

The Effect of Social Pressure on Expenditures in Malawi*

Jessica Goldberg[†]

July 28, 2016

Abstract

I vary the observability of a windfall payment to 294 members of agricultural clubs in rural Malawi in order to study the effect of social pressure on timing of expenditures. While other studies have documented that social pressure affects the quantity of income and consumption, I focus on timing because spending money quickly may be a strategy for reducing obligatory transfers. Such a shift in timing is welfare reducing if it reduces consumption smoothing or the ability to search for better prices, or leads to greater spending on temptation goods. Respondents who receive money in the presence of their agricultural club spend 30 percent more in the week immediately following the payment than those who receive equivalent transfers in private settings. There is no overall change in the composition of spending, but some evidence that social pressure to share windfall income has a larger effect on poorer households.

*An earlier draft of this paper was circulated under the title “The Lesser of Two Evils: The Roles of Social Pressure and Impatience in Consumption Decisions.” The research was supported by grants from the Rackham Graduate School at the University of Michigan. It was approved by the IRB committee at the University of Michigan. I thank Brian Jacob, Pamela Jakiela, David Lam, Jeff Smith, and Dean Yang for their helpful comments, and Santhosh Srinivasan and Lutamayo Mwamlima for their assistance with field work. All errors and omissions are my own.

[†]University of Maryland, goldberg@econ.umd.edu

1 Introduction

Social pressure to share income has been documented in developing and developed-country settings. Such pressure can be embodied in mutually-beneficial informal insurance networks (Townsend 1994) or generate unilateral contributions that reduce the welfare of net donors (Platteau 2000). In either case (though especially in the latter), there are incentives to hide income that can distort consumption (Kinnan 2014), investment (Jakiela & Ozier 2016), and borrowing (Baland, Guirkinger & Mali 2007). These distortions can affect timing as well as quantity of consumption and investment. Kinnan (2014) notes that consumption time-paths that would be suboptimal with perfect information can be used by households to hide income and therefore increase private utility when information is imperfect. Social anthropologists document similar a phenomena; for example, Maranz (2001) writes that pressure to share income means that individuals “often made wasteful or ill-considered expenditures just to keep friends from borrowing.”

While a growing literature recognizes the importance of social pressure, there are few direct estimates of how it affects individuals’ expenditures. Well-identified estimates of the effects of social pressure are difficult to obtain, because it is a cause and consequence of complex and often unobservable relationships between individuals and institutions. Baland, Guirkinger & Mali (2007) rely on observational data about borrowing and savings at Cameroonian credit cooperatives and surveys asking members about reasons for simultaneous borrowing and saving. Kinnan (2014) documents auto-correlation in observational consumption data that is consistent with the predictions of a model of imperfect information and hidden income.

Two studies experimentally vary observability of financial decision-making in order to study behavior under different information conditions. Ashraf (2009) studies intra-household bargaining by varying whether Filipino spouses’ allocations of one day’s wages are observed by their partners. Jakiela & Ozier (2016) offer Kenyan participants the opportunity to choose between stylized investment opportunities with different returns, when their decisions are either secret or announced to an audience including members of their extended family. Information matters in both contexts. Filipino husbands allocate more money to their private accounts when their decisions are secret, and Kenyan women forego profitable investments in order to hide returns from their extended families,

To my knowledge, Boltz, Marazyan & Villar (2015) is the only other study that combines experimental variation in the observability of income with data about spending or consumption outside the lab. That study, conducted in urban Senegal, measures willingness-to-pay (WTP) to hide income in the lab and sharing of income outside the lab. Wealthier men and women with higher positions in their extended family have higher WTP for income hiding.

When given the opportunity to hide some income, personal expenditures rise, and, for those with positive WTP for privacy, transfers to kin fall.

I vary whether members of Malawian agricultural clubs receive windfall income in public or private settings, and study the subsequent use of prize money. While the composition of spending is similar for recipients whose windfall was observed and those who received money secretly, the timing of spending was not: public recipients spent 30 percent more of their prize money in the week immediately following the transfer than private recipients. I describe the experiment in Section 2, the data in Section 3, and discuss the results in detail in section 4. Section 5 concludes.

2 Experimental Design

Individuals are exposed to pressure to share when income or consumption is observable. I manipulate the observability of income by making public and private windfall income payments in the form of raffle prizes to members of 155 agricultural clubs in central Malawi. These clubs of approximately 10 members each were formed in late 2007 for the purpose of receiving extension services and borrowing through group liability schemes.¹ In contrast to the samples in Jakiela & Ozier (2016) and Boltz, Marazyan & Villar (2015), these clubs do not include extended family members. The experiment described in this paper took place in May 2008.

Members assembled for regularly scheduled meetings (typically, at local primary schools) were given the opportunity to participate in a raffle that would award a cash prize to one winner. In this “public” raffle, the opportunity and value of the prize are announced to the group. Each member draws a ticket from a bag, and the member whose ticket is marked with a star is declared the winner. A staff member records the winner’s name and awards the cash prize in front of the whole group. In this way, everyone present knows that there was a raffle; the identity of the winner; and the value of the prize.

Immediately after the awarding of the public prize and conclusion of the club meeting, all club members are interviewed by our field team in private locations near the meeting spot. These private interviews provide the opportunity to award the second, private prize.

Using club rosters, I pre-select a second individual to receive the same amount of money secretly.² I refer to this person as the “private” winner. The club is not told about the second prize, and instead of awarding the money publicly, the private winner is given his cash while

¹Club members are participants in an experiment about using dynamic incentives to increase loan repayment rates (Gine, Goldberg & Yang 2016).

²In practice, I select a secret winner and several ranked alternates, in case the designated recipient is absent or was independently selected for the public prize.

responding to the baseline survey out of view of others in the community. He is told that no one else in the community has received money in secret, and that no one will be told that he (the private winner) received money. A short supplemental survey about expected use of the prize money is administered to both public and private raffle winners. Because the supplemental survey is brief and completing the baseline survey takes longer for some group members than others, it is unlikely that the time to complete the raffle questionnaire signals anything out of the ordinary to other group members.

Both the public and private raffle winners receive MK 2500 (\$US 17.86, at an exchange rate of MK 140 = \$1 US) paid in cash and on the same day. That sum is roughly equivalent to one-tenth of average annual per capita cash income in Malawi, and at the time would have purchased 25 kg of fertilizer or five chickens. Since the public and private raffle winners are randomly chosen, any differences between how they choose to use the money can be attributed solely to the effect of their communities' awareness – or lack thereof – of their windfall.

3 Data

My final sample is of 294 raffle winners, half of whom won in “public” settings and the other half of whom won under “private” conditions. All prize recipients were surveyed in May 2008; data include baseline characteristics and respondents' plans for how and when they would spend the windfall income.

Follow-up surveys were administered to the raffle winners in 81 of the initial 155 clubs in August 2008.³ At least one raffle winner was present in 77 of those clubs; in total, 114 prize recipients participated in follow-up surveys about their actual use of the prize money.

Table 1 presents summary statistics for baseline characteristics of the public and private raffle winners. Public and private winners do not differ significantly in their gender, age, years of education, land owned, household size, number of children, or house quality score. Including or excluding these baseline characteristics does not affect the sign or significance of subsequent results. However, public and private winners do differ substantially (though not significantly at conventional levels) in their likelihood of being resurveyed in August. I examine this apparently selective attrition in Table 2.

Not all clubs were resurveyed in August, but since each club has one private and one public winner, we anticipate equal response rates for the two types of raffle winners. In fact, public raffle winners (43.2%) were somewhat more likely than private raffle winners (34.5%) to appear in the August survey; the p-value for the test of equal attrition is 0.127. Higher

³The raffle experiment was embedded within the Gine, Goldberg & Yang (2016) study, and that project revisited only a subset of the initial clubs in August 2008.

attrition among private raffle winners could indicate fear of exposure: private winners might have been concerned that, despite earlier assurances, their windfall would be revealed to the group during subsequent encounters with the survey team. If the private winners who were most concerned about privacy (and therefore avoided the follow-up survey) were also those whose spending was most different from public winners, then this type of attrition would bias results towards zero.

Those raffle winners who did respond to the August survey were not observably different from the attriters in their gender, age, land ownership, household size, number of children, or house quality score. They do have an average of one year less schooling ($p=0.031$). Overall, the samples in both raffle questionnaires are well-balanced on observable characteristics.

I study two outcomes: timing of expenditures and sharing of prize money. Respondents were asked to report each planned (May) or realized (August) expenditure of prize money, whether the expenditure was for their own consumption or for someone else’s, and when the expenditure took place relative to the prize distribution.

To study timing, I aggregate spending by date: the same day as the raffle, within one week of the raffle, within the same month (in May), and in each of the three subsequent months. My primary measure of “immediate” spending is money spent within one week of the raffle, since the lotteries, surveys, and related activities occupied most of the day and gave prize recipients little time to spend. Also, market days happen once per week in most villages, so the week is a natural interval for measuring expenditure. Results are not sensitive to using the narrower same-day time frame.

To measure sharing, I aggregate expenditures into five mutually exclusive and exhaustive categories: consumption by the prize recipient him- or herself, consumption by others in the recipient’s household, consumption by others not in the recipient’s household, investment or purchase of durable goods for the household, and savings. I include purchase of agricultural inputs such as fertilizer and pesticides, purchase of livestock, and purchase of building materials in the “investment” category. Results for analysis of these categories are not sensitive to alternative definitions of investment, such as removing livestock.

In both cases, I run OLS regressions of the outcome on an indicator for whether the respondent received the prize in a public setting:

$$Y_{ic} = \alpha + \beta \text{Public}_i + \mathbf{X}_i + \epsilon_{ic} \tag{1}$$

Some specifications include baseline measures of the variables included in the balance tests from Table 1. Since the sample is comprised entirely of raffle winners, the coefficient β measures the effect of receiving the prize in a public setting relative to in private. Results

are robust to including club fixed effects.⁴

4 Results

4.1 Timing of expenditures

People likely face more requests for money from their friends and families when they are known to have cash on hand than otherwise. Then, randomizing whether raffle winners in my experiment received cash secretly or in a public setting generates variation in the social pressure they initially face over sharing their windfall with others. Over time, however, information about the private lotteries may become public, as private winners either reveal information to others or are observed spending in excess of their usual habits. Therefore, spending money quickly – before others come to know about it – may be a strategy for controlling expenditures and maximizing private welfare.

Rapid spending reduces welfare when individuals have preferences for smoothing consumption or when hasty decisions mean making purchases at higher prices than would be available with more opportunity to bargain or comparison shop. Differences in the timing of expenditures between those who receive money privately compared to publicly reflects constraints imposed by public information about income. I test whether these constraints are binding by comparing expenditures of public and private raffle winners in the week following prize disbursement.

Table 3 compares anticipated (measured in the April survey) and ex post (August survey) spending from the prize money in the week in which it was received. Respondents who received the public prize anticipated spending MK 1932 of the MK 2500 within a week. Those who received money in public anticipated MK 236 higher immediate expenditures. The difference persists when controlling for baseline characteristics (column 2) and the magnitude of the difference is slightly higher among the subsample who are resurveyed in August (column 3), though the difference is not statistically significant in the smaller sample. The advantage of using the August survey is that it contains retroactive reports of actual spending, rather than relying on anticipated spending. Private prize winners recall spending somewhat less than anticipated, MK 1251, during the week of the raffle. The difference between the private and public winners is MK 321, statistically significant at the 90 percent confidence level and equal to a 30 percent or 0.30 standard deviation(SD) increase relative to the private raffle winners. The difference between the groups is slightly smaller (MK 301) and not statistically significant when controlling for baseline characteristics.

⁴Results available upon request. Coefficients are slightly larger and t-statistics slightly smaller in these specification.

4.2 Income sharing

Social pressure may affect who consumes from a transfer as well as when it is spent. Respondents were asked who would consume each non-durable they reported; in Table 4, I use these classifications to examine spending on winners themselves (column 1), members of their households (column 2), and people outside of their households (column 5). I also report spending on durable goods or investments in column (3). I aggregate the first three categories to obtain a measure of total spending on winners and members of their households; this is reported in column (4). Finally, money that is saved in cash, a bank, or an informal savings institution is reported in column (6).

When surveyed immediately after the lotteries, private winners anticipate spending an average of MK 190 of the prize on themselves. Public winners anticipate an extra MK 27, a difference that is neither economically meaningful nor statistically significant. Public winners anticipate slightly smaller spending on members of their households or investments, for a total anticipated difference of MK 85 on consumption or investment for the household. Public winners report they will spend an extra MK 24 on people outside their household, and will save an average of MK 175 less than those who won private lotteries. None of the differences in anticipated spending are statistically significant, and all are small relative to the total value of the prize. While the sample is underpowered to detect differences, we can rule out that the public raffle *increased* spending in any category by more than one third of a standard deviation relative to the private winners (at the 95 percent confidence level). We can rule out increases of more than 0.06 SD for savings.

Table 5 replicates the analysis in Table 4, except it uses data from the endline survey administered in August. Not surprisingly, the combination of a smaller sample and longer recall period reduce the precision of the estimates, and there are no significant differences between the public and private raffle winners in any categories. Spending on people outside the household was low among both the private (MK 190) and public (MK 209) raffle winners. While public information about windfall income appears to have accelerated the timing of consumption, it does not seem to have changed its composition or led to explicit income sharing. We reject that the public raffle increased average income sharing by any more than MK 377 (calculated from column (5), $18.973 + 1.96 * 86.553$), or 15 percent of the total prize.

4.3 Heterogeneity

Social pressure to share income may have differential effects that depend upon position in the community wealth distribution. One hypothesis is that social pressure to share income is opportunistic – exerted when the expected benefit is positive because the target is expected to have resources that can be redistributed. While wealthier individuals will always face

greater pressure to share under this hypothesis, awarding money in public rather than secretly may disproportionately affect behavior of poorer individuals. Since wealthy individuals are always presumed to have money that could be shared, winning a sum that is small relative to permanent income or wealth (though meaningful relative to short term consumption) may not change their spending or sharing patterns. For poor individuals, however, the raffle may create a different social dynamic. In general, these individuals face less pressure to share simply because in expectation, they have less to contribute. If they receive windfall income in public, however, they are known to have liquidity, and therefore become targets for social pressure.

Land ownership is a proxy for long-term wealth. Therefore, I test for evidence of this sort of heterogeneity by estimating the effect of the public raffle treatment on winners whose land ownership is above the median for their farming club, compared to those whose land ownership is below their club median:⁵

$$Y_{ic} = \alpha + \beta_1 \text{Public}_i + \beta_2 \text{Above median land}_{ic} + \beta_3 \text{Public}_i \times \text{Above median land}_{ic} + \mathbf{X}_i + \epsilon_{ic} \quad (2)$$

The coefficient β_1 measures the effect of income observability on households with below-median land ownership. The difference between the effect of income observability on poorer and wealthier households is captured by β_3 , and the total effect on wealthier households equals $\beta_1 + \beta_3$.

Table 6 reports estimates of heterogeneous effects of the public prize on anticipated (column 1) and ex post (column 2) spending in the week after the raffle. The public prize significantly increases the amount that winners who own less land than the median member of their club expect to spend in the week following the raffle. The effect of the public raffle is significantly lower for wealthy winners. Ex post, the magnitudes are smaller and no longer significant, but follow the same pattern. One explanation consistent with this pattern is the hypothesis above: poorer winners anticipate that they will face pressure to share when they have atypical liquidity, but wealthier winners do not expect the prize to change their exposure to such redistributive pressure.

Next, I examine how poorer and wealthier club members spend their prize. Table 7 analyzes the same categories of spending used in Tables 4 and 5. The results are imprecise, but the change in private consumption is consistent with opportunistic social pressure. Winners who have below-median landholding reduce the amount they spend on themselves if they win public instead of private lotteries. However, the effect vanishes for winners who have above-median landholding.

⁵I use baseline data from the full sample of club members, not just raffle winners, to calculate club medians.

5 Conclusion

Observed income-sharing is low for Malawian farmers who received \$18 in windfall income, but the timing of expenditures depends on the observability of the prize. Those who received the prize in public – and consequently were exposed to greater pressure to share income – spent more of their money in the week immediately following the transfer than peers in the same communities whose prize was awarded secretly. In this experiment, subjects may have spent their windfall quickly to evade sharing obligations. Previous studies have documented that sharing norms can affect investment (Jakiela & Ozier 2016) and consumption (Kinnan 2014, Boltz, Marazyan & Villar 2015), and this study adds evidence that such pressure may change the timing of consumption even when it does not shift its composition.

There is suggestive evidence of opportunistic social pressure that responds to the expected availability of funds. First, social pressure in this experiment is triggered by the observability of income. Second, the pressure generated by public information has a bigger effect on expenditure timing and private consumption for poor individuals, who may face increased pressure to share when they are known to have cash compared to periods when they are less likely to have resources to share. In Malawi, many payments, especially those that target poor households, are easily observable. Wages for employment and the national public works scheme are made via highly visible “pay parades.” Agricultural clubs often disburse loans or proceeds from cash crops sold through club accounts in group meetings. These results suggest that the information environment for payments is a constraint to the timing of expenditures, so changes in payments, including utilizing mobile money to increase privacy without increasing transaction costs, may be welfare improving.

Tables

Table 1: Summary statistics

	Public	Private	P-value: public vs. private
Male	0.945 (0.019)	0.919 (0.023)	0.373
Age	43.356 (1.134)	44.885 (1.061)	0.325
Years of education	5.822 (0.304)	6.081 (0.279)	0.530
Land owned (acres)	6.724 (0.468)	6.871 (0.597)	0.847
HH size	5.753 (0.166)	5.743 (0.162)	0.965
Number of Children	1.521 (0.103)	1.561 (0.105)	0.784
House quality score (PCA)	0.054 (0.086)	-0.040 (0.080)	0.428
In follow-up sample	0.432 (0.041)	0.345 (0.039)	0.127
Observations	146	148	

Table 2: Test of differential attrition

	In follow-up sample	Not in follow-up sample	P-value: differential attrition
Male	0.947 (0.021)	0.922 (0.020)	0.406
Age	44.833 (1.266)	43.678 (0.983)	0.469
Years of education	5.395 (0.326)	6.306 (0.263)	0.031
Land owned (acres)	6.877 (0.532)	6.748 (0.521)	0.868
HH size	5.737 (0.181)	5.756 (0.151)	0.937
Number of Children	1.579 (0.118)	1.517 (0.093)	0.679
House quality score (PCA)	-0.018 (0.093)	0.022 (0.076)	0.740
Observations	114	180	

Table 3: Effect of income observability on spending within one week

	(1)	(2)	(3)	(4)	(5)
	Anticipated (baseline)			Ex post (endline)	
Public	236.097** (113.878)	232.312** (117.396)	293.754 (189.479)	320.794* (187.870)	301.334 (201.854)
Controls for baseline characteristics		X	X		X
Observations	294	294	114	114	114
R-squared	0.01	0.04	0.02	0.02	0.05
Mean of dep. var. for private winners	1815.07	1815.07	1715.29	1073.33	1073.33

Data in columns (1)-(3) come from the survey conducted immediately after the lotteries, and refer to anticipated spending in the next week. Data in columns (4) and (5) come from the follow up survey administered in a subset of villages in August, and are respondents' recollection of their actual spending in the week after the raffle. The sample in columns (1) and (2) includes all raffle winners surveyed at baseline. The sample in columns (3) to (5) includes only those winners resurveyed at endline. When covariates are included, they are the variables included in Table 3. OLS regressions. Robust standard errors are clustered at the village level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$.

Table 4: Effect of income observability on categories of anticipated expenditures

	(1) Self non-durable	(2) Household non-durable	(3) Investment	(4) Total HH (1)+(2)+(3)	(5) Non-household	(6) Save
Public	27.435 (66.619)	-60.268 (101.814)	-51.715 (106.414)	-84.548 (126.396)	24.598 (36.605)	-174.898 (115.299)
Observations	294	294	294	294	294	294
R-squared	0.03	0.04	0.04	0.04	0.07	0.03
Mean of dep. var. for private winners	179.32	838.38	596.62	1614.32	78.85	1960.74

Data come from the survey conducted immediately after the lotteries, and refer to total anticipated spending from the prize money. Column 1 includes expenditure on non-durables to be consumed by the winner himself; column 2 is expenditures on non-durables to be consumed by members of the winner's household; column 3 includes investment and purchase of non-durables for the household; column 4 is the sum of 1-3, column 5 is all items purchased for people outside the winner's household, and column 6 is money to be saved in formal or informal institutions or in cash. Baseline covariates are included in all specifications, and include all variables from Table 3. OLS regressions. Robust standard errors are clustered at the village level. * p<0.10, ** p<0.05, *** p<0.001.

Table 5: Effect of income observability on categories of expenditures, ex post

	(1) Self non-durable	(2) Household non-durable	(3) Investment	(4) Total HH (1)+(2)+(3)	(5) Non-household	(6) Save
Public	15.141 (137.845)	-205.202 (175.533)	176.714 (167.442)	46.271 (100.364)	18.973 (86.553)	-0.825 (127.365)
Observations	114	114	114	114	114	114
R-squared	0.09	0.14	0.05	0.02	0.02	0.03
Mean of dep. var. for private winners	515.69	979.61	559.02	2395.49	190.20	2205.29

Data come from the survey conducted in August 2008 in a subset of villages, and refer to total spending from the prize money. Column 1 includes expenditure on non-durables to be consumed by the winner himself; column 2 is expenditures on non-durables to be consumed by members of the winner's household; column 3 includes investment and purchase of non-durables for the household; column 4 is the sum of 1-3, column 5 is all items purchased for people outside the winner's household, and column 6 is money to be saved in formal or informal institutions or in cash. Baseline covariates are included in all specifications, and include all variables from Table 3. OLS regressions. Robust standard errors are clustered at the village level. * p<0.10, ** p<0.05, *** p<0.001.

Table 6: Effect of income observability on spending within one week, by wealth

	(1) Anticipate (baseline)	(2) Ex post (endline)
Public	781.836** (331.179)	554.096 (464.071)
Public * Above club median (land owned)	-683.595* (353.696)	-327.314 (493.559)
Observations	294	114
R-squared	0.06	0.05
Mean of dep. var. for private winners	1815.07	1073.33

See notes for Table 3. “Above club median (land owned)” is an indicator that equals 1 when a respondent’s land owned is above the median calculated for all surveyed individuals in his/her farming club, including those who did not win raffle prizes. Covariates are included in all specifications; these include all variables from Table 3, except substitute an indicator for above-median land ownership for the continuous measure. OLS regressions. Robust standard errors are clustered at the village level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$.

Table 7: Effect of income observability on categories of expenditures, by wealth

	(1) Self non-durable	(2) Household non-durable	(3) Investment	(4) Total HH (1)+(2)+(3)	(5) Non-household	(6) Save
Public	-489.191 (318.597)	-86.093 (400.768)	184.740 (341.578)	-93.418 (71.391)	-71.324 (216.352)	-25.758 (235.583)
Public * Above club median (land owned)	647.628* (377.746)	-161.582 (447.470)	-19.392 (412.230)	173.311 (122.629)	117.259 (261.645)	24.793 (284.096)
Observations	114	114	114	114	114	114
R-squared	0.11	0.14	0.03	0.03	0.01	0.02
Mean of dep. var. for private winners	515.69	979.61	559.02	2395.49	190.20	2205.29

See notes for Table 5. “Above club median (land owned)” is an indicator that equals 1 when a respondent’s land owned is above the median calculated for all surveyed individuals in his/her farming club, including those who did not win raffle prizes. Covariates are included in all specifications; these include all variables from Table 3, except substitute an indicator for above-median land ownership for the continuous measure. OLS regressions. Robust standard errors are clustered at the village level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$.

References

- Ashraf, N.** 2009. “Spousal Control and Intra-Household Decision Making: An Experimental Study in the Philippines.” *American Economic Review*, 99(4): 1245–1277.
- Baland, J, C Guirkingier, and C Mali.** 2007. “Pretending to Be Poor: Borrowing to Escape Forced Solidarity in Credit Cooperatives in Cameroon.” *Working Paper, Center for Research in Economic Development, University of Namur.*
- Boltz, Marie, Karine Marazyan, and Paola Villar.** 2015. “Income Hiding and Informal Redistribution: A Lab-in-the-Field Experiment in Senegal.” Working Paper, Paris School of Economics.
- Gine, Xavier, Jessica Goldberg, and Dean Yang.** 2016. “Credit Market Consequences of Improved Personal Identification: Field Experimental Evidence from Malawi.” *American Economic Review*, 8(1): 129–149.
- Jakiela, Pamela, and Owen Ozier.** 2016. “Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies.” *Review of Economic Studies*, 83(1): 231–268.
- Kinnan, Cynthia.** 2014. “Distinguishing Barriers to Insurance in Thai Villages.” Working Paper, Northwestern University.
- Maranz, D.** 2001. *African Friends and Money Matters: Observations from Africa*. Vol. 37 of *Publications in Ethnography*, SIL International.
- Platteau, JP,** ed. 2000. *Institutions, Social Norms, and Economic Development*. Harwood Academic Publishers.
- Townsend, R. M.** 1994. “Risk and Insurance in Village India.” *Econometrica*, 62(3): 539–591.