Methods of Consumption Smoothing: Coping with Pension Arrears in Post-Soviet Russia

Clair Null*

Abstract

The fact that households smooth their consumption has been well-documented in the empirical literature. However, relatively little evidence exists to explain how such smoothing is attained. In this paper we employ a rich dataset of financial transactions to investigate how households respond to exogenous income shocks.

Specifically, we consider a range of income sources and expenditures recorded by a panel of Russian households with pensioners between 1995 and 2000. Using pension arrears as a source of exogenous variation in income, we are able to observe participation in informal risk-sharing networks as well as formal insurance schemes, borrowing and lending, saving and asset depletion, and changes in income from labor, farming, and government transfers. We exploit a simple accounting identity that allows us to account for the endogenous nature of the complete set of financial transactions employed by a household. We begin by estimating a household-level fixed effect model and investigate whether transaction use differs across several household demographics. We then relax the assumption that all households follow the same decision-making process and consider a random coefficients estimator. We provide some evidence that our results are robust to attrition bias.

To the extent that households only allocate roughly $\frac{1}{6}$ to $\frac{2}{5}$ of pension income to consumption expenditures, they must be putting the remaining share into another type of transaction that helps to smooth consumption. However, our results do not conclusively identify transaction use other than to show that labor income is a substitute for pension income. A random coefficients analysis indicates that choice of smoothing mechanisms may to be more idiosyncratic than a fixed effects estimating framework is equipped to deal with.

*Department of Agricultural and Resource Economics, UC Berkeley. I would like to thank my advisor, Ethan Ligon, and my office-mate, Shanthi Nataraj, for their invaluable advice on this project. I am also grateful to participants of the ARE Development Workshop for helpful comments. Please do not cite or distribute.

1 Introduction

During the 1990's the countries of the Former Soviet Union underwent painful transitions from command economies to market economies. As GDP contracted and inflation soared, the region suffered what has been called "one of the most acute poverty and welfare reversals in the world" [2]. In this paper, we focus on Russia, where an ambitious data collection project carefully tracked household and individual responses to the unfolding crisis. As both the incidence and severity of poverty increased [5], the Russian Longitudinal Monitoring Survey (RLMS) documented household strategies of dealing with reductions in income, increases in prices, and widening variation in both. While the Soviet economy had left relatively few decisions to individuals (restictions on mobility prevented migration, work requirements regulated employment, and centralized distribution systems standardized prices and availability of goods), in the post-Soviet economy households faced myriad choices that involved painful trade-offs as they allocated increasingly scarce resources among expanding expenses on goods and services formerly provided by the government (e.g. natural gas and health care).

In this paper, we investigate a variety of coping methods adopted by households in response to a specific exogenous income shock: pension arrears. Ultimately, we hope to understand how household characteristics determine which methods a household employs and how heavily the household relies on that method. When a government is unable to fully insure its citizens against idiosyncratic income shocks, or is even the source of those income shocks as in the case of pension arrears, it may still be possible to facilitate effective consumption smoothing at the household level by ensuring that households have access to the sorts of mechanisms that can mitigate the painful consequences of an income shock.

The rest of this paper proceeds as follows. We begin with a brief survey of the literature on consumption smoothing in Section 2, followed by a description of models that underlie the research question in Section 3. Section 4 describes the estimation strategy, the data are outlined in Section 5, Section 6 presents results, and Section 7 concludes.

2 Literature Review

The consumption smoothing literature can be broadly divided into intertemporal and interpersonal models. The former, known as the permanent income hypothesis as described by Friedman ([3]), predicts that individuals choose consumption levels based on their expectations of permanent income rather than on realizations of current income. The latter, known as the full-insurance model, predicts that individuals with risk-averse preferences will maximize their welfare by eliminating idiosyncratic shocks to income by means of interpersonal transfers, assuming they have access to complete markets for contingent claims as characterized by Arrow & Debreu ([1]). A complete survey of applications of these models is beyond the scope of this paper; here we focus on several contributions to the literature which have particular relevance to the study at hand.

Paxson [9] was one of the first authors to present credible evidence that income shocks are treated differently than permanent income. Using weather as an instrument for income shocks faced by farmers in Thailand, she showed that these households saved a higher fraction of transitory income than permanent income, indicating that savings were being used as a means of smoothing consumption.

Townsend [12] offered the first empirical test of the full-insurance model and applied his method to ICRISAT data. Though the null hypothesis of full insurance is rejected by his the regressions, he presents a barrage of statistical evidence that household consumptions respond to village average consumption but are not as linked to idiosyncratic shocks, suggesting that some sort of insurance scheme is in operation at the village level.

In a theoretical paper that expands upon the test of full-insurance proposed by Townsend, Ligon [6] solves a permanent income model as well as a private information model. In order to test which model best fits the data, he shows that the three models can be nested based upon the conditions that derive from their corresponding maximization problems. Again using ICRISAT data, he concludes that the private information model is a better description of two of the three villages he considers. In the third, there is evidence that different households are best described by different models.

Despite the fact that a number of authors have followed Townsend's method and tested the full-insurance model in a variety of developing country settings, and still more have applied the permanent income model, very little comprehensive evidence exists on how consumption-smoothing is actually implemented. Many studies consider one particular method by which risk may be spread over people and / or time independently, without being able to relate this method to other means by which households might spread risk. In one of the few papers that incorporates and compares multiple methods of consumption-smoothing, Lim & Townsend [7] document interesting patterns among households in ICRISAT villages. They find that though none of the existing models is completely representative of the villages they study, to the extent that risk is shared among households and over time, buffer stocks of currency and crop inventory are the key mechanisms. Institutions that provide credit and insurance also contribute to the process of smoothing consumption, but the purchase and sale of real capital assets and livestock are not means by which households spread risk across individuals and / or time. Moreover, they find that the relatively rich use crop inventory more intensively than do their poorer counterparts, who rely more on currency inventories. Finally, users of crop inventories are more likely to pass tests of the full-insurance, permanent income, and private information models, whereas users of currency inventories are more likely to fail tests of the private information model.

Relatively few papers have made use of the rich data on financial transactions collected by the RLMS. Like this paper, Skoufias [10] also used panel data from the RLMS to investigate patterns of consumption smoothing. Following Townsend's method, he rejects the null hypothesis of full insurance for both food and non-food consumption using primary sampling units as potential insurance groups[1]. Categorizing households by various demographic characteristics, he finds that poor households are less insured relative to the average non-poor household in the sample, and that the degree of covariation between household income and consumption differs across regions and is higher for female-headed households but lower for households with younger children (perhaps as a consequence of government benefits for children). Finally, Skoufias estimates probit models that predict use of other consumption-smoothing mechanisms in response to various types of income shocks[2]. In general, he finds that there is no single method used most frequently, but rather, households employ a variety of techniques in combination to deal with these shocks to income.

Jensen & Richter [4] also focus on the Russian pension crisis of 1996 using data from the RLMS, but are primarily concerned with the health consequences of pension arrears. Using a differences-in-differences approach to identify the effects of unpaid pensions in 1996, the authors find that households whose pensions were in arrears significantly reduced their consumption of calories, protein, medications, and health services. However, the authors also find evidence that households were able to replace some of the lost income by increasing labor supply and selling assets.

This paper contributes a more sophisticated and detailed study of the coping methods used by Russian households with pensioners in the face of an exogenous income shock. We use a somewhat longer panel that allows for us to control for time effects and rather than simply estimating the effect of an arrears "treatment," we estimate the response of a variety of financial transactions to the amount of pension arrears,

[1] There are 38 primary sampling units represented in the RLMS data.

[2] These mechanisms used as dependent variables include: having a household member with a second paying job, having a member involved in informal economic activities, having positive net transfers received, having positive net debt, selling assets in the previous 3 months, selling poultry in the previous 30 days, cultivating land in the previous 12 months, and cultivating land at the time of interview. The three income shocks considered are being owed salary at the primary job, being on forced leave from primary job, and being unemployed.

acknowledging that such transactions are endogenously determined.

3 Models of Consumption Smoothing

3.1 Individual Utility Maximization

We consider a representative agent who is concerned with maximizing her utility in the presence of uncertainty. The rest of her life spans some definite interval of time divided into T periods, denoted by $t = 0, ..., T$. In any period t, some state of the world, $s_t \in \Omega_t$, is realized. For simplicity, let the set of possible states be finite with $\Omega_t = \{1, ..., \mathcal{S}_t\}$. In a given period t, the agent expects each state to occur with subjective probability $\pi(s_t) \in [0, 1] \forall\, s_t \in \Omega_t$ s.t. $\sum_{s_t=1}^{\mathcal{S}_t} \pi(s_t) = 1 \forall\, t$. We note that the agent is unable to influence which state of the world is realized, ruling out moral hazard type problems. Future utility is discounted at rate β. In each period her utility is a strictly concave function of her consumption in that period, which we write as $u(C(s_t))$, assuming that her utility function does not vary over time or over states of the world but that her consumption allocations may. In general, all random variables and allocations in this theoretical discussion are time- and state-dependent, which we express by writing $x(s_t)$ where s is specific to t, as denoted by the subscript, and x is a function of the particular s realized in period t. Then in a given period, the agent's problem is to maximize her discounted expected utility over all future periods

$$\max_{\{C(s_t)\}} \sum_{t=0}^{T} \beta^t \sum_{s_t \in \Omega_t} \pi(s_t) \cdot u(C(s_t))$$

The agent has access to credit markets so she faces an intertemporal budget constraint. Standardizing the price of the consumption good to 1, the constraint can be expressed as $\sum_{t=0}^{T} C(s_t) \leq \sum_{t=0}^{T} [W(s_t) - S(s_t)]$ where $W(\cdot)$ is her wealth and $S(\cdot)$ is the amount of that wealth that she sets aside as savings for the future. The temporal evolution of wealth can be described as $W(s_t) = (1 + r(s_t)) \cdot S(s_{t-1}) + Q(s_t)$, where $Q(\cdot)$ is total non-interest income and $r(\cdot)$ is the rate of interest on a risk-free asset, assuming the agent has access to functioning credit markets.

3.2 Credit Markets: Consumption-Smoothing Across Time

Since $u(\cdot)$ exhibits decreasing marginal utility, the agent will smooth her consumption across time[3]. The first order conditions of the agent's problem with respect to consumption allocations yield the familiar Euler equation which states that the ratio of expected marginal utilities in any two periods is constant. Formally,

$$\mathbb{E}\left[\frac{u'(C(s_{t+1}))}{u'(C(s_t))}|s_t\right] = \frac{1}{\beta(1 + r(s_t))} \tag{1}$$

The expectation operator in this equation comes from the fact that the agent does not know which states will be realized in future periods. Equation (1) governs consumption decisions in the two periods but does not mention income in those two periods. Thus, consumption is smoothed over time, rather than being strictly determined by income in any given period. Agents make use of credit markets, either by saving and dis-saving or by borrowing and lending. Information asymmetries between individuals may influence an agent's choice of whether to save / dis-save or borrow from / lend to other agents. At this point we abstract away from such frictions in the market, so for the purposes of the model the two approaches to smoothing consumption across time are equivalent.

3.3 Society: Risk-Sharing Across Individuals

By Jensen's inequality, $\mathbb{E}[u(C(s_t))] \leq u(\mathbb{E}[C(s_t)])$ since $u(\cdot)$ is concave. Thus, the agent is risk-averse and will be willing to pay to avoid uncertainty. We now place the agent

[3]This will be true even in the absence of uncertainty. To show this, suppose $\Omega_t = \{s_t\} \forall t$ so that there is no uncertainty, i.e. $\pi(s_t) = 1 \forall t$. This is not to say that the same state occurs in every period t, i.e. $s_t \neq s_j$ for two periods t and j, but rather that in each period only one state will occur. Then the first order conditions reduce to $\frac{u'(C_{t+1})}{u'(C_t)} = \frac{1}{\beta(1+r_t)}$.

in the context of her society, and consider the Pareto optimal solution to society's problem as a means of investigating ways that members of society might be able to act collaboratively to achieve higher utilities than they are able to individually. The end result is intuitive; by pooling idiosyncratic risk, so that only aggregate risk remains, agents can reduce the variance of their consumption allocations.

We can think of a social planner[4] who solves the following problem:

$$\max_{\{C_i(s_t)\}} \sum_{i=1}^{N} \lambda_i \sum_{t=0}^{T} \beta_i^t \sum_{s_t \in \Omega_t} \pi_i(s_t) \cdot u(C_i(s_t))$$

$$\text{s.t.} \sum_{i=1}^{N} C_i(s_t) \leq \sum_{i=1}^{N} [W_i(s_t) - S_i(s_t)] \, \forall \, t$$

where λ_i, $i = 1, ..., N$ are weights that reflect the N agents' reservation utilities. These weights only have meaning relative to one another, so we can standardize them such that $\sum_{i=1}^{N} \lambda_i = 1$ for simplicity. We let β, $\pi(\cdot)$, and $u(\cdot)$ vary by person. The first order conditions with respect to consumption for this new problem are

$$u_i'(C_i(s_t)) = \frac{\phi \mu(s_t)}{\lambda_i \beta_i^t \pi_i(s_t)}$$

where ϕ is the Lagrange multiplier on the resource constraint and we have used $\mu(\cdot)$ as short-hand for aggregate resources in society. This equation makes clear that each agent's consumption allocation is determined by aggregate resources, her reservation utility, her rate of time preference, and her subjective probability that state s_t will occur. Importantly, we see that consumption varies with income only via aggregate resources; individual realizations of W are irrelevant aside from their contributions to μ.

Rearranging these conditions for two arbitrary agents i and j, we see that marginal

[4] The social planner is an easy way of thinking about the problem, though by definition, the Pareto optimal solutions could be obtained by the agents acting directly themselves and employing the appropriate transfers.

utilities for any two agents are perfectly correlated since

$$\frac{u_i'(C_i(s_t))}{u_j'(C_j(s_t))} = \frac{\lambda_j \pi_j(s_t)\beta_j^t}{\lambda_i \pi_i(s_i)\beta_i^t} \tag{2}$$

This sort of risk-sharing across individuals can be achieved via private or government transfers, or by purchasing formal insurance.

Since agents in this framework also have access to credit markets, as in the previous section, equation (1) also holds. Note that in this discussion we have assumed that agents have full information about one another, so there are no impediments to the risk-sharing network being functional. In the case of private information, e.g. when an agent is unable to confirm another agent's realization of W, risk-sharing arrangements must be incentive-compatible in the sense that they must induce agents to disclose the truth about themselves. Ligon [6] has shown that with private information, (2) holds only in expectation whereas (1) does not hold since agents are unwilling to make loans to one another without full information.

3.4 Modeling Transaction Use

The models outlined thus far in this section have described why households engage in consumption smoothing and two broad classes of mechanisms by which they might do so: smoothing over time and smoothing over individuals. However, these models have made no predictions about the specific sorts of transactions through which these mechanisms are implemented. We now present a rudimentary framework for the decision-making process that determines which transactions a household employs and the relative degree to which a household relies on one transaction versus another.

Our data are described in detail in Section 5, but we preview that discussion by categorizing the particular transactions considered in this paper as intertemporal smoothing (saving/dissaving, collecting/paying interest on assests, and borrowing/lending)[5],

[5]We note that among the intertemporal smoothing transactions, differenting between sav-

interpersonal smoothing (interpersonal transfers, government transfers, or formal sector insurance)[6], or a third class of transactions that involve substitution between leisure and consumption (labor and farm income).

Facing a set of available transactions in the wake of an income shock[7], a household weighs the costs ($\mathcal{C}_x$) and benefits ($\mathcal{B}_x$) associated with each transaction, where subscript x denotes the transaction in question. We allow both the costs and benefits of a transaction to vary with household characteristics and the "intensity" of use such that $\mathcal{C}_x = \mathcal{C}_x(\Upsilon, Q_x)$ and $\mathcal{B}_x = \mathcal{B}_x(\Upsilon, Q_x)$ where Υ is the set of household characteristics and Q_x is the "quantity" of the transaction. Allowing the marginal costs and benefits of a transaction to vary depending on how intensity of use addresses the possibility that a household might choose to use multiple transactions. For example, a rural household might expand its garden to produce vegetables to sell in response to pension arrears, but since there is a limit to how much more intensively the household could farm, other transactions may be used in addition. For simplicity, let $\mathcal{C}_x = \mathcal{B}_x = 0$ for $Q_x = 0$. This is a full information assumption that rules out cases in which a household spends resources investigating a potential transaction but ultimately decides not to engage in that transaction.

A maximizing household must set Q_x such that $\frac{\partial \mathcal{C}_x}{\partial Q_x} = \frac{\partial \mathcal{B}_x}{\partial Q_x}$ if $Q_x > 0$. In the case of a positive income shock, the income above what is needed in order to achieve

ing/dissaving and borrowing/lending may be artificial (e.g. borrowing can be considered simply as dissaving from an initial balance of zero). However, it is also possible that the costs and or benefits of saving/dissaving and borrowing/lending differ when the two activities are not carried out at the same institution, as would be the case if savings institutions did not have lending facilities and moneylenders were the only sources of loans. Assuming that saving/dissaving involves only one's own personal assets and borrowing/lending involves the temporary transfer of assets between two parties, we allow for the possibility that households might simultaneously save/dissave and borrow/lend.

[6]While borrowing/lending could also be thought of as a type of interpersonal smoothing, we restrict this class to transactions that occur independently in each time period, thus ruling out borrowing/lending as an interpersonal smoothing mechanism.

[7]We abstract away from time complications since formal sector insurance and farm income are the only transactions we discuss which require that a household choose the transaction prior to the occurance of a shock. All other transactions can be implemented ex post as a means of coping with a shock, though borrowing/lending and interpersonal transfers may have implications for future time periods (in the first case this would be in the form of paying interest on assets).

the planned level of consumption is allocated over some subset of all transactions. In the case of a negative income shock, the shortfall in income below what is needed to achieve the planned level of consumption is drawn from some subset of all transactions[8]. Thus, if Q_x is measured as the monetary value of transaction x, with flows out of a transaction taking positive values and flows into a transaction taking negative values, the income shock must be equal to $-\sum_{x \in X^*} Q_x$, where X^* is the set of transactions the household has decided to use (i.e. assign a non-zero value). As a concrete example, suppose a household's pension of 700 rubles is not paid and it makes up for this shortfall by borrowing 300 rubles, accepting a gift of 200 rubles from friends, and working odd jobs to earn 200 rubles. Then we have $-700 = -\sum_{x \in B,\tau,L} Q_x$ where $Q_B = 300$ and $Q_\tau = Q_L = 200$. Moreover, we know that $\frac{\partial(\mathcal{C}_B - \mathcal{B}_B)}{\partial Q_B} = \frac{\partial(\mathcal{C}_\tau - \mathcal{B}_\tau)}{\partial Q_\tau} = \frac{\partial(\mathcal{C}_L - \mathcal{B}_L)}{\partial Q_L} = 0$ and $\frac{\partial(\mathcal{C}_x - \mathcal{B}_x)}{\partial Q_x} > 0$ for all other transactions x.

While we do not undertake a complete characterization of the cost and benefit functions, we do propose several testable assumptions on the composite function $\frac{\partial(\mathcal{C}_x - \mathcal{B}_x)}{\partial Q_x}$.

1. The costs of using farm income as a consumption-smoothing mechanism are higher for urban households than for rural households since urban households must commute to agricultural plots outside the city repeatedly to care for the crops, whereas rural households only need to make one trip into the city to sell their produce. The benefits of farm income, however, are the same for the two types of households. Thus,
$$\frac{\partial(\mathcal{C}_F - \mathcal{B}_F)}{\partial Q_F}(\text{rural}, Q_F) \leq \frac{\partial(\mathcal{C}_F - \mathcal{B}_F)}{\partial Q_F}(\text{urban}, Q_F) \quad \forall\, Q_F$$

2. Households comprising only elderly members face higher costs (i.e. have a more difficult time finding supplemental labor income) and lower benefits (i.e. are paid less for temporary work) when using labor income as a consumption-smoothing

[8]In either case, the subset may be the full set of transactions.

transaction than do households that also include working-age members. Thus,

$$\frac{\partial(\mathcal{C}_L - \mathcal{B}_L)}{\partial Q_L}(\text{young}, Q_L) \leq \frac{\partial(\mathcal{C}_L - \mathcal{B}_L)}{\partial Q_L}(\text{elderly}, Q_L) \quad \forall\, Q_L$$

3. Households comprising only elderly members have smaller informal networks to whom they can turn for assistance smoothing consumption via interpersonal transfers, and thus the benefits of such networks might be less if the members' incomes are more highly correlated than those of the networks to which younger households belong. The costs of participating in such networks, however, would not differ between elderly-only and younger households. Thus,

$$\frac{\partial(\mathcal{C}_\tau - \mathcal{B}_\tau)}{\partial Q_\tau}(\text{young}, Q_\tau) \leq \frac{\partial(\mathcal{C}_\tau - \mathcal{B}_\tau)}{\partial Q_\tau}(\text{elderly}, Q_\tau) \quad \forall\, Q_\tau$$

4. Similarly, rural households also likely belong to smaller informal networks with more highly correlated incomes than do urban households, though the costs of participation in such networks would not be expected to differ for the two types of households. Thus,

$$\frac{\partial(\mathcal{C}_\tau - \mathcal{B}_\tau)}{\partial Q_\tau}(\text{urban}, Q_\tau) \leq \frac{\partial(\mathcal{C}_\tau - \mathcal{B}_\tau)}{\partial Q_\tau}(\text{rural}, Q_\tau) \quad \forall\, Q_\tau$$

5. Rural households have less access to formal insurance than do urban households, making the costs of using formal insurance as a mechanism for smoothing consumption higher for rural than urban households. Assuming that insurance is actuarily fair, the benefits of formal insurance would be the same for the two types of households. Thus,

$$\frac{\partial(\mathcal{C}_I - \mathcal{B}_I)}{\partial Q_I}(\text{urban}, Q_I) \leq \frac{\partial(\mathcal{C}_I - \mathcal{B}_I)}{\partial Q_I}(\text{rural}, Q_I) \quad \forall\, Q_I$$

6. Similarly, there may be a steeper learning curve regarding formal insurance for elderly individuals who are more accustomed to social insurance being provided by the government compared to young adults who are more familiar with the market economy. Again, assuming insurance is actuarily fair,

$$\frac{\partial(\mathcal{C}_I - \mathcal{B}_I)}{\partial Q_I}(\text{young}, Q_I) \leq \frac{\partial(\mathcal{C}_I - \mathcal{B}_I)}{\partial Q_I}(\text{elderly}, Q_I) \quad \forall\, Q_I$$

In all of these cases, the implication is that a transaction will be used more by households for whom marginal costs minus benefits are less than for their counterparts.

Finally, we note that our estimation strategy, as described in the next section, does not require any specific functional form for utility. Rather, it takes as given that the relevant maximization problems have already been solved, and the appropriate financial transactions implemented, in order to achieve the optimal levels of consumption.

4 Estimation Strategy

4.1 Exogeneity of Pension Income

Pension arrears were widespread in Russia during the period under consideration. As documented in Figure 2, there was a drastic decline in pension payments in 1996 and 1998 relative to 1995 and 2000. In 1996 only 69% of households entitled to an old-age pension received their payments and in 1998 only 86% did, compared to 92% in 1995 and 98% in 2000. The assignment of pension arrears was based on permanent characteristics, namely geographic location[9] and pre-1996 pension level. At the time, Russia's pension system was funded on a pay-as-you-go basis by region; surplus funds from payroll taxes above what were needed to cover a region's pension obligations were sent to the federal

[9] Even the 356 households that move out of their original dwellings do not move out of their original survey area. It is possible, of course, that the households that were not observed for all 4 rounds of the survey disappeared precisely because they moved out of their dwelling and also out of their original survey area, but we have no way of knowing.

government to be redistributed to regions that were unable to fund themselves [4]. As the pension crisis progressed, this meant that arrears were more common in debtor regions that relied upon transfers from the federal government. In our sample, this is evidenced by the significant difference between the proportions of rural households between the non-arrears and arrears groups, as shown in Table 1. The federal pension administrative body instructed that priority should be given to paying those whose pensions were less than the minimum subsistence level, followed by non-working and single pensioners, on the premise that these groups were most reliant on their pension income [4]. However, the data indicate that these policies were not enforced. As shown in Table 1, pre-1996 pension levels were actually significantly smaller for households that eventually went into arrears, even conditional on location, and neither household size nor labor-force participation mattered significantly in determining whose pensions were paid.

Since pension receipts do not seem to be completely randomly assigned, we may be concerned that they are correlated with unobserved household characteristics that we can not include in the regressions. In order to allow for this possibility, in Section 6 we estimate a model with household level fixed effects (FE). As long as such characteristics are constant over time, the fixed-effects estimator is unbiased. Though such an assumption is fundamentally untestable, it seems more plausible than assuming there are no such correlations. With household fixed effects, identification of our parameters of interest comes from variation within each household over time, rather than across households.

4.2 Determination of Transactions

We exploit a simple accounting framework in order to identify the responses of particular types of financial transactions to changes in an exogenous source of income, specifically, pension receipts. We begin with the intertemporal identity

Table 1: Determinants of Arrears Status

	Non-Arrears Households		Arrears Households		
	Mean	**Std.Dev.**	**Mean**	**Std.Dev.**	**P-value**[1]
1996					
Rural	0.24	0.43	0.49	0.50	0.00
Pension Size[2]	931.08	615.58	746.79	585.17	0.00
| Urban	962.40	634.23	860.50	618.40	0.03
| Rural	831.35	541.16	625.16	521.89	0.00
Single	0.29	0.45	0.29	0.45	0.87
Non-working	0.52	0.50	0.56	0.50	0.21
Single and Non-working	0.26	0.44	0.28	0.45	0.30
N	1094		481		
1998					
Rural	0.30	0.46	0.43	0.50	0.00
Pension Size[2]	854.89	618.72	652.23	604.85	0.00
| Urban	894.80	645.01	768.63	640.34	0.05
| Rural	762.27	542.53	493.75	515.61	0.00
Single	0.29	0.45	0.26	0.44	0.43
Non-working	0.57	0.49	0.55	0.50	0.48
Single and Non-working	0.27	0.45	0.25	0.44	0.54
N	1211		198		

Notes:
1. For a t-test of the equality of means between households that did not experience arrears and those that did.
2. Pension levels in 1995.

$$K_{i,t+1} = (1 + r_t) \cdot K_{i,t} + L_{i,t} + \tau_{i,t} + B_{i,t} + G_{i,t} + F_{i,t} + I_{i,t} - C_{i,t} + P_{i,t} + Other_{i,t}$$

where $K_{i,t}$ are household i's period t assets, r_t is the interest rate, $L_{i,t}$ is labor income, $\tau_{i,t}$ are net transfers received, $B_{i,t}$ is net borrowing, $G_{i,t}$ are subsidies and benefits received from the government, $F_{i,t}$ is farm income, $I_{i,t}$ is insurance pay-outs net premiums, $C_{i,t}$ is consumption, $P_{i,t}$ is pension income, and $Other_{i,t}$ are other sources of net income. Define savings as

$$S_{i,t} \equiv K_{i,t+1} - K_{i,t} = r_t \cdot K_{i,t} + L_{i,t} + \tau_{i,t} + B_{i,t} + G_{i,t} + F_{i,t} + I_{i,t} - C_{i,t} + P_{i,t} + Other_{i,t} \quad (3)$$

We now distinguish between $P_{i,t}$, which is exogenous, and the set of k endogenous transactions denoted by the vector $y_{i,t}$, i.e.

$$y_{i,t} = (S_{i,t} \quad r_t \cdot K_{i,t} \quad L_{i,t} \quad \tau_{i,t} \quad B_{i,t} \quad G_{i,t} \quad F_{i,t} \quad I_{i,t} \quad C_{i,t} \quad Other_{i,t})'$$

In the context of the theoretical model presented in Section 3, the elements of $y_{i,t}$ are the solutions to the utility maximization problem according to each household's preferences and the markets available to it.

Acknowledging that shocks may affect various demographic groups differently, we can construct a vector of interactions between exogenous income, $P_{i,t}$, and a vector of l mutually-exclusive demographic dummy variables, $z_{i,t}$, such that $x_{i,t} = z_{i,t} \otimes P_{i,t}$. Using this notation, we can restrict $z_{i,t}$ to be the scalar 1, in which case all households are assumed to respond to pension income in the same way.

We can now write out a system of simultaneous equations that describe the allocation of pension income over the set of financial transactions. In matrix algebra

$$y'_{i,t} = y'_{i,t}\Gamma + x'_{i,t}\Pi$$

where Γ is a $k \times k$ matrix with zeros on the diagonal such that each transaction depends on all the other transactions and Π is an $l \times k$ matrix of coefficients that describes how each demographic group responds to pension income. As evident from this specification of the system, we are assuming that all household types follow the same decision-making process (since Γ does not vary by household characteristics), but are allowing for pension income to have different effects on the optimal level of transactions for different demographic groups. Rearranging terms[10] we have

$$y'_{i,t} = x'_{i,t}\Theta \tag{4}$$

where $\Theta = \Pi(\mathrm{I}_k - \Gamma)^{-1}$. In this paper we focus on the reduced form, since we can not identify Π independently without additional restrictions. We argue that the reduced form parameters Θ are of interest; though we are not able to identify the effect of pension income on each of the transactions independently of the others, the aggregate effect of pension income on each type of transaction is observable. Thus, while our results do not allow policy recommendations of the sort that would hold all other transactions constant, we can discuss the entire system of transactions and the effect of pension income on equilibrium outcomes for the various demographic groups.

In this case, since the set of control variables is the same for all k equations, system estimation reduces to estimating each equation independently.

[10]Note that $\mathrm{I}_k - \Gamma$ is invertible as long as no two transactions move one-for-one with each other and are identical in their responses to the other transactions. This is intuitive, since if that were the case, the two transactions really should not be treated as separate, but rather grouped together as one type of transaction.

4.2.1 Restrictions on Parameters

The accounting identity imposes an adding-up constraint on the parameter estimates. Writing the equation for each transaction separately, we have

$$y_{i,t} = P_{i,t}\beta_y$$

where $y \in \{S,\ rK,\ L,\ \tau,\ B,\ G,\ F,\ I,\ C,\ Other\}$. From Equation 3

$$S_{i,t} - r_t \cdot K_{i,t} - L_{i,t} - \tau_{i,t} - B_{i,t} - G_{i,t} - F_{i,t} - I_{i,t} + C_{i,t} - Other_{i,t} = P_{i,t}$$

and thus by substituting the estimating equations into the accounting identity, we find that

$$\sum_{y \in \{S,C\}} \beta_y - \sum_{y \in \{rK,L,\tau,B,G,F,I,Other\}} \beta_y = 1 \tag{5}$$

We can interpret the β parameters as the rates at which pension is converted into the other transactions, so intuitively, this restriction is simply saying that all pension income must be allocated across the set of transactions. We will provide evidence that this restriction is observed when we estimate the relationships between pension income and the other transactions in Section 6.

5 Data

Data for this paper is taken from the RLMS, a nationally-representative survey of the Russian population that has been conducted 13 times since 1992. The RLMS is funded by USAID and the National Institutes of Heath and is administered by the University of North Carolina's Population Center in collaboration with the Russian Institute of Nutrition and the Institute of Sociology at the Russian Academy of Sciences. The goal of the project was to measure the effects of reforms on the economic well-being of Russian households and individuals during the transition period. To that end, data include detailed information on household composition, living conditions, incomes, expenditures, and agricultural production as well as individual-level migration, labor supply, and personal opinions.

5.1 Survey Design

The RLMS is a panel of dwellings randomly selected using a multi-stage probability sample, with selection at each stage according to "probability proportional to size"[11]. The RLMS is technically a repeated cross-section of households; i.e. enumerators returned to the same address for each round of the survey, regardless of whether or not the occupants were the same. However, a panel of households can be constructed since the identity of the household within each dwelling is known. Moreover, beginning in round 7 (1996), data are included for households that had moved, provided they could be found at their new address, though they were given zero weight in the sample. For this reason, we restrict the analysis to data from survey rounds 6 and later, using the panel of households of which the round 6 sample was composed[12]. To alleviate bias due to attrition from the sample, we use only 4 rounds of data, covering the period 1995-2000, excluding 1997 and 1999, when no surveys were conducted.

5.2 Variable Definitions and Summary Statistics

Summary statistics for demographic variables are shown in the top panel of Table 2. Financial variables are summarized in Figure 1. The left panel shows unconditional distributions which depict how important the various transactions are on average in the sample, while the right panel shows distributions conditional on non-zero values to

[11] Details on the designation of sampling units at each stage of the survey design can be found on the RLMS website at `http://www.cpc.unc.edu/projects/rlms/project/sampling.html`

[12] Were there no attrition, this group would be representative of the Russian population in 1995, though it may not be representative of the actual population for subsequent years if the actual population differed from 1995.

Table 2: Summary Statistics

	Unbalanced Panel		Balanced Panel		
	Mean	**Std.Dev.**	**Mean**	**Std.Dev.**	**P-value**[1]
Demographic Characteristics					
Household Size[2]	2.88	1.45	2.95	1.47	0.00
Number of Elderly[3]	0.66	0.78	0.69	0.79	0.04
Adult Male H. of H.	0.61	0.49	0.62	0.49	0.50
Adult Female H. of H.	0.12	0.32	0.12	0.32	0.86
Elderly Male H. of H.	0.13	0.34	0.14	0.35	0.14
Elderly Female H. of H.	0.14	0.35	0.13	0.33	0.01
Rural	0.28	0.45	0.30	0.46	0.00
Poor[4]	0.45	0.50	0.46	0.50	0.18
Privatized Residence	0.63	0.48	0.64	0.48	0.10
N	12372		10076		
Use of Transactions[5]					
P	0.42	0.49	0.42	0.49	0.30
S	0.14	0.35	0.14	0.35	0.80
rK	0.11	0.32	0.11	0.31	0.46
L	0.59	0.49	0.59	0.49	0.33
τ	0.39	0.49	0.39	0.49	0.77
B	0.30	0.46	0.30	0.46	0.73
G	0.28	0.45	0.28	0.45	0.44
F	0.63	0.48	0.67	0.47	0.00
I	0.03	0.17	0.03	0.18	0.69
C	1.00	0.03	1.00	0.02	0.32
Other	0.99	0.08	1.00	0.07	0.43

Notes:
1. For t-test of the equality of means between the balanced and unbalanced panels.
2. Household size ranges from 1 to 12 with a standard deviation of 1.47.
3. Elderly for men is ages 60 and above; for women, 55 and older. Number of elderly household members ranges from 1 to 4 with a standard deviation of 0.53.
4. According to the All-Russia poverty index.
5. Transactions variables defined in the footnotes to Figure 1.

Figure 1: Distributions of Financial Transactions

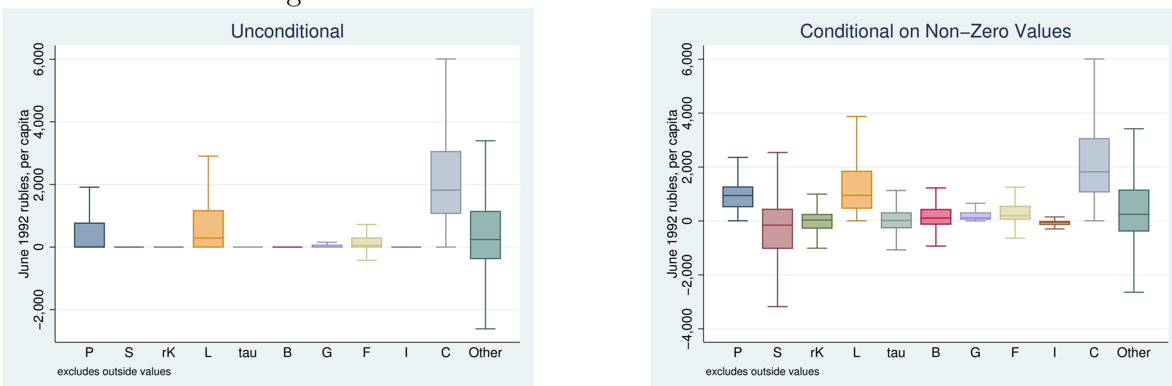

Definitions:
P = old-age pension
S = net savings (sum of expenditures on stocks, bonds, and other valuable papers and saving minus sales of personal property, jewelry, hard currency, and stocks, bonds, and other valuable papers)
rK = net earnings on assets (sum of interest on investments, repayment of loans outstanding and rental income minus payments on outstanding loans)
L = labor income
τ = net transfers[13] (sum of alimony and transfers received from relatives / friends / organizations minus alimony paid and transfers to relatives and friends)
B = net borrowing (loans taken out minus loans made)
G = government subsidies and benefits other than old-age pensions (sum of child, fuel, and apartment benefits/subsidies; unemployment benefits; student stipends; and other pensions/benefits)
F = farm income (net income, cash and non-cash, from home-produced goods plus income from sale of livestock / bees and goods harvested from the wild)
I = net insurance income (insurance pay-outs minus premiums paid)
C = consumption expenditures (including food, alcohol, tobacco, clothing, fuel, durable and luxury goods, services, rent & utility payments, and miscellaneous non-food expenditures)
Other = other net income

Figure 2: Use of Various Financial Transactions

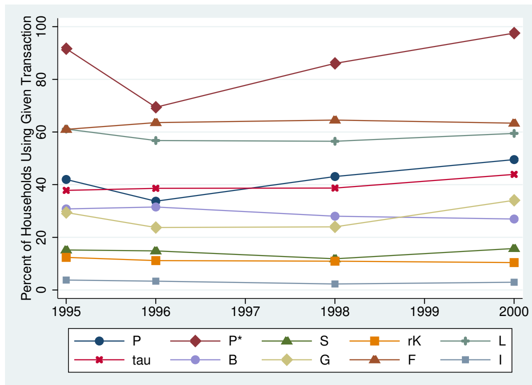

Notes: Over 99% of households in each year have nonzero values for both C and Other; these lines are omitted from the graph to reduce clutter.

give a sense of how important the various transactions are for those households that make use of them. The bottom panel of Table 2 gives percentages for how many observations in the panel make use of a particular financial transaction, in order to help relate the unconditional distributions to the conditional ones. Most income seems to come from pensions and labor, with consumption the most significant outflow. Farm income is also important to these households, but many households incur net expenses from farm activities[14]. Though relatively few households save or dis-save, those that do seem to make relatively large deposits and withdrawals. As we would expect from a representative sample, the distributions of earnings on assets, interpersonal transfers, and borrowing are all centered around zero. Formal insurance and government transfers other than pensions seem to play a relatively minor role in the budgets of these households. More detail on transaction use is shown in Figure 2 which breaks down the percentage of households that use a given transaction on an annual basis. Here

[14]Since surveys were conducted between October and December, our measure of farm income may be biased downwards. If farming households sell most of their harvest during the summer months and then save or invest their earnings, what we would like to categorize as farm income may be showing up as withdrawals from savings or earnings on assets. In this case, our estimates of the effect of an exogenous income shock will understate the response of farm income and overstate the responses of savings and/or earnings on assets. Our measure of farm income does include the imputed value of household production consumed during the past month.

we distinguish between the percentage of all households that received pension income (P on the graph) and the percentage of pension-eligible households that actually received pension income (P^* on the graph)[15]. The two lines clearly follow the same trend with significant arrears in 1996, though only around 50% of households had an elderly member who was eligible for pension payments in any given round.

5.2.1 "Other" Income

Since the "other" income transaction category seems to be relatively important, particularly in the unconditional distributions (see Figure 1), and is by definition unexplained, we might be concerned that these transactions are somehow correlated with household characteristics. To investigate this issue, we treat the data as a cross-section and regress "other" income on a variety of household characteristics. This kitchen-sink approach allows us to identify characteristics that are correlated "other" income, though of course it has nothing to say about causality.

In Column 1 of Table 3, we regress the levels of "other" income on known characteristics of the household's location, poverty status, composition, dwelling, and head. In this specification, we find no characteristics that are significantly correlated with the level of "other" income conditional on the other controls.

In Column 2, we use the absolute value of "other" income as the dependent variable, looking for correlations with simply the presence of "other" income, regardless of whether these were net flows in or out of the household's accounts. Here, we find that the survey seemed to do a relatively better job of collecting complete information on household finances in 2000 (relative to 1995) since "other" income was smaller in magnitude. We also find suggestive evidence that rural households have somewhat smaller values for "other" income than do urban households, implying that "other" income might not be a result of inaccurate tracking of agricultural activities. To the contrary, this result leads

[15]Were it not for pension arrears, P^* would be a horizontal line at 100%.

us to wonder if perhaps the financial transactions of households in urban areas are more complex than their rural counterparts, causing the survey to have a more difficult time eliciting complete financial information in cities. Households that live in privatized dwellings also appear to have larger magnitudes of “other” income. This could be due to such households being more “financially savvy,” and engaging in other transactions that are not tracked by the survey.

If we consider that “other” income might actually be due to households inaccurately reporting the components of their income, leading to accounting errors that show up as “other” income, we might be more concerned with whether or not certain types of households consistently report such discrepancies or if they seem to be random, averaging out to zero over several rounds for each household. To investigate this possibility, we regress the average level of “other” income for each household on the set of control variables in Column 3. The only marginally significant correlate in this specification is the number of elderly members of the household, suggesting that more elderly members might lead to a household having larger net expenses on average over the 4 rounds.

None of the three specifications are even relatively good fits of the data, however, and we conclude that there are no obvious patterns between observable household characteristics and “other” income. This of course does not necessarily mean that such correlations don’t exist for *unobservable* characteristics, but we proceed on that assumption.

5.3 Unbalanced Panel

As shown in Table 4, our data form an unbalanced panel since we do not observe the same set of households in every round. If attrition was not randomly assigned but rather was correlated with the assignment of pension arrears, then our estimates will no longer be consistent. While it is fundamentally impossible to prove that this is not a problem, in Section 6 we provide robustness checks that lend credibility to our claim

Table 3: Characteristics Correlated with "Other" Income

	Levels	Abs. Value	HH Mean
	(1)	(2)	(3)
1996	54.88	10.23	
	(307.09)	(310.83)	
1998	6571.08	6249.56	
	(6994.65)	(6992.87)	
2000	-361.07	-568.66**	
	(283.06)	(274.22)	
Rural	-519.40	-750.44*	-918.87
	(432.21)	(425.39)	(778.18)
Poor	-4168.49	-5513.73	1597.97
	(4566.64)	(4566.55)	(1741.53)
HH Size	889.20	814.48	622.01
	(1100.55)	(1102.66)	(799.63)
No. of Elderly	-821.86	-1032.58	-785.56*
	(707.43)	(710.47)	(405.93)
Privatized	132.61	444.97***	261.95
	(164.18)	(139.58)	(189.37)
Rental	644.59	394.02	-786.20
	(560.40)	(564.16)	(943.13)
Dorm	41233.25	41242.57	28929.85
	(39892.39)	(39911.25)	(28052.01)
Adult F. H.o.H.	-2761.19	-2943.65	-2160.57
	(2698.29)	(2684.46)	(2182.26)
Elderly M. H.o.H.	529.22	-208.18	2148.30
	(503.03)	(499.68)	(1765.98)
Elderly F. H.o.H.	553.51	288.06	1483.85
	(1115.52)	(1103.51)	(1777.11)
Obs.	11673	11673	3699
R^2	0.003	0.003	0.007

Notes: Huber-White standard errors are clustered by region. Single, double, and triple stars denote significance at the 10%, 5%, and 1% confidence levels, respectively.

Table 4: Observations by Round

Round	**Year**	**No. of HH's**	**Percent of Sample**
6	1995	3,781	30.56
7	1996	3,254	26.3
8	1998	2,818	22.78
9	2000	2,519	20.36
N		12,372	

that attrition plays a minor role in the parameter estimation.

6 Results

6.1 Parameter Estimation

6.1.1 Household Fixed Effects

We begin by estimating fixed effects models of Equation 4 in which identification comes from variation within each household over time. This approach controls for time-invariant unobserved heterogeneity that could otherwise bias coefficient estimates if correlated with the variables of interest. In these specifications, the idiosyncratic error term in the regression can be interpreted as the amount by which the household's optimal transaction level differed from the average societal rate of conversion from pension income into the transaction under consideration. Since we are assuming that all households follow the same decision-making process, these deviations must be due to problems implementing the desired optimal level of the transaction (i.e. a savings institution refuses to accept deposits, an employer pays an unexpected bonus, a check from a friend is lost in the mail, etc.).

We report three different fixed effects specifications: one in which all households are constrained to have the same response to pension shocks (Table 5), one in which responses are allowed to differ by household type depending on whether a household comprises only pensioners or includes other members (Table 7), and one in which responses are allowed to differ between urban and rural locations (Table 8). In all three specifications, we are most interested in the coefficients on pension income, though we note some interesting observations regarding the other control variables.

Our primary conclusion from these data is that while pension income does appear to increase consumption expenditures (Column 9), consumption does not move one-for-one with pension income. On the other hand, none of the other transactions except

Table 5: Regression Results – FE

	S	rK	L	τ	B	G	F	I	C	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
P	-.32	0.02	-.09***	-.08	-.005	-.04***	0.004	-.007***	0.28**	-.78**
	(0.44)	(0.02)	(0.03)	(0.06)	(0.04)	(0.01)	(0.03)	(0.003)	(0.14)	(0.36)
HH Size	1.13	-8.17	-38.08	26.41	-26.72	-3.70	-83.84***	-5.38	-404.37***	-275.25***
	(50.23)	(10.66)	(30.28)	(28.38)	(41.19)	(3.65)	(14.95)	(3.34)	(93.26)	(88.12)
No. of Elderly	24.23	10.46	-190.86***	-90.78**	-89.34	-75.50***	-47.80	7.54	-261.64*	177.35
	(185.29)	(17.07)	(43.10)	(41.43)	(109.16)	(8.90)	(35.64)	(5.05)	(135.14)	(214.88)
Privatized	22.65	-11.78	73.64	44.15	62.34	-8.29	-3.43	-6.05	-88.24	-205.57*
	(92.52)	(31.22)	(45.27)	(86.89)	(102.85)	(12.60)	(17.98)	(4.36)	(173.37)	(109.32)
1996	121.61	19.57	72.38	26.63	7.42	17.81**	-123.98***	4.93	-272.38**	-145.92
	(157.01)	(18.79)	(48.21)	(28.38)	(43.43)	(6.97)	(22.76)	(4.12)	(105.76)	(131.48)
1998	215.44	5.96	-251.67***	-12.83	-25.86	2.86	-48.66**	2.98	-993.26***	-420.94***
	(156.59)	(14.81)	(42.25)	(21.07)	(27.93)	(6.50)	(21.76)	(1.85)	(75.94)	(148.62)
2000	138.46	15.75	-29.77	9.15	-65.02	7.19	-81.74***	-.61	-607.11***	-305.88*
	(217.93)	(12.26)	(48.43)	(30.19)	(80.42)	(6.35)	(26.75)	(1.70)	(118.38)	(181.63)
Obs.	11336	11287	10789	11338	11252	11339	11339	11275	10789	11343
F statistic	0.96	0.92	12.53	1.97	1.31	12.08	8.95	1.87	40.34	11.59

Notes: Huber-White standard errors are clustered by region. Single, double, and triple diamonds denote significance at the 10%, 5%, and 1% confidence levels, respectively.

Table 6: Regression Results – FE, Balanced Panel

	S	rK	L	τ	B	G	F	I	C	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
P	-.40	0.03	-.10**	-.07	-.04	-.04***	0.004	-.008**	0.3**	-.83*
	(0.52)	(0.02)	(0.04)	(0.04)	(0.03)	(0.01)	(0.03)	(0.003)	(0.14)	(0.44)
HH Size	-6.61	-1.02	-19.20	-3.79	-12.08	-2.63	-85.74***	-6.11*	-298.24***	-181.76**
	(56.10)	(8.61)	(19.71)	(18.85)	(39.43)	(4.34)	(16.12)	(3.62)	(46.73)	(73.12)
No. of Elderly	29.27	11.21	-191.61***	-38.05	-120.03	-76.48***	-31.82	8.33	-270.56**	146.62
	(207.40)	(16.28)	(42.88)	(37.01)	(115.36)	(9.67)	(36.66)	(5.46)	(130.90)	(235.65)
Privatized	-42.99	6.24	76.83	76.50	62.64	-9.28	-6.02	-6.68	-75.74	-313.37**
	(72.35)	(15.31)	(53.15)	(82.96)	(90.08)	(11.94)	(19.75)	(4.88)	(164.41)	(128.86)
1996	186.25	30.70	44.45	34.27	-28.41	16.44**	-130.93***	5.92	-343.23***	-112.01
	(193.92)	(18.91)	(40.49)	(23.14)	(50.46)	(7.34)	(25.44)	(5.21)	(85.31)	(155.57)
1998	242.21	1.80	-260.19***	-12.31	-38.01	1.29	-50.82**	3.72**	-981.56***	-373.90**
	(183.92)	(16.11)	(41.01)	(21.04)	(33.40)	(7.04)	(24.18)	(1.88)	(78.13)	(166.97)
2000	176.28	17.16	-39.17	4.24	-76.47	6.39	-85.18***	-.08	-614.17***	-256.81
	(241.50)	(12.62)	(42.64)	(29.51)	(80.71)	(6.08)	(28.32)	(1.58)	(114.67)	(194.65)
Obs.	9345	9309	8950	9346	9282	9347	9347	9297	8950	9349
F statistic	0.64	1.66	13.5	1.29	1.36	10.47	8.04	1.81	33.21	8.05

Notes: Huber-White standard errors are clustered by region. Single, double, and triple diamonds denote significance at the 10%, 5%, and 1% confidence levels, respectively.

Table 7: Regression Results – FE by Household Type

	S	rK	L	τ	B	G	F	I	C	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
P-elderly	-.35	0.02	-.11***	-.10*	0.01	-.03***	-.002	-.005**	0.26*	-.82**
	(0.5)	(0.02)	(0.03)	(0.06)	(0.04)	(0.01)	(0.03)	(0.002)	(0.15)	(0.39)
P-nonelderly	-.18	0.02	-.008	-.02	-.07	-.10***	0.03	-.01***	0.38**	-.66*
	(0.26)	(0.03)	(0.07)	(0.08)	(0.12)	(0.03)	(0.04)	(0.006)	(0.16)	(0.36)
HH Size	-8.55	-7.96	-43.83	21.64	-22.37	0.17	-85.49***	-4.88	-411.49***	-284.23***
	(60.71)	(10.74)	(32.67)	(28.50)	(45.98)	(3.45)	(15.41)	(3.28)	(96.13)	(96.75)
No. of Elderly	22.71	10.49	-191.45***	-91.54**	-88.48	-74.88***	-48.06	7.63	-262.36*	175.88
	(182.84)	(17.03)	(42.76)	(41.32)	(107.60)	(9.05)	(35.69)	(5.06)	(135.27)	(214.17)
Privatized	24.50	-11.81	74.40	45.06	61.32	-9.03	-3.12	-6.13	-87.29	-203.84*
	(90.04)	(31.23)	(45.26)	(87.44)	(102.21)	(12.44)	(17.95)	(4.34)	(174.59)	(108.03)
1996	122.02	19.56	72.67	26.82	7.26	17.66**	-123.92***	4.91	-272.02**	-145.55
	(157.96)	(18.79)	(48.15)	(28.38)	(43.44)	(7.07)	(22.72)	(4.10)	(105.65)	(132.10)
1998	219.28	5.88	-249.29***	-10.93	-27.58	1.33	-48.00**	2.77	-990.33***	-417.38***
	(162.10)	(14.73)	(42.10)	(20.92)	(27.87)	(6.58)	(21.66)	(1.86)	(76.45)	(151.78)
2000	141.89	15.68	-27.53	10.84	-66.49	5.81	-81.16***	-.79	-604.34***	-302.69
	(222.72)	(12.25)	(48.04)	(30.46)	(81.63)	(6.41)	(26.70)	(1.71)	(118.05)	(185.31)
Obs.	11336	11287	10789	11338	11252	11339	11339	11275	10789	11343
F statistic	0.85	0.82	10.97	2.55	1.34	10.37	7.91	1.77	35.83	10.07

Notes: Huber-White standard errors are clustered by region. Single, double, and triple diamonds denote significance at the 10%, 5%, and 1% confidence levels, respectively.

Table 8: Regression Results – FE by Location

	S	rK	L	τ	B	G	F	I	C	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
P-rural	0.17*	0.03	-.05*	-.06	-.08	0.005	0.03	-.001	0.23	-.43*
	(0.09)	(0.02)	(0.03)	(0.11)	(0.05)	(0.007)	(0.09)	(0.005)	(0.16)	(0.22)
P-urban	-.51	0.01	-.11***	-.09	0.02	-.06***	-.008	-.009**	0.3*	-.93*
	(0.61)	(0.02)	(0.04)	(0.07)	(0.06)	(0.02)	(0.02)	(0.004)	(0.17)	(0.49)
HH Size	3.21	-8.10	-37.90	26.51	-26.99	-3.50	-83.71***	-5.36	-404.62***	-273.75***
	(49.43)	(10.63)	(30.27)	(28.42)	(41.37)	(3.58)	(14.92)	(3.33)	(93.25)	(87.89)
No. of Elderly	36.34	10.83	-190.01***	-90.17**	-91.06	-74.28***	-47.06	7.68	-262.81*	186.16
	(192.90)	(16.88)	(42.98)	(41.44)	(108.05)	(8.73)	(36.05)	(5.05)	(136.05)	(218.86)
Privatized	26.01	-11.68	73.88	44.32	61.85	-7.95	-3.23	-6.01	-88.57	-203.14*
	(91.41)	(31.20)	(45.21)	(86.79)	(102.87)	(12.49)	(17.91)	(4.36)	(173.19)	(108.85)
1996	132.10	19.90	73.21	27.15	5.88	18.87***	-123.34***	5.05	-273.53**	-138.34
	(163.70)	(18.75)	(48.27)	(27.73)	(43.27)	(7.18)	(22.89)	(4.16)	(106.91)	(135.71)
1998	209.82	5.76	-252.08***	-13.11	-25.01	2.30	-49.00**	2.91	-992.70***	-424.99***
	(151.84)	(14.81)	(42.33)	(21.17)	(27.97)	(6.49)	(21.61)	(1.84)	(75.55)	(145.69)
2000	132.23	15.54	-30.20	8.84	-64.12	6.56	-82.12***	-.69	-606.50***	-310.38*
	(213.20)	(12.23)	(48.46)	(30.17)	(80.82)	(6.32)	(26.54)	(1.73)	(117.64)	(178.85)
Obs.	11336	11287	10789	11338	11252	11339	11339	11275	10789	11343
F statistic	1.24	1.02	11.91	2.04	2.43	10.71	7.89	1.61	36.44	10.17

Notes: Huber-White standard errors are clustered by region. Single, double, and triple diamonds denote significance at the 10%, 5%, and 1% confidence levels, respectively.

labor income (Column 3), non-pension government transfers (Column 6), and "other" income (Column 10) seem very responsive to changes in pension income[16]. Despite the richness of this data, we are unable to identify many means by which consumption is protected from exogenous income shocks. We can interpret the very strong relationship between pension income and "other" income as evidence that the methods households are using to smooth consumption are not included in the list of transactions about which the survey inquired. This interpretation is consistent with the regression results to the extent that the estimated coefficients on "Other" are negative in all our specifications.

Allowing different types of households to have different responses to pension income changes the results somewhat. We observe that for households comprising only elderly members (Table 7), pension income reduces labor income (Column 3), as would be expected from traditional labor supply models in which leisure is a normal good[17]. For households that also include non-elderly members, the relationship between pension income and labor income is not as strong, in contrast to our hypothesis 2 in Section 3.4. This may be due to the fact that the ratio of pension to labor income for such households is smaller relative to households comprising only elderly members. In Table 8, we see that that pension income appears to reduce labor income more for urban households than their rural counterparts, perhaps because labor markets are more flexible in urban areas. However, neither of these differences are statistically significant[18].

As in the first specification which required all households to have the same responses to pension income, in Tables 7 and 8, Column 9, we see that a sizable portion of pension income is spent on consumption for all types of households, even when the coefficient on pension income is allowed to vary by household type. Surprisingly, the proportion devoted to consumption appears larger for households that include non-elderly mem-

[16]The coefficient on pension income in the formal insurance regression (Column 8) is also statistically significant but is so small as to be economically insignificant.

[17]Recall that old-age pensions were not means-tested and were unaffected by present earnings so there are no negative work incentives created by these government benefits.

[18]P-values of 0.20 and 0.26 in Wald tests of equality for point estimates on pension income in labor income regressions by household member type and location, respectively.

bers than for elderly-only households, though this difference is not significant. The consumption responses of rural and urban households to pension income are also not significantly different from one another[19].

Though the point estimates on pension income for elderly-only and counterpart households also appear to differ from one another in the interpersonal transfer, and "other" income regressions (Table 7, Columns 3, 4, and 10, respectively), we fail to reject the null hypotheses of equality using Wald tests[20]. Similarly, in the regressions broken down by location (Table 8), none of the point estimates differ significantly between rural and urban households[21]. From these results, we infer that households' allocations of pension income into the various financial transactions does not differ substantially based on the two household characteristics under consideration, namely, household composition and location.

Recall Equation 5 from Section 4.2.1, which required a certain relationship among the parameter estimates for all transactions. We calculated the sum of the parameters on pension income in each specification (taking into account the correct signs), and present these results in Table 9. While the identity does not hold exactly in any specification, in only one of our six specifications is the sum significantly different than 1 at a 95% confidence level, and even in that case, with a slightly larger confidence interval we would no longer reject the null hypothesis that the model restrictions hold. These small discrepancies from 1 could be the result of rounding errors due to lack of precision in the estimation process. We conclude that our structural model of pension allocation into transactions is respected in our estimation procedure.

Finally, we note some common trends among the other control variables. In all three

[19]P-values of 0.41 and 0.74 in Wald tests of equality for point estimates on pension income in consumption regressions by household member type and location, respectively.

[20]P-values of 0.23, and 0.62 in Wald tests of equality for point estimates on pension income in elderly-only versus counterpart households in the interpersonal transfer and "other" income regressions, respectively.

[21]P-values of 0.23, 0.23, and 0.33 in Wald tests of equality for point estimates on pension income in rural versus urban households in the savings, formal insurance, and "other" income regressions, respectively.

Table 9: Restrictions on Parameters

Specification	Sum of Coefficients on P	95% Confidence Interval[1]
FE	0.9596	0.9159-1.0033
FE, Balanced Panel	0.9530	0.9053-1.0007
FE, by Household Type		
Elderly	0.9433	0.8990-0.9876
Non-elderly	1.0262	0.9563-1.096
FE, by Location		
Rural	0.9523	0.8928-1.0118
Urban	0.9625	0.9071-1.0179

Notes:
1. Confidence intervals based on standard errors from 1000 non-parametetric bootstrap replications, clustered by region.

fixed effects specifications, household size seems to be correlated with average levels of farm income, consumption, and "other" income (Columns 7, 9, and 10, respectively), such that larger families have less of all three transactions. Households with more elderly members seem to have lower labor income, less interpersonal transfers in absolute terms, and less consumption. Living in a privatized dwelling, on the other hand, seems to have no relationship with transaction levels other than "other" income, which is significantly less for households that live in privatized dwelling (Column 5). The year dummy variables highlight that in general conditions were changing drastically over this period. The financial crisis of 1998 shows up in the very negative point estimates on 1998 in the labor and consumption regressions (Columns 3 and 9), with lingering effects into 2000 as shown by the negative point estimates on that variable.

6.1.2 Random Coefficients

We now relax the assumption that all households follow the same decision-making process and instead allow each household to allocate its pension income over the various transactions in an idiosyncratic manner. We use Swamy's random coefficient estimator for panel data, which provides a consistent and asymptotically efficient estimator for the mean coefficient in the sample and an unbiased and consistent estimator for the variance-covariance matrix [11]. Essentially, the random coefficient estimator is a weighted average of the household-level OLS coefficients, where the weights are inversely proportional to the covariance matrices of these coefficients. The model assumes that the idiosyncratic errors are structurally and temporally uncorrelated. Moreover, they are assumed to be normally distributed with mean zero and are heteroskedastic with different variances for each micro unit (in our case, households). The coefficient vectors are assumed to be random and uncorrelated across micro units, while following the same distribution. Our results for the unbalanced panel are presented in Table 10. Here, the only independent variables are pension income and the intercept. This is partially out of necessity, since we now have only three or four observations per individual regression, but also reflects the fact that all time-constant household characteristics are controlled for by means of the intercept since the regressions are run at the household level. We again interpret the idiosyncratic error term as a reflection of the household's inability to implement the optimal level of a particular transaction, assuming each household's decision-making process is constant over time.

As in the fixed effects analysis, we find that pension income leads to significantly and substantially increased consumption. We also find the common result that pension income significantly reduces labor income, though the effect does not seem to be as strong as the fixed effects analysis indicated. Interestingly, in the random coefficient analysis, pension income also appears to lead to significantly less net transfer income, non-pension government transfers, and net insurance income, though this last effect is quite small in magnitude.

The power of pension income to explain the various transactions is quite low in some of these regressions, but for those regressions in which the parameter estimates are significantly greater than zero, we consistently reject the null hypothesis that the

Table 10: Regression Results – Random Coefficients

Dependent Variable	β_P	**Std. Err.**	**# of Obs.**	**# of Groups**	**Wald P-value**	**Parameter Constancy P-value**
S	0.0241	0.0005	10898	2810	0.2748	0.0000
rK	0.0024	0.0000	10846	2805	0.7179	0.0000
L	-0.0367**	0.0003	10205	2673	0.0465	0.0000
τ	-0.051***	0.0002	10902	2811	0.0003	0.0000
B	-0.0906	0.0068	10802	2799	0.2727	0.0000
G	-0.0129***	0.0000	10903	2811	0.0000	0.0000
F	0.0077	0.0001	10903	2811	0.5053	0.0000
I	-0.0013***	0.0000	10832	2804	0.0049	1.0000
C	0.1576***	0.0018	10205	2673	0.0002	0.0000
Other	-0.0554	0.0045	10914	2814	0.4090	0.0000

Notes: Standard errors are calculated using 100 bootstrap replications. Single, double, and triple stars denote significance at the 10%, 5%, and 1% confidence levels, respectively.

regression does no more than explain the mean of the dependent variable (see the Wald P-values in Table 10). We also note that in all cases, a test of parameter constancy (across households) is strongly rejected. While the random coefficients analysis imposes some strong distributional assumptions, these results suggest that households may not follow the same decision-making process.

Note that the identity in Equation 5 will not necessarily hold for the results in Table 10, since the reported coefficients are the means of a distribution of household-level parameter estimates for each transaction. However, the identity will still hold at the household level[22].

6.2 Robustness Checks

Our primary concern is that attrition from the sample may have biased our estimates. If the factors that lead to attrition are time-constant, then our fixed effects estimators are consistent. On the other hand, if such factors are correlated with the idiosyncratic error terms, additional means of correcting for attrition are necessary. In this section

[22] An analysis of household-level coefficients is available from the author.

Table 11: P-Values in Test of Attrition Bias

S	0.53
rK	0.26
L	0.42
τ	0.38
B	0.08
G	0.85
F	0.71
I	0.12
C	0.77
ε	0.60

Notes: P-values are for standard t-tests that the coefficient on the attrition indicator is different from zero. With the exception of the addition of the attrition indicator, the estimated regression model is the same as in Table 5.

we first present a simple test of the assumption that attrition is uncorrelated with the idiosyncratic error terms and then discuss the results of estimating our model on the balanced panel of households for which we have observations in every round.

6.2.1 Testing for Attrition Bias

Wooldridge [13] suggests a modification of a test for attrition bias initially proposed by Nijman and Verbeek [8]. We generated an attrition indicator variable that switches from zero to one in the period before a household drops out of the survey. Under the null hypothesis of no attrition bias, the idiosyncratic error term is uncorrelated with the attrition indicator for all rounds. As a result, the attrition indicator should have no explanatory power. While this test does not prove that our results are free of attrition bias, we do fail to reject the null hypothesis of no attrition bias at reasonable confidence levels for all but two of our transactions as shown in Table 11.

6.2.2 Balanced Panel Results

If attrition is randomly assigned, then the unbalanced panel should not differ from the balanced panel in any significant way. Comparing the fixed effect estimation results from the balanced panel to those from the unbalanced panel presented above in Tables 5 through 8 thus provides an informal means of assessing the potential severity of attrition bias. We do not include all the balanced panel results here[23], but the point estimates are quite similar between the two panels. We find almost no differences between the unbalanced and balanced panels in the general fixed effects specification (Tables 5 and 6), and the specifications broken down by household type and location are equally good matches. The similarities between the two panels lend credibility to our assumption that attrition did not significantly bias our point estimates.

7 Conclusion

Whereas the socialist economic system had sought to reduce idiosyncratic risk to only aggregate risk faced by society as a whole, in the transition to a capitalist economy households had to adjust to dealing with increased income variability, some of which was even caused by the government itself, as in the case of pension arrears. We have attempted to explain the mechanisms that households used in order to smooth consumption during this period. However, despite the fact that our rich panel dataset allowed us to account for time-constant unobserved heterogeneity and the endogenous nature of the set of financial transactions that households use, we are able to draw only a few illuminating conclusions about how households dealt with shocks to pensions, an exogenous source of income. While both our fixed effects and random coefficients analyses provide evidence that not all pension income is allocated to consumption expenditures, the only other relatively robust finding is that labor income appears to be a substitute for pension income. As the previous literature has often concluded, households do smooth their consumption, though with a large set of financial transactions at their disposal to do so, choice of consumption-smoothing techniques appears to be idiosyncratic.

[23] Available from the author by request.

The lack of other clear patterns in these data lead us to wonder how a household determines which methods it will employ in order to stabilize consumption in the face of variable income. Future research could focus describing what types of households make use of a certain transaction and ultimately identifying the causes for these choices. Are all types of transactions available to all households? Which transactions are most effective at smoothing consumption? How much do households depend on the various smoothing mechanisms? As is often the case, this research paper has raised more questions than it has answered.

References

[1] Kenneth Arrow and Gerard Debreu. The Existence of an Equilibrium for a Competitive Economy. *Econometrica*, 22:265–290, 1954.

[2] Regional Bureau for Europe and the CIS, editors. *Poverty in Transition?* United Nations Development Programme, New York, NY, 1998.

[3] Milton Friedman. *A Theory of the Consumption Function.* Princeton University Press, 1957.

[4] Robert T. Jensen and Kaspar Richter. The Helath Implications of Social Security Failure: Evidence from the Russian Pension Crisis. *Journal of Public Economics*, 88:209–236, 2003.

[5] Jeni Klugman, editor. *Poverty in Russia: Public Policy and Private Responses.* The World Bank, Washington, DC, 1997.

[6] Ethan Ligon. Risk Sharing and Information in Village Economies. *Review of Economic Studies*, 65:847–864, 1998.

[7] Youngjae Lim and Robert M. Townsend. General Equilibrium Models of Financial Systems: Theory and Measurement in Village Economies. *Review of Economic Dynamics*, 1:59–118, 1998.

[8] Theo Nijman and Marno Verbeek. Nonresponse in Panel Data: The Impact on Estimates of a Life Cycle Consumption Function. *Journal of Applied Econometrics*, 7(3):243–257, 1992.

[9] Christina Paxson. Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand. *The American Economic Review*, 82(1):15–33, 1992.

[10] Emmanuel Skoufias. Consumption Smoothing in Russia: Evidence from the RLMS. *Economics of Transition*, 11(1):67–91, 2003.

[11] P.A.V.B. Swamy. Efficient Inference in a Random Coefficient Regression Model. *Econometrica*, 38(2):311–323, 1970.

[12] Robert M. Townsend. Risk and Insurance in Village India. *Econometrica*, 62(3):539–591, 1994.

[13] Jeffrey M. Wooldridge. *Econometric Analysis of Cross Section and Panel Data*. Cambridge: The MIT Press, 2002.