

American Economic Association

Giffen Behavior and Subsistence Consumption

Author(s): Robert T. Jensen and Nolan H. Miller

Reviewed work(s):

Source: *The American Economic Review*, Vol. 98, No. 4 (Sep., 2008), pp. 1553-1577

Published by: [American Economic Association](#)

Stable URL: <http://www.jstor.org/stable/29730133>

Accessed: 28/01/2012 12:36

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Economic Association is collaborating with JSTOR to digitize, preserve and extend access to *The American Economic Review*.

<http://www.jstor.org>

Giffen Behavior and Subsistence Consumption

By ROBERT T. JENSEN AND NOLAN H. MILLER*

This paper provides the first real-world evidence of Giffen behavior, i.e., upward sloping demand. Subsidizing the prices of dietary staples for extremely poor households in two provinces of China, we find strong evidence of Giffen behavior for rice in Hunan, and weaker evidence for wheat in Gansu. The data provide new insight into the consumption behavior of the poor, who act as though maximizing utility subject to subsistence concerns. We find that their elasticity of demand depends significantly, and nonlinearly, on the severity of their poverty. Understanding this heterogeneity is important for the effective design of welfare programs for the poor. (JEL D12, O12)

The “Law of Demand,” which holds that as the price of a good increases, consumers’ demand for that good should decrease, is one of the bedrock principles of microeconomics. Economists have long recognized, however, that the axioms of consumer theory do not guarantee that demand curves must slope downward, and that the Law of Demand, while descriptively valid in many situations, may not apply to very poor consumers facing subsistence concerns. Alfred Marshall first publicized this idea in the 1895 edition of his *Principles of Economics*:

As Mr. Giffen has pointed out, a rise in the price of bread makes so large a drain on the resources of the poorer labouring families and raises so much the marginal utility of money to them, that they are forced to curtail their consumption of meat and the more expensive farinaceous foods: and, bread being still the cheapest food which they can get and will take, they consume more, and not less of it (208).

Since Marshall’s time, a discussion of “Giffen” behavior has found its way into virtually every basic economics course, despite a lack of real-world evidence supporting Marshall’s conjecture.¹ Studies by George J. Stigler (1947) and Roger Koenker (1977) argue that demand for neither bread nor wheat was upward sloping in Britain during Marshall’s time. The standard textbook example of a Giffen good, potatoes during the Irish Potato Famine of 1845–1849 (Paul A. Samuelson 1964), has also been discredited (Sherwin Rosen 1999). Not only are there no data to support the claim, but at a more basic level it is unlikely that consumption of potatoes could

* Jensen: Watson Institute for International Studies, Brown University, 111 Thayer St., Providence, RI 02912, and National Bureau of Economic Research (e-mail: robert_jensen@brown.edu); Miller: John F. Kennedy School of Government, Harvard University, 79 JFK St., Cambridge, MA 02138 (e-mail: nolan_miller@harvard.edu). We thank Alberto Abadie, Chris Avery, Sebastian Bauhoff, Amitabh Chandra, Suzanne Cooper, Daniel Hojman, Brian Jacob, Elizabeth Lacey, Erzo Luttmer, Mai Nguyen, Albert Park, Rodrigo Wagner, Sangui Wang, Richard Zeckhauser, and two anonymous referees for valuable comments and discussions, and Frank Mou, Dulles Wang, and Fan Zhang for research assistance. We gratefully acknowledge financial support from the National Institute of Aging, the William F. Milton Fund at the Harvard Medical School, the Dean’s Research Fund at the John F. Kennedy School of Government, the Center for International Development at Harvard University, and the Hefner China Fund.

¹ We use the term “Giffen behavior” rather than “Giffen good” to emphasize that the Giffen property is one that holds for particular consumers in a particular situation and therefore depends on, among other things, prices and wealth. Thus, it is not the good that is Giffen, but the consumers’ behavior.

have increased when the price rose during the famine, at least in the aggregate, precisely because the price rise was caused by a blight that destroyed much of the crop.² While some laboratory studies have found evidence of Giffen behavior, these experiments have been far from removed from reality.³

In this paper we present data from a field experiment exploring the response of poor households in China to changes in the prices of staple food items, which provide the first rigorous, empirical evidence of real-world Giffen behavior. In fact, we find Giffen behavior with respect to *two* goods, rice and wheat. Further, these goods, and the populations that exhibit Giffen behavior, meet some basic but common conditions that suggest this behavior may be widespread in the developing world. Thus, the absence of previously documented cases most likely results from inadequate data or empirical strategies rather than from their nonexistence.

Giffen behavior has long played an important, though controversial, role in economic pedagogy, as well as in the history of economic thought.⁴ Finding convincing evidence of such behavior is important, however, for economic theory more broadly. The fact that there has, to date, been no convincing evidence of Giffen behavior stands as a minor embarrassment to economists (John H. Nachbar 1998), one that is reflected in the discussion of the Giffen phenomenon often being presented as a paradox of economic theory rather than as a real (or even possible) mode of behavior (e.g., Stigler 1947). This lack of evidence has prompted a range of reactions among economists. Some have interpreted it as support for the descriptive validity of the Law of Demand:

Perhaps as persuasive a proof [of the “Law of Demand”] as is readily summarized is this: if an economist were to *demonstrate* its failure in a particular market at a particular time, he would be assured of immortality, professionally speaking, and rapid promotion while still alive. Since most economists would not dislike either reward, we may assume that the total absence of exceptions is not from lack of trying to find them.

— Stigler (1966, 24).

Others’ reactions to the lack of validation for the Giffen phenomenon have been more extreme, interpreting it as an indictment of neoclassical consumer theory. Along these lines, Lawrence A. Boland (1977) points out that not only is the theory unable to rule out Giffen behavior, it is also unable to explain why it is not observed. Put another way, if the neoclassical model is correct, then under certain (albeit uncommon) conditions, Giffen behavior should exist. If it has not been observed, it is either because the appropriate conditions have not been satisfied, the appropriate data have not been available to measure it, or the theory is incomplete or flawed.⁵

² Gerald P. Dwyer Jr. and Cotton M. Lindsay (1984) present a summary of the basic case against the potato version of the Giffen paradox. See also Terrence McDonough and Joseph Eisenhauer (1995). In both the bread and potato cases, it is possible that poor individuals exhibited Giffen behavior but the market overall did not. However, the data to test this hypothesis do not exist.

³ Raymond C. Battalio, John H. Kagel, and Carl A. Kogut (1991) find evidence of upward sloping demand among rats given limited “budgets” and the choice between root beer and a quinine solution, and R. J. DeGrandpre et al. (1993) find in a laboratory setting that human smokers, given the choice between brands of cigarettes and a limited budget of “puffs,” can exhibit Giffen behavior.

⁴ The lack of verified examples has raised numerous concerns about the pedagogical role of the Giffen story: “Since the Giffen paradox is not useful for understanding the Irish Experience, is it asking too much for future writers of elementary texts to find another example? Fictions have no place in the teaching of economics” Rosen (1999); “We shall have to find a new example of the positively sloping demand curve, or push our discussion of it deeper into footnotes” Stigler (1947).

⁵ Others have argued that it is not our understanding of consumers that is flawed, but rather our understanding of markets. For example, William Dougan (1982) argues that markets with upward sloping demand curves are inherently unstable, and thus unlikely to be observed, while Nachbar (1998) shows in a general equilibrium framework that observing the equilibrium price and quantity of a good move in the same direction in response to a supply shock implies that the commodity is normal, not inferior, and thus not Giffen at all. Thus, economists looking for Giffen behavior at the level of the market are unlikely to find it. Yoram Barzel and Wing Suen (1992) argue that if consumers can exploit

Beyond documenting the existence of Giffen behavior, our field experiment also provides an opportunity to study more broadly the consumption behavior of the “extreme poor,” a population that worldwide includes more than one billion people living below the World Bank’s extreme poverty line of one dollar per person per day. These households, like Marshall’s “labouring families” and those in our sample, are often highly dependent on a single staple food for the bulk of their nutritional needs. Consequently, fluctuations in the prices of these staple foods can have large effects on real wealth and purchasing power. Anecdotally, such price fluctuations, even fairly large ones, are increasingly common in developing countries.⁶ And while there is a large literature examining household vulnerability and responses to income shocks, there is comparably little evidence with respect to price shocks. Our analysis, by focusing on the extremely poor and by introducing exogenous price changes for staple foods, is useful for understanding this vulnerability.

In an earlier study using panel data from the China Health and Nutrition Survey (Jensen and Miller 2002), we found suggestive evidence that poor households in China exhibited Giffen behavior with respect to their primary dietary staple (rice in the south; wheat and/or noodles in the north).⁷ However, because the study relied on possibly endogenous variation in market prices, we were unable to identify a causal relationship between price changes and consumption. To address this concern, for the present study we conducted a field experiment in which, for five months, randomly selected households were given vouchers that subsidized their purchases of their primary dietary staple. Building on the insights of our earlier analysis, we studied two provinces of China: Hunan in the south, where rice is the staple good, and Gansu in the north, where wheat is the staple. Our analysis in these provinces focused on households classified as the “urban poor,” a population that includes approximately 90 million individuals throughout China (Martin Ravallion 2007).

Using consumption surveys gathered before, during, and after the subsidy was in place, we find strong evidence that poor households in Hunan exhibit Giffen behavior with respect to rice. That is, lowering the price of rice via the experimental subsidy caused households to reduce their demand for rice, and removing the subsidy had the opposite effect. This finding is robust to a wide range of empirical specifications. In Gansu, the evidence for Giffen behavior is somewhat weaker, due to the partial failure of two of the basic conditions under which such behavior is expected; namely that the staple good has limited substitution possibilities, and that households are not so poor that they consume only staple foods. Focusing our analysis on those whom the theory identifies as most likely to exhibit Giffen behavior, we find stronger evidence of its existence.

We also provide important new insights into the consumption behavior of poor households. In particular, we find the consumption response to an increase in the price of a staple good follows a previously undocumented inverted-U pattern predicted by consumer theory in the presence of subsistence concerns, with the very poorest and the least poor of the poor responding by decreasing demand in response to an increase in the price of the staple in the standard way, while the group in the middle increases demand (i.e., exhibits Giffen behavior). These results have

intertemporal variations in prices, they will never purchase more of a good when its price increases. The very poor, however, are likely to be liquidity constrained and thus lack this ability.

⁶ For example, occasional reports from China note rice prices that double from year to year in some localities (“Surge in Consumer Prices Puts China on Guard,” *China Daily*, April 22, 2004). Jed Friedman and James Levinsohn (2002) note that the mean price of rice increased by almost 200 percent in Indonesia (where the typical household spends nearly 30 percent of its total household budget on rice) during the 1997/98 financial crisis.

⁷ Ours is not the first study to suggest rice as a likely candidate for Giffen behavior. Dwyer and Lindsay (1984) propose (but do not test) this possibility for Singapore, and John-Ren Chen (1994) finds suggestive evidence of positively sloped demand for rice in Taiwan.

important implications for the design, targeting, and evaluation of programs aimed at improving nutrition among the poor.

The paper continues in Section I, where we present a discussion of the consumption behavior of the poor that motivates Giffen behavior. Section II discusses the field experiment, data, and estimation strategy. Section III presents the results, and Section IV concludes.

I. Giffen Behavior and Consumption among the Poor

The conditions under which we would expect Giffen behavior can be demonstrated by elaborating Marshall's statement.⁸ Imagine an impoverished consumer near a subsistence level of nutrition, whose diet consists of only two foods, a "basic" or staple good (in Marshall's case, bread) and a "fancy" good (meat). The basic good offers a high level of calories at low cost, while the fancy good is preferred because of its taste but provides few calories per unit currency. A poor consumer will therefore eat a lot of bread in order to get enough calories to meet his basic needs and use whatever money he has left over to purchase meat. Now, if the price of bread increases, he can no longer afford the original bundle of foods. And if he increases his consumption of meat, he will fall below his required caloric intake. So, he must instead increase his consumption of bread (which is still the cheapest source of calories) and cut back on meat.

The Giffen phenomenon illustrates the potential significance of the wealth effects of price changes for extremely poor households. Although the price increase makes the staple less attractive in relative terms, the fact that it makes the consumer so much poorer (in real terms) forces him to consume more bread. Translating this to the language of consumer theory, the conditions under which Giffen behavior is likely to be observed, therefore, include that the good in question be strongly inferior and that expenditure on that good comprise a large portion of the consumer's budget. As can be seen from the elasticity version of the Slutsky equation, $\varepsilon_p = \varepsilon_p^h - b\varepsilon_w$, where ε_p is the observed price elasticity of demand, $\varepsilon_p^h (\leq 0)$ is the Hicksian compensated elasticity, ε_w is the wealth elasticity, and b is the budget share of the good, only then can the wealth effect of a price increase be large enough to offset the pure substitution effect.

In light of these observations, we can state a set of conditions under which Giffen behavior is most likely to be observed:⁹

- C1: Households are poor enough that they face subsistence nutrition concerns.
- C2: Households consume a very simple diet, including a basic (staple) and a fancy good.
- C3: The basic good is the cheapest source of calories available, comprises a large part of the diet/budget, and has no ready substitute.

When dealing with extreme poverty of the sort exhibited by the urban poor in China, another requirement becomes important. While consumers who are too wealthy will not exhibit Giffen behavior, those who are too poor also cannot exhibit Giffen behavior. To take an extreme example, consider a consumer who is so poor that he consumes only bread. When the price of bread increases, he has no choice but to consume less bread. Thus, it is critical to the Giffen story that the consumer be consuming at least some of the fancy goods (e.g., meat) that are more enjoyable but

⁸ Much of the theory of Giffen behavior has previously appeared elsewhere. The interested reader should see the online Appendix to this document (<http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.4.1553>) for a discussion of the theory underlying this behavior.

⁹ Some of these conditions have been noted before by, for example, Otis W. Gilley and Gordon V. Karels (1991).

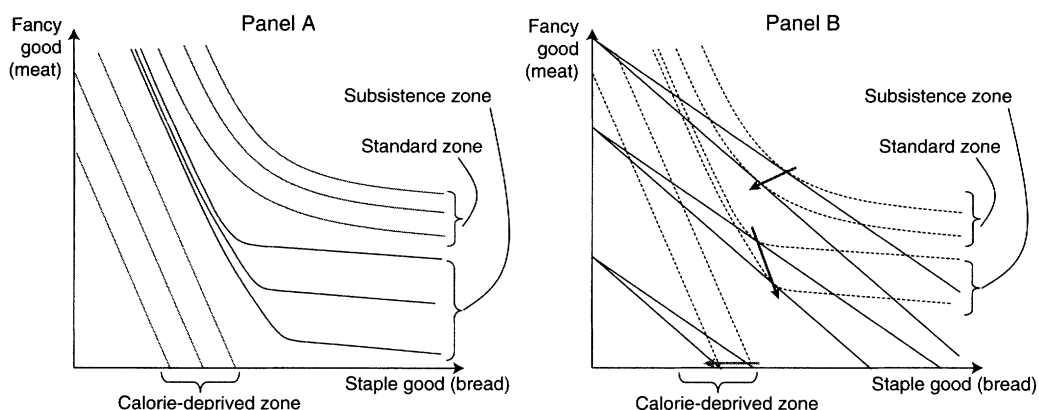


FIGURE 1. ZONES OF CONSUMER PREFERENCES

more expensive sources of calories so that he has something to substitute away from when the price of the staple increases. In light of this, we add the following requirement to the three stated above:

C4: Households cannot be so impoverished that they consume *only* the staple good.

The theory thus predicts that only consumers who are poor, but not too poor, will exhibit Giffen behavior. Panel A of Figure 1 depicts the indifference curves for a typical consumer choosing how much of the basic and fancy goods to consume. The basic or staple good is relatively high in calories, while the fancy good offers more “taste,” i.e., the enjoyable but nonnutritive aspects of food.¹⁰ The consumer’s indifference map can be divided into three regions. The outer set of indifference curves corresponds to the standard case, where the consumer’s calorie intake is well above subsistence. Over this range the consumer trades off between calories and taste (and thus between the basic and fancy goods) in an ordinary way. As the consumer’s calorie consumption decreases, he crosses into a “subsistence zone.” Over this range, caloric intake becomes much more important to the consumer. Consequently, the consumer’s indifference curves take on the familiar “elbow” shape associated with Giffen behavior.¹¹ Consumers in this range behave as if they maximize taste, subject to the constraint that they meet their minimum caloric needs. As the consumer’s calorie consumption decreases even further, he crosses from the subsistence zone to the calorie-deprived zone. In this region, the consumer’s calorie intake is below subsistence levels. Hence, his primary concern is maximizing calories, and the consumer’s indifference curves are, in effect, iso-calorie curves.

The consumer’s response to an increase in the price of the staple good will differ across the three regions of his indifference map. When the consumer is relatively wealthy, he will demand a bundle of goods in the standard zone. In this case, as illustrated in panel B of Figure 1, we expect the consumer to respond to an increase in the price of the staple good by consuming less of it. Thus, demand is downward sloping. As wealth decreases, the consumer’s demand moves into the

¹⁰ The substitution across goods with varying nutritional and nonnutritional attributes also motivates the literature concerned with the income elasticity of demand for calories (see John Strauss and Duncan Thomas 1995; Angus Deaton 1997).

¹¹ See the online Appendix for more discussion of the relationship between the shape of indifference curves needed to generate Giffen behavior and subsistence concerns.

subsistence zone, and the consumer focuses more on maintaining caloric intake as his primary goal. It is over this region that Giffen behavior arises, as the consumer responds to an increase in the price of the staple good by substituting toward the cheaper source of calories, which is still the staple good. Over this range, the consumer still trades off calories against taste, although caloric intake is given much greater importance. A consumer in the subsistence zone behaves, in effect, as if he maximizes taste subject to the constraint that calories reach a certain minimum requirement. Finally, as the consumer's wealth decreases even further, he is unable to afford to meet his subsistence calorie needs. Calories are maximized by consuming only the staple good, and so the consumer has no choice but to respond to an increase in the price of the staple good by consuming less of it.

This set of predictions provides important insight into the search for Giffen behavior and may also help explain why such behavior has not been previously detected. First, we should not expect to observe Giffen behavior at the market level; a subset of consumers might exhibit Giffen behavior with respect to a particular commodity while the overall market exhibits downward sloping demand. Thus, the search should not be for a "Giffen good," but for Giffen behavior. In addition, the nonlinear response suggests the search is even more nuanced than just focusing on the poorest households. Selecting just the very poorest, or even aggregating over a broader set of households that includes both those in the calorie-deprived and subsistence zones, may not be sufficient. Consumers in the intermediate, subsistence range must be isolated in order to find such behavior.

II. Empirical Strategy

A. *The Experiment*

A central problem in documenting Giffen behavior, and indeed in any analysis of demand, is finding both sufficient and exogenous price variation. As a practical problem, whether data are cross-sectional, time-series, or panel, there is often not a great deal of variation in prices for the kinds of goods likely to be candidates for Giffen behavior. This applies especially to cross-sectional data, as arbitrage should eliminate spatial price differences, especially for easily storable and nonperishable commodities such as grains. Further, the prices for staple goods might even be fixed by the government for the poorest households, such as under India's Public Distribution System, and any remaining price variation may be due to unobservable quality differences. A more serious concern is that, even with sufficient price variation, the source of that variation is often potentially endogenous, since price is the equilibrium of a system of simultaneous equations. A positive correlation between price and consumption could simply represent shocks to, or differences in, demand over space or time; higher demand leads to higher prices, which could be misinterpreted as Giffen behavior. Although instrumental variables could address this problem, finding instruments that shift supply but do not directly affect demand is difficult.¹²

¹² Most previous studies of Giffen behavior have failed to address this identification concern. A few cases have used instrumental variables, but with problematic instruments. For example, Anthony E. Bopp (1983) uses refinery utilization rates and the price of crude oil as instruments for the price of kerosene; however, both instruments likely also affect the price of substitute fuels, and are likely to be driven by other unobserved factors also affecting fuel demand, such as weather. Shmuel Baruch and Yakar Kannai (2001) use the lagged prime interest rate as an instrument for the price of a low-grade Japanese alcohol (*shochu*), which is likely to be a poor predictor of the price of *shochu*, or, to the extent that it does predict the price, will likely also affect the prices of substitutes (or income—and thus demand). Rainfall is commonly suggested as an instrument for price. However, rainfall will be an invalid instrument for the price of a given food item, since it likely also affects the prices of other foods, as well as wages and income. One exception to the endogeneity concern in the search for Giffen behavior is the study by David McKenzie (2002), which uses the elimination of tortilla subsidies and price controls as a natural experiment to test for (and ultimately reject) such behavior in Mexico.

To overcome these challenges, we conducted a field experiment in which we provided randomly selected poor households in two Chinese provinces with price subsidies for staple foods. In Hunan, a southern province, rice is the staple good, and in Gansu, a northwestern province, wheat is the staple good (consumed primarily as buns, a simple bread called *mo*, or noodles). These regional differences in preferences are primarily determined by geography, climate, and history, with wheat the dominant crop grown in Gansu and rice dominant in Hunan. Accordingly, we subsidized rice (only) in Hunan and wheat flour (only) in Gansu.

Within each sample cluster (described below), households were randomly assigned to either a control group or one of three treatment groups. Households in the treatment groups were given printed vouchers entitling them to a price reduction of 0.10, 0.20, or 0.30 yuan (Rmb; 1 Rmb \approx \$0.13) off the price of each *jin* (1 *jin* = 500g) of the staple good (the subsidy level stayed fixed for each household over the course of the study). These subsidies represented substantial price changes, since the average preintervention price of rice in Hunan was 1.2 yuan/*jin*, and the average for wheat flour in Gansu was 1.04 yuan/*jin*. The vouchers were printed in quantities of 1, 5, and 10 *jin*, and the month's supply of vouchers was distributed at the start of each month, with each household receiving vouchers for 750g per person per day (about twice the average per capita consumption). All vouchers remained valid until the end of the intervention. Households were told in advance they would receive vouchers for five months and that any unredeemed vouchers would not be honored afterward.

The vouchers could be redeemed at local grain shops. The merchants in these shops agreed to honor the vouchers in exchange for reimbursement and a payment for their participation. Households and merchants were told they were not permitted to exchange the vouchers for anything but the staple good, that there would be periodic auditing and accounting to make sure they were in compliance with the rules, and that any violations would result in their being removed from the study without any additional compensation. Households and merchants were explicitly told that selling the vouchers for cash or reselling rice or wheat bought with the vouchers would result in dismissal from the program.

There are several points about the intervention worth noting. First, all foods in China are sold in free markets, at market determined prices. A 1993 reform of the grain distribution system largely put an end to price controls, state food stores, or free rations. Second, the number of subsidized households in each sample site is trivial relative to the size of the population (all sites were county seats, most with populations over one million), so the intervention could not have affected market prices.¹³ Third, the experiment is predicated on the assumption that either households are limited in their ability to borrow and save, or they have short planning horizons; otherwise, the wealth effect of the five-month subsidy would be trivial, making Giffen behavior unlikely. To the extent the wealth effect of the price change can be smoothed over the lifetime, this will bias us against finding Giffen behavior. Fourth, limiting the quantity of vouchers to 750g/person/day might limit the potential demand response for the staple good (although the amount is still quite generous), but it should not induce Giffen behavior, as might be the case if we limited the vouchers to a quantity smaller than what they would prefer to consume.¹⁴ Finally, while staple foods such as rice can be found in varying qualities or varieties with different prices, because the households in our sample are extremely poor, our data show that they consume almost exclusively only the lowest-cost variety. Therefore, quality substitution in response to the

¹³ Similarly, because the samples were drawn from lists of the poor spread throughout large cities, we believe it is unlikely the various study participants knew each other or the benefits others were receiving.

¹⁴ One concern is that by limiting the potential increase in consumption in response to the price decline, we might skew the average consumption change toward a decline (i.e., Giffen behavior). However, in practice, almost no households even approached the voucher limit, most likely due to their extremely low incomes and a lack of access to credit, so this is unlikely to be a major concern.

price subsidy is not a concern for our analysis. Two final concerns with the experiment, namely whether there was cheating (in the form of cashing out or reselling) despite our rules against doing so or whether the vouchers might create a “salience” or signaling effect, are discussed with the results in Section III E.

B. Data

The survey and intervention were conducted by employees of the provincial-level agencies of the Chinese National Bureau of Statistics. The sample consisted of 100–150 households in each of 11 county seats spread throughout Hunan and Gansu provinces (Anren, Baoqing, Longshan, Pingjiang, Shimen, and Taojiang in Hunan, and Anding, Ganzhou, Kongdong, Qingzhou, and Yuzhong in Gansu), for a total of 1,300 households (650 in each province), with 3,661 individuals. Within each county, households were chosen at random from lists of the “urban poor” maintained by the local offices of the Ministry of Civil Affairs.¹⁵ Households on this list fall below a locally defined poverty threshold (the *Di Bao* line), typically between 100 and 200 yuan per person per month, or \$0.41–\$0.82 per person per day, which is below even the World Bank’s extreme poverty line of one dollar per person per day. It is estimated that about 90 million individuals throughout China live below the *Di Bao* threshold.

The questionnaire consisted of a standard income and expenditure survey, gathering information on the demographic characteristics of household members, as well as data on employment, income, asset ownership, and expenditures. A key component of the survey was a 24-hour food recall diary completed by each household member. Respondents were asked to report everything they ate and drank the previous day, whether inside or outside the home, by specifically listing the components of all foods eaten.¹⁶ These foods were recorded in detail in order to match with the 636 detailed food items listed in the 1991 Food Composition Tables constructed by the Institute of Nutrition and Food Hygiene at the Chinese Academy of Preventative Medicine. As we will see below, because households are very poor, most diets are very simple and consist of a small number of basic (not processed, prepared, or packaged) foods like rice, bean curd, or stir-fried cabbage, so concerns about coding the specific quantities of the various ingredients in a complex dish or meal are not significant.¹⁷

Data were gathered in three waves, conducted in April, September, and December of 2006. After completing the first survey, treatment households were told they would receive the subsidies for five months, from June through October. Thus, the initial interviews occurred before treatment households knew of or received the subsidies, the second occurred after the subsidy had been in place for slightly more than three months, and the final interviews were conducted one to two months after the subsidy had ended, by which time treatment households would likely have exhausted any stocks of rice or wheat flour they may have purchased with the subsidy, and would therefore again be purchasing at the full market price. Sample attrition was extremely low, since the three rounds occurred in a relatively short span. Only 11 of 1,300 households (less than 1 percent) in the first round did not appear in the second round. All households in the second

¹⁵ We chose urban areas because in smaller towns or rural areas many of the poorest households grew rather than purchased their staple food, and lower population density meant fewer households living in extreme poverty, which would have both required a greater number of sample clusters and prevented varying the treatment within clusters.

¹⁶ While it may seem difficult to recall or estimate how many grams of, say, rice were eaten with a meal, for the extreme poor who are on a very limited budget, food is often apportioned and accounted for much more carefully. Further, diets for these extremely poor households often vary little or not at all from day to day, except on special occasions, so recalling the quantity of specific food items is not as difficult.

¹⁷ Similarly, because households were so poor, almost all food (98 percent) was at-home consumption, so respondents were aware of the exact ingredients and quantities used.

C. Estimation Strategy

Given the random assignment of the price change and the panel nature of our survey, our basic strategy is to simply compare the household-level changes in dietary intake²⁰ of the staple good for treatment and control groups. Since assignment to treatment and control groups was randomized within sample counties, we add county*time fixed effects, so that we are in effect comparing the changes for households with different subsidy levels within the same community. This strategy controls for any county-level factors that change over time, such as the price of foods, labor market conditions, or the value of government transfer programs.

We regress the percent change in intake of the staple good for household i in period t on the change in the subsidy (in percent). The percent change formulation normalizes for factors such as household size, composition, and activity level and allows us to interpret the coefficients as elasticities. For each household, we observe two changes: the change between periods 2 and 1 ($t = 2$), capturing the effect of imposing the subsidy; and the change between periods 3 and 2 ($t = 3$), capturing the effect of removing the subsidy. Thus, we estimate

$$(1) \quad \% \Delta \text{staple}_{i,t} = \alpha + \beta \% \Delta p_{i,t} + \sum \gamma \% \Delta Z_{i,t} + \sum \delta \text{County*Time}_{i,t} + \Delta \varepsilon_{i,t},$$

where $\% \Delta \text{staple}_{i,t}$ is the percent change in household i 's consumption of the staple good, $\% \Delta p_{i,t}$ is the percent change in the price of the staple due to the subsidy (negative for $t = 2$ and positive for $t = 3$), $\% \Delta Z$ is a vector of percent changes in other control variables including income (split into earned and unearned (government payments, pensions, remittances, rent, and interest from assets) sources) and household size, and County*Time denotes a set of county×time dummy variables. We compute all changes as arc-percent-changes (i.e., $100 \times (x_t - x_{t-1}) / ((x_t + x_{t-1})/2)$).²¹ The percent change in the subsidy is computed as 100 times the change in the subsidy divided by the average (net of subsidy) price of the staple good in the two corresponding rounds. The results are robust to a wide range of alternative specifications, some of which we discuss in Section IIIB.

D. Refining the Test for Giffen Behavior

The discussion in Section I highlighted a nuanced prediction of the standard consumer model in light of subsistence concerns. The poorest of the poor should have a negative price elasticity, the poor-but-not-too-poor may have a positive price elasticity, i.e., exhibit Giffen behavior, and the relatively wealthy should once again have a negative price elasticity. Although our primary concern in this paper is documenting the presence of Giffen behavior, we are also interested in testing this broader prediction of the theory.

Unfortunately, classifying households or individuals directly by consumption zone is not possible. Not only is there no consensus on what constitutes a subsistence level of calories, but any such threshold would certainly vary widely by age, sex, height, weight, body fat, muscle composition, level of physical activity, health status, and a range of other factors. As a result, although we can compute caloric intake for each individual, identifying whether specific individuals are below, near, or above their subsistence level of caloric requirements is not possible. For the

²⁰ While we also gathered data on food purchases and expenditures, actual daily intake is likely to be a better measure of consumption or demand. This is due to the fact that food is storable, purchases are lumpy or infrequent, and households' recall of what they ate the day before the survey is likely to be more accurate than recall of purchases over the past month.

²¹ We prefer the arc-percent-change specification over the simple percent change because the subsidies represent large changes and because the arc formulation has the desirable property of being symmetric over time. The results are largely unchanged, however, if we use the simple percent change instead.

same reason, it is not possible to define these regions based on income or expenditure; individuals with different characteristics will require different amounts of expenditures or income to achieve nutritional sufficiency. Any such cutoffs would be imperfect, including some people who, because of high weight or activity levels, are unable to achieve maintenance nutrition with the specified income, and excluding others who have lower than expected nutritional (and thus income) needs because of small stature or low activity levels.

The method for classifying households we employ is based in the theory. Those who are so poor that they cannot achieve subsistence nutrition will consume a very high proportion of their food in the form of the staple good, regardless of size and activity level. Thus, splitting the data by the pre-intervention or initial share of caloric intake from consumption of the staple (Initial Staple Calorie Share, or ISCS) provides a more direct measure of whether a consumer or household is well off enough that they could, potentially, exhibit Giffen behavior. Consider panel B of Figure 1. Along the x-axis, 100 percent of calories come from the staple, while along the y-axis, the share is zero. In between these extremes, the share of calories from the staple is constant along rays from the origin, with the share decreasing monotonically as the slope of the ray increases. This provides a method of identifying whether a consumer is in the calorie-deprived, subsistence, or standard zone based on the share of calories received from the staple good. Consumers in the calorie-deprived zone will have a high ISCS, consumers in the standard zone will have a low ISCS, and consumers in the subsistence zone will have an intermediate ISCS.

While just using ISCS does not overcome the problem of identifying the exact threshold cutoff for moving from the calorie-deprived to subsistence zones, the advantage of this measure is that it is more “need neutral,” in that it normalizes for individual differences in caloric requirements. The measure also captures the simpler idea that if a household is so poor that it does not consume any of the fancy good, it cannot respond to a price increase by consuming less of it. While the ISCS may not be a perfect indicator of whether a household is near the subsistence zone (because of unobserved taste variation, for example), we believe it to be superior to other available measures.²²

Exploratory calculations using a simplified version of a minimum-cost diet problem (see, for example, the discussion in Robert Dorfman, Samuelson, and Robert M. Solow 1958) for China suggest that the ISCS associated with a minimum-cost, nutritionally sound diet (designed to ensure adequate consumption of calories and protein, and consisting of rice or wheat and bean curd) is much less variable than either required calories or required expenditure (details provided in the online Appendix). We compute the minimum cost diet for a range of weight/age/gender/activity level combinations, and find that the ISCS associated with the minimum-cost, nutritionally sound diet ranges between only 0.79 and 0.86 in Hunan and 0.78 to 0.85 in Gansu. Consumers or households wealthy enough to be consuming a diet with a lower ISCS would seem to be those who could, in principle, exhibit Giffen behavior. In light of this, our baseline specification splits households based on whether their ISCS is less than 0.80 (this corresponds

²² The broad conclusions of our analysis hold if we instead use, say, staple budget shares to parse the data (see the online Appendix). We believe, however, that the ISCS is a more appropriate measure. First, expenditure data are notoriously noisy, especially due to large but infrequent purchases such as durable goods that can skew budget shares. By contrast, the diets of households in our sample rarely contain more than three or four items and typically do not vary from day to day. Second, as shown, ISCS is not very need dependent, whereas budget share thresholds will vary considerably by household, due to differing housing, health care, education, and nutritional needs. Further, while some expenditures such as entertainment are highly discretionary, others categories such as housing, health care, and utilities are much less so. Thus, unlike the fairly precise ISCS cutoffs derived, clean cutoffs based on budget shares are difficult to derive, since the amount of truly discretionary income is difficult to measure. These thresholds would therefore contain many classification errors. For example, households with school-aged children would appear “richer” by the budget share measure because of expenditures on school fees; however, these expenditures are not really discretionary, and in fact make the household less “wealthy” than an identical household not facing these fees.

adds changes in income (earned and unearned) and household composition. Controlling for these other variables will help absorb any residual variation, and isolate the “pure” price effect of the intervention, as opposed to any behavioral effects the intervention may have on household size or either source of income (though in regressions for both provinces and for all population subgroups, the effect of the subsidy on these other variables is small and not statistically significant, suggesting the treatment had no such behavioral effects). Adding these other control variables changes the results only very slightly, increasing the coefficient and improving precision. While this coefficient is statistically significant only at the 10 percent level, it provides our first suggestive evidence of Giffen behavior in Hunan. As would be expected for households exhibiting Giffen behavior, the income effect is negative for unearned income, confirming that rice is an inferior good. The point estimate of the elasticity of unearned income is small, though there is likely to be significant measurement error in this variable, biasing the coefficient toward zero.²⁴

As we have emphasized, however, Giffen behavior is likely to be exhibited only by a specific subset of the poor. Therefore, in columns 3 through 6 we refine the test by parsing the data according to the theory, separating households by whether their preintervention staple calorie share suggests they are likely to be too poor to purchase something other than rice. For the group consuming at least some substantial share (20 percent) of calories from sources other than rice (column 3), i.e., the poor-but-not-too-poor, we find very strong evidence of Giffen behavior. A 1 percent price increase causes a 0.45 percent increase in consumption, and the effect is statistically significant at the 1 percent level (and little changed by adding in the other control variables). Thus, as theorized by Marshall and others, when faced with an increase in the price of the staple good, these households do, indeed, “consume more, and not less, of it” (Marshall 1895).

By contrast, but again consistent with the theory, the group consuming more than 80 percent of their total calories from rice (i.e., those still largely unable to consume meat), respond in the opposite direction (columns 5 and 6), with a large decline in rice consumption. Since these households consume essentially only rice, they have no choice but to respond to an increase in the price of rice by reducing demand. Thus, beyond finding evidence of Giffen behavior, the results also provide initial support for the subsistence model underlying such behavior. We find Giffen behavior where the theory predicts it, and downward sloping demand elsewhere. We explore the subsistence model further in Section IIIC below.

B. Robustness

The finding of Giffen behavior is robust to a wide range of alternate specifications, shown in Table 4. Since including the change in household size or either source of income rarely makes more than a marginal difference on our estimates of the price elasticity, for conciseness of presentation we show only the results with these additional control variables included. Columns 1 to 3 present results from a log-log specification, regressing the change in the log of household rice consumption on the change in the log of the net-of-subsidy price of rice and changes in the logs of the other control variables. The results again reveal Giffen behavior for households consuming less than 80 percent of their calories from rice, and downward sloping demand for those above this threshold. The point estimates of the elasticities are much greater here than for the arc percent changes in Table 3. However, this difference is largely attributable to the greater weight given to very low values with a log specification; for example, if we trim just the lowest

²⁴ The coefficient on earned income is positive (though also small); however, since greater caloric intake may improve productivity and earnings (Thomas and Strauss 1997), especially among those with very low nutritional status, this coefficient may be biased due to endogeneity. Unfortunately, we lack convincing instruments for changes in earned income. Dropping this variable does not change the results.

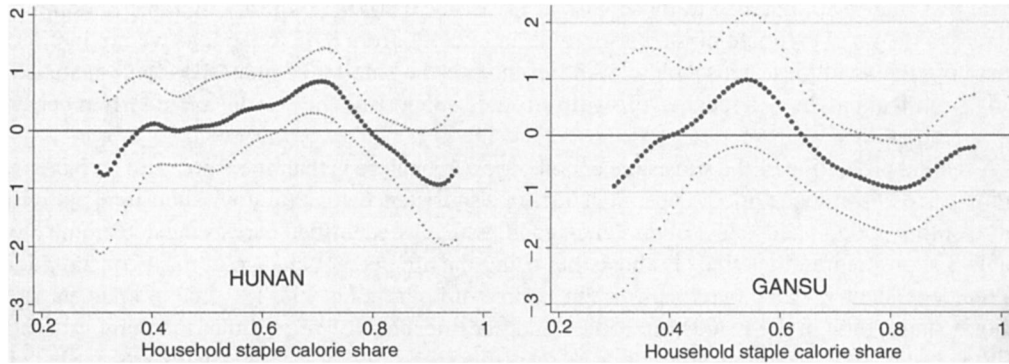


FIGURE 2. COEFFICIENT PLOTS

academics and policymakers. We have already seen that consumers with very high staple calorie shares (i.e., the poorest-of-the-poor) do not exhibit Giffen behavior. In addition, the model also predicts that once consumers are wealthy enough to pass beyond the subsistence zone into the standard consumption zone, staple demand should once again slope downward; in effect, we predict an inverted-U shape, with downward sloping demand (negative coefficients) for low and high values of staple calorie share, and Giffen behavior (positive coefficients) for intermediate values.²⁸ As stated, unlike the 80 percent calorie share, it is not possible to define a threshold beyond which households are likely to be in the standard or normal consumption zone, nor are we even guaranteed that our sample of the urban poor contains any such households. We therefore take a simple, flexible approach using a series of locally weighted regressions. At each staple calorie share point from 0.30 to 0.95 (there are few observations below 0.30 or above 0.95), we estimate equation (1) using a window of staple calorie shares of 0.05 on either side of that point; within that window we estimate a weighted regression, where observations closest to the central point receive the most weight (we use a biweight kernel, though the results are robust to alternatives). Figure 2 plots the resulting coefficients on the arc percent price change variable (i.e., the price elasticity) at each initial staple calorie share point for Hunan, along with the associated 95 percent confidence interval. The basic inverted-U shape in staple calorie share is clear. The elasticity is negative for the lowest and highest staple calorie shares, and positive in between. The Giffen range, where the point estimate of the elasticity is positive, reaches from 0.53 to 0.84 (which includes nearly two-thirds of the Hunan sample) though it is only statistically significant from 0.63 to 0.75. The peak of the curve reaches an elasticity of 0.85, at a staple calorie share of 0.70. And the threshold at which the elasticity turns negative is 0.80, which matches surprisingly well our simple minimum cost diet calculation. In general, the precision of these estimates is lower than those observed in Tables 3 to 5, since here we are restricting each regression to a band of ± 0.05 around a particular point, which reduces the sample size.

Not only does this figure support the theory in that Giffen behavior is most likely to be found among a range of households that are poor (but not too poor or too rich), it also guides us to a particular range when theory cannot provide a specific set of thresholds, as with the threshold between the subsistence and normal consumption zones. In particular, this curve suggests we restrict the range in which we test for Giffen behavior not only to those with a staple calorie

²⁸ Though, if we do not have enough households wealthy enough to fall into the normal consumption zone, we expect that the coefficients should at least decline as staple calorie share declines.

share less than 0.80, but also to those with at least, say, 0.60. In column 7 of Table 3, doing so increases the point estimate of the elasticity dramatically, from 0.47 to 0.64, as we are in effect removing the wealthiest households.²⁹ And even with the smaller sample, the effect is statistically significant at the 1 percent level, again strongly supporting the conclusion of Giffen behavior in Hunan.

A second prediction of the subsistence model we can explore is that, in response to an increase in the price of the staple good, consumers facing a subsistence constraint will consume not only more of that good, but also less of the fancy good, which we identified here as meat. Column 8 of Table 3 shows regressions like (1) above, but using the arc percent change in meat consumption as the dependent variable (we focus on the sample of households with less than 80 percent rice calorie share, though the results are robust to other thresholds). We find that the point estimate of the elasticity of meat consumption with respect to the price of rice is negative as predicted, though it is not statistically significant. However, one limitation of this analysis is that in Hunan, only about 45 percent of households reported meat consumption.³⁰ Therefore, in column 9 we focus on households that consume at least 50g of meat per person in round 1, which is still a very modest amount.³¹ Here, the results are more evident; a 1 percent increase in the price of rice leads to a large (1.13 percent), statistically significant decrease in meat consumption, as predicted by the model.

Thus, again, while our primary goal was to document the existence of Giffen behavior, these two results (the inverted-U shape of the response of rice consumption to a change in price and the decline in meat consumption in response to a change in the price of rice) support the subsistence model of consumption with a staple good and a taste-preferred but more expensive source of calories (such as meat) outlined above.

D. Gansu

As shown in Table 2, wheat-based foods (primarily buns, the simple bread *mo*, and noodles), are the staple good in Gansu. However, not all wheat-based foods are made at home from flour; most notably, noodles are often either consumed at restaurants or road-side food stalls, or purchased from shops as a prepared or packaged food. Since the subsidy we provided applied only to the purchase of wheat flour, for our analysis we use only the consumption of wheat foods typically produced at home from flour.³² And, as suggested by the calculations in the online Appendix, because there is some consumption of these other forms of wheat, our threshold staple calorie share for Giffen behavior based on wheat flour alone is closer to 0.70.³³ Table 6 presents the main results. In contrast to the case of Hunan, the coefficient is negative for the full sample in column 1. Even focusing on those below the staple calorie share threshold of 70 percent (column 2), while the coefficient is positive, it is extremely small and not statistically significantly different from zero. In addition, there is no evidence that wheat is even an inferior good in these cases.

²⁹ This coefficient differs slightly than the peak coefficient in the figure since the latter arises from a weighted regression, with more weight assigned to points closer to the peak of the curve.

³⁰ Although we condition on the staple calorie share in our regressions, the residual is not simply calories from meat.

³¹ While it may seem natural to have run all the specifications above stratifying based on meat consumption rather than staple calorie share, the latter is more general and does not rely on our ability to specifically identify meat as the (only) fancy good.

³² Over 90 percent of the consumption of wheat-based foods in Gansu was reported as "wheat flour," with most of the remainder reported as noodles. However, we cannot rule out that some noodles were made at home from flour but recorded as noodles, or that some consumers mistakenly reported purchased bread as wheat flour.

³³ Alternatively, we could use a staple calorie share of 0.80 based on consumption of all wheat foods, rather than just those produced at home from flour.

little consumption of the fancy good, meat.³⁴ As shown in Table 2, households in Gansu receive on average only 1 percent of their calories from meat, which is even less than the 7 percent observed in Hunan; further, only one-quarter of households reported any meat consumption in our first period consumption diary. The bulk of nonstaple calories comes largely from vegetables (especially potatoes, which themselves may potentially be a staple food) and oil, neither of which is likely to be considered a fancy good. With little consumption of the fancy good, it is perhaps not surprising that most households do not behave like Giffen consumers in Gansu. There is simply no way for them to finance additional purchases of rice by reducing meat, since they are consuming almost no meat to begin with.³⁵ This also suggests that the best place to find Giffen behavior is among those consuming a nontrivial amount of meat. Therefore, in the next-to-last column of Table 6, we consider only households that consume at least 50 grams of meat per person in the initial period. Though the sample shrinks considerably because meat consumption is so uncommon, we do find evidence of Giffen behavior among this group, with a 1 percent increase in the price of wheat causing a 1.3 percent increase in wheat consumption.

Gansu also departs from the ideal conditions for Giffen behavior in that wheat as a staple is consumed in a number of other forms that may act as substitutes for each other, many of which are not made directly by consumers at home from wheat flour. Unfortunately, our experimental design failed to account for this additional complexity.³⁶ In Hunan, the staple good, rice, is consumed typically only in its basic form. By contrast, in Gansu wheat is consumed as *mo* and buns made at home, plus noodles, and other wheat-based, prepared foods like bread, biscuits, or deep-fried dough purchased from shops or food stalls. While Table 2 showed that average pretreatment wheat consumption per capita in Gansu was 344g, typically about 34 grams, or 10 percent, of that wheat is from items other than *mo* or buns. If a household consumes their staple food in many forms and the price of one increases, they may not need to engage in Giffen behavior because they can reduce consumption of that one and increase consumption of the other, substitutable forms of the staple that did not experience the price increase. While this is unlikely to happen often in reality because the price of all the forms of the staple will be linked to the price of the raw ingredient (here, wheat), the unique structure of our subsidy did just that, subsidizing only the form of the staple prepared at home, and not the close substitutes purchased in stores. This may both explain why we do not find widespread evidence of Giffen behavior in Gansu, and also suggests we might find such behavior if we focus on those households where consumption of these other forms of wheat is small or zero in the initial period.³⁷ The final column of Table 6 provides some suggestive evidence of this possibility, focusing on the condition that the household consumes less than 50g of these alternative forms of wheat. Among this group there is again statistically significant evidence of Giffen behavior, with a very large elasticity.

Overall, then, while the results for Gansu do not yield as evident, robust evidence of Giffen behavior as was found for Hunan, we believe this is most likely due to our failure to recognize *ex ante* that for the majority of households in our sample, diets do not conform to two of the basic

³⁴ This result was unanticipated, since the northern provinces in our original paper (Jensen and Miller 2002), and our field test of the survey for the current study, revealed considerably more meat consumption in Gansu.

³⁵ While there is some consumption of pulses and, to a lesser extent, dairy, these goods are also unlikely to be regarded as fancy goods in the way that meat is, since most households turn to these goods only when they cannot afford meat. Further, there is no way to cut back consumption of these foods while maintaining protein intake; with meat, households can reduce consumption but switch to pulses as a less expensive source of protein.

³⁶ Though in selecting sample sites, the authors personally visited only two of the counties in Gansu (Anding and Yuzhong). These counties, both with significant Muslim populations who traditionally consume primarily the home made bread *mo*, fit the pattern better, with 88 percent of all wheat consumption coming from flour, compared to 74 percent in the other three counties. If we limit our analysis to just these two counties, we find a positive coefficient for all staple calorie share thresholds, though due to the smaller samples, the coefficients are not statistically significant.

³⁷ Some of this variation is geographic or based on religion, as noted above.

conditions under which we predict Giffen behavior (consumption of a fancy good, and a staple good for which there are no close substitutes). When we restrict our sample to take these factors into consideration, we do find evidence of Giffen behavior, though the samples are smaller, precisely because most households do not conform to the conditions in Gansu. It is possible that if we sampled a slightly wealthier group of households that consume more of the fancy good, and perhaps altered our experimental design (e.g., to subsidize all wheat foods, not just wheat flour), we might find stronger evidence of Giffen behavior.

E. Addressing Potential Alternative Explanations for the Results

The analysis thus far provides robust evidence of upward sloping demand in Hunan, with somewhat weaker evidence in Gansu. However, two explanations other than Giffen behavior need to be explored. First, the vouchers themselves may have induced a behavioral effect. For example, households may have increased demand for the staple because they interpreted being given vouchers as a signal of its value (e.g., that it had significant health benefits). Alternatively, the vouchers may have increased the salience of the staple, or households might have felt they should eat more of it in order to take advantage of the subsidy before it ran out. However, in these cases we would expect the vouchers to increase consumption, the opposite of what is observed as Giffen behavior. Alternatively, and perhaps less likely, households may view the vouchers as providing adverse information about the staple good; for example, they may view the attempt to sell more rice as an indication that there is something wrong with the current stock, in which case they might want to consume less of it (though consumers were told the subsidies were being provided by outside researchers rather than merchants, farmers, or the government). But since the effects vary by the staple calorie share, to explain our results it would have to be that the vouchers had a salience or signal effect only for some subset of households based on their calorie share, which seems unlikely.³⁸

A second concern is the possibility that households cheated,³⁹ for example by swapping vouchers for cash instead of using them to purchase the staple,⁴⁰ or by reselling rice or wheat purchased with the vouchers at a higher price. In the extreme case where a household simply sells all their vouchers, the program would be a pure wealth transfer, and consumption of inferior goods like rice or wheat would be expected to decline even though their effective price had not changed. In less extreme cases where only some vouchers are sold, even at less than face value, the program would still provide a pure wealth transfer that could in itself decrease consumption of the staple, which we might again misinterpret as Giffen behavior.⁴¹

³⁸ If the consumption of the treatment groups responded both to having received any subsidy at all (i.e., the signal or salience effect) and to size of the subsidy received, we could eliminate the former and identify the elasticity off of the size of the subsidy alone by running regressions excluding the control group. Doing so yields similar results, including continued evidence of Giffen behavior. For example, for Hunan, the price elasticity for the full sample is 0.33(0.22); for ISCS less than 0.80 it is 0.61(0.25), and for ISCS greater than 0.80 it is $-0.75(0.44)$. Similar results hold for Gansu, though, as above, the estimates are less precise and in some cases we cannot reject zero even though the point estimates have the correct sign. These results indicate that the Giffen effect is not driven by some common signaling or salience effect among the treatment groups. However, it is of course possible that larger subsidies create stronger signaling effects, so these results do not imply there were no such effects at all.

³⁹ Cheating, where shopkeepers do not provide the full subsidy to consumers (for example, those with poor math skills), effectively lowers the value of the subsidy, so the Giffen behavior we find would likely have been even stronger had such cheating not occurred.

⁴⁰ Most shopkeepers sold only grain, so most households could not have exchanged the vouchers for other foods.

⁴¹ If households bought rice at subsidized prices on behalf of (or as a gift to) their friends or relatives but did not make a profit from doing so, this does not affect the households' wealth and thus does not bias our experiment. We of course cannot rule out that treatment households shared some of their rice with others. However, this "leakage" would tend to diminish the likelihood of finding Giffen behavior, since the effective value of the subsidy is reduced.

Preventing cashing out of the vouchers or resale was one of our primary concerns in designing the intervention. In doing so, however, we also wanted to ensure that the process of redeeming the vouchers would be as much like an ordinary market transaction as possible, and to keep the administrative burden of the intervention manageable. In addition, we wanted to allow for the fact that a natural reaction to receiving access to discounted rice or wheat would be for households to build up their stores of these goods, which would look very similar to cashing out (i.e., the number of vouchers redeemed exceeds the amount of rice or wheat consumed).

With these concerns in mind, a number of safeguards were built into the experimental design. First, the consent scripts given to treatment households stated they were explicitly prohibited from selling the vouchers or the rice or wheat bought with the vouchers. They were also told that they would be monitored and that violators would be dismissed from the program. Our native Chinese implementation team, which was very familiar with the population from which our survey households were drawn, felt that in light of such a rule households would be very unlikely to cash out the vouchers.

Second, the vouchers were distributed only incrementally (one-fifth each month) over the course of the intervention (though all vouchers remained valid until the end of the intervention). As a result, households engaging in early cashing out would be limited in their ability to do so (since they would have only a small part of the vouchers in hand) and would risk losing the value of all future vouchers if they were caught. For our purposes, the crucial question is whether there was cashing out before the second round of the survey, since this is the only round for which the subsidy was in effect. At the time of the second survey, a significant amount of the benefit of the program still lay in the future, which reduced the incentive to cash out.

Third, in order to ensure that they were used only by treatment households and not resold, the vouchers (which were printed in Beijing in multicolor ink and bore a special stamp, making them difficult to counterfeit in the survey regions) had to be signed by the redeemer at the time they were used. Shopkeepers had an incentive to ensure the vouchers were used only by the appropriate households, since they were reimbursed only after the vouchers and signatures had been validated by our managers.⁴² Finally, shopkeepers were also given incentives to prevent cashing out through a lump-sum payment given at the end of the intervention if they were found to have complied with its guidelines, including preventing resale and cashing out vouchers.

The safeguards discussed above were accompanied by monitoring and auditing to check for compliance. These audits did not discover any such cheating, and our survey personnel, who visited the households, did not discover evidence of cashing out. Finally, comparing voucher redemptions relative to estimated consumption and storage suggests that to the extent it occurred at all, cashing out could not have been significant or widespread (details in the online Appendix).

IV. Conclusion

We find strong, clear evidence of Giffen behavior among poor households in Hunan, China, and somewhat less robust evidence in Gansu. To the best of our knowledge, this is the first rigorous, real-world empirical evidence of Giffen behavior. It is ironic that despite a long search, in sometimes unusual settings, we found examples in the most widely consumed foods for the most populous nation in history. However, the examples were found exactly where theory would predict they should occur: impoverished (but not too impoverished) consumers, heavily dependent on a staple good, with limited substitution possibilities. And while our experiment focused on two areas of China, our framework and analysis suggest that Giffen behavior may not be rare at

⁴² Our Chinese management team was the residual claimant on the value of unredeemed vouchers, and so they, themselves, had a strong incentive to enforce the rules of the intervention and prevent cashing out.

all. To begin, the socioeconomic class we are studying—the urban poor in China—is extremely large, representing an estimated 90 million individuals. Thus, to the extent that our experiment is representative of that population, in terms of number of people, Giffen behavior is probably quite common. Second, our empirical approach differs in important ways from previous attempts to identify Giffen behavior. Earlier approaches focused on market-level price and quantity data, which for many reasons are unlikely to exhibit Giffen properties. Our approach uses individual-level data and focuses on parsing the data by an appropriate measure of wealth. These techniques have never been applied to other likely candidates for Giffen behavior. We believe there is nothing unusual or unique about the Chinese case; anecdotal evidence suggests that similar dietary patterns to those observed here are found among much of the world's poor (with staples including maize, millet, sorghum, and cassava, in addition to rice and wheat).⁴³ Thus, it may be better to interpret our study as the first evidence of Giffen behavior rather than the only possible evidence. In fact, it is even possible that if the right data were available, the claims by Marshall and Samuelson that bread and potatoes were Giffen would be verified.

These results have important policy implications. The most obvious is that the patterns of substitution away from calorie-dense foods such as rice or wheat implied by Giffen behavior and observed in the data suggest that programs designed to improve nutrition may have little impact for subsistence-zone consumers; this is particularly important given that nearly two-thirds of our sample of urban poor households fall in this zone. While the subsidy would still have a positive impact on utility for these households, to the extent that the aim of the policy is to improve nutrition, it may have a much weaker effect than intended (theoretically, it is possible they will have no effect at all, or even reduce nutrition). A full, detailed theoretical and empirical analysis of the nutritional implications of price subsidies is beyond the scope of the present paper, and is instead addressed in a companion paper, Jensen and Miller (2008).

In identifying the relevance of the three consumption zones, the results also point to an important nonlinear effect of programs designed to improve nutrition. Such programs are unlikely to significantly improve nutrition among subsistence households unless they effectively move them out of the subsistence zone, since within this zone households will spend additional income on improving the taste of their meals, rather than their nutritional content. Only if the size of the benefit is large enough to move a subsistence-zone household into the standard zone will it substantially improve nutrition. Thus, programs that appear not to improve nutrition might have a significant effect on these households if their benefits were increased. However, the same nonlinearities have the opposite implications for households in the calorie-deprived zone. For these households, increasing the value of the benefits will improve nutrition only up to a point, after which some households will be in the subsistence zone and the nutrition improvement will level off, and possibly even decline. This creates a challenge for program targeting, as these nonlinearities suggest programs will be effective only if the size of the subsidy is tailored to the household's consumption zone. From a program evaluation perspective, this leveling off or decline in calories could make the average effect of a program appear small, masking larger changes among specific groups unless the heterogeneous response is taken into consideration.⁴⁴ Thus, recognizing the heterogeneity of poor households is important for the design, evaluation, and targeting of such programs. It is also important to note that the difficulties facing subsidies

⁴³ There is ample quantitative evidence of these patterns as well. For example, Shankar Subramanian and Deaton (1996) find that the poorest households in rural Maharashtra, India, on average receive 53 percent of their calories from sorghum (and spend on average 20 percent of their total household budgets on this food).

⁴⁴ In fact, the two effects could offset each other from an evaluation perspective; relatively small benefits improve nutrition among the calorie-deprived but decrease it among subsistence consumers, while larger benefits begin to reduce calories among the formerly calorie-deprived but increase them among the formerly subsistence consumers. Thus, the heterogeneous response could mask significant gains among subsets of consumers.

or price controls would apply to almost any other program whose goal is to improve nutrition. In particular, any program that results in an increase in real wealth, whether cash transfers, employment schemes, in-kind transfers, or food stamps, will generate similar substitution effects.

The results are also important in light of the large literature concerned with the income elasticity of demand for calories (see John Strauss and Duncan Thomas 1995; Deaton 1997). Those policymakers who wish to increase nutrition but believe the income elasticity is small might be inclined to suggest, as alternatives to cash transfers, subsidies or price controls that directly encourage increased consumption of nutritious foods. Similar policy prescriptions may arise if there is a concern, often stated, that simply giving cash is not desirable because households may spend it on other luxuries (food or nonfood). However, our results suggest even these price-based policies will face similar difficulties by virtue of the large wealth effects they create.

It is also worth noting that the results present somewhat more positive news about household vulnerability. Policymakers are often concerned that households highly dependent on staple goods may experience nutrition declines when the prices of those goods increase, either transitively during a crisis or permanently with the liberalization of markets, or as a result of trade agreements, and a number of programs are designed to protect against this very possibility. Our results suggest that in fact households in the subsistence zone are able to buffer their staple consumption against such changes quite well (unless perhaps the price changes are very large). However, there may still be a value in offering such assistance to households in the calorie-deprived zone, or perhaps even households in the subsistence zone (though the justification would have to be on general welfare, rather than nutritional, grounds). But here, again, the heterogeneity of consumers is important to take into consideration.

Finally, although historically the primary importance of the Giffen phenomenon has been pedagogical, we believe it also serves a more fundamental role in economic theory. The neoclassical model of the consumer is one in which the consumer maximizes stable preferences subject to a budget constraint, and in recent years this model has come increasingly under attack. These objections run from the simplistic, "people don't maximize" arguments to the sophisticated criticisms found in psychology and behavioral economics. The possibility of Giffen behavior is a clear, complex, nuanced prediction of the neoclassical model. Where naïve intuition suggests that consumers should respond to a price increase by consuming less of the good in question, consumer theory suggests that a sophisticated consumer, mindful of the interplay between nutritional and budgetary concerns, might increase consumption. Further, the theory tells us exactly when and where we should expect to find Giffen behavior, and where we should not. To our knowledge, no other theory predicts the pattern of behavior implicit in the Giffen phenomenon, i.e., the sign of the price elasticity changing from negative to positive and back to negative as very poor consumers' wealth increases. Thus, while economists' failure to document Giffen behavior in the past has been interpreted as a criticism of the approach, our finding of Giffen behavior provides a type of vindication. Giffen behavior is predicted by the neoclassical model and no other. While psychological and behavioral theories help to account for some areas of economic behavior, in the case of the Giffen phenomenon, and of the consumption behavior of the extremely poor more generally, the standard model appears to be the right one.

REFERENCES

- Baruch, Shmuel, and Yakar Kannai.** 2001. "Inferior Goods, Giffen Goods, and Shochu." In *Economics Essays, A Festschrift for Werner Hildenbrand*, ed. Gerard Debreu, Wilhelm Neuefeind, and Walter Trockel, 9–17. Heidelberg: Springer-Verlag.
- Barzel, Yoram, and Wing Suen.** 1992. "The Demand Curves for Giffen Goods Are Downward Sloping." *The Economic Journal*, 102(413): 896–905.

- Battalio, Raymond C., John H. Kagel, and Carl A. Kogut.** 1991. "Experimental Confirmation of the Existence of a Giffen Good." *American Economic Review*, 81(4): 961–70.
- Boland, Lawrence A.** 1977. "Giffen Goods, Market Prices and Testability." *Australian Economic Papers*, 16(28): 72–85.
- Bopp, Anthony E.** 1983. "The Demand for Kerosene: A Modern Giffen Good." *Applied Economics*, 15(4): 459–67.
- Chen, John-Ren.** 1994. "The Effects of Land Reform on the Rice Sector and Economic Development in Taiwan." *World Development*, 22(11): 1759–70.
- Deaton, Angus.** 1997. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. Baltimore: Johns Hopkins University Press for the World Bank.
- DeGrandpre, R. J., Warren K. Bickel, S. Abu Turab Rizvi, and John R. Hughes.** 1993. "Effects of Income on Drug Choice in Humans." *Journal of the Experimental Analysis of Behavior*, 59(3): 483–500.
- Dorfman, Robert, Paul A. Samuelson, and Robert M. Solow.** 1958. *Linear Programming and Economic Analysis*. New York: McGraw Hill.
- Dougan, William R.** 1982. "Giffen Goods and the Law of Demand." *Journal of Political Economy*, 90(4): 809–15.
- Dwyer, Gerald P., Jr., and Cotton M. Lindsay.** 1984. "Robert Giffen and the Irish Potato." *American Economic Review*, 74(1): 188–92.
- Friedman, Jed, and James Levinsohn.** 2002. "The Distributional Impacts of Indonesia's Financial Crisis on Household Welfare: A 'Rapid Response' Methodology." *World Bank Economic Review*, 16(3): 397–423.
- Gilley, Otis W., and Gordon V. Karels.** 1991. "In Search of Giffen Behavior." *Economic Inquiry*, 29(1): 182–89.
- Halbrendt, Catherine, Francis Tuan, Conrado Gempesaw, and Dimphna Dolk-Etz.** 1994. "Rural Chinese Food Consumption: The Case of Guangdong." *American Journal of Agricultural Economics*, 76(4): 794–99.
- Jensen, Robert T., and Nolan H. Miller.** 2002. "Giffen Behavior in China: Evidence from the China Health and Nutrition Survey." Harvard University, John F. Kennedy School of Government Faculty Research Working Paper RWP–02–014.
- Jensen, Robert T., and Nolan H. Miller.** 2008. "Do Consumer Price Subsidies Really Improve Nutrition?" Harvard University, John F. Kennedy School of Government Research Working Paper RWP–08–025.
- Koenker, Roger.** 1977. "Was Bread Giffen? The Demand for Food in England Circa 1790." *Review of Economics and Statistics*, 59(2): 225–29.
- Marshall, Alfred.** 1895. *Principles of Economics*. London: Macmillan.
- McDonough, Terrence, and Joseph Eisenhauer.** 1995. "Sir Robert Giffen and the Great Potato Famine: A Discussion of the Role of a Legend in Neoclassical Economics." *Journal of Economic Issues*, 29(3): 747–59.
- McKenzie, David.** 2002. "Are Tortillas a Giffen Good in Mexico?" *Economics Bulletin*, 15(1): 1–7.
- Nachbar, John H.** 1998. "The Last Word on Giffen Goods?" *Economic Theory*, 11(2): 403–12.
- Ravallion, Martin.** 2007. "Geographic Inequity in a Decentralized Anti-Poverty Program: A Case Study of China." The World Bank Policy Research Working Paper Series 4303.
- Rosen, Sherwin.** 1999. "Potato Paradoxes." *Journal of Political Economy*, 107(6): S294–313.
- Samuelson, Paul.** 1964. *Economics*. 6th ed. New York: McGraw Hill.
- Stigler, George J.** 1947. "Notes on the History of the Giffen Paradox." *Journal of Political Economy*, 55(2): 152–56.
- Stigler, George J.** 1966. *The Theory of Price*. 3rd ed. New York: The Macmillan Company.
- Strauss, John, and Thomas Duncan.** 1995. "Human Resources: Empirical Modeling of Household and Family Decisions." In *Handbook of Development Economics, Volume 3A*, ed. Jere R. Behrman and T. N. Srinivasan, 1183–2023. Amsterdam: North-Holland.
- Subramanian, Shankar, and Angus Deaton.** 1996. "The Demand for Food and Calories." *Journal of Political Economy*, 104(1): 133–62.
- Thomas, Duncan, and John Strauss.** 1997. "Health and Wages: Evidence on Men and Women in Urban Brazil." *Journal of Econometrics*, 77(1): 159–85.