Feeding America and the Second Welfare Theorem

Canice Prendergast*

December 8, 2017

*University of Chicago Booth School of Business. Email: canice.prendergast@chicagobooth.edu.
I am grateful to Carole Theus, Harry Davis, Mike Halligan, Melanie Nowacki and Bill Thomas for much help on this. Many thanks to Paul Kim and Wendy Wong for excellent research assistance, and to Sam Altmann, Ben Roth, Scott Kominars, Alex Teytelboym, and participants at the HCEO conference for helpful comments.

1 Introduction

The Second Welfare Theorem is wondrous advertising for allocating resources through the price system. By appropriately assigning budgets to consumers, it shows how a social planner can implement any Pareto Optimal outcome as a competitive equilibrium. In effect, efficiency and distribution can be simultaneously solved - not by messing with prices, or by avoiding market allocations - but solely through the redistribution of income.

Abstract as the Second Welfare Theorem may be, it is the (academic) shadow that looms over the problem faced by Feeding America. Every year, it receives donations of hundreds of millions of pounds of food from manufacturers, distributors, and retailers, and its task is to allocate them to over 200 food banks across the United States. Until 2005, it did so in a traditional way: it used a centralized assignment mechanism that gave all food banks the same amount and type of food per client. This had many merits: it was transparent, based on objective criteria, and was focused on getting most food to areas of seemingly greatest need. (The details are described in Prendergast, 2017.) Despite this, there was widespread dissatisfaction with it. Much of this was related to the absence of choice by the food banks: they were not getting what they would have chosen themselves. This was because Feeding America simply didn't know what they wanted most.

Given this, it is natural to ask whether there could be a way to garner the benefits of choice characterized by the welfare theorems. In early 2004, the author was part of a group charged with reconsidering how food should be assigned. At the end of a process that lasted over a year, Feeding America implemented the Choice System, where food banks bid each day on truckloads of food using a specialized currency that it designed. By most objective measures, the change has been a success. On the demand side, there is dispersion in outcomes across food banks in ways that are radically different from the old system. On the supply side, the amount of food in the system increased rapidly, and by a large amount.

The purpose of this work is not to rehash those details, which are described in Prendergast, 2017. Instead, we focus on how those on the ground quickly identified reasons why the benefits of a price based system could fail to materialize, and to describe the (many) design features of the Choice System aimed at mitigating such concerns. Specifically, when the possibility of a market based system was raised with other members of the group - most of its members were food bank directors - it was met with largely universal skepticism, or in some cases, hostility. This resistance was not because food bank directors did not recognize the value of being able to choose what they wanted. Instead, it was based on a fear that it would unfairly benefit some food banks (the "strong") to the detriment of others (the "weak"), thereby undermining the competitive equilibrium of our textbooks. Here is a narrative that would capture much of that concern:

> "So, how is this going to work? You have a large food bank like Los

> Angeles, which not only has warehouses of other food, a load of money, but also tons of workers on site. On the other hand, you have the Idaho Food Bank, which has very little other food, no money, and a couple of people in the building. Tell me how any system based on food banks paying for food, and using their limited personnel to do so, is not going to benefit Los Angeles relative to Idaho?"

Said another way, many food bank directors have chosen their careers to help those they feel have been disenfranchised by markets, and their radars were particularly sensitive to the possibility that they would put their imprimatur on a system that might do the same in their world.[1]

This narrative implicitly captures a number of obstacles to getting the kinds of benefits underlying the Second Welfare Theorem. First, *how can Feeding America get incomes right?* Los Angeles has much more money than Idaho. More generally, food banks vary enormously in their ability to fund raise. Yet in order to attain the desired Pareto Optimal outcome, each food bank needs a specific amount that it spends on food. This ultimately lead to the use of a "fake" currency used to buy food.

Second, *will some food banks work out how to play the game better than others?* Bidders in auctions sometimes do not understand their strategic environment (Kagel et al., 1987, Kagel and Levin, 1993), and do not bid according to our theories. A particular concern is that this strategic confusion would disadvantage smaller food banks. Relevant here is Hortacsu and Puller, 2005, who showed how smaller bidders were especially disadvantaged in the electricity auction market in Texas. Such issues about how strategic confusion causes inequities have also arisen in the market design literature, most notably in the work of Pathak and Sonmez, 2008, and Abdulkadiroglu, Pathak, Roth and Sonmez, 2006. They show that under the old Boston (non-strategy proof) school choice mechanism, some parents worked out how to manipulate the system while others did not.[2]

Third, *can food banks evaluate the merits of offerings correctly?* The Pareto optimal allocation of the welfare theorems assumes that consumers know what they are buying. Yet there are reasons why they may not. For example, Hastings et al., 2017, show how competition in the privatized Mexican social security system resulted in poorer citizens making considerably worse investment decisions than wealthier ones. In this context, assume that a small food bank is considering bidding for a truckload of peanut butter. It may have never bid for peanut butter before, may not know often loads of peanut butter are offered on the market, and isn't quite sure how much its money balances are worth. How does it construct a bid?

[1]Had these fears of inequality not been adequately resolved, it would likely have derailed the implementation of the Choice System. Feeding America is largely a democratic organization and was only willing the implement the change after a majority of the food banks voted in favor of it at their National Annual meeting.

[2]Similar evidence provided by Budish and Cantillon, 2012, shows how some MBA students at Harvard seemed better able to manipulate their old class assignment system than others.

Finally, *how can we guarantee that all food banks face similar prices?* Food banks ultimately care about food reaching its clients. However, bigger food banks have more clients. Economies of scale can imply that effective prices are lower for the larger food banks. The most example of this arises from the indivisibility of a truckload. Specifically, some food banks are too small to be able to use a full truckload of some foods, and may be disadvantaged as a result.[3]

These kinds of problems quickly became the central concern of the redesign exercise. To address these issues, the paper has three parts. We begin by offering a simple theoretical model to show how a decentralized market where food banks bid for food in a first price auction *can* give rise to the efficient outcome. Its purpose is not simply to show this, but rather to outline the key parameters for this to arise: the shadow price of currency, the fixed cost of a truck, and information about the location and slope of the demand curve. We then turn to describe how many of the features of the Choice System have been specifically aimed at reducing these problems. Finally, we offer empirical evidence on one of these - fractional or joint bidding - to show that it indeed seems to be operating as intended.

2 A Simple Model

We begin by offering a model with multiple goods (two for simplicity) and different kinds of food banks (again just two). In each period food arrives to the system, and must be assigned by Feeding America to a food bank. Its objective is to assign the food to the food bank whose marginal valuation is highest. After showing the optimal allocation, we then describe how it can be implemented through a bidding mechanism.

Food banks have demand for two goods, 1 and 2. There are n food banks (215 in reality), and are of two types k , big ($k = B$) and small ($k = S$). On any given day, $t = 1, ...\infty$, a per capita quantity of good j, q_{jt} needs to be assigned by Feeding America among n food banks, where per capita is relative to the entire set of food banks. Although donations are made every period, consumption occurs over a discrete "interval", which lasts T periods. (Think of this as a month or two.) Aggregate supply over the T periods is assumed to be deterministic: $Q_i = \sum_{t=0}^{T} q_{it}$ is the aggregate known per capita supply of good i. A food bank's utility is the sum of U_k over all intervals. There are an infinite number of these intervals.

There is no discounting. Food bank k's utility per client in the interval, U_k, is given by

$$U_k(.) = U_1(\sum_{t=1}^{T} q_{1kt}) - \kappa_{1k} D_{1k} + U_2(\sum_{t=1}^{T} q_{2kt}) - \kappa_{2k} D_{2k}. \qquad (1)$$

[3] As another example of how effective prices vary across consumers, Orhun and Palazzo (2017) document how the poor are less likely to buy many non-perishable goods in bulk than wealthier households, and effectively pay higher prices for the same products.

Let $U_i' > 0$, and $U_i'' \leq 0$.[4] The only unusual features of (1) are the $\kappa_{ik} D_{ik}$ terms. First, D_{ik} is a dummy variable equal to 1 if food bank k consumes good i, and 0 otherwise. κ_{ik} is a fixed cost per capita that food bank k incurs from consuming good i.[5] This term picks up both transportation costs and spoilage issues, and are included in the model to reflect the potentially higher returns to large food banks. Because large good banks have more consumers, the can use a full truckload of food more readily than smaller food banks, and need to store it for less time.

The objective of the principal in any interval is to maximize the sum of the expected utilities. Let $Q_{jk} = \sum_{t=1}^{T} q_{jkt}$ be the per capita food allocation over the T periods. Then if all food banks should consume both goods, efficiency is characterized by

$$U_j'(Q_{jS}) = U_j'(Q_{jB}) = U_j'(Q_j). \tag{2}$$

This is the usual condition of equating marginal rates of substitution across consumers, and determines the objective of the principal.

2.1 Decentralizing the Allocation

Given this, now consider how this can be decentralized through the price system: by a game where food banks bid for food. To do so, the principal allows food banks to bid on food using some - as yet unspecified - currency. Each food bank type k bids using this currency. To do so, they receive total income per capita over the T periods in that currency of y_k, $k = S, B$, where per capita here means for each food bank.

On each day t, they simultaneously place bids of b_{kjt} for a unit of good j with the possibility of an order limit (so, for example, they could bid \$2 for each unit up to 6 units of the good). If they win, they pay their bid - a first price auction - where the q_{jt} highest bids received the good for their bid price, and the other food banks receive nothing. If they lose, they retain their currency for the next period's auction. Ties are allocated randomly. At the end of period T, the principal reallocates all currency such that they receive and additional y_k above their period T balances for period 1 of the next interval.

We now turn to the bidding behavior of a food bank k. The concern here is that big food banks benefit relative to their smaller counterparts. To model this, we assume that all food banks of a type hold similar beliefs and act identically, but they could differ between the two groups. (This will be efficient in the equilibrium below.) Consider the case where all food banks choose to purchase both kinds of goods. This allows us to use standard marginal conditions to determine outcomes. Let the belief of food bank k of the probability of winning the good be given by $\Omega_k(b_k)$, where $\Omega_k'(b_k) \geq 0$. (We make this time independent for simplicity: this will be the case

[4] The absence of complementarities allows us to ignore the need for almost equal incomes, as in Budish, 2011.

[5] This could easily be modeled as a fractional spoilage cost, in which case it would lie inside the U terms, without changing the results.

in the equilibrium below.) Furthermore, let its belief of the marginal value (shadow price) of a unit of currency be given by λ_k, and its belief of the quantity of good i it will consume be $\hat{Q}_{ik}$.

Consider a stationary bidding strategy: $b_{ikt} = b_{ik}$ until the food bank has received its equilibrium allocation. Given this, the food bank chooses b^*_{ik} to maximize $\Omega_k(b_{ik})(U'_k(\hat{Q}_{ik}) - \lambda_k b_{ik})$, so that its optimal bid b^*_{ik} for a unit of good i will be characterized by

$$\Omega_k(b^*_{ik})\lambda_k = \Omega'_k(b^*_{ik})[U'_k(\hat{Q}_{ik}) - \lambda_k b^*_{ik}] \tag{3}$$

or

$$b^*_{ik} = \frac{U'_k(\hat{Q}_{ik})}{\lambda_k} - \frac{\Omega_k(b^*_{ik})}{\Omega'_k(b^*_{ik})}. \tag{4}$$

This should look familiar. First, marginal utility is normalized by the value of currency - which itself depends on other bidding opportunities. Because of diminishing marginal utility, this requires that the food bank hold an expectation of equilibrium consumption. Second, the bid is additionally shaded in the usual way by the slope of the demand curve. As this logic holds for both kinds of goods, one can then identify the ratio of the bids by

$$\frac{U'_k(\hat{Q}_{1k})}{U'_k(\hat{Q}_{2k})} = \frac{b^*_{1k} - \frac{\Omega_k(b^*_{1k})}{\Omega'_k(b^*_{1k})}}{b^*_{2k} - \frac{\Omega_k(b^*_{2k})}{\Omega'_k(b^*_{2k})}}. \tag{5}$$

This formulation of the problem allows us to see the value of having all actors on the same page, as seen in Proposition 1.

Proposition 1 *Assume that κ_{ik} is low enough so food banks buy both kinds of food. Then if $y_B = y_S = 1$, all food banks hold similar beliefs on $\frac{\Omega'_k(b^*_{ik})}{\Omega_k(b^*_{ik})}$, and their beliefs on quantities are correct: $\hat{Q}_{ik} = Q_i$, the efficient outcome is a Perfect Bayes Nash equilibrium of the bidding game.*

In words, this proposition shows how the Second Welfare Theorem can be implemented if all food banks (a) hold the same beliefs about the slope and level of the demand curve, (b) have the same income (here normalized to 1 per capita), and (c) have correct beliefs about how much food they will get. Its proof is trivial, because the bidding strategies that implement is are so simple. Specifically, market clearing prices are the solution to $\frac{U'_1(Q_1)}{U'_k(Q_2)} = \frac{p^*_1 - \frac{\Omega_k(p^*_1)}{\Omega'_k(p^*_1)}}{p^*_2 - \frac{\Omega_k(p^*_2)}{\Omega'_k(p^*_2)}}$ (which defines the ratio of the prices) and $p^*_1 Q_1 + p^*_2 Q_2 = 1$ (which determines price levels and ensures exact budget balance). Given that, the bidding strategies are (a) all food banks bid the market clearing price p^*_i until they receive Q_i and (b) $p^*_i - \epsilon$ beyond that.[6]

[6]Note also that the out of equilibrium behavior needed to hold this equilibrium behavior is far from perverse, as food banks are all bidding their marginal valuation of additional units of the good beyond Q_i. So, for example, if a food bank "forgot" to bid for food in period T, this would be the optimal strategy by any food bank.

The purpose of Proposition 1 is not primarily to show that actors who have equal access and information can implement the efficient outcome. Instead, it is to point to the key parameters that give rise to efficiency: λ_k, Ω_k, Ω'_k, $\hat{Q}_{ik}$, and κ_{ik} (which matters for food banks consuming both kinds of goods). Finding ways to align these is a central concern of the design features below.

Before turning to design features and the empirical evidence, it is worth noting in passing that the simplicity of the model abstracts from why we need choice over central assignment. Here all food banks consume the same quantities, so couldn't Feeding America just give everyone the same allocation? In reality, there is considerable variation in preferences of food banks, both temporary and permanent, often generated by the fact that some food banks have more food from other sources than do others. This is described in Prendergast, 2017. Making an extension in this model would be straightforward, and does not change the central features above. Indeed, in a setting where marginal utilities of food banks are unknown, all that is necessary for the result is that the parties know the equilibrium prices and bid accordingly.[7]

3 Where Can It Go Wrong?

The Proposition above is simply a theoretical construct showing the Second Welfare Theorem with equal information and opportunity. In reality, there are considerable concerns that this may not be the case, and many of the design features of the Choice System were to mitigate such concerns.

Above we distinguished between large and small food banks. Empirically, each food bank has what is called a *goal factor*, which roughly was the number of poor people in a food bank's area relative to its average across the US. Goal factor varies for two reasons across food banks: (i) some areas have bigger population, and (ii) some areas have higher poverty rates. Most of the variation in goal factor is the population size, so below when we talk about larger food banks, it refers to those with higher goal factors.

3.1 The Value of Money

Most markets use real money to reveal preferences. The concern with using real money here is that while is can be used to reveal relative preferences for different kinds of food, it may render its equity objective much more problematic. It is recognized that some food banks have more access to cash donors than others, and usually the smaller food banks have the least money. As a result, a bidding system using real

[7] As one example, consider a case where there is random variation in the marginal rate of substitution between goods 1 and 2, where the utility of food bank k is $U_k(.) = m_k U_1(\sum_{t=1}^{T} q_{1kt}) - \kappa_{1k} + (2 - m_k)U_2(\sum_{t=1}^{T} q_{1kt}) - \kappa_{2k}$, where m_k is drawn from a uniform distribution on $[0, 2]$. Then with a large number of food banks aggregate demand is unchanged from above, but each food bank has an unknown (to the planner) marginal rate of substitution between the two goods.

money would likely disadvantage smaller, poorer, food banks. To deal with this, the Choice system allows only a specialized currency designed by Feeding America called *shares* to be used to bid on food. These shares were initially assigned based on a food bank's goal factor, which in words means that each food bank gets the same number of shares per poor client. In this way, budgets per client are equalized, as in Proposition 1.

As supply arrives every day, it matters that shares can be saved, and do not depreciate. Only by this feature will the food banks withhold bidding for another day. Empirically, this is facilitated by the fact that the daily flow of food is high: about a million pounds of food a day. As a result, if a food bank does not find something it wants today, it likely will not have to wait long before a preferred item shows up. This is in marked contrast to other markets with fake currency, notable the barter markets with scrip currency (as one local example in Illinois, see artofbarter.com) where the absence of liquidity renders the shadow value of that currency low, causing these systems to have limited success.

The model focuses just one one interval. The reality is of an infinitely repeated game, with the outcome repeated over and over. The model is agnostic about how share balances are replenished, simply saying that they are reset by the time that the next interval begins. In reality, budgets are replenished daily. Specifically, suppose that in a given day 50,000 shares are spent. Those shares are reallocated every evening at midnight according to goal factor, where those with more clients in need receive more shares. In this way, the playing field is continuously leveled and the shadow price of currency equalized over time. This also avoids the possibility that food banks simply run out of shares before the next assignment date.

Caps: The system allows food banks to save their shares. One theoretical concern could be where a food bank saves its shares in an attempt to "corner the market" at some later point. This was (empirically) circumvented by imposing a limit to shares: when a food bank reached a pre-specified cap, they no longer receive shares in the nightly reallocation. This ceiling was set at about 2.5% of the total number of shares in the system, so any market power is very limited.

3.2 Corner Solutions

The model above assumes that food banks are not "priced out" of a market. Instead, all food banks buy both kinds of good. In reality, there are two relevant concerns. First, food is generally assigned by the truckload. Yet a food bank may not need or want a full truckload. This would be reflected in the model by the fixed cost κ_{ik}, and if that is large enough, it is not worth it for food bank k to bid on food i. Assume that a truckload has τ_k of food k, costs κ, but that big food banks have C_B clients that can consume it, while small food banks have $C_S < C_B$. Then if $C_S U'(\frac{\tau_k}{C_S}) < \kappa < C_B U'(\frac{\tau_k}{C_B})$, the large food banks will consume (at least some) of food k and the small food banks will consume none.

Even if the smaller food banks continue to purchase both goods, large food banks may still receive more utility from the allocative system than the small ones. In the formulation above, it will always be the case because smaller food banks cannot spread the truckload as broadly across its clients as can a larger one (formally, $C_iU'(\frac{\tau_k}{C_i})$ is decreasing in C_i because of decreasing marginal utility). With other consumption technologies, this need not be. A simple way to see this is to instead assume that a truckload can service C^* clients but not more (for example there could be τ_k cereal boxes in the truck, and each client gets one), then the fixed cost *per capita* is has a lower bound of $\frac{\kappa}{C^*}$. Then if $C_B \geq C^* > C_S$, the indivisibility of a truck has no effect on the large food banks but does disadvantage the smaller ones.

Joint Bidding: To mitigate this problem, joint or fractional bidding is allowed. The food banks *coordinate among themselves to submit a single bid*, where they transmit to Feeding America who pays for what. In the formulation above where a truck can serve C^* clients, if at least n^* small food banks combine to bid, where n^* is the smallest integer such that $n^*C_S > C^*$, the smaller food banks are now no longer at a disadvantage compared to their larger counterparts. Then the allocation system offers the same utility to all parties, a desirable outcome. Below we empirically explore the effectiveness of joint bidding.

Credit: The second discreteness problem in reality involves liquidity: what if a food bank simply does not have enough currency to bid on a desired truckload? This is abstracted from in the model above, where everything is deterministic, but in reality a smaller food bank may find that it does not have enough currency in its balances if a desired load comes along. This again could price the smaller food banks out of some markets, distorting the marginal conditions above. The system allows short term credit to smaller food banks (below the median size), where the amount of available credit is large enough to buy the most desired goods. This credit has a zero interest rate, but there is a minimum repayment required of half of all new currency that is granted to them. Elsewhere (Prendergast, 2017), we show that credit is extensively used (about 12% of all winning bids involve credit), and it has been used predominantly for the most expensive goods.

Note also that joint bidding can also alleviate liquidity problems: a food bank may not have access to enough currency to buy a full truckload of food, but can "afford" a partial truckload. If this is the motivation for fractional bids, it will be most likely for the more expensive goods, where these liquidity issues will bind more. We will see evidence on this below.

3.3 Information on Demand

The efficient outcome above requires bidders to share common information about the location and shape of the demand curve. In general, bidders can shade their bids in the hope of getting goods cheaply. This it itself does not distort allocations: for example, if all bidders reduce their bids by 1 share, then there is no allocative inefficiency. If, however, some bidders believe that they should shade by 1 share and others shade by 100 shares, then there is clearly allocative inefficiency, as those who shade more lose the good sometimes when they should get it. (This is also the logic that can cause efficiency losses from asymmetric first price auctions.) Such efficiency losses would also arise if the food bank misperceives (relative to others) the level of demand Ω.

Many of the design features of the system were intended to both provide extensive information on demand, and additionally to mitigate the fear that some are better informed than others.

The Website Bidding is done online. The website has been designed to reduce the likelihood of poorly informed choices. As one example, consider the problem of a food bank trying to work out a reasonable bid for a truckload of baby food. For some food banks, they may never have bid on baby food before. How do they know what baby food normally sells for? The website has been designed to render such exercises trivial, as a click on the product category quickly reveals the price history of this category of goods. As a result, food banks can quickly calibrate what is a reasonable bid for a given load. As such problems are more likely for smaller, less experienced, food banks, this helps to level the playing field.

A Trial Period The Choice System went live on July 1, 2005. However, before it went live, all participants carried out dummy bidding for 3 months, where they acted as if they were bidding on goods using the website, but where the allocations were done under the old system. In this way, the participants had three months to learn how to avoid mistakes.

Frequency of Loads Concerns about errors in bidding are likely to be most pervasive in settings where auctions are rare. This is not true here. Food banks bid and win loads frequently, so can relatively quickly iron out mistakes. Additionally, it allows them to identify equilibrium quantities that they can expect to receive relatively quickly. As one indicator, by the time of writing, the average food bank has won on average over 500 auctions: these are not the kind of rare, high stakes, auctions where such errors are likely to remain important.

Delegated Bidding Most food banks directors in 2005 had never bid online for anything. So how were they supposed to work out what they should be bidding for a truck of corn coming from North Dakota, in a currency that had no other value than for Feeding America food? An issue that worried some during the design phase was that some food banks would be sufficiently intimidated by the idea of bidding online that they could withdraw from the food allocation process entirely. For those, and for food banks who believed that they would err if left to bid for themselves, the system allows any food bank to simply delegate its bidding to a representative from Feeding America. In that instance, they would simply let Feeding America know what kind of food they would like. (None have ever done so.)

Sealed Bids The measures taken to level the playing field did more than just tell food banks where the demand curve lay. In one instance, the market design conceals the demand curve. Specifically, a concern was raised that smaller food banks would not be able to spend as much time checking auctions online. In a continuous auction, larger food banks could dedicate an employee or volunteer to only buy when items looked particularly cheap (a la Ebay). To alleviate these concerns, bidding is done by sealed bids, so the highest current bid is not known to participants. As a result, the system also conceals features of demand to bidders in order to increase equity.

4 Joint Bidding

It is beyond the scope of this short essay to document the extent to which each of the safeguards above served to alleviate inequality in outcomes across food banks. However, a flavor of their impact can be seen by taking one example: the use of joint bidding as a way of helping food banks. To address these issues, the data below derive from the roughly 65,000 auctions that occurred over the first five and a half years of the Choice System, from July 1, 2005 to December 31, 2011. The data refer to outcomes from winning bids.

The discussion above has been oriented around facilitating access for the smaller or poorer food banks. In that context, joint bidding can help two problems. First, to aid food banks who cannot use a full truckload, but would value some portion of one. Second, to aid food banks who may not have liquidity or credit for a full truck of an expensive food: here their share balance could be used to buy part of a truckload. To answer these issues, it is useful to document and interpret joint bidding through three different lens: (i) how often is it used?, (ii) is there a subset of food banks for whom it is particularly important, (iii) are these the smaller food banks?, and (iv) what food it used for?

How often is joint bidding used? The likelihood that the winning *bid* is a joint bid is shown in Figure 1 from 2005 to 2011, and varies annually from 1.2% to 2%. Yet on average there are three bidders in a joint bid so the fraction of times that *the winning food bank* of an auction uses a joint bid is 4-6%. While this is a low percentage, it is worth pointing out that over this time period, more than 3,000 winners had bid jointly.

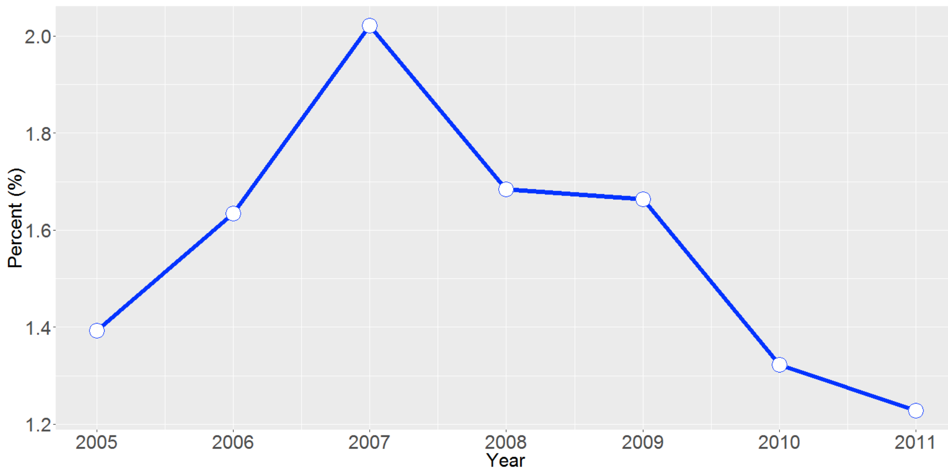

Figure 1: **The Percent of Winning Bids that are Joint**

How many food banks join to bid? The distribution of number of bidders in a winning bid is given in Figure 2. As mentioned above, the mean is 3, though the median is two food banks. Note that over 60 auctions had 5 bidders who won (the count is over 300 = 60 x 5), so for some food banks, even 20% of a truckload is valuable.

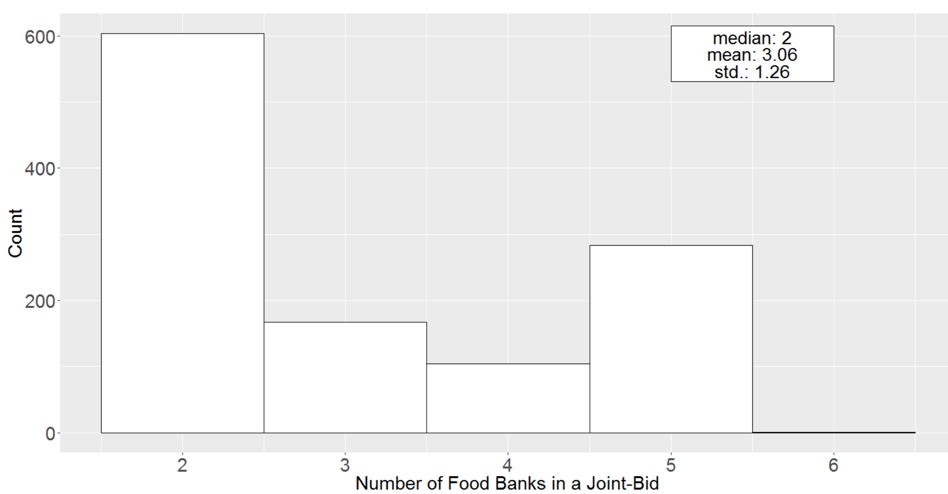

Figure 2: **Number of Bidders Among Joint Bids**

The concentration of joint bidding Much of the concern with the design of the Choice System was to ensure that some of the less advantaged food banks would not be left behind. As a result, many of the provisions above may not be used by all food banks, but instead would be concentrated among a small number of food banks. To answer this, we consider the concentration of the use of joint bids.

To do so, we compute the distribution of winning bids over an entire year by a food bank. So, for example, a unit of observation would be all winning bids by the Chicago Food Depository in 2007. First note that almost 70% of food bank observations never use joint bidding. So, in an average year, only 30% of food banks ever win with a joint bid. On the other hand, for many of those who use joint bidding, it is very important part of their bidding strategy. This is shown in Figure 3 which shows the distribution of the fraction of wins that are joint (among those who have used it), While most use it rarely, for 38 food bank years, it is over a quarter of the times that they win, and for 14 observations, a majority of their allocations involve joint wins. Hence, while joint bidding is small relative to the entire market, it is heavily used by a subset of food banks.

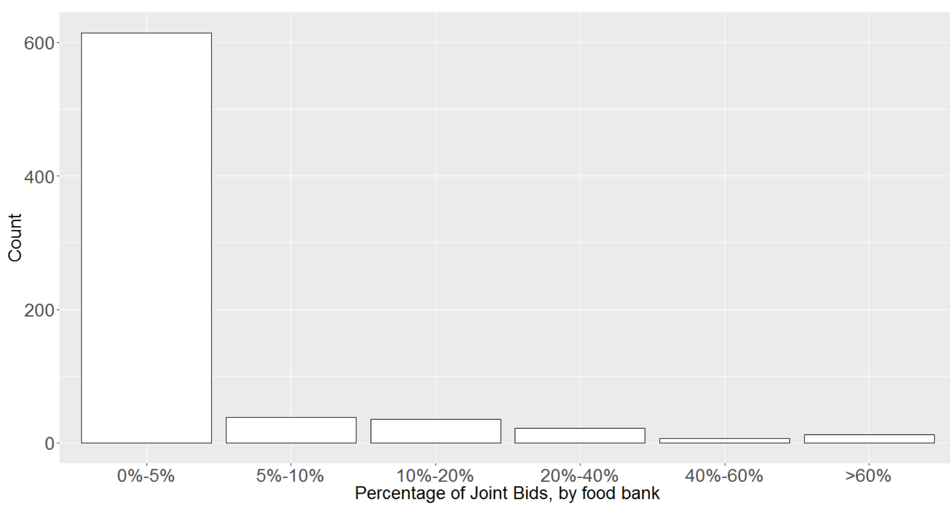

Figure 3: **% of Trades with Joint Bids (among those who use it)**

Networks So far, we have shown that not so many food banks use it, but some of those who do use it extensively. An additional characteristic of joint bidding is that who people bid with also shows considerable concentration. To see this, Figure 4 plots the network of joint bidding, where the numbers are the identity of the food banks. The size of the circles show the frequency of joint bidding for a food bank, and the depth of the arms measure the frequency of partnerships. Here is can be seen that there is a subset (food banks numbers 21, 149, 222, 226, and 277) who have created a network that they use over and over.

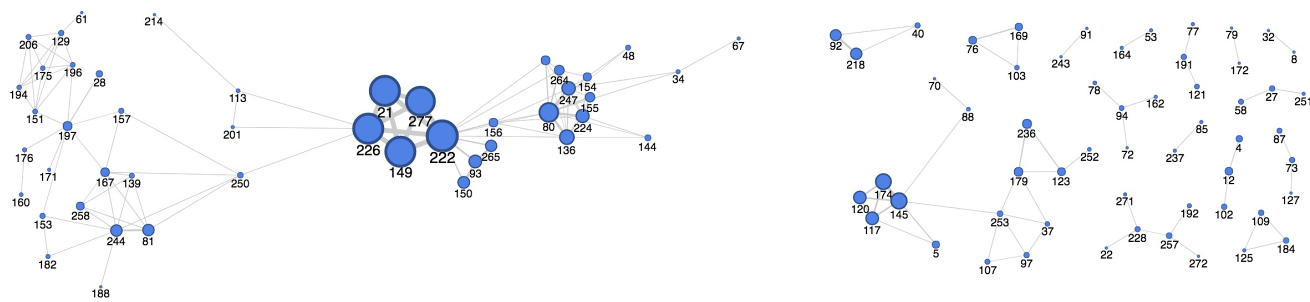

Figure 4: **Joint Bidding Network**

Which food banks joint bid? As suggested by the theory above, those who use joint bidding are, on average, smaller. The average goal factor of food banks who never joint bid is 0.81, while the average who bid at least once is 0.72, but for those who bid at least once every year it is 0.56. Figure 5 shows this relationship visually, where we plot the log of fraction of all expenditures that are on joint bids against the log of goal factor. The predominance of joint bids among the smaller goal factor food banks is visually apparent. In Table 1, we show that this relationship between food bank goal factor and the preponderance of joint bidding among small food banks remains true in a multivariate regression. The relationship is highly significant.

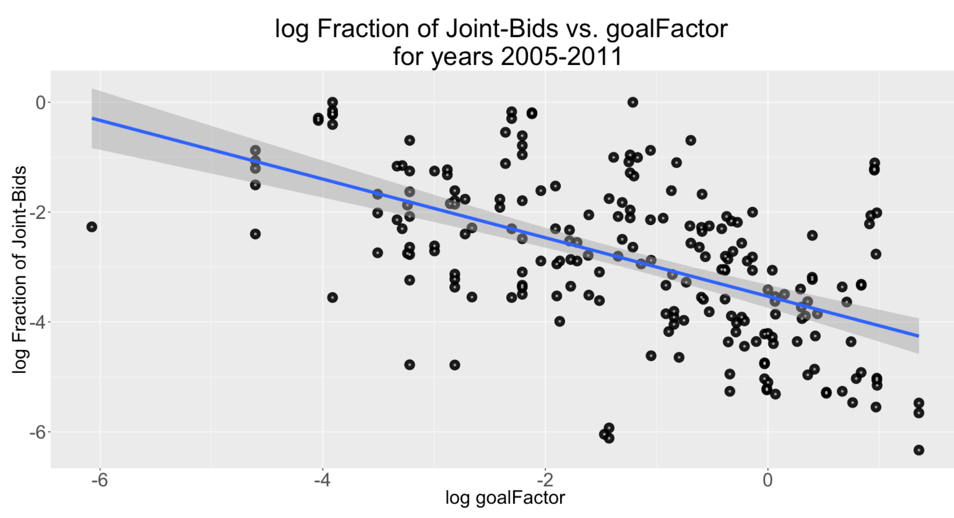

Figure 5: **The relationship between log of fraction joint bids and goal factor**

What kinds of goods get joint bids? A second value to joint bidding is in overcoming liquidity constraints, where a food bank may not want to spend its allocation (or credit) on a single item, but may prefer to only take a fraction of a truck. If this is the case, we are likely to see joint bidding arise primarily for more expensive goods. First, the average price per pound (in the specialized currency) of items bought with joint bids is 0.63, compared to 0.27 among those goods bought with a single bidder. Of course, this could suggest that joint bids "pay too much" rather than are used for more expensive goods.

To show that joint bidding is focused on more expensive foods, we consider the *kind* of food that is bid on jointly. Some categories of food are expensive (cereal, pasta) and some are cheap (produce, canned goods). Empirically, a bidder can get over 100 pounds of produce for a single pound of pasta. Figure 6 shows the relationship between the frequency of proportion of the time that a joint bid wins for that food type compare to the average price of that food type. So, the two foods farthest to the left in Figure 6 are produce and beverages, while on the far right are the two most expensive kinds of food, pasta and cereal. It can be clearly seen that the joint bidding is less likely for the cheap goods. This is robust to regression analysis in Table 2, where we predict the likelihood that a bid is a joint bid using the average price of the category of food being bid upon. The coefficient is positive and highly significant. Hence, it is not only smaller food banks who use joint bidding, but it is additionally used to buy more expensive goods.

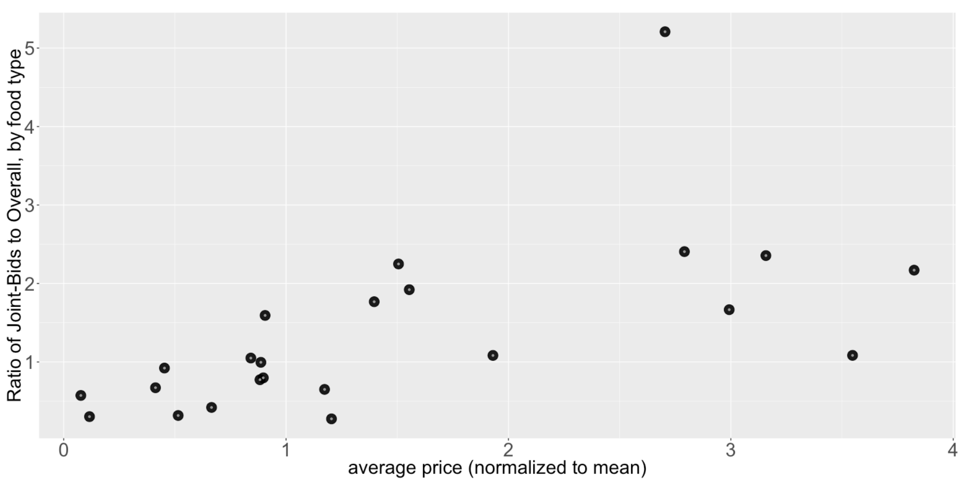

Figure 6: **Distribution of Joint Bids by Price of Food**

Many of the design features in the last section of the paper were meant to level the playing field, especially for small players. The data on joint bidding suggest that this is acting as intended.

5 Conclusion

Textbook characterization of markets usually emphasize their efficiency. Yet in the minds of many, they can seem unfair, offering benefits to the strong or informed at the expense of the weak or uninformed. Such issues have recently been taken on board by those involved in market design, with solutions proposed. Specifically, in both Pathak and Sonmez, 2008, and Budish and Cantillon, 2012, the preferred solution is to design mechanisms that took away strategic advantages: namely, by designing strategy proof mechanisms. In both case, welfare benefits have been gained. Such an elegant solution was not feasible in the food banking case. Here, it was largely carried out by better information (the design of the website for instance), by restricting some strategies (not allowing continuous bidding), and by easing the costs of being small through credit and joint bidding. While these solutions may not be as discrete and elegant as strategy proofness, it is hoped that they have served to get food to areas of greatest need, thereby rendering the benefits of the Second Welfare Theorem a little less abstract.

References

Abdulkadiroglu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sonmez (2006), "Changing the Boston School Choice Mechanism: Strategyproofness as Equal Access", National Bureau of Economic Research Working Paper 965.

Budish, Eric (2011), "The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes", *Journal of Political Economy*, Vol 119(6), pp 1061-1103.

Budish, Eric and Estelle Cantillon (2012), "The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard", *American Economic Review*, Vol 102(5), pp 2237-71.

Hastings, Justine, Ali Hortacsu, Chad Syverson (2017), "Sales Force and Competition in Financial Product Markets: The Case of Mexico's Social Security Privatization", forthcoming, *Econometrica*.

Hortacsu, Ali, and Steven Puller (2005), "Understanding Strategic Bidding in Restructured Electricity Markets: A Case Study of ERCOT", National Bureau of Economic Research Working Paper 11123.

Kagel, John H., Harstad, Ronald M. and Levin, Dan (1987), "Information impact and allocation rules in auctions with affiliated private values: a laboratory study", *Econometrica*, vol. 55 (November), pp. 1 2 75-304.

Kagel, John H., and Levin, Dan (1993), "Independent Private Value Auctions: Bidder Behaviour in First-, Second- and Third-Price Auctions with Varying Numbers of Bidders", *The Economic Journal*, Vol. 103, No. 419, pp. 868-879.

Orhun, A.Yesim, and Mike Palazzolo (2017), "Frugality is Hard to Afford", mimeo, University of Michigan.

Pathak, Parag A., and Tayfun Sonmez (2008), "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism", *American Economic Review*, 2008, 98:4, 1636?1652.

Prendergast, Canice (2017), "The Allocation of Food to Food Banks", mimeo.

Table 1: The fraction of all bids that are joint for each food bank in 2005-2011 (excluding negative prices)

	Dependent variable:	
	Log(% of bids that are joint bids)	
	(1)	(2)
Constant	0.103***	0.138***
	(0.010)	(0.019)
log(goalFactor)	−0.805***	−0.797***
	(0.105)	(0.103)
year dummies	No	Yes
Observations	207	207
R^2	0.046	0.054
Adjusted R^2	0.045	0.045
Note:	***p<0.01	

Table 2: The relationship between Average Food Price and Joint Bids for 2005-2011 (excluding negative prices)

	Dependent variable:		
	Probability that Bid is Joint		
	(1)	(2)	(3)
Constant	0.13***	0.016***	0.014***
	(0.001)	(0.002)	(0.003)
Average Price of Food Category	0.038***	0.038***	0.038***
	(0.002)	(0.002)	(0.002)
year dummies	No	Yes	Yes
Goal Factor x year	No	No	Yes
Observations	59,145	59,145	59,145
R^2	0.01	0.013	0.013
Adjusted R^2	0.01	0.013	0.013
Note:	***p<0.01		

Note: Unit of observation is food bank-auction where each food bank's participation in a joint auction is a separate observation.

Proof of Proposition 1 Consider a possible equilibrium where all food banks bid $b_{ikt} = p_i^*$ until they receive Q_i^* per capita and $p_i^* - \epsilon$ beyond that, where $\frac{U_1'(Q_{1k}^*)}{U_k'(Q_{2k}^*)} = \frac{p_1^* - \frac{\Omega_k(p_1^*)}{\Omega_k'(p_1^*)}}{p_2^* - \frac{\Omega_k(p_2^*)}{\Omega_k'(p_2^*)}}$ and $p_1^* Q_1 + p_2^* Q_2 = 1$. At this proposed equilibrium, $\Omega_k(b_{ik}) = 1$ for $b_{ik} > p_k^*$, $\Omega_k(b_{ik}) = 1$ for $b_{ik} < p_k^*$, and $\Omega_k(b_{ik}) \in [0, 1]$ for $b_{ik} = p_k^*$.
Note that at those prices paid, there is no deviation that the food bank prefers, as all its shares are spent, and it is equating marginal utilities across goods. To see this, he can either bid more or less. But in equilibrium, reducing the price causes demand discretely to 0, and the food bank is left with unused shares, which have no value in the current interval. If it continues to have unused shares at the end of period T, the marginal value of those shares is slightly below $\frac{U_k'(Q_{ik}^*)}{p_i^*}$ and the agent prefers to bid p_i^* today. Similarly, increasing the bid above p_1^* does not increase total quantity, and so there is no desire to deviate. As all food banks make this calculation, this is a Perfect Bayes Nash equilbrium.