

The Allocation of Food to Food Banks

Canice Prendergast*

October 21, 2016

Abstract

Feeding America distributes food to food banks across the United States. In 2005, it transitioned from the centralized allocation process, where they would make decisions based on their perception of food bank need, to one where local affiliates would bid for food items through an online auction mechanism. To do so, it constructed a specialized constructed currency called “shares”. The process by which this change came about, its necessary idiosyncrasies, and its outcomes are described. We both show that it exhibits desirable theoretical properties, and document considerable welfare gains. Supply of food rose by roughly 100 million pounds around the time of its introduction, and the choices of the food banks vary enormously from the allocations they received under the old system. A structural exercise estimates that the value of reallocated demand effectively meant that each pound of food allocated through this system increased efficiency by almost another additional pound.

*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 for excellent research assistance. Thanks also to Eric Budish, Scott Kominers, and Bob Topel for helpful comments and to seminar participants at the AMMA, NBER, MIT, University of Chicago, Duke University, and Rice University. Most of all, however, none of this could have been written without the other members of the America’s Second Harvest Allocation Task Force: John Alford, John Arnold, Al Brislain, Bill Clark, Phil Fraser, Maria Hough, Mike Halligan, Brenda Kirk, Rob Johnson, Susannah Morgan, Steve Sellent, Roger Simon, Harry Davis, Don Eisenstein, and Robert Hamada.

1 Introduction

Food banks throughout the U.S. provide nutrition to the poor and needy. The distribution of food to those in need typically occurs at a local and fragmented level, where food pantries and soup kitchens operate in churches, community centers, schools, and so on. Much of the food is donated by food producers or distributors.¹ The intermediary between these producers and the food pantry is typically a regional food bank: for example, the Chicago Food Depository. The subject of this paper is how a large not-for-profit organization, Feeding America, allocates food to these regional food banks across the United States. Specifically, in 2005 the author was part of a group that designed and implemented a transition from a centralized system, where Feeding America made assignments based on its perception of their needs, to a market-like system based on food bank choice, using a specialized currency constructed by the organization.

Feeding America (the third largest not-for-profit in the United States after the Red Cross and the United Way) is a national human services agency whose mission is “to feed America’s hungry through a nationwide network of member food banks”. It does so through sourcing donations of food across the country, both from large food manufacturers (such as Kraft) and distributors (such as Walmart), and from smaller entities such as local grocery stores, and allocating that food to roughly 210 regional food banks. There are, in general, two ways to do this: by centralized assignment of food - where Feeding America tells an individual food bank what it gets - or by allowing food banks to choose what they want, perhaps with some prices to guide that choice. Before 2005, the agency (like many not-for-profits) eschewed the use of choice and instead used an algorithm to centrally assign food based on its perception of need. Specifically, they would assign pounds of food in direct proportion to its metric of food need - based on poverty rates and size of area - called “goal factor”.

In 2004, a group of 14 - including the author - was charged with evaluating and appropriately changing the allocation mechanism used by Feeding America.² Nine members of the committee were directors of regional Food Banks, three were senior staff at Feeding America, and four were academics at the University of Chicago.³ That group recommended changing to an allocation system based on food bank choice, where individual food banks bid daily on loads of food. To do so they use a specialized currency called *shares* that was created by the organization.

Conceptually, Feeding America’s problem is not difficult: food should end up with the food bank whose need is greatest, after taking account of issues like spoilage and transportation costs. In practice, it is much more problematic. The difficulty is not

¹As a concrete example, a Tyson Foods plant in Kansas has an extra truckload of frozen chicken. How does this end up in a small food pantry far from Kansas?

²At the time, the organization was called America’s Second Harvest, but throughout this essay, we will refer to the organization as Feeding America.

³The members of the Task Force and their affiliations at the time are named at the end of the essay.

primarily in estimating a measure of aggregate need in a “service area”: one can construct measures of poverty at this level that reflect reasonably well aggregate food needs. Despite this, there remain considerable obstacles in identifying how much any given food bank needs a particular load that Feeding America has to offer.

First, food banks receive an average of 20% of their food from Feeding America, and it knows little about the other food it may have received. Some of this variation is transitory, where for example a food bank may already have received eggs this week from another source, and does not really need those being allocated by Feeding America. A second source of variation on the supply side reflects permanent differences, known as “food richness”. Some food banks have close ties with local manufacturers or distributors of food - these are called food rich - whereas others have little access to other food - these are the food poor. Because of these other sources of food, Feeding America typically knows little about what is sitting on the shelves of food banks, and hence, the residual demand curve that it faces. Second, Feeding America assigns a wide range of food: pasta, produce, frozen meat, baby food, peanut butter, and so on. Some foods are more valuable than others to food banks, so how does Feeding America trade off quality versus quantity in its allocations?

We begin by describing the old system, the process by which the change came about, and the exact choice of mechanisms. These were designed to be a tradeoff between theoretical considerations leading to efficient allocations, and a degree of simplicity that would allow the Choice System to be implemented on the ground. This is a combination of share budgets based on need, sealed bid first price auctions twice a day, access to credit for small food banks, the use of negative prices in some instances, an appeals procedure for unusual events, a simple monetary supply rule, and the share revenues generated in any given day being recirculated that evening.

There were a number of objectives in implementing this new market based allocation mechanism. The first was to allow food banks to sort based on their food preferences, both transitory and permanent. The second objective of the Choice System was to encourage more supply of food. Much of the market design literature in economics addresses how to better assign agents to a fixed supply of “slots”: children to schools, courses to students, kidneys to patients, and so on (see Roth, 2008, for a survey). By contrast, one of the motivations for this change was to encourage greater supply of food, (i) by encouraging more offers from donors, (ii) using the liquidity of the system to lead Feeding America to accept more donations, and (iii) by the food banks themselves. On the final issue, food banks receive food from many other sources than Feeding America. The new system allows them to sell *that* food through the allocation system and gain additional shares.

The immediate objective of the paper is to explain and describe these choices, and to document and interpret responses. Beyond showing how a fake currency can be used to generate the benefits of a market (which is rare), we focus on two themes that may be of more general interest for market design. First, we show how assigning a relatively small “fringe” through choice can affect the efficiency of the entire food

distribution system. The mechanism that was implemented allocates only 10% of a food bank's food, but the extreme positions taken by the food banks on that fringe can correct imbalances in the other 90%. As such, it may not be necessary to generate an entire market for a good to generate most of the welfare gains: instead, a small piece may be enough. Second, we study how our market design affects not just efficiency but also inequality of outcomes across participants. A characteristic of the setting here is that some food banks are poorer than others, not in their allocations from Feeding America, but from those that they receive elsewhere. As we will see below, the "food poor" have fared particularly well from this change, thereby mitigating some of these inequalities.

We then consider outcomes. Using data from the outcomes of 65,000 auctions between 2005 and 2011, and aggregates before then, a variety of indications suggest considerable gains from choice. First, the old system treated all pounds of food as equal. The prices (in shares) from the Choice System show enormous deviation from this rule, where at times food banks are willing to pay 50 pounds of one kind of food to get one pound of another kind. The most extreme example is that our estimates show that food banks trade 116 pounds of produce for a single pound of pasta.

As striking is the extent to which food banks sort over long periods. The data clearly show that a benefit of choice is not simply to redress temporary imbalances of supply but rather to allow food banks to permanently sort on the quantity-quality dimension. Some food banks hold back their shares for the high quality items, while others focus their bidding on the large quantities of less desired food. Compared to what would have happened under the old allocation system, 30% of food banks get more than twice as many pounds of food, while another 25% choose to get less than half as many. Also striking is wide variation in sorting on propensity to consume, even over the long run. Specifically, some food banks hold low balances and frequently use credit to buy, whereas others accumulate large permanent stocks of unused shares.

Another way to benchmark welfare gains is to see how allocations have changed under the new system. Our setting is one where - somewhat unusually in the market design field - we can identify and estimate the allocation process both before and after the regime change, which lends itself to measuring welfare changes *ex post*. Under the old system, a food bank that has a 10% greater perceived need (which usually maps into a larger food bank) indeed got 10% more pounds of food. After the changeover, 10% greater need only results in 4.3% more pounds: the Choice System has flattened the relationship between food received and Feeding America's notion of need.

Where does the rest go? Consistent with the theoretical framework that we offer, it is made up of larger food banks (a) buying higher quality food, and (b) simply not spending their shares. These arise because the larger food banks have greater alternative sources of food. By contrast, the smaller food banks both consume more and consume on lower quality foods. The behavior of the food rich has the general equilibrium impact of driving down prices for less expensive goods, thereby leveling the playing field between the rich and the poor in the system.

The smaller food banks have benefited in other ways. As one example, credit is used extensively. For the first two years after the Choice System was introduced, credit use was rare. However, by 2009 - 2011, 12% of all winning bids involved the credit card. Credit is primarily used to buy expensive goods. Finally, many of the safeguards that were introduced at the outset for the benefit of the smaller food banks have fallen into disuse due to lack of need.

An important problem for Feeding America is to manage donor relations. One piece of this entails Feeding America taking offerings from donors that food banks - to put it simply - don't want. Under the old system, arms were twisted to ensure that these would be placed. An innovation of the Choice System is that negative prices are allowed, where food banks' bids need not be positive. This has eased the ability to place these "hard to move" products. Between 2006 and 2011, 8% of all winning bids have involved transfers being made to the food bank rather than them paying. The frequency of negative prices has declined from 13% in 2006 to a more manageable 4% by 2011.

The paper focuses on two sources of gains from the Choice System. First, and most straightforward is to encourage more supply. The supply of food to the system increased by about 35% (or 100 million pounds) in the year around its introduction. To put this from the perspective of another lens, an increase in supply of 130 million pounds a year is equivalent to providing a full day's food for roughly an extra 60,000 people every day. Of this increase, 12 million pounds was food sourced by food banks elsewhere and then placed on the Choice System. To give some sense of causation, within a narrow time window around the change - the first seven months - supply of food rose by 50 million pounds on a base of about 140 million so the timing is well synchronized with the change.

The second source of welfare gain is through reallocation of demand. In order to put more structure on these benefits, a model is offered. Here we carry out two exercises. First we characterize the constrained efficient outcome, and how the mechanism chosen can implement this. To do so, we assume price taking behavior, which we justify given the large number of bidders and the ability of food banks to substitute across both food and time. The purpose of this exercise is to provide conditions under which the Choice System results in *all* food being assigned as the social planner would, not just the food offered by Feeding America. We first show that the constrained efficiency of the entire food distribution system arises by offering each food bank an equal number of shares per client, and recirculating the supply of shares daily in a symmetric way. As such, it outlines the ability to allocate only a subset of the market (a "fringe") through auctions to generate efficiency of the entire system.

The second purpose of the model is to offer testable welfare implications that can be tested with observed data. To do so, we make appropriate restrictions to the model, by aggregating consumption to composite low and high quality goods, and assuming that preferences are quadratic. Providing such an estimate has to overcome two difficulties: (i) that outside endowments are not observed, and (ii) that there is

no natural numeraire to measure welfare gains. In this section, we estimate that the gains in efficiency of the roughly 3 billion pounds in the food distribution system is 8.89%, equivalent to about 260 million pounds, or 150,000 people fed each day. In keeping with the results above, we estimate that over 80% of this gain occurs from the permanent sorting of food banks on the quality-quantity dimension. These stark result reflects the ability of food banks to correct distortions in their allocations from other sources by taking extreme positions in the Choice System. Allied to the direct gains from increased supply, this suggests strong evidence for the value of adding consumer choice to an atypical not-for-profit setting. It is hoped that these lessons may be of some value elsewhere.

Before beginning, it is worth noting that the use of specialized currency in market design settings is relatively rare. There are many unusual features of this setting that allow its use. First, this is far from a one-shot allocation to a consumer: instead, food banks are closer to participants in an infinitely repeated game. This implies that food banks are willing to forego consumption today (by saving) if their desired products are not currently available. As we will see below, there is enormous dispersion in demand across food banks, which makes waiting particularly attractive to those who do not want current offerings. Second, the flow of goods is large - over a million pounds a day - so that participants who do not find what they want today likely will not have to wait long for a preferred good to come along. For these reasons, our specialized currency becomes a store of value that facilitates its efficient use.

We begin in Section 2 by describing the old system, before discussing the specifics of the Choice System in Section 3. Section 4 describes a myriad of outcomes to suggest that these gains have indeed been realized. Section 5 describes a framework through which to interpret both the efficiency gains that can arise here, and offers testable implications on welfare gains. Section 6 offers an empirically testable model and estimates of welfare gains. This system involves the printing of a form of currency, and the monetary rule is discussed in Section 7. Section 8 concludes.

2 Allocations before 2005

The food that Feeding America receives are of two types. First, sometimes donors specify where their donation should go: usually their preference is for a local food bank.⁴ Second, many donors give directly to Feeding America, who then can allocate that food as it sees fit. The subject of this paper is an innovation in 2005 for allocating this second type of donation. At the time of the change, roughly 220 million pounds of food were allocated in this way.

The old system allocated this food based on a metric of need called *goal factor*.

⁴So for example, Tyson in Kansas may have a relationship to the Kansas Food Bank and Feeding America's role is little more than encouraging these relationships and matching these parties when food is available.

This is a weighted measure of (i) the relative poverty of a food bank's service area compared to the nation, and (ii) the relative population of the service area. This was then multiplied by the total number of pounds allocated by Feeding America to construct "goal pounds": the total number of pounds of food that an affiliate *should* receive. Affiliates were ranked on goal pounds relative to pounds received, with the affiliate furthest below its goal pounds ranked highest. Food was then offered to a food bank based on its rank. This mechanism was used since the late 1980s, and allocated 200-220 million pounds of food in 2000 to 2004.

At a concrete level, a food bank would receive a call or email from Feeding America letting them know that they had been assigned a "load". This sometimes had conditions, such as a required pickup date. Food banks were (and remain) liable for transportation costs. The choice of a food bank was to either say yes or no. If a food bank refused a lot, these counted against their need measure as if it had been accepted.⁵ In effect, they received no credit for what were known as "turn down" pounds, so that need of an affiliate was based not on pounds *delivered* but rather pounds *offered*.⁶

This allocation system was widely seen by food banks as representing Feeding America's commitment to fairness, allied to a desire to assign based on need. This sense of fairness was reflected in a number of ways. First, there was also an appreciation that the measures used were transparent. Second, Feeding America only reluctantly intervened to use discretion: while it may know that a given food bank was unlikely to accept a given donation (for example, tinned fruit in North Carolina being offered to a food bank in California), they typically stuck to the rules of the allocation system to avoid any perception of favoritism. (Such interventions did arise, but not often.) Despite these benefits, the allocation system had considerable drawbacks. Foremost among these is the absence of demand side indicators: Feeding America was deciding what was best for individual food banks without knowing what the food banks really wanted or needed. The role for individual choice was minimal other than a refusal to accept goods.

A second problem is that the assignment system treated all foods equally (subject to some minor modifications). A pound of potato chips was the same as a pound of peanut butter. Yet some food is preferred to others: some are nutritionally bet-

⁵An exception to the pounds offered calculation was that produce did not count against pounds offered. Produce is a difficult issue for the food banking industry, largely as it need to be moved quickly due to spoilage issues. This is particularly so as produce is sometimes only donated to Feeding America when it is close to spoiling anyway. If an affiliate was offered produce, it did not count against their "pounds offered" calculation. Produce was typically offered first to the nearest food bank to the donor.

⁶This may seem strange to the reader: why penalize a food bank for refusing to take food that it does not want? This ignores an important issue faced by Feeding America, namely to maintain donor relations. Donors typically want excess food removed from their warehouses for a variety of reasons - to free up storage space, for tax reasons, and so on. As such, there are pressures on Feeding America to remove food quickly, and that pressure is sometimes felt by the affiliates.

ter, whereas others involve higher transportation costs per pound (potato chips are particularly bad on both counts, whereas peanut butter is especially good). Feeding America did not delve into this issue, as it did not know enough about preferences to apply appropriate “weights”.⁷

Much of the discussion below addresses changes to the allocation system to reduce distorted allocations. But isn't there someone most places that needs the food? It is worthwhile deliberating a moment on the nature of these misallocations. The most obvious - but perhaps least important - is the scenario where the poor in one food bank transitorily fare better than those in another. For example, suppose that Feeding America assigns chicken to one food bank twice in a month, while another gets dairy twice. (Chicken is seen as more valuable than dairy.) A bigger concern is food that spoils and is not consumed. It is a feature of food banking that a considerable amount of food ends up in the trash, as landfill, or as animal feed. Some of this arises because donors often give food that is close to its expiration date. (Anyone who has volunteered in a food bank will know the experience of having the task of separating edible from inedible food.) This is exacerbated by capacity constraints on storage, particularly for foods that require refrigeration. Here not knowing the residual supply of food banks makes centralized allocation difficult. Take dairy products for example: sending eggs or cheese to a food bank that does not have excess refrigeration capacity - because its fridges are full - likely results in those products not being used. This is also a significant issue with produce.⁸ Another important component of this inefficiency is where Feeding America turns down donations that it feels will be difficult to place quickly and effectively due to either spoilage concerns or an inability to pick up within the donor's deadline.

All of the above points to problems that arise with a centrally administered system, one that does not incorporate unknown food bank preferences and constraints. Economists are used to extolling the virtues of consumer choice in allocation mechanisms, with appropriately determined prices guiding those choices. Why not then let the food banks choose what they want?⁹ The most natural - and common - way to create such a budget is to attach prices to goods, and let consumers choose goods subject to a budget constraint. Local food pantries are then required to “put their money where their mouth is” to better reflect the strength of their preferences.¹⁰

⁷It would occasionally intervene in an ad hoc way where if a food bank received a particularly good product (hamburgers, say), it would not get meat the next time its turn came around. While this subjectivity was mostly believed to be exercised in the interests of fairness, it was at times a concern for food banks who worried about how they fared through its exercise.

⁸One way around this was informal sharing between food banks. For example, if a food bank did not need an offered donation, it could give it to another who might.

⁹Such pricing occurs further down the supply chain of food banking. For example, soup kitchens and food pantries in Chicago pay to receive some food from the Chicago Food Depository, where different food carries different (subsidized) prices.

¹⁰In many assignment settings, budgets are created naturally through the inability to consume more than one of the good in question. For instance, our children can only go to one school at a

Consumer choice as an allocation mechanism is predicated on one key premise: that “willingness to pay” by consumers is aligned with the objectives of the organization. The issue here becomes whether an appropriate budget can both incorporate unknown food bank preferences and simultaneously meet their overall needs. The problem with using real money to buy food is that while it helps to identify whether a given food bank wants pasta or fish, it is less clear that it satisfies Feeding America’s desire to locate most food with the neediest food banks. For this to happen, the food bank in greatest need must have the biggest budget. There is little confidence that in reality it would: instead, there is a very real danger of the opposite. Food banks would rely on fund raising to pay for this food, and those food banks in the areas of greatest need may have the least access to fund raising, thereby exacerbating the problem. Because of this, Feeding America was reluctant to use the price system in any meaningful way and instead used centralized assignment, despite its warts.

In sum, centralized allocation fails to reflect food banks’ idiosyncratic demands, while pricing with real money fails to offer budgets based on need. Given this, how about consumer choice with fake money? In theory, these two problems are separable: prices can be used to orient choice, and a free hand in choosing budgets could potentially satisfy overall need. As such, it became a promising candidate to resolve both problems, particularly as food banks are “repeat customers”.

3 The Choice System

When the Task Force convened to discuss a redesign of the allocation mechanism, it became clear that there was considerable discontent at the misallocation of food, often leading to spoilage. The example that routinely cropped up was when the Idaho Food Bank was offered potatoes, even though they already had a warehouse full of potatoes. Despite this, when the idea of a “market” was introduced as an alternative, it met with considerable resistance in many quarters. Food banks exist to serve the marginalized, often those that the market economy has left behind. The preferences of food bank directors often reflect that concern for marginalization, and a fear that markets tend to benefit the strong or powerful. Consequently, while the group was open to change, the initial response to a consumer driven choice system was muted. As one food bank director told the author, “I am a socialist. That’s why I run a food bank. I don’t believe in markets. I’m not saying I won’t listen, but I am against this”.

The group met for over a year before converging on what is called the Choice System, using a currency called *shares* to bid on loads of food placed onto the system. Before describing its details, it is important to note that its ultimate introduction lay

time, so a choice system that allows parents to say they want school x can be used to reduce their chances of achieving school y . In this way, tradeoffs can be used to elicit consumer information in a useful way.

not in its broadest conceptualization. One indicator of this more generally is that specialized currencies are very rare in reality.¹¹ Instead, the success of this innovation lay in the myriad of tweaks and additional institutional details that were necessary both for buy-in from the relevant constituents and reflected important considerations on the ground. None of the academics involved in this redesign - the author included - understood the many pitfalls that could have derailed the implementation of this system successfully: for that they relied heavily and consistently on the food bank directors and the staff of Feeding America. The new system would not have occurred without a willingness to listen and adapt on both sides, and the patient and expert moderating of one of our members, Harry Davis.

The starting point of the new system was the creation of a currency called shares. These shares would be allocated to the food banks in proportion to their goal factor, thereby aligning budgets to Feeding America's perception of need. Food banks then logged onto a website (see Figure 1) on which were posted a set of offerings of truckloads of food: for example, a truckload of pasta from a food distributor in Tennessee. At the time of implementation, there were approximately 30 to 40 offerings a day. Food banks would then use its shares to bid, and the winner of the auction was the food bank who bid most.¹² The price paid was the bid of the highest bidder. That number of shares would then be subtracted from the winning bidder's balance. Any items that did not sell on a given day would be carried over to the following day for more bidding. Balances did not depreciate. All shares that were spent on a given day were redistributed at midnight. The rebalancing was done using the same formula as the initial allocation, where the most needy received the greatest fraction of the spent shares. Shares could not be traded for real money nor used for anything other than the items on the auction market described below.

This describes, in the broadest brush, the central details of the auction mechanism and the allocation of shares. However, fairness considerations dominated much of the group discussion, and many of the more precise details below reflect those considerations. In most cases, the concerns were not about who got how many shares, but rather how other potential inequities could result in a playing field which might favor some food banks over others:

- The first concern was for the “little guys”. Food banks vary in size and organizational sophistication, ranging from small banks with a couple of employees operating on a shoe string, to larger outlets in major cities with many employees. A concern that arose consistently was that the allocation process should

¹¹The example that has received most attention in the market design literature has been the use of bidding systems for business school classes at top universities, where students are given points to bid on courses (Budish, 2011, Budish and Cantillon, 2012).

¹²Shares could only be used to bid on what Feeding America calls “yellow pounds”. These are the donations that are made directly to Feeding America, and as mentioned above accounted for about 200 million pounds of food. The donations made to a specific food bank that were alluded to in the introduction (and called “blue pounds”) were not included in the new Choice System.

not harm these smaller food banks relative to their larger counterparts.

- Probably the greatest difficulty in designing the system concerned the issue of “food richness”. Areas that have a denser network of food producers and distributors likely have more sources of alternative food than those that have few. Leveling the playing field in favor of food poor areas was a consistent source of discussion.
- A broader characterization of the food rich issue is unmeasured need, factors affecting need that are not captured in the goal factor. For example, what happens if there is a natural disaster in an area? Or, more commonly, a plant closing in a town? The old system allowed for some discretion by Feeding America by bumping food banks up on the priority list. Finding some way to incorporate these unmeasured needs was a concern.

The precise details of the Choice System, to which we now turn, reflect these (and other) concerns.

3.1 Bidding Mechanisms

Bidding: Bidding occurs twice a day with sealed bids, with the winner paying the number of shares bid by the winning bidder.¹³ Bidding closes at noon and 4pm, with the outcomes being revealed immediately by email after bidding closes. All food for each bidding cycle is posted at least two hours beforehand. The technology allows a food bank to search only for items it might have an interest in, by excluding items based on either geographical constraints or certain kinds of food. A simple click also reveals the history of prices for similar items.

A worry that was raised early in the deliberations was that larger food banks could dedicate a staff person to the bidding process and if there was continuous bidding, those food banks could wait until the last minute and “snipe”. Smaller food banks, which may have one or two employees, could not do this, and would ultimately lose from a system that placed a return to frequently checking the website. This was partially averted by twice-a-day sealed bid auctions, with all food posted at least two hours beforehand.

Joint Bidding: Food banks have the opportunity to bid jointly for items.¹⁴ A concern for smaller food banks is through the indivisibility that a truck is needed

¹³The group went back and forth on what price would be paid by winning bids. A desire to minimize strategic considerations led to some members arguing for a second price auction, but the sense among the participants in the process was that the clarity of “you pay what you bid” was more important. As a result, a first price auction was chosen.

¹⁴Some food banks always bid together. These are known as clusters. They also received food jointly under the old system.

for transportation. Larger food banks can typically use a truckload of any offering, whereas their smaller counterparts may only be able to effectively distribute say a quarter of a truckload.¹⁵ The joint bidding provision was implemented to aid smaller food banks to fulfill their needs while overcoming indivisibilities. Here two (or more) food banks would agree to split a truckload offering, where they submit online the fractions of who pays what.

Delegated Bidding: Food banks can delegate bidding to Feeding America. To do this, they call a delegate at Feeding America and explain their needs, who bids on their behalf. This was offered as an option to a food bank that simply did not feel that it had the resources to effectively manage the process. Such delegated bidding could also be done temporarily, say when the food bank director is on vacation for some period of time.

Credit: Food banks can access credit. Credit was implemented because of a concern that the smaller entities might never receive the most desired products, as a truckload of the most desired goods could sell for more than their share balance. Food banks can use credit to increase their balances to the estimated cost of a highly desired item. They pay off those debts with at least half their future allocations until the debt is paid off. There is no interest rate on these debts. As the larger food banks typically hold a balance larger than this, credit is only available to the smaller and more needy food banks.

Hard-To-Move Product and Negative Prices: Under the old system, there was a degree of arm twisting that arose for product that the food banks did not want. This was done to maintain donor relations. To facilitate the movement of these goods, the new system allows for negative prices. These are called “bonus shares”.¹⁶ Food banks could bid negative shares for lots (up to a limit of -2,000 shares per load) and the good is assigned to the food bank that offers the smallest number of negative shares.

3.2 Unmeasured Hardship

As mentioned above, the old allocation system allowed a degree of discretion, where Feeding America could change rankings based on non-statistical metrics. In order to

¹⁵It is worth noting that some offerings on the market are Less Than Truckloads.

¹⁶When the new system was introduced, bonus shares were not offered on a good on the first day offered. This feature was introduced because there was a concern that if smaller food banks do not check offers every day, they might miss the opportunity to get bonus shares. The two day process gave them more time to see that they could gain shares from bidding on an item. A manifestation of the dissipation of concerns for smaller food banks is that the system has changed in the last few years such that now negative shares are possible on the first day that a product is offered.

allow such discretion in the Choice System, a Fairness and Equity Committee was instigated. This committee, whose members would be other food bank directors, would meet quarterly to review applications from individual food banks who make a claim that their allocation of shares should exceed those currently offered under the measures of poverty included in the goal factor. The committee could increase the goal factor of a food bank by up to 50% for up to three years, thereby entitling them to more shares.¹⁷

3.3 Supply

Market design solutions are often aimed at better matching unknown consumer demands to a fixed supply of “slots”: schools to children, kidneys to patients, classes to students, and so on. So far, our discussion of the Choice System has largely reflected these concerns. However, a significant issue throughout the deliberations was how to generate more supply of food for the poor, both from traditional food donors (producers and distributors) and also from food banks themselves. This could potentially occur in a variety of ways:

More and better supply from traditional donors: Much of Feeding America’s activities come in soliciting donations of food from manufacturers, distributors, grocery stores, and so on. The new system potentially allows further inducements to donate:

- The central objective of the new system is that food be used by those who need it most. One way in which donations could be more effectively solicited is with the message that any food given will be used to the best possible end, as the market will allocate food more efficiently.
- The ability of the Choice System to create liquidity (through many food banks bidding on food) could result in Feeding America accepting donations that were previously denied, due to a fear that they could not be quickly placed with a food bank.
- An auxiliary outcome of the Choice System is to identify those foods which are most valued. Previously there were no good indicators of what foods were most desired by end users: now there are prices. These prices could be used to focus solicitation on those donors who have the highest valued foods by users.

¹⁷This would be used for both temporary relief measures, such as with a natural disaster or a plant closing, but also potentially for more permanent issues such as documented food poorness or high cost of living areas. The Fairness and Equity Committee would fund these extra shares offered to food banks through an annual allocation.

Maroon pounds: Above, we described two kinds of offerings: yellow pounds (those donated to Feeding America) and blue pounds (those directly donated to specific food banks). The Choice System added another source of food called *Maroon Pounds*.

Maroon pounds are foods that an individual food bank already has, perhaps from another source, but for which it may not be the highest value user. The Choice System allows food banks to place this food on the market. The ability to resell food may, or example, lead a food bank to accept a donation when they cannot itself use, but which someone else can. These are bid on in exactly the same way as other product, but where here the winning bid is transferred to the seller rather than redistributed to all food banks.¹⁸

These Maroon pounds are treated differently to other offerings in two ways. First, they are not eligible for bonus shares (the negative prices), as their donor issues have already been resolved. Second, they are taxed. Specifically, a tax of 10% is imposed on the seller of any Maroon shares transacted.¹⁹ The issue of taxing donations from outside the system was the subject of considerable discussion. The ultimate decision to tax these revenues revolved around a revealed preference indicator of “food richness”. As mentioned above, there was a desire to level the playing field based on food richness.²⁰ It was felt that if a food bank was putting food onto the market for shares, they probably had more than enough for themselves. As a result, maroon pounds became the revealed preference metric for food richness, and this became the reason for taxing them.

3.4 The Website

As a summary, Figure 1 gives a screenshot of the website facing a food bank on a typical day. This lists the available offerings (for example, various flavors of Yoplait yoghurt in 4-6oz cups), its weight and location, the food bank’s share balance and credit limit, and so on. On the left of the screen are tools that offer historical price data, and also gives a food bank the option to offer Maroon pounds to the system. Bidders then can place a sealed bid. Two such screen will be seen each day, one

¹⁸These maroon pounds are designed to allow a final source of improved food distribution, which is through the ability to mix loads. Say a food bank has won three bids, with a truckload each of baby food, pasta, and tinned salmon. A smaller food bank may have no interest in an entire truckload of any of these, but would be interested in a truck that has a third of each. Mixing arises when a food bank takes these loads, reconfigures them, and then places them back on the market as Maroon pounds.

¹⁹The tax revenues were given to the Fairness and Equity Committee for disbursement.

²⁰The committee was largely unwilling to “tax” food richness using objective statistical measures of such richness, such as the presence of large food producers or distributors in a service area. While it was generally acknowledged that these were likely correlated with food richness, staff or food bank directors were quick to point out exceptions: for example a food bank, which though located close to a major food distributor, was actually food poor. These exceptions rendered it impossible to use such measures in computing goal factors. Instead, there was more comfort with dealing with the food rich issue through revealed preferences.

for the offerings at noon CST and the other for the 4pm auction. The offerings are usually available by the previous evening. As noted above, the outcomes will be sent by email to all bidders immediately at the close of the auction.

3.5 Other Rollout Issues

Technology: The clarity and simplicity of the technology used played a central role in the Choice System. Before it went live, food banks had played a demonstration version, designed by Don Eisenstein and implemented by Mike Halligan, for over three months and were familiar with its operation.

Buy-in: Feeding America is a democratic institution and the food banks voted on whether to pass the new proposal. The work for this was done by the food banking and staff members members of the committee, and it would never have been introduced without their commitment. At the end of their efforts, the new proposal passed resoundingly.²¹

4 Outcomes

In this section, we outline the outcomes of the change to the Choice System. Before providing the relevant data, it is worth noting that the transition has been perceived as a success by its participants. Food banks are engaged, bid frequently, and largely extol the merits of being able to choose what they want over what they are told to take. This has benefits both to overcome transitory shocks to demand and supply but also by allowing them to sort on the spectrum of quality-quantity dimension. Supply of food has increased significantly, and many of the safeguards that were put in place have fallen into benign neglect due to lack of use.

The data offered come in a number of forms. First, Feeding America monitored the engagement of food banks and their behavior seven months after its introduction, so we being with some data during that time period in order to garner some initial responses. Second, Feeding America has offered the author considerable additional data. The data used here come in two forms. For some exercises, aggregates will be provided both before and after the change to the choice system, from 1999 to 2011.

²¹The subject of this essay is to understand the transition of the allocation system from one that is administered to one that involves client choice. However, in the interests of completeness, it is worth noting that this group dealt with a number of other concerns. Two stood out. The first was a change in the definition of need, the “goal factor” used in calculating relative allocations. This is described in the Appendix. Second, due to a previous merger, Feeding America inherited some additional food banks which shared a service area with an exiting food bank. Previously these “Food Rescue Organizations” were not offered food through the allocation system, but they were successfully added as part of this process.

Those after the changeover derive from aggregating data on 64,570 auctions from 2005 to 2011.

Not all data is used for all the exercises below. In particular, some food banks bid as clusters, where one bids for a group of others. However, it is not clear from the data who is attached to which cluster, so much of the analysis below focuses on these non-clustered food banks. Sometimes food banks disappear in the data, which implies a merger. These rare events are also discarded in some analysis below. These restrictions reduce the number of usable auctions to closer to 42,000. The data have also been stripped of any geographic indicators so we cannot use distance or geography as controls.

The Distribution of Food: A primary driving force of the variation in prices below is the kind of food offered. We begin in Figure 2 by showing the distribution of food from 2005-2011. The wide variety of products can be readily seen: almost any kind of food can show up, as well as non food items such as health care or beauty products. Yet it is worth noting that almost 50% of pounds (40% of loads) are produce, snacks, or beverages. As will become clear a little below, two of these are problematic for food banks.²² By contrast, many of the most desired items - meals and meats for example - are not abundant, with meat and dairy jointly accounting for only 7% of loads. There is little year to year variation in the composition of food in the system.

Initial Responses: All food banks became quickly engaged in the bidding process. Within the first 7 months, 97% of food banks won at least one load. Over this period, each item received a mean of three bids. The range was from 1 to 29 bids. There was quickly a strong relationship between perceived quality of offerings and prices, and food banks quickly began to offer Maroon pounds. However, this is not to claim that the system converged immediately to its steady state. Instead, as we will see below, behavior in the first year or two looks considerably different from their later levels.

In the theoretical set-up below, welfare implications from the Choice System can be decomposed into two pieces: (i) the reallocation of demand from the “equal pounds per client for all” rule of the old algorithm, and (ii) changed supply. We address each in turn.

4.1 Reallocation of Demand

Under the old allocation mechanism, Feeding America treated a pound as a pound, and its objective was to offer each food bank an equal number of pounds per pseudo

²²First, beverages - usually carbonated drinks - typically have little nutritional value (though are often very valued by clients), and are expensive to transport. Second, produce is in such abundant supply that there is little willingness to pay. In addition to this, much produce needs refrigeration and is often close to expiration by the time it is donated to the system.

client. We address the reallocative aspects of the system in a number of steps. First, we show that it is quite clear that a pound is not a pound: there is huge variation in prices per pound. Second, we show that this variation in willingness to pay has results in considerable sorting of food banks on the quality-quantity dimension, and that the new allocation looks very different from the one where food banks' receipts of pounds of food was proportional to their goal factor. Third, we illustrate considerable permanent sorting on characteristics other goal factor. We finally address how these changes have affected the smaller food banks.

Variation in Prices: In Figure 3, we show the empirical distribution of price paid per pound from 2006 to 2011. This does not condition on the kind of food, when it was offered, and so on. There is enormous variation. The average price level - meaning the number of winning bid shares per pound - is 0.28, so that on average food banks were receiving three pounds of food for each share offered. The variance of the distribution is 0.61. Almost half of all offerings sell for 0.05 shares or less - a sixth of the mean - where a food bank can get 20 pounds of food for each share it spends. The distribution is right skewed, where the median purchase is 0.05. For more expensive items, 25% of loads sell for 0.3 shares or more, and in 10% of cases, the buyer gets two pounds of food per share or less. As one extreme observation, in 8% of cases the buyer receives less than one pounds of food per share spend, while in 25% of outcomes, the buyer received 100 pounds of food per share, or more.

The old algorithm treated all pounds as equal, and as a result there was randomness in the quality of goods received by each food bank: sometimes when your number came up, you got good food and other times less desired food. While we do not attempt to simulate the old algorithm to show these inefficiencies, the variance in prices suggest that they may be large.

Loads vary by many characteristics - location, transportation costs, likely spoilage, and so on - but the primary source of variation in their desirability is the kind of food. Figure 4 shows price variation by food type. The numbers here have been normalized such that the median good has a pseudo price of 1, so the numbers here reflect price relative to average price. At the lower end of the distribution, produce (marked as "fresh" in the language of the offerings) sells for 4% of the price of the average good, and beverages trade for 5%. In simpler terms, a food bank can get 30 pounds of produce for a single pound of the average good, or 20 pounds of beverages. Remember from the Figure 3 that these two categories represent more than a third of all offerings, so this points to a large fraction of offerings to the system that the (marginal) food bank is unwilling to pay for. By contrast, prices for rice, pasta and cereal are roughly three times the price of the median good.²³

In Table 1, we provide regression analysis predicting the log of the winning bid for a given load, using product characteristics, year and month dummies, and food bank

²³All else equal, these (though desired) are not as desirable as meat or meals, but the latter requires refrigerated trucks, which are considerably more expensive.

attributes.²⁴ The excluded category on products is cereal, the most desired good. The regression results show similar outcomes to the raw data. The levels are striking: a single pound of cereal gets 143 pounds of produce, or 10.91 pounds of juice.²⁵ The most desired goods are pasta, rice, diapers, paper plates, meat, cereal, and meals. The greatest rate of exchange between any of these is 2.17 pounds of meals for a pound of cereal. Produce and beverages stand out as the least desired.

It is useful to consider here why there is so much variation in prices by food type. Paying the same amount for a pound of pasta as for eight pounds of dairy may seem strange, given that in a store their relative prices might be closer to one to one. The reason why there can be such divergence in relative prices largely lies in the fact that prices are reflecting the residual demand of food banks given their outside endowments. So, for example, while a food bank may value pasta and dairy products close to equally overall, they already have enough dairy through other sources, and so relative price difference become amplified. Yet in a way this reflects the beauty of the Choice System, where allowing trade in a relatively small part of the market can lead to the entire market being efficient.

Sorting: The large variation in prices above gives rise to the possibility of food banks sorting on the quality-quantity dimension. We now turn to the extent to which this arises. For data reasons, it is simpler to measure dispersion in prices.²⁶

In Figure 5, we aggregate data on food banks from 2006 to 2011, and measure the average price per pound paid over that entire period. By taking all six years, we abstract from temporary rebalancing and instead measure how food banks permanently sort on the quality-quantity dimension. As a benchmark, under the old algorithm, each food bank would have received 3 pounds of food for each share. (Of course, shares were not in existence then, so this is simply dividing the number of pounds of food in 2008 by the outstanding stock of shares.) These data show enormous sorting on the quality dimension. Among those who restrict purchases to high end goods,

²⁴This regression is in logs to ease interpretation. However, this entails that all negative prices have been excluded from the regression. In Table 8, we provide a regression in levels, including all these negative outcomes, and the results do not change.

²⁵It is worth emphasizing one caveat. These results show that food banks are not willing to pay for produce. This does not necessarily imply that they do not value it, as prices reflect the willingness to pay of the second highest bidder. For produce, this distinction may be relevant. Take the following example. A food bank in Tennessee needs produce, and a load nearby is available. Produce is available just about every day, and needs refrigerated shipping. Because loads are so abundant, food banks know that a load close to their current location will be available soon. As a result, they may not bid on produce that is not located close. The Tennessee food bank realizes that it face little competition for the load and bids low. Hence care should be taken in assuming that - in the case of produce - low prices imply low marginal valuations.

²⁶Similar outcomes arise if we were to plot average number of pounds divided by goal factor. We are aware of course of the issue that measuring the variance of the inverse of prices is not the same as the inverse of the variance of prices. However the dispersion is so great that it is likely to swamp any of these Jensen's inequality or existence issues.

25% receive half the number of pounds per share or less than with the old algorithm. By contrast, there are a large number of “cheapskates” who are receiving vastly more food than under the assigned outcome: almost 30% of all food banks receive twice as many pounds per share as under the old regime. 10% get at least three times as much as before. Figure 6 offers a similar snapshot, but for a single year, 2007. While a similar conceptual story arises, there is more dispersion, reflecting the fact that some rebalancing on the quality-quantity dimension reflects transitory effects.

It is worth emphasizing here that these benefits underestimate the value of demand reallocation the Choice System for two reasons. First, by considering price as the metric for quality differences, it misses much rebalancing that occurs within price categories. For example, consider a food bank that already has a lot of yoghurt. It can rebalance by buying other goods that sell for the same price as yoghurt, such as milk. Such compositional rebalancing is surely important. The second source of underestimate is that we use annual data as the relevant time frame for rebalancing. This does not pick up rebalancing across intervals within a year. For example, if the relevant time period over which a food bank optimizes is two months, and in half of those months buys expensive cereal to rebalance and in the other half buys cheap potatoes, its price level will be the mean, yet there will be significant rebalancing benefits that are not estimated here.²⁷

Savings: Another margin on which food banks have responded to the new system is not only on the quality of food, but also on whether they consume. In the framework above, we addressed this by considering a case where the food rich do not value the shares on the margin, but the food poor do. To address this, consider the propensity to consume from shares received in a given year. This measured by total shares spent divided by goal factor (as shares received are proportional to goal factor). We offer two exercises here. First, in Figure 7 we plot the permanent marginal propensity to consume, by showing the distribution of total expenditure to total share allocations from 2006 to 2011. In words, of all the shares they get over a 6 year period, what fraction do they spend? Here the average food bank is normalized to 1. The dispersion in this measure shows that food banks clearly use the margin of not consuming, with some food banks spend half as much as other food banks. This strongly suggests that some food banks are permanently staying out of the market for some of their shares.

How different is it to the old system?: Under the old system, food was allocated based on its measure of need, *Goal Factor*. In this section, we address how the relationship between food received and goal factor has changed with the Choice System. We do this for two reasons. First, it addresses how close to right the old system was, and is a natural measure to reflect the welfare gains from Choice. Second,

²⁷The reason, of course, is that the interval over which they optimize (the interval I in the model) is a month, but the data aggregate over $12I$ and miss the relevant variance. In order to properly address this, we would need to posit a value for I .

much of the discussion in the implementation of the system involved taking care of the smaller food banks. A food bank's goal factor reflects pretty well its size, so that by addressing the changed relationship between goal factor and food received we can begin to address its impact on smaller food banks.

To begin, remember that under the old system, the algorithm offered food to food banks in direct proportion to their goal factor. One difficulty with many market design exercises is that researchers often do not know the *ex ante* algorithm for allocating goods. Here we not only know the algorithm used, but can also test for its veracity. We test the "before" period by regressing $\text{Log}(\text{Total Pounds})$ on $\text{Log}(\text{Goal Factor})$, where a coefficient of 1 would be the theoretical outcome of the old algorithm. We carry out this exercise for the period 2002 to 2004, with one caveat. The author has not been provided goal factor data before 2006, so that variable is used in this exercise.²⁸ Table 2 shows that the coefficient is 1.01 and highly significant. Table 2 also examines the stability of this by seeing if this relationship varies by year, but there is no statistically significant difference for the three years. This shows that the "pre" period is as advertised, where a 1% increase in relative need increases food allocation by 1%.

Tables 2 also shows this relationship after the change. For 2006 to 2011, the coefficient falls a highly statistically significant coefficient of 0.43. These are statistically different from each other.²⁹ To put this in simpler terms, a food bank with a 10% higher goal factor used to get 10% more pounds. Under the Choice System, they now get 4.3% more, so that the new system is reducing the relationship between perceived overall need and pounds received by the residual 5.7%.

Where does the remaining 5.7% go? In the theoretical description below, we note two likely responses to the new system by the food rich: to buy more expensive goods, or not to consume at all. In Tables 4 and 5, we address these two responses. We begin in Table 4 by regressing the average price paid by a food bank over a year on goal factor. In words, do they buy more expensive products over the year? Table 4 shows that average prices rise 0.37% for every percentage point increase in goal factor. So much of the remainder is made up of larger food banks responding by moving up the quality ladder. This arises because many of the cheaper staple goods suffer from diminishing returns - this is particularly so for produce - and as a result food banks that already have access to such staples will substitute into rarer more high quality goods such as pasta, cereal, or diapers which always find willing consumers.

In Table 5, we turn to the marginal propensity to consume, which is measured by the ratio of share expenditures to goal factor (shares issued are proportional to goal factor). We regress this on goal factor, time dummies, and interactions. Over the entire time period, this coefficient is mildly negative. However, this masks an issue that arose in the first six months (2005) when the smaller food banks were

²⁸Goal factors change slowly so this is not a major concern.

²⁹Regressions carried out for each year from 2006 to 2011 shows that this coefficient remains constant at that number.

still learning how to bid, and did not initially spend their shares. The second column shows that - after 2005 - a 10% increase in goal factor reduces the marginal propensity to consumer by 2.2%, explaining the residual unexplained by choosing higher quality food.

To summarize, the Choice System untethers the one-to-one relationship between pounds received and perceived need. Previously, increases in perceived need results in one-for-one increases in pounds. Under the choice system, the food banks choose to take this as 43% extra pounds, 35% higher quality food, and 20% unspent shares.

Sorting Again: The results above on Goal Factor show that one criterion on which food banks sort is their goal factor, where larger food banks choose to buy higher quality goods. However, goal factors are only a small portion of how food banks sort permanently across food types. Here we address this by asking how much of the variance in prices paid by food banks can be explained by their permanent revealed preference. To do so, we hazard a guess as to a reasonable interval over which food banks attempt to balance food type. We choose a two month period, over which the average food bank wins ten loads, and compute the average price per pound paid by a food bank over that interval.³⁰ Some of the variation in this measure is likely determined by factors that are not food bank specific, as aggregate demand and supply can change over time. In Table 6, in keeping with the preferred specification above, we begin by regressing the log of these two month average prices on year and month dummies to measure residual food bank specific variation. Adding goal factors to this regression improves the predictive power of the regression by 15%, so that a non-trivial amount of sorting is on goal factor. However, when we use food bank fixed effects in Table 6 rather than goal factor to see how much price variation can be explained by the average price paid by a food bank from 2006 to 2011 (its permanent preference for quality), we explain 40% of the residual variance.³¹ The marginal effect of goal factor conditional on the fixed effects is negligible. This provides strong evidence for permanent sorting of food banks on the quality-quantity dimension, but primarily on dimensions that are not easily observable: while food banks do sort on goal factor, it reflects a relatively small component of their permanent preference for quality.

This result points to the difficulty of making a simple tweak to the algorithm to

³⁰Two months is probably longer than the usual interval over which food banks would choose a portfolio of food, but we need food banks to win enough loads to avoid random variation in the type of food they buy. We compute the average price per pound, P_{ijk} for each food bank i in year j and two-month interval k : $\bar{P}_{ijk} = \frac{1}{N_{ijk}} \sum_{t \in I_{jk}} P_{it}$, where N_{ijk} is the total number of purchases food bank i makes in year j , interval k and P_{it} is the purchase food bank i makes at date t . I_{jk} denotes the time interval k in year j .

³¹Remember that regressing Log prices implies excluding negative prices, here from food banks whose two month averages are negative. Carrying out the regression in levels does not change the results.

improve efficiency. Given the results above on how goal factors predict behavior, a natural simple alternative to the old algorithm would be to assign fewer high quality goods to the high goal factor food banks, and plentiful low quality goods to those with low goal factors. The fact that goal factor is only a third of the permanent sorting of food banks on the quality dimension illustrates the limited value of such an improved algorithm over the revealed preference characteristic of the Choice System.

The Small Guys: The results above show that the larger food banks have responded by buying more expensive food and not spending some shares. This leaves more pounds of lower quality foods for the smaller food banks at cheap prices. In the framework below, we show how smaller food banks benefit from the choices of their larger counterparts by reducing the relative price of the cheaper goods. Put more simply, one of the benefits of the Choice System for the smaller food banks is that it has rendered large quantities of staple foods available at low prices. To a first approximation, if a food bank has access to a truck, food is always available. Of course, they do not get so much high quality food, but that is their choice.

Three other crutches were offered to the smaller food banks: the Fairness and Equity Committee, Credit, and Joint Bidding. First, as another indication that the smaller food banks have fared well from the change, note that at the design stage, two institutions were added to aid them: the ability by food banks to delegate bidding to Feeding America, and the Fairness and Equity Committee. No food bank has chosen to delegate bidding control to Feeding America, except for cases where the director is on vacation for a short period. More striking is that the Fairness and Equity Committee has *never* convened, due to the widespread satisfaction with the outcomes of the Choice System.

Much of the design of the Choice System was aimed at making the kind of tweaks that would make it run more efficiently, with a particular eye on leveling the playing field for smaller food banks. One example of this is allowing access to credit for the smaller food banks. Table 7 shows credit use over time by measuring the fraction of winning bids that involve credit. In the early stages of the Choice System, the use of credit was relatively rare, with only 4% of winning bids involving the use of credit shares in the first 18 months. However, over time food banks have learned to like credit, so that from 2008 to 2011, the fraction of winning bids has remained stable at roughly 11%. Remember that only about half of all food banks qualify for credit, so that among those food banks that qualify, almost a quarter of all the winners use credit. This suggests that the need for these smaller food banks to smooth transitory fluctuations in their endowments is important.

The above lets us know that credit is used, but does not tell us what it is used for. Much of the motivation for allowing food banks access to credit was to overcome the indivisibility induced by a truckload, where a small food bank would have share balances that were too low to win the more expensive items. This would suggest that credit would be used for expensive items. To address this, first see Figure 8, which

identifies the kind of goods for which credit is used relative to the overall distribution of food. These are ordered from cheapest to most expensive. It can be clearly seen that credit use is greater for more expensive products, as purchases of the cheap items are lower and the more expensive items are higher: credit is a third as likely to be used on produce or beverages, and 50% more likely on meals and pasta. This is formalized in Table 8, where we predict price per pound now including a dummy for the use of credit in the winning bid. (Unlike Table 1, this does not control for product as we wish to identify prices paid, not prices paid conditional on product bought.) The coefficient is positive, highly significant, and large: items bought with credit have prices that are 4.7 times those bought without credit. This strongly suggests the use of credit facilitating smaller food banks affording these more expensive items.

Finally, joint bidding was allowed to overcome the indivisibility that arises from a truckload: now food banks could share a truck. Table 7 shows the frequency of joint bids, which average between 1.5% and 2% of winning bids. On average, three bidders combine in these cases, so an alternative way to state the number above is that in 5-6% of cases, the winner is involved in joint bidding.

“Hard to Move” Products: Under the old system, it was difficult to place “hard to move” product, to use the parlance of Feeding America. Arms were twisted in order to keep donors happy. An innovation of this market was to allow negative prices, where food banks would be paid to take certain loads. Table 7 shows the frequency of such negative prices over time. These data show a striking decline in the need to bribe food banks into taking food. In the first two years of the Choice System, 11% of loads involved the need for “bonus” shares, yet this has declined considerably to only 5% in 2010 and 2011. The level of 5% is relatively small, and suggests that the need to keep donors happy involves relatively little distortion.

There is one important caveat to only considering negative prices, which is that the Choice System does not allow negative prices for produce.³² Table 7 also extends the analysis above by showing the frequency of non-positive prices. Similar to the last figure, a large decline can be seen here, from 45% of all loads in 2005 and 2006 to a stable 17% from 2007 to 2011. This again reflects the greater liquidity of the system.³³

These data would show how the use of negative shares (or very low positive bids on produce) have eased this ability of Feeding America to facilitate donor relations. However, some auxiliary data would suggest an issue may remain because in a large number of cases the lower bound on prices is hit. First, conditional on a negative

³²This was a decision made because produce is so abundant in the system that there was a concern that on some days the average price paid could be negative, which would result in the reallocation of shares at midnight reducing balance from one day to the next, which was seen as politically infeasible.

³³In addition, this may reflect a change in bidding strategies by some food banks. Over time, some food banks seem to have worked out how to bid slightly greater than 0 in order to beat other food banks bidding 0.

price arising, in 43% of cases, the price is the limit of -2000. Second, in 27% of cases, produce sells for its minimum value of 0. These data suggest it remains the case that even with these low bids, a considerable amount of product is not desired. If no bids arise in the feasible range, the donations are turned down or Feeding America resorts to twisting arms again. If these bounds are set to reflect the shadow value of these donor relations, then this is of course the efficient outcome, but it remains unclear of some donations should be encouraged beyond the current low levels that are used.

4.2 Changed Supply

The second conceptual benefit of the Choice System is possible increased supply of food, though either greater willingness of donors to give food, or that Feeding America becomes more likely to accept donations. We begin in Figure 9 by comparing the amount of food in the system before 2005 to the amount of on the auction afterwards. Supply supply by 2006 to 2011 averaged 350 to 370million pounds, up from 200 to 220 million before the change. As a result, the increase in total supply of pounds is between 130 to 150 million pounds. While some of this food surely would have ended up being used elsewhere, this is a large and rapid change.³⁴

An increase of say 130 million pounds a year remains an abstraction. What does this mean for people fed? Ignoring all the welfare benefits associated with better allocation of demand across food banks, consider the value of an extra 100 million pounds of food for people fed. The average person eats 4 pounds of food per day, so allowing for 20% spoilage (which is large), this implies that the increased supply allows approximately an additional 80,000 people to receive those 4 pounds *every day*. As such, it represents a considerable endorsement for this market-like allocation mechanism.

These are time series relations, of course, so some caution should be taken in assuming causality. However, a clearer picture may be possibly seen by considering a narrow window around the time of its introduction. While the supply of food to the system was relatively constant before the change to the Choice System, the number of pounds of food on the system rose by 50 million pounds in the first seven months after its introduction. Specifically, after 7 months, 192 million pounds of food had been “sold” though the new allocation system, compared to roughly 140 million pounds by that time in a normal year. This sharp increase at the time of introduction lends credence to the idea that this increase was indeed causal.

Some of the increased supply is Maroon pounds. These are given in Figure 10. These add approximately 12 million pounds to supply each year from 2005 to 2012. However, it is worth nothing that maroon pounds sell for about twice the average price per pounds on the system. As a result, if one adjusts for quality, the impact of these Maroon pounds would be closer to 25 million pounds. It is difficult to calibrate whether the amount of Maroon pounds is high or low at 12 million pounds, but it

³⁴It is also worth pointing out that this arose before the onset of the financial crisis.

is worth pointing out that it is only about 5% of the total of the system. There are at least two reasons why this is probably so. First, these are typically goods that already are located at the donating food bank, and so there is a clear wedge between the value of the goods to the donor and to any recipient, which has to transport it. Second, this option is arising in the shadow of a system where food banks already share with each other, and an enunciated concern with the Maroon pounds by some food banks is that it is seen as “unfriendly”, to gain shares at the expense of other food banks, when they could simply be given to someone. A reasonable interpretation of the supply of Maroon pounds may be that they are used when there are few takers in the usual network through which a food bank shares, and so it is an alternative to the goods not being efficiently used at all.³⁵ It is also important to bear in mind that some of this increase would likely have found its way into the charitable system through some other means, so that no claim is being made that this is the net increase in food going to the poor.

Before concluding a discussion of outcomes, it is worth noting that one value of the new system is that it allows dispersed outcomes for food banks while maintaining transparency. Before the changeover, Feeding America could probably have used its discretion in an ad hoc way to steer food to a food poor area over a food rich one, but they were reluctant to do so as it could be construed as favoritism. A major benefit of the Choice System is its transparency: all offerings are broadcast beforehand and anyone can bid. This ex ante transparency has rendered the concerns of ex post variation in allocations much more palatable. As one example, prices are often zero or negative, yet there are few complaints about this as anyone could have bid on them.

5 A Model

The purpose of this section is to offer a framework to interpret the facts above.³⁶ Before doing so, it is worthwhile to consider a range of welfare implications of the Choice System based on the description of the setting above:

1. Much of the discussion above has addressed food richness, where the food rich

³⁵This change in supply is not quality weighted. Theoretically, it is not clear how quality will be affected. First, knowing prices will help Feeding America solicit donations: as food banks value peanut butter so much more than produce it may orient their solicitation efforts towards those high quality goods. Second, the average quality of maroon pounds is much higher than the average good on the system. Yet this must be balanced against the donation decisions that Feeding America makes: which offers will they accept? One feature of the Choice System is that the market has become more liquid and product easier to place. This is not a concern for placing chicken or peanut butter. Instead, it is likely that Feeding America becomes more likely to accept marginal quality donations: tinned vegetables for example. If this effect is dominant, then average quality could fall. Unfortunately, no data is available on the composition of food donations before the Choice System, so we cannot determine how the new allocation system has affected quality.

³⁶See Alexandrov et al, 2015, for another theoretical analysis of food banking.

- focus their acquisitions on high quality goods, allowing the food poor more access to the lower price staples.
2. Outside endowments of food also vary in their composition. For example, a food bank could have great contacts with dairy producers to access such products, and so a primary interest of their Feeding America allocations would be to avoid further dairy products. Allowing choice can redress such compositional imbalances.
 3. A possible benefit of the Choice System is that food rich banks that do not need food will not access the market, but rather leave their share balances permanently unspent. By leaving this food for those who need it more, additional welfare is attained.
 4. Increased supply of food.
 5. Under the old algorithm, Feeding America would often offer loads to food banks which were not geographically close. The Choice System facilitates focusing their bidding on more proximate offerings, allowing them to cut down on transportation costs.
 6. The old algorithm treated a pound as a pound. Yet food varies in its quality, and the algorithm could randomly allocate too much high quality food to one food bank and too little to another. With diminishing returns, this can entail an imbalance across food banks which the market can redress.

Given this, it is useful to address the role of the framework below. Its objective is to provide a structure to understand issues 1., 2., 3., and 4., above, but not 5. and 6. This is largely because the empirical exercises we carry out can only address these issues. We carry out two exercises. First, we show how the policy of “equal incomes” - namely, an equal number of shares per client and nightly reallocation of the shares - leads to the implementation of the constrained efficient outcome. Second, we simplify the model to allow empirical testing: under the assumption of linear demand curves (quadratic preferences), we can provide welfare implications of the Choice System through two variables, the increase in supply of food, and the variance of pounds per client.³⁷

Somewhat surprisingly, economics does not have an “off the shelf” workhorse model for empirically evaluating the welfare implications of moving from central assignment to a choice-based market system. Here we offer such a simple framework. It is based on a number of assumptions. The most important of these is price taking,

³⁷We do not address the distortions caused by the randomness of the algorithm, as it requires simulating the algorithm and making a guess as to how Feeding America dealt with geographic considerations. Nor do we address transportation costs benefits as the data have been stripped of all geographic information.

where strategic behavior of participants in their bidding is minimal.³⁸ The case for price taking relies on the nature of the auction. First, previous theoretical research has shown how the benefits of non-price taking behavior rapidly become small as the number of bidders rise (Rustichini et al., 1994). Here $n = 210$ food banks.³⁹

Second, there are enormous substitution opportunities for any given load. The first source of substitution is across foods. For example, a food bank will see tinned peas as a close substitute for tinned beans, or will see rice as a close substitute for pasta or potatoes. The existence of close substitutes on the food end implies less variation in willingness to pay (per unit of quality) across food banks. Another source of substitution is over time: food banks are often largely indifferent about whether they receive a load of usable food this week or next week. The objective of food banks is largely to ensure a sufficient range of foods over say a month. We formalize this below through what we call “intervals” below, where they seek a range of foods over the length of that interval. Such opportunities for inter-temporal substitution also restrict market power as all food bank who will serve say pasta within two weeks from now will bid on it today. For these reasons, strategic bidding is ignored below.

5.1 Framework

There are m goods that can be used by food banks. On any given day $t = 1, \dots, \infty$, a quantity of good j , q_{jt} needs to be assigned among n food banks. (n here is 210.)

Demand Although donations are made every period, consumption occurs over a discrete “interval”, which lasts T periods. (Think of this as a month.) I refer to each T period length as an interval I . For an interval I , the food bank has preferences defined over total quantities of the goods that they consume in the interval. Let x_{jI} be the quantity of good j consumed over the interval I . Then the utility of a food bank from good consumption *per client* in that interval is $U(x_{1I}, x_{2I}, \dots, x_{mI})$, where $U^i(x_{1I}, x_{2I}, \dots, x_{mI}) > 0$, and $U^{ij}(x_{1I}, x_{2I}, \dots, x_{mI}) \leq 0$, for all i, j , where superscripts refer to derivatives.⁴⁰ (We model this as food per client to cut down on notation. Allowing food banks to vary in size does not affect the outcomes.)

³⁸Note that strategic behavior by bidders cannot be ruled out purely on theoretical grounds. This is a first price auction, and only under special assumptions will this generate the usual efficiency outcome where the most valued bidder always wins and does so at the price of the second highest bidder. Specifically, revenue equivalence can fail when bidders have asymmetric valuations for goods. Such asymmetries are a central part of the welfare gains from the market below, and as such - theoretically at least - strategic bidding cannot be ruled out.

³⁹Of course, for any given load in a specific location, the number of active bidders will be much lower, but this auction is far from the small numbers case one often sees.

⁴⁰Note that we are ignoring negative prices here. This is for simplicity, as it allows us to choose the principal’s preferences as the aggregate of the food banks’. Negative prices would require adding donor preferences, which add little.

A food bank's consumption has two components: the food that food bank k receives from Feeding America, $\sum_{t=0}^T q_{jkt} = q_{jkI}$ for all $j = 1, \dots, m$, and its endowments of food from other sources. We assume that food bank k 's other supply of food of good j per client is given by $\mu_k D_j$, where μ_k is unknown to Feeding America. Each food bank's μ_k is independently drawn from a distribution $F(\mu)$ with density $f(\mu)$, where without loss of generality $E\mu = 1$.⁴¹ Hence those with higher μ are food richer. We assume that μ is a permanent characteristic, and that the D_j do not vary across intervals.

Food bank k 's utility in - say - the first interval U_{k1} is then given by⁴²

$$U_{k1}(\cdot) = U\left(\sum_{t=0}^T q_{k1t} + \mu_k D_1, \sum_{t=0}^T q_{k2t} + \mu_k D_2, \dots, \sum_{t=0}^T q_{kmt} + \mu_k D_m\right). \quad (1)$$

This describes the utility of a single food bank in a single interval. The principal's objective (i.e., Feeding America's preferences) is to maximize this summed over all food banks and all intervals from 0 to ∞ :

$$V = \sum_{I=0}^{\infty} \sum_{k=0}^n U_{kI}, \quad (2)$$

where U_{kI} is the utility of food bank k in interval I .

Supply In each period t , a supply of goods arrives. While the supply in a given period is potentially stochastic, we assume that in every T periods, total supply is fixed and deterministic, where q_j is the supply of good j . (This assumption is meant to capture the notion that a food bank may have little idea what will be coming on the market today, but over a month or a quarter they have a good idea of the supply of different kinds of food.) We assume that q_j is sufficiently divisible across the food banks.

The model is set up such that the supply of goods is identical across intervals. As preferences are convex, the constrained efficient solution will be a repeated outcome of the interval game.⁴³ We begin by solving the interval game. To do so, we make one (realistic) assumption about preferences in this setting.

Assumption 1: There are $l \geq 1$ goods for which $U^j = k_j$, where k_j is a constant.

⁴¹Realistically, Feeding America is aware that some foods banks are likely endowed with more food from elsewhere than others, but its need to operate transparently implies that it cannot act on this soft information.

⁴²Adding transportation costs does not change the results.

⁴³One might imagine an alternative where the principal could learn over time about μ and respond to this. However, the food bank directors were very much against using previous history of purchases to affect future allocations of shares, so we ignore it here.

In words, there are some goods which do not suffer from diminishing returns. Without loss of generality, order these so that all goods from 1 to $m - l - 1$ exhibit diminishing returns, while goods $m - l$ to m do not. The role of this assumption is to allow at least one good to transfer utility costlessly. Its plausibility is described below.

5.2 The Constrained Efficient Outcome

Begin by considering the constrained efficient outcome (we ignore the I subscript as all intervals are identical). The food banks have private information on their type μ . In equilibrium, the principal offers $\sum_{t=0}^T q_{kjt}(\hat{\mu})$ of good k to a food bank who reports itself to be type $\hat{\mu}$. Let $U(\mu; \hat{\mu})$ be the utility of a food bank that has type μ and reports type $\hat{\mu}$. The constrained efficient outcome chooses $\sum_{t=0}^T q_{kjt}(\hat{\mu})$ for all k, j and t to maximize

$$V_I = E_{\mu} \sum_{k=1}^n U(\mu; \mu), \quad (3)$$

subject to $\sum_k \sum_{t=0}^T q_{kjt}(\mu) \leq q_j$, $\sum_k q_{kjt}(\mu) \leq q_{jt}$, and

$$U(\mu; \mu) \geq U(\mu; \tilde{\mu}), \quad (4)$$

for all $\tilde{\mu}$.

Proposition 1 *The constrained efficient outcome $q_j^*(\mu)$ has the following features:*

1. For all goods $j < m - l$, $q_j^*(\mu) = (1 - \mu)D_j + \frac{Q_j}{n}$.
2. For goods $j \geq m - l$,

$$\frac{d\left(\sum_{j \geq m-l} \kappa_j q_j^*(\mu)\right)}{d\mu} = \sum_{j < m-l} D_j U^j(x_j^*). \quad (5)$$

where $x_j^* = D_j + \frac{Q_j}{n}$.

This outcome, which offers first best welfare despite the food bank's private information, has a simple interpretation. For those goods that exhibit diminishing returns, all food banks should have the same consumption per client. This implies consumption of $q_j^*(\mu) = (1 - \mu)D_j + \frac{Q_j}{n}$. This has the intuitive feature that poorer food banks (μ low) should get more of this food. All else equal, this would imply that rich food banks would claim to be poor. To overcome this, they have to be offered more of the goods that are characterized by no diminishing returns. As the change in utility on the first $m - l - 1$ goods from marginally claiming a higher μ is $\sum_{j < m-l} \frac{dq_j^*}{d\mu} U^j(x_j^*) = \sum_{j < m-l} D_j U^j(x_j^*)$, they must be compensated with an additional $\frac{d\left(\sum_{j \geq m-l} \kappa_j q_j^*(\mu)\right)}{d\mu}$ equal to this to compensate them.

The Distinction Between the Goods: We have used a distinction between those goods which exhibit diminishing returns and those that do not. This serves two purposes. The first is theoretical: to separate redistribution from efficiency. The second is empirical. The kind of goods that do not exhibit diminishing returns are typically low supply, high value foods, such as cereal, peanut butter, or rice. The reason for this is both its low supply and an indivisibility. For all food banks, there is almost never a case where all its clients get the highest valued foods. In effect, they are rationed. In a setting where these goods are rationed, welfare gains are simply the number of people served times the valuation of the first unit of the good. It is for this reason we assumed the absence of diminishing returns. This is not true for abundant, lower quality foods, where spoilage issues become more important, and clients are receiving multiple units (some of which may be thrown out, unused). This is meant not simply as a justification for the theoretical apparatus above, but in addition to make the empirical point that the food rich are likely to focus on the high quality, rare goods, and the food poor will focus on larger quantities of the lower quality goods, as we saw in the data above.

5.3 Equal Incomes and Symmetric Redistribution

Having defined efficient equilibrium for the interval game, we now show it can be implemented a Bayes Nash equilibrium of a bidding game, where food banks bid on lots of food each period using a currency which only has value within the system. In effect, this illustrates the Second Welfare Theorem in this setting.

Budgets and the Money Supply At the beginning of each interval I of the game (period 0 of that interval), food bank i is endowed with a supply of shares of s_{i0I} per capita. These shares are chosen by the principal, and can depend on the identity of the food bank but must be history independent.⁴⁴ These share can be used to bid on the quantities that are available each day. At each period t , each of the N food banks can place a bid for a unit of food, and makes a bid b_{jtI} for a unit of good j .⁴⁵ The winner of the auction is the food bank that makes the highest bid, and a price of b_{itI} in shares is paid by the winner (a first price auction).

At the end of each period, all shares spent are reimbursed to the food banks. Let p_{jtI} be the price of good j in period t (the highest bid) of interval I . The total revenue generated in period t is $R_{tI} = \sum_j p_{jtI} q_{jtI}$. At the end of each period, each food bank's stock of shares is changed by $s_{itI}(R_{tI})$. So for example, more shares could be given to one food bank rather than another. Once again, these are history independent. This

⁴⁴This was a constraint of the process where the food banks signed up for the Choice System: attempts to make share allocations depend on purchase history proved infeasible. While such reallocation could be used for redistribution purposes, it was politically infeasible.

⁴⁵We are ignoring within day variation in prices. As the equilibrium has constant prices not just within day but across days, this is to avoid yet more notation

process continues until the end of the interval. The food banks carry over their final balances to the next interval with additional transfers $s_{i0(I+1)}$. The next interval then begins, with the same steps, except that the principal can choose a different level of shares and a different way of reallocating them if she so desires.

Timing The timing of the model is as follows. At the beginning of the game, the principal commits to s_{ijI} for all i, I , and $t = 0, 1, \dots, T$. (These time periods are not indexed by I simply to cut down on notation.) In any interval I , all parties know q_j and food bank i begins with s_{i0I} shares. At each period $t \geq 0$ in that interval, all participants observe q_{jIt} and on the basis of this, they bid b_{ijtI} . If bids are tied, the winner is determined randomly. The shares are then reassigned according to s_{itI} . This ends period t , and period $t + 1$ begins. At the end of period T , consumption occurs and the next interval $I + 1$ begins. This continues ad infinitum.

Equilibrium arises here where - given s_{ijI} for all I - food bank i prefers bid b_{ijtI} for all j and t , given the equilibrium strategies of the other parties, and budget balance, where the banks can borrow at zero interest rate over time. In that setting, consider the static replication of an interval game where for all I , $s_{i0I} = 1$, $s_{itI} = \frac{R_{tI}}{N}$ for all i and t . Call these respectively “initial equal incomes” and “symmetric redistribution”.

Proposition 2 *The efficient outcome of the interval game above is an equilibrium of the bidding game with equal initial incomes and symmetric redistribution.*

This shows how it is possible to implement the efficient interval outcome, and represents the implementation of the Second Welfare Theorem.⁴⁶ As this is repeated in all intervals, it implements the constrained efficient outcome. It has the standard features of a competitive equilibrium, where equilibrium bid prices p_i^* are characterized by the solution to the set of equations $\frac{p_k^*}{p_j^*} = \frac{U_k^*(x_k^*)}{U_j^*(x_j^*)}$ and $\sum_k p_k^* q_k^*(\mu) = 1$.⁴⁷

Before turning to the testable implications of the model, it is worth taking stock here. The implication of Proposition 2 is not solely that the food allocated by Feeding

⁴⁶We have assumed above not only that food banks are price takers but also that equilibrium prices are known by all. The Choice System could have a downside if food banks held different beliefs about equilibrium prices. This unlikely for a number of reasons. First, the website offers a history of prices for similar goods, so it is relatively easy to see what the ballpark price is likely to be for a given load. Second, there are roughly 60 auctions a day, for 200 days a year, so that the food banks quickly developed a lot of experience bidding. By the time of writing, the average food bank has won over 500 auctions. This is far from a setting where auctions are infrequent, where such misperception would likely arise.

⁴⁷In this equilibrium, each food bank bids the equilibrium price $b_j = p_j^*$ until it reaches its equilibrium competitive allocation and then bids $b_k = p_k^* - \epsilon$, for ϵ small. This is simply so that those who are still buying at the end of an interval do not have any strategic power. A similar outcome would arise if food was storable across intervals without this assumption. We are not claiming that equal shares and symmetric redistribution is the uniquely optimal allocation rule. For example, one would be to choose $R_{ijt} = 0$ so the funds are never redistributed during an interval, but where everyone is given a new budget of 1 at the beginning of each interval.

American is efficiently allocated but rather that *all* food is assigned competitively. As such, the role of this section is to show that a “fringe” that is allocated through prices and consumer choice can overcome all allocative inefficiencies in the rest of the system. This section therefore points to the benefits of the Choice System as more wide reaching than if there was no outside food.

5.4 Extensions

The framework offered above has a number of assumptions which are addressed below.

Non-negativity constraints So far we have (implicitly) assumed that net supplies to a food bank can be negative, as the Choice System allows food banks to place its outside endowment on the market. If we dropped this possibility, we would have the possibility of a non-negativity constraint on net supplies. Two conditions are necessary then. First, there needs to be enough supply of the goods $j < l$ to allow equation of marginal utilities. This is straightforward. Second, there needs to be enough goods l through m to compensate the high μ types. The minimum amount of goods l through m that is needed is given by combinations of supplies q_j such that $\sum_{j \geq l} \kappa_j q_j \geq \sum_j D_j U^j(x_j^*)$. Remember from the data above that food banks equate one pound of a high quality good (such as cereal) for over 100 pounds of produce, so that relatively little physical supply of these high quality goods may be needed.

The Propensity to Consume: A characteristic of the data is that some food banks choose not to spend their shares. This arises not just temporarily but over long periods of time, and reflect not savings in the sense of holding back resources for a rainy day, but rather that their balances are not worth spending in steady state. With the existing model this is not possible, as food banks would drive up the price of that good such that balances are always exhausted. This is easily rectified by noting that food banks must pay for transportation in real money. In the Appendix, we consider a case with marginal transportation costs, where in equilibrium the food rich hold excess shares and the food poor spend their entire budgets.⁴⁸

6 Empirically Testing the Model

The model above cannot be used for empirical purposes as it has little structure. Here we simplify the model to carry out two empirical exercises. First, the data

⁴⁸To intuitively see this, consider a hypothetical outcome for the food rich if they could buy food at zero share cost. Their choice of quantities is given by the equation of the marginal utility of food and marginal transport cost. If, for those quantities, at equilibrium prices, the food rich can afford to purchase those quantities with slack in their share budget constraint, they will permanently save. This is because the marginal value of shares is 0 to them.

above show enormous dispersion in food bank choices. We begin by showing how measures of dispersion are correlated with welfare gains. Second, we use the model to offer a more precise measure of the welfare gains caused by reallocation of demand.

To do this, we simplify in two ways. First, the empirical measure requires that a demand curve be estimated. With the wide variety of goods here, this involves estimating more elasticities - own and cross-price - than is feasible. Accordingly, we assume that there are two composite goods, A and B , where good A is the lower quality good and good B the higher quality good. The lower quality goods $i = 1, \dots, m-l-1$ are characterized by diminishing returns while the high quality goods $i = m-l, \dots, m$ do not. We compute the composite by attaching linear weights to each good. Specifically, let $x_{AkI} = \sum_{i=1}^{m-l-1} \phi_i x_{ikI}$ be *per client* consumption of the composite low quality good by food bank k in interval I , and $x_{BkI} = \sum_{i=m-l}^m \phi_i x_{ikI}$ be consumption of the composite high quality good. We compute composite q_A and q_B equivalently.

Second, we assume that the return to (composite) good A is quadratic, while the return to (composite) good B is linear. Utility from food bank k having total consumption of the two goods x_{AkI} and x_{BkI} in interval I is then given by

$$U(x_{AkI}, x_{BkI}) = \tau x_{AkI} - \gamma \frac{x_{AkI}^2}{2} + x_{BkI}. \quad (6)$$

For notational simplicity, from now on the subscripts I and k are dropped unless necessary. Let the ratio of supply of good B to good A is given by δ . Consumption x_i has two parts: an endowment $D = (D, \delta D)$, where D is the supply of good A and δD the supply of good B , and consumption generated through quantities attained through the auction q_A and q_B . The relative supply of goods on the auction market is also given by δ , though this symmetry is simply to cut down on notation, and the total quantity of good A is Q .

Begin by considering constrained efficiency. First consider good A . Remember that each food bank has an outside endowment given by μ . Equating marginal rates of substitution implies that

$$q_A(\mu) = \frac{Q}{n} - \frac{(\mu - 1)d}{2}, \quad (7)$$

so that the food poor receive more of the low quality good. In equilibrium, the two parties trade off good A and good B in the ratio of κ , where $\kappa = \tau - \gamma x_A^*$, where x_A^* is first best consumption. Then all available supply of B is exhausted when $q_B(\mu) = \frac{\delta Q}{n} + \frac{(\mu-1)d\kappa}{2}$. (In the Appendix, we also compute absolute prices that would be used for bidding.)

6.1 Dispersion and Welfare

The purpose of the framework is to identify measurable implications of allowing choice over the previously used assignment mechanism. The assumption of linear demand

curves allows us to do this analytically.⁴⁹ To allow for supply responses, let the supply of food under the old system be Q' . Then as supply changes from Q' to Q , welfare would change by $\Delta U_Q = (\mu + \tau - \frac{\gamma(\delta+1)d}{2})\frac{Q-Q'}{n} - \frac{\gamma}{2}[(\frac{Q}{n})^2 - (\frac{Q'}{n})^2]$. The second benefit is through the reallocative benefit of better matching of demand to marginal valuations, and is given by

$$\Delta U_R = \frac{\gamma d^2 \sigma_\mu^2}{4}, \quad (8)$$

where σ_μ^2 is the variance of the distribution of endowments. Then the overall change in welfare is given by $\Delta U = \Delta U_Q + \Delta U_R$.

Non-negativity constraints: We are ignoring the non-negativity constraint. If supplies must be non-negative, two conditions must hold. First, Q must be high enough to allow equation of marginal utilities. Second, the competitive equilibrium implies that the food poor receive more than their share of good A . But this means that the food rich must be “bribed” by getting more of good B : otherwise they would claim to be food poor. In an online appendix (link), we characterize the equilibrium with non-negativity constraints. Here the dispersion measure of welfare gain is not variance but rather standard error.

The Composition of Food: A simplifying feature of the model above is that food banks vary in their endowments, but the variation is vertical, in the sense that some have more of everything than others do. Yet an important potential benefit of choice is that it allows food banks to react to variation in the composition of their food: so for example, a food bank may have a lot of dairy products but little cereal. In the Appendix we show that a similar exercise can be carried out where the benefits of re-equilibrating an imbalanced composition of foods can also be parameterized through a variance measure. However, it is the case that we cannot distinguish between these benefits and those outlined above, so we need to be somewhat agnostic about the sources of gain.

6.2 Estimating Welfare Gains from the Reallocation of Demand

Up to now, we have used the modeling sections of the paper only as a loose guide to understanding empirical outcomes. Here we use the model above to estimate a more precise welfare gain from firms sorting on the quality-quantity dimension. The objective is to estimate (8). While this is conceptually simple, a number of steps necessary for implementation are carried out below.

⁴⁹An alternative would be to create a hedonic for food demand, and as in Abdulkadiroglu et al, 2015, and then use that hedonic to compute a distance from an agent’s preferred point using parametric assumptions on preferences. One other possibility would be to survey participants about their perceptions of its merits, as in Budish and Kessler, 2015.

First, we characterize observed food types into high or low quality categories. To do this, we rank goods by price over the entire sample period, and then draw a line at the 50th percentile of expenditure, where high quality goods are those covered by the top half, and low quality goods are in the bottom half of expenditures. Using this ranking, low quality goods are produce, beverage, juice, dairy, protein, canned fruit, canned vegetables, and snacks. High quality goods are the remaining goods. Remember that good i has weight ϕ_i in the composite good, and we do not directly observe ϕ_i . However, in equilibrium it must be the case that $\frac{\phi_i}{\phi_j} = \frac{p_i}{p_j}$. As a result, we can compute the two goods as price weighted indexes where the quantity of good A for food bank k in interval I is given by $q_{AkI} = \sum_{i=1}^{m-l-1} \frac{\bar{p}_i q_{ikI}}{\bar{p}_A}$, where $q_{BkI} = \sum_{i=1}^{m-l-1} \frac{\bar{p}_i q_{ikI}}{\bar{p}_B}$, and where \bar{p}_i is the average historical price of good i from 2005 to 2011 and \bar{p}_A is the average quantity weighted price of good A from 2005 to 2011. We also need to choose an interval over which food banks seek a balanced portfolio of food. Here we choose two months.⁵⁰

Given this, the objective can be seen visually in Figure 12, where demand for the low quality good is plotted. The efficient outcome is a price for good A given by p_A^* , where all clients consume x_A^* in a two month interval. However, under the old system, the food poor consume less \underline{x}_A and the food rich consume more \bar{x}_A . The welfare gain from the change is then given by the areas $A + B$.

There are a number of challenges here.

1. First, it should be intuitively clear from the figure that the welfare gain depends on the dispersion in outside endowments, as in (8). Unfortunately, we cannot directly observe d or σ_μ^2 . However, we do observe something highly correlated, namely the choice of purchased food in the Choice System. Let the *total* quantity of food obtained be $q_k = q_{Ak} + q_{Bk}$ (we drop the interval subscript for simplicity). In the Appendix, we show that

$$\Delta U_R = \frac{\gamma \text{Var}(q_k)}{2(1 - \frac{p_A}{p_B})^2}. \quad (9)$$

As a result, the value of better allocations from choice is linear in the variance of average quantities.⁵¹ Hence we can relate welfare gains to measures of

⁵⁰This is probably longer than the usual food bank would use. However, we wish to be conservative because food comes in the indivisible unit of a truck and shorter intervals could induce variance simply on the dimension of whether the food bank won an auction or not. Choosing a longer interval smooths this out.

⁵¹In the analysis above, we have simply posited that the old allocation system offered an equal number of pounds to all food banks per client. This allowed us to measure the value of the Choice System by the variance of quantities. In reality, food banks could turn down a food offering under the old system, and sometimes did so even if it meant going to the bottom of the line. This is relevant as any deviations from the algorithm changes the analysis above in one simple way: with linear demand curves, the welfare gain through better allocations in the Choice System are no longer parameterized by $\text{Var}(q)$ but rather $E(\nu^2)$, where ν is the difference in average quantity by a food bank between the Choice System and the old system.




Figure 1: Welfare
35

dispersion.

2. The second challenge is that we do not have a numeraire that is easily intelligible. Unlike usual welfare exercises in industrial organization where one can interpret these in dollars, here they are denominated in shares, which are hard to interpret. To overcome this, we normalize by first best per client welfare, $U^* = \tau x_A^* - \frac{\gamma x_A^{*2}}{2} + x_B^*$. A potential problem is that we do not observe τ (γ will be estimated as the slope of the demand curve). However for the quadratic, we can recover τ as in equilibrium the relative price level is the marginal rate of substitution $\frac{p_A}{p_B} = \tau + \gamma x_A^*$, and hence $U^* = \frac{p_A}{p_B} x_A^* + \frac{\gamma x_A^{*2}}{2} + x_B^*$. Hence we compute welfare gains as % changes in per client welfare relative to the welfare they receive in the efficient outcome.
3. Our numeraire is in terms of consumption x_i but we do now observe outside endowments. We do not have a precise estimate of the fraction of all food received by food banks that comes from Feeding America. The fraction of all food that food banks receive from Feeding America yellow pounds which we call f is not precisely know, but the usual estimates provided are in the range of 10% to 15%. In order to be conservative here, we assume that they receive 10% of all food through this mechanism.
4. The parameter γ measures diminishing returns to low quality food. In equilibrium, $\gamma = \frac{d(\frac{p_A}{p_B})}{dq_A}$. We begin by estimating γ . To do this, consider the following relationship predicting winning bid price:

$$p_{it} = \beta X_{it} + \gamma_{Ah} E[Q_{At}] + \gamma_{Al} E[Q_{At}] D_l + \gamma_{Bh} E[Q_{Bt}] + \gamma_{Bl} E[Q_{Bt}] D_l + \epsilon_{it}, \quad (10)$$

where E is the expectations operator, Q_{it} is the quantity of composite good i in the interval to which time period t belongs, D_l is a dummy variable equal to 1 only if the good is low quality, and X is a set of other regressors. Note however, that the model is denominated in outcomes per client. We do not have data at the client level, nor indeed do we know the number of clients. Hence this estimates $\hat{\gamma} = \gamma C$ where C is the number of clients. This allows us to recover how prices change with expected supply, where $\hat{\gamma} = \frac{p_A + \gamma_{Al}}{p_B + \gamma_{Ah}} - \frac{p_A}{p_B}$. To overcome the per client issue, we normalize by first best welfare which also is linear in clients.

The identification assumption here is that variation in Q_{it} is independent both of demand variation and also of other endowments.⁵² We use one period lagged supplies (i.e., total number of pounds in the previous two months) to measure expected supply. The novelty of this over the previous regressions is through the supply terms. This OLS regression is carried out in Table 9, and shows

⁵²Of the two, the second is likely the more restrictive. We discuss robustness below.

that the prices of high quality goods (low quality goods) rise (fall) when there is more supply of low quality goods, though the effect on high quality goods is only significant at the 10% level.⁵³

As a result of the issues above, let q_{iI} be the *per client* number of pounds for food bank i in interval I (i.e., two month period I) and let q_j be the average per period total number of pounds of type $j, j = A, B$, received by food banks from 2005 to 2011. We show in the appendix that

$$\frac{\Delta U_R}{U^*} = \frac{\hat{\gamma} E_T (Var(q_{iI}))}{2(1 - \frac{p_{At}}{p_{Bt}})^2 \left(\frac{p_A}{p_B} \frac{q_A}{f} + \frac{\hat{\gamma}(\frac{q_A}{f})^2}{2} + \frac{q_B}{f} \right)}, \quad (11)$$

where $f = 0.1$ is the fraction of all food received that derives from Feeding America.

We do not observe per client allocations, so to empirically implement this, let G_{iI} be a food bank i 's goal factor in interval I , and let \bar{G}_I be the average goal factor in interval I . Further let Q refer to the total number of pounds in an interval. We show in the Appendix that (11) can be implemented as

$$\frac{\Delta U_R}{U^*} = \frac{\hat{\gamma} E_T \left(\bar{G}_I^2 \sum_{i=1}^N \frac{G_{iI}}{\bar{G}_I} \left(\frac{Q_{iI}}{\bar{G}_I} - \frac{Q_I}{\bar{G}_I} \right)^2 \right)}{2(1 - \frac{p_{At}}{p_{Bt}})^2 \left(\frac{p_A}{p_B} \frac{Q_A}{f} + \frac{\hat{\gamma}(\frac{Q_A}{f})^2}{2} + \frac{Q_B}{f} \right)}. \quad (12)$$

Using (12), we compute the welfare gains from reallocation of demand at 8.89% of first best surplus from the 3 billion pound food distribution system, equivalent to roughly 260 million pounds of food. (The details are in the Appendix.) This reflects the conceptual point above about how allowing choice on a fringe (here 10%) can generate large allocative welfare gains by allowing these marginal choices to redress large imbalances elsewhere. In this way, these pounds can be leveraged to have impacts far beyond simply their own pounds. Here the multiplier on efficiency is about 0.89, where the value of the 300 million pounds in the Choice System is an additional 260 million pounds. This is equivalent to a complete full day's food for an additional 150,000 people every day. This large additional benefit arises because food banks are so far from their desired points - the huge variance in outcomes across food banks reflects this - that the ability to take extreme positions through the auction helps to reverse these inefficiencies in empirically substantive ways.

Much has been made above about how food banks take advantage of choice to permanently sort on their other endowments of food, rather than to manage temporary shocks. To see the importance of this in this context, we recompute these welfare

⁵³These magnitudes are very small - we compute $\hat{\gamma} = 4.6e^{-7}$. This number is to interpret, but translates to a price elasticity of -0.8. This is robust to alternative assumptions on expected supply, such as current supply, one month lagged, and previous quarter.

gains taking out the temporal variation within a food bank’s choices. In effect, this measures the welfare gain from the average choice made by a food bank from 2005 to 2011 compared to a situation where they simply consume the population mean. The welfare gain here is 82% of the number above, at 7.29% of first best surplus.⁵⁴ This once again shows the importance of food banks’ ability to sort on the quality-quantity dimension.⁵⁵

Identification and Potential Bias Here we address the robustness of the results above. First, we consider the impact of mis-specifying the slope of the demand curve. The estimate above assumes that changes in observed supply of the composite goods are uncorrelated with other relevant factors. The most relevant case where this could be untrue is where the supply of food that Feeding America receives is correlated with supply of food from other sources.⁵⁶ This is problematic as it inflates the estimates of the slope of the demand curve in non-trivial ways. As an example, with our assumption that $f = 0.1$ (so that 90% of food comes from elsewhere), let the correlation between food received from Feeding America and elsewhere be say 0.3. Then the demand curve estimate above is too high by a factor of $1 + 9(0.3) = 3.7$. In most circumstances, this would cause serious concerns about the credibility of the results. This is not so here. The reason is that we are estimating surplus gain relative to first best welfare, and both change similarly with γ . Specifically, in the case where the correlation coefficient is 0.3, this would reduce the welfare gains from 8.89% above only to 8.77%. Hence this estimate is relatively robust to even quite severe mis-specification.

Another possible source of potential bias is that the benchmark that we are using is the rule of the old algorithm, where each food bank gets an equal number of pounds per client. Yet while the correlation between $\log(\text{pounds})$ and $\log(\text{goal factor})$ was 1 in Table 2, some food banks received more food than that and some received less. If those food banks who received more pounds than predicted before 2005 are also those who have got more pounds afterwards under the old regime, our estimates could be inflated. A natural way to deal with this is to predict what would have happened after 2005 with the old algorithm as an extrapolation of their previous allocations.

⁵⁴One small tweak is necessary here. Permanent consumption is measured as the average over the 33 two month intervals from 2005 to 2011. Then if a food bank temporarily buys one pound of food in an interval, it increases the permanent average by 0.03 pounds. As a result, the 82% number above should be reduced to something closer to 79%.

⁵⁵It is also worth noting the sensitivity to changing the assumption on f . If the Choice system’s 300 million pounds only allocated 5% of all food received by food banks, the change in welfare is 2.26% of 6 billion pounds or 135 million pounds of food, while if the Choice System allocated 20%, the welfare gain would be 23.6% of 1.5 billion, or 355 million pounds of food.

⁵⁶While much food receipt is likely idiosyncratic, a natural example where there is correlation is where a large crop of produce causes not only donations to Feeding America to rise, but also those from elsewhere. Note also that this is also a reason why we do not instrument for supply by seasonality. While this may potentially be a good instrument for overall supply, it surely is not an effective way to isolate supply to Feeding America.

In words, those who got more pounds before 2005 would get more pounds after that. While this is a potentially appealing approach, it faces a data constraint, namely that we do not observe the kind of food that food banks receive before 2005. As a result, we cannot compute the composite goods in the way we do above. It is possible to carry out a separate exercise, however, which is to model the food banks' preferences where diminishing returns are independent of food type for the low quality foods, in which case all we need to know to estimate the demand curves and welfare is total number of pounds in each composite. When we carry out this exercise, we compute a welfare gain of 10.70%, in the similar ballpark to the results above.

Remember that in the modeling sections we have ignored strategic bidding. Yet even in dynamic auction settings where there are many potential consumers, there is some evidence of strategic bidding, where bidders shade their bids below their true valuations (Backus and Lewis, 2016). This has the impact here of underestimating true consumer surplus as bids do not reflect marginal valuations, but rather something lower. Such benefits are ignored here.

One final point is worth making here: the only benefits modeled above from reallocation of food received is the substitution between a high quality composite good and a low quality composite. This allows no gains from reallocating within goods of similar price. Yet diminishing returns likely arise at the individual good level. For example, *Yoghurt* and *Milk* are both dairy products in the categories below and sell for similar prices. In the model above - and the results below - they would be bundled as a composite good with perfect substitution between them. However, in reality a food bank that already has milk benefits from the ability to reallocate demand within categories by focusing its purchases on yoghurt. Such reallocative benefits within narrow categories which sell for similar prices will not be picked up here.

7 Money Supply

This system operates using the constructed currency of shares. Feeding America controls the supply of shares, and an important issue is what governs that supply. The view of the design group was that share issues should facilitate price transparency. A concern of a system such as this, with constructed currency, is that participants may find it difficult to know how much to bid for an item. Individual food banks typically know nothing about the aggregate supply of shares in the system: all they can see are their balances and the prices of transacted lots. In order to aid bidders in making bids, it was felt that the historical price of a particular good should provide strong information about a reasonable price now. As a result, the desire was to choose share supply to generate zero inflation for a given good if demand and supply conditions are

unchanged.⁵⁷ To that end, the system was designed such that the historical record of previous prices would be a strong reflection of current valuations, all else equal. In order to do this, Feeding America adjusts the money supply every year proportionate to the number of pounds supplied to the system. This does not control for changes in the velocity of transactions, nor changes in the quality of food being offered to the Choice System, but would at least allow some adjustments based on total donation of pounds to the system.⁵⁸

The Price Level Figure 11 shows the average number of shares paid for a pound of food over time. It has shown considerable fluctuation over time, with considerable variation around its mean of 0.29. Notable is the impact of the financial crisis: in 2008, the price of a pound of food jumped to 0.4, and has gradually reduced back to “normal” levels since. These effects also arose in the multivariate regressions in Table 1. In 2008, the price of a pound of food was 130% higher than in 2006. Also of relevance below is the interaction between goal factors and the year dummies, which show that for 2008, it was the large food banks that were paying highest prices. It is difficult to evaluate the monetary objective of transparent prices from this: instead, it represented the kind of very unusual event where prices should change. Specifically, this picks up the large impact that the financial crisis had on the food banking industry: demand rose, and supply fell, so that prices for goods rose as economics would predict. This is not a case where we believe that the supply of shares should change to eliminate this, as this is the kind of inter-temporal shock that should be reflected in higher prices.

One possible point is worth making here, however. A price rise for a good implies that the value of that good increases relative to some alternative. So if the price of a TV rises, it implies that the value of TVs relative to alternatives goes up. Yet remember that shares have no value outside the system, so what is the alternative? The most natural alternative is purchases in the future, where food banks changed the timing of their purchases. One way in which this arose was the use of credit. Yet

⁵⁷This was felt to be particularly pertinent in the context of leveling the playing field for the “little guy”. The reason for this is that smaller food banks may bid on a particular item (bread in Massachusetts, for example) quite infrequently compared to larger food banks, and this may give larger food banks an advantage in bidding as they know better how to compute reasonable bids.

⁵⁸Keeping prices constant “all else equal” is non-trivial. Consider the simplest Quantity Theory of Money $MV = PT$, where M is the money supply (here shares), V is the velocity with which a share is transacted, P is the price level, and T the quantity of transactions. The desire here was to try to ensure that $\dot{P} = 0$, all else equal. For example, suppose that T doubles from one year to the next. If the number of shares is left unchanged, prices would likely deflate by 50%. Hence, the number of shares would need to be scaled by the supply of goods on the market. Yet T is value weighted, so it may not be enough simply to scale the supply of shares by the number of pounds in the system. Finally, there may be variation in V , the velocity with which the shares are trading. Velocity changes for many reasons: as one example, there is considerable variability in savings rates over time. As another example, food banks are liable for all transportation costs, and changes in gasoline prices can have a first order effect on the marginal propensity to consume from shares.

credit is so short term in the Choice System (it is paid off typically in a month) that it likely has no effect on aggregate prices. Instead, the likely effect of the financial crisis was to run down savings.

Daily Reallocation of Shares All shares that are spent in a given day are reallocated at midnight. The shares are reallocated according to the same goal factor formula, where those in greatest need are topped up at a greater rate than those who are less needy.

There were a number of reasons for this. First, to maintain constancy of average share balances.⁵⁹ Second, one of the biggest conceptual hurdles faced in this process was to inculcate in the minds of the food banks that they are the owners of the food being donated, and not Feeding America. In effect, they are not only the buyers of the food, but also the sellers. This became particularly pertinent when concerns were raised about the danger that only the large, food rich food banks would receive the most desirable items, as they would bid more than any other food bank. This was seen by many on the task force as inherently unfair: that the large food rich banks would get the “good stuff”, leaving the rest for the others. This concern became mitigated when it was pointed out that the beneficiary of these high priced sales was not Feeding America but rather the rest of the food banks. This is because those shares would be reallocated to everyone else at midnight. The author remembers one of the other food bank directors on the committee joyfully pointing out “so if Los Angeles bids us out of the market by paying a fortune for a truckload of frozen chicken, we really get their shares that night?” That sense of ownership through the reallocation of shares indirectly helped buy-in across the food bank network.

8 Conclusion

Seen from afar, the idea that a specialized currency could be used to allocate food more efficiently while simultaneously respecting the relative level of need in an area may seem straightforward. However, despite the conceptual simplicity of the solution, it is worth pointing out that it is very rare to observe these kind of “Monopoly money” solutions being used to allocate resources in real world settings. There are, of course, a large number of barter markets which involve the trading of scrip, but these are sparse and characterized by rampant illiquidity. Indeed, as one of the only examples offered of such mechanisms is bidding for business school courses, this surely points to the limited empirical importance of these kind of solutions.

So why did this use of fake money seem to work? It should be clear that the setting in which food is assigned lends itself to specialized currency. First, it is an infinitely

⁵⁹This was to avoid the problem of food banks budgeting over any discrete time interval, with the danger of running out of money at the end of the month, or spending too much at the beginning of the month, much like the evidence on Food Stamp use (Shapiro, 2005, Hastings and Washington, 2010).

repeated setting, which matters as it leads food banks who do not currently value a good efficiently to hold onto their fake money for better offerings in the future. In a one shot setting, this would not arise as the shadow value of the currency becomes zero. Second, the flow of offerings is high: the average food bank that turns down an offering today does not have to wait long to get something that it likely truly wants at market prices. This would be difficult to implement on say an annual cycle where an entity has to wait a year for its next chance. The frequency of offerings also leads to quick learning by the bidders: by now, the average food bank has won 500 auctions so the bidding process has become second nature to them. Third, there is wide variation in marginal valuations both across goods and participants. Fourth, there are a large number of players, which renders some strategic issues largely moot. For all these reasons, its is a conceptual terrain which lends itself to such a setting.

However, it may be that it is not the broad match of the concept of fake currency to the problem that generated its likely success, but rather the myriad of small details that got it over the line. Here these details involved a series of tweaks - simple bidding mechanisms, credit, negative prices, the opportunity to delegate bidding, a fairness committee, the ability to bid jointly, the daily reallocation of shares, the use of a fully functioning demonstration game, and so on - that made the difference. That some of these buttresses were not ultimately necessary may hardly be the point, as much of the implementation of this system was political.

The apparent success of the Choice System raises other possibilities. First, could it be extended to other parts of the food distribution chain? For example, the Chicago Food Depository distributes food to many parts of the city, some of which are blighted with greater poverty than others. Would it be possible to set up a system of fake currency to do better than charging food pantries real money? As one possibility, could they give credit cards denominated in fake currency with which to distribute food from its warehouses? While this has its challenges - not least the fact that many clients of food banks commute from where they live to a food bank in another neighborhood - the outcomes of the Choice System may open some possibilities.

Second, one objective of the system is to rebalance the field in favor of the food poor. This has arisen not through any explicit system of redistribution - other than the 10% tax on maroon pounds - but rather by the sorting of the market. However, one thing that seems clear is that the food rich are sorting into the most expensive products. This raises the possibility of a luxury tax on high priced goods, where the revenues of high quality goods would be redistributed not to all food banks based on goal factors, but disproportionately redistributed to those food banks who buy low priced goods.

The issue of nutrition may also be worth further exploration. The use of bidding has also likely facilitated a new form of sorting that has occurred over the last decade. A new trend among some food banks is to be less focused on volume of food for the poor, but instead nutrition has become the focus of many food banks. (It is, of course, ridiculous to claim that any food bank is not concerned with nutrition:

instead, a recent trend is for some food banks to focus much more intensively on this issue.) For these food banks, their bidding is now more focused on a set of foods which Feeding America has labeled “Foods to encourage”, those with the highest nutritional value. Yet this trend is far from universal, and for many food banks their priority remains the alleviation of hunger through a wide variety of foods. This divergence in preferences would have been very difficult to administer under a centralized assignment system. The Choice System allows this divergence to be naturally reflected in different offerings to the poor across geographic areas.

Despite the apparent success of this allocation system - with more food being better allocated across the country - the Task Force was far from omniscient. The most substantive problem remains produce, which traded on the Choice System exactly like any other good. Produce is problematic as it spoils quickly.⁶⁰ As a result, it is a relatively low value food to food banks, especially as transportation costs can be large. The Choice System does take time; at least a day to sell, and then it needs to be transported to a food bank, and from there to a pantry or soup kitchen. As I write, a decision has been made at Feeding America to take produce from the Choice System, and reallocate it to a new platform that will allow it to move more quickly, where the food is simply given to the food bank that can collect it fastest. Perhaps it would have been valuable to adapt the Choice System for goods that need to transact rapidly. As such, the Choice System has not been a panacea for all ills. Despite this, we believe that its architecture has led to some robust successes that may be valuable for other possible applications in the not-for-profit sector.

⁶⁰Produce is also a problem as it is not clear how much is actually used by end users. Sometimes this is because they do not know well how to cook certain kinds of food. This part of the reason that some food pantries have moved towards prepared meals over the last decade.

REFERENCES

- Abdulkadiroglu, Atila, Nikhil Agarwal, and Parag A. Pathak (2015)**, “The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match”, mimeo, MIT.
- Aleksandrov, Martin, Haris Aziz, Serge Gaspers and Toby Walsh (2015)**, “Online Fair Division: Analysing a Food Bank Problem”, mimeo, UNSW, Sydney.
- Backus, Matthew and Greg Lewis (2016)**, “Dynamic Demand Estimation in Auction Markets”, NBER Working Paper 22375.
- 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.
- Budish, Eric and Jud Kessler (2015)**, “Experiments as a Bridge from Market Design Theory to Market Design Practice: Changing the Course Allocation Mechanism at Wharton”, mimeo, Booth School of Business.
- Hastings, Justine and Ebonya Washington (2010)**, “The First of the Month Effect: Consumer Behavior and Store Response”, *American Economic Journal: Economic Policy*, Vol 2: 142-162.
- Prendergast, Canice (2016)**, “Bureaucratic Responses” *Journal of Labor Economics*, Vol 34, S183-S215.
- Roth, Al (2008)**, “What Have We Learned From Market Design?,” *Economic Journal*, Vol 118, 285-310.
- Rustichini Aldo, Mark A. Satterthwaite and Steven R. Williams (1994)**, “Convergence to Efficiency in a Simple Market with Incomplete Information”, *Econometrica*, Vol 62(5), pp 1041-1063.
- Shapiro, Jesse (2005)**, “Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle?,” *Journal of Public Economics*, Vol 89, 303-325..

TASK FORCE MEMBERS (AND THEIR AFFILIATIONS AT THE TIME):

John Alford, Chair, Mississippi Food Network, Jackson, MS
John Arnold, Second Harvest Gleaners FB, Grand Rapids, MI
Al Brislain, A2H Senior Vice President, Affiliate Services
Bill Clark, Philabundance, Philadelphia, PA
Phil Fraser, SHFB of Santa Clara/San Mateo, San Jose, CA
Maria Hough, A2H Managing Director of Logistics
Mike Halligan, A2H Senior Vice President, Business Development & Logistics
Brenda Kirk, Houston Food Bank, Houston, TX
Rob Johnson, Atlanta Community Food Bank, Atlanta GA
Susannah Morgan, Food Bank of Alaska, Anchorage, AK
Steve Sellent, Great Plains Food Bank, Fargo, ND
Roger Simon, The Idaho Food Bank, Boise, ID
Advisors:
Harry Davis, Professor, University of Chicago
Don Eisenstein, Professor, University of Chicago
Robert Hamada, Professor, University of Chicago
Canice Prendergast, Professor, University of Chicago

Changes to the Goal Factor Under the old system, this was a weighted average of population and poverty of an area relative to the national average for food bank i via

$$[(Population_i/USPopulation) + (PovertyPopulation_i/USPovertyPopulation)]. \quad (13)$$

This was changed to one which applied empirical weights based on usage, as many over the national poverty level use food pantries and soup kitchens. The new definition has three components: those under the poverty line, those between the poverty line and 185% of the poverty line, and those above 185% of the poverty lone, using usage weights for the three groups. The formula is now given by

$$\frac{0.73(Pop < 100\%Poverty)_i + 0.22(Pop[> 100\%but < 185\%])_i + 0.05(Pop > 185\%)_i}{0.73(U.S.Pop < 100\%) + 0.22(U.S.Pop[> 100\%but < 185\%]) + 0.05(U.S.Pop > 185\%)}. \quad (14)$$

Table 1: (Log) price per pound paid by a Foodbank (2005-2011)

	<i>Dependent variable:</i>	
	log(pricePerPound)	
	(1)	(2)
Constant	-0.975*** (0.068)	-0.978*** (0.068)
log(goalFactor)	0.281*** (0.008)	0.224*** (0.043)
typeDiapers	-0.092 (0.108)	-0.083 (0.108)
typePasta	-0.143 (0.116)	-0.148 (0.116)
typePHSE	-0.195*** (0.068)	-0.194*** (0.068)
typeMeal	-0.405*** (0.054)	-0.405*** (0.054)
typeRice	-0.179 (0.181)	-0.181 (0.181)
typeMeat	-0.773*** (0.070)	-0.780*** (0.070)
typeBP	-1.011*** (0.069)	-1.011*** (0.069)
typeClean	-1.128*** (0.055)	-1.135*** (0.055)
typeProtein	-1.446*** (0.106)	-1.453*** (0.106)
typeNF	-2.569*** (0.097)	-2.559*** (0.097)
typeFruit	-1.026*** (0.167)	-1.039*** (0.167)
typeCond	-1.632*** (0.067)	-1.630*** (0.067)
typeSnack	-1.521*** (0.046)	-1.518*** (0.046)
typeOther	-1.680*** (0.045)	-1.678*** (0.045)
typeHBC	-1.907*** (0.067)	-1.901*** (0.067)
typeVeg	-2.094*** (0.069)	-2.095*** (0.069)
typeNutri	-2.335*** (0.075)	-2.338*** (0.075)
typeBaby	-2.747*** (0.115)	-2.741*** (0.115)
typeDairy	-2.320*** (0.060)	-2.316*** (0.060)
typeJuice	-2.392*** (0.058)	-2.396*** (0.058)
typeBev	-3.978*** (0.049)	-3.974*** (0.049)
typeFresh	-4.964*** (0.040)	-4.961*** (0.040)
year dummies	-0.207*** (0.054)	-0.198*** (0.055)
year2007	yes	yes
month dummies	yes	yes
log(goalFactor):year2006		0.064 (0.048)
log(goalFactor):year2007		0.113** (0.047)
log(goalFactor):year2008		0.151*** (0.048)
log(goalFactor):year2009		0.078 (0.048)
log(goalFactor):year2010		-0.002 (0.048)
log(goalFactor):year2011		-0.087* (0.049)
Observations	47,850	47,850
R ²	0.411	0.412
Adjusted R ²	0.411	0.412
Residual Std. Error	2.110 (df = 47809)	2.108 (df = 47803)

Note:

*p<0.1; **p<0.05; ***p<0.01

FEEDING AMERICA CHOICESystem® Logout | User Guide

Welcome

- Bidding Session
- Reports
- Maroon Load
- Tools
- Offers & Receipts
- Bulletin Board
- Support

Time: 8:36 CT

- Morning session closes at 12:00 CT
- Afternoon session closes at 16:00 CT

FEEDING AMERICA

	Load #	Loads	Description	Subsidy	Weight	Point Of Origin	Your Bid	
	L144900	1	Bev, Gold Peak, MinuteMaid	\$0	27347	Wac,TX-76732		Bid
	L144778	3	Dairy, Silk Cashew 6/64 oz.	\$0	43808	Mount Crawford,VA-22841		Bid
	L144780	1	Dairy, Silk Cashew 6/64 oz.	\$0	43054	Mount Crawford,VA-22841		Bid
	L144779	1	Dairy, Silk Cashew 6/64 oz. & Ture Almond	\$0	43498	Mount Crawford,VA-22841		Bid
	L144786	1	Dairy, Silk, Internatinal Delight Ice Coffee	\$0	24532	Portland,OR-97230		Bid
	L144785	1	Dairy, Silk, Internatinal Delight Ice Coffee	\$0	43394	Portland,OR-97230		Bid
	L144902	1	DAIRY, YOPLAIT Yogurt various flvrs,4-6oz cups	\$0	35925	Allentown,PA-18106-8803		Bid
	L144901	1	DAIRY, YOPLAIT Yogurt various flvrs,4-6oz cups	\$0	27981	Allentown,PA-18106-8803		Bid
	L144899	1	HBC, Asst P&G HBC Damaged and Disco	\$0	8188	West Branch,IA-52358		Bid
	L144861	1	HBC, EQUATE COLD & HOT PATCH	\$0	28372	Chicago,IL-60638		Bid
	ML7691	1	NF, Emoty Banana Boxes	\$0	3900	Colorado Sorinos.CO-80915		Bid

Available Shares **60699** Available Credit Shares **57000** Total Shares Available **117699** [Printable View](#) [Export To Excel](#)

My Current Bids

Load #	Load Description	Gross Weight	Bid Amount	Shares Used	Credit Used	Transportation	Multiple Ship To	Ship Comments

Figure 2: Website

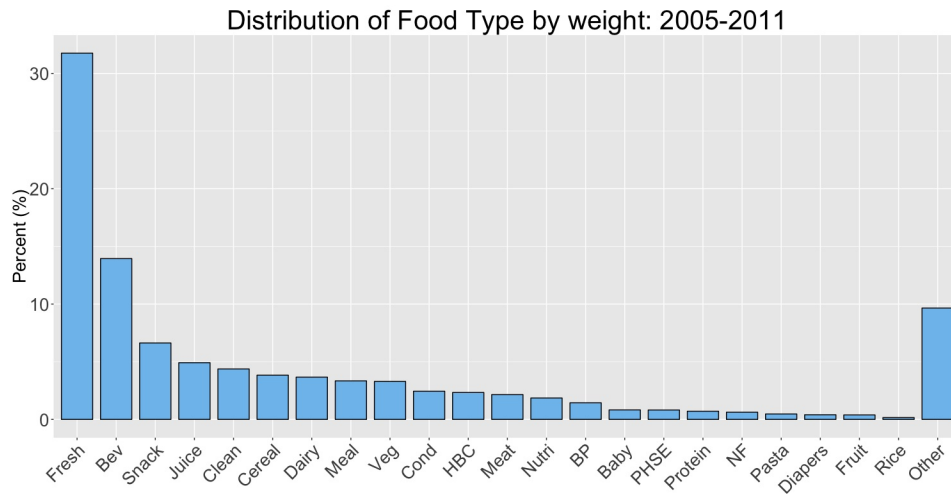


Figure 3: **The Distribution of Food**

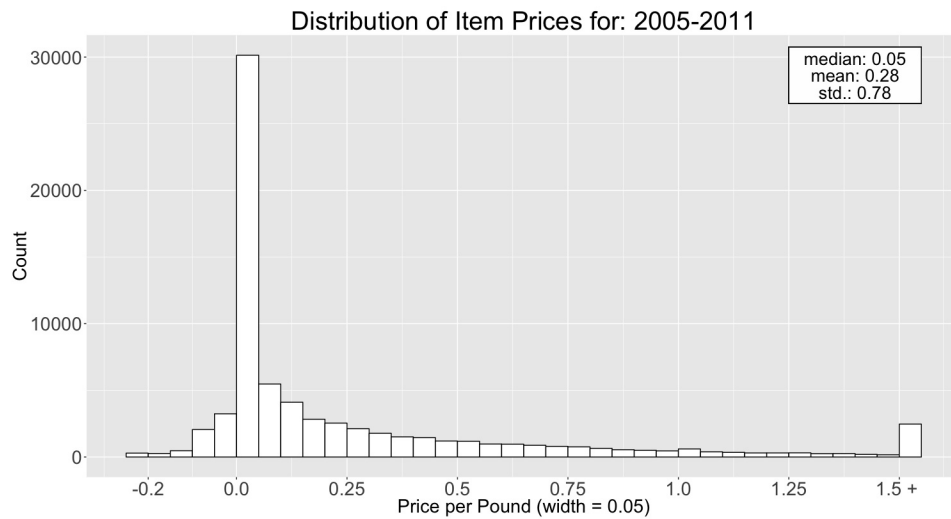


Figure 4: **Price per pound**

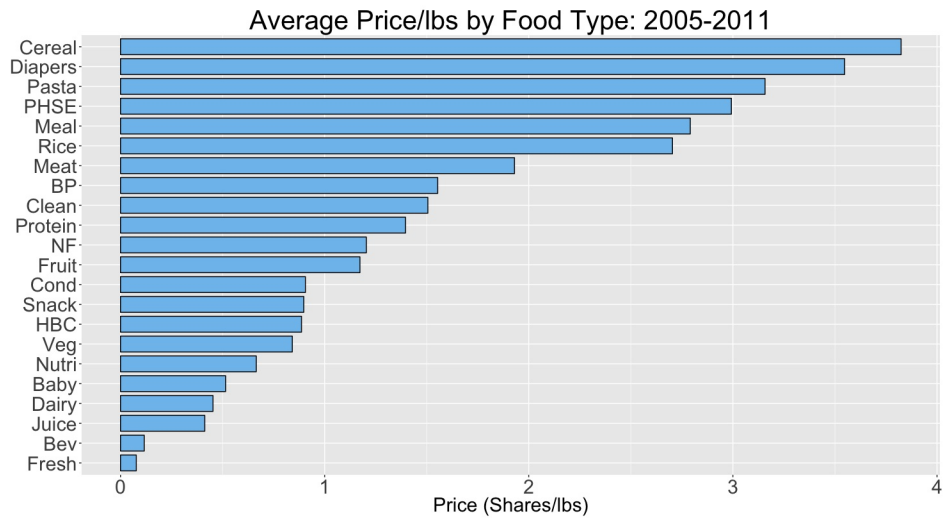


Figure 5: Prices of Different Foods




Figure 6: The Distribution of Average Prices Paid by Food bank

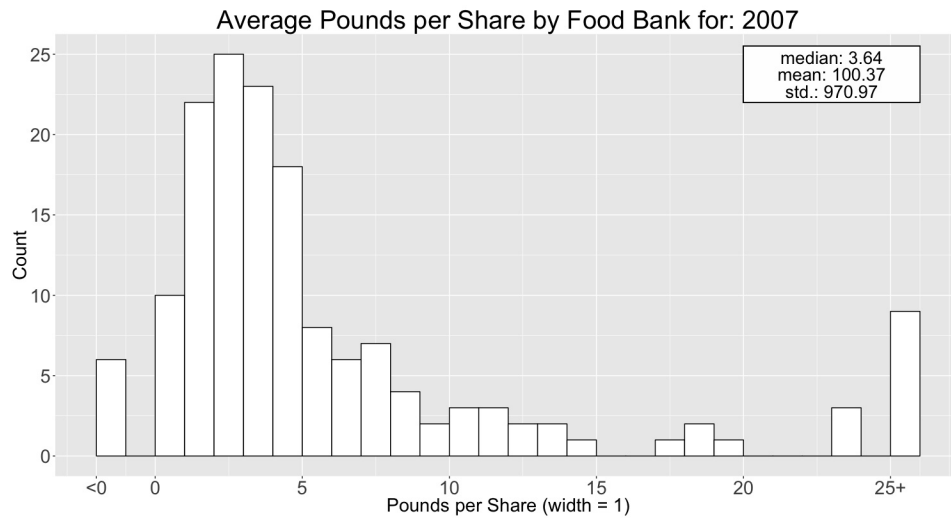


Figure 7: The Distribution of Average Prices Paid by Foodbank in 2007




Figure 8: Ratio of Total Expenditures to Total Shares from 2006-11

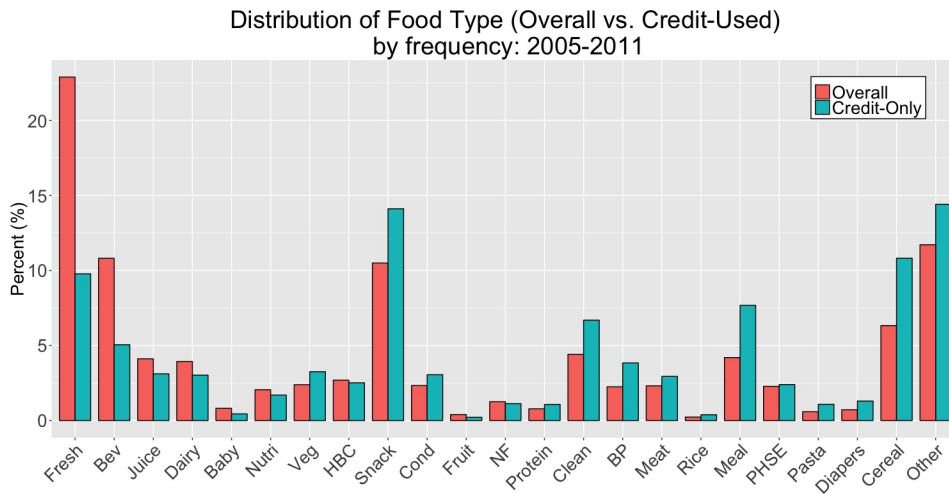


Figure 9: Products bought with Credit




Figure 10: Total Pounds in the System

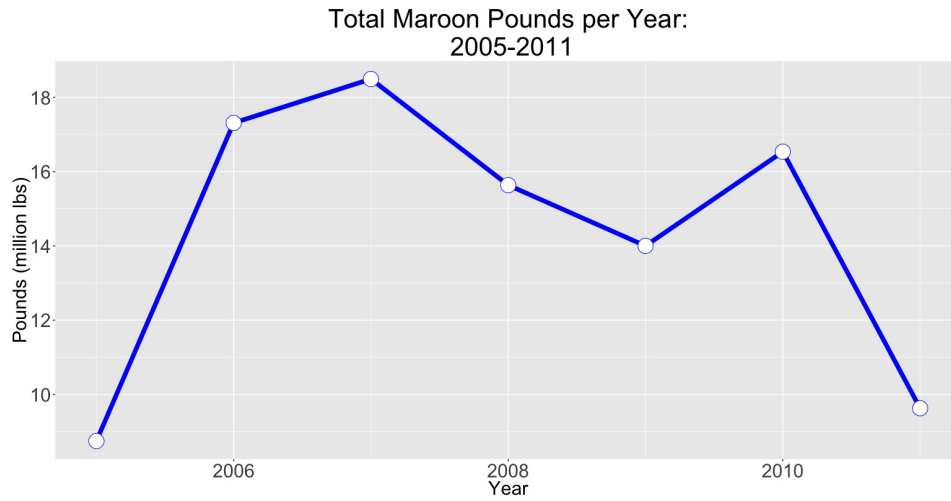


Figure 11: Maroon Pounds



Figure 12: Price per pound

Table 2: Log pounds for a food bank: 2002 to 2004⁶¹

	<i>Dependent variable:</i>		
	log(yellowPounds)		
	(1)	(2)	(3)
Constant	14.148*** (0.067)	14.112*** (0.106)	14.136*** (0.117)
log(goalFactor)	1.011*** (0.047)	1.011*** (0.047)	1.044*** (0.082)
year2003		-0.004 (0.140)	-0.013 (0.165)
year2004		0.111 (0.140)	0.048 (0.165)
log(goalFactor):year2003			-0.013 (0.115)
log(goalFactor):year2004			-0.084 (0.115)
Observations	293	293	293
R ²	0.536	0.537	0.538
Adjusted R ²	0.534	0.532	0.530
Residual Std. Error	1.068 (df = 291)	1.071 (df = 289)	1.073 (df = 287)

Note:

*p<0.1; **p<0.05; ***p<0.01

⁶¹To compare between before (2002-2004) and after (2005-2011), we further subset the data to exclude food banks with zero/missing yellowPounds and goalFactors in any given year from 2002-2011. This reduces the number to 100 food banks

Table 3: Log pounds for a food bank: 2005 to 2011⁶²

	<i>Dependent variable:</i>		
	log(yellowPounds)		
	(1)	(2)	(3)
Constant	14.188*** (0.043)	14.226*** (0.098)	14.282*** (0.111)
log(goalFactor)	0.432*** (0.030)	0.434*** (0.029)	0.501*** (0.070)
year2006		0.210 (0.134)	0.105 (0.158)
year2007		0.105 (0.134)	0.024 (0.159)
year2008		0.064 (0.134)	0.014 (0.158)
year2009		-0.121 (0.134)	-0.189 (0.159)
year2010		-0.257* (0.134)	-0.327** (0.159)
year2011		-0.254* (0.134)	-0.273* (0.159)
log(goalFactor):year2006			-0.126 (0.099)
log(goalFactor):year2007			-0.099 (0.105)
log(goalFactor):year2008			-0.057 (0.110)
log(goalFactor):year2009			-0.083 (0.109)
log(goalFactor):year2010			-0.086 (0.107)
log(goalFactor):year2011			-0.019 (0.107)
Observations	693	693	693
R ²	0.236	0.259	0.262
Adjusted R ²	0.235	0.252	0.248
Residual Std. Error	0.953 (df = 691)	0.942 (df = 685)	0.945 (df = 679)

Note: *p<0.1; **p<0.05; ***p<0.01

⁶²Same as ⁶⁰

Table 4: Average annual (log) price paid by a food bank (2005-2011)

	<i>Dependent variable:</i>	
	log(avgPrice)	
	(1)	(2)
Constant	-1.526*** (0.110)	-1.531*** (0.128)
log(goalFactor)	0.396*** (0.029)	0.390*** (0.091)
year2006	0.024 (0.147)	0.067 (0.176)
year2007	0.335** (0.147)	0.425** (0.177)
year2008	0.844*** (0.146)	0.829*** (0.176)
year2009	0.889*** (0.147)	0.874*** (0.177)
year2010	0.808*** (0.146)	0.778*** (0.177)
year2011	0.694*** (0.146)	0.645*** (0.177)
log(goalFactor):year2006		0.045 (0.113)
log(goalFactor):year2007		0.098 (0.119)
log(goalFactor):year2008		-0.016 (0.120)
log(goalFactor):year2009		-0.015 (0.120)
log(goalFactor):year2010		-0.031 (0.119)
log(goalFactor):year2011		-0.053 (0.119)
Observations	783	783
R ²	0.258	0.262
Adjusted R ²	0.252	0.249
Residual Std. Error	1.059 (df = 775)	1.061 (df = 769)

Note: *p<0.1; **p<0.05; ***p<0.01

Table 5: (Log) annual propensity to consume by a food bank (2006-2011) (Joint bids excluded)

	<i>Dependent variable:</i>	
	log(consumption_gf)	
	(1)	(2)
Constant	12.758*** (0.064)	12.887*** (0.073)
log(goalFactor)	-0.014 (0.017)	0.159*** (0.053)
year2006	0.082 (0.086)	-0.074 (0.101)
year2007	0.425*** (0.085)	0.303*** (0.102)
year2008	0.802*** (0.085)	0.725*** (0.101)
year2009	0.730*** (0.085)	0.572*** (0.101)
year2010	0.651*** (0.085)	0.468*** (0.101)
year2011	0.367*** (0.085)	0.181* (0.101)
log(goalFactor):year2006		-0.203*** (0.068)
log(goalFactor):year2007		-0.166** (0.068)
log(goalFactor):year2008		-0.115* (0.069)
log(goalFactor):year2009		-0.205*** (0.069)
log(goalFactor):year2010		-0.232*** (0.068)
log(goalFactor):year2011		-0.236*** (0.069)
Observations	686	686
R ²	0.234	0.238
Adjusted R ²	0.228	0.226
Residual Std. Error	0.708 (df = 679)	0.708 (df = 674)

Note: *p<0.1; **p<0.05; ***p<0.01

Table 6: Predicting Log (Average Two Month Prices) by food bank fixed effects and goal factor (2005-2011)

	<i>Dependent variable:</i>			
	log(avgPrice)			
	(1)	(2)	(3)	(4)
Constant	-2.098*** (0.109)	-1.832*** (0.101)	-1.424*** (0.177)	-1.624*** (0.180)
year2006	-0.040 (0.107)	0.042 (0.099)	0.079 (0.077)	0.094 (0.076)
year2007	0.328*** (0.106)	0.433*** (0.098)	0.488*** (0.076)	0.484*** (0.076)
year2008	0.825*** (0.106)	0.913*** (0.097)	0.999*** (0.076)	0.976*** (0.076)
year2009	0.816*** (0.106)	0.918*** (0.098)	1.007*** (0.076)	0.990*** (0.076)
year2010	0.749*** (0.107)	0.843*** (0.098)	0.954*** (0.076)	0.944*** (0.076)
year2011	0.458*** (0.107)	0.546*** (0.099)	0.707*** (0.077)	0.692*** (0.077)
intervalMar/Apr	-0.047 (0.080)	-0.048 (0.073)	-0.030 (0.057)	-0.028 (0.057)
intervalMay/Jun	0.133* (0.079)	0.153** (0.073)	0.165*** (0.056)	0.168*** (0.056)
intervalJul/Aug	0.165** (0.078)	0.174** (0.072)	0.202*** (0.056)	0.206*** (0.055)
intervalSep/Oct	0.241*** (0.077)	0.256*** (0.071)	0.270*** (0.055)	0.274*** (0.055)
intervalNov/Dec	0.209*** (0.078)	0.210*** (0.072)	0.232*** (0.055)	0.234*** (0.055)
log(goalFactor)		0.418*** (0.016)		0.342*** (0.059)
Food Bank FEs	no	no	yes	yes
Observations	4,043	4,043	4,043	4,043
R ²	0.053	0.198	0.533	0.537
Adjusted R ²	0.051	0.195	0.518	0.522
Residual Std. Error	1.404 (df = 4031)	1.293 (df = 4030)	1.000 (df = 3917)	0.996 (df = 3916)

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 7: Frequency of credit, joint bidding, negative prices, and non-positive prices.

Year	Credit	Joint-Bid	Negative Price	Non-Positive Price
2005	2.67%	1.39%	8.97%	53.38%
2006	6.11%	1.63%	13.59%	35.81%
2007	9.33%	2.02%	9.24%	18.16%
2008	11.61%	1.68%	8.75%	15.50%
2009	11.17%	1.66%	9.71%	18.23%
2010	12.05%	1.32%	6.08%	15.79%
2011	10.74%	1.23%	3.76%	15.65%

Table 8: The use of credit

	<i>Dependent variable:</i>
	log(pricePerPound)
Constant	−3.612*** (0.188)
log(goalFactor)	0.346*** (0.017)
Credit Dummy	1.646*** (0.045)
year2007	−0.008 (0.057)
year2008	0.629*** (0.059)
year2009	0.063 (0.062)
year2010	−0.413*** (0.065)
year2011	−1.012*** (0.066)
Month FEs	yes
Observations	29,508
R ²	0.091
Adjusted R ²	0.090
Residual Std. Error	2.746 (df = 29488)
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01

Table 9: The impact of expected supply on winning bids (2005-2011)

	<i>Dependent variable:</i>
	pricePerPound
Constant	0.722*** (0.027)
log(goalFactor)	0.042*** (0.002)
E[High Quality Composite Good Supply]	7.1123e-07** (2.8837e-07)
E[Low Quality Composite Good Supply]	1.4577e-07* (8.1134e-08)
E[High Quality Composite Supply] x Dummy for Low Quality Good	-3.3931e-07 (3.128e-07)
E[Low Quality Composite Supply] x Dummy for Low Quality Good	-2.3338e-07*** (8.7095e-08)
Year dummies	yes
Food type dummies	yes
Month dummies	yes
Observations	53,226
R ²	0.310
Adjusted R ²	0.310
Residual Std. Error	0.507 (df = 53182)

Note: *p<0.1; **p<0.05; ***p<0.01

Appendix

Constrained efficiency: We use the revelation principal to induce the constrained efficient equilibrium. Let a food bank of type μ make a report $\hat{\mu}$ to the principal, where the equilibrium utility of the agent is $U(\mu; \hat{\mu})$. Then the objective of the principal in a given interval is to

$$\max_{q_k(\mu)} E_{\mu} U(\mu; \mu) \quad (15)$$

subject to $U(\mu; \mu) \geq U(\mu, \hat{\mu})$ for all $\hat{\mu}$, $\sum_k q_{kj} \leq q_j$, and $\sum_k q_{kjt} \leq q_{jt}$.

It suffices to show the first best here, and to describe how it can overcome the private information constraint. The first best is characterized by

$$U_k^j(\mathbf{x}) = U_l^j(\mathbf{x}), \quad (16)$$

for all food banks k and l , and all goods j , where \mathbf{x} is the vector of consumption. For goods $j < m - l$, it requires that $x_{kj} = D_j + \frac{Q_j}{n}$, for all food banks k . For this to occur, it must be the case that for all goods $j < m - l$, $q_j^*(\mu) = (1 - \mu)D_j + \frac{Q_j}{n}$. The remaining issue is how to satisfy the food bank's incentive compatibility constraint. Note that in the first best, for these goods $\frac{dq_j}{d\mu} = -D_j$. Hence for these goods, the return to marginally increasing a report of μ is $\sum_{j < m-l} \frac{dq_j^*}{d\mu} U^j(x_j^*) = \sum_{j < m-l} D_j U^j(x_j^*)$. To satisfy truth telling, a food bank with higher μ must be compensated with an additional $\frac{d(\sum_{j \geq m-l} \kappa_j q_j^*(\mu))}{d\mu}$ for all goods $j \geq m - l$. All that remains is to show that there are no non-local deviations that are preferred to local deviations. This is ensured by the strict concavity of preferences for all goods $j < m - l$. This completes the proof.

Competitive Equilibrium with Equal Incomes and Symmetric Redistribution: We begin by assuming that all food banks play a static repeated interval game, where their expenditure of shares exactly equals their initial budgets $s_0 = 1$. (The interval notation is not relevant and is not included here, as all intervals are identical in the proposed equilibrium.) We then verify that food banks wish to play this repeated static outcome. In this interval game, food bank j bids b_{ijt} for a unit of food i in period t , with b_{ijt}^* its equilibrium choice. Let \mathbf{b}_{-ijt}^* be the equilibrium vector of bids by other food banks for food j in period t , and p_{ijt}^* be the winning price in period t .

Consider a suggested equilibrium outcome where $s_0 = 1$, and $s_t = \frac{R_t}{N}$, and for all i , $\sum_j p_{ijt}^* q_{ijt} = 1$. In words, this implies that in each interval, they spend their initial endowment of 1, but save all the redistributed income, which by the end of the period becomes their endowment for the beginning of the next interval. If $\sum_j \sum_i p_{ijt}^* q_{ijt} = N$, this implies that each food bank begins the next interval with an endowment of 1 share, and the static outcome can be repeated.

Let $x_j^* = D_j + \frac{q_j}{n}$. Then let p_j^* to be the solution to this following set of equations:

1. For all i, j and k , $\frac{p_k^*}{p_i^*} = \frac{U^k(\mathbf{x}^*)}{U^l(\mathbf{x}^*)}$, and

2. $\sum_j p_j^* q_{ij}^*(\mu_i) = 1$,

where μ_i is the type of food bank i , and $q_{ij}^*(\mu_i) = (1 - \mu_i)D_j + \frac{q_j}{n}$. A unique set of p_j^* always exist as weakly diminishing returns implies that the ratio $\frac{p_k^*}{p_i^*}$ is unique, and the level of each p_k^* in those ratios can be scaled such that $\sum_j p_j^* q_{ij}^*(\mu_i) = 1$.

Given this, consider the following bidding strategy:

1. For all i and j , food bank i bids p_j^* if $q_{ij} < q_{ij}^*(\mu_i)$, and

2. For all i and j , food bank i bids $p_j^* - \epsilon$ if $q_{ij} \geq q_{ij}^*(\mu_i)$.

This is a Bayes Nash equilibrium for all firms: conditional on \mathbf{b}_{-ijt}^* , it is in the interest of each firm to do bid as above. No firm has an incentive to undercut as firms that have already reached their equilibrium level bid slightly below that and so underbidding implies not receiving the good. Hence this is an equilibrium, under the assumption that no firm wishes to spend more than 1 in aggregate over the interval. To see that they spend no more than this, consider the marginal utility of a share. Waiting until the next interval offers marginal utility of income of $\frac{U^k(\mathbf{x}^*)}{p_k^*} = \frac{U^l(\mathbf{x}^*)}{p_l^*}$, which is equal to than the marginal value of spending an incremental share in the current interval for the marginal share, and higher for inframarginal shares. Hence this is an equilibrium of the bidding game.

The Propensity to Consume: Consider a case where all food banks incur a transportation cost denominated in “real” money, given by $C(q_A, q_B)$, where $C_i > 0$, $C_{ii} \geq 0$, and $C_{ij} \geq 0$, and where the food banks are of two types: food poor, with endowment D_j and food rich, with endowment δD_j , where $\delta > 1$. Let \bar{q}_j be the consumption of the food rich, and \underline{q}_j be the consumption of the food poor. Let the marginal utility of such “real” money be λ . Consider equilibrium quantities by the food rich if the goods were sold for free (in shares). This is given by \bar{q}_i^{**} where

$$U_i(\bar{q}_i^{**}) = \lambda C_i(\mathbf{x}^*), \quad (17)$$

Then consider an equilibrium in which the food rich, at prices p_j^* can afford to purchase \bar{q}_i^{**} with slack, so that $\sum p_i \bar{q}_i^{**} < 1$. If this is the case, then the marginal value of shares is 0 to the food rich and they will choose not to consume $1 - \sum p_i \bar{q}_i^{**}$. At these quantities, there are residual supplies of $q_j - \bar{q}_j^{**}$ of good j . Then, for the food poor with endowments D_j . consider an equilibrium price ratio $\frac{p_k^{**}}{p_l^{**}}$ of

$$\frac{p_k^{**}}{p_l^{**}} = \frac{U_k(x(\underline{\mathbf{q}}_j^{**})) - \lambda C_k(\underline{\mathbf{q}}_j^{**})}{U_l(x(\underline{\mathbf{q}}_j^{**})) - \lambda C_l(\underline{\mathbf{q}}_j^{**})}, \quad (18)$$

where in addition, the levels of prices are such that $\sum_j p_j^* q_j^{**} = 1$. Then the equilibrium is where the food poor consume as before, equate marginal utilities and spend all their income, while the food rich do not value the marginal share and simply hold it as balances in steady state.

The Quadratic Case The utility of type μ in an interval is given by

$$U(q_A, q_B; \mu) = \tau(\mu D + q_A(\mu)) - \gamma \frac{(\mu D + q_A(\mu))^2}{2} + q_B(\mu) + \mu \delta D. \quad (19)$$

The equilibrium consists of $q_A(\mu)$ and $q_B(\mu)$ subject to these being being feasible, and each type weakly preferring its allocation over the other's. The objective of each party is to maximize (19). The first order condition for the food bank of type μ is given by $\frac{p_A}{p_B} = \tau - \gamma(D + q_A(\mu))$. If the budget is balanced with total expenditures of 1, then

$$q_B(\mu) = \frac{\mu Q}{n} + \frac{(\mu - 1)d \frac{p_A}{p_B}}{2}. \quad (20)$$

To see how these can be implemented, consider the case where there are two types, where \bar{q}_i is the consumption of the food rich, and \underline{q}_A is the consumption of the food poor. These quantities are only feasible if they satisfy budget balance, which requires $p_A \bar{q}_A + p_B \bar{q}_B = 1$, and $p_A \underline{q}_A + p_B \underline{q}_B = 1$, subject to $\bar{q}_A + \underline{q}_A = Q$, and $\bar{q}_B + \underline{q}_B = \mu Q$, and the equilibrium quantities above. Substituting for relative prices yields absolute prices as

$$p_A^* = \frac{1 - \frac{\mu Q - \kappa d(\delta - 1)}{\mu Q + \kappa d(\delta - 1)}}{\Omega} \quad (21)$$

where κ is the relative price level as above, and $\Omega = \left(\frac{Q + (\delta - 1)d}{2} - \left(\frac{\mu Q - \kappa d(\delta - 1)}{\mu Q + \kappa d(\delta - 1)} \right) \left(\frac{Q - (\delta - 1)d}{2} \right) \right)$, where $p_B^* = \kappa p_A^*$.

The average welfare of the old system per client was $\underline{U} = (\tau + \delta)(d + \frac{Q'}{n}) - \frac{\gamma}{2} E \left(\mu d + \frac{Q'}{n} \right)^2$. The per food bank utility from the Choice System is given by $U^* = (\tau + \mu)(d + \frac{Q}{n}) - \frac{\gamma}{2} \left(d + \frac{Q}{n} \right)^2$. Supply changes from Q' to Q . If there were no other benefits, welfare would change by $\Delta U_Q = (\mu + \tau - \gamma d) \frac{Q - Q'}{n} - \frac{\gamma}{2} \left[\left(\frac{Q}{n} \right)^2 - \left(\frac{Q'}{n} \right)^2 \right]$. The second benefit is through the reallocative benefit of better matching of demand to marginal valuations, and is given by $\Delta U_R = \frac{\gamma d^2 \sigma_\mu^2}{2}$. Then the overall change in welfare is given by $\Delta U = \Delta U_Q + \Delta U_R$. Note that the *total* quantity of food obtained $q(D) = q_A(D) + q_B(D)$ is given by $q(\mu) = \frac{(1 + \delta)Q}{n} - \frac{(\mu - 1)d(1 - \frac{p_A}{p_B})}{2}$. Then as $Var(q) = (\sigma_\mu^2 d(1 - \frac{p_A}{p_B}))^2$, this implies that

$$\Delta U_R = \frac{\gamma Var(q)}{2(1 - \frac{p_A}{p_B})^2}. \quad (22)$$

The Composition of Food: Here we consider a case where the composition of food changes, holding richness constant. Holding “wealth” constant implies that when a food bank has more of A , they have less of B , where the amount “less” is determined by relative prices. The endowment of goods A and B for food bank i are now D_{iA} and D_{iB} but where $D_{iA} = D_i + \lambda_i$, and $D_{iB} = \mu D_i - \frac{p_A}{p_B} \lambda_i$, where $D_i \in \{\underline{d}, \bar{d}\}$. Varying food in proportion to its price level is simply done to separate composition effect from income effects. Let the distribution of λ_i for food bank i as binary: $\lambda_i \in \{z, -z\}$, with equal probability and independently distributed across food banks. Assume that the number of food banks is large enough that this does not affect the aggregate quantity of food. First note that relative prices remain at κ , as the noise washes out in equilibrium with large n . Given those prices, the quantity of good A purchased by food bank i of type D is now simply given by $\bar{q}_{Ai} = \frac{Q}{n} - \lambda_i - \frac{(\delta-1)d}{2}$ and $\underline{q}_{Ai} = \frac{Q}{n} - \lambda_i + \frac{(\delta-1)d}{2}$. While these transitory shocks do not affect *total* quantities x_i^* consumed under the Choice System, they do affect the welfare from the old system. Welfare under the old assignment rule is now given by

$$\hat{U} = (\tau + \mu) \left(\frac{\bar{d} + \underline{d}}{2} \right) \frac{Q'}{n} - \frac{\gamma}{8} \left((\bar{d} + z + \frac{Q'}{n})^2 + (\underline{d} + z + \frac{Q'}{n})^2 + (\bar{d} - z + \frac{Q'}{n})^2 + (\underline{d} - z + \frac{Q'}{n})^2 \right) \quad (23)$$

The value of the choice system over the market then becomes $\Delta \tilde{U} = U^* - \hat{U} = \Delta U + \frac{\gamma}{8} z^2$, and welfare gains are parameterized by a measure of dispersion.

Computing a structural estimate of the returns to better allocations: In equilibrium, if C is the number of clients, $\frac{p_A}{p_B} = \tau - \gamma \left(\frac{Q_A + \frac{d(\delta+1)}{2}}{C} \right)$, so that $\frac{d(\frac{p_A}{p_B})}{dQ_A} = -\frac{\gamma}{C}$, where Q_A is total supply of good A . (In the model above, we use the notation n but that is because of the assumption that the model is in terms of per capita consumption.) However, when we estimate $\hat{\gamma}$ in (10), we do not estimate the sensitivity per capita. Instead, we estimate $\hat{\gamma} = \frac{\gamma}{C}$. Then consider the change in utility

$$\Delta U_R = \frac{\gamma(1 - \frac{p_A}{p_B})^2 \text{Var}(q)}{2} = \frac{C \hat{\gamma} (1 - \frac{p_A}{p_B})^2 \text{Var}(q)}{2}. \quad (24)$$

But as per capita allocation are $q_i = \frac{Q_i}{C_i}$ where C_i is the number of clients for food bank i , this simplifies to $\Delta U_R = \frac{C \hat{\gamma} (1 - \frac{p_A}{p_B})^2 \text{Var}(\frac{Q_i}{C_i})}{2}$. As $C_i = \frac{G_i}{\sum_{j=1..n} G_j} C$, this becomes

$$\Delta U_R = \frac{(\sum_{j=1..n} G_j)^2}{2C} \hat{\gamma} (1 - \frac{p_A}{p_B})^2 E \left(\frac{Q_i}{G_i} - \frac{Q}{\sum_{j=1..n} G_j} \right)^2. \quad (25)$$

This measures welfare per client. However, we do not know the number of clients, and so this cannot be estimated. However, we can normalize it by the first best

welfare level per client, which is given by $U^* = \tau x_A - \frac{\gamma}{2} x_A^2 + x_B$, where x_i is per capita consumption of good i . Let the fraction of all consumption that derives from Feeding America be given by f . Then

$$U^* = \tau \frac{q_A}{f} - \frac{\gamma}{2} \frac{q_A^2}{f} + \frac{q_B}{f}. \quad (26)$$

However, we only observe total quantities $Q_i = Cq_i$ and estimate $\hat{\gamma}$, this implies that we can estimate

$$U^* = \tau \frac{Q_A}{Cf} - \frac{C\hat{\gamma}}{2} \frac{Q_A^2}{Cf} + \frac{Q_B}{Cf} = \frac{\tau \frac{Q_A}{f} - \frac{\hat{\gamma}}{2} \left(\frac{Q_A}{f}\right)^2 + \frac{Q_B}{f}}{C}. \quad (27)$$

Finally, note that $\frac{p_A}{p_B} = \tau - \gamma \left(\frac{Q_A}{fC}\right)$, so this simplifies to

$$U^* = \frac{\frac{p_A}{p_B} \frac{Q_A}{f} + \frac{\hat{\gamma}}{2} \left(\frac{Q_A}{f}\right)^2 + \frac{Q_B}{f}}{C}, \quad (28)$$

and so

$$\frac{\Delta U_R}{U^*} = \frac{\frac{(\sum_{j=1..n} G_j)^2}{2} \hat{\gamma} \left(1 - \frac{p_A}{p_B}\right)^2 E\left(\frac{Q_i}{G_i} - \frac{Q}{\sum_{j=1..n} G_j}\right)^2}{\frac{p_A}{p_B} \frac{Q_A}{f} + \frac{\hat{\gamma}}{2} \left(\frac{Q_A}{f}\right)^2 + \frac{Q_B}{f}}, \quad (29)$$

as required.

Welfare gains for all food banks: This is then computed using the following basic statistics, where $n = 113$, and the average goal factor is 0.7175816:

- The average squared sum of goal factor: 6575.796
- The average total number of low quality composite pounds in a two month period: 26,416,743 (lbs)
- The average total number of high quality composite pounds in a two month period: 7,829,824 (lbs)
- The goal-factor weighted variance of composite total pounds received by a food bank over each two months divided by its goal factor: 6.874e+11
- The average price of low quality goods from 2005 to 2011: 0.105788502 (shares/lbs)
- The average price of high quality goods from 2005 to 2011: 0.7055 (shares/lbs)

All that remains to carry out the computation is to choose f , for which we choose 0.1, and an estimate of $\hat{\gamma}$. From Table 9, the change in the price of the high quality goods (low quality goods) by a unit increase in low quality supply is $1.4577e - 07$ ($-2.3338e - 07$). Then $\hat{\gamma}$ is computed as $\frac{0.105788502 - (2.3338e - 07)}{0.7055422989 + (1.4577e - 07)} - \frac{0.105788502}{0.7055422989} = 0.46e - 07$. This yields the welfare gains of 8.89%.