varpi_t\right) \text{ and } W_{B,t} = R\left(\widetilde{W}_{B,t-1} - C_{B,t-1} - p_H H_{B,t-1} - \varpi_t\right), \tag{20}$$

where each member's initial wealth endowment $W_{i,1}$ is given. This intertemporal budget constraint includes a stochastic uninsurable shock to the relative wealth of the household members: $\varpi_t \sim N(0, \sigma_t)$. This captures shocks to the market value of the human capital of each member, stochastic inheritances, or an unpredictable component to savings returns. The shock can also be thought of as providing a reduced-form representation of shocks to relative consumption needs that may arise, for example, from unexpected illness. Consistent with the setting in which wealth is shared, I abstract from shocks to the total wealth of the household, which will not meaningfully change the analysis, focusing entirely on shocks to the relative wealth of each member. The presence of a relative wealth shock captures one key benefit to household membership: risk sharing.[23] Shocks to initial wealth are subsumed into the analysis by studying any initial combination of $W_{A,1}$ and $W_{B,1}$, and thus, without loss of generality, I set $\sigma_1 = 0$. I also assume that the standard deviation of the relative wealth shock is constant over the rest of the household's life $\sigma_t = \sigma$ for $t \geq 2$.

The objective function of each member is unchanged from before as per (3). Since the objective of each member is the same and the resources of the combined household are unaffected by the use of separate accounts, it follows that the Pareto-efficient full-commitment consumption allocation is the same as when household

[23]Evidence of the importance of risk sharing in households is provided by papers such as Kotlikoff and Spivak (1981), Rosenzweig and Stark (1989), and Hess (2004).

members share wealth.

I limit the analysis to the case in which $T = 3$, which is sufficient to capture the key forces in this setting. In general, it is not possible to find a closed-form analytical solution to characterize equilibrium choices when members save in separate accounts. For this reason, I present numerical solutions to the model.

B. Equilibrium Choices at $t = 3$ for the Log Cobb-Douglas Case

Intuition for the numerical solutions discussed below can be obtained from the analytical characterization of equilibrium consumption and transfer choices in the final period for the case of log Cobb-Douglas preferences ($\varepsilon = z = \xi = 1$).[24] This is qualitatively representative of the equilibrium that obtains under more general preferences. Since the household is symmetric each period, apart from the level of the state variables $W_{A,t}$ and $W_{B,t}$, the equilibrium in any period can be fully described without loss of generality by limiting attention to the scenario in which $W_{i,t} \geq W_{j,t}$. The equilibrium choices of the household at $t = 3$ fall into one of three regions, depending on the relative magnitude of $W_{i,3}$ and $W_{j,3}$:

$$\begin{cases} H^*_{i,3} > H^*_{j,3} = 0, \quad \Psi^*_{i,3} > \Psi^*_{j,3} = 0 \quad \text{when } W_{i,3} > W_{j,3} W_3^{Trans} \\ H^*_{i,3} > H^*_{j,3} = 0, \quad \Psi^*_{i,3} = \Psi^*_{j,3} = 0 \quad \text{when } W_{i,3} \in \left[W_{j,3} W_3^{Joint}, W_{j,3} W_3^{Trans}\right] \\ H^*_{i,3}, H^*_{j,3} > 0, \qquad \Psi^*_{i,3} = \Psi^*_{j,3} = 0 \quad \text{when } W_{i,3} \in \left[W_{j,3}, W_{j,3} W_3^{Joint}\right], \end{cases} \tag{21}$$

where W_3^{Trans} and W_3^{Joint} are constants such that $W_3^{Trans} > W_3^{Joint} > 1$.[25] In words, when both members start the period with wealth that is sufficiently close

[24]Derivations for this scenario are provided in Section III of the Internet Appendix.
[25]Both constants are defined in the Internet Appendix in (IA133) and (IA117), respectively.

in magnitude ($\frac{W_{i,3}}{W_{j,3}} \leq W_3^{Joint}$), they both endogenously contribute to the purchase of the public good. In the analysis that follows, I refer to W_3^{Joint} as the "joint-contribution threshold." When the relative wealth of i exceeds this threshold, she will be the only one who contributes to the purchase of the public good. If the discrepancy in wealth is sufficiently large ($\frac{W_{i,3}}{W_{j,3}} \geq W_3^{Trans}$), i will be the sole provider of the public good *and* endogenously elect to transfer wealth to j, who will combine this with his own wealth to purchase the private good for himself. In the analysis that follows, I refer to W_3^{Trans} as the "transfer threshold."

The resulting equilibrium value function for member i across all possible relative wealth combinations is

$$V_{i,3}\left(W_{i,3}, W_{j,3}\right) = \left\{ \begin{array}{cc} \ln\left(W_{i,3} + W_{j,3}\right) + v_3^I & \text{when } W_{i,3} > W_{j,3} W_3^{Trans} \\ b_{i,3}^{II} \ln\left(\widetilde{W}_{i,3}\right) + b_{j,3}^{II} \ln\left(\widetilde{W}_{j,3}\right) + \widetilde{v}_3^{II} & \text{when } W_{i,3} \in \left[W_{j,3} W_3^{Joint}, W_{j,3} W_3^{Trans}\right] \\ \ln\left(W_{i,3} + W_{j,3}\right) + v_3^{III} & \text{when } W_{i,3} \in \left[\frac{W_{j,3}}{W_3^{Joint}}, W_{j,3} W_3^{Joint}\right] \\ b_{i,3}^{IV} \ln\left(\widetilde{W}_{i,3}\right) + b_{j,3}^{IV} \ln\left(\widetilde{W}_{j,3}\right) + \widetilde{v}_3^{IV} & \text{when } W_{i,3} \in \left[\frac{W_{j,3}}{W_3^{Trans}}, \frac{W_{j,3}}{W_3^{Joint}}\right] \\ \ln\left(W_{i,3} + W_{j,3}\right) + v_3^V & \text{when } W_{i,3} < \frac{W_{j,3}}{W_3^{Trans}} \end{array} \right], \tag{22}$$

where $\left\{b_{i,3}^{II}, b_{j,3}^{II}, b_{i,3}^{IV}, b_{j,3}^{IV}\right\}$ and $\left\{v_3^I, \widetilde{v}_3^{II}, v_3^{III}, \widetilde{v}_3^{IV}, v_3^V\right\}$ are constants defined in the Internet Appendix.[26] The key feature of (22) is that the value function of both members is determined by the *joint wealth of both members* $(W_{i,3} + W_{j,3})$ when their wealth is sufficiently close (and they both contribute to the public good) and when their wealth is sufficiently different (and the wealthy member anticipates making a transfer to the other member). Anticipating this value function in prior periods, the private incentive to save of both members will be affected by the

[26]Note that $V_{i,3}\left(W_{i,3}, W_{j,3}\right)$ and its first derivatives are continuous in both arguments, which ensures that the marginal effect of a change in the wealth of either member is captured purely by its direct effect on value function in each region and not by its effect on shifting the cutoff between regions.

possibility of either of these scenarios arising.

When both members contribute to the public good, each member takes into account the contribution the other will make given their level of wealth. As a result, when household members make saving decisions anticipating this scenario in the future, their private incentive to save is diminished because they recognize that an additional dollar of saving will partially offset the contribution to the public good from the other member.

Similarly, when member i makes a transfer to j, she does so taking into account the relative wealth of both members. As a result, anticipating this scenario in the future will reduce the private incentive to save for both members. For the member with more wealth, they will recognize that a portion of each dollar saved will go toward the consumption of the other member, decreasing their private incentive to save. Similarly, for the member with less wealth, they will recognize that an additional dollar of savings will offset the transfer they will receive from the other member, decreasing their private incentive to save as well.

As a result, at $t < 3$, despite having separate accounts, both household members will anticipate that their wealth will in effect be shared in some future states of the world and thus remains partially subject to the same dynamic commons problem present when wealth is shared. Provided that the support of the relative wealth shock is sufficiently large, both members will place strictly positive probability on reaching these scenarios and, as I show below, will overconsume for the same reason they did when wealth was shared. Notice also that only one of these two channels (joint contribution to the public good or the possibility that one member makes a voluntary transfer to the other) is needed to generate dependence on joint wealth and hence produce overconsumption through the dynamic commons problem.

C. Equilibrium Choices at $t = 1$

C.1. Consumption and Transfer Choices as a Function of the Initial Wealth Distribution

Equilibrium transfer and consumption choices at $t < T$ share the same qualitative pattern as described for $t = T$. This is illustrated in Figure IA5, which shows equilibrium outcomes at $t = 1$.[27] Panel A shows that both members contribute to the public consumption good when their wealth is sufficiently close (labeled "joint-contribution region") and that the household member with more wealth endogenously transfers wealth to the other when the wealth imbalance is sufficiently large (labeled "transfer region").[28] Panel B shows that over both these regions, the value function of each member is unaffected by the relative wealth of each member and is therefore determined only by the combined level of household wealth.

C.2. Time-Inconsistency with Separate Accounts

Each period, household members make consumption and savings decisions anticipating the possibility that they will either jointly contribute to public consumption or make (receive) transfers to (from) each other in the future. Anticipation of each of these scenarios reduces the private incentive to save in prior periods. The resulting time-inconsistency is illustrated in Panel A of Figure 2, which shows the degree to which total household expenditure in the first period exceeds the

[27]The same qualitative pattern holds at $t = 2$.

[28]Note that outcomes are shown for the relative wealth range in which $\frac{W_{1,A}}{W_1} \in \left[\frac{1}{2}, 1\right]$. Equilibrium outcomes when $\frac{W_{1,B}}{W_1} \in \left[\frac{1}{2}, 1\right]$ are symmetric to those described here.

full-commitment optimum and compares this to the extent of overconsumption when wealth is shared (and hence invariant to relative wealth). The combined effect of the suboptimal equilibrium consumption decisions is summarized in the value of commitment shown in Panel B of Figure 2, which compares its level to the case of shared savings. The degree of overconsumption and the value of commitment is lower with separate accounts than with shared savings. This occurs because members internalize more of the benefits of their savings. However, for the parameters used in Figure 2, roughly 85% or more of the original value of commitment remains. This is due to the combination of two factors. First, anticipating the possibility of ongoing financial connection between members continues to produce time-inconsistency and overconsumption. In addition, there is sub-optimal risk sharing under separate accounts, a problem that is perfectly solved with a joint account.[29] The value of commitment is also nonmonotonic in relative wealth, decreasing slightly for intermediate wealth imbalances where members save anticipating a lower probability of being financially linked in the future. The value of commitment problem is most severe at an extreme wealth imbalance, when one member consumes anticipating a high likelihood of more transfers in the future.

These results contain two crucial lessons about time-inconsistency within the household. First, the dynamic commons problem that leads to systematic undersaving does not rely on household members saving in a joint account - sharing public consumption and/or the possibility of endogenous transfer due to altruism is sufficient to generate this problem. Second, the use of separate accounts does lower the severity of the time-inconsistency by creating future states of the world

[29] For evidence of imperfect risk sharing in households, see, for example, Dercon and Krishnan (2000) and Attanasio et al. (2015)).

Figure 2. Time-Inconsistency with Separate Accounts

This plot shows the degree of overconsumption (Panel A) and value of commitment (Panel B) at $t = 1$ when household members save in separate accounts. Without loss of generality, both plots show outcomes over the range of relative wealth where $\frac{W_{1,A}}{W_1} \in \left[\frac{1}{2}, 1\right]$, since the equilibrium outcomes below this range are symmetric with this case. For comparison, the value of commitment under joint accounts is also shown (invariant to the distribution of relative wealth between the members). Each plot is drawn using the following parameters: household members have an exponential discount factor of $\delta = 0.95$, the gross interest rate is $R = 1/0.95$, the household exists for $T = 3$ years, the elasticity of substitution between goods is unity $\varepsilon = 1$, the intertemporal elasticity of substitution is unity $z = 1$, the degree of risk aversion is unity $\gamma = 1$, the elasticity between each member's private utility is unity $\xi = 1$, the standard deviation of the relative wealth shock is $\sigma = 0.1$, the degree of self-interest within the household is $\Delta = 0.4$, the utility weight on the public good is $\mu_H = 0.4$, and the relative price of the public good is one $p_H = 1$.

Panel A. Time-Inconsistency of Total Household Consumption

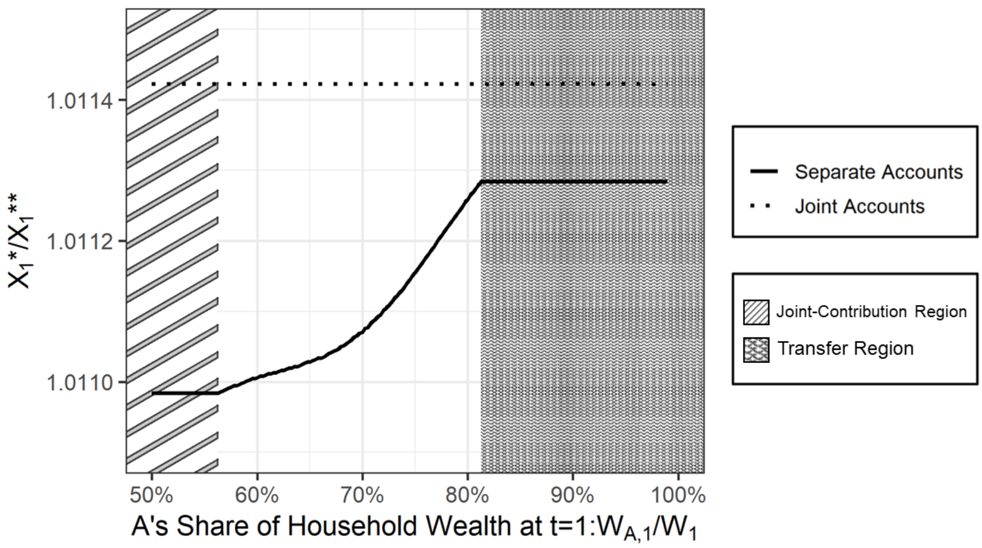

Panel B. Value of Commitment

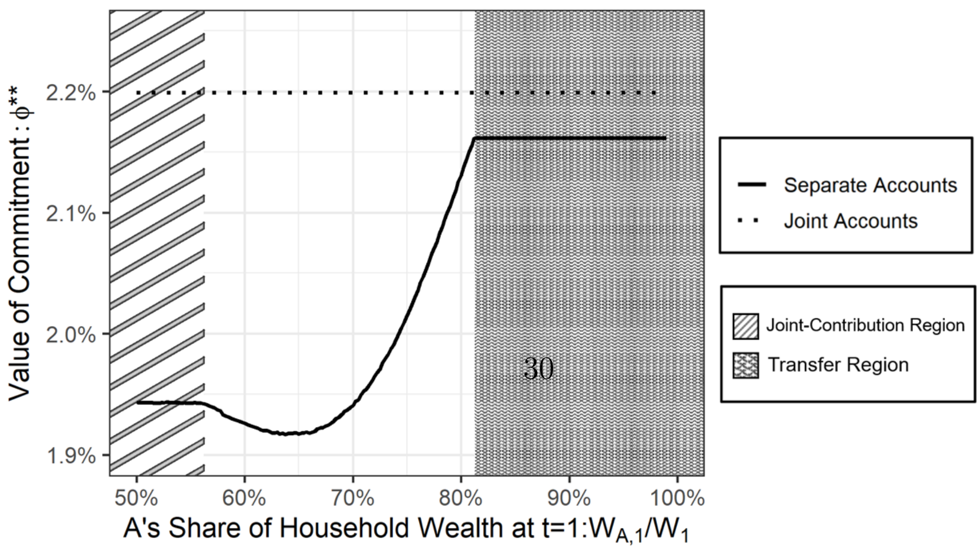

where the members are financially autonomous. This is consistent with empirical evidence that use of additional individual savings accounts can increase savings by protecting wealth from other immediate and extended family members.[30] It is also consistent with several papers that document household preference to keep savings separate from other household members.[31]

C.3. Equilibrium Allocation of Consumption within the Household

Figure 3 decomposes total household overconsumption. Panel A shows that private consumption exceeds the full-commitment optimum, but by less than when the household saves in a joint account. The panel also shows that total private consumption is lower for intermediate levels of relative wealth imbalance, in the range over which future financial independence is more likely. Panel B shows that this effect comes from a reduction in private consumption from the less wealthy member who spends less anticipating the need to save for their private consumption in the future. A striking result is that beyond a point, and holding total household wealth constant, the member with less wealth actually spends *more* on private consumption when their relative wealth is *lower*. This occurs because a larger wealth discrepancy implies a higher probability of receiving a transfer in the future. Since any additional dollar of saving will reduce this transfer, this lowers the marginal value of saving.

Panel C shows that, by contrast, expenditure on the public good is unambiguously increasing in the degree of relative wealth imbalance as the wealthy member

[30] See, for example, Anderson and Baland (2002), Dupas and Robinson (2013a), Dupas and Robinson (2013b), Ambler (2015), Hoel (2015), Jakiela and Ozier (2016), Schaner (2017), and Boltz et al. (2019).

[31] See, for example, Castilla and Walker (2013), Castilla (2015), Castilla (2019), Afzal et al. (2022), and D'Exelle and Ignowski (2022).

Figure 3. Time-Inconsistency with Separate Accounts - Decomposed

These plots decompose the time-inconsistent consumption choices of the household with separate accounts at $t = 1$. Panel A shows the total level of private consumption relative to the full-commitment Pareto optimum. Panel B shows the level of individual private consumption of each member measured relative to total household wealth at $t = 1$. Panel C shows the total level of public consumption relative to the full-commitment Pareto optimum. Without loss of generality, both plots show outcomes over the range of relative wealth where $\frac{W_{1,A}}{W_1} \in \left[\frac{1}{2}, 1\right]$ since the equilibrium outcomes below this range are symmetric with this case. For comparison, each plot also contains the equilibrium outcome with joint accounts (invariant to the relative wealth of each member). Each plot is drawn using the following parameters: household members have an exponential discount factor of $\delta = 0.95$, the gross interest rate is $R = 1/0.95$, the household exists for $T = 3$ years, the elasticity of substitution between goods is unity $\varepsilon = 1$, the intertemporal elasticity of substitution is unity $z = 1$, the degree of risk aversion is unity $\gamma = 1$, the elasticity between each member's private utility is unity $\xi = 1$, the standard deviation of the relative wealth shock is $\sigma = 0.1$, the degree of self-interest within the household is $\varDelta = 0.4$, the utility weight on the public good is $\mu_H = 0.4$, and the relative price of the public good is one $p_H = 1$.

Panel A. Total Private Consumption

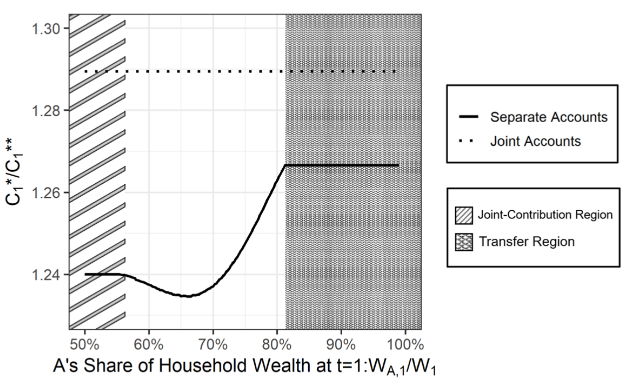

Panel B. Individual Private Consumption

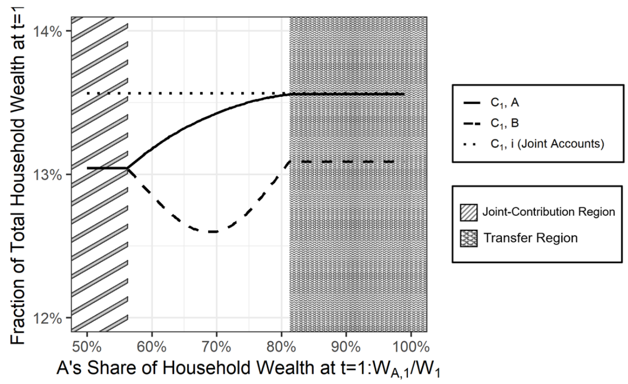

Panel C. Public Consumption

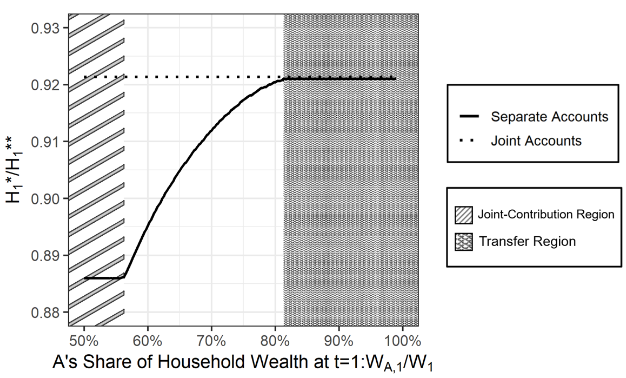

effectively becomes the planner for the household. Panel C also shows that allocative inefficiency, whereby the public good is underprovided, is more acute than in the case with shared wealth, and is most severe when the relative wealth of each member is most equal. A testable implication of these results is that transfer programs that reduce intra-household wealth inequality might have the unintended effect of decreasing expenditure on public goods such as child education and health.[32]

C.4. Explaining the Decision to Transfer Wealth at $t = 1$

The finding that in equilibrium the wealthy member will transfer wealth at the beginning of the life of the household is not immediately obvious. Waiting to transfer preserves the option value of adapting any transfer to the realization of future relative wealth shocks and does not constrain the ability of the less wealthy member to consume early in life since she can always elect to borrow. The incentive to transfer early comes from increasing the probability that both members will be financially independent in the future and hence decreasing the incentive for the less wealthy member to overconsume. Panel A of Figure IA6 illustrates this incentive by showing the private consumption of the less wealthy member with and without the transfer in the first period. The private consumption of the less wealthy member is reduced as a result of receiving the transfer at $t = 1$. Making a transfer credibly commits the wealthy member to being less likely to transfer wealth in the future by lowering the expected future wealth imbalance. As a result, the less wealthy member places a higher marginal value on savings

[32]This relies on the assumption that household members place the same utility weight on public consumption.

since an additional dollar of savings is less likely to crowd out a future transfer from the other member. The wealthy member benefits from this increased savings because it lowers the expected transfer they will make in the future. For the household as a whole, the effect of the transfer is to lower total expenditure at $t = 1$ and thereby lower total overconsumption (as shown in Panel B of Figure IA6).

C.5. Comparative Statics in Δ, μ_H , ε, and z

The fundamental lesson from this section is that despite holding wealth in separate accounts, the household remains partially subject to the same fundamental dynamic commons problem as when wealth is shared. As a consequence, most of the comparative static relationships that characterize the household with joint wealth continue to hold quantitatively.

The welfare cost of the time-inconsistency problem, as captured by the value of commitment, remains monotonically increasing in the degree of self-interest (Δ) within the household (Panel A of Figure IA7). As when wealth is shared, this is primarily because, as self-interest increases, each member places less private value on savings, since, with some probability, this wealth will effectively be shared with the other member. As a result, in most cases, the degree of overexpenditure relative to the Pareto first-best increases in the degree of self-interest (Panel B of Figure IA7 shows the degree of overconsumption at $t = 1$). An exception to this pattern arises when the difference in wealth between the two members is high and $\Delta \to 1$ (Panel B of Figure IA7). In this case, increased self-interest credibly lowers the expectation of any transfer that will occur in the future and therefore lowers the resulting overconsumption. Despite this, the value of commitment continues to rise because any benefit this creates in terms of decreasing overconsumption is outweighed by inferior risk sharing (as confirmed in Panel A of Figure IA7).

As in the case of shared wealth, the value of commitment, for most parameters, is increasing in the utility weight placed on the public good (μ_H) (as shown in Panel A of Figure IA8). As before, an increased concern for the public good unites the objective of each member, which reduces the dynamic commons problem. Also similar to the case when wealth was shared, this monotonic relationship breaks down when a public and private good are near-perfect substitutes (ε is large). In this case, the allocative inefficiency, whereby members overspend on the private good, is pronounced at intermediate values of μ_H. As a result, when ε is large, the value of commitment is highest at intermediate values of μ_H (as shown in Panel B of Figure IA8 where $\varepsilon = 5$).

The comparative statics related to the elasticity of intertemporal substitution (z) are qualitatively unchanged from the case where members share wealth. The value of commitment is increasing in most values of z (Panel A of Figure IA9) because a higher elasticity unambiguously exacerbates the degree of overconsumption (as shown in Panel B of Figure IA9). However, once the elasticity becomes sufficiently high, the utility cost of this intertemporal inefficiency and the associated value of commitment fall in z.

C.6. Comparative Statics in σ, γ, and ξ

When household members save in separate accounts, a new set of comparative static relationships exist that are not present with joint accounts. These relate to the role of the relative wealth shock, which is diversified away when wealth is shared, and the curvature of the objective function, which determines risk aversion and the utility consequence of intra-household consumption differences.

The value of commitment is monotonically increasing in the standard deviation of the relative wealth shock σ (as shown in Panel A of Figure IA10). In part, this follows because full commitment provides perfect risk sharing while the equilibrium with separate accounts leaves the consumption outcomes of each member subject to the relative wealth shock. Consistent with this logic, the value of commitment is increasing in the degree of risk aversion γ (as shown in Panel B of Figure IA10). The elasticity of substitution between the individual consumption of each member, as captured by ξ, also captures the curvature of each member's objective function (with respect to the distribution of private consumption within the household). Accordingly, the value of commitment is decreasing in this elasticity for most parameter values (as shown in Panel C of Figure IA10).

Despite the importance of imperfect risk sharing with joint accounts, the relative wealth shock has a second, and often larger, effect on the time-inconsistency problem that is independent of risk aversion. When the standard deviation of the relative wealth shock σ is high, this alters the expected probability that members will share responsibility for purchasing the public good and the probability that one member will elect to transfer wealth to the other. For most parameter values and initial levels of relative wealth, an increase in σ increases the likelihood that the household members will endogenously be financially linked in the future. This is illustrated in Panel D of Figure IA10, which is drawn for the case in which the curvature of the objective function is low $\gamma = 0.1$ and $\xi = 5$. As before, the value of commitment is increasing in the amount of relative wealth uncertainty. Further, the range over which a transfer occurs at $t = 1$, as captured by the flat region on the right-hand side of each value-of-commitment curve, is monotonically decreasing in σ.

C.7. Robustness of Relative Wealth Shock Distribution

In the analysis so far, the relative wealth shock has been assumed to follow a normal distribution under which arbitrarily large relative wealth shocks are possible. As a result, in expectation, there is always a strictly positive probability that either member will make or receive a transfer in the future. However, the central results of the model are qualitatively unchanged if the relative wealth shock is instead generated by a distribution with finite support. To illustrate this, I solve for equilibrium household consumption choices assuming that $\varpi_t \sim U\left(-\frac{\sqrt{12}}{2}\sigma, \frac{\sqrt{12}}{2}\sigma\right)$ for $t = 2, 3$.[33] The equilibrium consumption choices of the household remain time-inconsistent and are qualitatively similar to the case in which the shock is normally distributed. To illustrate this, in Figure IA11 (Panels A and B, respectively), I show how the value of commitment varies with the degree of self-interest within the household and the standard deviation of the relative wealth shock under this alternate distributional assumption.[34]

D. Marginal Propensity to Consume

Many transfer programs are designed to change the relative wealth of household members, often increasing the relative wealth of women versus men. Well-studied examples include the 1979 U.K. Reform of Child Benefits analyzed by Lundberg et al. (1997) and Ward-Batts (2008), the Mexican PROGRESA program (Attanasio and Lechene (2002, 2014)), or the reform of the South African social pension

[33]The bounds of the uniform distribution are set to ensure that σ remains the standard deviation of the distribution.

[34]Compare the results to Panel A of IA7 and Panel A of IA10, respectively.

program studied by Duflo (2003). A consistent theme in these papers is that the way a transfer is spent depends in part on who within the household receives the additional funds. One possible explanation, employed to understand this within the context of efficient bargaining, is that the allocation of funds may alter the bargaining power of each member. However, the possible link between the allocation of funds and bargaining power (suggested, for example, by Lundberg et al. (1997)) is only suggestive and relies on the assumption that households negotiate to efficient outcomes, a prediction that many empirical studies reject. Moreover, many of these programs often target specific household members to increase public consumption (such as spending on child healthcare and education), relying on the argument that one member (typically a woman) places a higher utility weight on these goods than the other. Even if this is the case, other factors besides differences in utility weights are likely to play a role.[35] The model in this paper with separate accounts offers a noncooperative framework for understanding how the recipient of the additional dollar affects the way a transfer is spent. To illustrate this, Panel A of Figure 4 shows the marginal propensity to consume, measured in terms of total household expenditure, for all possible initial intra-household distributions of wealth. Three values of the marginal propensity are shown based on whether member A or B receives the additional dollar or whether it is received by both in proportion to their initial wealth (and hence does not alter the distribution of wealth within the household).

The marginal propensity to consume varies based on who receives the additional dollar when the distribution of wealth is between the joint-contribution and

[35] For example, Ringdal and Sjursen (2021) suggest that time preferences, not gender, better explain how a transfer is spent.

Figure 4. Marginal Propensity to Consume with Separate Accounts

This plot shows the marginal propensity to consume for total household expenditure (Panel A) and expenditure on the public good (Panel B) at $t = 1$. The marginal propensity is shown based on whether an incremental dollar is given to member A, member B, or both household members in proportion to their initial wealth. Without loss of generality, this plot is shown over the range of relative wealth where $\frac{W_{1,A}}{W_1} \in \left[\frac{1}{2}, 1\right]$ since the equilibrium outcomes below this range are symmetric with this case. Each plot is drawn using the following parameters: household members have an exponential discount factor of $\delta = 0.95$, the gross interest rate is $R = 1/0.95$, the household exists for $T = 3$ years, the elasticity of substitution between goods is unity $\varepsilon = 1$, the intertemporal elasticity of substitution is unity $z = 1$, the degree of risk aversion is unity $\gamma = 1$, the elasticity between each member's private utility is unity $\xi = 1$, the standard deviation of the relative wealth shock is $\sigma = 0.1$, the degree of self-interest within the household is $\Delta = 0.4$, the utility weight on the public good is $\mu_H = 0.4$, and the relative price of the public good is one $p_H = 1$.

Panel A. Marginal Propensity to Consume - Total Household Expenditure

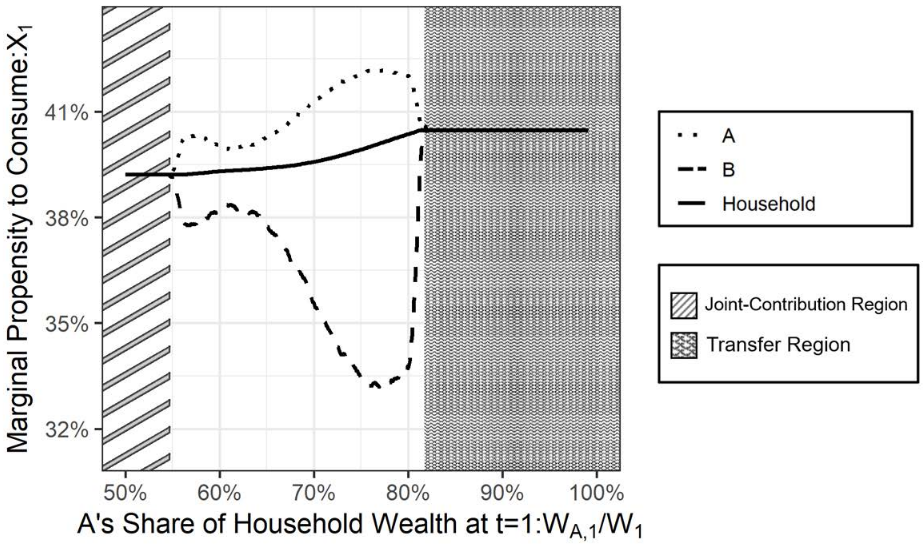

Panel B. Marginal Propensity to Consume - Public Consumption

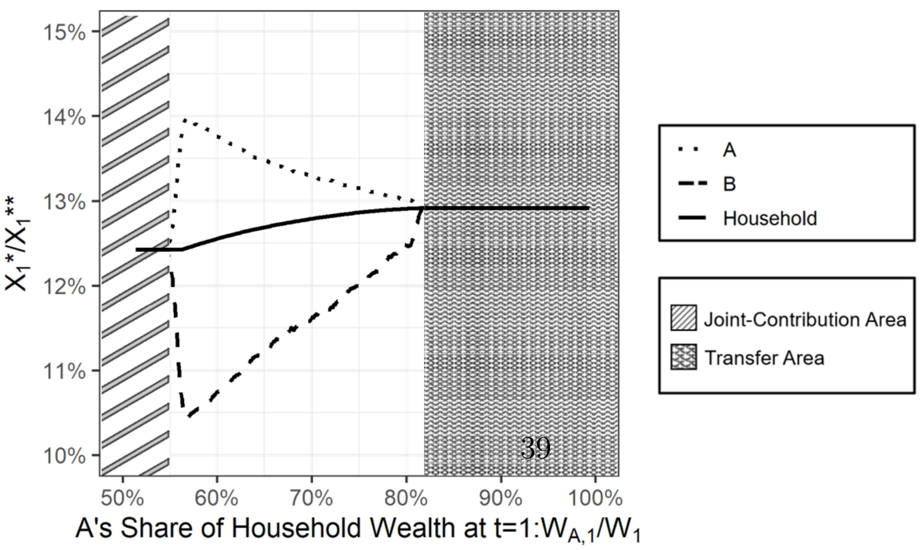

transfer threshold. In this region, giving an additional dollar to A (the relatively wealthy member) further exacerbates the relative wealth imbalance and thereby increases the probability that A will transfer wealth to B in the future. The increased anticipation of a wealth transfer in the future reduces the private incentive to save for both members, thereby leading to a higher marginal propensity to consume. In contrast, giving an additional dollar to B lowers the wealth disparity and, as a result, leads to a smaller marginal propensity to consume. Panel B of Figure 4 shows that the same forces generate a qualitatively similar pattern for the household's marginal propensity to consume the public good. This paper's framework therefore suggests that if a transfer program is designed to promote household savings (or limit overconsumption), then assuming symmetric preferences, directing funds to the member with less wealth will be effective. Conversely, if the program is designed to produce a short-term increase in total consumption (for example, to stimulate aggregate consumption) or to increase immediate consumption of public goods, then all else equal, a transfer directed to the member with the most wealth will have a larger effect.

Figure 4 also shows that targeted transfers will be effective only in some scenarios. When the initial wealth distribution is close to equal (below the joint-contribution threshold) or highly unequal (above the transfer threshold), the marginal propensity to consume is unaffected by who receives the additional dollar. In both of these scenarios either the relative contribution to the public good or the transfer between members will endogenously adjust to undo any effect the additional dollar has on the allocation of wealth within the household.[36]

[36]This logic also applies to the marginal propensity to consume measured in terms of the individual components of household consumption (as illustrated in Figure IA12).

When considering how a transfer affects the private consumption of each member (shown in Panel A and B of Figure IA12), the predominant force is intuitive: i's private consumption will increase more if i is given an additional dollar, as compared to the scenario in which the dollar is received by j. However, a striking counterexample exists when the wealth distribution is close to the transfer threshold. In this case, illustrated in Panel B of Figure IA12, B will increase her private consumption *more* in response to an additional dollar given to A than if she receives it herself. This occurs because giving an additional dollar to B also lowers the expected level of any future transfer and hence increases the private incentive of both members to save.

An implication of these results is that the effect of government transfer programs to people in households with separate savings (such as extended families) will depend on the way wealth is distributed within the family and who receives the transfer. Predicting how members of an extended family will react to a transfer or income shock should therefore take into account the entire wealth distribution across all household members, including the extended family. The model also implies that a larger proportion of a positive wealth shock will be saved by an individual with a wealthier extended family. This would therefore have the potential to explain why wealth differences between households are persistent.

IV. Choice of Separate or Joint Accounts

Having studied household consumption and savings decisions with separate and joint accounts, I now examine the question of which savings arrangement a household will select. To do so, I extend the model of the household by allowing household members to decide which savings method they prefer each period.[37]

A. Timing of Events with Endogenous Account Choice

The timing of events within the first period is as follows. Members start life at the beginning of $t = 1$ with separately held individual wealth of $W_{A,1}$ and $W_{B,1}$. Each member sends a message $A_{i,1} \in \{0, 1\}$ indicating whether they consent to combining all of their wealth in a joint account ($A_{i,1} = 1$) or not ($A_{i,1} = 0$). A joint account is adopted if and only if both members consent: $A_1 = \min\{A_{A,1}, A_{B,1}\}$. If both members do consent to combining wealth in a joint account ($A_1 = 1$), they then make simultaneous noncooperative consumption decisions from their shared wealth as per the framework described in Section I. I assume that if $A_t = 1$ at any t, then the household will continue to save and consume using a joint account for the rest of their life.[38] Alternatively, if both members do not consent to share wealth ($A_1 = 0$), then the sequence of events within the period follows the framework described in Section III: members simultaneously choose transfers and then make simultaneous noncooperative consumption decisions. Crucially, if members elect to adopt separate accounts at $t-1$, then wealth evolves between $t-1$ and t according to the individual intertemporal budget constraint in (20) and is therefore affected by the relative wealth shock. For any period $t > 1$, where $A_{t-1} = 0$, the timing of

[37]The ability of household members to separate their finances may face legal limitations. For example, following the federal Uniform Marriage and Divorce Act (UMDA) of 1970, U.S. states moved toward equitable distribution of household assets (see Golden (1983), Turner (2005), and Voena (2015)). Thus, even with separate accounts, if both members of a married household recognize that divorce is possible with some probability, then they will anticipate that their wealth is shared with some probability in the future.

[38]It is also possible to consider the endogenous choice to evenly divide joint wealth into separate accounts. However, abstracting from this comes with no loss of generality since I have not found a numerical solution for any set of parameters in which the household would elect to share wealth at t and then endogenously reverse this decision at a later period.

events is identical to that described for the first period. By the standard argument of subgame perfection, members rationally adapt their expected value function to incorporate the scenarios under which joint or separate accounts will be adopted in the future. For analytical tractability, I continue to limit analysis to the case in which $t = 3$ and present numerical solutions of the model.

B. *Equilibrium Account and Consumption Choice*

I start by describing equilibrium account and consumption choices for any initial division of wealth within the household. For brevity, I focus on choices at $t = 1$.[39] Two fundamental forces determine the decision to adopt a joint account. First, sharing wealth provides complete insurance against relative wealth shocks. In contrast, separate accounts implement only partial risk sharing through voluntary transfers. Second, maintaining separate accounts improves the private incentive to save and thus lowers the utility cost of overconsumption. Since choosing to save with joint accounts requires the approval of both members, the decision will hinge on the choice of the member who has more wealth at the start of the period.[40] As a result, the equilibrium account choice of the household can be characterized by a "joint-account threshold," which is the maximum wealth difference at which a joint account will be selected for a given set of parameters.

If parameters are such that the household adopts a joint account, equilibrium consumption choices will fall into one of three regions, depending on the initial allocation of wealth within the household.[41] These are illustrated in Figure 5. When

[39] The same qualitative trade-off characterizes the decision at $t = 2$. Trivially, there is no incentive to change the type of accounts used by the household in the final period.

[40] Technically, this is assured because with separate accounts and endogenous transfers, $V_{i,t}(W_{i,t}, W_{j,t})$ is non-decreasing in $W_{i,t}$.

[41] Trivially, if parameters are such that household members do not elect to adopt a joint account

the allocation of wealth is sufficiently equal (below the joint-account threshold), both members will agree to adopt a joint account. The resulting consumption choices and degree of overconsumption will be identical to that described in Section II. When the initial intra-household allocation of wealth is sufficiently unequal (above the joint-account threshold), the member with the most wealth will not consent to adopting a joint account. Instead, both members will consume out of their individual accounts and only the wealthiest member will purchase the public good.[42] Around the joint-account threshold, there is a discrete change in the private consumption of both members (Panel A of Figure 5) and total household expenditure (Panel B of Figure 5). This discrete change occurs because the use of separate accounts increases the private incentive of each member to save. Despite this, for most parameter values, the value of commitment *increases* discretely at the joint-account threshold (Panel C of Figure 5), reflecting the utility loss of imperfect intra-household risk sharing that comes with using separate accounts. Above this threshold, the degree of overconsumption increases in the degree of wealth inequality in the household. This is due to the increased anticipation of a transfer from the wealthier member in the future. As in the case of separate accounts, if the initial allocation is sufficiently unequal to be above the transfer threshold, the wealthiest member will endogenously transfer wealth to the other member at $t = 1$.

for any initial wealth distribution, most notably when the initial wealth allocation is equal, then equilibrium consumption choices will be identical to the case with separate accounts (see Section III).

[42]If parameters are such that a joint account is adopted at some wealth allocation, this eliminates the scenario in which in equilibrium both members contribute to the public good from separate accounts.

Figure 5. Equilibrium with Endogenous Account Selection

This plot shows equilibrium outcomes at $t = 1$ when household members can endogenously choose each period whether to save using a joint or separate account. Without loss of generality, this plot is shown over the range of relative wealth where $\frac{W_{1,A}}{W_1} \in \left[\frac{1}{2}, 1\right]$ since the equilibrium outcomes below this range are symmetric with this case. Each plot indicates the threshold at which the household elects to adopt a joint account as well as the threshold at which the member with more wealth transfers wealth to the other at $t = 1$. Panel A shows the equilibrium private consumption of each member, Panel B shows total household expenditure measured relative to the first-best full commitment solution, and Panel C shows the value of commitment. Each plot is drawn using the following parameters: household members have an exponential discount factor of $\delta = 0.95$, the gross interest rate is $R = 1/0.95$, the household exists for $T = 3$ years, the elasticity of substitution between goods is unity $\varepsilon = 1$, the elasticity of intertemporal substitution is unity $z = 1$, the degree of risk aversion is unity $\gamma = 1$, the elasticity between each member's private utility is unity $\xi = 1$, the degree of self-interest within the household is $\Delta = 0.4$, the utility weight on the public good is $\mu_H = 0.4$, the standard deviation of the relative wealth shock is $\sigma = 0.15$, and the relative price of the public good is one $p_H = 1$.

Panel A. Equilibrium Private Consumption

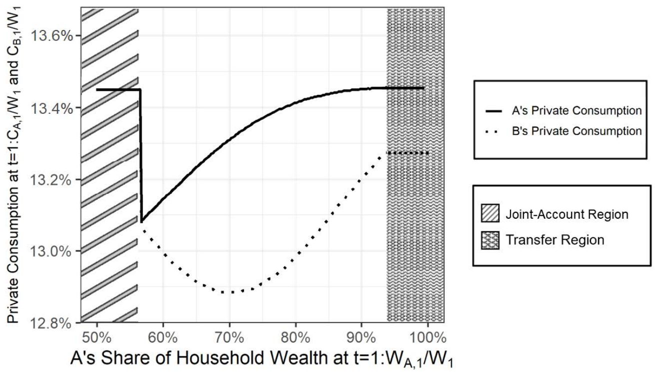

Panel B. Equilibrium Total Household Expenditure

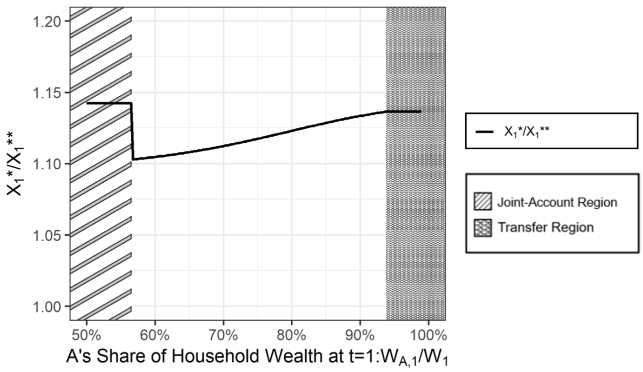

Panel C. Value of Commitment

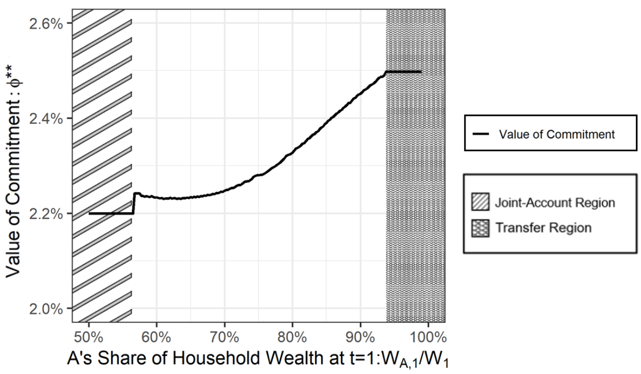

C. Endogenous Account Choice - Comparative Statics

I now present comparative statics for the adoption of joint accounts. To do so, I show how the joint-account threshold varies with model parameters.

The insurance motivation to adopt a joint account is illustrated by the way equilibrium account choice varies with the standard deviation of the relative wealth shock σ (shown in Figure 6).[43] Joint accounts are selected only when σ is sufficiently high, and the range of initial-wealth allocations over which this is selected is monotonically increasing in the standard deviation of the shock. This result comes from two channels. First, increased intra-household risk improves the benefit of the insurance provided by joint accounts. Second, when members save separately, increased intra-household risk increases the probability that members will make endogenous transfers in the future and hence, in anticipation, reduces the private incentive to save. As a result, the benefit of saving separately is lower when σ is larger.

[43] Qualitatively similar comparative statics are exhibited with respect to the curvature of the period utility function as captured by γ and ξ. However, the magnitude of variation is small and hence the results are omitted from the paper.

Figure 6. Account Choice and Intra-Household Risk

This plot shows the threshold of the relative wealth distribution below which the household will elect to save in a joint account at $t = 1$. The threshold is shown for different levels of the standard deviation of the relative wealth shock σ. Without loss of generality, the threshold is shown in terms of the relative wealth of member A $\frac{W_{1,A}}{W_1} \in \left[\frac{1}{2}, 1\right]$ and a symmetric threshold exists for B below this range. Each plot is drawn using the following parameters: household members have an exponential discount factor of $\delta = 0.95$, the gross interest rate is $R = 1/0.95$, the household exists for $T = 3$ years, the elasticity of substitution between goods is unity $\varepsilon = 1$, the intertemporal elasticity of substitution is unity $z = 1$, the degree of risk aversion is unity $\gamma = 1$, the elasticity between each member's private utility is unity $\xi = 1$, the degree of self-interest within the household is $\varDelta = 0.4$, the utility weight on the public good is $\mu_H = 0.4$, and the relative price of the public good is one $p_H = 1$.

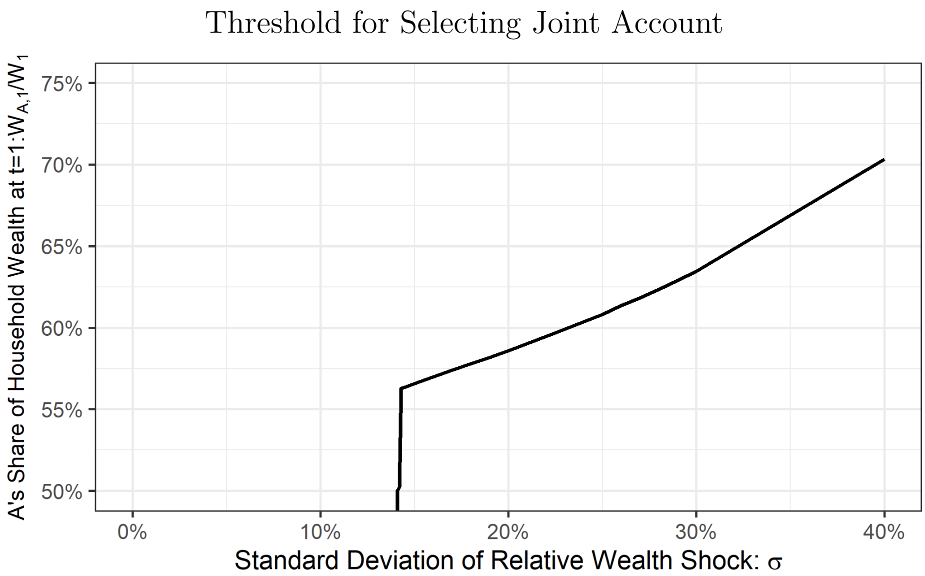

Differences in the alignment of household preferences, captured by either the degree of self-interest Δ or the utility weight on public consumption μ_H, have offsetting effects on household account choice. Low self-interest reduces the incentive to overconsume and in turn lowers the relative benefit of separate accounts for improving the private incentive to save. Conversely, increased self-interest lowers the probability and magnitude of transfers between members, thereby limiting the partial insurance that the household achieves with separate accounts. In combination, these forces lead to a nonmonotonic relationship between the degree of self-interest and account choice (shown in Panel A of Figure IA13) whereby joint accounts are selected at both extremes of the possible level of self-interest.

By a similar argument, increased utility weight on the public good aligns incentives and reduces the relative benefit of separate accounts for the private incentive to save. Conversely, increased concern for the public good limits the effect of relative wealth shocks under separate accounts, because public consumption is determined mainly by total household wealth, and hence limits the relative insurance benefits achieved by adopting a joint account. These forces also lead to a non-monotonic relationship between concern for the public good and account choice (shown in Panel B of Figure IA13).

The household will choose to save in a joint account when the elasticity of intertemporal substitution is low (as shown in Panel C of Figure IA13). A low elasticity of intertemporal substitution partially mitigates the private incentive of each member to overconsume and therefore decreases the relative advantage of holding separate accounts.

V. Discussion and Empirical Implications

The canonical theory of household financial behavior supposes that a household can be modeled as a single optimizing agent. Considerable evidence rejects this premise: household members have distinct preferences and are self-interested.[44] This paper introduces a model of consumption and saving in a multi-person household in which members have shared wealth, consume private and public goods, and

[44]See, for example, Lundberg et al. (1997), Browning and Chiappori (1998), Phipps and Burton (1998), and Ashraf (2009).

have imperfectly aligned altruistic preferences. The central finding of the paper is that the household is time-inconsistent, consuming a larger fraction of wealth each period than under an optimal plan. The model highlights that in a multi-person household, savings are subject to a dynamic commons problem. The tendency to undersave stems from each member's ability to deviate unilaterally from the optimal household plan and direct more resources to their own private consumption at the cost of shared wealth. The extent of time-inconsistency is larger when members are more self-interested and when they place less weight on shared public consumption goods such as children and housing.

Unlike standard theories based on the unitary model of the household, this paper is able to rationalize the way that control of assets within the household can affect consumption and savings decisions. Most notably, holding wealth separately increases the private incentive to save and reduces the extent of time-inconsistency. Despite this, even with separate accounts, members are linked by their altruism and concern for public consumption. A relatively wealthy member may voluntarily transfer funds to the other at some point in the future. Alternatively, if both members have sufficiently equal levels of wealth in the future, they will both contribute to public consumption goods taking into account the contribution of the other member. The possibility of either scenario in the future partially recreates the dynamic commons problem, despite having separate accounts, because members anticipate that their savings may be transferred to or offset a transfer from the other member at some point later in life. This result also suggests that extended family members (for example, adult children) may also be subject to the same time-inconsistency. Finally, I show that household members may optimally choose to share wealth in order to pool risk even though doing so exacerbates the time-inconsistency problem highlighted above.

A saving technology suggested by the analysis in this paper is for household members to save in the form of assets that require joint approval for withdrawals. Joint approval removes the ability of household members to act unilaterally and therefore limits the possibility of overconsumption. As such, the paper provides a framework that can rationalize why some household saving assets require joint approval to withdraw or borrow against. As a primary example, the 1984 U.S. Retirement Equity Act revised the rules governing all retirement plans covered by the 1974 Employee Retirement Income Security Act to require exactly this form of joint approval. This covers all assets held by married households in 401(k) plans, IRA accounts, and defined benefit plans, and thus accounts for the bulk of U.S. retirement savings outside of housing. Aura (2005) shows that the passage of these laws did in fact increase savings for households affected by this legislative change. Similarly, joint ownership of a house prevents a household member from borrowing against home equity savings without the approval of their spouse. Joint approval may come at the cost of significant inflexibility, for example, limiting the ability to adapt consumption choices to privately observed shocks to the marginal utility of each member. Further consideration of this trade-off is left for future work.

Finally, the model suggests that household members may use punishment strategies to mitigate the temptation to overconsume. Evidence that households do not exhibit dynamic efficiency (see, for example, Udry (1996), Duflo and Udry (2004), Mazzocco (2007), De Mel et al. (2009), Robinson (2012), and Lise and Yamada (2019)) suggests that in practice, these strategies are of limited effectiveness. This may be because shocks to marginal utility are unobserved or because households are unable to credibly commit to renegotiate planned punishment strategies.

A more detailed theoretical consideration of savings-related punishment strategies within the household is left for future work.

REFERENCES

Adams, Abi, Laurens Cherchye, Bram De Rock, and Ewout Verriest, 2014, Consume now or later? Time inconsistency, collective choice, and revealed preference, *American Economic Review* 104, 4147–4183.

Afzal, Uzma, Giovanna d'Adda, Marcel Fafchamps, and Farah Said, 2022, Intrahousehold consumption allocation and demand for agency: A triple experimental investigation, *American Economic Journal: Applied Economics* 14, 400–444.

Ainslie, George, 2001, *Breakdown of Will* (Cambridge University Press).

Alexy, Allison, 2007, Deferred benefits, romance, and the specter of later-life divorce, *Contemporary Japan [Japanstudien]* 19, 169–188.

Almås, Ingvild, Alex Armand, Orazio Attanasio, and Pedro Carneiro, 2018, Measuring and changing control: Women's empowerment and targeted transfers, *The Economic Journal* 128, F609–F639.

Amador, Manuel, 2008, Sovereign debt and the tragedy of the commons, Stanford Working Paper.

Ambler, Kate, 2015, Don't tell on me: Experimental evidence of asymmetric information in transnational households, *Journal of Development Economics* 113, 52–69.

Anderson, Siwan, and Jean-Marie Baland, 2002, The economics of roscas and intrahousehold resource allocation, *Quarterly Journal of Economics* 117, 963–995.

Apedo-Amah, Marie, Habiba Djebbari, and Roberta Ziparo, 2020, Gender, information and the efficiency of household production decisions: An experiment in rural Togo, IZA Discussion Paper, No. 12948.

Ashraf, Nava, 2009, Spousal control and intra-household decision making: An experimental study in the Philippines, *American Economic Review* 99, 1245–1277.

Ashraf, Nava, Erica Field, and Jean Lee, 2014, Household bargaining and excess fertility: An experimental study in Zambia, *American Economic Review* 104, 2210–2237.

Ashraf, Nava, Erica M. Field, Alessandra Voena, and Roberta Ziparo, 2020, Maternal mortality risk and spousal differences in the demand for children, NBER Working Paper 28220.

Ashraf, Nava, Dean Karlan, and Wesley Yin, 2006, Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines, *Quarterly Journal of Economics* 121, 635–672.

Attanasio, Orazio, and Valerie Lechene, 2002, Tests of income pooling in household decisions, *Review of Economic Dynamics* 5, 720–748.

Attanasio, Orazio, Costas Meghir, and Corina Mommaerts, 2015, Insurance in extended family networks, NBER Working Paper 21059.

Attanasio, Orazio P., and Valérie Lechene, 2014, Efficient responses to targeted cash transfers, *Journal of Political Economy* 122, 178–222.

Aura, Saku, 2005, Does the balance of power within a family matter? The case of the Retirement Equity Act, *Journal of Public Economics* 89, 1699–1717.

Baseler, Travis, 2023, Hidden income and the perceived returns to migration, *American Economic Journal: Applied Economics* 15, 321–352.

Beshears, John, James J. Choi, David Laibson, Brigitte C. Madrian, and Jung Sakong, 2011, The demand for commitment: Evidence from a savings experiment, NBER Working Paper.

Bloch, Francis, and Vijayendra Rao, 2002, Terror as a bargaining instrument: A case study of dowry violence in rural India, *American Economic Review* 92, 1029–1043.

Bobonis, Gustavo J., 2009, Is the allocation of resources within the household efficient? New evidence from a randomized experiment, *Journal of Political Economy* 117, 453–503.

Boltz, Marie, Karine Marazyan, and Paola Villar, 2019, Income hiding and informal redistribution: A lab-in-the-field experiment in Senegal, *Journal of Development Economics* 137, 78–92.

Browning, Martin, Francois Bourguignon, Pierre-André Chiappori, and Valerie Lechene, 1994, Income and outcomes: A structural model of intrahousehold allocation, *Journal of Political Economy* 102, 1067–1096.

Browning, Martin, and Pierre-Andre Chiappori, 1998, Efficient intra-household allocations: A general characterization and empirical tests, *Econometrica* 66, 1241–1278.

Browning, Martin, Pierre-Andre Chiappori, and Valerie Lechene, 2006, Collective and unitary models: A clarification, *Review of Economics of the Household* 4, 5–14.

Castilla, Carolina, 2015, Trust and reciprocity between spouses in India, *American Economic Review* 105, 621–624.

Castilla, Carolina, 2019, What's yours is mine, and what's mine is mine: Field experiment on income concealing between spouses in India, *Journal of Development Economics* 137, 125–140.

Castilla, Carolina, and Thomas Walker, 2013, Is ignorance bliss? The effect of asymmetric information between spouses on intra-household allocations, *American Economic Review* 103, 263–268.

Cesarini, David, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Östling, 2017, The effect of wealth on individual and household labor supply: Evidence from Swedish lotteries, *American Economic Review* 107, 3917–3946.

Chen, Joyce J., 2013, Identifying non-cooperative behavior among spouses: Child outcomes in migrant-sending households, *Journal of Development Economics* 100, 1–18.

Chen, Zhiqi, and Frances Woolley, 2001, A Cournot-Nash model of family decision making, *The Economic Journal* 111, 722–748.

Chiappori, Pierre-André, 1988, Rational household labor supply, *Econometrica: Journal of the Econometric Society* 56, 63–90.

Chiappori, Pierre-André, Bernard Fortin, and Guy Lacroix, 2002, Marriage market, divorce legislation, and household labor supply, *Journal of Political Economy* 110, 37–72.

Choukhmane, Taha, Lucas Goodman, and Cormac O'Dea, 2022, Efficiency in household decision-making: Evidence from the retirement savings of U.S. couples, MIT Working Paper, 31195.

Collins, Daryl, Jonathan Morduch, Stuart Rutherford, and Orlanda Ruthven, 2009, *Portfolios of the Poor: How the World's Poor Live on $2 a Day* (Princeton University Press).

Conlon, John J., Malavika Mani, Gautam Rao, Matthew W. Ridley, and Frank Schilbach, 2021, Learning in the household, NBER Working Paper 28844.

Dauphin, Anyck, Bernard Fortin, and Guy Lacroix, 2018, Is consumption efficiency within households falsifiable?, *Review of Economics of the Household* 16, 737–766.

De Mel, Suresh, David McKenzie, and Christopher Woodruff, 2009, Are women more credit constrained? Experimental evidence on gender and microenterprise returns, *American Economic Journal: Applied Economics* 1, 1–32.

Dercon, Stefan, and Pramila Krishnan, 2000, In sickness and in health: Risk sharing within households in rural Ethiopia, *Journal of Political Economy* 108, 688–727.

Donni, Olivier, and Pierre-André Chiappori, 2011, Nonunitary models of household behavior: a survey of the literature, in *Household Economic Behaviors*, 1–40 (Springer).

Duflo, Esther, 2003, Grandmothers and granddaughters: Old-age pensions and intrahousehold allocation in South Africa, *World Bank Economic Review* 17, 1–25.

Duflo, Esther, and Christopher Udry, 2004, Intrahousehold resource allocation in Cote d'Ivoire: Social norms, separate accounts and consumption choices, NBER Working Paper 10498.

Dupas, Pascaline, and Jonathan Robinson, 2013a, Savings constraints and microenterprise development: Evidence from a field experiment in Kenya, *American Economic Journal: Applied Economics* 5, 163–192.

Dupas, Pascaline, and Jonathan Robinson, 2013b, Why don't the poor save more? Evidence from health savings experiments, *American Economic Review* 103, 1138–1171.

D'Exelle, Ben, and Liz Ignowski, 2022, Intra-household resource allocation in rural Tanzania: Why women care about disclosure, *Journal of Development Studies* 58, 2021–2043.

Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Epstein, Larry G., and Stanley E. Zin, 1991, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis, *Journal of Political Economy* 99, 263–286.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue, 2002, Time discounting and time preference: A critical review, *Journal of Economic Literature* 40, 351–401.

Golden, Lawrence J., 1983, *Equitable Distribution of Property* (Shepard's/McGraw-Hill, New York).

Harris, Christopher, and David Laibson, 2001, Dynamic choices of hyperbolic consumers, *Econometrica* 69, 935–957.

Hess, Gregory D., 2004, Marriage and consumption insurance: What's love got to do with it?, *Journal of Political Economy* 112, 290–318.

Hoel, Jessica B., 2015, Heterogeneous households: A within-subject test of asymmetric information between spouses in Kenya, *Journal of Economic Behavior & Organization* 118, 123–135.

Jackson, Matthew O., and Leeat Yariv, 2015, Collective dynamic choice: The necessity of time inconsistency, *American Economic Journal: Microeconomics* 7, 150–178.

Jakiela, Pamela, and Owen Ozier, 2016, Does Africa need a rotten kin theorem? Experimental evidence from village economies, *Review of Economic Studies* 83, 231–268.

Kotlikoff, Laurence J., and Avia Spivak, 1981, The family as an incomplete annuities market, *Journal of Political Economy* 89, 372–91.

Laibson, David, 1997, Golden eggs and hyperbolic discounting, *Quarterly Journal of Economics* 112, 443–477.

Laibson, David, Andrea Repetto, and Jeremy Tobacman, 2003, A debt puzzle, in *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 228 (Princeton University Press).

Lancaster, Kelvin, 1973, The dynamic inefficiency of capitalism, *Journal of Political Economy* 81, 1092–1109.

Levhari, David, and Leonard J. Mirman, 1980, The great fish war: An example using a dynamic Cournot-Nash solution, *Bell Journal of Economics* 11, 322–334.

Lise, Jeremy, and Ken Yamada, 2019, Household sharing and commitment: Evidence from panel data on individual expenditures and time use, *Review of Economic Studies* 86, 2184–2219.

Lundberg, Shelly, and Robert A. Pollak, 1993, Separate spheres bargaining and the marriage market, *Journal of Political Economy* 101, 988–1010.

Lundberg, Shelly, and Robert A. Pollak, 2003, Efficiency in marriage, *Review of Economics of the Household* 1, 153–167.

Lundberg, Shelly, and Robert A. Pollak, 2007, The American family and family economics, *Journal of Economic Perspectives* 21, 3–26.

Lundberg, Shelly, Robert A. Pollak, and Terence J. Wales, 1997, Do husbands and wives pool their resources? Evidence from the United Kingdom Child Benefit, *Journal of Human Resources* 32, 463–480.

Mazzocco, Maurizio, 2004, Saving, risk sharing, and preferences for risk, *American Economic Review* 94, 1169–1182.

Mazzocco, Maurizio, 2007, Household intertemporal behaviour: A collective characterization and a test of commitment, *Review of Economic Studies* 74, 857–895.

Phipps, Shelley A., and Peter S. Burton, 1998, What's mine is yours? The influence of male and female incomes on patterns of household expenditure, *Economica* 65, 599–613.

Ratha, Dilip, Christian Eigen-Zucchi, and Sonia Plaza, 2016, *Migration and Remittances Factbook 2016* (World Bank Publications).

Ringdal, Charlotte, and Ingrid Hoem Sjursen, 2021, Household bargaining and spending on children: Experimental evidence from Tanzania, *Economica* 88, 430–455.

Robinson, Jonathan, 2012, Limited insurance within the household: Evidence from a field experiment in Kenya, *American Economic Journal: Applied Economics* 4, 140–164.

Rosenzweig, Mark R., and Oded Stark, 1989, Consumption smoothing, migration, and marriage: Evidence from rural India, *Journal of Political Economy* 97, 905–926.

Samuelson, Paul A., 1974, Complementarity: An essay on the 40th anniversary of the Hicks-Allen revolution in demand theory, *Journal of Economic Literature* 12, 1255–1289.

Schaner, Simone, 2015, Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings, *American Economic Journal: Applied Economics* 7, 135–174.

Schaner, Simone, 2017, The cost of convenience? Transaction costs, bargaining power, and savings account use in Kenya, *Journal of Human Resources* 52, 919–945.

Shapiro, Jesse M., 2005, Is there a daily discount rate? Evidence from the food stamp nutrition cycle, *Journal of Public Economics* 89, 303–325.

Smith, Tom W., Peter Marsden, Michael Hout, and Jibum Kim, 2011, General social surveys, 1972–2010, *Sponsored by National Science Foundation. NORC, Chicago: National Opinion Research Center* .

Thaler, Richard H., and Shlomo Benartzi, 2004, Save more tomorrow: Using behavioral economics to increase employee saving, *Journal of Political Economy* 112, S164–S187.

Thaler, Richard H., and Hersh M. Shefrin, 1981, An economic theory of self-control, *Journal of Political Economy* 89, 392–406.

Tornell, Aaron, and Andres Velasco, 1992, The tragedy of the commons and economic growth: Why does capital flow from poor to rich countries?, *Journal of Political Economy* 100, 1208–1231.

Treas, Judith, 1993, Money in the bank: Transaction costs and the economic organization of marriage, *American Sociological Review* 58, 723–734.

Turner, Brett R., 2005, *Equitable Distribution of Property* (Thomson/West, Eagan).

Udry, Christopher, 1996, Gender, agricultural production, and the theory of the household, *Journal of Political Economy* 104, 1010–1046.

Varian, Hal R., 1990, Goodness-of-fit in optimizing models, *Journal of Econometrics* 46, 125–140.

Voena, Alessandra, 2015, Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples?, *American Economic Review* 105, 2295–2332.

Vogler, Carolyn, Michaela Brockmann, and Richard D. Wiggins, 2006, Intimate relationships and changing patterns of money management at the beginning of the twenty-first century, *British Journal of Sociology* 57, 455–482.

Ward-Batts, Jennifer, 2008, Out of the wallet and into the purse: Using micro data to test income pooling, *Journal of Human Resources* 43, 325–351.