

8-2022

Three Essays on Consumption Smoothing

Logan James Miller
University of Arkansas, Fayetteville

Follow this and additional works at: <https://scholarworks.uark.edu/etd>



Part of the [Behavioral Economics Commons](#), and the [Macroeconomics Commons](#)

Citation

Miller, L. J. (2022). Three Essays on Consumption Smoothing. *Graduate Theses and Dissertations*
Retrieved from <https://scholarworks.uark.edu/etd/4617>

This Dissertation is brought to you for free and open access by ScholarWorks@UARK. It has been accepted for inclusion in Graduate Theses and Dissertations by an authorized administrator of ScholarWorks@UARK. For more information, please contact scholar@uark.edu, uarepos@uark.edu.

Three Essays on Consumption Smoothing

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy in Economics

by

Logan James Miller
University of Arkansas
Bachelor of Science in International Business, 2013
University of Arkansas
Master of Arts in Economics, 2019

August 2022
University of Arkansas

This dissertation is approved for recommendation to the Graduate Council.

Dongya Koh, Ph.D.
Dissertation Chair

Sherry Li, Ph.D.
Committee Member

Peter James McGee, Ph.D.
Committee Member

Abstract

Macroeconomic theory has established that consumption smoothing leads to higher standards of living. A stable consumption path can lead to more stability and less uncertainty between periods of high and low income. However, there is a wide body of literature that shows people do not consistently smooth their consumption when exposed to adverse income shocks. This dissertation uses experimental and empirical methods to better understand the obstacles people face when trying to smooth their consumption over time. It looks to understand the differences in pairs and individuals' ability to smooth consumption. It also explores how the household's level of income affects the level of consumption smoothing in response to an income shock and what constraints households face when trying to access the different insurance mechanisms.

Dedication

This work is dedicated to my parents, Chip and Kathy Miller, who gave everything for my education and taught me to adapt and overcome in every situation that life throws at me. I never would have made it to this point without them. I would also like to dedicate this work to my loving wife, Catherine Christine Miller, for her unending support and understanding throughout my PhD program.

Acknowledgements

For all of his guidance and support through this process, I want to acknowledge my advisor, Dr. Don Koh, who has made me into the researcher I am today.

I also want to acknowledge the guidance from my committee members, Dr. Sherry Li and Dr. Pete McGee, who were both instrumental in teaching me how to conduct good economic research and using it as a means to better understand the world around us; my cohort, Nathaniel Burke, Ahmad Shah Mobariz, and James Willbanks, for their constant feedback, encouragement, and never-ending support.

I also want to acknowledge my grandparents, Jim and Sharon Miller, for teaching me to be myself and to be proud of the man I have become.

Contents

1	Introduction	1
2	Chapter 1: Joint vs. Individual Performance in a Dynamic Choice Problem	3
	Introduction	3
	Literature Review	5
	Theory	7
	Experimental Design	9
	Results	12
	Textual Analysis	18
	Consumption Heuristics	24
	Conclusion	26
	References	30
	Appendices	34
3	Chapter 2: Market Completeness by Income Group	61
	Introduction	61
	Data	63
	Specification	65
	Results	67
	Conclusion	72
	References	73
	Appendices	75

4 Chapter 3: Consumption Insurance Constraints	82
Introduction	82
Empirical Evidence on Insurance Constraints	85
Data Construction	85
Estimation Methods	94
Empirical Results	98
Conclusion	112
References	112
Appendices	115
5 Conclusion	119

Introduction

The modern world is full complicated financial securities, commodities, and insurances available to investors and consumers alike. In 1954 Kenneth Arrow and Gerard Debreu introduced the concept of the Arrow-Debreu market, or complete market, to help provide some insight into such complicated financial structures. Complete market theory asserts that there is a market for every good where the concept of a “good” encapsulates any state, date, or environment in which a good can or may be consumed and sold. What this means for the consumers within a complete market is that they should be able to smooth their consumption no matter what the state of the world was by selling any assets, collecting on insurance, or using personal savings. In other words, any idiosyncratic income shock would not affect their consumption, consumption can only be affected by aggregate shocks across the entire economy. The three papers in this dissertation come together to better understand the obstacles people face when trying to smooth their consumption over time, and why the existing literature has only found evidence of partial insurance.

This dissertation is the result of trying to understand why households in one of the world’s most developed economies still only have partial insurance against adverse income shocks. This open-ended question led to the examination in the second paper in this work. I do a relatively standard evaluation of market completeness using Panel Survey of Income Dynamics (PSID) data, but I separate the sample into quartiles based on household income. I show that the bottom half of American earners struggle to smooth their consumption in some capacity. The question becomes whether or not this is due to characteristics endogenous to the household, or if it is the result of some form of constraint faced by the household. This finding and new line of questioning lead to the investigations contained in the first and third papers.

The first paper looks to strip away any complicating factors and asks a very simple question about household structures. Are there significant differences in the ability to con-

sumption smooth between individuals and pairs? This is similar to, but not the same as, the consumption decisions faced by married or single households. In order to better understand differences between households, it is necessary to bring the convoluting factors that the first paper abstracts away from back into the fold. The third paper looks to discern if households face barriers to using available insurance mechanisms. These barriers are directly related to household characteristics unique to each household in the data.

Logan Miller¹ Ryan Rholes²

1 Introduction

Modern macroeconomic theory typically models decision-makers as rational individuals capable of solving complex, dynamic choice problems without influence from, and independent of, other economic agents. Indeed, this representative agent assumption critically influences how monetary and fiscal policy operate in these models. Macroeconomists have taken this representative agent assumption seriously when designing experiments to test the micro-foundations of macroeconomic theory (Duffy, 2016).

However, using individuals as representative decision-makers in the lab may not fully capture the behavior modeled in macroeconomic theory. Among the macroeconomic models that do assume a representative household or firm, most do not explicitly rule out the notion of joint decision making within that household or firm. In fact, theory sometimes explicitly models this. For example, many theories include the assumption of multiple decision roles in the household, like the worker-shopper pair introduced by Lucas and Stokey (1983).

Evidence from Panel Study of Income Dynamics (PSID) data seems to support this distinction between individual and joint-decision households. The data show that among the lower income quartiles in America, married households better smooth their consumption in response to negative income shocks than do single households.³ An important question is why this is true. This may be due to structural differences (e.g. dual earners) but may also be due to the fact that married households form joint decisions.

This paper examines whether the latter explanation has any merit by revisiting dynamic

¹University of Arkansas:lmiller@walton.uark.edu

²University of Oxford: ryan.rholes@economics.ox.ac.uk

³We provide details about this suggestive evidence in Section 7.4 of the appendix.

optimization in the laboratory to investigate whether pairs forming joint decisions outperform individuals in a dynamic optimization task. We do this by having either individuals or pairs solve a finite-horizon dynamic optimization problem in a setting with stochastic income and that allows for both borrowing and saving. Because pairs and individuals face identical settings, structural differences cannot account for any treatment-level differences we observe.

We find that pairs significantly outperform individuals relative to the rational, representative benchmark. This is true whether we measure performance along the unconditionally-optimal or conditionally-optimal consumption paths. On average, pairs earn about 40% more than individuals after accounting for fixed show-up fees paid to all subjects.

Pairs in our experiment engage in unrestricted communication via a chat window to form joint decisions. Analysing this chat data provides valuable insight into how subjects in those treatments think through dynamic choice problems. Subjects almost exclusively frame discussions in terms of spending even though saving and borrowing are important components of earnings maximization. Further, subjects develop simple, invariant heuristics that can lead to persistent and compounding errors. This textual analysis corroborates our classification of individuals and pairs into different consumption heuristics. Further, we use textual analysis to classify each student's subjective outlook on debt using free-form answers provided in a post experiment survey-of-decisions. Though the majority of subjects view debt as inherently bad, we find that a non-trivial subset of subjects have a more nuanced view of debt. We provide suggestive evidence that these subjects outperform their counterparts in our optimization task.

Finally, we show that consumption heuristics are not stable over the lifecycle. As play progresses, the proportion of subjects employing relatively simple spending rules increases while the proportion of students using more complex rules decreases.

2 Literature Review

There is an extensive literature, thoroughly discussed in Duffy (2016) and summarized here, that studies the ability of individuals to solve dynamic stochastic optimization problems. Generally, subjects deviate considerably from the optimal consumption path.

We are not the first to use stochastic income to study dynamic optimization in the lab. Hey and Dardanoni (1988) show subjects fail to optimize in response to a stochastic income, a no-borrowing constraint, and a constant rate of return on savings. Carbone and Hey (2004) and Carbone (2006) simplify this design by eliminating discounting and by simplifying the stochastic income process and find these reductions in the complexity of the lifecycle problem do not move subjects to rational consumption path. Carbone and Infante (2015) study dynamic optimization under certainty, risk, and ambiguity and find that subjects significantly under consume when faced with ambiguity relative to risk and certainty. Carbone, Hey, and Neugebauer (2021) study consumption smoothing in a Lucas Tree model where subjects trade consumption claims via a long-lived asset, with an alternative solution, where agents can trade short-lived consumption claims between periods. They find the exchange economy with short-lived assets is more efficient in encouraging consumption smoothing.

Others have studied various types of learning in dynamic optimization by allowing individual decision-makers to interact in various capacities. Ballinger, Palumbo, and Wilcox (2003) provide evidence in support of inter-generational learning in the context of dynamic choice via a 60-period life-cycle problem under income uncertainty. The authors grouped subjects into three-member "families" and randomly assigned each family member to either the first, second, or third generation. Members of the first generation had no opportunity to learn. However, members of subsequent generations could both observe and communicate with members of the previous generation for several periods before beginning to make their own decisions. This generational transmission of information improves decisions of subsequent generations. Our study differs from theirs in that subjects in our Pairs treatments do not pass

along knowledge but instead work together to generate knowledge, and our Pairs subjects form joint decisions and share the payoff of this joint decision.

Brown et al. (2009) show that allowing for social learning improves the speed of own-learning compared to rates of own-learning from subjects in private-learning treatments. In contrast, Carbone and Duffy (2014) show that revealing the average level of past consumption causes subjects to deviate further from both the conditionally- and unconditionally optimal consumption path. Bao, Duffy, and Hommes (2013) show that pairing subjects together and having each subject either forecast or optimize leads to quicker convergence to the rational expectations equilibrium than does having a single subject perform both tasks. Duffy and Orland (2021) test a buffer stock model in the lab and show that imposing liquidity constraints does not increase savings but higher income variation does.

Ubiquitous across these previous studies is the use of individual decision-makers. However, there are also studies comparing the behavior of groups in macroeconomic settings. For example, Blinder and Morgan (2005) show that groups outperform individuals setting monetary policy to maintain to stabilize an experimental economy around inflation and employment targets. This finding was corroborated by Lombardelli et al. (2005) who also show that groups outperform individuals as policymakers because groups can strip out the effect of bad play in a given period, and because group members are able to share information and learn from each other's interest rate decisions. Similarly, Rholes and Petersen (2020) show in a learning-to-forecast experiment that aggregating over group expectations produces more stable inflation dynamics than do individual expectations.

Most closely related to our work are Carbone and Infante (2015), and Carbone, Georgalos, and Infante (2019), which both study differences between pairs and individuals in a dynamic optimization setting. Carbone and Infante (2015) conclude that stable pairs perform no differently than individuals in solving the life cycle problem once experienced and that pairs with rematching perform worse than individuals. We find the opposite – stable pairs in our

experiment consistently outperform individuals as planners, even after gaining experience. Carbone, Georgalos, and Infante (2019) compare group and individual performance in an optimization task while facing either risk or ambiguity and find that groups are better planners under ambiguity but individuals are better planners under risk. Because the support and distribution of the stochastic income process in our environment are known to our subjects, our setting best matches decisions under risk. Thus, our results again differ in that pairs in our environment consistently outperform individuals.

Finally, we also contribute to the extensive literature that studies differences between groups and individuals. Examples are Cooper and Kagel (2005) who find that teams play more strategically than individuals and Kugler, Bornstein, Kocher, and Sutter (2007) who have show groups are less trusting than individuals but are equally trustworthy. Charness and Sutter (2012) note that group choices better align with standard game-theoretic predictions and Kagel and McGee (2016) show that, with experience, teams cooperate more than individuals in prisoner’s dilemma games.

3 Theory

Subjects in both our Individuals and Pairs treatments maximize their discounted lifetime utility, subject to an intertemporal budget constraint:

$$\max \mathbb{E}_0 \sum_{t=1}^{t=T} \beta^t U(c_t) \tag{1}$$

$$s.t. \sum_{t=1}^{t=T} c_t \leq \sum_{t=1}^{t=T} w_t + a_0 \tag{2}$$

where c_t is consumption, a_0 is initial wealth, and w_t is an i.i.d. per-period stochastic income

with $w \sim U\{\underline{w}, \bar{w}\}$.⁴ Subjects in our experiment save freely and borrow up to \underline{w} in all but the final decision period. We denote saving and borrowing throughout as s_t .

We induce the quadratic utility function

$$U(c_t) = \phi c_t - \frac{1}{2} c_t^2. \quad (3)$$

This functional form is useful for several reasons. First, it allows subjects to consume zero in any period without incurring negative utility. Second, it is concave across the action space, which induces a consumption smoothing motive.⁵ Finally, combining this functional form with equations (1) and (2) above yields Hall's (1978) stochastic equation:

$$c_t = (1 - \kappa)\phi + \kappa\mathbb{E}_t c_{t+1} \quad (4)$$

where $\kappa \equiv \beta(1 + r)$. We set $\beta = 1$, $r = 0$ in order to reduce the complexity of our choice problem, which reduces Equation (4) to the consumption Euler equation:

$$c_t = \mathbb{E}_t c_{t+1}. \quad (5)$$

Solving by backward induction yields our unconditionally-optimal consumption path ⁶

$$c_{T-j} = \begin{cases} y_{T-j} + s_{T-j-1}, & j = 0 \\ \frac{j}{j+1}\mu + \frac{j-1}{j+1}(y_{t-j} + s_{t-j-1}), & j \in (1, 2, \dots, T-1) \end{cases}$$

⁴Income is drawn from a discrete uniform distribution so that per-period income is always an integer value.

⁵Restrictions on ϕ are such that, across the feasible action space, the first derivative of $u(c_t)$ is strictly positive and the second derivative is strictly negative. This means that subjects in our experiment can never consume beyond the bliss point regardless of how much wealth they accumulate.

⁶Notice that if $r > 0$ then per-period consumption is lower and per-period savings are higher in most periods. This might lead to behavior similar to that found in Carbone and Infante (2015).

This solution indicates that optimal consumption is a linear function of the mean of the income distribution, μ , and period wealth. Intuitively, subjects should focus less on the income distribution and more on wealth as the game nears completion. We plot the unconditionally-optimal consumption path alongside the income processes used in all experimental sessions in Figure 1. The unconditionally-optimal path is the same for all subjects because we hold the stochastic income processes constant across all subjects.

We also consider subjects' decisions relative to a conditionally-optimal level of consumption, \hat{c}_t^* , which accounts for past consumption errors by recalculating optimal consumption for each remaining period conditional on past mistakes.⁷

$$\hat{c}_t^* = c_t^* + \frac{(y_t - c_t^*) + s_{t-1}}{T - (t - 1)}, \forall t \in \{2, \dots, T - 1\} \quad (6)$$

4 Experimental Design

We use a simple 2×1 , between-subjects design built around a standard learning-to-optimize (LTO) framework where we compare the ability of Individuals and Pairs to solve two different twenty-period decision blocks of the finite-period smoothing problem outlined in Section 3. We set $\phi = 1,600$, $\bar{w} = 80$, $\underline{w} = 60$, $\beta = 1$, $r = 0$ for all sessions. We eliminated discounting and rates of return to reduce the complexity of the choice problem. We choose ϕ so that subjects could never consume beyond the bliss point for any possible income draw.

The consumption smoothing motive in our setting comes from the concavity of the induced quadratic utility function. Subjects spent, saved, and borrowed per-period income, allotted as experimental credits (ECs) according to two pre-drawn stochastic income processes. Importantly, subjects received consumption points in each period equivalent the consumption

⁷We do not plot the conditionally-optimal path here since it depends on individual deviations from the unconditionally-optimal consumption path.

Unconditional Optimal Consumption and Income

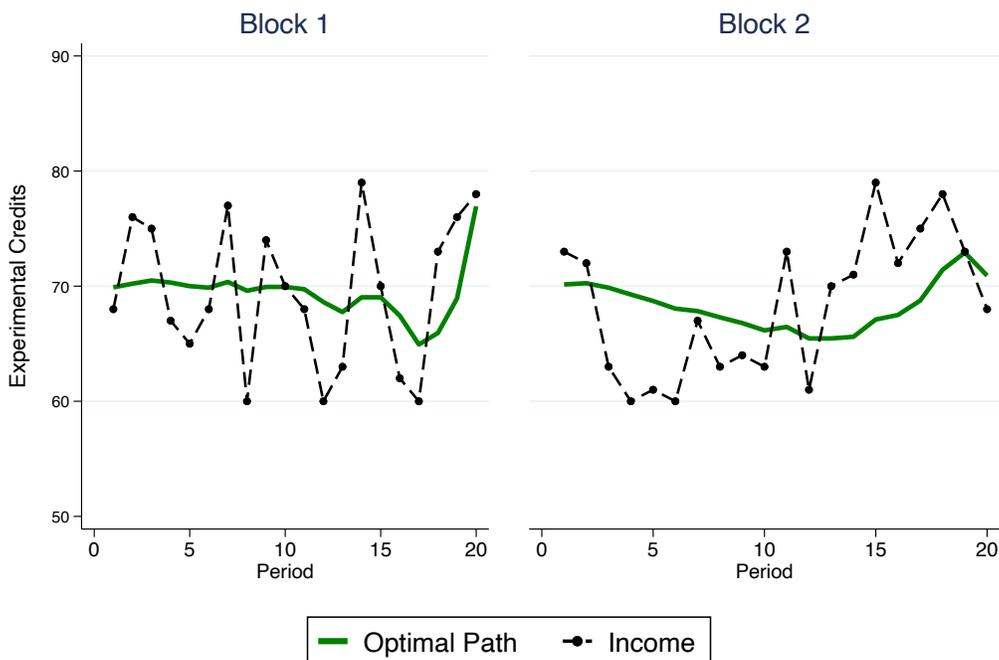


Figure 1. This figure shows the unconditionally-optimal consumption path for decision blocks 1 and 2 of all experimental sessions. The graph also includes the pre-drawn stochastic income processes used for blocks 1 and 2 in all experimental sessions.

utility resulting from their consumption decision in that period. Using pre-drawn income processes allowed us to hold the income process constant across treatments for each decision block.

Sessions began with a 6-question, individual-level Cognitive Reflection Test (CRT) introduced by (Frederick, 2005), also adopting questions from the Cognitive Reflection Test-Long (CRT-L) developed by (Primi et al., 2016). Subjects had 90 seconds to answer each CRT question and earned \$.25 for each correct answer. We followed this with an individual-level Eckel-Grossman test of risk preferences (Eckel and Grossman, 2002). Following these two tasks, subjects in the Individuals treatment worked alone to solve both lifecycle problems. For the Pairs treatment, we randomly matched subjects into stable pairs and allowed them to engage in unrestricted chat to solve the lifecycle problems. Pairs had to reach a consensus to enter a consumption decision. Subjects were not time-constrained when solving the

lifecycle problem in either treatment. All subjects were students recruited at the University of Arkansas.⁸ We ended each session with a demographic survey that also included a survey of attitudes toward debt and saving.

Instructions provided detailed information about the utility function, income process, lifecycle duration, and borrowing and saving so that they had sufficient information to fully solve the lifecycle problem. Further, we provided subjects with information about their per-period income, and their current bank account balance to help them keep track of their borrowing/savings. We also provided subjects with a consumption smoothing tool to reduce the cognitive complexity of the problem. To use the tool, subjects could propose a hypothetical level of consumption and learn the corresponding levels of utility (we called these consumption points in the game), savings or debt, and the marginal utility of consumption (we called this the ‘marginal increase’ in the game). Subjects could use this tool as many or as few times as desired. We provide an example of the decision screen for an individual in Figure 8 and for pairs in Figure 9 in Section 7.1 of the Appendix.

	Individuals	Pairs
Instructions & Comprehension Quiz	Individual	Individual
Cognitive Reflection Test	Individual	Individual
Eckel-Grossman Risk Assessment	Individual	Individual
Two rounds of decision-making	Individual	Joint
Eckel-Grossman Risk Assessment	Individual	Joint
Demographics & Survey of Decisions	Individual	Individual

Table 1. This table describes the order of events when conducting a session and indicates whether the task was completed individually (Individual) or in a pair (Joint).

For the Individuals treatment, we converted consumption points to U.S. dollars at 50 points per \$1. For the Pairs treatment, we converted consumption points at 25 points per \$1.⁹ This conversion scheme holds subject-level incentives constant across treatments. Subjects also received a \$10 show-up fee. We conducted all sessions at the University of Arkansas’s

⁸IRB protocol #: 1908210566

⁹We rounded payoffs to the neat highest point. For example, a score of 51.4 points would earn an individual \$1.04.

Behavioral Business Research Laboratory. We have 26 observations in the Individuals treatment and 27 observations in the Pairs treatment for a total of 80 unique subjects.¹⁰ We implemented our experiment using zTree (Fischbacher, 2007).

5 Results

We show treatment-level mean absolute unconditional and conditional consumption errors by period and treatment in panels (a) and (c) of Figure 2, respectively.¹¹ We also show the difference in treatment-level mean absolute unconditional and conditional consumption errors in panels (b) and (d) of the same figure. For panel (b) and (d), observations above the x-axis denote an instance where Pairs outperformed Individuals. Visually, it appears that pairs outperform individuals in solving the finite life-cycle problem along both optimal consumption paths (we also show this using medians rather than averages in Figure 7, which is located in Section 7.1).

Also worth noting in panel (c) of Figure 2 is the gradual buildup of absolute conditional errors toward the end of each decision block. This is likely due to the adoption of simple consumption heuristics primarily focused on income (we discuss this in more detail in Section 5.2). Because the conditionally-optimal path assumes subjects will account for previous mistakes in remaining decisions, these heuristics are increasingly penalized when moving along the conditionally optimal consumption path.

We confirm this results in two ways. First, we use a two-sample t-test to assess differences in mean absolute errors using all decisions from both treatments under the null hypothesis that the difference in mean absolute consumption errors along both consumption paths is equal across treatments. We report directional p-values, based on Figure 2, from this test

¹⁰We note our sample size is a bit smaller than the general rule-of-thumb of $N=40$. However, we feel this is okay given that our results are both stark and highly-significant despite our sample sizes.

¹¹We show average consumption and per-period consumption heterogeneity in Figure 10 located in Section 7.2 of the Appendix. We also show granular consumption data in Section 7.1 of the Appendix.

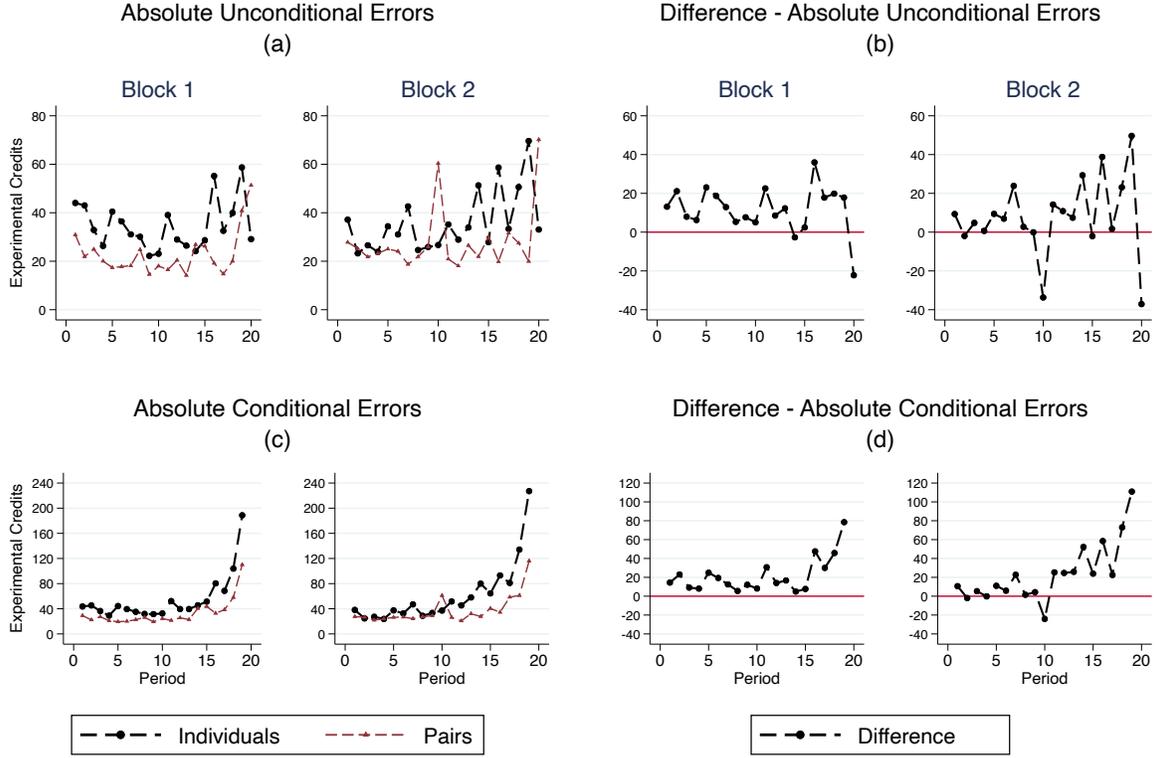


Figure 2. This figure depicts treatment-level average absolute consumption errors (panels (a) and (c)) and differences between the absolute consumption errors by treatment (panels (b) and (d)). For panels (b) and (d), values above zero indicate that pairs outperformed individuals in that period.

for both decision blocks independently and pooled in Table 2. These tests confirm that Pairs outperform Individuals along both consumption paths in both decision blocks and when pooling data.

Because observations in our experiment are potentially serially correlated, we perform a second set of statistical tests on data collapsed to the observation level. To do this, we calculate the root-mean-square deviation (RMSD) of unconditional and conditional consumption errors for each subject (or pair of subjects). Specifically, we calculate $RMSD_i^U = \sqrt{\frac{\sum_{t=1}^{t=T} (c_{i,t} - c_t^*)^2}{T}}$ for unconditional absolute consumption errors and then simply replace c_t^* with $\hat{c}_{i,t}^*$ from Equation (6) for $RMSD_i^C$. We then test for mean differences in the RMSD by treatment using a two-sample t-test. We report p-values from this exercise by decision block

	Summary of Statistical Tests					
	Unconditional Absolute Error			Conditional Absolute Error		
	Block 1	Block 2	Pooled	Block 1	Block 2	Pooled
Mean Error - I	34.62	35.93	35.27	54.70	61.43	58.02
Mean Error - P	22.94	28.03	25.49	32.96	37.55	35.25
2-Sample t-test	0.00	0.02	0.00	0.00	0.00	0.00
Mean RMSD - I	36.29	76.50	54.10	54.47	118.70	83.93
Mean RMSD - P	24.04	55.59	39.31	31.68	69.72	49.30
2-Sample t-test	0.03	0.09	0.09	0.01	0.02	0.02

Table 2. this table reports p-values (rows 3,6) from a series of two-sample t-tests for differences in mean absolute consumption errors and mean RMSD of consumption errors across treatments.

and using pooled data in Table 2. Again, results indicate that Pairs outperform Individuals along both consumption paths in both decision blocks and when pooling data.

We next estimate a series of mixed-effects regression models of the form

$$Y_{i,t} = \beta_0 + \beta_{i,t}X_{i,t} + \mu_i + \epsilon_{i,t}. \quad (7)$$

where our outcomes of interest, $Y_{i,t}$ are either absolute unconditional consumption errors or absolute conditional consumption errors and $X_{i,t}$ contains a set of independent variables described below. Both outcomes are measured as absolute deviations from the respective optimal path in terms of ECs. We estimate Equation (7) while restricting our data by treatment and also for our full data sample.¹² We report results of these estimation exercises in Table 3.

Columns 2-5 of this table report results using unconditional absolute consumption errors while columns 5-9 report results using conditional absolute consumption errors. Columns labeled Individuals use only data from our Individuals treatment, Pairs use only data from

¹²Though random effects models are common in this literature (examples are Carbone and Duffy (2014), Ballinger et al. (2011)), a Hausman test indicates the need to control for potential fixed effects, which perhaps result from static session effects (Fr chet te, 2012).

our Pairs treatment, and Pooled uses data from both. We can compare coefficient estimates in the Individuals column to its counterpart in the Pairs column for a given error type to learn about how equivalent information differently influences consumption errors for pairs and individuals. MaxCRT refers to the highest CRT score within a pair and MinCRT refers to the lowest CRT score within a pair. For individuals, MaxCRT simply refers to the individual's CRT score. Wealth refers to subjects' accumulated savings, Income represents current-period income (in ECs), and Block is a dummy variable takes a value of 1 if data comes from the second finite lifecycle in a session.

	Regression Results - Mixed Effects Estimations							
	Unconditional Absolute Error				Conditional Absolute Error			
	Individuals	Pairs	Pooled	Pooled	Individuals	Pairs	Pooled	Pooled
Pairs			-9.78*** (2.32)	-18.00*** (3.17)			-22.77*** (3.50)	-23.07*** (4.01)
MaxCRT	-2.41 (5.88)	-5.56*** (1.01)		-4.37*** (0.97)	-3.42 (7.87)	-7.21*** (1.20)		-4.61*** (1.24)
MinCRT		-2.70* (1.55)		-3.10** (1.46)		-3.55* (1.89)		-4.88*** (1.85)
Wealth	0.03** (0.01)	0.010 (0.01)		0.03*** (0.01)	0.24*** (0.04)	0.06*** (0.02)		0.20*** (0.03)
Income	0.43 (0.30)	0.35** (0.18)		0.40** (0.17)	2.61*** (0.43)	1.32*** (0.21)		2.03*** (0.25)
Block	1.91 (4.06)	5.65*** (2.16)		3.99* (2.29)	9.28* (5.56)	6.34** (2.51)		8.84*** (3.10)
Constant	4.14*** (0.16)	3.62*** (0.14)	3.97*** (0.12)	3.95*** (0.12)	4.45*** (0.12)	3.71*** (0.13)	4.35*** (0.10)	4.22*** (0.10)
N	1040	1080	2120	2120	988	1026	2014	2014

Table 3. This table shows of mixed effects regressions. Column 1 lists variable names, where maxCRT (min) refers to the highest (lowest) CRT score in the pair. For individuals, maxCRT refer to the individual's CRT score. Columns 2 thru 5 report results using the absolute unconditional consumption error as the dependent variable and columns 6 thru 9 report results using the absolute conditional consumption error as the dependent variable. Columns labeled as 'Individuals' or 'Pairs' use only the data from their corresponding treatment. Columns labeled 'Pooled' use all data. We report robust standard errors in parentheses. Note that the difference in N arises because there is no conditional error in the first period of either decision block.

We start by comparing between Individuals and Pairs. First, we note that CRT score significantly impacts neither unconditional nor conditional absolute errors in our Individuals

treatment. However, CRT score is a highly significant predictor of performance for Pairs regardless of outcome. In particular, increasing the maximum CRT score within a pair leads to decisions that are, on average, about 5.5 to 7 ECs closer to the optimal path. Interestingly, increasing the maximum CRT score improves performance along both pathways by about twice as much as increasing the minimum CRT score by the same amount.

A common finding in the LTO literature is that consumption overreacts to income. Our results align with this finding. We see that larger per-period income draws lead to larger conditional and unconditional errors for Pairs and larger conditional errors for Individuals. Interestingly, we see that the conditional errors in the Individuals treatment react twice as strongly as conditional errors in our Pairs treatment.

Additionally, we see that consumption errors are increasing in the accumulation of savings. Focusing on conditional errors, we see that consumption errors from Individuals are four times as large as consumption errors for Pairs.

Our primary coefficient of interest in our Pooled columns is that associated with our indicator variable for Pairs decisions. As expected based on results in Table 2, we see that joint decisions from our Pairs treatment are significantly closer to optimal along both the unconditionally- and conditionally-optimal paths.

Regression results indicate that subjects making joint decisions in our Pairs treatment were, on average, more than 15 ECs closer to conditionally-optimal, and 9 ECs closer to the unconditionally-optimal, level of consumption. Additionally, we see that absolute consumption decisions increase whenever accumulated wealth increases and with higher income. Both effects are considerably larger along the conditional than along the unconditional path.

Since subjects in our experiment are concerned with earnings maximization, it perhaps makes the most sense to consider average earnings differences between subjects in our Pairs and Individuals treatments. Subjects in the Individuals treatment earned \$20.20 on average,

while subjects in the Pairs treatment earned an average of \$24.34. Because we are concerned with earnings differences that result from differences in decisions, we subtract from these averages the fixed show-up fee of 10. We see that subjects in the Pairs treatment earned approximately $\frac{\$14.34 - \$10.20}{\$10.20} = 40.59\%$ more, on average, than subjects in the Individuals treatment. Without making this adjustment, earnings differences are still quite large: Pairs earn approximately 20% more than Individuals. A 2-sample t-test confirms this difference is significant at the 1% level.

We also quantify differences between the Pairs and Individuals treatments by comparing the performance of pairs to synthetic pairs formed using subjects in our Individuals treatment. Our interest is in how much we must improve the performance of these synthetic pairs before their decisions are no longer statistically distinguishable from real pairs at a 10% level of significance. To do this, we randomly match individuals into synthetic pairs and assume each pair consumed in a given period the average of what the two individuals consumed in that period. We repeat this matching process for all possible pairings and average results over all observations.¹³ We find that we can reduce the conditional consumption error of synthetic pairs by approximately 37%, on average, before the performance of real and synthetic pairs becomes indistinguishable. In level terms, this reduces the average conditional consumption error of synthetic pairs from 57.94 to approximately 36.5 experimental credits.

There are at least two possible explanations for the superior performance of subjects in our Pairs treatment. First, subjects in a pair are able to discuss strategies and exchange ideas, and sometimes balance preferences in order to form a joint decision. Second, subjects in the Pairs treatments may have to more carefully consider the optimization problem in order to communicate with an assigned partner. Thus, one could question if pairs do better because they are making a joint decision or instead because they are forced to more carefully consider their spending, saving, and borrowing decisions.

¹³With 26 individuals, we have $c(26, 2) = 325$ possible pairings.

We attempt to distinguish between these two explanations by implementing a third treatment, which we call the Ledger treatment. The Ledger treatment is identical to the Individuals treatment, except that subjects in the Ledger treatment have access to the same chat window as do subjects in Pairs treatments, which they can use as a sort of journal to articulate the logic of their individual decisions. In order to be consistent with the Pairs treatment, we neither require subjects in the Ledger treatment to use the ledger nor do we allow them access to ledger entries from previous periods.

A two-sided t-test shows that mean absolute consumption errors along both consumption paths (unconditional and conditional) are not significantly different across Ledger and Individuals treatments. However, the same test shows that decisions from subjects in the Pairs and Ledger treatments are highly significantly different. We take this as evidence that superlative Pairs performance results from the process of joint decision making rather than more carefully considering the optimization problem.

5.1 Textual Analysis

Because subjects in the Pairs treatment of our experiment engaged in unrestricted chat to make joint decisions, we are able to use textual analysis to gain deeper insight into how subjects frame the dynamic optimization problem and develop heuristics.

Following Cooper and Kagel (2005), we establish a set of categories we use to classify the language used by subjects in our Pairs treatment, which we describe in Table 4. These categories are neither exhaustive nor mutually exclusive. Rather, the categories are complementary, which allows for some nuance in classification despite the binary coding system. We trained two research assistants (RAs) who then worked independently to classify language into our pre-selected categories. As an example, if a pair discussed how to allocate resources in terms of spending but never in terms of savings, the research assistants would likely code 'Discuss Savings' as a zero and 'Discuss Spending' as a one.

We use these codings from our RAs to construct a measure that captures, on average, how often chat aligns with a given category. We construct this measure by first summing over all periods, sessions, and pairs for both research assistants and then dividing this sum by two times the total number of periods times the total number of pairs. Thus, we report a number bounded between zero and one where a value of one means all pairs used language compatible with that category in all periods. Anything less than one means that there is at least one pair who does not use that language in at least one period. To measure classification agreement, we divide the number of times the RA’s disagree by the number of opportunities to code a discussion category, subtract this from one, and then convert to percentage terms. We report both measures in Table 4.

Category	Description	Mean	Agreement(%)
Discuss Saving	Pair frames discussion in terms of saving	.070	97.69
Discuss Spending	Pair frames discussion in terms of spending	.873	97.87
Save More	Someone proposes saving more relative to previous suggestion/period	.025	99.91
Save Less	Someone suggests saving less relative to previous suggestion/period	.004	100.00
Spend More	Someone proposes spending more relative to previous suggestion/period	.054	99.35
Spend Less	Someone proposes spending less relative to previous suggestion/period	.046	97.41
Nominal Target	Pair discusses a nominal target (i.e. consumption points)	.091	93.70
Real Target	Pair discusses a real target (i.e. total dollar earnings)	.017	99.72
Marginal Target	Pair targets a 'marginal increase' target	.048	100.00
Savings Target	Pair tries to maintain a certain amount of savings	.006	99.35
Period Earnings Target	Pair discusses a per-period earnings target	.014	97.41
Total Earnings Target	Pair discusses a lifetime earnings target	.011	98.43
Proportional Spender	Pair discusses spending a proportion of income or total wealth	.053	99.24
Borrow	Pair discusses borrowing against future income	.045	99.63
Constant Spending	Pair discusses spending a constant amount	.038	97.13
Save & Binge	Pair discusses saving heavily to spend a large lump sum later	.044	99.91

Table 4. This table provides information regarding our textual analysis. The first two columns define the categories used by two research assistants (RAs) who worked independently to classify the language used by Pairs when forming joint decisions. The third column provides a measure of how frequently Pairs used language consistent with each category. The fourth column provides a measure of the level of classification agreement between our two RAs. We construct values in column three by summing over all periods, sessions, and pairs for both RAs, and dividing this sum by two times the total number of periods times the total number of pairs. We construct our agreement measure by dividing the number of times the RAs disagree about a given classification by the number of opportunities to code a discussion category, subtracting this from one, and then converting to percentage terms.

Notice in Table 4 the relatively high frequency of the "Discuss Spending" category (87% of Pairs interactions), which indicates that pairs mostly frame discussions around spending rather than saving or borrowing. Though subjects must spend credits to earn money, the

stochastic income process, coupled with the consumption smoothing motive, makes saving and borrowing important components of earnings maximization. We also see that subjects, explicitly or implicitly, discuss spending strategies that fix consumption either in levels or as a proportion of wealth or income. This aligns with Figure 4, which shows that about half of our Pairs use a constant marginal propensity to consume heuristic.

These sorts of simple heuristics greatly reduce the cognitive load of the optimization task but might fail subjects whenever saving or borrowing is necessary for optimization. For example, Pairs spending a fixed proportion of the per-period endowment would not borrow whenever necessary to spend at the unconditionally- or conditionally-optimal level. This aligns with Carbone and Hey (2001) and Hey and Knoll (2011) who find that subjects are more likely to develop simple decision criteria and adopt strategies aimed at reducing the cognitive complexity of the choice task.

The tendency of Pairs to develop simple heuristics leads to considerable under borrowing in our experiments. However, we do not see in our chat data that Pairs openly express disdain for borrowing. Thus, under borrowing may result from subjects developing simple heuristics (i.e. proportional spending rules) that overlook borrowing and not from the fact that Pairs are actively averse to debt.

To better understand why Pairs do not discuss borrowing, we evaluate responses to two questions we included in our post-experiment survey-of-decisions that asked subjects to provide their subjective outlook on debt and savings.¹⁴ We find that about 60% of all subjects (Bad) view debt as expressly bad with the remaining subjects (Nuanced) taking a more nuanced view or debt as either good or sometimes good. These subjects argue that the goodness or badness of debt depends on who benefits from the debt, the magnitude of debt, and how one uses debt.

Table 5 evaluates the impact of debt outlook on performance in our optimization task. We

¹⁴We include the full survey in Section 7.7.

restrict our sample to Individuals and compare the RMSD of consumption errors along both the unconditionally- and conditionally-optimal paths by block using an independent-samples t-test. Results suggest that subjects with a nuanced outlook (row 2) on debt outperform subjects who view debt as strictly bad (row 1) in both blocks along both consumption paths. We report p-values for each block independently and consumption path in row 3.

	Impact of Debt Outlook on Performance			
	Unconditional		Conditional	
	Block 1	Block 2	Block 1	Block 2
Bad	44.17	58.82	71.72	103.36
Nuanced	24.99	36.57	34.15	51.4
2-Sample t-test	.12	.18	.08	.12

Table 5. This table provides suggestive evidence of the impact of subjective debt outlook on performance in our optimization task. We do this by comparing the RMSD of both absolute unconditional and conditional errors across subject debt outlooks for subjects in our Individuals treatment. We have a total of 16 observations split evenly between the two debt outlooks. Unfortunately, we lost survey responses for nine subjects from our first Individuals treatment session due to a network failure.

This finding aligns with Meissner (2016) and Ahrens, Bosch-Rosa, and Meissner (2022), which both demonstrate that individuals perform worse when solving dynamic optimization problems that require borrowing relative to saving.

We also note that the adoption of these simple time- and wealth-invariant consumption heuristics might explain the upward trend in conditionally-optimal consumption errors that we do not see in the unconditionally-optimal consumption errors. This is because a heuristic that leads to an absolute error in one period will, on average, lead to a similar absolute error in a later period. The invariant nature of the heuristic could prevent subjects from avoiding current-period errors and adjusting for past errors.

We also observe that subjects more often frame discussions in nominal rather than real terms. This is not surprising, given that our pairs tend toward simple heuristics that reduce the complexity of dynamic choice. Since nominal and real earnings are isomorphic, it might

be the case that subjects prefer nominal framing because it avoids the added complexity of real framing. This aligns with Petersen and Winn (2014), who find that nominal inertia arising in a choice task results from cognitive complexity and that money illusion exerts only second-order effects in the same task.

Finally, we see that our pairs discuss saving and binging as a strategy with surprising frequency. It is easy to assume that such behavior, first documented by Noussair and Matheny (2000), is reactionary since it demonstrates a misunderstanding of the consumption smoothing motive. However, we see here that this behavior can be thoughtful, planned, and forward-looking.

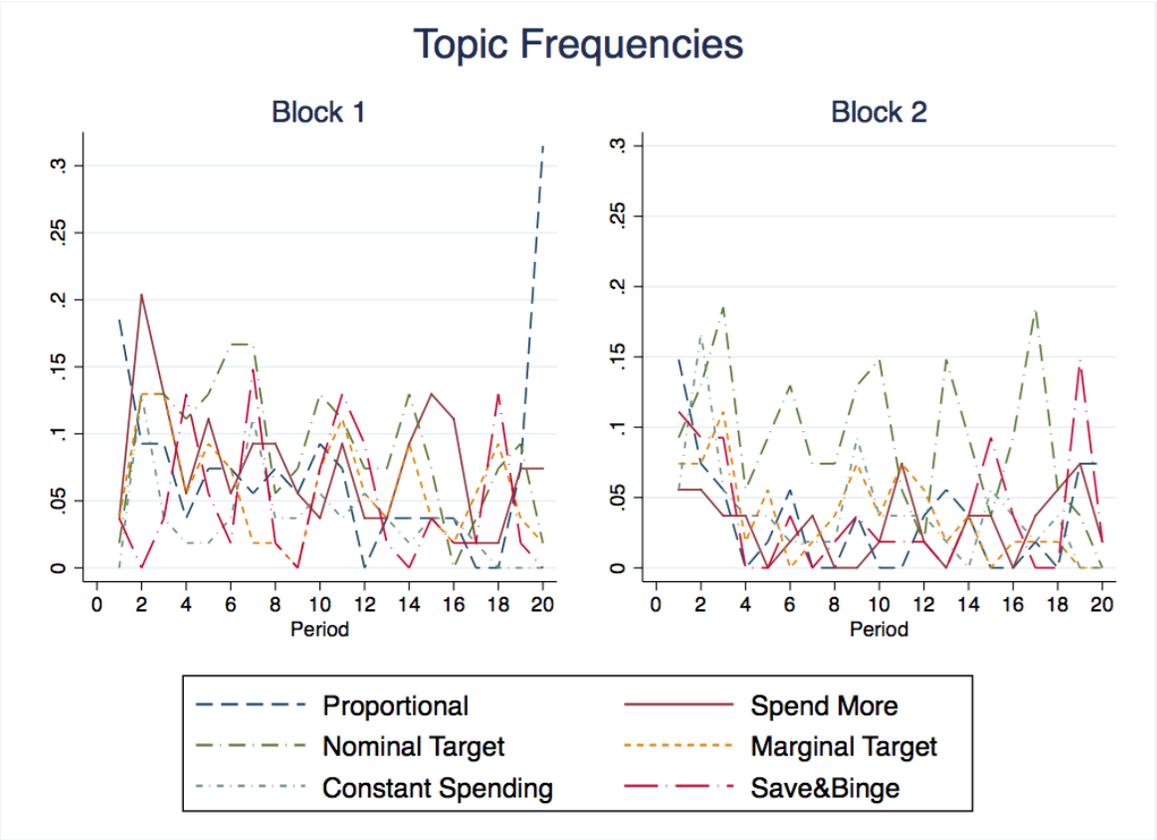


Figure 3. This graph plots, by period and block, the six language categorizations used most frequently by subjects in our Pairs treatment to reach a joint decision. The graph includes only those categories discussed by more than one pair within a single period at least once during the experiment. We construct this frequency measure by summing observations for a given category for all pairs within a period and by using observations from both RAs. We then divide this sum by two times the number of pairs.

Next, we consider textual classifications by period, depicted in Figure 3, in order to better understand how communication evolves. To do this, we focus on six thematic categories and eliminate categories not discussed by two or more pairs at least once in a single period. The proportional strategy indicates that a pair discussed spending or saving a proportion of their per-period endowment or accumulated wealth. The nominal and marginal strategies indicate a pair uses a nominal target or a marginal target to guide its consumption decision. The constant spending strategy emerges when pairs discuss spending a fixed level of EC's in each period. Finally, the save and binge strategy involves pairs saving EC's with the intention of 'binging' a large amount of EC's later in a single period. Similar to what we did in Table 4, we first sum the occurrences of a given category for all pairs and both research assistants and divide this number by two times the number of pairs. We then plot this number for each of our six categories by period.

The cyclical pattern of these frequencies suggests that not all pairs discuss a strategy in all periods. Instead, pairs discuss a strategy, follow it for some time, and then reaffirm or discuss the strategy again after a few additional periods. Second, we note that the frequency of discussion for all strategies, excluding the nominal target strategy, falls in block 2 relative to block 1, which could indicate that pairs settle into a heuristic as they gain experience. In particular, we see that pairs gradually think less about the marginal benefit of consumption when making decisions. Finally, we note that the discussion of proportional strategies spikes at the end of block 1, but not block 2. This is likely because pairs realized at the end of block 1 that they needed to spend all remaining savings. The absence of this same spike at the end of block 2 matches the marked decrease in bingeing behavior at the end of block 2 relative to block 1.

5.2 Consumption Heuristics

We now consider the heuristics used by individuals and pairs to make consumption decisions. To do this, we construct a set of 5 heuristics that may possibly describe consumption decisions in our experiment (See Carbone (2005) and Tasneem and Engle-Warnick (2018) for other examples of consumption heuristics).

Table 6. Forecasting heuristics

Model	Heuristic Name	Abbreviation	Model
H1	Hand-to-mouth	H-to-M	$C_t = Y_t$
H2	Unconditional Optimizer	U. Opt.	$C_t = C_t^*$
H3	Conditional Optimizer	C. Opt.	$C_t = C_t^* + \frac{(Y_t - C_t^*) + S_{t-1}}{T - (t-1)}$
H4	Constant Spending	ConSpend	$C_t = C_{t-1} = \dots = C_1$
H5	Constant M.P.C.	ConMPC	$C_t = Y_t \frac{\gamma}{30}, \gamma = \{1, 2, 3, \dots, 29\}$

Consumption heuristics.

H1 assumes that a subject consumes all of her income in each period. This is equivalent to having a fixed marginal propensity to consume (MPC) of 1 in each period. A real-world equivalent is an individual or family that lives paycheck-to-paycheck. H2 assumes that subjects optimize perfectly along the unconditionally-optimal path. This heuristic captures the behavior of a fully rational agent in the context of our finite lifecycle problem. H3 assumes that subjects optimize along the conditionally-optimal path. H4 supposes that a subject spends a constant value in each period regardless of income. H5 assumes that a subject i (or pair i) spends a fixed proportion α_i of income in each period. Note that Hand-to-Mouth would be equivalent to Constant MPC whenever $n = 30$ since this yields $MPC = 1$.

For each period, we calculate what a subject i (or pair of subjects i) would consume according to each consumption heuristic, $C_{i,t}^H$ and the corresponding error $C_{i,t} - C_{i,t}^H$. We then calculate the RMSD for each heuristic for each subject (or pair of subjects) as $RMSD_i^H =$

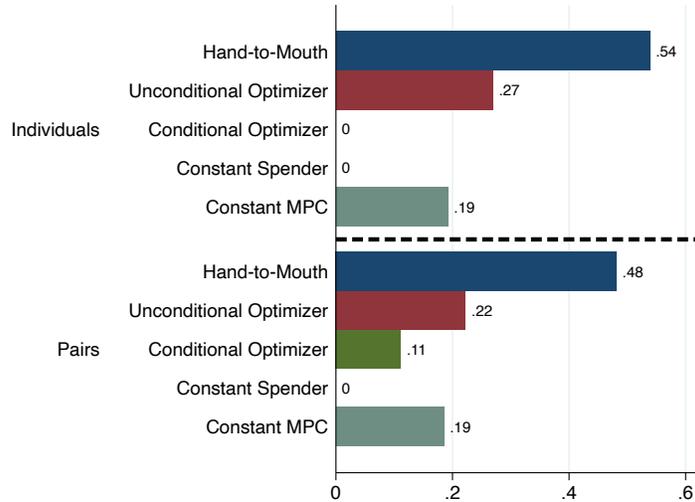


Figure 4. Consumption heuristics for all decision periods.

$\sqrt{\frac{\sum_{t=1}^{t=T} (C_{i,t} - C_{i,t}^H)^2}{T}}$. We then classify a subject (or pair) to whichever heuristic produces the smallest RMSD.

We show results from this classification exercise using all periods from both decision blocks in Figure 4. First, we note the relatively large proportion of unconditional optimizers – balanced across treatments – in both the Pairs and Individuals treatments. This is possibly due to the provision of the consumption calculator in our experiment, which reduces the complexity of the optimization problem.

Additionally, we see that in both treatments the majority of subjects are using some version of proportional spending, where many have an MPC that is either close to or equal to one. This aligns with results from our textual analysis section where we find that subjects in our Pairs treatment typically frame decisions in terms of spending and develop consumption heuristics based on proportional spending, which aligns with results from Carbone (2005).

We also note that a meaningful proportion of Pairs – and no Individuals – classify as conditional optimizers. This aligns with our results that Pairs significantly outperform Individuals along the conditionally-optimal consumption path.

We also consider whether and how heuristics change as subjects progress through a lifecycle. To do this, we follow the same classification exercise described above but classify subjects (or pairs of subjects) into a heuristic in five-period intervals. We show results from this exercise in Figure 5, where several interesting patterns emerge.

First, we see that most subjects initially adhere to some variant of ConMPC but that the proportion of subjects using this heuristic decreases as the lifecycle progresses. In both treatments the proportion of ConMPC starts around 60% for block 1 and around 40% in block 2 and decreases to about 20% - 30% in both blocks of both treatments. Conversely, we see that the proportion of subjects adhering to H-to-M increases with time. One possibility is that ConMPC decreases because some ConMPC subjects shift their MPC upward such they they become H-to-M subjects. Though both heuristics decrease the complexity of the choice problem relative to calculating the optimal consumption path, H-to-M consumption removes a layer of complexity from ConMPC since subjects don't need to calculate consumption as a proportion of income.

Second, we note that the proportion of U. Opt. subjects is relatively stable in Individual treatments but less so in Pairs treatments, where the proportion peaks at roughly 40% in block 1 and experiences a low of near 0% in block 2.

Finally, we see that the proportion of C. Opt. subjects is relatively stable at near 0% in our Individuals treatment but increases over both decision blocks in our Pairs treatment. This suggests that at least some subjects in our Pairs treatment may begin to converge to optimal play as the life cycle progresses.

6 Conclusion

This paper revisits the learning-to-optimize literature to study the relative ability of Pairs and Individuals to solve a finite-period, dynamic optimization problem. We find that

joint decision making leads Pairs to significantly outperform Individuals along both the unconditionally- and conditionally-optimal consumption paths. This performance gap, on average, leads to subjects in our Pairs treatment earning about 40% more than subjects in our Individuals treatment.

This demonstrates convincingly that simple household differences – joint decision making vs. individual decision making – can lead to systematic differences in budgetary decisions and, as a consequence, systematic differences in welfare. Though we abstract considerably from the complexity of the real world, our experimental design sheds light on why we observe in observational data that married households in America’s bottom income quartiles better smooth negative income shocks than do single households. We provide suggestive evidence of this in Section 7.4.

This might be the result of structural differences across household types (dual earners, etc.). However, it may be at least partially driven by the fact that households are forming either joint or individual decisions. We show in our experiment that Pairs outperform Individuals even in the absence of structural differences, which supports the idea that joint decision making helps explain real-world differences in how these types of households smooth negative income shocks. This is especially true for lower-income households. A possible implication is that increasing access to financial and budgetary planning services might be a reasonably cheap and affordable way to increase welfare for lower-income, single households.

We use textual analysis from Pairs chat data and from a post-experiment survey-of-decisions to try and understand why we observe these performance differences and also how people approach solving dynamic optimization problems.

Chat data suggests that Pairs often negotiate joint consumption decisions by updating toward one another. This is corroborated by responses to our survey-of-decisions question “What was your strategy for overcoming disagreements?” where the overwhelming majority of pairs indicated that they used mutual compromise to reach a joint decision. This suggests that

at least one benefit of forming joint decisions is a sort of 'wisdom of the crowd' effect. This moderation of more extreme decisions reduces boom-bust consumption cycles leading to less extreme errors, which could offer substantial benefit if one assumes that the cost of errors is convex rather than linear.

Additionally, Pairs almost exclusively frame discussions in terms of spending even though the stochastic per-period income process, coupled with the consumption smoothing motive, makes saving and borrowing important components of earnings maximization. Further, we see that Pairs develop simple heuristics that can lead to persistent errors that compound over time, which helps explain why absolute conditional errors are larger, on average, than absolute unconditional errors. Finally, we see that saving and binging can be the outcome of forward-looking behavior rather than the result of extreme myopia or lack of a strategy entirely.

We also provide suggestive evidence that having a more nuanced outlook on debt leads to better performance in a consumption smoothing problem where borrowing is a necessary component of optimal behavior. A potential implication of this is that financial education focused on the potential benefits and safe use of debt could improve budgetary decisions insofar as it eases strictly negative outlooks on debt.

Classifying subjects into heuristics reveals that a substantive proportion of subjects in both the Individuals and Pairs treatments are best categorized as unconditional optimizers. This is likely due to the inclusion of a consumption tool that reduces the complexity of our optimization task. If so, this suggests that providing increased access to budgetary tools and/or advice may lead real-world households to behave in a more theory consistent way.

We also show that consumption heuristics are not necessarily stable over time. This might be because heuristics evolve with experience or perhaps that income dynamics influence heuristics. One concrete pattern emerges – in both decisions blocks in both treatments, the proportion of subjects using a Hand-to-Mouth heuristic increases. This is possibly because

more subjects default to overly simple consumption rules as the lifecycle progresses.

Our results differ from the few other studies that compare the performance of pairs and individuals in a dynamic optimization task. At least one possible explanation for why this is true is that differences in our experimental design leads to a different level of problem complexity. In our experiment, for example, we do away with nominal interest rates and also provide an optimization calculator to subjects. Both Carbone and Infante (2015), and Carbone, Georgalos, & Infante (2019) consider an environment that includes a positive interest rate and provides no optimization tool. It is reasonable to think that a choice problem can be either sufficiently easy that there is no room for performance differences or sufficiently complex that forming joint-decisions is unlikely to matter. If so, it is possible our design lies somewhere between these two extremes. Further, neither of these works allow for borrowing within a period and both feature a bimodal income distribution. The confluence of these design choices yields an environment in which it is optimal for subjects to accumulate wealth and increase spending toward the end of the lifecycle. This path of optimal behavior coincides with behavior typically observed in these experiments, which may help explain the different outcomes that we observe.

References

- Ahrens, S., Bosch-Rosa, C., Meissner, T. (2022). Intertemporal consumption and debt aversion: a replication and extension. arXiv preprint arXiv:2201.06006.
- Ballinger, T. P., Hudson, E., Karkoviata, L., Wilcox, N. T. (2011). Saving behavior and cognitive abilities. *Experimental Economics*, 14(3), 349-374.
- Ballinger, T. P., Palumbo, M. G., Wilcox, N. T. (2003). Precautionary saving and social learning across generations: an experiment. *The Economic Journal*, 113(490), 920-947.
- Bao, T., Duffy, J., Hommes, C. (2013). Learning, forecasting and optimizing: An experimental study. *European Economic Review*, 61, 186-204.
- Blinder, A. S., Morgan, J. (2005). Are two heads better than one? Monetary policy by committee. *Journal of Money, Credit and Banking*, 789-811.
- Brown, Alexander L., Zhikang Eric Chua, and Colin F. Camerer. (2009). Learning and Visceral Temptation in Dynamic Savings Experiments. *Quarterly Journal of Economics*, 2009, 124(1): 197-231.
- Carbone, E. (2005). Demographics and behaviour. *Experimental Economics*, 8(3), 217-232.
- Carbone, E. (2006). Understanding intertemporal choices. *Applied Economics*, 38(8), 889-898.
- Carbone, E., Duffy, J. (2014). Lifecycle consumption plans, social learning and external habits: Experimental evidence. *Journal of Economic Behavior & Organization*, 106, 413-427.
- Carbone, E., Georgalos, K., Infante, G. (2019). Individual vs. group decision-making: an experiment on dynamic choice under risk and ambiguity. *Theory and Decision*, 87(1), 87-122.

- Carbone, E., Hey, J. D. (2001). A Test of the Principle of Optimality. *Theory and Decision*, 50(3), 263-281.
- Carbone, E., Hey, J. D. (2004). The effect of unemployment on consumption: an experimental analysis. *The Economic Journal*, 114(497), 660-683.
- Carbone, E., Infante, G. (2015). Are groups better planners than individuals? An experimental analysis. *Journal of Behavioral and Experimental Economics*, 57, 112-119.
- Charness, G., Sutter, M. (2012). Groups make better self-interested decisions. *Journal of Economic Perspectives*, 26(3), 157-76.
- Cooper, D. J., Kagel, J. H. (2005). Are two heads better than one? Team versus individual play in signaling games. *American Economic Review*, 95(3), 477-509.
- Duffy, J. (2016). Macroeconomics: A Survey of Laboratory Research. *Handbook of Experimental Economics*, Volume 2, JH Kagel and AE Roth.
- Duffy, J., Orland, A. (2021). Liquidity constraints and buffer stock savings: Theory and experimental evidence. Available at SSRN 3672126.
- Fischbacher, U. z-Tree: Zurich toolbox for ready-made economic experiments. *Exp Econ* 10, 171–178 (2007). <https://doi.org/10.1007/s10683-006-9159-4>.
- Fréchette, G. R. (2012). Session-effects in the laboratory. *Experimental Economics*, 15(3), 485-498.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, 19(4), 25-42.
- Hey, J. D., Dardanoni, V. (1988). Optimal consumption under uncertainty: An experimental investigation. *The Economic Journal*, 98(390), 105-116.

- Hey, J. D., Knoll, J. A. (2011). Strategies in dynamic decision making—An experimental investigation of the rationality of decision behaviour. *Journal of Economic Psychology*, 32(3), 399-409.
- Kagel, J. H., McGee, P. (2016). Team versus individual play in finitely repeated prisoner dilemma games. *American economic Journal: microeconomics*, 8(2), 253-76.
- Kugler, T., Bornstein, G., Kocher, M. G., Sutter, M. (2007). Trust between individuals and groups: Groups are less trusting than individuals but just as trustworthy. *Journal of Economic psychology*, 28(6), 646-657.
- Lombardelli, Clare, James Proudman and James Talbot (2005), ‘Committees versus Individuals: An Experimental Analysis of Monetary Policy Decision Making’, *International Journal of Central Banking*, 1(1), 181–205.
- Lucas Jr, R. E., Stokey, N. L. (1983). Optimal fiscal and monetary policy in an economy without capital. *Journal of monetary Economics*, 12(1), 55-93.
- Meissner, T. (2016). Intertemporal consumption and debt aversion: an experimental study. *Experimental Economics*, 19(2), 281-298.
- Carbone, E., Hey, J., Neugebauer, T. (2021). An experimental comparison of two exchange economies: Long-lived asset vs. short-lived asset. *Management Science*, 67(11), 6946-6962.
- Noussair, C., Matheny, K. (2000). An experimental study of decisions in dynamic optimization problems. *Economic Theory*, 15(2), 389-419.
- Petersen, L., Winn, A. (2014). Does money illusion matter? Comment. *American Economic Review*, 104(3), 1047-62.
- Primi, C., Morsanyi, K., Chiesi, F., Donati, M. A., Hamilton, J. (2016). The development

and testing of a new version of the cognitive reflection test applying item response theory (IRT). *Journal of Behavioral Decision Making*, 29(5), 453-469.

Rholes, R., Petersen, L. (2020). Should central banks communicate uncertainty in their projections?. Available at SSRN 3529832.

Tasneem, D., Engle-Warnick, J. (2018). Decision Rules for Precautionary and Retirement Savings. Center for Interuniversity Research and Analysis on Organizations.

7 Appendix

7.1 Tables and Figures

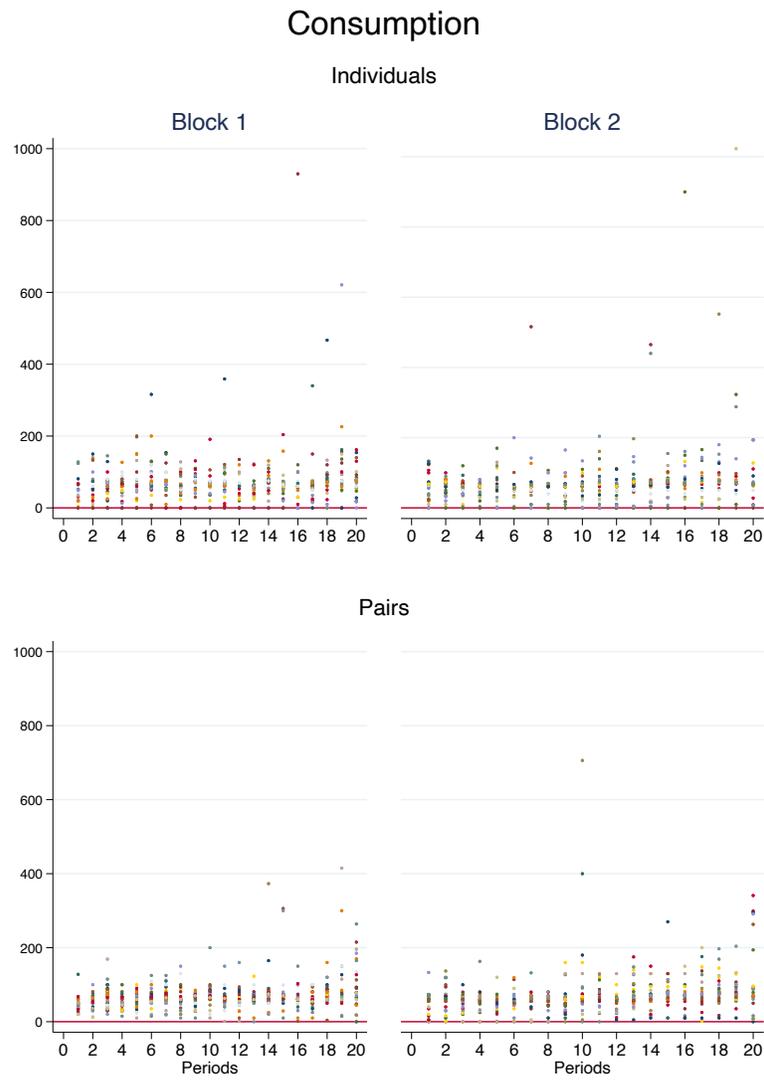


Figure 6. This figure shows the consumption decisions made in all periods and in both blocks for Individuals and Pairs.

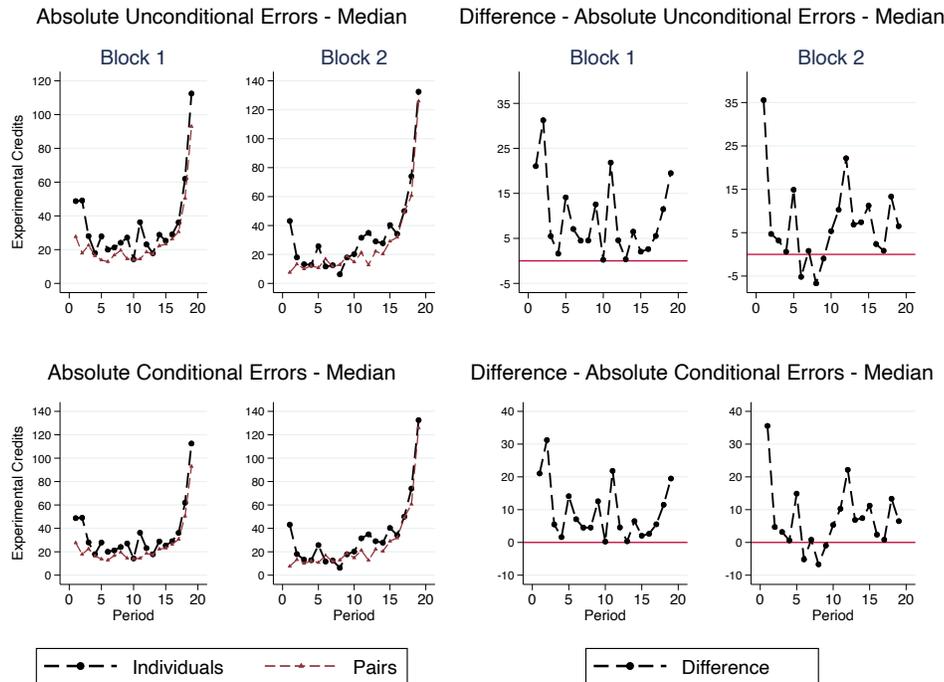


Figure 7. This figure depicts treatment-level median absolute consumption errors and their differences. Values above zero in the differences panels indicate that Pairs outperformed Individuals in that period.

Period 1 of 40 Remaining time [sec]: 39

Income this period 68
Bank account balance 68

Potential consumption spending 0 ————— max
Potential consumption spending

Consumption this period 40
Saving/Borrowing 28
Bank account balance 28
Consumption points 6
Marginal increase 1.200

Consumption spending

Figure 8. Decision screen for Individuals treatment.

Balance Table		
	Individuals	Pairs
Gender (% Male)	57.69	61.11
Avg. Age	23.29	23.29
Avg. Outside Debt (\$)	7429.41	11,900.00

Table 7. This table reports the balance across treatments.

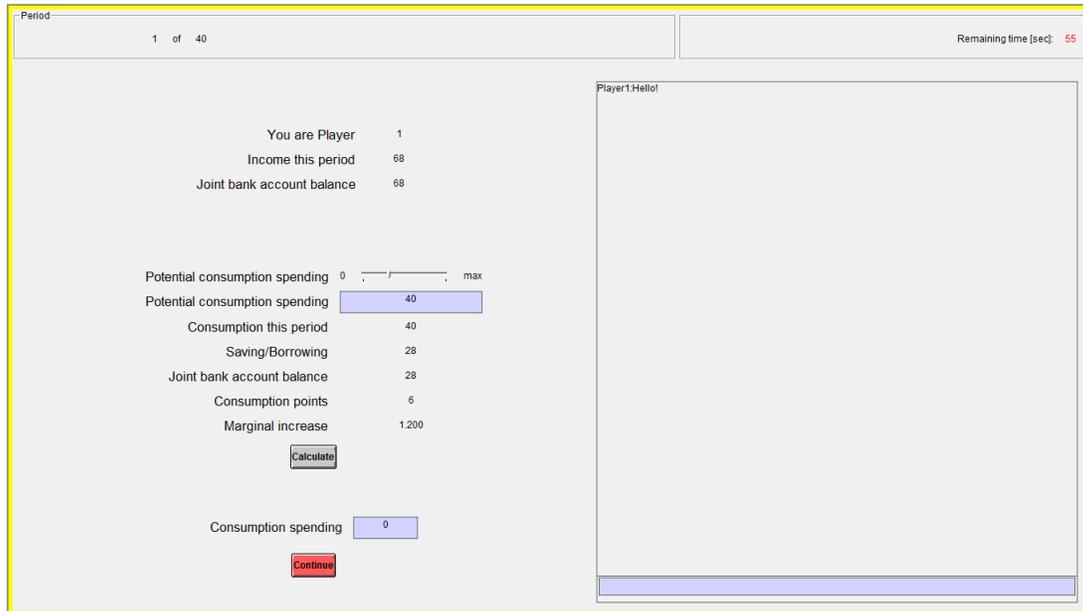


Figure 9. Decision screen for Pairs treatment.

We provide an example, corresponding to Figure 8, that explains how an individual might use the consumption tool and the available information to play this game.

Notice under ‘Income this period’ that our hypothetical subject has received an endowment of 68 experimental credits (ECs) in period 1. This is reflected in the “Bank account balance,” which updates each period to account for per-period and previous saving/borrowing. The subject may then explore the outcome of all possible consumption decisions using the ‘Potential consumption spending’ slider or by entering hypothetical levels of consumption in the gray box labeled ‘Potential consumption spending’.

For this example, our subject could spend between 0 and 128 ECs, since subjects could borrow

up to 60ECs in all but the final period of a lifecycle. Moving the slider or entering a value in the box and clicking calculate will update all other variables. In Figure 8, our hypothetical subject has selected a potential consumption value of 40. Notice that all available information has been updated to reflect this. "Consumption this period" is updated to reflect the chosen value of 40.

The "Saving/Borrowing" field updates to 28 to reflect the 28 ECs that would remain in the subject's bank account after spending 40 of the available 68 ECs¹⁵ This balance is also shown in the "Bank account balance" field within the consumption calculator.¹⁶ Further, the "Consumption points" field updates to show the consumption points earned under a choice of spending 40 ECs on consumption, which is 6.

The subject is also shown the marginal utility from using one more EC on consumption in the "Marginal increase" field, which is 1.200. The subject then enters their chosen value for consumption in the 'Consumption spending' box and presses the red button labeled 'Continue' to proceed.

7.2 Average Consumption and Consumption Heterogeneity

This section of the appendix provides details on the average consumption and consumption heterogeneity by period for both Pairs and Individuals. We measure consumption heterogeneity as the cross-sectional standard deviation of consumption decisions within a period. We graph both in Figure 10.

Worth noting in Figure 10 is that average consumption increases over time for both individuals and pairs in both decision blocks. Our results replicate a common finding in this literature that subjects under consume in early periods and over consume in later periods. This is especially true in the first half of our second lifecycle where subjects must borrow

¹⁵This number would be negative if the subject decided to spend more than 68 ECs.

¹⁶These numbers match because this is period 1. They would not necessarily match in later periods.

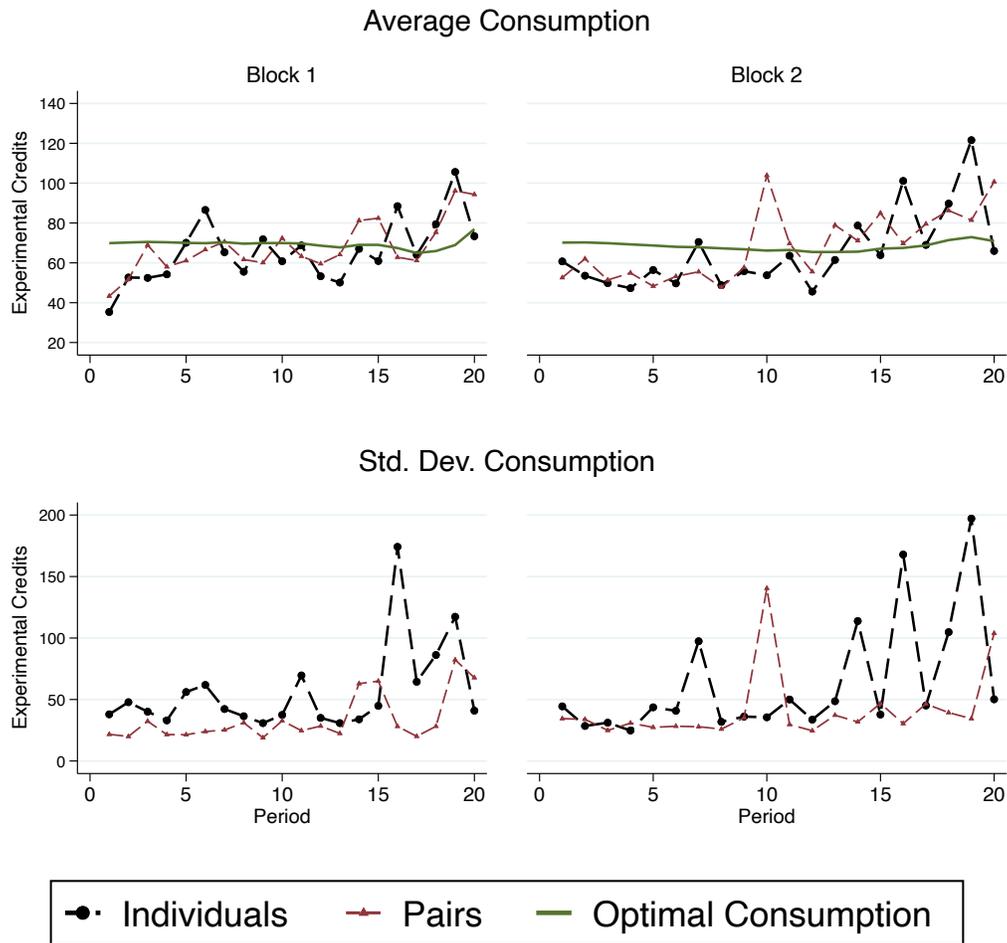


Figure 10

to consume along the unconditionally-optimal path. To see this, refer to the stochastic income process depicted in Figure 1, and note that per-period income was consistently below the unconditionally-optimal level of consumption. Finally, we note that participants in our Individuals treatment seem to exhibit more heterogeneity in consumption than do subjects in our Pairs treatment.

7.3 Empirical Data

We analyze data from the Panel Study of Income Dynamics (PSID) at the family-level, or household level, to evaluate consumption smoothing. The PSID is a longitudinal household

survey that has been conducted by the Institute for Social Research at the University of Michigan since 1968.¹⁷

Administrators of the PSID survey ask participants to report on their consumption expenditure totals across a large range of items over the course of the previous year. Examples of these items include food, utilities, transportation, education, childcare, and health care. Because respondents self-report data, people may over- or under-report their incomes and consumption expenditures for personal reasons, or due to memory lapse. Additionally, the PSID over-samples low-income families. However, we control for this in our analysis.

We restrict our sample to the years 1999 - 2017. We do this to account for two major changes in PSID data collection that came in 1999. First, surveyors began collecting data biannually instead of annually. Second, the PSID became a richer data source as surveyors began collecting additional information about household consumption and income.

For the purposes of this exercise, we made certain sample selection decisions when cleaning the PSID data. We restrict the sample to household heads aged 20 to 65. We used the OECD-modified adult equivalence scale to adjust for the increase that is proportionate per adult necessary to maintain some standard of living given a change in demographic circumstances, like the birth of a new child. We then adjusted all consumption and income measures by the personal consumption expenditure (PCE) index, to account for changes in prices, and by the OECD-modified equivalence scale. We drop all observations from the original Survey of Economic Opportunity (SEO) sample and the branches of this original sample to avoid the bias that would be introduced from the over sample of poor households, restricting our sample to just the Survey Research Center (SRC) sample. We drop observations where the household head reported working more than 5,200 hours or the household head reported working more than 520 hours at half of the minimum wage. We also drop observations where consumption expenditures are reported to be zero or negative. Thus, we restrict the sample

¹⁷Surveyors collect data on a range of topics including education, employment, income, wealth, and expenditures, which makes it well-suited for the study of consumption smoothing.

to observations that only report positive consumption expenditures. Finally, we restrict the sample to the lowest income quartile.

7.4 Empirical Motivation

This section provides some suggestive empirical evidence that further motivates our laboratory experiment. Table 8 reports the results of regressing four types of consumption growth (food, non-durable, durable, and total) on a number of household characteristics, conditional on the household head receiving an income shock of spending some months unemployed.

Table 8. Empirical Evidence

Variables	Food	Non-durable	Durable	Total
Single	-0.1781*** (0.0719)	-0.0934** (0.0536)	-0.0697 (0.0613)	-0.1112** (0.0440)
Year FE	✓	✓	✓	✓
Y Group FE	✓	✓	✓	✓
Controls	✓	✓	✓	✓
Observations	848	857	830	857
R ²	0.0806	0.0179	0.0322	0.0538

We do this using feasible GLS and include both year and income group fixed effects. Additionally, we control for the reported sex of the head of household, the number of children living in the household, the number of adults living in the household, educational level of the head of household, and the reported race of the head of household. Columns two through four in Table 8 show the effect of a negative income shock on consumption across three categories: food consumption, consumption of non-durable goods, and durable goods. Column five reports the effect of a negative income shock on total consumption.

The coefficients reported in Table 8 show the percentage point response of consumption growth, by category, to an unemployment shock. Thus, relative to a married household,

a single household experiences a 17.81 percentage point decrease in food consumption, a 9.34 percentage point decrease in non-durable consumption, and an 11.12 percentage point decrease in total consumption. If markets were complete, meaning consumers are able to insure against all possible states of the world, these coefficients should take a value of zero, which would indicate no change in consumption growth in response to an unemployment shock.

7.5 Instructions for Individuals

Overview:

Welcome! You are here today to participate in an economic experiment involving the experimental simulation of an economy. If you read these instructions carefully and make appropriate decisions, you may earn a considerable amount of money that will be paid to you in cash immediately following the experiment.

We will pay each participant \$10 for attending this experimental session. Throughout the experiment you can accrue additional earnings based on the decisions and predictions you make. You will earn points for each decision you make. Every 50 points you earn is worth an additional \$1.

You are not allowed to communicate with other participants during this experiment. If you have any questions, the experimenter will be glad to answer them privately. If you have not done so already, please turn off your cell phone now. If you do not comply with these instructions, you will be excluded from the experiment and deprived of all payments aside from the minimum payment of \$10 for attending.

Today's experiment consists of 2 sections.

Section 1 Instructions:

The first section has two parts. The first part of section one requires you to choose among a set of possible gambles. We will implement whichever gamble you choose and pay you based on the outcome of this gamble. The second part of section 1 requires you to answer a series of questions. We will pay you \$.25 for each question you answer correctly. We will provide further instructions for section 1 on your screen whenever necessary.

The second section of today's experiment involves two 'sequences' of decision making. Each sequence consists of 20 periods. You will make a new decision in each of these periods. You will make these decisions using an experimental program displayed on the screen at your terminal. Your goal during the second section of today's experiment is to convert income into consumption points. Your income in this game is valueless until you convert it into consumption points. We will convert these consumption points into cash and pay you at the end of this experiment.

Section 2 Instructions:

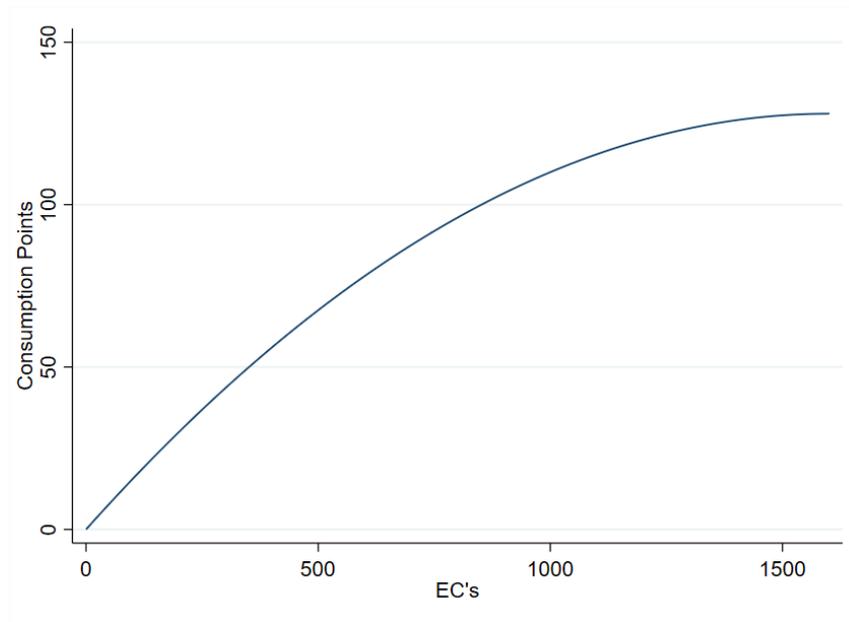
You are endowed with experimental credits (ECs) at the beginning of each period. We refer to these experimental credits as income. The amount of income you receive in each period is determined randomly and will always be an amount between 60 and 80, inclusively. You may receive as income any number of ECs between 60 and 80 with equal probability. Income in each period is independent of whatever income you received before.

After randomly determining your per-period income, the program will display this amount to you and deposit this money automatically into your bank account. The program will also display bank account balance (see Figure 1). This amount in your bank account represents your total wealth.

You must decide in each period how much of your total wealth to convert into consumption points for that period. You will earn points for consuming. Specifically, the number of points you earn in a single period is given by:

$$u(c_t) = [1600 * c_t - \frac{1}{2}c_t^2] \frac{1}{10,000}$$

Graphically:

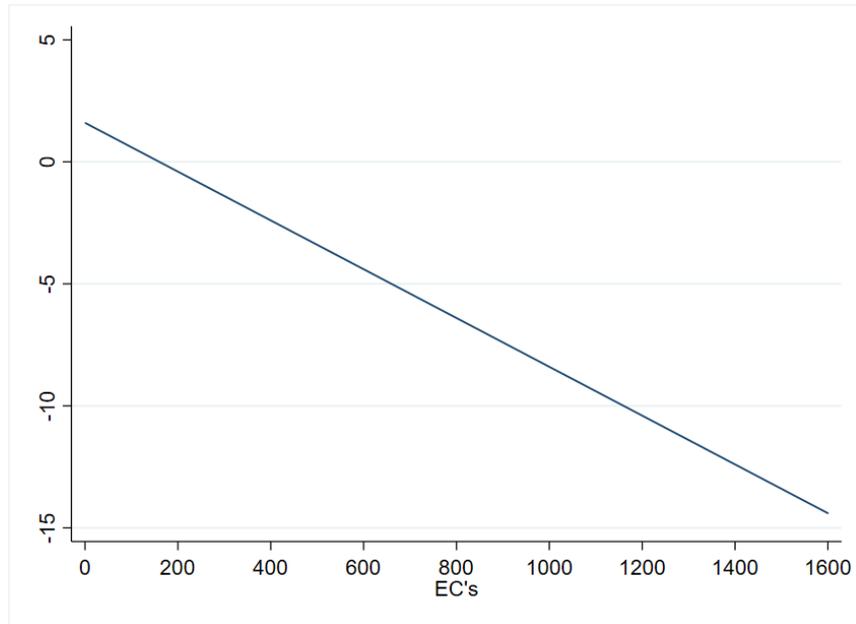


You can see from the graph above that each EC you spend on consumption (X-axis) earns a positive, but diminishing, number of consumption points (Y-axis). Each EC that you spend within a period will earn you less consumption points than the previous EC spent in the same period. This is known as diminishing marginal returns.

Specifically, the rate at which you can convert wealth into consumption points is given by:

$$u'(c_t) = 1.6 - \frac{c_t}{100}$$

Graphically:



This graph shows you how many additional points you receive within a period (Y-axis) for spending a certain amount of wealth (X-axis).

ECs have no value in this experiment. Only consumption points have value. We convert consumption points to U.S. dollars at the rate of 50 points for \$1.

Saving and borrow:

Saving:

You may save money in this experiment. Saving occurs automatically. If you spend an amount of ECs that is less than the amount of ECs in your bank account at the beginning of a period, this is called saving. Since we automatically deposit your per-period into your bank account and all of your available income is stored in your bank account, saving requires no additional actions.

Any wealth that you do not use in a period for consuming will remain in your bank account and will be available for consuming in later periods. Note that your bank balance does not

earn interest. Any money left in your account at the end of the 20th period of a sequence becomes worthless.

Borrowing:

You may borrow up to 60 credits in all periods except the last period. You cannot borrow in the last period because you are not allowed to end this game with a negative balance. Borrowing is also straightforward. If you wish to borrow money for consumption, simply add the amount of money you wish to borrow for consumption to your consumption decision. The program will always allow you to spend (except in the final period) an amount equal to whatever is in your bank account at the beginning of a period plus 60 ECs.

Saving and Borrowing example:

Suppose you have 100 ECs in your bank account at the beginning of period 2:

1. Suppose you spend 75 ECs on consumption. Then your bank account balance at the end of period 2 will be 25 ECs. Your bank account balance at the start of period 3 will be 25 ECs plus whatever endowment you receive for period 3.
2. Suppose you decide you want to spend 130 ECs. To do this, simply submit 130 ECs as your consumption decision (we discuss how to do this later in instructions). The program will allow you to spend the 130ECs and your bank account balance at the end of period 2 will be -30 ECs. Your bank account balance at the beginning of period 3 will be -30 ECs plus whatever endowment you receive for period 3.

Making a consumption decision:

We discuss two things in this section of the instructions. First, we discuss a tool available to you that will aid your consumption decision. We call this tool the consumption calculator. Second, we discuss how to submit a consumption decision.

Consumption Calculator:

We provide you with a consumption calculator to assist you when making a consumption decision. This is shown in Figure 1 below.

Label	Value
Income this period	68
Bank account balance	68
Potential consumption spending	40
Consumption this period	40
Saving/Borrowing	28
Bank account balance	28
Consumption points	6
Marginal increase	1.200

Figure 11. Decision screen for Individuals treatment.

The consumption calculator allows you to select a potential level of income you'd like to spend on consumption and shows you how much money you would save or borrow based on that decision, your resulting bank account balance, and the number of consumption points you would earn for spending that amount of income on consumption in that period.

You can choose a potential level of consumption income in two ways. First, you can move the slider (top line of the middle section of the screen in Figure 1) to some potential level of consumption spending. Doing this will cause all information to update automatically. Second, you can type a level of potential consumption spending into the box in the same section. Next, clicking the 'calculate' button in this section will cause all information to update based on whatever number you entered into the box.

Additionally, this calculator will show you the additional amount of consumption points you would earn if you decided to spend an additional EC in that period. This is called the

marginal return to consumption. Recall, Each EC that you spend within a period will earn you less consumption points than the previous EC spent in the same period.

Information:

As shown in Figure 1 above, you will always have information about your current period endowment and bank account balance whenever making a consumption decision. Furthermore, you will always have the consumption calculator available to help you understand how a potential level of consumption spending would impact your earnings and change your available bank account balance for spending in future periods.

Additionally, we will complete each period (after you make a consumption decision) by providing a review screen that reminds you of how much income you spent on consumption in that period, your bank account balance at the end of that period, the amount of consumption points you earned in that period, and your total earnings. This is shown in Figure 2.

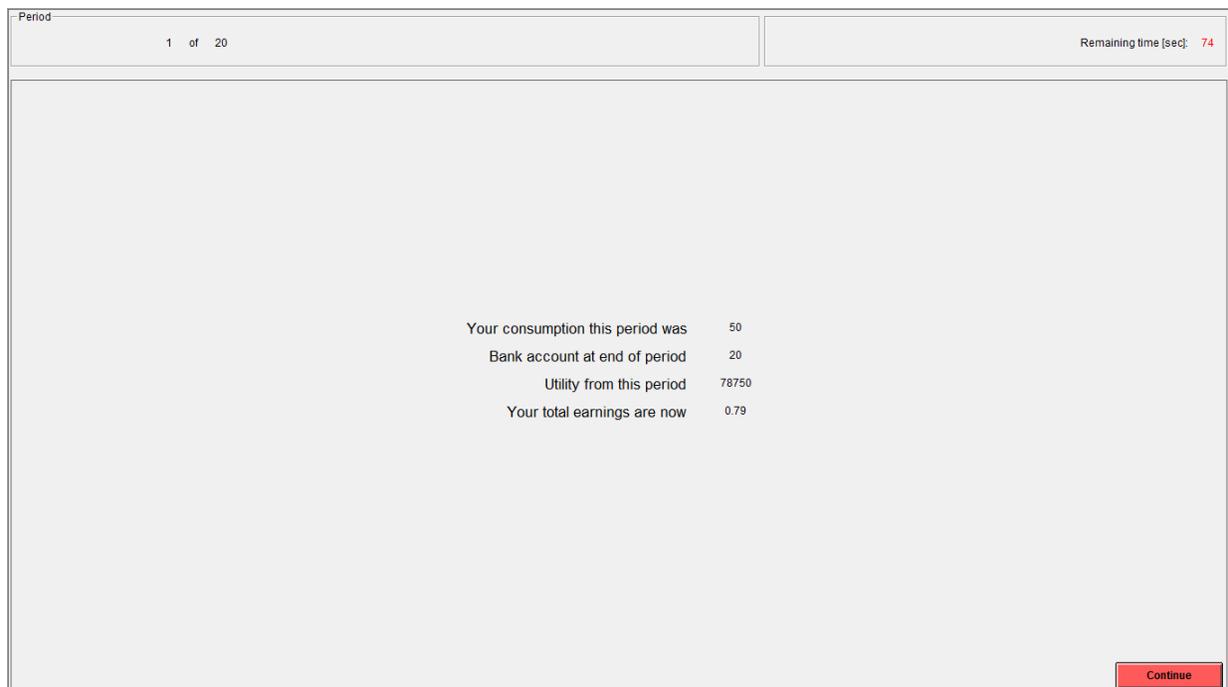


Figure 12. Review screen.

Once all subjects complete the first 20-period sequence, we will begin another 20-period

sequence. The only difference between the first and second 20-period sequence is that the sequence of endowments (the income you receive at the beginning of each period) will be different. This is because the sequence is randomly drawn with equal probability from the closed interval of $[60,80]$.

Payment:

Your payment today will consist of your \$10 show-up fee, your earnings from the initial questionnaire (where you earn \$.25 for each correct question), whatever you earn from your randomly implemented gamble, and your earnings from the two, 20-period sequences of decisions.

Questions?

Now is the time for questions. If you have a question, please raise your hand and the experimenter will answer your question in private.

Quiz:

Before continuing on to the experiment, we ask that you complete the following quiz. You can use the instructions to help answer these questions. Your performance on this quiz does not affect your payoff. Write or circle your answers to the quiz questions as indicated. Do not put your name on this quiz. If any questions are answered incorrectly, we will go over the relevant part of the instructions again.

1. In part one you will earn _____ for each correct answer in the quiz. 2. Suppose it is period 5. Does the endowment you receive in period 5 depend on the endowment you received in period 4? _____.

Does it instead depend on an endowment received in some earlier period (1, 2 or 3)? _____.

3. Suppose you have 100 ECs in your bank account at the beginning of a period. Does this

include your endowment for that period? _____.

4. Suppose you have 100 ECs in your bank account at the beginning of a period. What is the maximum amount you can spend on consumption this period? _____. What will be your bank account balance at the end of the period if you spend this maximum amount? _____.

5. True or False: We will pay you for the decisions you make in both sequences? _____

6. Suppose you earn 200 consumption points total. How much money do you earn? _____.

7.6 Instructions for Pairs

Overview:

Welcome! You are here today to participate in an economic experiment involving the experimental simulation of an economy. If you read these instructions carefully and make appropriate decisions, you may earn a considerable amount of money that will be paid to you in cash immediately following the experiment.

We will pay each participant \$10 for attending this experimental session. Throughout the experiment you can accrue additional earnings based on the decisions and predictions you make. You will earn points for each decision you make. Every 25 points you earn is worth an additional \$1.

You are not allowed to communicate with other participants during this experiment. If you have any questions, the experimenter will be glad to answer them privately. If you have not done so already, please turn off your cell phone now. If you do not comply with these instructions, you will be excluded from the experiment and deprived of all payments aside from the minimum payment of \$10 for attending.

Today's experiment consists of 3 sections.

Section 1 Instructions:

The first section has two parts. The first part of section one requires you to choose among a set of possible gambles. We will implement whichever gamble you choose and pay you based on the outcome of this gamble. The second part of section 1 will require you to answer a series of questions. We will pay you \$.25 for each question you answer correctly. We will provide further instructions for section 1 on your screen whenever necessary.

The second section of today's experiment involves two 'sequences' of decision making. Each sequence consists of 20 periods. You will make a new decision in each of these periods. You will make these decisions using an experimental program displayed on the screen at your terminal. Your goal during the second section of today's experiment is to convert income into consumption points. Your income in this game is valueless until you convert it into consumption points. We will convert these consumption points into cash and pay you at the end of this experiment.

You will make your consumption decisions in each period with a partner. We will randomly assign you a partner during this experiment. You will be able to communicate with your partner using a chat feature. Your partners are fixed for the entirety of this experiment. That is, you will work with the same partner for both 20-period sequences.

The third section again requires you to choose among a set of possible gambles. However, you will be working with the same partner to make this decision. You will be able to communicate with your partner using a chat feature. We will implement whichever gamble you and your partner choose and pay you based on the outcome of this gamble. We will provide further instructions for section 3 on your screen whenever necessary.

Section 2 Instructions:

You and your partner are jointly endowed with experimental credits (ECs) at the beginning of each period. We refer to these ECs as income. The amount of income you and your partner receive in each period is determined randomly and will always be an amount between 60 and 80, inclusively. You may receive as income any number of ECs between 60 and 80 with equal probability. Income in each period is independent of whatever income you received before.

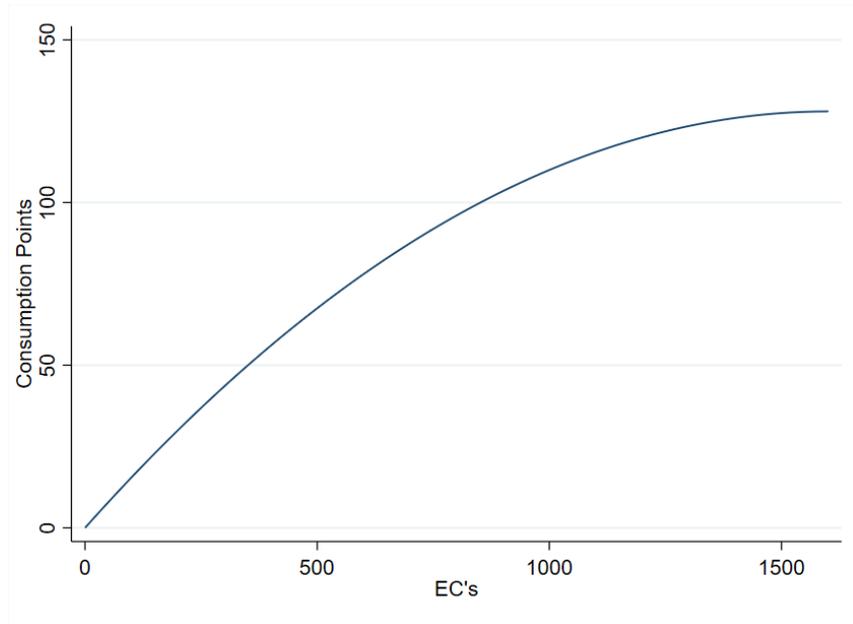
After randomly determining you and your partner's joint per-period income, the program will display this amount to you both and deposit this money automatically into your joint bank account. The program will also display the joint bank account balance (see Figure 1). This amount in your bank account represents your total wealth.

For example, suppose your joint endowment for a period is 70 ECs. You and your partner will both see this number. This means that together you must decide how to spend use these 70 ECs. To be clear, this would not mean that you have jointly gained 140 ECs.

You and your partner must decide in each period how much of your total wealth to convert into consumption points that period. Specifically, the number of points you and your partner earn in a single period is given by:

$$u(c_t) = [1600 * c_t - \frac{1}{2}c_t^2] \frac{1}{10,000}$$

Graphically:

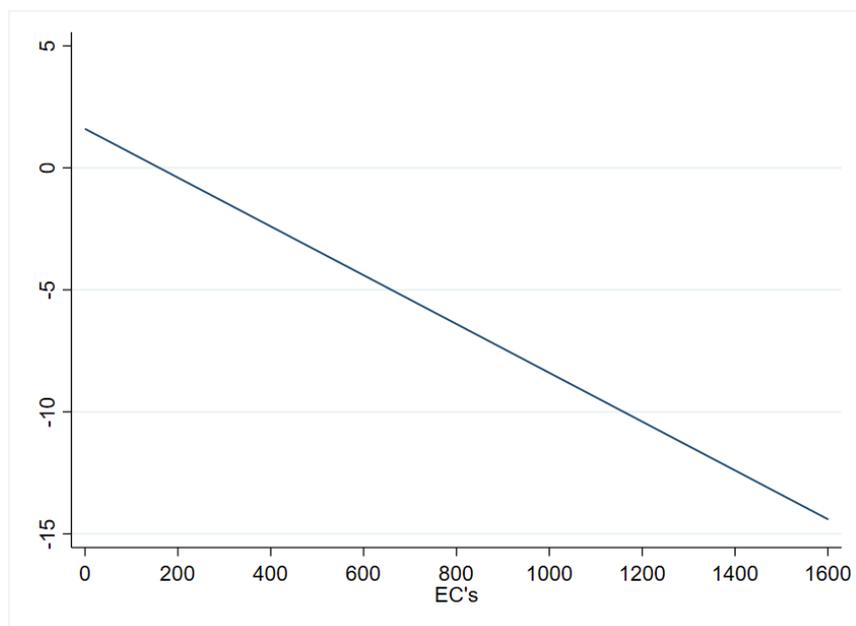


You can see from the graph above that each EC spent on consumption (X-axis) earns a positive, but diminishing, number of consumption points (Y-axis). Each EC that you spend within a period will earn you less consumption points than the previous EC spent in the same period. This is known as diminishing marginal returns.

Specifically, the rate at which you can convert wealth into consumption points is given by:

$$u'(c_t) = 1.6 - \frac{c_t}{100}$$

Graphically:



This graph shows you how many additional points you receive within a period (Y-axis) for spending a certain amount of wealth (X-axis).

ECs have no value in this experiment. Only consumption points have value. We convert consumption points to U.S. dollars at the rate of 25 points for \$1.

You and your partner will splint income evenly. For example, if your joint consumption decisions lead to a payoff of \$25 total, then you both receive \$12.50.

Saving and borrow:

Saving:

You may save money in this experiment. Saving occurs automatically. If you spend an amount of ECs that is less than the amount of ECs in your bank account at the beginning of a period, this is called saving. Since we automatically deposit your per-period income into your bank account and all of your available income is stored in your bank account, saving requires no additional actions.

Any wealth that you do not use in a period for consuming will remain in your bank account and will be available for consuming in later periods. Note that your bank balance does not earn interest. Any money left in your account at the end of the 20th period of a sequence becomes worthless.

Borrowing:

You may borrow up to 60 credits in all periods except the last period. You cannot borrow in the last period because you are not allowed to end this game with a negative bank account balance.

Borrowing is also straightforward. If you wish to borrow money for consumption, simply add the amount of money you wish to borrow for consumption to your consumption decision. The program will always allow you to spend (except in the final period) an amount equal to whatever is in your bank account at the beginning of a period plus 60 ECs.

Saving and Borrowing example:

Suppose you have 100 ECs in your bank account at the beginning of period 2:

1. Suppose you spend 75 ECs on consumption. Then your bank account balance at the end of period 2 will be 25 ECs. Your bank account balance at the start of period 3 will be 25 ECs plus whatever endowment you receive for period 3.
2. Suppose you decide you want to spend 130 ECs. To do this, simply submit 130 ECs as your consumption decision (we discuss how to do this later in instructions). The program will allow you to spend the 130 ECs and your bank account balance at the end of period 2 will be -30 ECs. Your bank account balance at the beginning of period 3 will be -30 ECs plus whatever endowment you receive for period 3.

Making a consumption decision:

We discuss two things in this section of the instructions. First, we discuss a tool available

to you and your partner that will aid your consumption decision. We call this tool the consumption calculator. Second, we discuss how to submit a consumption decision.

Consumption Calculator:

We provide you with a consumption calculator to assist you when making a consumption decision. This is shown in Figure 1 below.

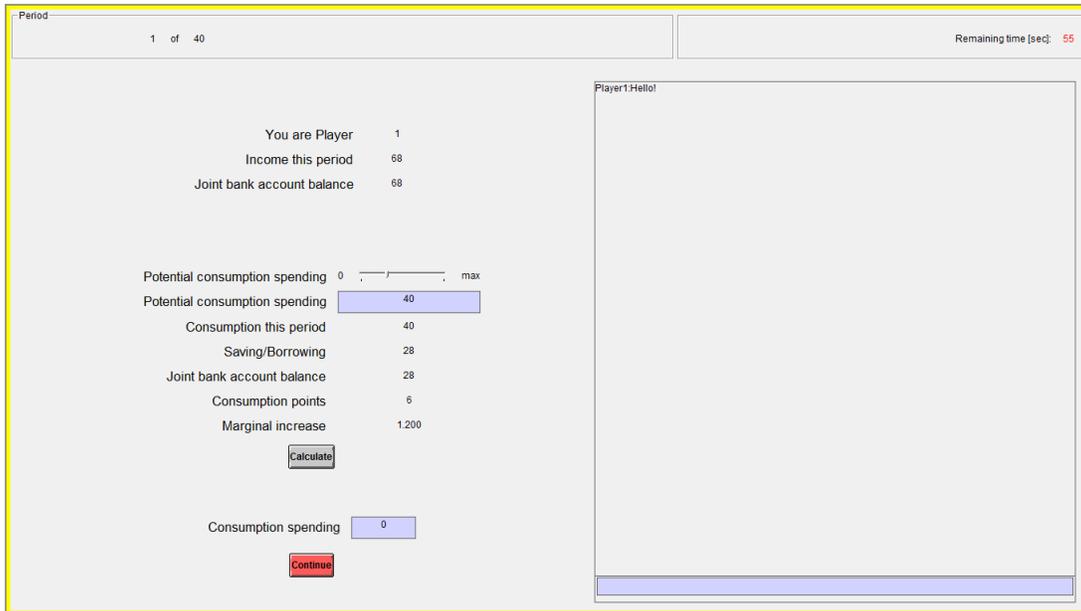


Figure 13. Decision screen for Pairs treatment.

The consumption calculator allows you to select a potential level of income you'd like to spend on consumption and shows you how much money you would save or borrow based on that decision, your resulting bank account balance, and the number of consumption points you would earn for spending that amount of income on consumption in that period.

You can choose a potential level of consumption income in two ways. First, you can move the slider (top line of the middle section of the screen in Figure 1) to some potential level of consumption spending. Doing this will cause all information to update automatically. Second, you can type a level of potential consumption spending into the box in the same section. Next, clicking the 'calculate' button in this section will cause all information to

update based on whatever number you entered into the box.

Additionally, this calculator will show you the additional amount of consumption points you would earn if you decided to spend an additional EC in that period. This is called the marginal return to consumption. Recall, Each EC that you spend within a period will earn you less consumption points than the previous EC spent in the same period.

Both you and your partner have independent consumption calculators. This means that your partner does not automatically see information for potential levels of consumption spending that you check using your calculator and vice versa.

You and your partner can chat freely using the chat box picture on the right side of Figure 1. You should use this chat box to jointly agree upon a decision about how much of your joint income you should spend on consumption in each period.

Once you have reached an agreement using the chat box, you should both input your consumption spending decision and click continue. If you both input the same number, the program will proceed and you will jointly earn whatever amount of consumption points corresponds to your joint decision. If the numbers do not match, the program will not continue forward. You will receive a notification from the program whenever you input a number that does not match your partners.

Information:

As shown in Figure 1 above, you will always have information about your current period endowment and bank account balance whenever making a consumption decision. Furthermore, you will always have the consumption calculator available to help you understand how a potential level of consumption spending would impact your earnings and change your available bank account balance for spending in future periods. This is shown in Figure 2.

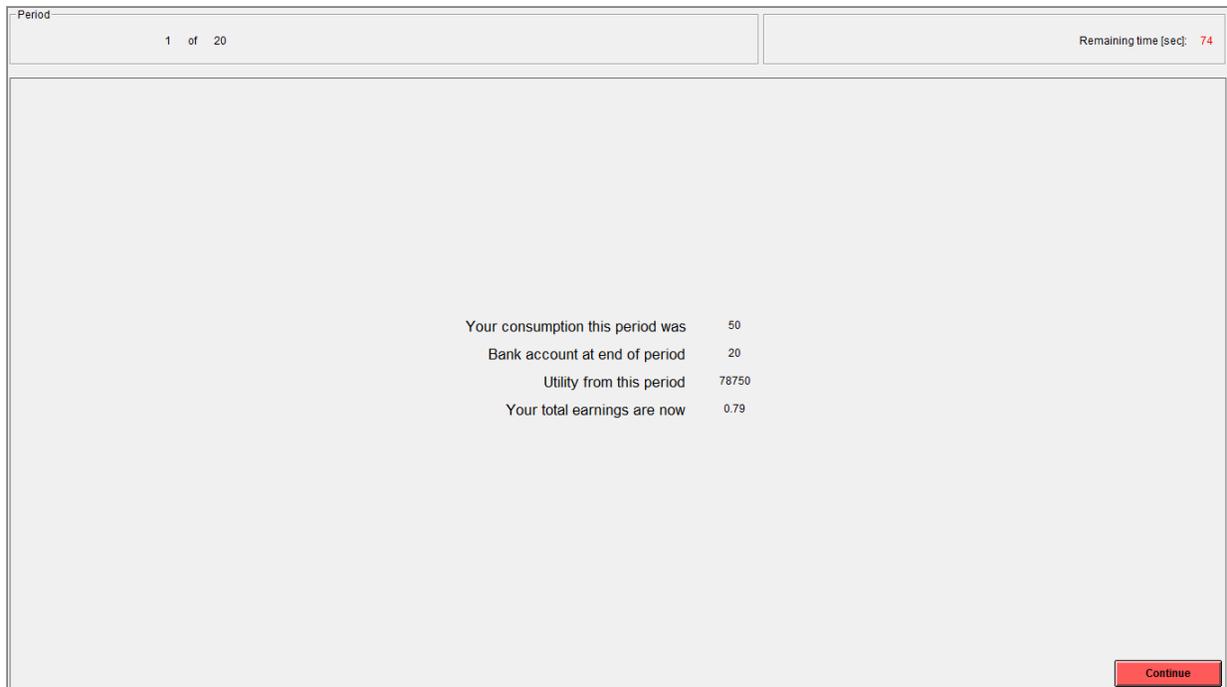


Figure 14. Review screen.

Additionally, we will complete each period (after you make a consumption decision) by providing a review screen that reminds you of how much income you spent on consumption in that period, your bank account balance at the end of that period, the amount of consumption points you earned in that period, and your total consumption points.

Once all subjects complete the first 20-period sequence, we will begin another 20-period sequence. The only difference between the first and second 20-period sequence is that the sequence of endowments (the income you receive at the beginning of each period) will be different. This is because the sequence is randomly drawn with equal probability from the closed interval of $[60,80]$.

Payment:

Your payment today will consist of your \$10 show-up fee, your earnings from the initial questionnaire (where you earn \$.25 for each correct question), whatever you earn from both of your randomly implemented gamble, and your earnings from the two, 20-period sequences

of decisions.

Questions?

Now is the time for questions. If you have a question, please raise your hand and the experimenter will answer your question in private.

Quiz:

Before continuing on to the experiment, we ask that you complete the following quiz. You can use the instructions to help answer these questions. Your performance on this quiz does not affect your payoff. Write or circle your answers to the quiz questions as indicated. Do not put your name on this quiz. If any questions are answered incorrectly, we will go over the relevant part of the instructions again.

1. In part one you will earn _____ for each correct answer in the quiz.
2. Suppose it is period 5. Does the endowment you receive in period 5 depend on the endowment you received in period 4? _____. Does it instead depend on an endowment received in some earlier period (1, 2 or 3)? _____.
3. Suppose you have 100 ECs in your bank account at the beginning of a period. Does this include your endowment for that period? _____.
4. Suppose you have 100 ECs in your bank account at the beginning of a period. What is the maximum amount you can spend on consumption this period? _____. What will be your bank account balance at the end of the period if you spend this maximum amount? _____.
5. True or False: We will pay you for the decisions you make in both sequences? True False.
6. If you and your partner together earn \$30 for your joint consumption decisions, then you will personally earn how much? _____.

7. Suppose you earn 200 consumption points total. How much money do you and your partner earn? _____.
8. Does the marginal increase from an EC spent within a period earn you more or less consumption points than the previous EC spent in the same period?

7.7 Other Materials

CRT Questions:

1. A bat and a ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How much does the ball cost in cents?
2. If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets, in minutes?
3. In a lake there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake, in days?
4. In an athletics team, tall members are three times more likely to win a medal than short members. This year the team has won 60 medals so far. How many of these have been won by short athletes?
5. If John can drink one barrel of water in 6 days, and Mary can drink one barrel of water in 12 days, how long would it take them to drink one barrel of water together?
6. Jerry received both the 15th highest and the 15th lowest mark in the class. How many students are in the class?

Demographics Survey:

1. Select your gender. (Male, Female, Other?)

2. What is your age?
3. Which year in school are you? (Freshman, Sophomore, Junior, Senior, Graduate)
4. What is your major?
5. To the best of your knowledge, what is your GPA?
6. Approximately how much student loan debt do you have?
7. Approximately how much other debt do you have?
8. What income class were you in growing up, e.g. lower, middle, upper?
9. What is your current political affiliation?

Survey of Decisions:

1. What information did you use in making your consumption decisions?
2. Did you have a decision rule, if so, what was it?
3. Did you feel like you had enough time to make your decisions?
4. Do you believe it is good or bad to have debt?
5. Do you believe it is good or bad to have savings?
6. How well do you believe you performed on the consumption task? 25th percentile? 50th percentile? 75th percentile? 99th percentile?

Extra survey of decisions questions for Pairs treatment:

7. What was your communication strategy with your partner?
8. Did you tend to agree or disagree with your partner?
9. What was your strategy for overcoming disagreements?

Chapter 2: Market Completeness by Income Group

Logan Miller

1 Introduction

The modern era is awash with complicated financial securities, commodities, and insurances available to investors and consumers alike. In 1954 Kenneth Arrow and Gerard Debreu introduced the concept of the Arrow-Debreu market, or complete market, to help provide some insight into such complicated financial structures. Under the auspice of complete market theory there is a market for every good where the concept of a “good” encapsulates any state, date, or environment in which a good can or may be consumed and sold. What this means for the consumers within a complete market is that they would be able to smooth their consumption no matter what the state of the world was by selling any assets, collecting on insurance, or using personal savings. In other words, any idiosyncratic income shock would not affect their consumption, consumption can only be affected by aggregate shocks across the entire economy. There have been two major complete market tests done on the U.S. market to date. One was done using Consumer Exchange Survey data (Mace 1991) and the other was done using Panel Study of Income Dynamics data (Cochrane 1991). Both studies concluded that the U.S. market is nearly complete, but the average consumer cannot insure against prolonged idiosyncratic shocks like spending months absent from work. By evaluating market completeness per income group we can begin to learn more about which part or parts of the market are uninsured and develop a deeper understanding of the incomplete market situation Americans are facing. In this paper I show that nondurable, durable, food, and total consumption growth is not statistically significantly different from zero for the two upper income quartiles when these income quartiles receive an income shock of spending several months unemployed. These two quartiles represent the upper half of American earners. In

contrast, the full insurance hypothesis is rejected for the lowest income classes.

I build on what is already a large collection of literature on complete markets, incomplete markets, and partial insurance. Two tests of complete markets in the United States were mentioned earlier, one using the Consumer Expenditure Survey (Mace 1991) and the other using the Panel Study of Income Dynamics (Cochrane 1991). Mace (1991) finds that for exponential utility types the overall results are somewhat indecisive. Full insurance is rejected for some goods, but not others. Once the change in aggregate consumption is accounted for, the change in household income does not explain changes in consumption. She also explores power utility types and finds that full insurance is rejected for most goods in this specification. Cochrane (1991) finds that full insurance is rejected for long illness and involuntary job loss, but short spells of illness, unemployment, loss of work due to strike, and involuntary moves are well insured. Involuntary job loss has the largest and most significant impact on consumption growth. Townsend (1994) conducted another complete market test on villages in India that showed poor villages in the semi-arid tropics of India appear to operate under complete markets, but his model is rejected statistically. I add to this stream of literature by evaluating market completeness at the income quartile level.

There are also the publications that discuss partial insurance in the U.S. economy. Ehrlich and Becker (1972) develops a theory of demand for insurance by emphasizing the interaction between market insurance, self-insurance, and self-protection. All of these are forms of insurance, or protection, I allow for in my model. Aiyagari (1994) provides an exposition of models with many agents, rather than the standard representative agent model, and allows for precautionary savings motives and liquidity constraints. I borrow from his standard growth model to help develop the model I use. Blundell et al. (2008) details partial insurance and how it relates to the persistence of income shocks over the 1980s in the United States. He finds full insurance against transitory shocks except among poor households which I build on in this paper to show it is particularly persistent shocks that the poor are unable to insure against. Finally, Kaplan and Violante (2010) assess the degree of consumption smoothing

implicit in the incomplete markets model.

This previous literature provides the foundation for a better understanding of the incomplete market situation the average American is facing which I hope to improve on. It also details some of the methods American consumers are using to build out their portfolios to begin to insure against income shocks and smooth consumption, and the degree to which the average consumer is insured against said shocks. My contribution to this body of literature will be an explanation of which income level is too low for the American consumer to fully insure.

2 Data

In this paper I use Panel Study of Income Dynamics (PSID) family-level data to evaluate market completeness by income group. The PSID is a longitudinal household survey that has been conducted since 1968 by the Institute for Social Research at the University of Michigan. It collects data on a range of topics including education, employment, income, wealth, and expenditures which makes it well-suited for my study of market completeness by income group. While there are more comprehensive surveys of consumption expenditures, the Consumer Expenditure Survey for example, the detailed data on income in the PSID is necessary for my study.

Specifically, I use family-level data after 1999 for household heads aged between 20-65. I restrict my sample to the years 1999 - 2017 to maintain consistency in the sample as there were changes in the survey administered by the PSID. The main change in regards to my project was the expansion of consumption expenditure items the surveyors asked participants to report on. The additional items allow for a more complete measure of durable, nondurable, and total consumption.

As the PSID is survey data, the administrators of ask participants to report on their consumption expenditure totals across items like food, utilities, transportation, education, childcare, health care, etc... over the course of the last year. The family-level survey is

administered every two years and reports data from the year prior to the time of the interview.

As with most surveys, there are few drawbacks to using this type of data. Due to the data being self-reported, there are potential biases that could be introduced. People may over- or under-report their incomes and consumption expenditures for personal reasons or due to incorrect reporting. The PSID was formed to study the dynamics of income and poverty. As a result, the original sample drawn in 1968 was formed from two independent samples: the Survey of Economic Opportunity (SEO sample) and the survey from the Survey Research Center at the University of Michigan (SRC sample). The SEO sample consisted of 1,872 low income families and the SRC sample consisted of 2,930 families that was originally designed by the Survey Research Center. This oversampling of low income families could also introduce biases which I control for in my data cleaning which will be described in more detail below.

For the purposes of this exercise I made certain sample selection decisions when cleaning the PSID data. I restrict the sample to household heads aged 20 to 65. I used the OECD-modified adult equivalence scale to adjust for the increase that is proportionate per adult necessary to maintain some standard of living given a change in demographic circumstances, like the birth of a new child. I then adjusted all consumption and income measures by the personal consumption expenditure (PCE) index, to account for changes in prices, and by the OECD-modified equivalence scale. I drop all observations from the original SEO sample and the branches of this original sample to avoid the bias that would be introduced from the oversample of poor households, restricting my sample to just the SRC sample. I drop observations where the household head reported working more than 5,200 hours or the household head reported working more than 520 hours at half of the minimum wage. I also drop observations where consumption expenditures are reported to be zero or negative. Thus, I restrict the sample to observations that only report positive consumption expenditures.

3 Specification

I use an ordinary least squares (OLS) model that regresses nondurable, durable, total, and food consumption growth on a binary indicator for whether or not the household head received a certain negative income shock for each income class. In addition, I include controls for income group effects, income growth, and year fixed effects. Income shocks are instrumented for using a binary indicator that takes a value of 1 when the household head reports having spent anywhere from 1 to 12 months out of work due to an idiosyncratic shock. These results are reported in Table 1. I use another OLS model that includes controls for whether or not the household received unemployment benefits or food stamps when the household head was unemployed, the results of which are reported in Table 2. The specifications are generated in order to test the proposition under full consumption insurance, consumption growth should be independent of negative, idiosyncratic income shocks.

$$\log\left(\frac{c_{t+1}^i}{c_t^i}\right) = \alpha + \beta_j y_{it} + \beta_j y_{it} * x_{it} + \beta_j \log\left(\frac{y_{t+1}^i}{y_t^i}\right) * y_{it} + z_t + \epsilon_{it} \quad (1)$$

where c is household consumption, y is the income groups, x is the idiosyncratic income shock, and z is the year fixed effects. More specifically, my dependent variable is non-durable/durable/total/food consumption growth which is the log difference of nondurable, durable, total, food consumption expenditures. Nondurable consumption is constructed by summing total food, total utility, gas, parking, bus and train, other transportation, education, childcare, and health care expenditures. Food consumption includes food at home, food away from home, and food delivered expenditure. Durable consumption includes total transportation which is made up of the transportation expenditures from nondurable as well as vehicle loan, lease, down payment and repair expenditures. It also includes total housing that accounts for all expenditure associated with housing: rent/mortgage, property tax, and home insurance. Total consumption is just the summation of the durable and nondurable consumption measures. My independent variables are income growth which is the log dif-

ference of household head labor income, a binary indicator for whether or not the family received an income shock, and a categorical variable for income groups. Household head is determined by relation to the original group of families surveyed in 1968. The income shocks I used include months the head was unemployed, months the head missed work due to others being sick, months the household head was sick, months the household head was on strike, and months the household head was laid off. The income shock variable shows a value of 1 if the household head missed anywhere from 1 to 12 months due to the above income shocks. I only include months long idiosyncratic income shocks as previous studies have shown that only the more persistent shocks lead to rejections of full insurance.

This OLS model was inspired by Cochrane's (1991) model which is just a cross-sectional analogy of Hall's (1978) regression to test the permanent income hypothesis. These are, however, distinct models as any model testing consumption insurance is looking at a household's ability to smooth their consumption over states of nature. Whereas, a model testing the permanent income hypothesis would be looking at that household's ability to smooth their consumption over time. My specification looks at a household's ability to smooth consumption over states of nature, employed or unemployed, and over time. One issue at the forefront of this regression is the likelihood that income, a right-hand side variable for me, is not uncorrelated with variation in preferences. In this sense, a preference shift that affects the household's desired allocation of consumption may also affect its allocation of leisure and labor hours (income) and there is likely some correlation. I will have to try an alternative approach in extensions of this work to attempt to get around this problem.

The parameters of interest are the β_j 's on the interaction between the income groups, y , and the negative income shocks, x . These β 's should not be significantly different than zero if consumption growth is in fact independent of negative income shocks as a result of full consumption insurance. My hypothesis is the β corresponding to the interaction between the lowest (poorest) income quartile and the negative income shock will be negative and significant. This would indicate that the lowest quarter of American earners do not have full

consumption insurance.

4 Results

Table 1 presents the results for tests of consumption insurance across nondurable, durable, food, and total consumption growth respectively when the household head spends several months unemployed. The coefficients reported are from an interaction between households in their respective income groups and the binary indicator for when households receive an income shock. I focus on the months unemployed shock in particular as it seems to have the largest impact on consumption smoothing. I show the effects of other types of idiosyncratic income shocks in the appendix. The lowest and second lowest income quartiles both experience significant, negative consumption growth across all of the models except (2), where the second income quartile is negative but it is not significant.

Table 1. A Test of Full Insurance Hypothesis by Consumption Items

Variables	(1)	(2)	(3)	(4)
Income Group 1	-0.0624*** (0.0201)	-0.1800*** (0.0279)	-0.0834*** (0.0247)	-0.1132*** (0.0183)
Income Group 2	-0.1218*** (0.0404)	-0.0114 (0.0554)	-0.0939* (0.0494)	-0.0770** (0.0369)
Income Group 3	-0.0366 (0.0579)	-0.0620 (0.0788)	0.0061 (0.0708)	-0.0128 (0.0529)
Income Group 4	-0.0303 (0.0579)	0.0370 (0.0788)	-0.1281* (0.0708)	0.0081 (0.0529)
Income Group Effects	X	X	X	X
Year Fixed Effects	X	X	X	X
Observations	28,635	28,371	28,531	28,635
R ²	0.0224	0.0403	0.0147	0.0369

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis. Column (1) is an OLS regression of nondurable consumption growth; column (2) is a regression of durable consumption; column (3) is a regression of food consumption growth; and column (4) is a regression of total consumption growth. Nondurable consumption is constructed by summing total food, total utility, gas, parking, bus and train, other transportation, education, childcare, and health care expenditures. Food consumption includes food at home, food away from home, and food delivered expenditure. Durable consumption includes total transportation which is made up of the transportation expenditures from nondurable as well as vehicle loan, lease, down payment and repair expenditures. It also includes total housing that accounts for all expenditure associated with housing: rent/mortgage, property tax, and home insurance. Total consumption is just the summation of the durable and nondurable consumption measures.

The first income quartile shows a 6.2 percentage point decrease in nondurable consumption growth when households are unemployed for several months, a 18 percentage point decrease in durable consumption growth, a roughly 8.3 percentage point decrease in food consumption growth, and approximately a 11.3 percentage point decrease in total consumption growth. These are all significant at the one percent level. If low income households had full consumption insurance against income shocks we would expect the coefficients to not be significantly different than zero, however, this is not the case. This means that the lowest income quartile in America is unable to smooth their consumption when they receive a negative income shock. Intuitively, it makes sense that the consumption type most affected for this income quartile is durable consumption. When the lowest earners receive a negative income shock they cannot consume much lower than the near subsistence levels they already consume at. Therefore, the consumption category taking the largest hit is the one with the most expensive goods: durables.

There are a range of possible explanations for why the first income quartile is unable to insure against income shocks. It could be the case that these agents are cognitively constrained and do not have the cognitive resources required to make tough consumption smoothing decisions as shown by Ballinger et al. (2011). It has been shown in previous literature by Carbone and Hey (2004) that consumption often too closely tracks income for agents to adequately smooth their consumption. It is also possible that the explanation for this group of people is simply mechanical. In other words, they simply do not make enough money to allocate a portion of that income to different forms of consumption insurance. It is even possible that these agents do not have knowledge of how to consumption smooth. Ballinger et al. (2003) show that the subjects that perform best at dynamic optimization are those that are able to watch their predecessors smooth consumption. If the agents in the lowest income quartile are not able to learn this behavior they themselves may not engage in it.

The second income quartile sees roughly a 12.2 percentage point decrease in nondurable

consumption growth, a 9.4 percentage point decrease in food consumption growth, and a 7.7 percentage point decrease in total consumption growth. These are significant at the one, ten, and five percent levels respectively. While it is unlikely that the reason for this is mechanical as these people are mostly lower middle to middle class, all of the reasons listed above are still valid explanations.

In Table 2, I present the results for tests of consumption insurance across nondurable, durable, food, and total consumption growth respectively with government benefits included when the household head spends several months unemployed. The coefficients reported are a triple interaction between households in their respective income quartiles, the indicator for when they receive an income shock, and an indicator that signifies they received some government benefit, unemployment benefits or food stamps, in the face of the income shock.

In all of the odd numbered models I am accounting for when the household receives unemployment benefits. The lowest and second lowest income quartiles again both experience significant, negative consumption growth with a few exceptions. In model (1) the first income quartile shows a roughly 7.2 percentage point decrease in nondurable consumption growth that is significant at the one percent level. When accounting for the unemployment benefits the household received, there is a 2.1 percentage point increase, though it is not significant, in nondurable consumption growth. The positive sign is what is important for interpretation. This suggests that while unemployment benefits do a little to help households smooth consumption, they do not completely cover the decrease in nondurable consumption growth. The results are very similar for models (3), (5), and (7) which are durable, food, and total consumption growth respectively. In model (3) the first income quartile is experiencing an 18 percentage point decrease in durable consumption growth. However, the sign and magnitude for household receiving unemployment benefits is barely positive. This suggests that, similar to the results discussed from Table 1, durable consumption is the category that takes the largest hit when a household in the first income quartile experiences a significant, negative income shock. In model (5) we see the first income quartile is showing a 12.57 percentage

point decrease in food consumption growth. Interestingly, we see that the use of benefits increases food consumption growth by almost ten percentage points and this is significant at the ten percent level. In model (7) the first income quartile loses 12.6 percentage points of total consumption growth, and while it is not significant there is a positive bump from the use of unemployment benefits.

The even numbered models correspond to when the household receives food stamp assistance. What is occurring here is a little less clear. In models (4), (6), and (8) in which the dependent variables are durable, food, and total consumption growth respectively, we see the first income quartile experience significant decreases in consumption growth. Curiously, when I account for food stamp assistance I see further, sometimes significant, decreases in consumption growth. However, the PSID does not include the value of the food stamps in the food expenditure category. In this sense, I would not see food stamps playing a role as consumption insurance but rather as a signal that the household already lacks consumption insurance. Data limitations keep me from explicitly testing this proposition as the PSID stopped collecting data on the value of the food stamps the household received after the 2007 survey. This is something that will have to be evaluated further with a different dataset or across a different selection of years.

Table 2. Consumption with Government Benefits

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Income Group 1	-0.0716*** (0.0253)	-0.0325 (0.0237)	-0.1880*** (0.0353)	-0.1800*** (0.0329)	-0.1257*** (0.0313)	-0.0551* (0.0290)	-0.1260*** (0.0231)	-0.1042*** (0.0217)
Income Group 2	-0.0763 (0.0588)	-0.1519*** (0.0432)	0.1352* (0.0806)	-0.0204 (0.0593)	-0.1235* (0.0719)	-0.1104** (0.0526)	0.0066 (0.0537)	-0.0922** (0.0395)
Income Group 3	-0.0265 (0.0971)	-0.0370 (0.0578)	0.0985 (0.1320)	-0.0623 (0.0788)	0.0218 (0.1186)	0.0047 (0.0705)	0.0289 (0.0887)	-0.0131 (0.0529)
Income Group 4	-0.0555 (0.0865)	-0.0313 (0.0570)	0.0115 (0.1176)	0.0367 (0.0788)	-0.1140 (0.1057)	-0.1299* (0.0706)	0.0312 (0.0790)	0.0074 (0.0529)
Group 1 Benefits	0.0213 (0.0391)	-0.1549*** (0.0418)	0.0074 (0.0543)	-0.0203 (0.0585)	0.0990** (0.0481)	-0.2406*** (0.0518)	0.0289 (0.0358)	-0.0718* (0.0382)
Group 2 Benefits	-0.0866 (0.0802)	0.2282* (0.1205)	-0.2719** (0.1099)	0.0594 (0.1643)	0.0538 (0.0980)	0.1198 (0.1469)	-0.1546** (0.0732)	0.1146 (0.1101)
Group 3 Benefits	-0.0161 (0.1205)	0 (0)	-0.2497 (0.1639)	0 (0)	-0.0244 (0.1473)	0 (0)	-0.0658 (0.1101)	0 (0)
Group 4 Benefits	0.0472 (0.1159)	0 (0)	0.0451 (0.1577)	0 (0)	-0.0247 (0.1417)	0 (0)	-0.0409 (0.1059)	0 (0)
Income Group Effects	X	X	X	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X	X	X	X
Unemployment Benefits	X		X		X		X	
Food Stamps		X		X		X		X
Observations	28,635	28,635	28,371	28,371	28,531	28,531	28,635	28,635
R ²	0.0226	0.0245	0.0403	0.0404	0.0150	0.0221	0.0372	0.0381

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis. Columns (1) and (2) are regressions of nondurable consumption growth; columns (3) and (4) are regressions of durable consumption growth; columns (5) and (6) are regressions of food consumption growth; columns (7) and (8) are regressions of total consumption growth.

These results motivate a variety of robustness checks that I include results for in the appendix. I regress each type of consumption growth (nondurable, durable, food, and total) across each of the different types of negative income shock: months unemployed, months off work for a sick family member, months off work for illness, months spent on strike, and months laid off. I also run these same regressions including the same government benefits from above. The results reported on are robust and can be summarized as follows. The only income shock that is consistently significant is when the head of household spends several months unemployed, and the first income quartile is consistently unable to insure their consumption against this type of shock.

5 Conclusion

I show that the first income quartile shows a 6.2 percentage point decrease in nondurable consumption growth when households are unemployed for several months, a 18 percentage point decrease in durable consumption growth, a roughly 8.3 percentage point decrease in food consumption growth, and approximately a 11.3 percentage point decrease in total consumption growth. These are all significant at the one percent level. This suggests that the lowest income quartile in America is unable to smooth their consumption when they receive a negative income shock and therefore are not operating under complete markets.

In order to check if certain government benefits, foods stamps and unemployment compensation, help bridge the gap for the lowest income quartile in run regressions including controls for the benefits. I show again that the lowest and second lowest income quartiles again both experience significant, negative consumption growth with a few exceptions. While unemployment benefits appear to help households make up some of the consumption they lose, food stamps do not play the same role. Indeed, it appears that food stamps instead serve as a signal that households already lack full consumption insurance.

References

- [1] Aiyagari, S. R. (1994). "Uninsured idiosyncratic risk and aggregate saving." *The Quarterly Journal of Economics*, 109(3), 659-684.
- [2] Ballinger, T. P., Palumbo, M. G., & Wilcox, N. T. (2003). "Precautionary saving and social learning across generations: an experiment." *The Economic Journal*, 113(490), 920-947.
- [3] Ballinger, T. P., Hudson, E., Karkoviata, L., & Wilcox, N. T. (2011). "Saving behavior and cognitive abilities." *Experimental Economics*, 14(3), 349-374.
- [4] Carbone, E., & Hey, J. D. (2004). The effect of unemployment on consumption: an experimental analysis. *The Economic Journal*, 114(497), 660-683.
- [5] Cochrane, J. H. (1991). "A simple test of consumption insurance." *Journal of Political Economy*, 99(5), 957-976.
- [6] Ehrlich, I., & Becker, G. S. (1972). "Market insurance, self-insurance, and self-protection." *Journal of Political Economy*, 80(4), 623-648.
- [7] Hall, R. E. (1978). "Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence." *Journal of Political Economy*, 86(6), 971-987.
- [8] Kaplan, G., & Violante, G. L. (2010). "How much consumption insurance beyond self-insurance?." *American Economic Journal: Macroeconomics*, 2(4), 53-87.
- [9] Krueger, D., Mitman, K., & Perri, F. (2016). "Macroeconomics and household heterogeneity." In *Handbook of Macroeconomics* (Vol. 2, pp. 843-921). Elsevier.
- [10] Mace, B. J. (1991). "Full insurance in the presence of aggregate uncertainty." *Journal of Political Economy*, 99(5), 928-956.

- [11] Townsend, R. M. (1994). "Risk and insurance in village India." *Econometrica: Journal of the Econometric Society*, 539-591

6 Appendix

Table 3. Nondurable Consumption

Variables	(1)	(2)	(3)	(4)	(5)
Income Group 1	-0.0624*** (0.0201)	0.0457 (0.0799)	0.0812** (0.0334)	-0.2028 (0.3286)	0.0522 (0.0412)
Income Group 2	-0.1218*** (0.0404)	0.1115 (0.0913)	0.0060 (0.0384)	-0.1166 (0.3287)	-0.0440 (0.0455)
Income Group 3	-0.0366 (0.0579)	0.0650 (0.0896)	0.0911** (0.0416)	0.1648 (0.2684)	-0.0383 (0.0629)
Income Group 4	-0.0303 (0.0579)	0.0444 (0.1290)	0.0153 (0.0533)	0.2522 (0.2684)	0.0001 (0.1471)
Income Group Effects	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X
Months Unemployed	X				
Months Others Sick		X			
Months Head Sick			X		
Months on Strike				X	
Months Laid Off					X
Observations	28,635	28,635	28,635	28,635	28,635
R ²	0.0224	0.0218	0.0221	0.0218	0.0219

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis.

Table 4. Total Consumption

Variables	(1)	(2)	(3)	(4)	(5)
Income Group 1	-0.1132*** (0.0183)	0.1000 (0.0730)	0.0186 (0.0305)	-0.3615 (0.3004)	0.0844** (0.0376)
Income Group 2	-0.0770** (0.0369)	0.0637 (0.0834)	0.0559 (0.0351)	0.0502 (0.3004)	-0.0262 (0.0416)
Income Group 3	-0.0128 (0.0529)	0.0334 (0.0819)	0.0788** (0.0380)	0.2059 (0.2453)	0.0144 (0.0575)
Income Group 4	0.0081 (0.0529)	0.0650 (0.1179)	-0.0730 (0.0487)	0.3536 (0.2453)	-0.0838 (0.1344)
Income Group Effects	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X
Months Unemployed	X				
Months Others Sick		X			
Months Head Sick			X		
Months on Strike				X	
Months Laid Off					X
Observations	28,635	28,635	28,635	28,635	28,635
R ²	0.0369	0.0356	0.0358	0.0356	0.0357

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis.

Table 5. Food Consumption

Variables	(1)	(2)	(3)	(4)	(5)
Income Group 1	-0.0834*** (0.0247)	-0.0298 (0.0976)	-0.0072 (0.0410)	0.0171 (0.4017)	0.0088 (0.0507)
Income Group 2	-0.0929* (0.0494)	-0.0614 (0.1116)	-0.0411 (0.0471)	0.0634 (0.4017)	-0.0591 (0.0559)
Income Group 3	0.0061 (0.0708)	0.1787 (0.1095)	0.1190** (0.0509)	0.2644 (0.3280)	0.0310 (0.0769)
Income Group 4	-0.1281* (0.0708)	-0.0273 (0.1577)	-0.0214 (0.0651)	0.0643 (0.3280)	-0.0011 (0.1797)
Income Group Effects	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X
Months Unemployed	X				
Months Others Sick		X			
Months Head Sick			X		
Months on Strike				X	
Months Laid Off					X
Observations	28,531	28,531	28,531	28,531	28,531
R ²	0.0147	0.0142	0.0143	0.0141	0.0141

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis.

Table 6. Durable Consumption

Variables	(1)	(2)	(3)	(4)	(5)
Income Group 1	-0.1800*** (0.0279)	0.0625 (0.1104)	-0.0532 (0.0458)	-0.2480 (0.4473)	0.1310** (0.0569)
Income Group 2	-0.0114 (0.0554)	0.0076 (0.1243)	0.0392 (0.0529)	0.1104 (0.4473)	-0.0266 (0.0622)
Income Group 3	-0.0620 (0.0788)	0.0417 (0.1220)	0.0651 (0.0569)	-0.2237 (0.3652)	0.0607 (0.0856)
Income Group 4	0.0370 (0.0788)	-0.0351 (0.1756)	-0.0463 (0.0725)	0.2290 (0.3653)	-0.0456 (0.2001)
Income Group Effects	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X
Months Unemployed	X				
Months Others Sick		X			
Months Head Sick			X		
Months on Strike				X	
Months Laid Off					X
Observations	28,371	28,371	28,371	28,371	28,371
R ²	0.0403	0.0388	0.0389	0.0389	0.0390

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis.

Table 7. Nondurable Consumption with Government Benefits

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Income Group 1	-	-0.0325	0.0213	0.1177	0.0784**	0.0979***	0.1039	-0.2195	-0.0087	0.0783*
	0.0716***	(0.0253)	(0.0865)	(0.0992)	(0.0351)	(0.0384)	(0.4647)	(0.3283)	(0.0653)	(0.0451)
Income Group 2	-	0.1091	0.1291	0.1291	0.0239	0.0318	-0.1184	-0.1163	-0.0717	-0.0367
	0.1519***	(0.0432)	(0.0913)	(0.0949)	(0.0406)	(0.0395)	(0.3287)	(0.3284)	(0.0851)	(0.0478)
Income Group 3	-	-0.0370	0.0644	0.0642	0.0993**	0.0921**	0.1642	0.1664	-0.0929	-0.0466
	(0.0971)	(0.0578)	(0.0896)	(0.0895)	(0.0425)	(0.0422)	(0.2684)	(0.2681)	(0.1201)	(0.0634)
Income Group 4	-	-0.0313	0.0453	0.0450	0.0162	0.0147	0.2527	0.2514	-0.2212	-0.0005
	(0.0865)	(0.0570)	(0.1290)	(0.1289)	(0.0533)	(0.0532)	(0.2684)	(0.2681)	(0.2079)	(0.1469)
Group 1 Benefits	0.0213	-	0.1531	-0.2555	0.0093	-0.1410*	-0.6174	0	0.0953	-
	(0.0391)	0.1549***	(0.2250)	(0.1666)	(0.1121)	(0.0764)	(0.6572)	(0)	(0.0835)	0.2644**
Group 2 Benefits	-0.0866	0.2282*	0	-0.2424	-0.1862	-	0	0	0.0358	-0.0912
	(0.0802)	(0.1205)	(0)	(0.3417)	(0.1229)	(0.1598)	(0)	(0)	(0.1002)	(0.1544)
Group 3 Benefits	-0.0161	0	0	0	-0.2190	-0.0509	0	0	0.0744	0.4130
	(0.1205)	(0)	(0)	(0)	(0.2120)	(0.2358)	(0)	(0)	(0.1407)	(0.4685)
Group 4 Benefits	0.0472	0	0	0	0	0	0	0	0.4439	0
	(0.1159)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0.2939)	(0)
Income Group Effects	X	X	X	X	X	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X	X	X	X	X	X
Months Unemployed	X	X	X	X	X	X	X	X	X	X
Months Others Sick			X	X						
Months Head Sick				X	X	X	X	X		
Months on Strike									X	X
Months Laid Off									X	X
Unemployment Benefits	X		X		X	X	X			
Food Stamps		X		X		X		X		X
Observations	28,635	28,635	28,635	28,635	28,635	28,635	28,635	28,635	28,635	28,635
R ²	0.0226	0.0245	0.0220	0.0239	0.0224	0.0244	0.0220	0.0239	0.0222	0.0240

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis.

Table 8. Total Consumption with Government Benefits

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Income Group 1	-	-	0.0826	0.1894**	0.0256	0.0443	-0.0598	-0.3742	0.0463	0.1000**
	0.1260***	0.1042***								
	(0.0231)	(0.0217)	(0.0790)	(0.0907)	(0.0321)	(0.0351)	(0.4247)	(0.3002)	(0.0597)	(0.0412)
Income Group 2	0.0066	-0.0922**	0.0634	0.0666	0.0660*	0.0659*	0.0499	0.0506	-0.0920	-0.0212
	(0.0537)	(0.0395)	(0.0835)	(0.0868)	(0.0371)	(0.0362)	(0.3004)	(0.3002)	(0.0777)	(0.0437)
Income Group 3	0.0289	-0.0131	0.0321	0.0328	0.0897**	0.0816**	0.2048	0.2072	-0.0882	0.0107
	(0.0887)	(0.0529)	(0.0819)	(0.0819)	(0.0388)	(0.0386)	(0.2453)	(0.2451)	(0.1098)	(0.0580)
Income Group 4	0.0312	0.0074	0.0653	0.0655	-0.0727	-0.0734	0.3534	0.3530	-0.3388*	-0.0841
	(0.0790)	(0.0529)	(0.1179)	(0.1179)	(0.0487)	(0.0487)	(0.2452)	(0.2451)	(0.1900)	(0.1343)
Group 1 Benefits	0.0289	-0.0718*	0.0950	-0.2922*	-0.1094	-	-0.6110	0	0.0547	-0.1764*
	(0.0358)	(0.0382)	(0.2056)	(0.1523)	(0.1025)	0.1585**	(0.6006)	(0)	(0.0763)	(0.0993)
Group 2 Benefits	-0.1546**	0.1146	0	-0.0453	-0.0967	-0.1720	0	0	0.0911	-0.0619
	(0.0732)	(0.1101)	(0)	(0.3124)	(0.1123)	(0.1461)	(0)	(0)	(0.0916)	(0.1412)
Group 3 Benefits	-0.0658	0	0	0	-0.3005	-0.1029	0	0	0.1390	0.1721
	(0.1101)	(0)	(0)	(0)	(0.1938)	(0.2156)	(0)	(0)	(0.1286)	(0.4284)
Group 4 Benefits	-0.0409	0	0	0	0	0	0	0	0.5103*	0
	(0.1059)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0.2686)	(0)
Income Group Effects	X	X	X	X	X	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X	X	X	X	X	X
Months Unemployed	X									
Months Others Sick				X						
Months Head Sick					X					
Months on Strike							X	X		
Months Laid Off									X	X
Unemployment Bene- fits	X		X		X		X		X	
Food Stamps		X		X		X		X		X
Observations	28,635	28,635	28,635	28,635	28,635	28,635	28,635	28,635	28,635	28,635
R ²	0.0372	0.0381	0.0358	0.0370	0.0361	0.0373	0.0359	0.0370	0.0363	0.0371

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis.

Table 9. Food Consumption with Government Benefits

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Income Group 1	- 0.1257*** (0.0313)	-0.0551* (0.0290)	-0.0519 (0.1057)	0.1010 (0.1209)	-0.0152 (0.0432)	0.0447 (0.0468)	0.2351 (0.5681)	-0.0223 (0.4002)	-0.0679 (0.0807)	0.0774 (0.0552)
Income Group 2	-0.1235* (0.0719)	-0.1104** (0.0526)	-0.0641 (0.1116)	-0.0469 (0.1157)	-0.0319 (0.0498)	0.0111 (0.0484)	0.0613 (0.4017)	0.0641 (0.4002)	0.0613 (0.1057)	-0.0047 (0.0585)
Income Group 3	0.0218 (0.1186)	0.0047 (0.0705)	0.1784 (0.1095)	0.1763 (0.1091)	0.1233** (0.0519)	0.1241** (0.0515)	0.2642 (0.3280)	0.2679 (0.3268)	0.1261 (0.1468)	0.0304 (0.0773)
Income Group 4	-0.1140 (0.1057)	-0.1299* (0.0706)	-0.0276 (0.1577)	-0.0253 (0.1571)	-0.0217 (0.0651)	-0.0221 (0.0649)	0.0635 (0.3280)	0.0631 (0.3268)	-0.0908 (0.2541)	-0.0019 (0.1790)
Group 1 Benefits	0.0990**	-	0.1329	-	0.0537	-	-0.4416	0	0.1205	-
	(0.0481)	0.2406*** (0.0518)	(0.2750)	0.4915** (0.2031)	(0.1371)	0.3920*** (0.0946)	(0.8033)	(0)	(0.1029)	0.6841*** (0.1350)
Group 2 Benefits	0.0538	0.1198	0	-0.2201	-0.1086	-	0	0	-0.1684	-
	(0.0980)	(0.1469)	(0)	(0.4165)	(0.1503)	0.8785*** (0.1948)	(0)	(0)	(0.1240)	0.5984*** (0.1883)
Group 3 Benefits	-0.0244 (0.1473)	0 (0)	0 (0)	0 (0)	-0.1109 (0.2592)	-0.2141 (0.2874)	0 (0)	0 (0)	-0.1310 (0.1720)	-0.0878 (0.5709)
Group 4 Benefits	-0.0247 (0.1417)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0 (0)	0.1784 (0.3592)	0 (0)
Income Group Effects	X	X	X	X	X	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X	X	X	X	X	X
Months Unemployed	X	X	X	X	X	X	X	X	X	X
Months Others Sick				X						
Months Head Sick					X					
Months on Strike							X	X		
Months Laid Off									X	X
Unemployment Benefits	X		X		X		X		X	
Food Stamps		X		X		X		X		X
Observations	28,531	28,531	28,531	28,531	28,531	28,531	28,531	28,531	28,531	28,531
R ²	0.0150	0.0221	0.0143	0.0218	0.0145	0.0226	0.0142	0.0217	0.0144	0.0223

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis.

Table 10. Durable Consumption with Government Benefits

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Income Group 1	-	-	0.0232	0.0456	-0.0486	-0.0357	-0.0662	-0.2549	0.0832	0.1329**
	0.1880***	0.1800***	(0.1198)	(0.1383)	(0.0482)	(0.0527)	(0.6324)	(0.4473)	(0.0907)	(0.0626)
Income Group 2	0.1352*	-0.0204	0.0076	0.0106	0.0396	0.0593	0.1104	0.1096	-0.1535	-0.0269
	(0.0806)	(0.0593)	(0.1243)	(0.1293)	(0.0561)	(0.0545)	(0.4473)	(0.4473)	(0.1157)	(0.0654)
Income Group 3	0.0985	-0.0623	0.0395	0.0413	0.0831	0.0678	-0.2256	-0.2230	-0.0767	0.0636
	(0.1320)	(0.0788)	(0.1220)	(0.1220)	(0.0580)	(0.0578)	(0.3652)	(0.3652)	(0.1634)	(0.0864)
Income Group 4	0.0115	0.0367	-0.0354	-0.0347	-0.0467	-0.0464	0.2276	0.2289	-0.3487	-0.0457
	(0.1176)	(0.0788)	(0.1756)	(0.1756)	(0.0725)	(0.0725)	(0.3652)	(0.3652)	(0.2828)	(0.2001)
Group 1 Benefits	0.0074	-0.0203	0.2027	0.0246	-0.1390	-0.1003	-0.3828	0	0.0604	-0.0569
	(0.0543)	(0.0585)	(0.3070)	(0.2289)	(0.1527)	(0.1050)	(0.8943)	(0)	(0.1155)	(0.1484)
Group 2 Benefits	-0.2719**	0.0594	0	-0.0562	-0.0044	-0.3823*	0	0	0.1770	-0.0123
	(0.1099)	(0.1643)	(0)	(0.4655)	(0.1676)	(0.2302)	(0)	(0)	(0.1366)	(0.2105)
Group 3 Benefits	-0.2497	0	0	0	-0.4946*	-0.0955	0	0	0.1852	-0.1824
	(0.1639)	(0)	(0)	(0)	(0.2886)	(0.3213)	(0)	(0)	(0.1914)	(0.6382)
Group 4 Benefits	0.0451	0	0	0	0	0	0	0	0.6046	0
	(0.1577)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0.0173)	(0)
Income Group Effects	X	X	X	X	X	X	X	X	X	X
Year Fixed Effects	X	X	X	X	X	X	X	X	X	X
Months Unemployed	X	X	X	X	X	X	X	X	X	X
Months Others Sick										
Months Head Sick					X	X				
Months on Strike							X	X		
Months Laid Off									X	X
Unemployment Bene- fits	X		X		X		X		X	
Food Stamps		X		X		X		X		X
Observations	28,371	28,371	28,371	28,371	28,371	28,371	28,371	28,371	28,371	28,371
R ²	0.0407	0.0404	0.0393	0.0391	0.0395	0.0392	0.0393	0.0391	0.0399	0.0392

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are in parenthesis.

Chapter 3: Consumption Insurance Constraints

Logan Miller¹ Don Koh² Teng Ma³

1 Introduction

Various consumption insurance channels are available for households to insure their consumption against income risks. An extensive literature has already shown that there are various mechanisms playing a key role in insuring consumption against income risks.⁴ On the other hand, another important literature shows that a full-insurance hypothesis is rejected in both developed and developing countries (Mace, 1991; Cochrane, 1991; Townsend, 1994; Attanasio and Davis, 1996). In explaining the absence of full insurance, the previous research both in theory and empirics sheds light on the absence of insurance mechanisms, while little attention has been paid to the barriers of using available insurance mechanisms. In this paper, we empirically demonstrate that households are constrained in using spouse's labor supply—both extensive and intensive margins—and self-insurance in response to household head's income shocks. Further, given these empirical evidence of insurance constraints, we quantify the welfare loss of each constraint in a two-good, two-earners life-cycle model.

Testing the consumption responses to income shocks requires household-level panel data. In this paper, we use Panel Survey of Income Dynamics (PSID), which is often used in this literature due to its long and detailed panel structure. Using the detailed and disaggregated information about household's economic variables, we examine the response of consumption expenditures by items, of saving and borrowing, of wife's labor supply, and of public and private transfers when household head receives an income shock.

¹University of Arkansas: lmiller@walton.uark.edu

²University of Arkansas: dkoh@walton.uark.edu

³Emporia State University:

⁴For example, extended family network (Kotlikoff and Spivak, 1981; Hayashi et al., 1996; Altonji et al., 1992), durable purchases (Browning and Crossley, 2009), public transfers (Engen and Gruber, 2001; Blundell and Pistaferri, 2003), the progressive income tax Kimball and Mankiw (1989); Auerbach and Feenberg (2000), and added worker effects (Stephens, 2002; Blundell et al., 2016, 2018)

Regarding the identification of income shocks, it is more common to use permanent and transitory components of income shocks from actual income changes in the literature. But we instead exploit unemployment experiences of household heads as an income shock due to its nature of income shocks to households and time consistency in data collection. In PSID, labor income and working hours of household head and wife are reported retrospectively in the previous year, while consumption is for the current interview year. On the other hand, households' employment status is asked both for the current and retrospective years. Therefore, to match the timing of consumption, working hours, and wealth, unemployment experiences are more suitable as an income shock. Although the number of unemployment experiences per household isn't abundant, households respond more sensitively to their unemployment experiences and any minor income changes.

An overview of our empirical finding is as follows. Regarding consumption insurance channels, there are three findings: (1) Households can completely smooth "food-at-home" expenditure against household head's unemployment; (2) public and private transfers are effectively used to insure against household head's unemployment shocks for both married and non-married households; and (3) consumption items with high income elasticity is another important insurance channel for food consumption.

In PSID, we show that non-married household cut down consumption expenditures of non-essential items without using their self-insurance. In contrast, married households do not significantly cut down consumption expenditures of non-essential items. In this respect, married households are successful smoothing their consumption expenditures. In principle, married households have wife's labor supply as an additional insurance, while non-married households don't. Surprisingly, however, their self-insurance and wife's labor supply of extensive and intensive margins are not significantly responding to unemployment shocks on average. This result indicates that (some of) married households are constrained in using these available insurance channels. Regarding the insurance constraints, there are three

findings: (1) number and age of kids are constraints for mom's extensive margin of labor supply; (2) current wife's working hours are a constraint for the marginal increase of wife's working hours; and (3) self-insurance is effectively used only for low-saving households. The first result is consistent with the previous literature (Blundell et al., 2018), while the second result hasn't been highlighted in the previous research. Also, the idea of the third finding is in line with the saving-constraint in Miranda-Pinto et al. (2020) but contradicts to the model implication in Ortigueira and Siassi (2013) where they demonstrate that two-earners model of incomplete market model shows that wealth-poor households rely more on spousal labor supply than saving. We empirically test if the use of spousal labor supply and saving differ by wealth level of households.

The main contribution of our paper is to shed light on the sources of constraints on various consumption insurance channels and the welfare loss of those constraints. There is an extensive research on the test of full insurance models (Mace, 1991; Cochrane, 1991; Townsend, 1994; Attanasio and Davis, 1996). This literature shows that the full insurance models are rejected in both developed and developing countries. But, it is beyond the scope of this literature to identify the source of market incompleteness. On the other hand, the recent interest shifted to the permanent income hypothesis (Hall and Mishkin, 1982; Altonji and Siow, 1987; Deaton, 1992; Dynarski and Gruber, 1997) This literature empirically tests the sensitivity of household consumption to different persistence of wage shocks (Blundell et al., 2008; Kaplan and Violante, 2010). But again, this literature cannot identify the source of the sensitivity. The aim of this paper is to pin down the sources of constraints on various insurance channels.

2 Empirical Evidence on Insurance Constraints

2.1 Data Construction

For the empirical analysis of insurance constraints, we use the PSID as our main data source. The PSID contains a long panel survey of representative U.S. households since 1978. Respondents were surveyed annually prior to 1999 and have been surveyed biannually since. The only consumption item the PSID includes before 1999 is food expenditure, but it has extended household expenditure items beyond food since 1999. Our interest is to explore the presence of consumption insurance constraints which inhibit households from fully using insurance mechanisms.

2.1.1 Consumption Expenditure by Items

Consumption items with different income elasticity and intertemporal elasticity of substitution may respond differently to an income shock. Since the 1999 survey, PSID data has incorporated detailed features of household consumption expenditures that extend beyond food expenditure. We take advantage of this detailed expenditure data to estimate household consumption responses to income shocks. More specifically, consumption items reported on by the PSID can be split into 10 broad categories: food, housing, utilities, transportation, education, childcare, health care, clothing, trips and other recreation. Most of these categories are made up of a number of more specific expenditures however, education, childcare, clothing, trips, and other recreation do not have any smaller expenditure items comprising these broader categories. Given the detailed information about household expenditures, we can split consumption into durable and nondurable. Durable consumption expenditures, housing and vehicles, will be reported in the illiquid asset category as they are expenditures on an illiquid asset (a house or car). Then, we further disaggregate nondurable consumption into food and nonfood expenditures when constructing the budget component.

Food consumption is the sum of expenditures on food at home, food away from home, and food delivered. Non-food expenditures are made up of the following items. First, we include utilities, which is the sum of expenditures on gas for the home, electricity, water/sewer, telephone/internet, and any other utilities. Next, we include components of the transportation category that are not expenditures on an illiquid asset. These include gasoline, parking, bus/train, taxicabs, and any other transportation expenses. Then, we include any expenditures on health care, which is the sum of expenditures on a hospital visit/nursing home, doctor, prescriptions, and health insurance. Finally, we include education, childcare, clothing, trips, and other recreation expenditures.

Durable consumption consists of expenditures on housing and vehicles. Housing expenditures include mortgage, rent, property tax, home insurance, house repairs, and house furnishings. Vehicle expenditures include loan payments, down payment, lease payments, auto insurance, additional vehicle costs, and vehicle repairs. Therefore, we include any monthly or annual fixed costs of the durable goods which hardly change as a result of temporary income shocks but are rather smooth over time and are nonzero – such as, mortgage, homeowner’s insurance, rent, property tax for housing, vehicle loan payment, vehicle down payment, lease payment, and insurance for vehicles.

Households report their consumption during interviews conducted in the current year. Surveyors ask respondents how much they spend per week, per month, or per year on expenditure items. The survey questions take different forms depending on the consumption item in question. The survey question regarding food consumed at home asks, ”How much do you spend on that food in an average week? How much do you (and everyone else in your family) spend on food that you use at home in an average week?” Whereas the survey question for gasoline expenditures (transportation) asks, ”Altogether for the month of (LAST MONTH), how much did you (and your family living there) pay for each of these transportation related expenses. (Do not include costs associated with trips and vacations, which we will ask about

Table 1. Cumulative Aggregation of Household Consumption by Item

Food Consumption	Non-food Consumption	Durable Consumption
Food at home + Food away from home + Food delivered	Utility + Transportation + Health + Clothing + Education + Child care + Recreation	Housing + Vehicle

later.) Gasoline?” Then, the surveyor would ask about other transportation related expenditures in the same manner. Other items, like childcare, are an annual total. All missing items are imputed by the administrators. It should be noted that all expenditures reported are for the current year.

Despite the advantageous features of household consumption expenditures in the PSID, there are several caveats to keep in mind. First, the PSID survey on household consumption is measured in a market-value expenditure. So, we denominate the household expenditures by the Consumer Price Index (CPI) of each year, which of course is not identical to the actual quantity of consumption. Then, we adjust the expenditures using the OECD-modified adult equivalence scale. Second, the consumption survey is retrospective of the consumption from a week, a month, or year before the interview, but still within the “current” year. The income survey is a retrospective report from the previous year. Therefore, there is a timing inconsistency between consumption and income variables. Nevertheless, our empirical exercise is intact from this time inconsistency because we use unemployment episodes as an income shock to households, and these are surveyed in a consistent year with the consumption survey. Third, as we disaggregate household consumption into items, the data is subject to measurement errors, noise, and zero expenditures. To get around these problems, we subsequently aggregate items as shown in Table 1 and as we explain below.

Despite our aggregation, some consumption items may still include zero expenditures for

certain households. For example, households with no kids would spend nothing for child care. Also, some households may spend nothing for medical services for a particular year. The presence of zero-values in the data however, causes problems in our empirical estimation (see Section 2.1.2). To prevent zero consumption from affecting our results, we start our empirical analysis with basic consumption items that rarely contain zero expenditures. For example, we choose “food at home” as a basic food consumption measure because there are almost no households who report zero consumption on “food at home.” Given this basic consumption, we subsequently add “food away from home” and “food delivered” to estimate the effect of income shocks on these items. This way, we can minimize the number of households with zero food consumption. In the same manner, we take “utility” expenditure as a basic non-food consumption measure and subsequently add expenditures on “transportation,” “health,” “clothing,” “education,” “child care,” and “recreation.” For durable consumption, we choose “housing expenditures” as a benchmark and add “vehicle expenditures.”⁵

2.1.2 Source of Income Shocks

In the recent literature, there are two main approaches used to identify idiosyncratic income shocks. One is to take a specific shock that directly affects household income, such as fiscal stimulus, changes in health, or unemployment. Another, more common way, of characterizing an income shock is by distinguishing a shock by its persistence. (Meghir and Pistaferri, 2004; Blundell et al., 2008). However, identifying labor income processes using the variance-covariance of income residuals has a problem. The residual income removes observable household characteristics and reduces the shock to single fluctuations of labor income that cannot be explained by changes in observable characteristics. These changes in residual income are then assumed to be an exogenously driven income shock by econometri-

⁵Despite subsequently adding consumption items to the benchmark items, we cannot completely eradicate zero consumption expenditures from our consumption items. Therefore, we implement two additional techniques to mitigate the zero-value effect: using an inverse hyperbolic sine function and a median regression for our estimation. Both are described in detail in the following sections.

cians, while they are more likely endogenous for households (Low et al., 2010; Huggett et al., 2011).

In this study, we focus on unemployment experiences of household heads as an income shock, as in Cochrane (1991) and Mace (1991), which are precisely identified in the PSID data. It should be noted that several problems with using unemployment shocks as a source of income shocks have already been pointed out in the literature. First, unemployment shocks may not be pure unanticipated shocks to households, even though they are considered a shock by econometricians. If a layoff is, to a certain degree, anticipated by households they can plan for the outcomes of that shock. Our biannual household survey data can mitigate this effect because the moment households receive information about their future layoff to the moment their consumption responds as a result of that information could belong to a two-year period contained within the household survey. Next, unemployment episodes that households experience are not abundant in the data. In our data after sample selection, the average unemployment experiences per household is 0.15 times during 1999-2017. We use 483 total unemployment experiences by households to analyze the consumption and insurance responses. Finally, the most important problem with using unemployment shocks as a source of income shocks is that the persistence of the shock is ambiguous. Some households experience very short-term (less than 1 month) unemployment, while others may experience much longer unemployment. Nevertheless, unemployment shocks are considered as temporary shocks in the literature because household heads find a new job in 5.9 months after unemployed in our data.

Timing of unemployment shocks We identify the income shock from unemployment experiences by household heads. The survey asks if a household head is currently unemployed in the current interview year as well as in the previous year. The current employment status coincides with the current consumption expenditures that households report in the survey, while the previous year's employment status coincides with household's retrospective

income in the data. Therefore, we can exploit unemployment experiences as an income shock to identify the effect on consumption, income, and insurance mechanisms. We focus on household heads who are currently unemployed but were employed in the previous interview (E-U). Then we test a complete market hypothesis by checking if E-U households have significantly negative consumption growth between the previous and current interview. EE or UU households have also experienced unemployment but we do not use those observations to identify income shocks. The survey asks if a household was unemployed in the previous year, but the effect of unemployment in the previous year on consumption growth would depend on the months the household spent after finding a job. We find that both the months after finding a job and unemployment in the previous year appear irrelevant for the consumption growth between the previous interview and the current interview.

Duration of unemployment In PSID, households are asked if the head is unemployed in each month of previous year. We use this information to compute the average duration of unemployment, up to 12 months. Our data shows that households are unemployed 5.76 months on average, which indicates a lower bound from the data limitation. Given the duration of unemployment, it is evident that unemployment shocks are interpreted as transitory shocks.

2.1.3 Consumption Insurance Channels

Households use a set of consumption insurance channels to insure consumption upon receiving an income shock. We select household consumption insurance channels from the household budget components: wife's labor supply, saving/borrowing, and public/private transfers.

Wife's Working Hours (retro) Of all the possible insurance channels, one of the most important channels that households exploit as an insurance is wife's labor income (Blundell et al., 2016). We separate the extensive margin of wife's labor supply from the intensive margin to check which margin responds more strongly to household head's unemployment. For the extensive margin, we exploit wife's employment status in the current interview year comparing to the status in the previous interview year. For the intensive margin, wife reports her working hours retrospectively in the previous year. For the time consistency, we use household head's unemployment status in the previous year instead of current interview year.

Public transfers (retro) Transfers to the household from government sources are categorized as public transfers. The household survey asks respondents to indicate how much government benefits the household head, spouse, and other family members each received, retrospectively, in the previous year. Typically, the main insurance for unemployed household head must be unemployment insurance. We construct a binary variable based on the information if household head received unemployment benefits while they are unemployed.

Private transfers (retro) Household private transfers are transfers received from non-governmental sources. The household survey asks respondents how much transfers the household head, spouse, and other family members each received, retrospectively, in the previous year. For example, when reporting child support transfers for the household head in 2013, this value would be the sum of all child support transfers to the household head that would have taken place over the course of 2012. These transfers are: child support, alimony, income from trusts, retirement account payouts, annuities, pensions, help received from relatives, help received from others (not relatives), individual retirement accounts (IRAs), miscellaneous transfers, and any other income. Child support amounts are reported for the household head, spouse, and other family members occupying the household. Alimony amounts are

reported for the household head and spouse. Trust income is reported for the household head and spouse. Retirement account payouts are reported for the household head, spouse, and other family members. Income from annuities is reported for the household head, spouse, and other family members. Income from pensions is reported for the household head and spouse. Help received from relatives is reported for the household head, spouse, and other family members. Help received from others is reported for the household head and spouse. Income from IRAs is reported for the household head and spouse. Miscellaneous transfer amounts are reported for the household head, spouse, and other family members. Finally, any other income values are reported for the household head and spouse. In our exercise, we focus on help received from relatives and others as main sources of private transfers.

Saving/borrowing We also consider saving and borrowing (debt) as important insurance channels for households. We separate household saving into liquid assets, debts, and illiquid assets. Liquid assets are the sum of the amount of money in checking, savings, CDs. Illiquid assets are the sum of profit from investment real estate, bonds/insurance, pension/annuity/ira, business/farm profit, non-ira stocks. Household debts are the sum of credit cards, student loans, medical debt, legal debt, debt owed to relatives, and remaining mortgage amount. We are evaluating the growth rate of these variables from one survey to the next. However, many households contain zero values on these variables, particularly debts and illiquid assets. Therefore, we first take liquid assets as our benchmark measure of self-insurance and then subsequently subtracting debts from liquid assets to construct net-saving. Further, we add the amount of illiquid assets to net-saving to construct the household asset position and its response to unemployment shocks.

Summary statistics Table 2 shows the summary statistics of our sample after sample selection (see Appendix for sample selection). We compare the mean and median of all households' characteristics, consumption, assets & debts, labor supply, transfers, and unem-

ployment duration for all households with those households who received an unemployment shock. While there are overall 22,216 households in our sample, there are only 483 households who experience unemployment in the interview year.⁶

Regarding the households characteristics, all the household heads are male by sample selection. The average age of household heads are 44.7 years old, but unemployed households are 42.8 years old. The average years of schooling is 13.9 years, which is slightly above high school graduates, and unemployed households also have quite similar years of schooling of 13.3 years. On average, 82% of households are married, while only 69% are married among unemployed households. The size of family unit is close to 3 including both married and non-married. The number of children is slightly less than 1. The difference in household characteristics between all households and unemployed households are minor.

We also calculate the consumption and assets/debts in \$ thousands. The average annual food consumption is around \$5,520, while the unemployed households consume \$4,840. Non-food and durable consumption also show that unemployed households consume approximately \$100-200 less than the average households. The difference indicates that on average, unemployed households (either married or non-married) cut down their consumption expenditures. Similarly, unemployed households exhibit on average \$765 lower (\$2,661 - \$1,896) in liquid assets and \$6,435 lower (\$1,1574 - \$5,139) in illiquid assets. Also, they have \$600 higher debts than the average households. It is not evident from the summary statistics that unemployed households have used saving and borrowing as an insurance, or they originally had a lower saving and more debts before unemployed.

Next, we compute the mean and median of labor supply for household head, spouse, and other family members. The average annual hours of work is 2,010 hours for the head and 1,170 hours of work for wives. For the unemployed households, wives work an average

⁶Retrospectively, the number of households who experienced at least one month of unemployment in the previous year is 466. We use either current year unemployment or previous year's unemployment depending on the response variables.

1,000 hours. In terms of labor income, household heads earn \$63,120 on average and wives earn almost more than half of household heads, \$28,380. Wives in unemployed households earn even less than the average, \$22,420. Other family members add only \$2,980 to the households income.

In terms of transfers or help received from relatives and others, 11% of households and 18% of unemployed households receive transfers. Also, 44% of unemployed households receive unemployment benefits, while only 13% of unemployed households receive food stamps. Finally, the unemployed households experience on average 5.9 month of unemployment duration and 5 months are the median.

2.2 Estimation Methods

Using the data constructed in the previous section, we examine if household's response to unemployment shocks varies by a set of consumption insurance to which they have access. In this study we focus on three particular consumption insurance mechanisms. The first insurance is wife's labor supply. Blundell et al. (2016) and Krueger and Wu (2021) argue that spouse's labor supply play a major role in insuring against permanent income shocks. However, it is trivial that single households with no partners have no access to this insurance. This implies that marital status is one of constraints for a household to use spouse's labor supply as an insurance. Therefore, we compare responses of non-married households who have no access to this particular insurance (constrained group) with the responses of married households (non-constrained group).⁷

Even among married households who have access to wife's labor supply *in principle*,

⁷Unfortunately, in PSID, we cannot distinguish single households with no partners from single households with partners. Non-married households include "never married," "widowed," "divorced, annulled," and "separated." In fact, the number in family unit allows us to identify a single household. However, there are only 24,672 observations for the households with only one family unit, as opposed to 99,532 observations for non-married households. Therefore, we decided to use non-married as a constrained group in our exercise to ensure enough number of observations.

Table 2. Summary Statistics

	All Households		Unemployed Households	
	Mean	Median	Mean	Median
Characteristics of HH Head				
Gender (1: male, 2: female)	1	1	1	1
Age	44.7	45	42.8	42
Years of schooling	13.9	14	13.3	13
Marital status (1: married, 0: single)	0.82	-	0.69	-
Size of family units	2.99	3	2.87	3
Number of children	0.92	0	0.93	0
Consumption (\$ thousands)				
Nondurable				
Food	5.52	5.01	4.84	4.24
Non-food	11.16	9.12	10.20	8.28
Durable	15.82	13.19	13.96	11.26
Assets & Debt (\$ thousands)				
Liquid assets	26.61	5.65	18.96	2.83
Illiquid assets	115.74	0.00	51.39	0.00
Debt	6.03	0.00	6.66	0.00
Labor Supply				
Annual Hours of Work (thousands)				
HH head	2.01	2.08	-	-
HH spouse	1.17	1.39	1.00	0.81
Annual Labor Income (\$ thousands)				
HH head	63.12	51.64	-	-
HH spouse	28.38	19.27	22.42	6.52
Other HH members	2.98	0.00	2.88	0.00
Transfers (1: received, 0: not received)				
Private transfers	0.11	-	0.18	-
Public transfers				
Unemployment benefits	0.07	-	0.44	-
Food stamps	0.04	-	0.13	-
Unemployment				
Duration	-	-	5.90	5
No. of HH	22216	22216	483	483

Note: We calculate the mean and median of each variable weighted by the inverse of PSID longitudinal sample weights. Consumption, net assets, and income are denominated by consumer price index of each year. “Unemployed Households” on the third and fourth columns indicates households whose head receives an unemployment shock in the current year, while he is employed in the previous year. The median of binary variables is not reported.

some married households are constrained in using the spouse’s labor supply in response to unemployment shocks. For example, there are some factors that hinder wife’s labor participation (extensive margin of labor supply): family care-giving, less work opportunities

than men, or health issues. Likewise, wife’s intensive margin could also be constrained by similar factors, which make the marginal cost of increasing additional working hours very costly. For example, a full-time job cannot flexibly change the working hours for additional salary. Even with a part-time job, the marginal cost of increasing an additional hour can be costlier when the current working hours are high. We aim at identifying these constraints that hinder the use of wife’s labor supply as an insurance.

Another effective consumption insurance that households count on is a current amount of self-insurance—savings (either liquid or illiquid assets) and borrowing. Nevertheless, not all the households take advantage of self-insurance in the first place when unemployed. Again, there are many reasons for the households not to use their self-insurance. For example, due to uncertain unemployment duration, households become precautionary to not dissave right away. Also, households can be credit-constrained in borrowing and hence keep their saving for the longer unemployment duration. In this exercise, we investigate the responses of households by household characteristics and current asset positions to see which households rely on self-insurance when unemployed. Therefore, in this exercise, we split the sample into “constrained” and “non-constrained” groups to contrast their responses to unemployment shocks. And then we delve into actual impediments that hinder households to use particular insurance channels.

Inverse hyperbolic sine function Household data on consumption, assets, and earnings (not surprisingly) contain a large number of zeros. With the zero values, however, taking a log difference as a measure of household response to unemployment shock generates outliers on both sides of distributions, which in turn distort our estimation. To mitigate the bias from outliers on our estimates, we apply an inverse hyperbolic sine (IHS) function to compute the growth rate:

$$\Delta x_{i,t+1} = \log(x_{i,t+1} + (x_{i,t+1}^2 + 1)^{0.5}) - \log(x_{i,t} + (x_{i,t}^2 + 1)^{0.5})$$

While both IHS and log functions are monotonic, the size of growth rate differs because log function transforms low values to a large negative number. Nevertheless, the ordering of consumption growth across households do not change. Also, IHS difference narrows the distribution of consumption growth and hence standard errors of estimates narrow down. Therefore, the statistical inference of estimates should also change between IHS difference and log difference.

Empirical model A recent common practice in estimating the degree of consumption insurance against income shocks is a partial insurance testing (Blundell et al., 2008). However, this practice is not well suited for our exercise because the linear approximation of intertemporal optimality condition that Blundell et al. (2008) use to derive a partial insurance equation assumes that liquidity constraints are not binding. In contrast, our interest is to examine the role of a particular insurance for two different groups, the one who are constrained in accessing particular insurance channels and the other who are not constrained. The constrained households may not have an equality in their optimality conditions. Therefore, we run a simple regression model to compare the household responses of consumption, asset accumulation, labor supply, and public/private transfers to unemployment shocks between constrained and non-constrained groups, $\{\mathcal{G}_0, \mathcal{G}_1\}$.

Our model specification is the following:

$$\Delta Y_t^i = \alpha_t + \alpha_c + \gamma' X_t^i + \sum_{c=0}^1 \beta_c (E_t^i \times \mathbf{I}(i \in \mathcal{G}_c)) + \varepsilon_t^i \quad (1)$$

where the dependent variable, ΔY_t^i , takes consumption by item, liquid and non-liquid asset, wife's labor supply, and public and private transfers ($Y_t^i \in \{c_t^i, a_t^i, h_{2,t}^i, T_{g,t}^i, T_{p,t}^i\}$) for a household i from year t to $t + 2$ in IHS difference; α_t indicates a year fixed effect; α_c indicates a constrained-group fixed effect; X_t^i indicates observable characteristics of household i in year t which includes age and age squared, gender of household head, a marital status, the number

of family units, the number of kids, and a residential area; E_t^i is a binary indicator that flags 1 if the head of household i was employed in year t and become unemployed in year $t + 2$; and ε_t^i denotes unobserved shocks and measurement errors from an iid normal distribution. Our interest is a parameter β_c that can be interpreted as a response of household in a group c given an unemployment shock in year $t + 2$.

Although our sample selection removes inexplicable data observations, household responses to unemployment shocks are widely dispersed around zero. This raises a concern that a few outliers may critically affect the estimation results even after taking IHS differences. Therefore, we use median regression in this exercise to mitigate the effect of outliers on the estimates.⁸ In addition, even though we control for year effect, there is a concern about heteroskedasticity. We report bootstrapped standard errors.⁹

2.3 Empirical Results

2.3.1 No Access to Wife’s Labor Supply

Spouse’s labor supply—both extensive and intensive margins—is one of consumption insurance channels that households heavily count on (Blundell et al., 2016). However, a non-married (single) household by nature has no access to this particular insurance. In this section, we examine how non-married households respond to unemployment shocks differently from married households.

A stark difference between married and non-married households in their responses to unemployment shocks is shown in Table 3. For this estimation, we run a median regression

⁸To check if a few outliers do not affect our results, we estimate a parameter β_c by using both mean regression and a median regression in Table 5.

⁹There are two concerns about statistical inference. First, there will be heteroskedasticity not only over time but also by group. Variance of consumption growth should be larger for higher income groups than lower income groups. Therefore, homoskedasticity does not hold. In this case, we need to use Eicker-White standard errors. Second, the household consumption growth might be related within the income group, but not across income groups. Therefore, data could be clustered and hence we need to use clustered standard errors. This clustering is determined by the PSID sampling design.

for all non-binary variables.¹⁰ For the binary variables (e.g., extensive margin of wife's labor supply, public & private transfers) we run a linear probability model with the same model specification as equation (1), but with a mean regression.¹¹ To begin with, "food at home" does not significantly respond to the shock for both married and non-married households. Note that the "food at home" variable in the data is measured by a market value denominated by the price index. This result supports the presence of a set of consumption insurance that smooths out not only the quantity of consumption but also the market-value expenditures of necessary good under income changes. In other words, both married and non-married households would not substitute their food items with cheaper ones when unemployed.

Married households can also insure other food and non-food consumption expenditures when household head is unemployed, even though the signs of estimates are all negative. This represents evidence that married households exploit a set of consumption insurance and can successfully smooth out those consumption expenditures. In contrast, non-married households show that their "food away from home" and "food delivered" expenditures are significantly cut down by 16.3% and 16.8%, respectively. Not only those but also expenditures on "transportation," "education," "child care," and "recreation" are significantly reduced by 10-11% upon receiving an unemployment shock. One commonality of these items is that they are non-essential consumption items relative to food consumption. This implies that non-married household not only use consumption insurance channels to smooth out consumption but also use an expenditure cut on non-essential items to smooth out food consumption. Durable consumption is again non-responsive to unemployment shocks for married and non-married households.

The third column of Table 3 shows a one-side testing of the estimates of married and

¹⁰Robustness checks with mean regression, individual fixed effect, and no controls of individual characteristics are in the appendix. Mean regression are highly sensitive to outliers in the data and therefore, the estimates are larger in absolute value than median regression. Individual fixed effects also alter the significance of estimates. This is due to the fact that unemployment experiences per household is too few that we may not have enough observations to run a credible estimation.

¹¹Robustness checks with a probit regression in the appendix.

Table 3. Household Responses to Unemployment Shocks by Marital Status

	Non-married (1)	Married (2)	Hypothesis Testing $H_0: (2) - (1) \leq 0$
Food Consumption			
Food at home	-0.037	0.023	0.060
+ Food away from home	-0.163**	-0.038	0.125*
+ Food delivered	-0.168**	-0.021	0.146**
Non-food Consumption			
Utility	-0.058	-0.011	0.047
+ Transportation	-0.107*	-0.010	0.097**
+ Health	-0.075	-0.006	0.068
+ Clothing	-0.092	0.004	0.096*
+ Education	-0.102**	-0.007	0.095**
+ Child care	-0.105***	-0.037	0.069*
+ Recreation	-0.105**	-0.007	0.097**
Durable Consumption			
Housing	-0.054	0.001	0.056
+ Vehicle	-0.032	-0.017	0.017
Asset Accumulation			
Liquid	-0.028	-0.002	0.027
- Debt	-0.013	0.016	0.029
+ Illiquid	-0.025	0.018	0.043
Spouse's Labor Supply			
Intensive Margin	-	0.004	-
Extensive Margin	-	0.028*	-
Public & Private Transfers			
Unemployment Benefit	0.353***	0.445***	0.092**
Transfers from Relatives	0.123***	0.063***	-0.060

Note: Table shows the responses of consumption by items, assets, spouse's labor supply, and public/private transfers to unemployment shocks by married and non-married households. The third column shows statistical difference of married and non-married estimates (column (2) minus column (1)). The null hypothesis is that the decline of items for married is larger than the decline for non-married ($(2) \leq (1)$). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

non-married households. The null hypothesis is to claim that the consumption drop for non-married households (who have no access to wife's labor supply) is no larger than the consumption drop of married households (who have access to wife's labor supply). Since married households have one more important insurance channel through their spouse's working hour, we conjecture that the response to unemployment shock of married households should be less stark than the response of non-married. The size of consumption drop is overall larger

for the non-married households than married. However, the consumption drop of “food away from home,” “food delivered,” “transportation,” “clothing,” “education,” “child care,” and “recreation” are the ones that are significantly larger for the non-married than married. To summarize, both married and non-married have an ample set of consumption insurance to smooth out their consumption on necessary goods. On the contrary, when households are constrained in the access of spouse’s labor supply, they tend to cut down the expenditures on items with high income elasticity.

Next, we examine the response of available consumption insurance channels: saving/borrowing, spousal labor supply, and public/private transfers. First, both married and non-married households receive unemployment benefits from the government when unemployed. The estimates tell us that 44.5% of married households whose head is unemployed receive the benefits, while 35.3% of non-married households receive the benefits.¹² The difference of estimates between married and non-married are 9.3% significant higher for married households than non-married. Any help from relatives and others are also significant for both married and non-married when they get unemployed. Though non-significant, the size of transfers from relatives for non-married is approximately twice as large as that for married households.

We now examine the response of self-insurance to unemployment shocks. “Asset Accumulation” in the table shows no significant use of saving/borrowing for both married and non-married households. Although non-married households use 2.8% of their liquid assets from the previous period, 1.3% net of debt, and 2.5% including illiquid assets when unemployed, all these values are not statistically different from zero. For married, none of their response in these assets are significant. There are several hypotheses for this result. First, saving/borrowing could be the last resort for households to insure their consumption. Sec-

¹²The reason for this low reception rate of unemployment benefits could be partly due to application lags. Some households report that they are unemployed and applied for the benefits but not approved at the time of interview. Another reason could be that the households are actively searching for a new job and may consider the unemployment duration to be short.

ond, unemployment duration is not long enough to necessitate the use of assets to supplement income losses for both households. Third, married households would become precautionary to save for longer duration of unemployment. We will test these hypotheses later in this section by disaggregating the households into subsamples.

Finally, spouse's labor supply is known to be one of the most important insurance channels, but only available to married households. Surprisingly, however, the intensive margin of spouse's labor supply does not respond significantly to the unemployment shock, while the extensive margin increases by 2.8%. One hypothesis behind this result is that even if married households have access to spouse's labor supply, the use of the spouse's labor supply is constrained. To summarize our findings in this table:

Fact 1: *Households can completely smooth "food-at-home" expenditure against household head's unemployment.*

Fact 2: *Public and private transfers are effectively used to insure against household head's unemployment shocks for both married and non-married households.*

Fact 3: *Consumption items with high income elasticity are another important insurance channel for food consumption.*

2.3.2 Constraints on the Use of Wife's Labor Supply

We will now investigate any potential constraints on the spouse's labor supply in the later section. Wife's labor supply is one of the important insurance channels for married households (Blundell et al., 2016). However, we found no significant use of wife's intensive margin and little use of extensive margin as an insurance when household head is unemployed in Table 3. This result implies that married households could have been restricted to use either

extensive or intensive margin of wife's labor supply despite the potential access to wife's labor supply. In this section, we investigate which households are particularly constrained in using the wife's labor supply as an insurance by disaggregating the married household into subsamples.

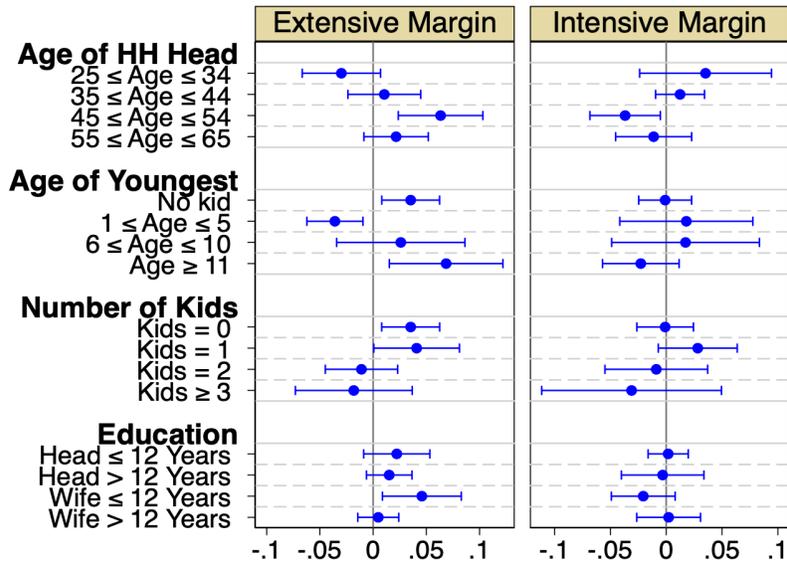
Figure 1 shows the response of wife's intensive and extensive margins of labor supply by household characteristics. For the estimation of effects by subsamples, we run a regression described in Section 2.2 in two stages. At the first stage, we control household characteristics with year/group fixed effects. Then, at the second stage, we run a regression on the residuals of unemployment effects by subsamples to estimate the response of wife's labor supply.

Note that the extensive margin of labor supply is a binary variable that household wives are either employed or non-employed. Since the estimation of effect with a binary variable is less affected by outliers, we apply a mean regression to run a linear probability model. The left figure in Figure 1 shows that the extensive margin responds significantly for a household with the household head age between 45 to 54. More notable result is that wives with no kids or the youngest kids above age 11 are more likely to participate the labor market than wives with younger-age kids. In terms of the number of kids, wives are more likely to start working when the number of kids are less than two. The kid's effect can match with the age of household head around 45 to 54. Regarding the education level of household head and wife, wives with less than 12 years of education tend to find a work when the head is unemployed. This result is consistent with our understanding that the number and the age of kids are one of the most critical constraints for moms to start working.

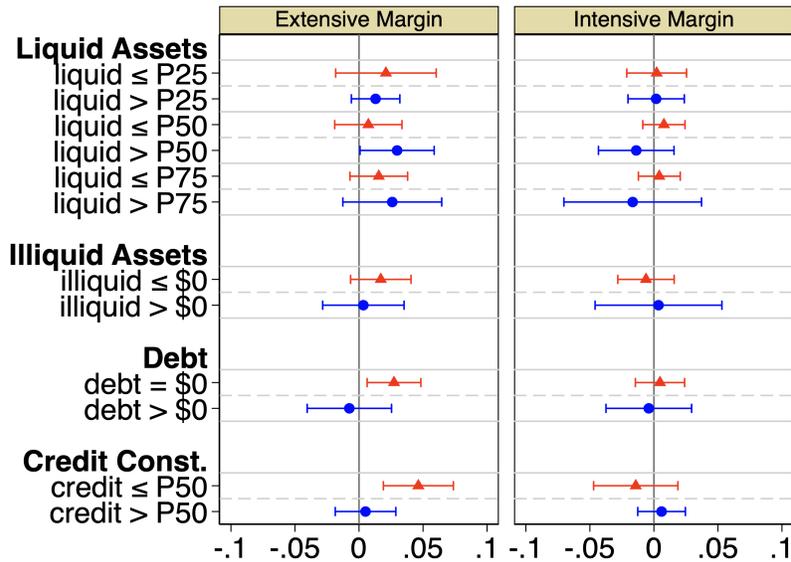
In contrast, the intensive margin of wife's labor supply on the right column of Figure 1(a) shows no significance in any of the household characteristics except the age of household head. It shows that the intensive margin rather reduces when household head aged between 45-54 is unemployed. Unlike the extensive margin, however, the number and age of kids do not affect the use of intensive margin as an insurance.

Figure 1. Responses of Wife's Extensive and Intensive Margins of Labor Supply

(a) By Household Characteristics



(b) By Wealth Level



Note: Figure shows the response of extensive margin (left column) and intensive margin (right column) of wife's labor supply by household characteristics when the household head is unemployed. The dots indicate the point estimates, and bars indicate 90% confidence intervals. A mean regression is used for the response estimation of extensive margin, while a median regression is used for the response estimation of intensive margin.

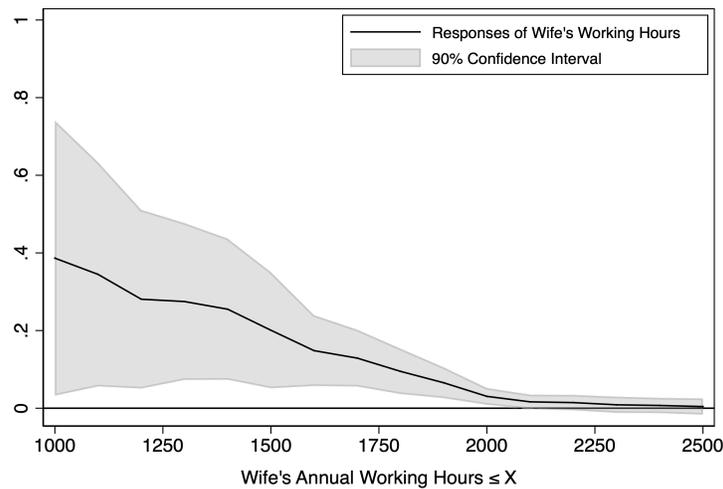
Further, we examine if household's wealth level plays a role as a constraint on wife's labor supply. Ortigueira and Siassi (2013) argue from their model that the wealth-poor households has a tendency to use wife's labor supply, while the wealth-rich tend to use saving. In Figure 1, we show the response of extensive and intensive margin from time t to $t+2$ conditional on the current household asset positions at time t measured by the amount of liquid assets, illiquid assets, debts, and credit constraints. The construction of asset variables are described in Section 2.2. The figure shows that the response of extensive and intensive margin for the households with the current amount of liquid assets below 25th, 50th, and 75th percentiles (red line with a triangle marker) all non-significant. In contrast, households with current liquid assets above each percentile also shows non-significant responses of extensive margin. Liquid assets above median show a marginally significant estimate, but not for the 75th percentile. The same non-significance is also shown for the intensive margin.

We also conduct the same analysis based on the household's current illiquid assets. Again, a statistical significance cannot be found. When we condition on the amount of debts, we find that household with no current debt tend to increase the extensive margin of wife's labor supply. Finally, we construct whether or not households are credit constrained by imputing the probability of being constrained from SCF data.¹³ When the probability of being credit constrained is low (below median), the wife's labor participation significantly rises. The results that the extensive margin of wife's labor supply hinges on the current debt holdings and credit constraint and that the intensive margin independently non-significant to household asset positions contradict to Ortigueira and Siassi (2013)'s model prediction.

We finally conjecture on the non-responsiveness of intensive margin that the response of wife's intensive margin is constrained by the wife's current working hours. When a wife is already working enough hours, the marginal cost of working additional hours must be high. To verify the validity of our conjecture, we estimate the response of intensive margin of wife's labor supply conditional on the wife's current hours of work in Figure 2. When the wife's

¹³Details about the imputation of credit constrained variable can be found in Appendix.

Figure 2. Responses of Wife’s Intensive Margin: Conditional on the Current Working Hours



Note: Figure shows the response of wife’s annual working hours to unemployment shocks conditional on the wife’s current working hours less than X hours. In other words, the value 2000 on the x-axis indicates the response of wives who are working less than 2000 hours a year. The gray shade indicates the 90% confidence interval.

current annual working hours is less than 1,000 hours, then the response of intensive margin is 40% higher for the wives with unemployed household head than those with employed household head. The response gradually declines as the current hours of work rises and is no longer significantly different from the response of wives with employed household head if wife’s current working hours exceed 2,000 hours a year. This figure clearly states that the wife’s intensive margin is strictly constrained by the current working hours. Another interpretation of this result is that the annual 2000 hours of working can be considered as a full-time worker. When a wife works full-time, there is little room to increase her working hours for an extra salary unless she finds another part-time job. Therefore, our result informs that the intensive margin of part-time wives can respond to unemployment shocks but not of full-time wives.

Let’s summarize the constraints on wife’s labor supply:

Constraint 1: *The number and age of kids in the household are important constraints for a mom's response of labor supply extensive margin.*

Constraint 2: *The current wife's working hours significantly affect the response of intensive margin. The wife's current working hours proxies wife's part-time/full-time working style.*

Fact 4: *Wife's labor supply is weakly responsive to an unemployment shock for the median households because of the two constraints stated above.*

2.3.3 Constraints on the Use of Saving/Borrowing

Another puzzling result that we obtained in Table 3 is the non-significant use of self-insurance by households whose head is unemployed. In this section, we disaggregate the sample by household characteristics to examine if the non-significance is driven by any particular subgroups of households.

First, we focus on saving response measured by the difference between current liquid assets and previous liquid assets. When saving is used to supplement an income loss due to unemployment, the response of saving needs to be negative. Saving in the household survey is very noisy that the estimates are sensitively affected by outliers. Therefore, we use a median regression to estimate the insurance effect. In panel (a) of Figure 3, we plot the estimates of saving responses by household characteristics, splitting the sample into married (left column) and non-married (right column). Both married and non-married households have no significant effect observed in these estimates.

Next, we examine the use of borrowing/debt to supplement income losses. Debts in our construction include credit debts, student debts, medical debts, legal debts and borrowing from relatives. In the household survey, more than 75% of households have zero debts. With

these observations, a median regression gives meaningless estimate and taking the difference has little information about households' response to shocks. Therefore, we subtract debt from saving to estimate the effect of net saving. If debts are effectively used to insure consumption when unemployed, either debt has to increase or net saving has to decrease. In panel (b) of Figure 3 demonstrates the response of net saving. Compared to the effect in saving, some households rather increase the net saving. For example, household head aged between 35-44 increases the net saving which attributes to the decrease in debt. Also, household head with an education more than 12 years increase the net saving. Note that these sub-groups of household rather reduce the amount of debt when unemployed. For non-married households, the effect on net saving is still non-significant.

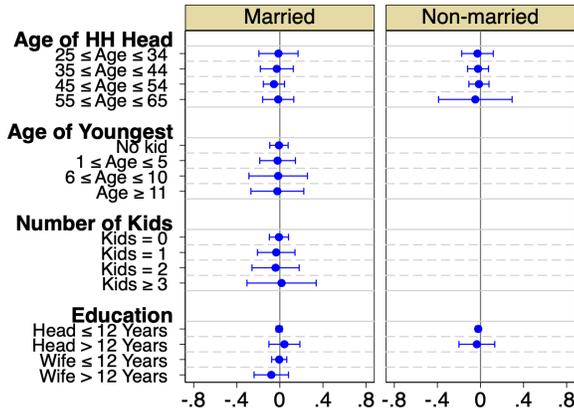
Finally, we explore the response of illiquid assets against unemployment shocks. More than 50% of households have zero illiquid assets in our household data. Therefore, we again add illiquid assets to our net saving and estimate the response of net saving plus illiquid assets to unemployment shocks in panel (c) of Figure 3. If households utilize illiquid assets to insure consumption, the change in the amount of illiquid assets must be negative, or the response of all assets must become negative. Compared to the responses of saving and net saving, all assets become non-significant for all the household subgroups except for the household head age between 35-44. The responses of all assets for non-married households are all non-significant. Disaggregating household by their characteristics do not provide an idea what makes the use of assets constrained. Some may argue that households liquidate some of their illiquid assets and use liquid assets as an insurance. Our result rules out this scenario because if the scenario is true, either liquid assets or liquid assets plus illiquid assets should decline by significant amount. We do not observe this response in our finding.

Finally, we check the response of saving/borrowing by the level of household wealth. In Figure 4 (a), we show the response of liquid assets by household's current amount of liquid assets, illiquid assets, debts, and credit constraints. Interestingly, married households with

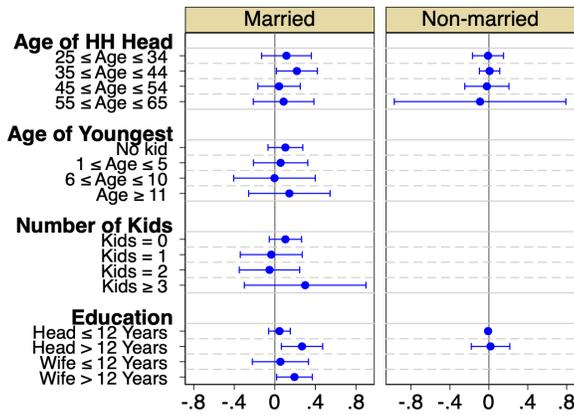
their liquid assets below 25th and 50th percentiles exhibit significant use of their saving. On the contrary, the level of illiquid assets, debts, and credit constraints do not show any significance. This effect still remains even for the net saving and net saving plus illiquid assets. This result seems a bit contradictory to our economic intuition since a standard consumption-saving model tells us that households with a large amount of saving should use the saving as an insurance. This effect is named “saving-constrained” households by Miranda-Pinto et al. (2020). They argue that low-income households are consuming at the consumption threshold under which they pay extra utility cost. Despite the small amount of saving, the households have to use their saving to maintain the threshold level of consumption. This result needs to be examined in more detail.

Figure 3. Response of Assets by Household Characteristics

(a) Saving



(b) Net Saving: Saving – Debts



(c) All Assets: Saving – Debts + Illiquid Assets

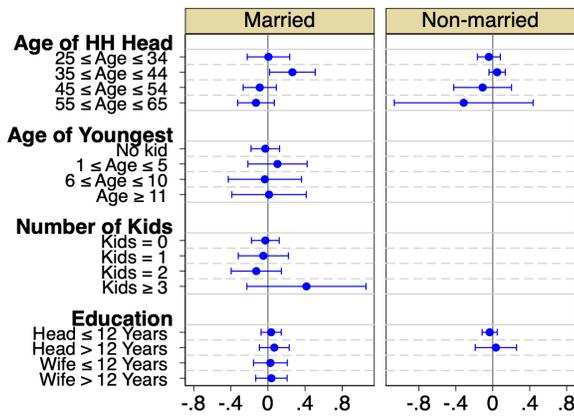
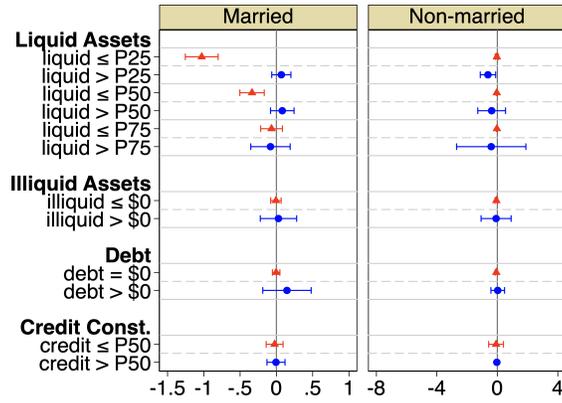
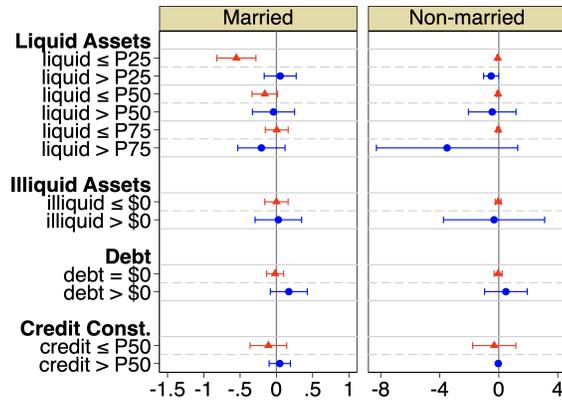


Figure 4. Response of Assets by the Level of Previous Assets

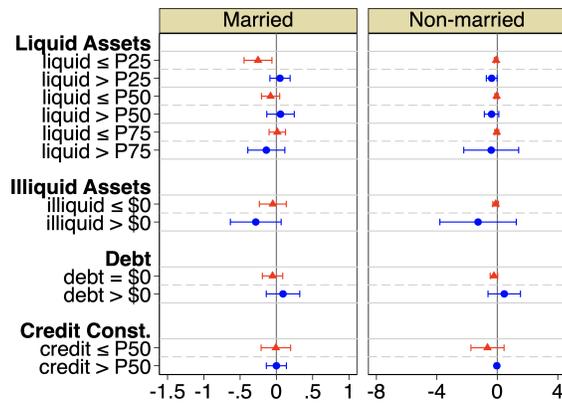
(a) Saving



(b) Net Saving: Saving – Debts



(c) All Assets: Saving – Debts + Illiquid Assets



3 Conclusion

In this study, we shed light on constraints that hinder households from using their potential consumption insurance channels, such as self-insurance, wife's labor supply, and public/private transfers. Understanding the source of constraints is important. In response to any particular income shocks households are unable to smooth out their consumption not because the economy is underdeveloped in providing a full set of insurance mechanisms. But instead it's because households are constrained in using their insurance. Therefore, a policy implication from our results is that the government must focus on removing or mitigating the constraints that households face in their everyday life by a policy.

References

- Altonji, B. J. G., Hayashi, F., and Kotlikoff, L. J. (1992). Is the Extended Family Altruistically Linked ? Direct Tests Using Micro Data. *The American Economic Review*, 82(5):1177–1198.
- Altonji, J. G. and Siow, A. (1987). Testing the Response of Consumption to Income Changes with (Noisy) Panel Data. *Quarterly Journal of Economics*, 102(2):293–328.
- Attanasio, O. and Davis, S. (1996). Relative Wage Movements and the Distribution of Consumption. *Journal of Political Economy*, 104(6):1227–1262.
- Auerbach, A. J. and Feenberg, D. (2000). The significance of federal taxes as automatic stabilizers. *Journal of Economic Perspectives*, 14(3):37–56.
- Blundell, R. and Pistaferri, L. (2003). Income volatility and household consumption: The impact of food assistance programs. *Journal of Human Resources*, 38(SUPPLEMENT):1032–1050.
- Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.
- Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). Consumption Inequality and Family Labor Supply. *American Economic Review*, 106(2):387–435.
- Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2018). Children, time allocation, and consumption insurance. *Journal of Political Economy*, 126:S73–S115.

- Browning, M. and Crossley, T. F. (2009). Shocks, Stocks, and Socks: Smoothing Consumption over a Temporary Income Loss. *Journal of the European Economic Association*, 7(6):1169–1192.
- Cochrane, J. H. (1991). A Simple Test of Consumption Insurance. *Journal of Political Economy*, 99(5):957.
- Deaton, A. (1992). *Understanding Consumption*.
- Dynarski, S. and Gruber, J. (1997). Can Families Smooth Variable Earnings? *Brookings Papers on Economic Activity*, 1997(1):229–303.
- Engen, E. M. and Gruber, J. (2001). Unemployment insurance and precautionary saving. *Journal of Monetary Economics*, 47(3):545–579.
- Hall, R. E. and Mishkin, F. S. (1982). The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. *Econometrica*, 50(2):461 – 81.
- Hayashi, F., Altonji, J., and Kotlikoff, L. (1996). Risk-Sharing between and within Families. *Econometrica*, 64(2):261–294.
- Huggett, M., Ventura, G., and Yaron, A. (2011). Sources of Lifetime Inequality. *American Economic Review*, 101(7):2923–2954.
- Kaplan, G. and Violante, G. (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics*, 2(October):53–87.
- Kimball, M. S. and Mankiw, N. G. (1989). Precautionary Saving and the Timing of Taxes. *Journal of Political Economy*, 97(4):863–879.
- Kotlikoff, L. J. and Spivak, A. (1981). The Family as an Incomplete Annuities Market. *Journal of Political Economy*, 89(2):372–391.
- Krueger, D. and Wu, C. (2021). Consumption Insurance Against Wage Risk: Family Labor Supply and Optimal Progressive Income Taxation. *American Economic Journal: Macroeconomics*, 13(1):79–113.
- Low, H., Meghir, C., and Pistaferri, L. (2010). Wage risk and employment risk over the life cycle. *American Economic Review*, 100(September):1432–1467.
- Mace, B. J. (1991). Full Insurance in the Presence of Aggregate Uncertainty. *Journal of Political Economy*, 99(5):928.
- Meghir, C. and Pistaferri, L. (2004). Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32.
- Miranda-Pinto, J., Murphy, D., Walsh, K. J., and Young, E. R. (2020). A Model of Expenditure Shocks.

- Ortigueira, S. and Siassi, N. (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics*, 60(6):650–666.
- Stephens, M. (2002). Worker Displacement and the Added Worker Effect. *Journal of Labor Economics*, 20(3):504–537.
- Townsend, R. M. (1994). Risk and Insurance in Village India. *Econometrica*, 62(3):539–591.

A Robustness Checks

A.1 Individual Characteristics and Fixed Effects

Table 4. Individual Characteristics and Fixed Effects by Marital Status

	Non-married				Married			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
Food Consumption								
Food at home	-0.000	-0.019	-0.128	-0.066	0.004	0.000	-0.071	-0.048
+ Food away from home	-0.134*	-0.145*	-0.121	-0.028	-0.041*	-0.043*	-0.111	-0.096
+ Food delivered	-0.134*	-0.152**	-0.124	-0.027	-0.020	-0.024	-0.107	-0.093
Non-food Consumption								
Utility	-0.045	-0.040	-0.331	-0.192	-0.009	-0.000	-0.046	-0.033
+ Transportation	-0.057	-0.072	-0.253	-0.197	-0.007	-0.005	-0.002	-0.008
+ Health	-0.036	-0.056	-0.163	-0.109	-0.003	0.007	-0.006	0.008
+ Clothing	-0.068	-0.082	-0.159	-0.157	-0.007	0.006	-0.009	0.013
+ Education	-0.097**	-0.106**	-0.143	-0.159	-0.000	-0.006	-0.015	0.002
+ Child care	-0.091*	-0.116**	-0.138	-0.156	-0.024	-0.026	-0.027	-0.011
+ Recreation	-0.101	-0.113*	-0.113	-0.128	-0.002	-0.013	-0.020	-0.004
Durable Consumption								
Housing	-0.066*	-0.059**	-0.678**	-0.812***	-0.020	-0.009	-0.031	-0.026
+ Vehicle	-0.100	-0.127	-0.420*	-0.444**	-0.023	-0.017	-0.041	-0.029
Asset Accumulation								
Liquid	-0.000	-0.033	-0.721**	-0.383	-0.013	-0.011	-0.205	-0.257
- Debt	-0.046	-0.028	-0.436	-0.836	0.000	0.013	-0.819**	-1.565***
+ Illiquid	-0.048	-0.073	-0.103	-0.054	0.000	0.005	-0.445	-1.120***
Spouse's Labor Supply								
Intensive Margin	-	-	-	-	0.000	0.003	0.000	-0.033
Mean/Median Regression	Med	Med	Mean	Mean	Med	Med	Mean	Mean
Year & Group FE	✓	✓	✓	✓	✓	✓	✓	✓
Individual Characteristics		✓	✓	✓		✓	✓	✓
Individual FE				✓				✓

Note: Table shows the robustness checks of our benchmark results. Column (1) runs a regression with year FE but without controlling for individual characteristics. Column (2) shows our benchmark result where we control for observable individual characteristics, e.g., age, age-squared, education, family size, number of kids, race, and residence. Column (3) shows estimation results with individual FE. For this regression, outliers of consumption growth are no longer an issue; therefore, we took mean regression with individual FE, year FE, and individual characteristics.

A.2 Robustness Checks on the Regressions of Binary Variables

Table 5. Mean Regression vs. Probit Regression

	Non-married			Married		
	Linear Prob.	Probit	XXX	Linear Prob.	Probit	XXX
Spouse's Labor Supply						
Extensive Margin	-	-	-	0.009	0.144	-
Public & Private Transfers						
Unemployment Benefit	0.332***	1.385***		0.426***	1.899***	
Transfers from Relatives	0.060	0.406***		0.027	0.364***	
Year & Group FE	✓	✓		✓	✓	
Individual Characteristics	✓	✓		✓	✓	
Individual FE	✓			✓		

Note: Table shows the robustness checks of our benchmark results. We compare the significance of our benchmark estimates with different estimation methods: median regression, mean regression, or probit regression.

A.3 Symmetric Responses to Unemployment Shocks

Table 6. Testing of Symmetric Responses to Unemployment Shocks

	EU Shock		EU & UE Shock	
	Non-married	Married	Non-married	Married
Food Consumption				
Food at home	-0.019	0.000	0.008	-0.000
+ Food away from home	-0.145*	-0.043*	-0.062	-0.059***
+ Food delivered	-0.152**	-0.024	-0.071	-0.045**
Non-food Consumption				
Utility	-0.040	-0.000	-0.041	-0.063**
+ Transportation	-0.072	-0.005	-0.090**	-0.031*
+ Health	-0.056	0.007	-0.062	-0.087***
+ Clothing	-0.082	0.006	-0.068	-0.084***
+ Education	-0.106**	-0.006	-0.074*	-0.074**
+ Child care	-0.116**	-0.026	-0.070	-0.090***
+ Recreation	-0.113*	-0.013	-0.058	-0.075**
Durable Consumption				
Housing	-0.059**	-0.009	-0.030	0.011
+ Vehicle	-0.127	-0.017	-0.040	0.015
Asset Accumulation				
Liquid	-0.031	-0.000	-0.007	-0.021
- Debt	-0.029	0.019	-0.013	-0.056
+ Illiquid	-0.066	0.002	-0.026	-0.022
Spouse's Labor Supply				
Intensive Margin	-	0.003	-	0.001
Extensive Margin	-	0.020	-	0.018*
Public & Private Transfers				
Unemployment Benefit	0.378***	0.469***	0.224***	0.260***
Transfers from Relatives	0.110***	0.061***	0.084***	0.042***

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Conclusion

This dissertation is comprised of three papers that explore the households' ability to smooth consumption in the face of adverse income shocks. It looks at various aspects of why and how a household makes its decisions. The first paper compares the relative ability of Individuals and Pairs to solve a finite, stochastic lifecycle problem in order to test the pervasive practice of using individuals as representative decision-makers in dynamic choice experiments. Pairs are significantly better aligned with the rational, representative-agent benchmark than Individuals; subjects forming a joint decision earn about 40% more, on average, than subjects making individual decisions. Chat data reveals that Pairs mostly discuss spending, rather than saving or borrowing, and adhere to simple consumption heuristics that are largely invariant to past errors.

The second paper shows that the first income quartile shows a 6.2 percentage point decrease in nondurable consumption growth when households are unemployed for several months, a 18 percentage point decrease in durable consumption growth, a roughly 8.3 percentage point decrease in food consumption growth, and approximately a 11.3 percentage point decrease in total consumption growth. This suggests that the lowest income quartile in America is unable to smooth their consumption when they receive a negative income shock and therefore are not operating under complete markets. In order to check if certain government benefits, food stamps and unemployment compensation, help bridge the gap for the lowest income quartile I run regressions including controls for the benefits. I show again that the lowest and second lowest income quartiles again both experience significant, negative consumption growth with a few exceptions. While unemployment benefits appear to help households make up some of the consumption they lose, food stamps do not play the same role. Indeed, it appears that food stamps instead serve as a signal that households already lack full consumption insurance.

In the third paper, we empirically demonstrate that households are constrained in us-

ing two main consumption insurance channels—wife’s labor supply and self-insurance—in response to household head’s unemployment. We document three empirical facts about the household responses to household heads’ unemployment experiences: (1) Households can completely smooth “food-at-home” expenditure; (2) public and private transfers are effectively used to insure against household head’s unemployment shocks; and (3) non-married households cut down consumption items with high income elasticity to insure food consumption. Additionally, we document three facts about constraints on the consumption insurance: (1) the number and age of kids are constraints for mom’s extensive margin of labor supply; (2) current wife’s working hours are a constraint for the marginal increase of wife’s working hours; and (3) self-insurance is effectively used only for low-saving households. Given these empirical evidence, we quantitatively analyze the welfare loss due to the constraints.