Limited Insurance within the Household: Evidence from a Field Experiment in Kenya[†]

By Jonathan Robinson*

In developing countries, unexpected income shocks are common but informal insurance is typically incomplete. An important question is therefore whether risk-sharing within the household is effective. This paper presents results from a field experiment with 142 married couples in Kenya in which individuals were given random income shocks. Even though the shocks were small relative to lifetime income, men increase private consumption when they receive the shock but not when their wives do, a rejection of efficiency. Such behavior is not specific to the experiment—both spouses spend more on themselves when their labor income is higher. (*JEL* D14, D81, G22, O12, O16)

Individuals in developing countries are subject to considerable risk but most lack access to formal mechanisms that would allow them to insure themselves against unexpected income shocks. Instead, households tend to use informal systems of gifts and loans to pool idiosyncratic risk. While these informal networks do provide some protection against shocks, they also face substantial problems of asymmetric information and payment enforceability, and existing evidence suggests that inter-household risk sharing networks are rarely, if ever, efficient (for example, see Townsend 1994; Udry 1994; Fafchamps and Lund 2003).

In the absence of effective formal or informal inter-household insurance mechanisms, a natural place for individuals to choose to cope with risk is within the household. Though such arrangements will be somewhat limited because income shocks are likely to be correlated within households, whether these mechanisms are effective in insuring the idiosyncratic risk that remains is an important question. If risk is not insured even within the household, despite the substantial incentives household members should have to insure each other in the absence of other risk-coping strategies, then programs which impact the ability of individuals to cope with risk will

likely have large welfare impacts (such as formal savings accounts or microinsurance programs).

This paper presents results from a field experiment in Kenya designed to directly test whether intra-household risk-sharing arrangements offer full insurance. The experiment followed 142 married couples for 8 weeks. Every week, each individual had a 50 percent chance of receiving a 150 Kenyan shilling (KS) (US $2.14) income shock, equivalent to roughly 1.5 days' income for men and 1 week's income for women. Information about the shocks was public knowledge—both spouses were told what their partner received. As these shocks are, by definition, random, transitory, and idiosyncratic, and since the payout of the shocks is public knowledge, the experimental design makes it possible to directly and simply test for allocative efficiency, by comparing the difference in the responsiveness of private consumption to shocks received by an individual to those received by his spouse.

The empirical approach is based on the assumption that, even though men and women may have very different preferences, the experimental shocks are too small (relative to lifetime income) to affect intra-household bargaining power. This is in contrast to larger income shocks, which may well affect bargaining power and, by extension, consumption decisions.[1] If the shocks are small, however, and so long as household members are risk averse, failing to insure the shocks would leave potential gains from trade unexploited and would constitute a rejection of the collective model of the household (i.e., Chiappori 1992; Browning and Chiappori 1998; Browning et al. 1994), which is based on the assumption that even if spouses have different preferences and bargain over outcomes, they are still able to achieve a Pareto efficient outcome.

In the context of this experiment, if the household pools risk efficiently, increases in private consumption should be the same for shocks received by an individual and those received by his spouse. However, I find that husbands increase their expenditures on privately consumed goods in weeks in which they receive the shock but do not change their expenditures in weeks in which their wives do, a rejection of Pareto efficiency. I do not detect statistically significant responses to either type of shock for women. These general results are robust to examining changes over several weeks rather than to just the week in which the shock was received. Note also that since the experimental shocks are fully observable by both spouses, behavior cannot be attributed to information asymmetries.[2]

While the experimental design allows for straightforward inference, the setup admittedly comes at some cost as well. In particular, the results come from a stylized experiment in which all shocks were positive. If people spend windfall income differently than their regular labor income, the results may not generalize. I attempt to test this by examining how private expenditures respond to weekly fluctuations in labor income, and I find that both men and women increase private expenditures in weeks in which their labor income is higher (this increase in response to income

---

[1] Many studies have shown that household decisions are sensitive to ostensibly exogenous changes in relative intra-household incomes. Examples include Duflo (2003); Thomas (1990); Lundberg, Pollak, and Wales (1997); and Haddad and Hoddinott (1995). Similarly, Anderson and Baland (2002) argue that intra-household conflict over savings/expenditures is a reason that so many women join Rotating Savings and Credit Association (ROSCAs) in Kenya.

[2] See Ashraf (2009) for evidence that information can affect intra-household savings decisions.

shocks is similar to that found by Duflo and Udry (2004) with respect to harvest income shocks). While the changes in labor income that I observe here are not necessarily exogenous and so should be interpreted with some caution, they are at least very suggestive that the overall findings are robust. A second issue is that while I have detailed data on each household in the sample, there are relatively few households (142) and all of them were sampled from daily income earners in one part of Western Kenya.

While the experimental shocks are not very large in themselves, the costs of not being able to cope with such shocks can actually be substantial. For example, in this same part of Kenya, Robinson and Yeh (2011) show that even daily health shocks affect the decision of sex workers to engage in unprotected sex. Similarly, Dupas and Robinson (2012) show that 31 percent of a sample of ROSCA participants encountered a health shock over the previous three months that they could not afford to fully treat. In the latter case, providing even basic savings products mitigated such vulnerability, and the demand for such products was substantial. The findings in this paper, which suggest that risk is uninsured even within the household, are therefore complementary, and suggest programs that provide more formal risk coping mechanisms could improve welfare.[3]

I. Relationship to Existing Literature

This paper contributes to a large and growing literature testing for intra-household efficiency that builds off the work of Pierre-André Chiappori and others. A number of early papers (most focused in developed countries) find evidence consistent with efficiency. In addition to others, these include Browning and Chiappori (1998) for Canada; Chiappori, Fortin, and Lacroix (2002) in the United States; Bourguignon, et al. (1993) for France; and Thomas and Chen (1994) for Taiwan.

There have also been a number of papers testing for efficiency in developing countries. These studies typically test either for productive efficiency (that households maximize profits) or for allocative efficiency (by testing whether allocation decisions are sensitive to transitory income shocks). The most notable study in the former category is Udry (1996), which rejects efficiency by showing that inputs could be profitably reallocated from male-controlled plots to female-controlled plots in Burkina Faso. Another relevant study is Schaner (2012), in which a field experiment is conducted to show that households are unable to maximize returns on savings because spouses have different discount rates.

This study is one of a number of papers to instead test whether couples are able to insure each other against income shocks. All of these studies require the identification of exogenous, idiosyncratic shocks that affect income realizations but do not affect preferences or intra-household bargaining power. Thus, while the shocks must be substantial enough to be economically meaningful, they must not be so large as

---

[3] An important question is why insurance is limited in this setting. In an earlier version of this paper, I find some suggestive evidence that insurance is constrained by limited commitment. However, the results are not conclusive. For evidence of limited commitment in risk sharing agreements, see Coate and Ravallion (1993); Ligon, Thomas, and Worrall (2002); Foster and Rosenzweig (2001); and Wahhaj (2007).

to affect bargaining weights. Typically the shocks which are used are rainfall or weather shocks among agricultural households (Duflo and Udry 2004; Dubois and Ligon 2009; Doss 2001; Bobonis 2009), health shocks (Dercon and Krishnan 2000; Goldstein 2004), or agricultural shocks, such as pests or plant disease (Goldstein 2004). With the exception of the Bobonis (2009) study in Mexico, each of these papers rejects efficiency.

This paper contributes to this literature in four main ways. First, this paper is one of the first to examine the consequences of relatively small weekly shocks among people who earn most of their income through daily labor. Most existing work focuses instead on relatively major illness or agricultural shocks among farming households, so relatively little is known about the effect of week-to-week shocks on outcomes. The little evidence that does exist suggests that such shocks might actually have important effects. For example, Robinson and Yeh (2011) show that daily health shocks affect the willingness of sex workers in this part of Kenya to have unprotected sex. Over a long enough time period, these shocks can affect the likelihood of getting HIV or other sexually transmitted diseases.

While these types of responses might be very important, it is almost impossible to measure them without precise data collected at high frequency, and with good measures of shocks. This paper is able to do this because the shocks are experimentally manipulated, and because the weekly monitoring survey was designed specifically to measure individual expenditures (and other variables) as precisely as possible. Furthermore, any evidence that households cannot insure each other against even small shocks strongly suggests that they cannot insure each other against bigger shocks.

Along the same lines, focusing on daily income earners is relevant since many people in developing countries earn at least some income from such occupations (i.e., Banerjee and Duflo 2007). This paper is therefore complementary to the much larger literature focusing on agricultural households.

Second, even relative to income from a single agricultural season, the shocks here are extremely small and so almost certainly of no consequence of household bargaining weights. This study should therefore be able to rule out this alternative hypothesis for the main results.

Third, there is evidence that the ability of households to achieve efficiency is context-specific. For example, Akresh (2008) performs similar tests to Udry (1996) in Burkina Faso, and finds that households achieve efficiency in some regions but not others. He also finds that efficiency is more likely to be obtained in poor rainfall years, presumably because the costs of inefficiency are particularly high when yields are low. Similarly, Rangel and Thomas (2005) use production and consumption data to test for efficiency among households in three West African countries (Burkina Faso, Ghana, and Senegal). Although tests using production data suggest inefficiency, the authors cannot reject efficiency with regard to consumption allocations. To the extent that the question of household efficiency is not settled, it is useful to collect more evidence on how and in what situations households are able to insure each other effectively.

This paper also contributes to a small experimental literature on risk sharing. However, the current paper is the only one (to my knowledge) to work with real-world, risk-sharing networks and to observe outcomes outside of a laboratory

or other controlled setting.[4] In the laboratory, Charness and Genicot (2009) examine risk sharing among UCLA undergraduates. Those studies that work with preexisting insurance networks include Barr (2003), Iversen et al. (2006) and Attanasio et al. (2012), and all look at behavior within a controlled experiment among households that share risk outside the experiment (in Zimbabwe, Uganda, and Colombia, respectively). Similarly, Chandrasekhar, Kinnan, and Larreguy (2010) test for limited commitment and for the role of access to savings within a controlled experiment in India. The closest study to this one is likely Ashraf (2009), which examines how observability and communication possibilities affect intra-household savings decisions in the Philippines. This experiment differs, however, both in that information about the shocks is always public and that the focus is on risk rather than on savings.

II. Theoretical Framework

In this section, I lay out a brief motivating framework for interpreting the main results (this follows from Browning and Chiappori 1998 and related papers, as well as Duflo and Udry 2004). Under the Pareto efficient collective model of the household, the household's optimization problem can be written as maximizing the following utility function:

$$(1) \qquad \max_{\{q_{mt}, q_{ft}, Q_t, L_{mt}, L_{ft}\}} E_t \sum_{t=0}^{T} \beta^t (u_m(c_{mt}, C_t, L_{mt}) + \lambda u_f(c_{ft}, C_t, L_{ft})),$$

subject to a budget constraint in which all sources of household income are pooled:

$$(2) \qquad Y_t = rY_{t-1} + w_m H_{mt} + w_f H_{ft} + S_{mt} + S_{ft} - p_{1t}(c_{mt} + c_{ft}) - p_{2t} C_t.$$

For all variables, the subscript m refers to the male and f to the female. The vectors c_{mt} and c_{ft} refer to private consumption, while C_t refers to shared consumption. p_{1t} and p_{2t} are prices for private and shared consumption, respectively, which I assume do not vary across spouses within the household.[5] L is leisure hours, H is labor hours, and w is the wage rate. Y_t is household wealth, which earns a gross return r in any period.

The key variables for this experiment are S_{mt} and S_{ft}, the experimental shocks. The key assumption is that $\frac{d\lambda}{dS_{mt}} = \frac{d\lambda}{dS_{ft}} = 0$: receiving the income shocks has no effect on the bargaining weight. This seems plausible given that the shocks represent only a day and a half's worth of income for men and a week's for women. If so, consumption should be allocated such that shocks should have the same effect whether they are received by the male or female.

---

[4] Another paper that makes use of random variation in incomes and observes real-world outcomes is Angelucci, De Giorgi, and Rasul (2011), which evaluates the effect of the Mexican Programa de Educación, Salud y Alimentación (PROGRESA) program on risk sharing and investment within extended families.

[5] Note that even if prices do differ across spouses (for instance, because travel costs are lower for men than for women), the basic implication of the model (that the responsiveness of private consumption should be the same for shocks received individually and shocks received by the spouse) should still hold.

To test the model empirically, I assume that leisure, private consumption, and shared consumption are additively separable. While these assumptions are restrictive, they seem to fit the data reasonably well, since neither labor supply nor shared consumption respond to the experimental shocks in the data (as I will show later).[6]

With this setup, the key test in this paper will be whether the household insures idiosyncratic risk. In particular, under the collective model, the household should allocate consumption such that $\frac{\partial u_m}{\partial c_m} = \lambda \frac{\partial u_f}{\partial c_f}$. In the context of this experiment, this implies that income shocks received by the husband should be spent in the same way as income shocks to the wife. To keep the analysis straightforward, I focus principally on private consumption. I do this because it is impossible to determine in the data who benefits from household public goods. For example, if male income shocks are spent more on shared food and female shocks more on children, it is difficult to evaluate the welfare impact of one type of spending against another. While it's still the case that income shocks should be spent similarly no matter who receives them, inference is much less crisp than in the case of private consumption, which can be directly assigned to one spouse or the other. Ultimately, the test of efficiency therefore boils down to testing whether private consumption is more responsive to own income shocks than to spousal income shocks.

Note, however, that the experimental income shocks are not only idiosyncratic but also transitory. Thus, the household should smooth consumption over these shocks (as has been tested in, for example, Paxson 1992). There is reason to suspect that households in Kenya are unable to do this, however, in large part because access to formal consumption smoothing mechanisms are limited and people have difficulty saving as much as they would like informally (Dupas and Robinson 2011). If, however, households manage to fully smooth consumption across weeks, the tests here will have no power.

A similar issue is that spouses need not actually pool risk with each other for full insurance to be obtained. It is possible that each spouse has her own informal insurance network in which she pools risk, but for these networks to not overlap. For example, each spouse might pool risk only with her own friends and relatives. In such a case, insurance will appear to be complete even though spouses do not share risk with one another.[7] Thus, as with intertemporal consumption smoothing, efficiency can only be rejected if inter-household informal insurance is incomplete.[8]

III. Empirical Specification

The main prediction of the collective model is that shocks received by the husband should have the same effect on the ratio of marginal utilities as equally sized shocks

---

[6] See Blundell, Chiappori, and Meghir (2005) for a formal extension of the collective household model with public goods. Identifying the model requires either separability of public consumption (as I assume here) or the identification of a distribution factor that affects intra-household bargaining weights but not preferences.

[7] See Goldstein (2004) for an example of this in Ghana.

[8] There are other ways of testing for efficiency. In particular, Bourguignon, Browning, and Chiappori (2009) derive testable conditions even if consumption categories are not assignable to a particular individual. However, those tests require a distribution factor which affects allocations but not preferences. I do not have an exogenous distribution factor in this dataset.

received by the wife. Testing this empirically therefore requires some assumption about preferences. In this paper, I assume that preferences are exponential such that they exhibit constant absolute risk aversion (CARA).[9] With this assumption, the basic empirical specification can be written as follows:

$$(3) \qquad y_{ht}^i = \gamma S_{ht}^i + \delta S_{ht}^j + \nu_i + \mu_t + \varepsilon_{ht}^i,$$

where h indexes the household and t time. The regression is run separately for both genders, i (where j indexes the spouse) and y_{ht}^i are the outcomes of interest. I will focus principally on private consumption (since it is impossible to know what fraction of shared consumption is consumed by an individual household member), but I will also present results for other consumption categories, as well as labor supply and savings. S_{ht}^i and S_{ht}^j are the amount each spouse received in experimental shocks (in Kenyan shillings). Finally, μ_t is a fixed effect for the week of the interview, and ν_i is an individual fixed effect (note that since separate regressions are run for each gender, the fixed effect can be thought of either as an individual or household fixed effect). Identification is based on the assumption that weeks in which a given household receives the shock are randomly determined.[10]

The test of Pareto efficiency is simply that the shocks only affect private expenditures through their effect on the pooled budget constraint, or that

$$(4) \qquad \gamma = \delta.$$

As the money might not be spent immediately, I run another specification in the online Appendix that includes current and lagged shocks. Nevertheless, as discussed in the previous section, the test is only powered to the extent that intertemporal consumption smoothing and inter-household insurance is incomplete.

IV. Experimental Design

A. *Sampling*

This project was conducted between April and October 2006 among a sample of 142 couples, drawn from a group of daily income earners (men who work as bicycle taxi drivers, called *boda bodas,* in Kiswahili, and women who sell produce and other items in the marketplace) in three towns in Western and Nyanza Provinces, Kenya.[11] Daily income earners were targeted because the project is focused on transitory shocks to weekly income, which are more commonly encountered among daily income earners than in a sample of, for instance, farmers. The sample is similar to

---

[9] An alternative approach would be to use CRRA preferences, which would amount to using logs rather than levels as dependent variables. I do not do this because there are weeks in which expenditures in a given category are equal to zero so that logs are undefined. Results do, however, look very similar with logs for the subsample of weeks in which expenditures are nonzero.

[10] If the shocks are truly random, then they should have no effect on outcomes in the weeks before they are received. Online Appendix Table A1 implements this regression and, reassuringly, finds no effects from this placebo test.

[11] The towns were Busia, Sega, and Ugunja.

that in Dupas and Robinson (2011), though drawn from different market centers. Also, the sample in this paper includes the spouses of all participants.

The towns targeted in this study are semi-urban areas located along a major highway from Nairobi, Kenya to Kampala, Uganda. Though many people in the area earn their living from agriculture, a substantial fraction earn at least some income from self-employment, as is common in the developing world (i.e., Banerjee and Duflo 2007). Many of these individuals work in town during the day but live in the surrounding rural areas.

To recruit individuals into the study, a trained enumerator conducted a census in the market centers of the three towns selected for the study. For the screening interview, the enumerator approached an individual at his place of work and asked to meet with him individually for a few minutes. The enumerator first asked the individual if he was married, and all those who were single were not interviewed further. For those who were married, the enumerator then asked the respondent if he would be interested in participating in a project that would take approximately eight weeks to complete, and that would require the administration of weekly monitoring surveys to both the respondent and his spouse. A precondition for participation was that the enumerator be allowed to visit the spouse at home without the primary respondent's supervision. Individuals were told that the weekly monitoring survey would take approximately one hour per week to complete, and that they would be compensated if they agreed to participate. If the individual was interested in the project, the enumerator took the respondent's name and contact information, and told the respondent that he would return later to begin the project. The spouse's consent was obtained later, at the first monitoring interview.

In total, 181 married individuals were interviewed during the screening interview. Of these, 142 couples enrolled in the full study (78.5 percent). Of the 39 couples who did not participate, 22 were not interested in participating, 6 could not be included because the spouse was often away and couldn't be traced for interviews, 6 were never found after the initial interview, 2 had moved, 2 were sick, and 1 person's spouse died shortly after enrolling in the study.

B. *Experimental Income Shocks*

To test whether intra-household insurance was complete, this project randomly provided 150 Kenyan shilling (about US $2.14)[12] income shocks to participants at the end of the weekly monitoring visit. The probability of receiving the shock in a given week was 50 percent for each spouse. Thus, there are weeks in which both spouses got the shock, weeks in which only one or the other got the shock, and weeks in which neither did. To make the payment of the shocks as transparent as possible, each enumerator carried with him a black plastic bag containing 56 slips of paper with the numbers 1–56 on them. Each number corresponded to a payment for *both* spouses. For each spouse, the drawing of 28 of the slips resulted in payment,

---

[12] The exchange rate was about 70 Kenyan shillings (Ksh) to US $1 during the study.

while the drawing of the other 28 resulted in no payment.[13] To make sure people understood the payouts, enumerators explained the experiment in great detail to people and showed them an example of how the drawing of the slips would affect payouts. At the conclusion of the visit, the enumerator answered any questions the respondents had, and went over any parts that were confusing. After the training, all respondents reported that they understood how the experiment was going to work.

The shocks were announced to each spouse, so that each knew what the other had gotten. Payments were made privately, however, and individuals were told that they could spend the money however they chose.

The experimental design has several key features. First, while the shocks are small compared to total lifetime income, they are not trivial either—they are equivalent to approximately 1.5 days' income for men and 7 days' income for women. Second, the shocks were announced to both spouses and thus publicly observable. While making the shocks fully observable does not fully mimic the shocks people receive in their normal lives (since many of those shocks are only partially observable), the advantage of providing full information is that any observed inefficiency is not attributable to the information available to the spouse.[14] Third, through the data collected with the monitoring surveys, it is possible to compare the experimental results with real world responses to fluctuations in weekly labor income.

An important disadvantage of the study, which is important to acknowledge, is that (for ethical and practical reasons) the income shocks provided were always positive, unlike real-world shocks, which can of course be either positive or negative. Thus, it's possible that people may have treated these payments as "windfall" income. I will attempt to address this in the empirical section by testing whether private expenditures respond to more natural labor income fluctuations. I find qualitatively similar results from that approach.

V. Data

There are three main data sources in this paper. First, a background survey was administered which included basic questions on demographics, credit, savings, asset ownership, and related issues. An important note is that the background survey was conducted at the end of the study and some individuals were not traced for it. For this reason, I only perform a check of randomization balance for characteristics that could not have possibly been affected by treatment (i.e., demographic characteristics). Second, a separate survey was administered to measure risk aversion. The survey followed Charness and Genicot (2009) and asked respondents to choose how much of a given amount they would like to invest in a risky asset that paid off 2.5 times the amount invested 50 percent of the time, but for which the

---

[13] The reason that there are 56 slips of paper, rather than just 4, is that the experiment was originally designed to vary the correlation in shocks between spouses to test for limited commitment. However, as the sample size was too small to make strong inferences on limited commitment, I do not report the results of that part of the experiment here.

[14] I do not attempt to examine the impact of hidden information on risk sharing in this experiment. See, for example, Ligon (1998) for evidence that information plays an important role in risk sharing arrangements.

amount invested was completely lost the other 50 percent of the time. To encourage truth-telling, respondents were told (before the survey) that one question would be picked later and actually paid out. After the survey ended, a question was randomly picked, a coin was flipped to determine if the amount invested would be multiplied by 2.5 or would be lost, and payouts were made.

The most important source of data, however, were the weekly monitoring surveys. For approximately eight weeks, a trained enumerator separately visited both spouses each week and administered a detailed monitoring survey that included questions on expenditures, income (and income shocks), and labor supply over the previous seven days. The survey also included information on transfers given and received, both to the spouse and to all other individuals. These transfers include cash as well as all other in-kind payments of goods or services (respondents were asked to value these transfers themselves). Thus, these surveys should give a comprehensive summary of all financial transactions for each individual in every week.

The surveys were conducted privately and confidentially, and information was not shared with the spouse.[15] If one of the spouses could not be found on the day of the survey, the enumerator tried again for the next several days. If this individual was eventually traced, the enumerator asked about the same time period that was asked of the spouse (the seven days prior to the scheduled meeting). If the individual could not be traced that week, the spouse's survey was also dropped, so the analysis includes only those weeks in which information is available for both spouses over the same time period.

Due to some problems with some enumerators, particularly toward the beginning of the data collection activities, the database is trimmed of the top and bottom 1 percent of responses for individual and household expenditures, as well as savings outliers. In addition, some surveys were missing information on one of the key dependent variables necessary for the main regressions and were therefore dropped. This leaves 898 visits for 142 couples.

A. Background Statistics

Table 1 presents summary statistics from the background survey, as well as a check that the randomization was implemented properly.[16] First, means are reported in columns 1 (men) and 4 (women). From panel A (which presents demographic information), 84 percent of the men in the sample are bicycle taxi drivers, while the rest are distributed among various other jobs. Fifty-three percent of women report having no job. The sample is predominantly of the Luo tribe, and the remainder is Luhya.[17] The average man in the sample is 30.6 years old and has received 7.7 years of education, while the average woman is younger (24.5) and less educated (with 7.0 years of

---

[15] In most cases, the primary respondent was interviewed at work and the spouse at home.

[16] Table 1 includes information on 136 men and 131 women, out of 142 in the sample. The remainder could not be traced for this survey (as mentioned previously, the background survey was conducted at the end of the project).

[17] The Luo are the most populous tribe in Nyanza Province (making up 53 percent of the Province's population), and the Luhya are the most populous in Western Province (making up 84 percent of the Population). Overall, the Luo make up 12 percent of the Kenyan population and the Luhya 15 percent (Central Bureau of Statistics 2004).

TABLE 1—SUMMARY STATISTICS AND RANDOMIZATION CHECK

	Males			Females		
		Coefficient of regression of dependent variable on average number of shocks received by:[a]			Coefficient of regression of dependent variable on average number of shocks received by:	
	Mean (1)	Respondent (2)	Spouse (3)	Mean (4)	Respondent (5)	Spouse (6)
Panel A. Demographic information[b]						
Occupation:						
Bicycle taxi driver	0.84	0.20 (0.17)	−0.13 (0.18)	— —	— —	— —
Market stall	0.05	0.10 (0.09)	−0.09 (0.10)	0.31	0.24 (0.22)	−0.30 (0.21)
No job	0.02	−0.07 (0.05)	0.09 (0.06)	0.53	0.06 (0.23)	0.41 (0.23)*
Other	0.09	−0.29 (0.13)**	0.11 (0.14)	0.15	−0.29 (0.17)*	−0.11 (0.16)
Luo tribe	0.88	−0.06 (0.15)	0.21 (0.16)	0.86	0.00 (0.17)	0.03 (0.16)
Age	30.57 (8.71)	−0.54 (3.93)	4.53 (4.15)	24.47 (6.83)	1.20 (3.23)	−3.33 (3.14)
Education	7.72 (2.41)	−0.96 (1.10)	1.29 (1.16)	7.02 (2.07)	−1.25 (1.01)	−1.03 (0.96)
Literate (Swahili)	0.85 (0.36)	0.12 (0.16)	−0.04 (0.17)	0.72 (0.45)	−0.03 (0.21)	0.23 (0.21)
Number of children[c]	2.45 (1.75)	−0.49 (0.81)	−0.06 (0.83)	2.45 (1.75)	−0.06 (0.83)	−0.49 (0.81)

(Continued)

schooling). The average couple has 2.5 children. Though not shown in this table, most respondents live in the surrounding rural areas and travel to town for work.

Panel B presents statistics on access to savings and credit. As is common in rural Kenya, access to formal savings and credit is very rare: just 2 percent of men and 1 percent of women have savings accounts. An equal number received a formal loan in the past year. Informal savings and credit are common, however, 63 percent of men and 44 percent of women participate in ROSCAs.[18] Men and women are about equally connected to informal credit (92 percent of men received a loan from a friend or family member in the past year and 89 percent gave a loan, compared to 91 percent and 80 percent of women, respectively). Panel C presents statistics on asset ownership. As expected, men are richer than women. They own 0.79 acres of land, compared to 0.15 acres for women. Similarly, women control a bit less than

[18] That men are more likely than women to participate in ROSCAs is in contrast to, for instance, Anderson and Baland (2002). This is likely because so many women do not have regular jobs in this sample.

TABLE 1—SUMMARY STATISTICS AND RANDOMIZATION CHECK (*Continued*)

	Males			Females		
		Coefficient of regression of dependent variable on average number of shocks received by:[a]			Coefficient of regression of dependent variable on average number of shocks received by:	
	Mean (1)	Respondent (2)	Spouse (3)	Mean (4)	Respondent (5)	Spouse (6)
Panel B. Savings and Credit						
Has formal savings account	0.02 (0.12)	—	—	0.01 (0.09)	—	—
Received formal loan in past year	0.02 (0.15)	—	—	0.01 (0.09)	—	—
Participates in ROSCA	0.63 (0.48)	—	—	0.44 (0.50)	—	—
Amount saved in ROSCAs (for those in ROSCAs)	3,097 (4,733)	—	—	2,035 (3,200)	—	—
Received gift or loan in past year	0.92 (0.27)	—	—	0.91 (0.29)	—	—
Amount received in gifts and loans in past year	2,393 (2,593)	—	—	1,589 (2,083)	—	—
Gave gift or loan in past year	0.89 (0.32)	—	—	0.80 (0.40)	—	—
Amount given in gifts and loans in past year	1,806 (2,944)	—	—	930 (1,428)	—	—
Panel C. Asset ownership						
Acres of land owned	0.79 (1.64)	—	—	0.15 (0.50)	—	—
Value of durable goods owned	2,708 (4,570)	—	—	797 (1,652)	—	—
Value of animals owned	2,914 (15,635)	—	—	145 (838)	—	—
Amount invested (out of 100 Ksh) in risky asset[d]	46.98 (22.17)	—	—	44.57 (21.87)	—	—
Observations	136			131		

Notes: All figures are self-reported means (at the individual level). There are fewer observations than in the monitoring surveys (in which there are 142 couples) because the background survey was administered after the project started and some could not be traced for this survey. All monetary figures in Kenyan shillings. Exchange rate was roughly 70 Kenyan shillings to $1 US during this time period. Columns 1 and 4: standard deviations in parentheses. Columns 2–3 and 5–6: standard errors in parentheses.

[a]These are coefficients of a regression of the dependent variable on the probability that the respondent received the experimental shock over the 8 weeks of the experiment (the total number of shocks divided by the number of weeks the couple could be traced). The probability is used rather than the total number of shocks because some respondents weren't traced in some weeks.

[b]Randomization verification is done only for the variables in Panel A because the other characteristics could potentially have been affected by treatment (since the background survey was collected after the study).

[c]The number of children must be the same within the household. In cases where responses differ, the wife's response is used.

[d]The risky asset paid off 2.5 times the amount invested with probability 50 percent, and 0 with probability 50 percent.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

950 Ksh (US $14) worth of animals and other durable goods, compared to more than 5,600 Ksh (US $80) for men.[19]

Taken together, these results suggest major differences among many dimensions between men and women in this sample. As such, differences in behavior between genders may be attributable to any number of observable or unobservable characteristics. For this reason, the purpose of this paper is not to highlight level differences between genders. Instead, it takes these differences as given and examines how small, transitory income shocks affect household allocations.

B. *Randomization Check*

Table 1 also presents regressions to check that the shocks were random. As will be discussed below, the specification to test for efficiency will utilize household fixed effects. The identifying assumption is thus that, *within* the household, weeks in which a shock is received by a given individual are randomly determined. However, a stronger test is that the total number of shocks received over the entire experiment should be random across households. Table 1 tests this by running the following regression:

$$(5) \quad characteristic_h^i \;=\; \beta_0 \;+\; \beta_1 \frac{\sum_{t=1}^{8} shock_{ht}^m}{\sum_{i=1}^{8} traced_{ht}} \;+\; \beta_2 \frac{\sum_{t=1}^{8} shock_{ht}^f}{\sum_{i=1}^{8} traced_{ht}} \;+\; \varepsilon_h,$$

where the dependent variable is a given individual background characteristic for spouse i in household h. $shock_{ht}^m$ and $shock_{ht}^f$ are indicator variables for the male and female in household h receiving the experimental shock in week t, and $traced_{ht}$ is an indicator for the household being traced for the survey in week t (recall that observations are dropped if either spouse could not be traced so that households only appear if both spouses completed the survey that week). The independent variables are therefore the empirical probabilities that an individual received the shock in a given week. If treatment were truly randomized, the coefficients β_1 and β_2 should be small and statistically insignificant for most variables.

Note that since the background survey was conducted at the end of the project, some of these variables are potentially endogenous. I therefore perform the randomization check only for those that are clearly unaffected by treatment (namely demographic characteristics).[20]

The coefficients are reported in columns 2–3 (men) and 5–6 (women) in Table 1. There are few statistically significant differences across households. Men who received more shocks were less likely to have occupations other than a bicycle taxi driver or market vendor. Women who received more shocks were less likely to have an occupation other than market vendor or housewife. Also, women whose husbands received more shocks were more likely to be housewives. On the whole, how-

---

[19] Durable goods include beds, sofas, tables, chairs, cookers, radios, TVs, mobile and landline phones, clocks, watches, sewing machines, irons, bicycles, and bednets.
[20] However, almost none of the other variables are related to the number of shocks received either (results available upon request).

ever, there appear to be minimal differences even across households and the results appear consistent with random chance.

Finally, given the fixed effects empirical approach, another more direct test is that the shocks should not affect outcomes *before* they are received. I find no effects from these placebo regressions (see online Appendix Table A1), which suggests again that randomization was implemented effectively.

C. *Summary Statistics from the Monitoring Surveys*

Table 2 provides some summary information from the weekly monitoring visits. Panel A presents summary statistics on weekly labor income and hours (not including agriculture). Here, income for those selling produce or other items (who are mostly female), is calculated as the difference in sales and money spent restocking. Of the couples sampled for the survey, men make about 719 Ksh per week (just over US $10) and women about 143 Ksh (about US $2). For men, this income comes primarily from their regular job. For women, income comes largely from informal sources, such as occasional sales of agricultural produce, rather than regular labor income. Even women without regular jobs earn some money: average income for such women is 53 Ksh (US $0.70) per week, compared to 231 Ksh (US $3.30) for women with jobs. In relative terms, then, the experimental income shocks are relatively large, especially for women: the $2 shock is equivalent to roughly 1.5 days' income for men and over a week's income for women. To put this in terms of a developed country equivalent, for men, the shock is equivalent to roughly $200 for a worker making $50,000 per year. For women, the shock is much larger, equivalent to roughly $950.

Though consumption was recorded in the surveys, expenditures will be used in the main specifications, for several reasons. First, to reduce the length of the monitoring survey, the consumption questions were asked only at the household level, so that I do not have specific measures of individual consumption shares, and thus they would have to be imputed. Second, the main test of efficiency is the consumption of private goods (alcohol, cigarettes, soda, clothing and shoes, hairstyling, entertainment, newspapers, own meals in restaurants, transportation and various other items), and expenditures on these items are equal to (the monetary value of) consumption in most cases. Any allocation of such items to others would have been recorded as in-kind transfers and, while some items could in principle be saved for future use or be consumed over multiple weeks (such as clothing, for example), most categories are consumed immediately (such as food in restaurants, alcohol, soda, and cigarettes).

Panel B presents the expenditure data. The first row of panel B show total expenditures: men spent about 820 Ksh a week, compared to 369 Ksh for women. Total household expenditures are therefore around $2.42 per day, indicating how poor these households are. The next few rows break expenditures into various broad categories: shared food, spending on children,[21] medical expenses, other shared expenses,[22] and total private expenditures. Though shared food and other shared expenses are the biggest categories, both men and women spend substantial sums on private items:

---

[21] This includes clothing, school fees, and school supplies.
[22] Other shared expenditures include cleaning supplies, rent, water, household bills, and other related expenses.

TABLE 2—SUMMARY STATISTICS FROM MONITORING SURVEYS

	Male (1)	Female (2)
Panel A. Income		
Total labor income	718.64	143.01
	(746.15)	(573.68)
Total hours worked	55.35	16.47
	(65.42)	(33.04)
Panel B. Expenditures		
Total expenditures	820.05	369.21
	(525.34)	(397.01)
Shared food	380.51	192.67
	(274.09)	(203.02)
Children	18.77	16.61
	(71.10)	(54.54)
Medical	42.59	25.34
	(103.42)	(90.75)
Other shared	126.72	59.92
	(228.13)	(119.09)
Transportation	107.98	34.75
	(121.14)	(113.29)
Total private	143.71	39.92
	(161.32)	(92.32)
Private categories		
Clothing	21.41	21.87
	(85.65)	(77.54)
Meals in restaurants	71.75	5.33
	(76.08)	(24.28)
Alcohol, soda, cigarettes	28.04	4.39
	(51.52)	(17.97)
Other private expenditures	22.49	8.34
	(74.95)	(25.11)
Panel C. Transfers and savings		
(Net) transfers to spouse	59.46	−59.46
	(147.44)	(147.44)
(Net) transfers outside HH	11.03	6.28
	(371.85)	(326.65)
Savings	−91.05	−92.51
	(774.02)	(602.51)
Observations	898	898
Number of IDs	142	142

Notes: In panel B, "Total private" expenditures include the subcategories listed in the bottom of the panel. The "other private expenditures" category includes hairstyling, entertainment, newspapers, transportation, mobile phone airtime, and similar items. Shared food includes all food consumed jointly at home. Spending on children includes school fees, school supplies, and clothing. Other shared expenditures includes cleaning supplies, rent, water, household bills, and related expenses. In panel C, transfers are defined as positive for outflows and negative for inflows and include cash and in-kind transfers. Savings is imputed as the sum of total income (including the experimental shocks) minus total expenditures. Standard deviations in parentheses.

private expenses make up about 18 percent of total expenditures for men and 11 percent for women.

The bottom part of the panel breaks down private expenditures into their primary components.[23] Men spend much more on meals in restaurants (usually lunch in town when they are working) and on alcohol, soda, and cigarettes. However, women also spend relatively sizeable amounts (given their income) on clothing for themselves and on other private items.

Panel C presents summary statistics on transfers (which are defined as positive for outflows and negative for inflows, and which include cash and in-kind transfers) between spouses and with individuals outside of the household, and on imputed savings (estimated as the difference between total cash flows and total expenditures). In total, women receive an average of 59 Ksh per week from their husbands, the vast majority of which are gifts rather than loans. Both men and women regularly send and receive transfers, and overall savings levels are quite low (average savings are actually negative, which might reflect some underreporting of income as is common in surveys of this type). However, so long as underreporting is constant across weeks, this type of bias should difference out over the panel.

VI. Experimental Results

A. *Main Specification*

The results from estimating the reduced form specification (3) by fixed effects are presented in panels A (men) and B (women) of Table 3. For ease of interpretation, the shock is measured as the number of shillings received that week (either 0 or 150). Thus, the coefficients in the Table can be interpreted as a propensity to consume out of a shilling's worth of shock.[24]

From panel A, the only statistically significant increase in expenditures for men are private expenditures (which are significant at 1 percent). The estimated propensity to spend on private items out of own income is 0.169. Interestingly, private expenditures do not change in weeks in which the wife receives the shock (the sign is actually negative). Consequently, the null hypothesis for efficiency (that these marginal propensities are equal) can be rejected at the 5 percent level. Though the other expenditure categories are less easily interpretable as a test of efficiency (since they are shared), there is little evidence of differences in expenditure responses to own and spouse shocks.[25]

By contrast, for women, private expenditures do not respond to the shocks (received either by herself or her husband). Private expenditures are actually slightly lower in such

---

[23] "Other" private expenditures include hairstyling, entertainment, newspapers, transportation, mobile phone airtime, and related items.

[24] The general pattern of the results look similar when conditioning on labor income, or when including an interaction between the two shocks.

[25] Another possible concern with the estimation strategy is that expenditures are censored from below at zero. To examine how important this issue is, the last row of each panel reports the percentage of observations for which the dependent variable is equal to zero in that week. Men have zero private expenditures on 12 percent of weeks, while women have zero expenditures on 61 percent of weeks. Thus, it seems unlikely that the relatively small number of zeroes for men is a major concern. While this could in principle be addressed directly by running a fixed effects Tobit, I do not do that here as there is evidence that the fixed effects Tobit estimator yields biased estimates of estimator variance (see Greene 2004).

TABLE 3—EXPERIMENTAL SHOCKS AND EXPENDITURES

	Expenditures						
	Total (1)	Private (2)	Shared food (3)	Medical (4)	Children (5)	Other shared (6)	Transport (7)
Panel A. Men							
Shillings received in experimental shock by respondent	0.190 (0.194)	0.169 (0.064)***	−0.025 (0.089)	0.048 (0.041)	−0.012 (0.032)	−0.096 (0.102)	0.102 (0.068)
Shillings received in experimental shock by spouse	−0.163 (0.192)	−0.027 (0.069)	−0.016 (0.087)	0.057 (0.045)	−0.019 (0.030)	−0.086 (0.111)	−0.069 (0.060)
Observations	898	898	898	898	898	898	898
Number of households	142	142	142	142	142	142	142
p-value for *F*-test of equality	0.21	0.05**	0.93	0.84	0.88	0.95	0.09*
Mean of dependent variable (Ksh)[a]	889.32	135.66	413.77	56.95	24.09	144.77	114.55
SD of dependent variable (Ksh)	557.30	122.24	298.74	143.25	84.40	250.88	106.76
Proportion of weeks dependent variable = 0	0.00	0.12	0.03	0.52	0.86	0.12	0.18
Panel B. Women							
Shillings received in experimental shock by respondent	0.180 (0.148)	−0.020 (0.042)	0.056 (0.067)	0.079 (0.041)*	0.032 (0.026)	0.041 (0.059)	−0.007 (0.047)
Shillings received in experimental shock by spouse	−0.058 (0.123)	−0.026 (0.039)	−0.051 (0.064)	0.015 (0.034)	−0.025 (0.024)	0.050 (0.041)	−0.021 (0.039)
Observations	898	898	898	898	898	898	898
Number of households	142	142	142	142	142	142	142
p-value for *F*-test of equality	0.14	0.91	0.23	0.07*	0.1*	0.88	0.77
Mean of dependent variable (Ksh)	428.51	47.28	227.98	28.43	18.25	68.51	38.07
SD of dependent variable (Ksh)	482.65	123.77	262.65	94.87	65.80	119.21	101.60
Proportion of weeks dependent variable = 0	0.03	0.61	0.08	0.64	0.84	0.28	0.72

Notes: All regressions are estimated by fixed effects with controls for the week of the interview. The experimental shock is measured in terms of Kenyan shillings (150 Ksh when the shock is received and 0 Ksh otherwise). See Table 2 for explanations of the various expenditure categories. Standard errors clustered at the household level in parentheses.

[a] The mean and standard deviation reported here is for weeks when neither spouse receives the shock.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

weeks, though statistically insignificant. Women do spend more on medical expenses when they receive a shock (significant only at the 10 percent level), but the effect is weak. There is also no discernible effect on other categories that have been associated with female preferences in other studies (for instance, spending on children).

Table 4 examines transfers, labor supply, and savings. Columns 1 and 2 show transfers to the spouse (these results are symmetric across spouses by definition, as every shilling sent by one spouse is received by the other). Men transfer 7.7 percent of the shock to their wives (which is insignificant), while women transfer 16.3 percent to their husbands (significant at 1 percent). Both men and women also appear to transfer some outside the household in such weeks (though the results are statistically insignificant). Columns 3 and 4 show that there is no discernible effect on weekly labor supply.

Finally, column 5 shows savings in bank accounts or in ROSCAs, while column 6 shows overall savings.[26] As can be seen, savings in banks and ROSCAs do not much

[26] Total savings is imputed as income minus expenditures.

TABLE 4—EXPERIMENTAL SHOCKS, TRANSFERS, LABOR SUPPLY, AND SAVINGS

	Net transfers to:		Labor supply		Savings[b]	
	Spouse (1)	Outside household (2)	Hours (3)	Labor income (4)	Bank/ ROSCA savings (5)	Total savings (6)
Panel A. Men						
Shillings received in experimental shock by respondent	0.077 (0.065)	0.090 (0.202)	0.018 (0.017)	0.139 (0.366)	0.020 (0.159)	0.782 (0.393)**
Shillings received in experimental shock by spouse	−0.163 (0.060)***	−0.133 (0.157)	−0.036 (0.035)	−0.145 (0.312)	−0.244 (0.154)	0.314 (0.319)
Observations	898	898	898	898	898	898
Number of households	142	142	142	142	142	142
p-value for *F*-test of equality	0.01***	0.31	0.27	0.48	0.35	0.31
Mean of dependent variable (Ksh)[a]	76.78	2.81	52.18	698.56	127.20	−270.34
SD of dep. var. (Ksh)	159.89	436.18	24.14	852.24	222.35	885.11
Proportion of weeks dependent variable = 0	0.62	0.21	0.08	0.07	0.50	0.00
Panel B. Women						
Shillings received in experimental shock by respondent	0.163 (0.060)***	0.050 (0.190)	−0.031 (0.020)	−0.020 (0.185)	0.082 (0.088)	0.586 (0.239)**
Shillings received in experimental shock by spouse	−0.077 (0.065)	−0.010 (0.160)	0.009 (0.011)	0.031 (0.195)	−0.154 (0.099)	0.175 (0.234)
Observations	898	898	898	898	898	898
Number of households	142	142	142	142	142	142
p-value for *F*-test of equality	0.01***	0.63	0.14	0.86	0.06*	0.17
Mean of dependent variable (Ksh)	−76.78	−11.15	16.77	165.33	116.64	−175.25
SD of dependent variable (Ksh)	159.89	549.09	24.88	604.19	376.07	698.14
Proportion of weeks dependent variable = 0	0.62	0.20	0.60	0.61	0.60	0.01

Notes: All regressions are estimated by fixed effects with controls for the week of the interview. The experimental shock is measured in terms of Kenyan shillings (150 Ksh when the shock is received and 0 Ksh otherwise). See Table 2 for explanations of the various expenditure categories. Standard errors clustered at the household level in parentheses.
  [a] The mean and standard deviation reported here is for weeks when neither spouse receives the shock.
  [b] Savings is imputed as the sum of total income (including the experimental shocks) minus total expenditures. Bank/ROSCA savings are withdrawals/payouts minus deposits/contributions.
     ***Significant at the 1 percent level.
      **Significant at the 5 percent level.
       *Significant at the 10 percent level.

respond to the shocks (as might be expected given that so few people have bank accounts and that ROSCA contribution cycles cannot typically be altered once a ROSCA is formed). However, the overall propensity to save in column 6 is quite high. In total, men save 78.2 percent of the shock and women 58.6 percent. This suggests that money is saved informally, in cash, at home.[27] Overall, given the standard errors, I cannot reject the Permanent Income Hypothesis for either spouse (that the propensity to save is equal to one). However, since I do observe statistically significant increases in consumption for men, the failure to reject is due to imprecision in the estimates.

---

[27] The dataset does not include a specific measure of savings at home or in cash. This is because people are reticent to report this information in a survey.

To summarize the main results, for every shilling men receive, they increase labor income (insignificantly) by 0.139 Ksh, spend 0.190 Ksh on expenditures, transfer 0.077 Ksh to their wife and 0.090 Ksh outside the household, and save the remaining 0.782 Ksh. Women decrease labor income by 0.02 Ksh, increase expenditures by 0.18 Ksh, transfer 0.163 Ksh to their husband and 0.05 Ksh outside the household, and save the remaining 0.586 Ksh.

B. *Lagged Shocks*

Since these regressions include only current outcomes on the current realization of shocks, it is possible that they do not fully capture the dynamics of household spending (for example, it is possible that people save the shocks over a week and spend the shocks later on). To examine this, I run specifications that also include measures for whether the respondent and his spouse received a shock the previous week. The cost of doing this is that I can only include observations that were tracked in successive weeks. This reduces the total number of observations to 618 (from 898) and the number of households from 142 to 140.[28]

The results are presented in online Appendix Tables A2 (men) and A3 (women). For men, the current week increase in private expenditures persists. The propensity to spend is 0.215 out of own current shock income (significant at 1 percent) and 0.067 out of the wife's. Though this difference is no longer statistically significant due to the decreased sample size, the pattern is very similar of the main results in Table 3. Again, there are few statistically significant changes in other outcomes (though there is a small decrease in labor hours that is significant at 10 percent). None of the lagged shocks on own income are significant for men.

Online Appendix Table A3 presents results for women. Again, there is no discernible effect on private expenditures. Though women increase total expenditures, this is mostly in shared categories. Labor income also appears to go down somewhat for women after the receipt of shocks, though the effect is imprecisely estimated. This could be evidence, however, that women treat own income shocks differently than spouse's income shocks in determining labor supply (which would itself be a rejection of efficiency). However, the effect appears to be too weak to make definitive conclusions.[29]

In sum, the overall patterns from the lagged shocks are generally supportive of the main results.

---

[28] Though all households were tracked for a minimum of four weeks, some were not found in consecutive weeks.
[29] Another specification to deal with the possibility that money is not spent immediately is to compare total expenditure levels over the entire experiment on the total number of shocks received. The general results look similar from such a specification, but the power is low since there is only one observation per household. Thus, given that online Appendix Tables A2 and A3 suggest that most private spending is immediate, I do not report these results here.

VII. External Validity and Alternative Hypotheses

A. *Behavior Outside of Experiment*

While the experimental approach adopted in this paper provides a clean test of intra-household efficiency within the experiment, a drawback is that the environment is somewhat stylized. In particular, the shocks are always positive and the experimental payout is akin to a small "windfall" separate from an individual's normal income source.[30] While this cannot be an issue if preferences are standard and people treat all sources of income similarly, they could be relevant if windfall income is treated differently. I attempt to address this possibility in this section.

Ideally, there would be an instrumental variable that would affect labor income but not preferences or bargaining power (rainfall, for instance). If exogenous labor income changes could be identified with this instrument, it would be possible to causally test for efficiency. Unfortunately, I do not have such an instrument (those that are potentially available, such as sickness or other shocks are either not strong enough to predict income or may directly affect preferences for private expenditures).

Thus, I have to rely directly on week to week changes in labor income. To do this, I run the following regression:

$$(6) \qquad y_{ht}^i = \gamma L_{ht}^i + \delta L_{ht}^j + \nu_h + \mu_t + \varepsilon_{ht}^i,$$

where L indexes labor income. I also control for the experimental shocks in this specification. Identification requires that weekly labor income for a given household is uncorrelated with preferences.[31] As this assumption is difficult to verify with this data, the results should be taken with some care.

With that caveat in mind, the results are very supportive of the main experimental findings. The results are presented in Tables 5 (expenditures) and 6 (transfers and savings). As the standard errors in these regressions are smaller than in the experimental section (given that there is more variation in income than in the experimental shocks), tighter inference is possible. Most notably, both men and women spend significantly more on private expenditures when they earn more labor income. While the propensities are not very large (0.025 for men and 0.023 for women), efficiency is rejected in both cases (at the 5 percent level for men and the 10 percent level for women). Again, the majority of these income fluctuations are saved, which suggests that they are indeed considered transitory shocks.

While these results are speculative given the possible endogeneity of weekly labor income, they do at least suggest that the experimental findings were not necessarily specific to the experiment.

---

[30] A related issue is that people may treat gains differently than losses, for example, because they are loss averse (i.e., Kahneman and Tversky 1979). If so, they will tend to be risk averse over gains and risk loving over losses. As the experiment involves only gains, loss averse individuals should have been more likely to insure each other than they would have been for losses. Thus, loss aversion seems unlikely to explain the results.

[31] Results look similar when controlling for hours worked, and when controlling for other shocks (such as sickness).

TABLE 5—LABOR INCOME FLUCTUATIONS AND EXPENDITURES

	Expenditures						
	Total (1)	Private (2)	Shared food (3)	Medical (4)	Children (5)	Other shared (6)	Transport (7)
Panel A. Men							
Respondent labor income	0.184 (0.039)***	0.025 (0.013)**	0.041 (0.016)**	0.012 (0.017)	0.002 (0.003)	0.074 (0.046)	0.029 (0.011)***
Spouse labor income	−0.005 (0.031)	−0.004 (0.005)	−0.014 (0.008)*	0.008 (0.015)	−0.001 (0.003)	0.013 (0.018)	−0.007 (0.006)
Observations	898	898	898	898	898	898	898
Number of households	142	142	142	142	142	142	142
p-value for F-test of equality	0.01***	0.03**	0.01***	0.87	0.59	0.21	0.01***
Mean of dependent variable (Ksh)[a]	820.05	143.71	380.51	42.59	18.77	126.72	107.98
SD of dependent variable (Ksh)	525.34	161.32	274.09	103.42	71.10	228.13	121.14
Proportion of weeks dependent variable = 0	0.00	0.12	0.03	0.52	0.86	0.12	0.18
Panel B. Women							
Respondent labor income	0.124 (0.043)***	0.023 (0.006)***	0.055 (0.030)*	0.007 (0.011)	0.009 (0.006)	0.027 (0.009)***	0.004 (0.006)
Spouse labor income	0.037 (0.013)***	0.009 (0.004)**	0.011 (0.007)	0.002 (0.003)	0.001 (0.002)	0.013 (0.006)**	0.000 (0.004)
Observations	898	898	898	898	898	898	898
Number of Households	142	142	142	142	142	142	142
p-value for f-test of equality	0.07*	0.05**	0.19	0.65	0.17	0.15	0.54
Mean of dependent variable (Ksh)	369.21	39.92	192.67	25.34	16.61	59.92	34.75
SD of dependent variable (Ksh)	397.01	92.32	203.02	90.75	54.54	119.09	113.29
Proportion of weeks dependent variable = 0	0.03	0.61	0.08	0.64	0.84	0.28	0.72

Notes: All regressions are estimated by fixed effects with controls for the week of the interview and the experimental shocks. See Table 2 for explanations of the various expenditure categories. Standard errors clustered at the household level in parentheses.

[a]The mean and standard deviation reported here is over all weeks.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

B. *Alternative Hypothesis: Differences in Risk Preferences*

Recent work has shown that men and women have different preferences for risk. In particular, women tend to be more risk averse than men (Croson and Gneezy 2009). Such differences are important for the structure of risk-sharing arrangements. In particular, the less risk averse individual could insure the more risk averse individual by accepting more consumption variance in exchange for a higher average level of consumption. Mazzocco and Saini (2012) find evidence for such heterogeneity across households in the ICRISAT dataset used by Townsend (1994), and show that accounting for this makes an important difference in empirical inferences.

I address this by making use of the experimentally elicited risk preferences in which individuals were asked how much of 50 or 100 Ksh they wanted to invest in a risky asset that would pay out 2.5 the amount invested half the time, but nothing the other half of the time. I then regress this measure on an indicator for the gender of the respondent. To be as transparent as possible, I do not include any other controls.

TABLE 6—LABOR INCOME FLUCTUATIONS, TRANSFERS, AND SAVINGS

	Net transfers to:		Savings[b]	
	Spouse (1)	Outside household (2)	Bank/ROSCA savings (3)	Total savings (4)
Panel A. Men				
Respondent labor income	−0.006	0.089	−0.043	0.733
	(0.010)	(0.041)**	(0.065)	(0.075)***
Spouse labor income	0.008	−0.011	−0.001	0.008
	(0.014)	(0.020)	(0.023)	(0.039)
Observations	898	898	898	898
Number of households	142	142	142	142
p-value for *F*-test of equality	0.54	0.02**	0.57	0.01***
Mean of dependent variable (Ksh)[a]	59.46	11.03	109.90	−91.05
SD of dependent variable (Ksh)	147.44	371.85	355.49	774.02
Proportion of weeks dependent variable = 0	0.62	0.21	0.50	0.00
Panel B. Women				
Respondent labor income	−0.008	0.081	−0.011	0.803
	(0.014)	(0.023)***	(0.016)	(0.049)***
Spouse labor income	0.006	−0.024	0.013	−0.018
	(0.010)	(0.014)*	(0.044)	(0.019)
Observations	898	898	898	898
Number of households	142	142	142	142
p-value for *F*-test of equality	0.54	0.01***	0.60	0.001***
Mean of dependent variable (Ksh)	−59.46	6.28	92.36	−92.51
SD of dependent variable (Ksh)	147.44	326.65	307.46	602.51
Proportion of weeks dependent variable = 0	0.62	0.20	0.60	0.01

Notes: All regressions are estimated by fixed effects with controls for the week of the interview and the experimental shocks. See Table 2 for explanations of the various expenditure categories. Standard errors clustered at the household level in parentheses.

[a] The mean and standard deviation reported here is over all weeks.

[b] Savings is imputed as the sum of total income (including the experimental shocks) minus total expenditures. Bank/ROSCA savings are withdrawals/payouts minus deposits/contributions.

*** Significant at the 1 percent level.

** Significant at the 5 percent level.

* Significant at the 10 percent level.

Results are presented in online Appendix Table A4 (note that I have information here on only 129 couples). Women invest 20.4 Ksh and 44.6 Ksh of the 50 Ksh and 100 Ksh amounts, respectively, in the asset (the constant in this regression). Men invest a bit more (Ksh 2.1 and 2.4, respectively), but these differences are insignificant and very small. For example, the standard deviation of the amount invested out of 100 Ksh is 22, so the difference between genders is equivalent to only 0.1 of a standard deviation. I further check that these differences are not driving the results by re-running equation (3) for spouses with similar risk preferences (those with no more than a 10 or 20 Ksh difference in the amount invested).[32] While couples with similar risk preferences are a selected subsample, and while restricting attention to

[32] In total, 43.4 percent of couples have no more than a 10 Ksh difference in the amount invested, and 62.8 percent have no more than a 20 Ksh difference.

a subset of the sample increases the standard errors, the main findings remain, suggesting that differential risk preferences are not the explanation.

VIII. Conclusion

Any test of intra-household risk coping must identify exogenous shocks that affect relative incomes but do not affect bargaining parameters or preferences. The contribution of this paper is to provide random shocks in a controlled experiment among married couples in Western Kenya. The experimental shocks are well suited for testing efficiency—they are randomly determined, transitory, idiosyncratic, and small relative to lifetime income. They are also perfectly observable (because they were announced to both spouses), so that information asymmetries are not relevant. Thus, the experiment represents a particularly direct and easily interpretable test of Pareto efficiency.

The results suggest that risk sharing is incomplete and that efficiency is not achieved. More speculative evidence further suggests that even outside of the experiment, these couples do not achieve efficiency over weekly labor fluctuations. Despite the prevalence of income shocks in this part of Kenya, it appears that spouses do not fully insure each other.

Understanding the effectiveness of intra-household risk coping is important because numerous other studies have shown that both inter-temporal and inter-household risk mechanisms are only partially effective (including several studies in this part of Kenya).[33] If potentially insurable individual risk is not insured even within the household, then it strongly suggests that the provision of more formal risk coping devices (at the individual level) could have large effects. For example, Dupas and Robinson (2012) show how providing even the most basic savings products decreases vulnerability to health shocks. Similar interventions seem well worth exploring given the incompleteness of informal risk sharing, both within and across households.

REFERENCES

Akresh, Richard. 2008. "(In)Efficiency in Intra-Household Allocations." http://netfiles.uiuc.edu/akresh/www/Research/Akresh_IntraHH.pdf.

Anderson, Siwan, and Jean-Marie Baland. 2002. "The Economics of Roscas and Intra-Household Resource Allocation." *Quarterly Journal of Economics* 117 (3): 963–95.

Angelucci, Manuela, Gacomo De Giorgi, and Imran Rasul. 2012. "Resource Pooling within Family Networks: Insurance and Investment." Stanford University. http://www.stanford.edu/%7Edegiorgi/progresa_insurance.pdf.

Ashraf, Nava. 2009. "Spousal Control and Intra-Household Decision Making: An Experimental Study in the Philippines." *American Economic Review* 99 (4): 1245–77.

Attanasio, Orazio, Abigail Barr, Juan Camilo Cardenas, Garance Genicot, and Costas Meghir. 2012. "Risk Pooling, Risk Preferences, and Social Networks." *American Economic Journal: Applied Economics* 4 (2): 134–67.

Banerjee, Abhijit V., and Esther Duflo. 2007. "The Economic Lives of the Poor." *Journal of Economic Perspectives* 21 (1): 141–67.

Barr, Abigail. 2003. "Risk Pooling, Commitment, and Information: An Experimental Test of Two Fundamental Assumptions." http://www.econ.yale.edu/conference/neudc03/papers/7c-barr.pdf.

---

[33] See, for example, Robinson and Yeh (2011) and Dupas and Robinson (2011).

Blundell, Richard, Pierre-André Chiappori, and Costas Meghir. 2005. "Collective Labor Supply with Children." *Journal of Political Economy* 113 (6): 1277–1306.

Bobonis, Gustavo J. 2009. "Is the Allocation of Resources within the Household Efficient? New Evidence from a Randomized Experiment." *Journal of Political Economy* 117 (3): 453–503.

Browning, Martin, Francois Bourguignon, Pierre-André Chiappori, and Valérie Lechene. 1994. "Income and Outcomes: A Structural Model of Intra-Household Allocation." *Journal of Political Economy* 102 (6): 1067–96.

Browning, Martin, and Pierre-André Chiappori. 1998. "Efficient Intra-Household Allocations: A General Characterization and Empirical Tests." *Econometrica* 66 (6): 1241–78.

Bourguignon, François, Martin Browning, and Pierre-André Chiappori. 2009. "Efficient Intra-Household Allocations and Distribution Factors: Implications and Identification." *Review of Economic Studies* 76 (2): 503–28.

Bourguignon, François, Martin Browning, Pierre-André Chiappori, and Valérie Lechène. 1993. "Intra-Household Allocation of Consumption: A Model and Some Evidence from French Data." *Annales d'Economie et de Statistique* 7 (29): 137–56.

Central Bureau of Statistics. 2004. *Kenya Demographic and Health Survey 2003*. Central Bureau of Statistics (CBS). Calverton, July.

Chandrasekhar, Arun G., Cynthia Kinnan, and Horacio Larreguy. 2010. "Informal Insurance, Social Ties, and Financial Development: Evidence from a Lab Experiment in the Field." www.aae.wisc.edu/mwiedc/papers/2011/Kinnan_Cynthia.pdf.

Charness, Gary, and Garance Genicot. 2009. "Informal Risk Sharing in an Infinite-Horizon Experiment." *Economic Journal* 119 (537): 796–825.

Chiappori, Pierre-André. 1992. "Collective Labor Supply and Welfare." *Journal of Political Economy* 100 (3): 437–67.

Chiappori, Pierre-André, Bernard Fortin, and Guy Lacroix. 2002. "Marriage Market, Divorce Legislation, and Household Labor Supply." *Journal of Political Economy* 110 (1): 37–72.

Coate, Stephen, and Martin Ravallion. 1993. "Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements." *Journal of Development Economics* 40 (1): 1–24.

Croson, Rachel, and Uri Gneezy. 2009. "Gender Differences in Preferences." *Journal of Economic Literature* 47 (2): 448–74.

Dercon, Stefan, and Pamela Krishnan. 2000. "In Sickness and in Health: Risk Sharing within Households in Rural Ethiopia." *Journal of Political Economy* 108 (4): 688–727.

Doss, Cheryl. 2001. "Is Risk Fully Pooled within the Household? Evidence from Ghana." *Economic Development and Cultural Change* 50 (1): 101–30.

Dubois, Pierre, and Ethan Ligon. 2009. "Nutrition and Risk Sharing within the Household." University of California at Berkeley Department of Agricultural and Resource Economics (CUDARE) Working Paper 1096.

Duflo, Esther. 2003. "Grandmothers and Granddaughters: Old-Age Pensions and Intra-Household Allocation in South Africa." *World Bank Economic Review* 17 (1): 1–25.

Duflo, Esther, and Christopher Udry. 2004. "Intra-Household Resource Allocation in Cote D'Ivoire: Social Norms, Separate Accounts, and Consumption Choices." National Bureau of Economic Research (NBER) Working Paper 10498.

Dupas, Pascaline, and Jonathan Robinson. 2011. "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya." National Bureau of Economic Research (NBER) Working Paper 14693.

Dupas, Pascaline, and Jonathan Robinson. 2012. "Why Don't the Poor Save More? Evidence from Health Savings Experiments." National Bureau of Economic Research (NBER) Working Paper 17255.

Fafchamps, Marcel, and Susan Lund. 2003. "Risk Sharing Networks in the Rural Philippines." *Journal of Development Economics* 71 (2): 261–87.

Foster, Andrew D., and Mark R. Rosenzweig. 2001. "Imperfect Commitment, Altruism, and the Family: Evidence from Transfer Behavior in Low-Income Rural Areas." *Review of Economics and Statistics* 83 (3): 389–407.

Goldstein, Markus. 2004. "Intrahousehold Efficiency and Individual Insurance in Ghana." London School of Economics and Political Science Discussion Paper 38.

Greene, William. 2004. "Fixed Effects and Bias Due to the Incidental Parameters Problem in the Tobit Model." *Econometric Reviews* 23 (2): 125–47.

Haddad, Lawrence, and John Hoddinott. 1995. "Does Female Income Share Influence Household Expenditures? Evidence from Cote d'Ivoire." *Oxford Bulletin of Economics and Statistics* 57 (1): 77–96.

Iversen, Vegard, Cecile Jackson, Bereket Kebede, Alistair Munro, and Arjan Verschoor. 2006. "What's Love Got to Do With It? An Experimental Test of Household Models in East Uganda." Center for the Study of African Economics (CSAE) Working Paper 14132.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47 (2): 263–92.

Ligon, Ethan. 1998. "Risk Sharing and Information in Village Economics." *Review of Economic Studies* 65 (4): 847–64.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2002. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Review of Economic Studies* 69 (1): 209–44.

Lundberg, Shelly J., Robert A. Pollak, and Terence J. Wales. 1997. "Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit." *Journal of Human Resources* 32 (3): 463–80.

Mazzocco, Maurizio, and Shiv Saini. 2012. "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences." *American Economic Review* 102 (1): 428–68.

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

Rangel, Marcos, and Duncan Thomas. 2005. "Out of West Africa: Evidence on the Efficient Allocation of Resources within Farm Households." University of Chicago Harris School Working Paper 05.15.

Ravallion, Martin, and Martin Dearden. 1988. "Social Security in a 'Moral Economy': An Empirical Analysis for Java." *Review of Economics and Statistics* 70 (1): 36–44.

Robinson, Jonathan, and Ethan Yeh. 2011. "Transactional Sex as a Response to Risk in Western Kenya." *American Economic Journal: Applied Economics* 3 (1): 35–64.

Robinson, Jonathan. 2012. "Limited Insurance Within the Household. Evidence from a Field Experiment in Kenya: Dataset." *American Economic Journal: Applied Economics.* http://dx.doi.org/10.1257/app.4.4.140.

Schaner, Simone. 2012. "Intra-Household Preference Heterogeneity, Commitment, and Strategic Savings." Dartmouth College Working Paper.

Thomas, Duncan. 1990. "Intra-Household Resource Allocation: An Inferential Approach." *Journal of Human Resources* 25 (4): 635–64.

Thomas, Duncan, and Chien-Liang Chen. 1994. "Income Shares and Shares of Income: Empirical Tests of Models of Household Resource Allocations." Research and Development (RAND) Labor and Population Program Working Paper 94-08.

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica* 62 (3): 539–91.

Udry, Christopher. 1994. "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria." *Review of Economic Studies* 61 (3): 495–526.

Udry, Christopher. 1996. "Gender, Agricultural Production, and the Theory of the Household." *Journal of Political Economy* 104 (5): 1010–46.

Wahhaj, Zaki. 2007. "A Theory of Household Bargaining in the Presence of Limited Commitment and Public Goods." http://www.kent.ac.uk/economics/staff/staffhost/zaki-wahhaj/Public_Goods.pdf.