Choice by Design:

Evidence from Feeding America's Food Allocation Problem

Sam M. Altmann*

Click Here For The Most Recent Version

Abstract

Feeding America, an organisation responsible for feeding 130,000 Americans every day, distributes food among a nationwide network of food banks. Their allocation mechanism, known as the 'Choice System', uses auctions and a virtual currency to give food banks choice over the food they receive. This paper examines the consequences of enabling this choice. I apply a dynamic auction model to food bank bidding data, estimating the distribution of food banks' heterogeneous and time-varying needs. The central challenge is that I do not observe food banks' inventories — a key determinant of bidding behaviour. I overcome this difficulty using variation in food banks' winnings (observed shifters of these unobserved stocks) to identify the model, which I estimate using a Gibbs Sampler. I then compare welfare under the Choice System to Feeding America's previous allocation mechanism which gave food banks very limited choice. I estimate that the Choice System increased welfare by the equivalent of a 32.7% increase in the quantity of food allocated. Most of this gain arises because food is allocated in batches, rather than sequentially.

---

*School of Economics and Finance, Queen Mary University of London, s.altmann@qmul.ac.uk.
The author benefited from many conversations about Feeding America and food banking with Canice Prendergast, Mike Loeffl, Sarah Pennel, Dhanu Sherpa, and Ann Sheppard. The author also wishes to thank Nikhil Agarwal, Matthew Backus, Ian Crawford, Francis DiTraglia, Shihang Hou, Paul Klemperer, Ellen Lees, Robert Miller, Luke Milsom, Anh Nguyen, Howard Smith, Alex Teytelboym, Neil Thakral, and Daniel Waldinger for useful comments and conversations, as well as seminar participants at Berkeley, Carnegie Mellon, CREST, Manchester, Maryland, Minnesota, Monash, Oslo, Oxford, Paris School of Economics, Sciences Po, Surrey, Tokyo UTMD, Toulouse School of Economics, and Queen Mary University of London. Finally, thank you to Feeding America for giving me access to their data. This paper previously circulated under the title "Choice, Welfare, and Market Design: An Empirical Investigation of Feeding America's Choice System."

1 Introduction

Organisations are regularly faced with the problem of allocating scarce resources as efficiently and as equitably as possible. Governments must decide how to allocate contracts to contractors, local authorities must allocate school places to students, and health officials must allocate kidneys to transplant patients. Feeding America, a not-for-profit responsible for feeding 130,000 people every day, must decide how to allocate truckloads of donated food among its network of regional food banks.

The efficient and equitable allocation of food is a priority for Feeding America, to ensure that food banks can meet the ever-increasing demand for their services. Like many food relief organisations around the world, Feeding America previously employed a mechanism that allowed food banks very little choice in the food they received. Under this mechanism, referred to as the 'Old System', food banks would queue until they were offered an essentially random truckload of food. This mechanism was unpopular as food banks were rarely offered the types of food they needed. Efficient central planning is difficult because of unobserved heterogeneity in food banks' needs: Different food banks need different types of food at different times.[1] This heterogeneity arises because food banks in different parts of the country have access to different types of food from their local donors, and these types of food are liable to change over time. Instead, Feeding America's current allocation mechanism, the 'Choice System', consists of an auction market in which food banks receive an amount of virtual currency to bid on loads of donated food (Prendergast, 2017). This gives food banks a large degree of choice, giving them control over their allocations.

In this paper I use a rich model of food bank bidding behaviour to investigate welfare under the Choice System, compared to alternative mechanisms that allow food banks varying degrees of choice. I develop a novel empirical strategy to estimate food banks' demand functions despite not observing their inventories, applying a dynamic auction model with storable goods to detailed Choice System data. I exploit the panel dimension of the data to allow demand to vary across food banks and over time, as different food banks have different storage capacities, cater to different

[1] I use the term 'needs' to capture both what a food bank has a preference for, on behalf of their clients, as well as what they have room for in their warehouse. In this way, the term is intended to capture the determinants of a food bank's demand function — a food bank with a warehouse full of cornflakes may still have positive marginal utility of additional cornflakes, but due to capacity constraints will not demand additional cornflakes.

numbers of clients, and receive different types of food at different times from their local donors. I then use these estimates to evaluate equilibrium allocations under a number of alternative allocation mechanisms. Counterfactual simulations demonstrate that, relative to the Old System, the Choice System is extremely effective at achieving Feeding America's goals: Welfare is 32.7% higher under the Choice System than the Old System. Furthermore, despite fears that a market mechanism could lead to a more unequal distribution of food, on average 87.9% of food banks are estimated to be better off under the Choice System.

In order to investigate food banks' needs, and so evaluate welfare under various allocation mechanisms, I first develop a structural model of food banks bidding for food on the Choice System. The structural model combines the storable goods framework of Hendel and Nevo (2006) with the dynamic multi-object auction model of Altmann (2022). Descriptive evidence demonstrates the need for this framework: Conditional on winning a load, food banks are less likely to bid on similar loads on subsequent days even when the price is essentially zero. They then return to bidding some time later, having given out this food over the course of several days. This suggests food banks treat loads as a storable good subject to storage costs. This dynamic linkage emphasises the need for a model that accounts for the dynamic environment. Meanwhile, when multiple similar loads are auctioned simultaneously food banks are less likely to bid on any given load. This suggests that similar loads are substitutable, and requires a multi-object model to account for the simultaneous auction environment.

The importance of choice depends on the degree of unobserved heterogeneity in food banks' preferences and storage costs, as well as the degree of substitutability of different types of food. The model incorporates this in three key ways. First, food is classified by how it is stored (capturing storage costs), and how it is used. Second, the long panel (around 900 days) allows me to estimate distinct parameters for each food bank, allowing for permanent heterogeneity across food banks. Finally, I allow for time-varying unobserved heterogeneity, which I attribute to the fact that I do not observe food banks' stocks of various types of food. This captures how a food bank's clients irregular take food from their local food bank, and food banks irregularly receive food from their local donors.

The central challenge, for both identification and estimation, is that I do not observe food banks' stocks. Current stocks are a key determinant of demand — if a food bank stops bidding it might be because, unobserved by the econometrician, they recently received food from a local donor. Building on Hu and Shum (2012) I prove that the model is non-parametrically identified. Key to the argument is variation in food banks' choice sets and winnings, and how this subsequently drives variation in bidding behaviour. Winnings in particular are key, as these are essentially observed changes in the unobserved stocks.[2] A methodological contribution of this paper is to develop a procedure to estimate bidders' values in a dynamic auction environment when individual state variables (stocks) are unobserved. I overcome this problem using a Gibbs Sampling procedure, employing a data-augmentation step to draw the unobserved stocks from their conditional posterior distribution. To the best of my knowledge this is the first paper to estimate a model of this type.

I employ the three step estimation procedure introduced in Altmann (2022). In the first step I estimate equilibrium beliefs by estimating the conditional distribution of winning bids. I then invert food banks' first order conditions for optimal bidding, obtaining an inverse bidding system as in Guerre et al. (2000) and Gentry et al. (2023). In the second step, using the inverse bidding system, I estimate the distribution of food banks' 'Pseudo-Static' payoffs from winning combinations of lots. This means I estimate the sum of bidders' flow payoff and their discounted continuation value — essentially estimating the model as though food banks were myopic. During this step I also estimate the transition process for food banks' stocks. Finally, in the spirit of Jofre-Bonet and Pesendorfer (2003), the continuation value can be written as a function of observed bids, beliefs, and this pseudo-payoff function. Therefore, in the third step I evaluate the estimated continuation value, before backing out the distribution of flow payoffs from the definition of the pseudo-payoffs.

I find significant evidence of demand heterogeneity both across food banks and over time. I estimate large differences in access to local donors and variability of local donations, varying by a factor of 30 across food banks. Meanwhile, I estimate that variation in stocks account for 93% of the unexplained variation in bidding behaviour. Food banks go through extended periods with high stocks, during which they rarely place bids, and periods with low stocks, during which they bid frequently.

[2]How bidding behaviour varies with the number and composition of available lots identifies food banks' storage capacities: A food bank facing high storage costs will only bid on a small subset of available lots, to avoid winning more than it can afford to store. Meanwhile, after winning a lot, the length of time before food banks return to their average bidding propensity enables identification of the unobserved state transition process: If it takes them a long time to return to bidding on a particular type of food, this suggests they generally have access to that food from their local donors.

Using the estimated model I then simulate equilibrium allocations under the Old System. This enables me to consider the welfare and distributional consequences of Feeding America's transition to the Choice System, building on evidence presented in Prendergast (2017) and Prendergast (2022). Welfare is 32.7% higher under the Choice System than under the Old System. This is roughly equivalent to an additional 86.5 tons of food allocated each day, enough to support an additional 35,100 people. This arises because, under the Old System, food banks are roughly three times more willing to accept any load they are offered than under the Choice System, as they do not know when they will next be offered food. They accept food that does not directly meet their needs, food that may be used more effectively by another food bank at that point in time. Meanwhile, on average 87.9% of food banks are better off under the Choice System, and no individual food bank is statistically significantly worse off.

I then use additional simulations, varying the different aspects of these mechanisms, to tease out the most important features of the Choice System. Features we can then take to other food relief organisations around the world. In particular, by comparing the Old System to running sequential auctions, we learn about the importance of allowing food banks to signal the intensity of their preference for each load. This establishes the relative importance of ensuring food goes to the food bank who values it most, rather than essentially being offered out at random. Then, by comparing sequential auctions to the Choice System, which uses simultaneous auctions, we learn about the importance of simultaneous versus sequential allocation. Under sequential allocation every donation is allocated before the next arrives. Whereas allocating food in batches ensures food banks have information about all the food being allocated on a given day when making decisions, giving them more control over their allocations. Decomposing the welfare gain from transitioning to the Choice System, 36% of this improvement comes from the signalling effect, while 48% comes from the batching effect. This is an important result in practice as the majority of other food relief organisations around the world allocate food sequentially.

Finally, I explicitly consider the efficacy of several mechanisms used by other food relief organisations. For example, a mechanism that offers food only to the nearest food bank, aiming to minimise transportation costs, achieves only 20% of the welfare under the Choice System. At least in the U.S. setting, transportation costs are not so important. Many food networks, including the Trussell Trust in the U.K., implicitly use this mechanism by linking food banks up with additional nearby donors instead of allocating food centrally.

1.1 Related Literature

This paper contributes to the literatures on empirical market design, empirical auction econometrics, and storable goods.

Empirical market design is a growing literature analysing preferences and allocations in centralised assignment markets. The papers most similar to this one are Prendergast (2017) and Prendergast (2022) who also study Feeding America's Choice System. My structural approach is complementary to their predominantly descriptive approaches, enabling detailed welfare analysis and investigation of which features of the Choice System are the most important. I employ richer data that is disaggregated at the auction level and includes information on losing bids. By studying the exact timing of bids I gain a detailed understanding of how food banks make inter-temporal substitutions. This allows me to simulate alternative dynamic allocation mechanisms and shed light on the key features of the Choice System. Papers studying other allocation settings include Agarwal et al. (2020) and Agarwal et al. (2021) on deceased donor kidney waitlists, Waldinger (2021) on public housing, and Verdier and Reeling (2022) on hunting licenses. Similar to this paper, they assess the value of giving agents choice over allocations, considering trade-offs between efficiency and other concerns of policy makers. While heterogeneity in match values is an important theme of all these papers, this is the first to consider the role of heterogeneity over time.

I also build on the dynamic multi-object auction model of Altmann (2022), which combines the models of Jofre-Bonet and Pesendorfer (2003) and Gentry et al. (2023). Unlike these papers reservation prices and endogenous entry are important in my application, with the average bidder only bidding on around 2% of auctions. The focus on a large auction market is similar to Backus and Lewis (2024) who introduce a framework for analysing dynamic bidding in large single-unit second-price auctions, which is employed in Bodoh-Creed et al. (2021) among others. In empirical auction studies it is standard to impute bidder's state variables, such as backlogs of contracts or stocks. This is the first auction study to allow state variables to be unobserved.

Finally, I contribute to the econometric literature on storable goods and identification of unobserved states. One distinction between my model and those of Hendel and Nevo (2006) and Erdem et al. (2003) (among others) is the role local donations, a key driver of heterogeneous behaviour across food banks and across time. To my knowledge, this is the first paper to formally prove identification of a storable goods model. Non-parametric identification results are important for ensuring that identification is not purely driven by restrictive parametric assumptions, and that these can be considered (potentially restrictive) finite-sample approximations. My proof builds on the methodology of Hu and Shum (2012) and Hu and Schennach (2008), leveraging observed shifters of the unobserved state for identification.

1.2 Overview

Section 2 describes the institutions and data being studied. Section 3 provides descriptive evidence of heterogeneous needs and presents several key stylised facts behind bidding behaviour. Section 4 outlines the model of food bank behaviour and discusses identification. Section 5 describes the parametrisation and estimation procedure, while 6 details the estimation results. Section 7 describes the counterfactual mechanisms and presents the results from several simulations.

2 Institutional Background and Data

This section details Feeding America and their allocation mechanisms. Details come from Prendergast (2017). In Section 2.2 I describe the data.

2.1 Feeding America

Feeding America began in 1976 as a collection of food banks that would solicit donations from local grocery stores and farms. As additional food banks joined their network it became necessary to co-ordinate resource sharing. In 2005, at the recommendation of a task force consisting of economists and food bank managers, they replaced the Old System with the Choice System. Many of Feeding America's associated food banks operate as standard food pantries — directly giving out food to those in need. However, the majority of food banks act as food distributors; themselves responsible for storing and sending out food to *hundreds* of local food pantries.

2.1.1 The Old System

Under the Old System any truckload of food donated to Feeding America was offered to the head of a queue. The potential recipient had a few hours to accept or decline before it was offered to the next food bank. This meant that each load could only be offered to around ten food banks before being returned to the donor. To discourage rejections, food banks would return to the back of the queue regardless of whether they accepted. A food bank's relative position in the queue was determined jointly by whether they had recently been offered food, and their 'Goal Factor', a measure of poverty in their catchment area. A higher Goal Factor implies more mouths to feed, so these food banks should be offered more food. Transportation costs were paid by the food banks, many of whom have fleets of trucks and lorries for this purpose.

The type of food in each truckload was essentially random, so that on average food banks received the same quantities of food per mouth. This would have been optimal if food banks all had the same preferences and capacities. In reality, different food banks needed different types of food at different times. Food banks use food from Feeding America to substitute for food they do not receive from their local donors. Feeding America wanted to improve welfare by taking account of differing needs, so decided to use a market mechanism to give food banks control over their allocations.

2.1.2 The Choice System

The Choice System consists of simultaneous first-price sealed-bid auctions. Two rounds of auctions occur each day, five days a week, with around 30 lots auctioned each day. Bidders observe the previous winning bids for a particular type of food, making it easier for food banks to know how to bid. Outcomes of auctions that occur simultaneously are independent, and bidders cannot place combination bids. Winners generally pay to transport their winnings.

Food banks bid with a virtual currency called 'shares'. Other than storage and transportation costs, the only opportunity cost a food bank faces when bidding is that they will have fewer shares to bid on other lots. Feeding America can ensure that food banks with larger Goal Factors are allocated more shares and, consequently, receive more food. Shares are redistributed each night. Food banks can save shares from one day to the next. Those with less than the median allocation of shares have access to interest-free credit, so that food banks can smooth their consumption over time. The money supply is set to ensure that prices remain constant (on average) over time, reacting to changes in the supply of food.

Food banks can bid negative amounts, down to a reserve price of −2000 shares. This incentivises food banks to accept undesirable loads, helping Feeding America maintain good relations with their donors by ensuring that every lot is graciously accepted. On average 34% of lots are sold at weakly negative prices, and 10% are sold at the reservation price.[3] Negative prices occur because food banks face storage costs. Partly these are physical costs, but also the opportunity cost of volunteers' time and effort to make sure every load is packaged and stored properly to ensure that as little food spoils as possible.

The introduction of a market mechanism had the potential to disadvantage smaller food banks. Feeding America incorporated several features to alleviate this risk, including credit use, a fairness committee, and joint bidding.[4] For this reason, food banks report great satisfaction with the mechanism. The Choice System incorporates several additional features designed to tackle various intricacies of the food allocation problem, including allowing multiple homogeneous loads to be allocated using discriminatory auctions and allowing food banks to sell excess local donations.[5]

2.2 The Data

Two main sources of data were used for this paper. I use proprietary bidding data from the Choice System, which was provided by Feeding America. I also use publicly available data on food bank demographics and catchment areas.

[3]While 9% of lots are not sold right away, most are sold the following day. Lots not sold right away are predominantly multiple loads of bottled water and juice. The numbers are skewed by 130 loads of 8 litre bottles sold over several months.

[4]Several small food banks place the majority of their bids jointly with other small food banks as they are unable to use an entire truckload of food on their own. Meanwhile, those responsible for the majority of consumption do not bid jointly. For this reason I generally ignore this decision.

[5]The structural model correctly models the discriminatory auctions distinctly from the standard simultaneous auctions — the only difference is that one cannot lose a unit for a high bid yet win a lot for a lower bid. Meanwhile, food sold by food banks only makes up 4.5% of lots. The food banks who rely most on the Choice System (and those included in estimation) almost never sell their excess, despite their storage costs. This is due to additional distortions in this market, including taxes and additional transportation costs. For this reason I generally ignore the decision to sell food.

2.2.1 Choice System data

The Choice System dataset contains information on 26,617 individual auctions held from January 2014 to October 2017, covering 165 food banks. The data includes winning and losing bids, as well as information on the food composition and locations.[6]

The sheer volume of types of food being auctioned makes categorisation necessary. I split food into 15 categories, largely the same categories used in Prendergast (2017). To capture different types food being imperfectly substitutable and having different uses I further split food into 152 subcategories. To capture storage costs I categorise food into four storage types: Dried, Tinned/Bottled, Refrigerated, and Non-Food.[7] Many loads contain multiple types of food. Around 30% of the food being auctioned is fresh produce. However, eighteen months into my sample Feeding America stopped allocating produce centrally, and began sending it to one or two (urban) food banks that previously consumed almost all the produce. For this reason, I drop data on produce, and instead treat it as local donations. The previous version of this paper included produce and found extremely similar results.

Figure 1 presents descriptive statistics on the auctioned lots, split by storage method. Several things are evident: First, many lots are allocated simultaneously, and lots come in variable sizes. Second, only a small number of bidders bid on any given lot, and a large proportion sell for negative prices — particularly low quality beverages (included in Tinned). This suggests low demand for these types of food.[8]

---

[6] I do not observe whether a given auction happened in the morning or afternoon, so assume all auctions in a day happen at the same time. This is a potential weakness of this analysis. However anecdotal evidence suggests that most food banks bid in only one auction round per day. This was suggested by Canice Prendergast, one of the Choice System's designers. If food banks are optimally choosing not to bid on any auction in a given round then the inaccuracy of my results will be minor.

[7] Refrigerated includes anything that needs to be stored in a fridge or freezer, such as meat and dairy. Tinned and Bottled food includes anything with a long shelf-life that is tinned or bottled, ranging from baked beans to bottled water. Dry food captures long shelf-life food such as cereal, pasta, or cookies. Non-food items includes non-edible items, predominantly cleaning products and baby food. See Appendix A for additional discussion of how food was categorised.

[8] This could be explained by bidders colluding. In practice collusion is highly unlikely, given how this harms non-colluding food banks and that most food bank managers are extremely prosocial.

Figure 1: Descriptive Statistics, across lots

	Dried	Tinned	Fridge	Non-food	Mixed	Total
Daily lots						
(mean)	10.74	10.83	5.11	3.3	4.78	32.74
(std)	8.35	10.8	3.76	3.21	3.94	16.8
Pounds per lot						
(mean, 000s)	22.5	34.3	28	20.4	27	26.5
(std, 000s)	9.7	8.3	10.2	12.2	10.6	10.9
Winning bid						
(mean)	2106	1085	2688	2967	2515	2134
(std)	5329	6414	6176	6436	5207	5818
No. bidders						
(mean)	2.95	2.7	2.5	3	2.8	2.8
(std)	3.14	3.5	3.12	3.26	3.08	3.21
% Allocated	93	83	91	91	95	91
% Negative prices	35	47	29	28	27	34

Note: Mixed loads are presented separately for this figure only. Winning bids includes the reservation price when no bids are received. 'Allocated immediately' refers to the percentage of lots that receive at least one bid above the reserve. Negative prices include loads allocated for 0 shares.

2.2.2 Auxiliary Data

I construct food banks' Goal Factors using the formulae in Prendergast (2022), freely available location and catchment data for food banks, and food insecurity data.[9] Figure 2 summarises the relevant demographic information and bidding behaviour of food banks. The key take-way is that characteristics and behaviour differ drastically across food banks, suggestive of their heterogeneous needs.

3 Descriptive Evidence

In section 3.1 I present suggestive evidence of heterogeneity, highlighting the likely value of choice. Then, in section 3.2 I investigate the key determinants of bidding, putting together several stylised facts motivating my model's key features.

[9]These are available from https://www.feedingamerica.org/find-your-local-foodbank and from Feeding America's 'Hunger in America' on-line tool https://map.feedingamerica.org

Figure 2: Descriptive Statistics, across food banks

	Mean	p10	p25	p50	p75	p90
Population (000s)	1913	384	676	1270	2543	4385
Poverty (000s)	284	64	99	191	373	645
Goal Factor	1	0.16	0.36	0.62	1.19	2.46
Bids Placed	356	9	35	137	390	810
Average Bid	4140	742	1322	2760	4684	9539
Lots Won	137	3	18	55	147	328
Average Payment	4571	541	1332	3033	5881	9684

Note: Statistics are calculated by food bank, then quantiles are evaluated across food banks. The mean Goal Factor is normalised to 1. Population and Poverty figures refer to the number of people in a food bank's catchment area.

3.1 Evidence of Heterogeneity

Under the Old System every food bank was, ex ante, offered the same allocation. If food banks have heterogeneous needs, unknown to the social planner, welfare might be increased by allowing food banks greater choice in their allocations. Therefore heterogeneity is a key determinant of the value of choice. To investigate heterogeneity I demonstrate how average bids or winning bids vary over types of food, over food banks, and within food banks over time. I consider a series of simple reduced form fixed effects regressions, controlling for lot composition, distance, and the censoring caused by the reserve prices and entry behaviour using a Tobit specification.

3.1.1 Differences Across Lots

Figure 3 plots average winning bids across subcategories, demonstrating that different types of food attract significantly different bids. These differences cannot only be explained by differences in supply, requiring demand side factors to explain them also. For example, Cereal is abundant and sells for relatively high prices, while Health and Beauty products are rare but sell for lower prices. It is clear there is a great deal of heterogeneity between lots, and that these lots cannot be substituted one for one.

3.1.2 Differences Across Food Banks

Food banks differ vastly in terms of their total consumption: Five food banks receive the same amount of food as 122 food banks who receive the least food from Feeding

Figure 3: Heterogeneity in Lots

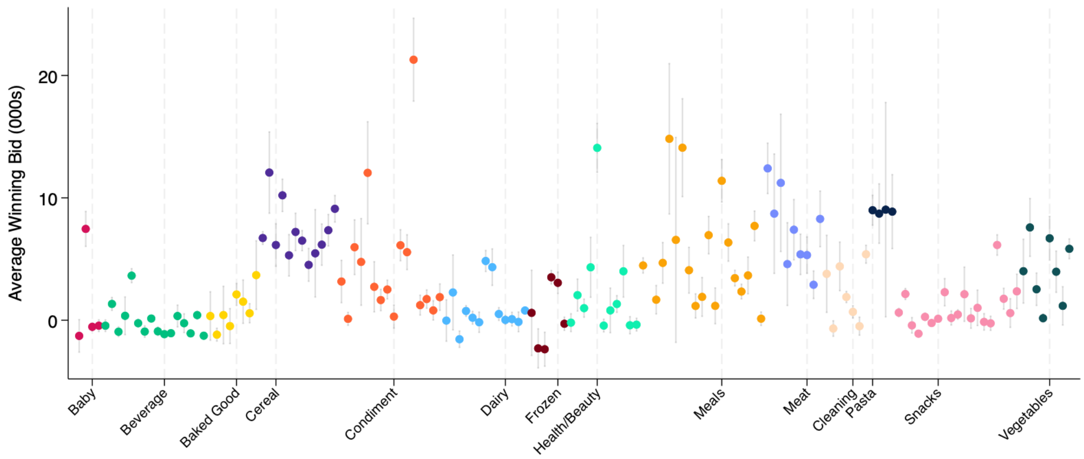

Note: Plots mean winning bids, and 95% robust confidence intervals, across subcategories, controlling for censoring and lot composition. Coefficients are ordered and coloured according to category.

America. However, these food banks are also choosing very different types of food. These 122 food banks, in total, spend 4 times as much as the five high consumption food banks. Therefore these five food banks are choosing to receive much cheaper food. This is likely because they rely on Feeding America for their staples, having fewer local donors than the other 122 food banks.

Figure 4 plots average bids across food banks and across different types of food according to how they are stored. Estimates are negative and large due to the degree of non-entry — the average food bank bids on only 2% of lots. There are three main takeaways: First, there is significant heterogeneity in average bids across food banks. Second, there is heterogeneity in average bids within food banks, across types of food. Third, that these two types of heterogeneity are not perfectly correlated: For some food banks average bids on Tinned food are higher than average bids on Dried food, but for other food banks this relationship is reversed.

3.1.3 Differences Over Time

To investigate the variation in bidding behaviour over time I run the same Tobit specification as above, considering how average bids vary from month to month. I focus on only those food banks who win at least 100 lots over my sample period. I then consider the degree of variation in my estimated parameters. A likelihood

Figure 4: Heterogeneity Across Food Banks

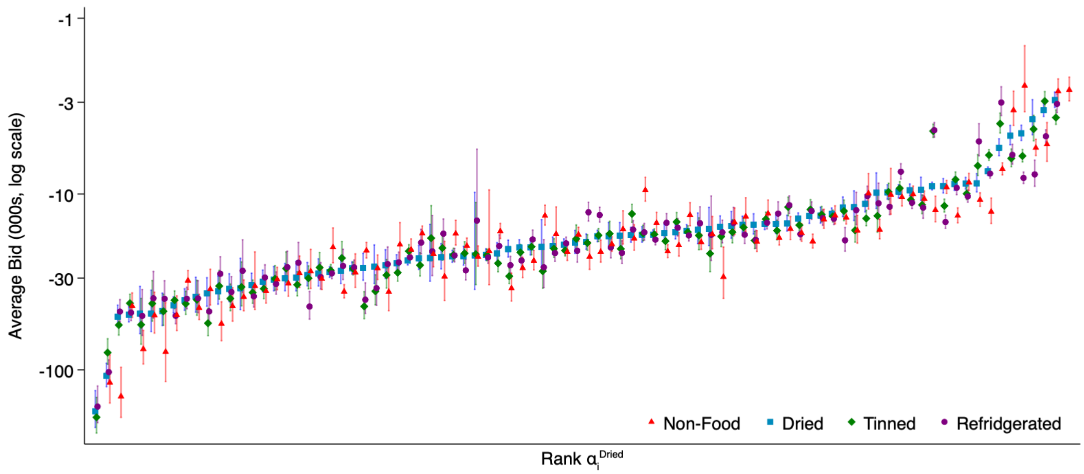

Note: The figure plots coefficients and 95% robust confidence intervals from regressing food bank × storage type on bids, controlling for distance and censoring/non-bidding. Food banks are ordered by average bids on Dried food (the blue points monotonically increase from left to right). I include results for those who won at least 50 lots.

ratio test that parameters are constant over time is rejected at 5% significance level for 96% of food banks. This is indicative of systematic heterogeneity in food banks' needs over time. Additional results are reported in Appendix B.

3.2 Stylised Facts

I now investigate several stylised facts which point towards key determinants of bidding behaviour, motivating the model's key features. I have already emphasised the role of several types of heterogeneity that will become important features of my model.

3.2.1 Negative Bidding

Negative bidding is common: 34% of bids are negative. Furthermore, with a negative reservation price non-entry only happens when food banks have negative marginal valuations, when food banks must be paid to accept certain loads. This occurs in 98% of bidder × lot combinations. Negative valuations likely occur because of limited storage capacity, as emphasised in Prendergast (2017).[10] They cannot throw away

[10]Storage costs include the opportunity cost of volunteer time and effort from repackaging and properly storing the food. Transportation costs are also a key factor in these negative valuations.

excess (non-expired) food as this sends a bad signal to donors.

3.2.2 Dynamic Complementarities

Figure 5 panel (A) demonstrates that, conditional on winning food of a particular type at time 0, the probability of bidding on lots of the same type falls by 35% (3 pp) on subsequent days.[11] As food banks win more of a particular type of food, the less they are willing to pay for an additional lot from that type. Because food banks are known to be forward looking, this evidence of 'dynamic linkages' highlights the need to model dynamics. Food banks treat this food as a storable good, working through their current stocks before returning to bidding on the Choice System.

Figure 5: Evidence of Dynamic and Static Complementarities

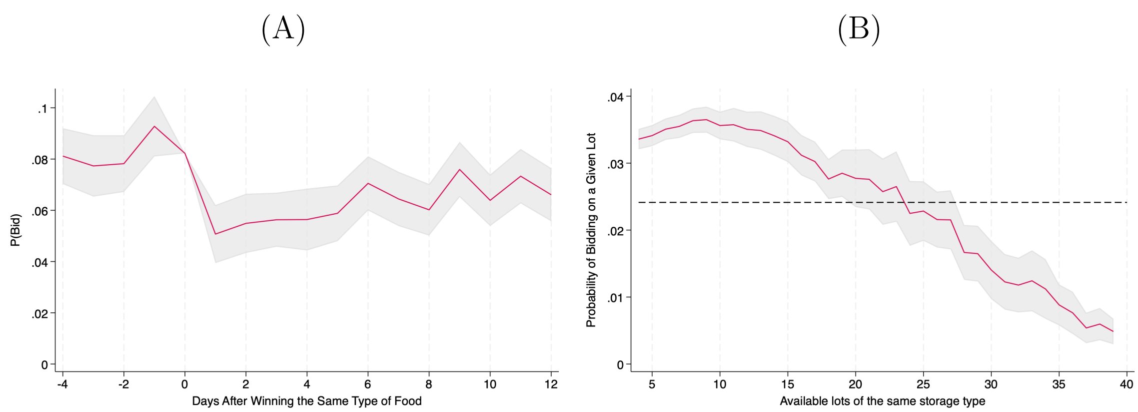

Note: The figure demonstrates evidence of both dynamic and static complementarities. Panel (A) plots the probability of placing at least one bid on a particular type of food, conditional on winning a load of that type of food at time zero. The probability at zero is normalised to the long-run average to demonstrate scale. Panel (B) plots the probability of bidding on a given lot, as a function of the number of lots available from the same storage type. The dotted line gives the unconditional probability of bidding on a given lot. Both plots include food bank × storage type fixed effects, and condition on the type of food being available. The grey area plots 95% robust confidence intervals.

[11] This could be caused by transportation costs or budget constraints. The same can also be said for panel (B) discussed shortly. However, conditional on winning one type of food we see only a 3% drop in the bidding probability on different food. If it were about budget constraints, we should see the same relationship irrespective of the type of food.

3.2.3 Static Complementarities

Figure 5 panel (B) demonstrates that, for a particular type of food, as the number of lots auctioned on a given day increases, food banks bid on a smaller proportion of lots. If payoffs were additively separable we would see a horizontal line. This suggests that lots exhibit a negative complementarity (substitutes) within a storage type — they do not want to win more food than they can afford to store. I cannot treat auction pay-offs as additively separable, and must instead take a multi-object approach, accounting for the simultaneous auction environment.

4 The Model

I now present the model of food banks bidding on the Choice System. Section 4.1 introduces the environment and primitives then 4.2 outlines the agents' dynamic optimisation problem. Section 4.3 discusses equilibrium and 4.4 considers identification.

4.1 Market and Primitives

Each period t over an infinite horizon, N food banks compete in up to L first-price auctions. Food banks are denoted by i and lots by l. a denotes the combination outcome from a round of auctions, i.e. the combination of lots each food bank won.

4.1.1 Auction Environment

Players simultaneously choose which lots to enter and what to bid. Entry decisions consist of an L dimensional vector $\mathbf{d}_{it}$. Entry $d_{itl} = 1$ if they enter lot l, $d_{itl} = 0$ otherwise. Each player plays an L dimensional vector of bids, denoted $\mathbf{b}_{it}$, with $b_{itl} = \emptyset$ if $d_{itl} = 0$. Bids must weakly exceed the reservation price, so that $b_{itl} \geq R_{tl}$ if $d_{itl} = 1$. Auctions are costless to enter. Winners are announced simultaneously. Winners pay their bids, and every player observes the identities and bids of winners.

Each lot is characterised by a row-vector of characteristics $\mathbf{c}_{tl}$, consisting of the the location, size, categories (c), subcategories (h), and storage method (g) of the lot. The pounds in each lot from each category/subcategory/storage method is denoted by $\left\{\mathbf{z}^c_{tl}, \mathbf{z}^h_{tl}, \mathbf{z}^g_{tl}\right\}$, so that if i wins lot l their stock of food from each category increases by $\mathbf{z}^c_{tl}$. For notational convenience I absorb these variables into the common state variable $\mathbf{s}_{0t}$. I make the following assumption about the common state variables:

Assumption 1. $\mathbf{s}_{0t}$ *follows an exogenous Markov process, drawn from* $F^0(.|\mathbf{s}_{0t-1})$.

This assumption ensures that supply and lot characteristics are exogenous. This requires that supply does not react to prices in the Choice System.

4.1.2 Primitives: States and Transitions

Food bank i begins the period in state $\mathbf{s}_{it} \in \mathbb{S}$. This represent the food bank's current stock of food. I primarily focus on their stocks from each storage method, so that the individual state has 4 dimensions.[12] This captures the dynamic costs of holding storable goods. If the outcome from period t is a they end the period in state $\mathbf{s}^a_{it}$. $\mathbf{s}_{it} = \mathbf{s}^a_{it}$ if and only if the player does not win a single lot. Writing $\mathbf{w}_{it}$ as i's winnings from period t I make the following assumptions about how stocks transition:

Assumption 2. *(i) Each period* $\mathbf{s}_{it+1}$ *is drawn from distribution* $F^{\mathbf{s}}_i(\mathbf{s}_{it+1}|\mathbf{s}_{it} + \mathbf{w}_{it})$

(ii) $E[\mathbf{s}_{it}] = 0$

(iii) For any winnings $\mathbf{w}_{it}$ *we have that:* $E[g(\mathbf{s}_{it})|\mathbf{w}_{it}, \mathbf{s}_{it+1}] = 0$ *implies* $g(\mathbf{s}_{it}) = 0$ *for any real bounded function* g.

Individual states are not observed. Day-to-day variation in stocks is likely a major source of variation in bidding behaviour. Food banks supplement their stocks of one type of food they have not recently received from local donors with food from Feeding America. The random variable $\mathbf{s}_{it}$ is observed each morning before items are posted on the Choice System, and depends on the stocks and winnings from the previous period. The random component of the stock process can be interpreted as the net daily change in food banks stocks — the food received from local donors, less the food taken by clients. This transition process incorporates the crucial assumption that food received from Feeding America, and food from local donors, are perfect substitutes. This is a necessary scale normalisation. Part *(ii)* of the assumption is a location normalisation, as only changes in stocks are identified. Part *(iii)* is a

[12] I also focus on their stock of each subcategory h in order to capture food banks' preferences over how the food is used. However, I will assume that payoffs are affine in subcategory stock (not subject to diminishing returns — food banks always have people to feed), meaning that the level of the stock of each subcategory is neither identified nor welfare relevant (up to normalisation).

high-level 'completeness' condition needed for identification, akin to an instrument relevance condition. It essentially requires that, given $\mathbf{w}_t$, the distribution of stocks at $t+1$ can be inverted to give the distribution of stocks at t.

4.1.3 Primitives: Payoffs

Following Altmann (2022) and Gentry et al. (2023) I decompose the flow payoff into a stochastic lot-specific component and a deterministic function of stocks:

Assumption 3. *(i) If food bank i ends the period with stocks $\mathbf{s}_{it}$, they receive payoff $\pi_i(\mathbf{s}_{it})$. The deterministic function $\pi : \mathbb{S}_i \to \mathbb{R}$ is finite, with $\pi(0)$ normalised to 0.*

(ii) If i wins lot l in period t they receive payoff υ_{ilt}, where (stacking over l) $\boldsymbol{\upsilon}_{it} \sim F_i^{\upsilon}$ is a random variable, known privately, observed before entry, and drawn independently across i and t, with $E[\boldsymbol{\upsilon}_{it}|\mathbf{s}_t] = 0$.

(iii) Payoffs are quasi-linear in shares (virtual currency).

The flow payoff function π captures both the costs of storing food, and the utility from holding various types of food to be able to distribute them to clients. Part (*ii*) embeds two assumptions. Assuming the privately known $\boldsymbol{\upsilon}_{it}$ is conditionally independent across individuals imposes independent private values. Assuming conditional independence across time is a standard assumption in most dynamic models. The mean independence assumption is required to separately identify F_i^{υ} and π_i, where the mean on $\boldsymbol{\upsilon}_{it}$ is treated as an element of π. The assumption that π has finite range is predominantly for mathematical convenience, while the normalisation is required as only marginal payoffs are identified. Quasi-linearity is standard in empirical auction studies.[13] However, I allow food banks to differ in their marginal value of wealth, given by $\lambda_i > 0$, capturing that some food banks receive more shares than others, and that some food banks rely on the Choice System more than others.[14]

[13] Altmann (2024) shows that quasi-linearity is observationally equivalent to a model with an inter-temporal budget constraint in a stationary environment. In this case, day-to-day fluctuations in budgets or stocks do not significantly impact expectations about how valuable accessing food from Feeding America will be in future, so the marginal value of wealth is constant over time.

[14] This is important for analysing the distributional consequences of choice. Usually, λ_i are equalised across bidders, enabling inter-bidder welfare comparisons. Instead, I assume the variance of υ_{ilt}, conditional on lot characteristics, is constant across food banks, equalising variance in the unmodelled variation in lot specific attributes. This is similar to the decision of whether to normalise the variance or one coefficient in binary choice models.

I also assume players have temporally additively separable preferences, and make forward looking decisions with discount parameter $\beta = 1$, so that food banks are extremely patient.[15] I assume F, π, $\mathbf{s}$, and β are common knowledge.

4.2 The Agent's Problem

A pure strategy consists of a mapping from a player's type and the state of the world onto entry decisions and bids $(\mathbf{d}_{it}, \mathbf{b}_{it})$. Ex-ante a player's strategy, Λ_i, admits a distribution of bids according to F_i, π_i and $\mathbf{s}$.

4.2.1 Beliefs

Denote $\Gamma_{il}(\mathbf{b}, \mathbf{d}; \Lambda_{-i})$ player i's belief about the marginal probability of winning lot l, given their bid and entry decision and the strategies of other players. Denote $P_{ia}(\mathbf{b}, \mathbf{d}; \Lambda_{-i})$ i's belief about the joint probability, conditional on $(\mathbf{b}, \mathbf{d}, \Lambda_{-i})$, that the combination outcome from the round of auctions is a. Γ and P constitute food banks' beliefs about other players' behaviour. In section 5 I make assumptions about these beliefs to make estimation feasible.

4.2.2 Value Function and Continuation Value

Assuming risk neutrality the Bellman equation is given by:

$$W(\boldsymbol{v}, \mathbf{s}; \pi, \Lambda_{-i}) = \max_{\mathbf{b},\mathbf{d}} \left\{ \bar{W}(\mathbf{b}, \mathbf{d}; \boldsymbol{v}, \mathbf{s}, \pi, \Lambda_{-i}) \right\}$$

$$\text{Where } \bar{W}(\mathbf{b}, \mathbf{d}; \boldsymbol{v}, \mathbf{s}, \pi, \Lambda_{-i}) =$$

$$\sum_l \underbrace{\Gamma_l(b_l, d_l; \Lambda_{-i})(v_l - \lambda_i b_l)}_{\text{lot specific}} + \sum_a \underbrace{P_a(\mathbf{b}, \mathbf{d}; \Lambda_{-i})[\pi(\mathbf{s}_i^a) + \beta \overbrace{\int_{\tilde{\mathbf{s}}} \int_{\tilde{\boldsymbol{v}}} W(\tilde{\boldsymbol{v}}, \tilde{\mathbf{s}}; \pi, \Lambda_{-i}) dF_i^{\boldsymbol{v}}(\tilde{\boldsymbol{v}}|\tilde{\mathbf{s}}) dF^{\mathbf{s}}(\tilde{\mathbf{s}}|\mathbf{s}^a)]}^{\text{continuation value}}}_{\text{combination specific}}$$

The continuation value gives the expected pay-off from the start of the following

[15] This was motivated by conversations with food bank managers. $\beta = 1$ requires a (large) finite dependence assumption, implied by stationarity of the equilibrium process, to ensure the difference between the value function at any two distinct states converges. Therefore, just as we have to do for the payoff function, the value function for one state must be normalised to zero.

period having ended the current period in state $\mathbf{s}^a$. This can be written as follows:

$$V(\mathbf{s}^a;\Lambda_{-i}) = \int_{\tilde{\mathbf{s}}}\int_{\tilde{\boldsymbol{\upsilon}}} W(\tilde{\boldsymbol{\upsilon}},\tilde{\mathbf{s}};\pi,\Lambda_{-i})dF_i^{\boldsymbol{\upsilon}}(\tilde{\boldsymbol{\upsilon}}|\tilde{\mathbf{s}})dF^{\mathbf{s}}(\tilde{\mathbf{s}}|\mathbf{s}^a)$$

A further important object is the sum of the deterministic payoff function and the discounted continuation value, denoted $\kappa(\mathbf{s};\Lambda_{-i}) = \pi(\mathbf{s}_i) + \beta V(\mathbf{s};\Lambda_{-i})$. I refer to κ as the 'Pseudo-Payoff' function. This is essentially what we would estimate if we incorrectly assume bidders are myopic. Estimating this function is key to my estimation procedure. The importance of this object arises because the value function (and hence the continuation value) can be written as functions of these pseudo-payoffs:

$$W(\boldsymbol{\upsilon},\mathbf{s};\pi,\Lambda_{-i}) = \max_{\mathbf{b},\mathbf{d}}\left\{\sum_l \Gamma_l(b_l,d_l;\Lambda_{-i})(\upsilon_l-\lambda_i b_l) + \sum_a P_a(\mathbf{b},\mathbf{d};\Lambda_{-i})\kappa(\mathbf{s}^a;\Lambda_{-i})\right\} \tag{1}$$

4.3 Equilibrium

I focus on symmetric Markov Perfect Equilibria (MPE), defined as follows:

Definition 4.1. : *An MPE consists of a set of strategies $\boldsymbol{\Lambda}^*$ and beliefs $\Gamma(\boldsymbol{\Lambda}^*)$, such that for any $(\boldsymbol{\upsilon},\pi,\mathbf{s})$:*

Optimality: $(\mathbf{b}_i^{\boldsymbol{\Lambda}^*},\mathbf{d}_i^{\boldsymbol{\Lambda}^*}) = arg\ \max\left\{\bar{W}(\mathbf{b},\mathbf{d};\boldsymbol{\upsilon},\mathbf{s},\pi,\Lambda_{-i}^*)\right\}$

Consistency: $\Gamma_{il}(b_{il},d_{il};\Lambda_{-i}^*) = \mathbb{I}[d_{il}=1]P(b_{il} > \underset{i'\neq i}{max}\left\{b_{i'l}\right\}|\Lambda_{-i}^*)$

The optimality condition requires that agents maximise the net present value of payoffs. The consistency condition requires that bidders' beliefs are consistent with the observed distribution of winning bids. This also requires bidders' beliefs about P are consistent with observed joint probabilities. Symmetry requires that bidders with the same 'type', and the same beliefs, place the same bids. This allows us to write the equilibrium strategies as a function of the state: $\boldsymbol{\Lambda}^* = \boldsymbol{\Lambda}(\mathbf{s})$.

Altmann (2022) proved that, conditional on existence of a symmetric Pure Strategy Nash Equilibrium in the bidding game conditional on entry, such an equilibrium exists.[16] I make the following assumptions about equilibrium:

[16]Many papers have also studied sufficiently complex auction games in which neither existence

Assumption 4. *(i) The data are generated by strategy profile* $\boldsymbol{\Lambda}^*$*, a symmetric MPE of the dynamic auction game.*

(ii) $\forall$ *i, l, and* $b_{il} > R_l$, $\Gamma_{il}(b_{il}, 1|\mathbf{s})$ *is strictly increasing and differentiable in* b_{il}*.*

(iii) $\forall$ *i and* $\mathbf{s}_i$ *the Hessian of the pseudo-payoff function* κ *has full rank.*

Part (i) is standard, ensuring that the observed data are generated by a stationary process. This embeds the stronger assumption that, in equilibrium, food banks' states do not trend upwards or downwards.[17] Part (ii) ensures that standard first order conditions are necessary for optimality, ensuring point identification. I allow for ties at the reservation price, which imply non-differentiability of Γ at R. Part (iii) is needed for identification: Conditional on the pseudo-payoff function κ it allows the first order conditions to be inverted for $\mathbf{s}_i$.

4.4 Identification

Altmann (2022) proves that π_i and F_i^{υ} are non-parametrically point identified from the observed distribution of equilibrium bids conditional on state variables. The difficulty, when we do not observe stocks, is that we must identify a non-linear function of this unobserved variable π, as well as the transition process for the unobserved variable $F^{\mathbf{s}}$. Previous work on storable goods models have only discussed identification informally (see, for example, Hendel and Nevo (2006)).

Instead, I prove that the model is non-parametrically point identified building on the framework developed in Hu and Shum (2012). This approach employs an argument based on the spectral decomposition of linear operators, building on Hu and Schennach (2008)'s work on identification of measurement error models. In this setting, the central idea is that the joint distribution of bids and winnings act as a noisy signal of the unobserved states, in that winnings are an observed change in the unobserved stocks. The joint distribution of this noisy signal over time can then be decomposed to separately identify the conditional bid distribution and the

nor uniqueness of equilibrium can be guaranteed, including Gentry et al. (2023) on simultaneous first-price auctions, Fox and Bajari (2013) on simultaneous ascending auctions, and Jofre-Bonet and Pesendorfer (2003) on dynamic auctions. Therefore, I do not consider this a first-order problem.

[17] I require that the distribution of local donations and food sent to clients is constant over my 3 year period. Feeding America's 'Hunger in America' resource shows that food bank usage and food insecurity remains stable in this time. Meanwhile, Altmann (2024) demonstrates that the distribution of bidding behaviour does not vary over the long term, only over the short term.

stock process, from which Altmann (2022)'s results apply. I require two additional invertibility assumptions akin to "instrument relevance" conditions:

Assumption 5. *(i) For any set of available lots* $\mathbf{s}_{0t}$ *we have that:* $E[g(\mathbf{s}_t)|\mathbf{s}_{0t}, \mathbf{b}_t] = 0$ *implies* $g(\mathbf{s}_t) = 0$ *for any real bounded function* g*.*

(ii) For any $\mathbf{s}_{0t+1}$ *there exists a pair* $(\mathbf{s}_{0t}, \mathbf{w}_t)$ *and a neighbourhood* $(\bar{\mathbb{W}}, \bar{\mathbb{S}}_0)$ *around* $(\mathbf{w}_t, \mathbf{s}_{0t})$ *such that for any* $(\bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}) \in (\bar{\mathbb{W}}, \bar{\mathbb{S}}_0)$ *we have:* $E[g(\mathbf{b}_{t+1})|\mathbf{s}_{0t+1}, \bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}, \mathbf{w}_{t-1}] = 0$ *implies* $g(\mathbf{b}_{t+1}) = 0$ *for any real bounded function* g*.*

These assumptions are standard 'completeness' conditions, the non-parametric analogue of the rank conditions required in linear regression and instrumental variable models. For example, Newey and Powell (2003) and Berry and Haile (2014) make similar assumptions for the identification of non-parametric instrumental variable and demand models respectively. Part i) requires that, for any set of available lots, bids are associated with 'enough' variation in current stocks to enable inversion of the conditional bid distribution. Part ii) requires that, for some $\mathbf{w}_t$, variation in past winnings $\mathbf{w}_{t-1}$ still create 'enough' variation in future bids $\mathbf{b}_{t+1}$ to pin down any function of these bids. Parts ii) is testable in principle, and is essentially the same variation plotted in Figure 5 panel (A).

These assumptions enable proof of the following proposition:

Proposition 1. *Under Assumptions 1-5 the distribution of idiosyncratic payoffs* F_i^v*, the flow-payoffs* π_i*, and the stock process* $F_i^{\mathbf{s}}$ *are non-parametrically point identified.*

Proof is given in Appendix C. The proof is omitted for brevity as it requires formal definition of linear operators and their spectral decomposition. The proof follows a simplified version of Hu and Shum (2012)'s argument, exploiting the exclusion restriction that, conditional on $\mathbf{s}_t$, $\mathbf{b}_t$ is independent of $\mathbf{w}_{t-1}$. This means I require weaker completeness conditions than their argument, as well as weaker normalisations because of how my signal (winnings) impacts the unobserved state in a known, additive, way. Because stocks are continuous and multi-dimensional, the argument requires the same of bids, so requires the (simultaneous) auction environment. For this reason, their argument has not been applied in other storable goods models.

4.4.1 Intuition Behind Identification

It is valuable to consider intuitively how we separately identify π_i and $F_i^{\mathbf{s}}$. Identification of the remaining objects is standard. The argument rests on the reduced form relationships plotted in Figure 5 and requires a panel of data for each food bank.

First, variation in the size and composition of lots available (i.e. choice sets), and how this impacts bidding behaviour, identifies the flow payoff function π. If a food bank generally receives high payoffs from holding a particular type of food, e.g. cereal, then the food bank will generally bid frequently on cereal. However, if the food bank also faces convex storage costs for dried food, then when there are many loads of cereal available the food bank may bid on only few of them. They do not want to win too many loads of cereal or face excessive storage costs. This is precisely the behaviour seen in practice in Figure 5 panel (B): When there are ten lots of dried food available food banks bid on a given lot with 3.5% probability (0.35 bids per round). But, when there are thirty lots available they only bid with around 1.3% probability (0.4 bids per round). We can also identify cross substitution effects from variation in available lots of one type on bids of a different type.

Second, variation in a food banks' winnings, and how this affects their subsequent bidding behaviour, identifies their stock process. This is the relationship plotted in Figure 5 panel (A). If a food bank's clients generally take much more cereal than the food bank receives from local donors, then after winning cereal their propensity to bid will not change substantially. They know they will give out this extra load quickly and immediately need more. On the other hand, if they take several days to give out a load of cereal then they may stop bidding on cereal and similar dried foods while their stocks are high, to avoid increased storage costs. Furthermore, the persistence of this change in behaviour identifies how much influence the food bank has over their net donations. If they have a lot of influence then after winning cereal they can give out more cereal to clients and request less from donors. Their stocks quickly return to normal levels, and so bidding behaviour also bounces back quickly. However if they have only limited control then they are unable to shift the extra cereal any faster than usual, keep receiving the same types of food from local donors, so their bidding behaviour takes longer to recover. In practice we see this type of persistence.

5 Estimation

I now describe my estimation procedure, which builds on Altmann (2022) to allow for the unobserved stocks. Section 5.1 outlines this three step procedure. While Altmann (2022) proposed a non-parametric procedure, this is inapplicable in the presence of unobserved stocks and so I take a fully parametric approach. Section 5.2 discusses parametrisation and estimation of beliefs. Section 5.3 discusses the second estimation step, in which I simultaneously estimate the stock transition process, the distribution of lot-specific values, and the pseudo-payoff function. In section 5.4 I discuss disentangling the flow payoff π and the discounted continuation value from the pseudo-payoff κ. Additional details are given in Appendix F.[18]

5.1 The 3-Step Procedure

The standard approach to estimating dynamic auction games, from Jofre-Bonet and Pesendorfer (2003), relies on the ability to write the value function as a function of the distribution of bids only. This requires that bid functions are invertible, just like how Hotz and Miller (1993) and Bajari et al. (2007) require that policy functions are invertible. Invertibility fails in the multi-object context because of an order problem: Bids are L dimensional, while values, and continuation values are 2^L dimensional.

Instead, Altmann (2022) introduces an estimation procedure that does not require invertibility. They demonstrate that we can write the value function as a function of the distribution of bids *and* 'pseudo-static' payoffs, as we did in equation 1. If we 'know' these pseudo-payoffs we can evaluate the value function, and hence the continuation value, which then allows us to back out the flow payoff π from the definition of the pseudo-payoff: $\kappa = \pi + \beta V$. To estimate the pseudo-payoff function we estimate the model almost as if we were estimating a misspecified static model.

In a single-object environment this procedure collapses down to Jofre-Bonet and Pesendorfer (2003), while in a single-agent discrete choice environment it is numerically equivalent to Hotz and Miller (1993). This is because, when the policy function is invertible there is a one-to-one relationship between the policy function and the pseudo-payoff function. Both these methods, as well as Bajari et al. (2007)'s procedure, involve first estimating non-primitive objects (conditional choice probabilities, policy functions, and bid distributions), before using these objects to back out the continuation value. Altmann (2022)'s procedure is similar, but uses estimates of the

[18] Due to computational requirements I focus my analysis on the 90 food banks that each won at least 50 lots. These food banks consume 94% of the food from the Choice System. Because heterogeneity is an important theme of the model I generally estimate separate parameters for each food bank. However, I lack sufficient identifying variation for each individual food bank. I use a Bayesian Hierarchical framework to flexibly introduce information pooling across bidders.

pseudo-payoff function for this purpose.

5.2 Step 1. Beliefs

Assumption 4 ensures food banks form beliefs consistent with observed play. Therefore, we can estimate beliefs using the observed distribution of winning bids without solving the model for equilibrium (Athey and Haile, 2007).

I make three additional assumptions to simplify estimation. First, without additional assumptions the continuation value for food bank i depends on the state of every food bank, creating an infeasibly large state-space. However, $\mathbf{s}_{-i}$ only enters the continuation value of player i through $\Gamma(.|\mathbf{s}_{t+1})$. As the number of bidders grows the probability of any individual and their stocks influencing the distribution of prices falls to zero. To formalise this, I assume that beliefs do not depend on the states of individual food banks. Instead they only depend on aggregate statistics of $\mathbf{s}_t$ through a demand index $\boldsymbol{\vartheta}_t$, detailed shortly.[19] This assumption is tested on the empirical equilibrium winning probabilities in Appendix H.1.

Second, I assume that, in equilibrium, food banks believe winning one lot is conditionally independent of winning any other lot. Equivalently, that winning bids are conditionally independent across auctions, simplifying estimation considerably. In Appendex H.1 I test and present support for this simplification. This allows me to write the combination win probabilities P as products of the marginal probabilities.

Finally, I make flexible parametric assumptions about Γ to facilitate estimation. Because winning bids can be considered the maximum of several conditionally independent variables I assume winning bids follow a generalised extreme value distribution, censored at the reservation price:

$$\Gamma_{il}(.|\mathbf{s}) = GEV(.;\xi_c,\zeta_c,\mathbf{c}_{lt}^T\mu + \vartheta_{lt}) \qquad \text{where} \qquad \vartheta_{lt} = \mathbf{s}_{0t}^T\vartheta \tag{2}$$

Where the shape and scale parameters ξ and ζ are category specific. $\mathbf{c}_{lt}$ gives a vector of lot specific location shifters, such as the subcategory composition. ϑ_{lt} describes how the distribution varies with the state of the world, forming a demand index to be

[19]This assumption is similar to the large market Oblivious Equilibrium (Weintraub et al., 2008) and Moment-based Equilibrium (Ifrach and Weintraub, 2017). Backus and Lewis (2024) employ a similar assumption. They argue that because there are many competitors it is unlikely that bidders follow the identities of which other bidders are likely to bid at any given time, and their states. It is unlikely that any given food bank keeps track of competitors' states.

estimated. This index is a linear function of the quantity of food, by type, auctioned at t and also over the previous 30 days up to $t-1$. This captures the competitive pressures on prices: If little food has been auctioned over the previous month then on average food banks' stocks will be low, and we expect higher prices. Full details of how I estimate beliefs are included in Appendix F.1, including additional covariates and how I estimate the probability of tieing at the reserve price.

5.3 Step 2. The Pseudo-Static Model

I now describe the second part of the estimation procedure, in which I jointly estimate $F^{\mathbf{s}}$, F^{v}, and κ for each food bank. The central difficulty concerns the unobserved state and the unobserved bids, as bids are unobserved when food banks do not enter a particular auction. I estimate the model using a Gibbs Sampler. I use data augmentation to iteratively sample both unobserved bids and unobserved states from their conditional posterior distributions, before updating my parameter estimates given the augmented data. Full details of the estimation procedure, including assumptions on prior distributions, are given in Appendix F.2.

5.3.1 Stock Process

Stock transitions follow $\mathbf{s}^g_{it+1} = \mathbf{s}^g_{it} + \mathbf{w}^g_{it} + \mathbf{x}_{it+1}$, where $\mathbf{x}_{it+1}$ gives the daily local donations minus food taken by clients. I assume $\mathbf{x}_{it+1}|\mathbf{s}^g_{it} \sim N(\boldsymbol{\delta}_i[\mathbf{s}^g_{it} + \mathbf{w}^g_{it}] + \boldsymbol{\mu}_i, \Sigma_i)$. This process incorporates only a simplified form of feedback from a food banks' stocks to their net donations, as this process is difficult to identify. The feedback parameter $\boldsymbol{\delta}_i$ controls the responsiveness of a food banks' net donations to their previous stocks, while $\boldsymbol{\mu}_i$ controls the unconditional average. I impose $\boldsymbol{\delta}_i$ to be a diagonal matrix, with entries $\delta_{ig} \in [-1, 0]$ to ensure stationarity. When $\delta_{ig} = 0$ the food banks' net donations are strictly exogenous, whereas when $\delta_{ig} < 0$ they have some control over their net donations: The higher their stocks, the more of that type of food they send out to clients, and the less of that type of food they procure from local donors.[20]

This stock process is therefore essentially an Auto Regressive process, with autoregressive coefficient $(I+\boldsymbol{\delta}_i)$, drift $\boldsymbol{\mu}_i$, and noise $\tilde{\mathbf{x}}_{it+1} \sim N(0, \Sigma_i)$. Unlike other storable

[20] The previous version of this paper imposed $\delta = 0$, so that net donations were strictly exogenous. I can reject $\delta = 0$ for most food banks. As predicted, allowing $\delta < 0$ leads us to estimate larger welfare benefits of choice as food banks can focus their consumption from Feeding America on the types of food they cannot so easily procure from local donors.

goods models I assume stocks do not exogenously decay over time.[21] Seasonality is indirectly captured through variation in stocks over time, as net donations may be lower in winter when demand for food banks is high.[22] The normality assumption is reasonable for these large food distributors receiving many donations from many different sources, and sending out food to many different food pantries. I estimate food bank $\times$ storage type specific feedback, mean and variance parameters $(\delta_{ig}, \mu_{ig}, \Sigma_{ig})$.

5.3.2 Lot-Specific Payoffs

I assume $\upsilon_{ilt} \sim N(\alpha_i distance_{ilt}, \sigma_l^2)$, so that the mean lot specific payoff depends linearly on the distance between food bank i and lot l. The variance σ_l^2 is category combination specific. Assumption 3 requires that the lot specific pay-offs are uncorrelated across t and i. To simplify estimation I also assume these variables are conditionally uncorrelated across lots l.

5.3.3 Combinatorial Payoffs

I fit a parametric form to the pseudo-payoff function $\kappa(\mathbf{s}_i, \mathbf{s}_0)$. To make estimation feasible I impose a quadratic form, ensuring that inverse-demand is affine in the unobserved stocks. This is similar to standard assumptions of quadratic storage costs, albeit closer to an opportunity costs of storage. Allowing for non-linear demand would require use of multidimensional particle filters, which exponentially increase the computation and memory requirements of estimation. Within these constraints, I choose a parametric function to reflect how food banks gain benefits from food according to how the food is used (depending on the subcategory) and how they face costs of storing the food (depending on the storage method). I assume the following:

$$\kappa_i(\mathbf{s}_i) = \Phi_i \mathbf{s}_i^h - \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g \tag{3}$$

[21]I would not be able to separately identify a decay parameter from δ. This assumption was motivated by discussions with food bank volunteers. Most of the donated food, even fresh produce, have long shelf lives, so that any daily decay parameter is close to 1.

[22]This process does not capture how food banks may begin buying extra food in advance of winter, predicting how demand for their services will change. This behaviour will instead be rationalised by stocks falling in advance of winter. Explicitly modelling seasonality is infeasible as it drastically increases the state space, since behaviour depends on how long is left until winter. Ignoring precautionary consumption should bias my results in favour of the Old System, since precautionary consumption is far easier under the Choice System.

Where Φ_i is an 1×152 row vector, and Ψ_i is a 4×4 dimensional matrix (constrained to be positive), $\mathbf{s}_i^g$ gives the food bank's stock of each storage type and $\mathbf{s}_i^h$ their stock of each subcategory. This form imposes that the benefit of holding food to give out to clients, which comes through $\mathbf{s}_i^h$, does not exhibit decreasing marginal returns. Meanwhile the opportunity cost of storing food is increasing and convex in $\mathbf{s}_i^g$, so that Ψ controls both the dynamic and static complementarities across lots.[23]

5.3.4 Estimation procedure

The model presented above leads to necessary optimality conditions for bidding which can be inverted for the Inverse Bid System, $\xi_{ilt}(\mathbf{b}, \mathbf{d}|\mathbf{s}_i, \mathbf{s}_0)$. Derivation of this system is given in Appendix D. This gives us the following three equation model, consisting of a 'Transition Equation', an 'Observation Equation', and a 'Censoring Equation':

$$\begin{aligned}
&\mathbf{s}_{it}^g = (I + \boldsymbol{\delta}_i)[\mathbf{s}_{it-1}^g + \mathbf{w}_{it-1}^g] + \boldsymbol{\mu}_i + \tilde{\mathbf{x}}_{it} && \rightarrow \textbf{Transition Eq.} \\
&\lambda_i y_{ilt} = \Phi_i \mathbf{z}_{tl}^h - \mathbf{z}_{tl}^{gT} \Psi_i (\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2\sum_{m \neq l} \Gamma_m(b_{itm}) \mathbf{z}_{tm}^g) + \upsilon_{ilt} && \rightarrow \textbf{Observation Eq.} \\
&y_{itl}^* = \begin{cases} b_{itl} + \frac{\Gamma_l(b_{itl})}{\nabla_b \Gamma_l(b_{itl})} & if \;\; b_{itl} > R \\ R + \frac{\Gamma_l(R+1)}{\Gamma_l(R+1) - \Gamma_l(R)} & if \;\; d_{itl} = 1, b_{itl} = R \\ R & if \;\; d_{itl} = 0 \end{cases} && \rightarrow \textbf{Censoring Eq.}
\end{aligned} \tag{4}$$

The observation and censoring equations come from the inverse bid system, while the transition equation was defined in Section 4. Importantly, the Observation Equation is affine in the unobserved state $\mathbf{s}_{it}^g$. Therefore the model is a case of a Censored Linear Gaussian State-Space model.[24]

[23] κ should depend on $\mathbf{s}_0$, capturing how the continuation value depends on beliefs about future supply. I treat κ as independent of $\mathbf{s}_0$ for several reasons. First, the supply of shares varies with supply to ensure prices remain approximately constant. However, relative prices may still vary. Second, I show in section 6.1 that the relationship between supply of different types of food and their prices is not economically significant. Finally, in Appendix H.2 I present results from a specification that includes the demand indices ϑ_{ltg} as inputs to κ and find that these generally have no impact.

[24] b_{itm} may be correlated with υ_{itl}, creating endogeneity in the Observation Equation. When υ_{itl} is large i may prefer to win lot l instead of lot m, lowering their bid on lot m. Results in Altmann (2022) and simulations suggest the resulting bias is very small, as $\Gamma_{im}(b_{itm})$ is generally unresponsive to υ_{itl}, depending much more on $\upsilon_{itm}, \mathbf{z}_{itm}$ and even $\mathbf{z}_{itl}$. In Appendix H.2 I use the instrumental

The likelihood for this model intractable. Instead, estimation is performed using a Gibbs Sampler, which consists of repeating the following steps:[25]

1. Draw beliefs Γ from their posterior distribution using Metropolis Hastings
2. Given Γ, the parameters of the pseudo-static model $\{k_i, F_i^v, F_i^{\mathbf{s}}\}_N$, and states $\{\mathbf{s}_{it}^g\}_{T,N}$, draw censored values of $\{y_{ilt}\}_{NTL}$ using the Censoring Equation
3. Given Γ, $\{k_i, F_i^v, F_i^{\mathbf{s}}\}_N$, and $\{y_{ilt}\}_{NTL}$, use the Carter-Kohn Algorithm to draw $\{\mathbf{s}_{it}^g\}_{T,N}$ using the Transition and Observation equations.
4. Given Γ, $\{y_{ilt}\}_{NTL}$ and $\{\mathbf{s}_{it}^g\}_{T,N}$, draw $\{k_i, F_i^v, F_i^{\mathbf{s}}\}_N$ by running bayesian regressions on the Transition and Observation equations.

5.4 Step 3. The 'Dynamic' Game

At this point we have draws of beliefs, $\{\kappa_i, F_i^v, F_i^{\mathbf{s}}\}_N$, and $\{\mathbf{s}_{it}^g\}_{T,N}$ from the posterior distribution. In order to evaluate the continuation value $V(\mathbf{s}_i^g, \mathbf{s}_0)$ I make use of the following proposition:

Proposition 2. *The ex-ante Value Function can be expressed as:*

$$E[W(\boldsymbol{\upsilon}_{it}, \mathbf{s}_i, \mathbf{s}_0)|\mathbf{s}_i, \mathbf{s}_0] = \frac{E[q_t(\mathbf{s}_i^g)\tilde{W}(\mathbf{b}_{it}, \mathbf{d}_{it}|\mathbf{s}_i^g, \mathbf{s}_0)|\mathbf{s}_0]}{E[q_t(\mathbf{s}_i^g)|\mathbf{s}_0]}$$

Where $q_t(\mathbf{s}_i^g)$ gives the posterior probability that $\mathbf{s}_{it}^g = \mathbf{s}_i^g$ and

$$\tilde{W}(\mathbf{b}, \mathbf{d}|\mathbf{s}_i, \mathbf{s}_0) = \sum_l \lambda_i \frac{\Gamma_l(b_l, d_l; \mathbf{s}_0)^2}{\nabla_b \Gamma_l(b_l, d_l; \mathbf{s}_0)} + \sum_{m \neq l} \Gamma_l(b_l, d_l; \mathbf{s}_0)\mathbf{z}_l^{gT} \Psi_i \mathbf{z}_m^g \Gamma_m(b_m, d_m; \mathbf{s}_0) - \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g$$

This proposition is proven in Appendix E.[26] The identity $\tilde{W}(\mathbf{b}, \mathbf{d}|\mathbf{s}_i, \mathbf{s}_0)$ follows from Altmann (2022), substituting the first order conditions back into the maximand, writing the ex-ante value function as a function of bids and the pseudo-payoffs. The

variable procedure of Altmann (2022) to demonstrate evidence of negligible bias.

[25] Recognise how this procedure builds on the identification argument presented in 4.4. In step 3. I use variation in winnings and the effect on bidding behaviour to infer changes in stocks, pinning down the distribution of net donations. In step 4. I use variation in $\mathbf{z}_t$ as well as winnings (through the sampled states), and how these impact bidding, to pin down κ.

[26] The version of the proposition I prove also includes an adjustment for binding reservation prices. This is excluded from the text for ease of exposition.

main proof extends results from Arcidiacono and Miller (2011) to the continuous choices. The sample counter-part to this object is then easily found. See Appendix F.3 for full details of this procedure. I evaluate the value function across a 30^4 grid of stocks, evaluated evenly across points from the posterior sampled states.[27]

Having evaluated the ex-ante value function for a parameter draw, I evaluate the continuation value using $V(\mathbf{s}_i, \mathbf{s}_0) = \int\int E[W(\boldsymbol{\upsilon}, \tilde{\mathbf{s}}_i, \tilde{\mathbf{s}}_0)|\tilde{\mathbf{s}}_i, \tilde{\mathbf{s}}_0]dF(\tilde{\mathbf{s}}_0|\mathbf{s}_0)dF(\tilde{\mathbf{s}}_i|\mathbf{s}_i)$. Finally I back out $\pi(\mathbf{s}_i) = \kappa(\mathbf{s}_i, \mathbf{s}_0) - \beta V(\mathbf{s}_i, \mathbf{s}_0)$.

6 Estimation Results

This section discusses the results from the three stages of estimation described in section 5 as well as how well the model fits the data. Only a small number of key results are reported in the text, focusing on the theme of heterogeneity. Additional results are reported in Appendix G, including the Gelman-Rubin convergence statistics. When discussing statistical significance I focus on 95% credible intervals.

6.1 First Stage Results

The key first stage parameters are the shape, scale, and location parameters that describe the generalised extreme value distribution. The shape parameters lie significantly within the interval $(-0.1, 0.5)$, with none of the parameters significantly below zero. The scale parameters are all between 2000 and 5000. The implied variance is much higher than the variance of winning bids. This variation is needed to rationalise the relatively high likelihood (≈ 0.3) of winning at the reservation price.

The estimated subcategory fixed effects are precisely estimated, widely dispersed, and strongly correlated with the average winning bids across subcategories presented in Figure 3 ($\rho = 0.829$). The previous 30 days supply of food has a significant negative effect on prices for most types of food. Each additional increase in the previous 30 day supply by one hundred loads ($\approx 1sd$) decreases the winning bid by 350 shares. This magnitude, while statistically significant, is not economically significant (around

[27] Such a large grid is feasible in this context as I only need to perform the procedure once. However, storing $85,000$ grids, one for each food bank $\times$ parameter draw, is not. I use a quadratic approximation of the ex-ante value function. In Appendix H.3 I evaluate the fit of this approximation by considering the R^2 of the approximation regression. 100% of these R^2s lie between 0.9 and 1. The fit is strong due to the quadratic term that appears in the ex-ante value function.

0.017 standard deviations), relative to the variation seen across different types of food through the subcategory parameters. Given the money supply varies with the food supply to ensure prices remain stable over time, this is unsurprising.

6.2 Second Stage Results

Unobserved State

Figure 6: Estimated Standard Deviation of Net Donations

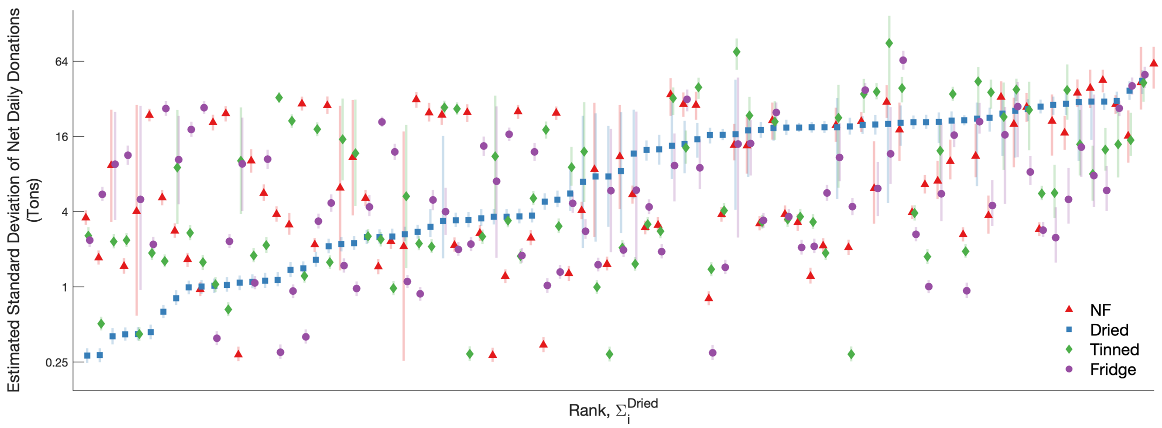

Note: The figure plots posterior means of standard deviations of net local donations, as well as 95% credible intervals. Results are sorted according to the estimates for the Dried storage type.

Figure 6 plots estimates of the standard deviation of food banks' net donation $\sqrt{\Sigma_i}$ parameters. 95% credible intervals are also plotted. Estimates are sorted according to the estimates for the Dried food type.

There are two key takeaways from this plot: First, the extent and significance of the heterogeneity across food banks. Some food banks have much more variable net donations than others, with estimates varying by a factor of 30. We see both vertical and horizontal variation. The second takeaway concerns the scale. The average load of food is around 10 tons, and around half of food banks' net donations vary by more than this each day. Therefore, there is a lot of heterogeneity in the types of food each food bank wants from Feeding America on any given day. Decomposing the unexplained variation in bidding behaviour between F^{υ} and $F^{\mathbf{s}}$, around 93% of the unexplained variation in bidding behaviour is from variation in unobserved stocks.

Results for mean net donations $\boldsymbol{\mu}_i$ and the feedback parameters $\boldsymbol{\delta}_i$ are presented in the same format in Appendix G. While precisely estimated, only 18% of the $\boldsymbol{\mu}_i$ parameters are significantly different from zero (95%), predominantly for food banks who are observed almost always bidding on the same types of food.

Meanwhile, around half of the feedback parameters $\boldsymbol{\delta}$ are estimated to be above -0.1, suggesting that most food banks have very little control over their net donations, and that their stocks exhibit a large degree of persistence.[28]

Lot-Specific and Pseudo-Payoffs

Distance coefficients vary across food banks from a cost of 5 to 120 shares per mile, with an average of 39. These figures are higher than the figure from Prendergast (2022) of around 0.16 shares per mile. However the figures are not directly comparable, as Prendergast takes this figure as the coefficient from a regression of distance on the observed winning bid. When I perform this exercise I find a coefficient of 0.14, which is extremely similar. The difference arises as my analysis includes losing bids and also accounts for both bid shading and endogenous entry.

I estimate significant heterogeneity in the estimated willingness to pay for food ($= \kappa_i(\mathbf{s}_{it} + \mathbf{z}_{ilt}) - \kappa_i(\mathbf{s}_{it})$) across different types of food, across different food banks, and within food banks across time. The posterior mean average willingness to pay is -50,600 ($\pm$ 2,400). This figure varies from -71,500 to -14,800 across different types of food, from -195,000 to 30,200 across food banks, and from -116,000 to 46,600 over time. It is difficult to compare these measures to previous estimates of demand elasticities for food given the specific setting of large food banks using fake money to pay for food, rather than consumers purchasing food. However, the figures are similar in size and magnitude to the average bids presented in Figure 4.

The marginal value of a share λ_i is estimated to vary across food banks by a factor of 4, with λ_i for the food bank with the median consumption normalised to 1. Parameters are negatively correlated with a food bank's Goal Factor, which is sensible since a higher Goal Factor implies more shares. However this relationship is very weak, stressing the importance of unobserved food wealth.

[28] I also correlate my estimates with observable characteristics of food banks, such as population density in their catchment area, and agricultural rents. This analysis is omitted as I do not find any particularly striking results. Correlations are in the expected directions but small. For example, food banks in areas with higher population density or lower agricultural rents are estimated to have lower average net donations.

6.3 Third Stage Results

Figure 7: Estimated Storage Costs

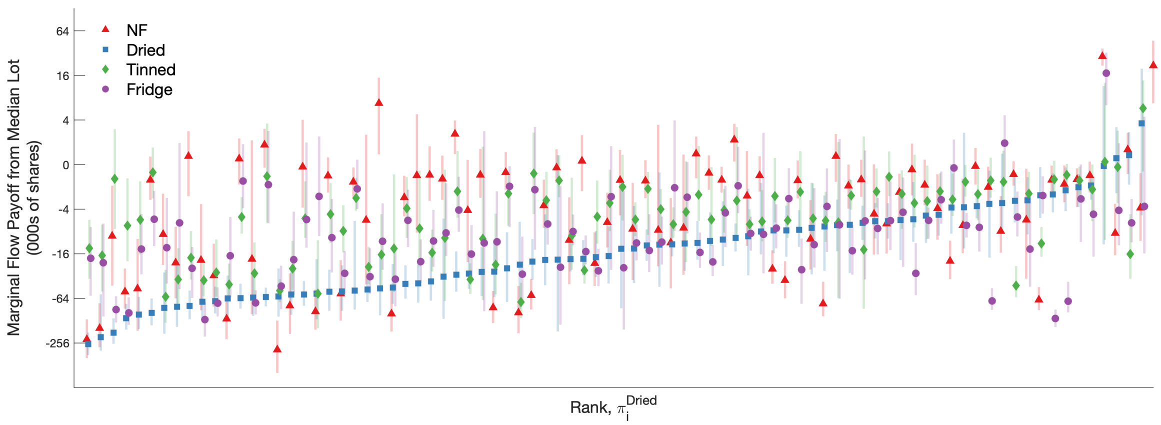

Note: The figures plots posterior mean marginal flow pay-offs (i.e. negative of the marginal storage costs) for a 20,000 load for each storage type. Estimates are ordered according to the estimates for Dried loads. Marginal payoffs are evaluated when stocks are at their long run mean.

Figure 7 plots posterior means of the marginal flow-payoffs from receiving 10 tons of each type of food, evaluated when stocks are at their long run average. Estimates are sorted according to the estimate for the Dried storage type.

I estimate significant differences across food banks, as well as across types of food. Marginal flow pay-offs are generally negative, indicative of storage costs. This suggests different food banks face different storage costs. Positive marginal pay-offs suggest that food banks also benefit from not having an empty warehouse.

6.4 Model Fit

Figure 8 displays the true and predicted moments for various key measures. The model matches the mean and standard deviation of bids conditional on entry, as well as the average number of bids placed by each food bank on a given day. I slightly under predict the probability of bidding on any given lot and the average distance each load of food travels. Figure 9 recreates Figure 5 for simulated data, plotting shaded 95% credible intervals against the estimates from the true data in red. I do not plot standard errors for the true data. In Appendix G I present Gelman-Rubin convergence statistics. Generally, the data converged well.[29]

Figure 8: Model Fit

Measure	Mean True	Mean	q0.025	q0.975	Std True	Mean	q0.025	q0.975
Average Bid given entry	1950	1860	1720	2040	4290	4350	4270	4470
Probability of bidding	0.0275	0.0239	0.0232	0.0252	0.164	0.153	0.151	0.157
Number of Bids per day	0.676	0.683	0.664	0.721	1.66	1.33	1.3	1.36
Distance to lot given won	805	748	735	765	634	587	576	606

Note: This table presents several observed and model simulated moments of the data.

7 Counterfactuals

Feeding America introduced the Choice System, replacing the Old System, to give food banks choice over the food they received. The importance of choice depends on the extent of the heterogeneity in what food banks want from Feeding America. Heterogeneity across types of food, across food banks, and across time. The estimates in the previous section demonstrate evidence of such heterogeneity. Therefore, we have reason to believe that choice is likely to be very important in this setting.

I now use counterfactual simulations to consider the welfare and distributional consequences of Feeding America's transition from the Old System to the Choice System. I then simulate several additional allocation mechanisms to tease out the most important features of the Choice System, features that are useful for other food relief organisation around the world and in other market design settings.

7.1 The Old System

I described the Old System in detail in Section 2. I model the Old System in continuous time, assuming that food banks could receive either an offer of food from Feeding

[29] 100% of first stage parameters are below the standard threshold of 1.1, while 91% of second stage parameters are below the 1.1 threshold and 94% are below the more lenient 1.2 threshold. The model did not converge for 5 food banks, which I then removed due to their implausible parameter estimates. These food banks only consume 1.5% of the food on the system.

Figure 9: Model Fit

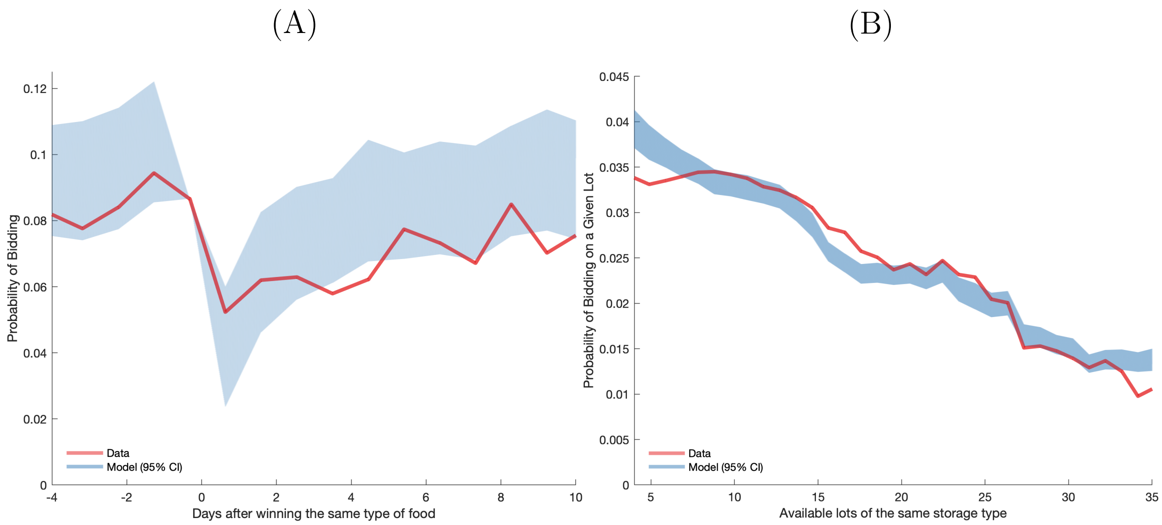

Note: This figure recreates Figure 5 from simulated data, compared to the true data. The shaded area gives 95% credible intervals of the simulated data, while the red line gives mean estimates from the true data. Standard errors for the true data are not plotted. The red lines differ slightly from Figure 5 as I do not include fixed effects or other controls.

America, or a shipment of food from their local donors / to their clients, at any point in time. Continuity of time ensures that multiple of these events occurs simultaneously with probability almost surely zero. I assume food banks do not observe offers made to, nor decisions of, other food banks. They do not know their place in the queue; only their own Goal Factor and when they were last offered a load. They form beliefs about the rate they receive calls from Feeding America, and then the probability of being offered a load with various characteristics. I allow each load of food to be offered to 10 food banks before it is returned to the donor.

I assume a Markov Perfect Equilibrium in symmetric strategies, as defined in section 4. This requires that food banks make optimal acceptance decisions given their beliefs, and that beliefs are consistent with the observed realisation of acceptance rates. Appendix I.2 details how equilibrium beliefs and equilibrium value functions are formed. Given beliefs I find each food banks' value function by numerically solving the Hamilton-Jacobi-Bellman differential equation. I then simulate the mechanism and update beliefs using observed offer rates, repeating until convergence.[30]

[30] While I cannot rule out equilibrium multiplicity, I did not find evidence of multiplicity in simulations. I varied the starting values, varying the initial 'pickyness' of food banks' acceptance decisions: Either as picky as under the Choice System, or as accepting as to take any food offered.

7.1.1 Welfare Measurement

My counterfactuals produce welfare measures in terms of consumer surplus, measured in shares. While this cardinal measure enables inter-food bank welfare comparisons, shares have no value outside the Choice System. Instead, similar to Agarwal et al. (2021), I report welfare as the equivalent change in the supply of food on the Choice System that would have the same total value in shares. This measure is valid under competitive equilibrium because the money supply adjusts to ensure prices are constant, given changes to the supply of food. If the supply of food doubles, the money supply adjusts so that expenditure doubles. Therefore, if consumer surplus under the Old System is double that under the Choice System, I liken this to double the nominal expenditure, which equates to double the supply of food.[31]

Importantly, the 'level' of welfare is not identified because the levels of both stocks and flow payoffs are not identified. I normalise the level of welfare using welfare when Feeding America allocates no food at all. Therefore, welfare results are reported on a scale of zero (food is allocated no better than if no food was allocated) to 320 tons (the daily average food supply under the Choice System).

I report utilitarian welfare and a weighted sum using Goal Factors as priority weights. I also report the total distance travelled and amount of food allocated. These are important measures for policy makers given the significant transport costs as well as the political cost to Feeding America of wasting food.[32]

7.2 Results

Figure 10 presents the headline results, plotting simulated welfare under the Choice System, the Old System, and No Allocation mechanism for each draw of the model parameters from their posterior distribution. Welfare is on average 32.7% higher under the Choice System than the Old System. The transition to the Choice System lead to an increase in welfare equivalent to increasing the supply of donated food by 105 tons per day, which is enough to provide an additional 35,100 meals. These figures

[31] This is a lower bound on the value of the Choice System: If supply and expenditure changes by a factor X, consumer surplus changes by a factor $< X$ due to the concavity of payoffs (storage costs are convex). Therefore doubling consumer surplus requires more than doubling the food supply.

[32] I cannot account for endogeneity in the food supply with respect to the allocation mechanism. Since Prendergast (2017) reports the Choice System caused more food to be donated to Feeding America, we should again consider these results a lower bound on the true benefit.

Figure 10: Welfare

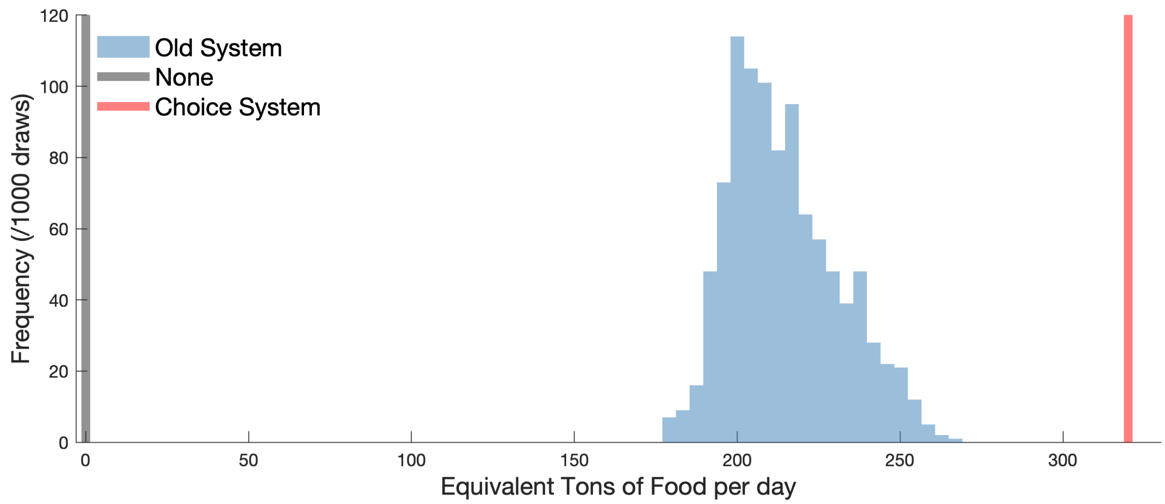

Note: This plot shows the posterior distribution of welfare under each mechanism, evaluated over 1000 draws from the posterior distribution of parameters. Welfare is weighted by Goal Factor and measured relative to no allocation mechanism and the Choice System. On average, welfare increased by the equivalent of increasing the supply of distributed food by 105 tons per day.

are statistically significant at the 0.1% level. When welfare is not weighted by Goal Factor, this figure falls to 27% higher. My results are similar to Prendergast (2022), who finds a welfare improvement of around 21%. Given that my fully structural approach incorporates greater heterogeneity, it is unsurprising I find a larger effect.

Additional outcome measures are reported in Figure 13. Only 227 tons of food are successfully allocated each day, compared to 300 tons under the Choice System, so wastage is reduced by around a quarter. Furthermore, food banks sort into consuming closer lots — the average ton of food travels 79 miles under the Old System, compared to 60 miles under the Choice System. Total storage costs are only 7.4% higher under the Choice System, despite 25% more food being stored, suggesting the food banks are able to tailor their consumption towards the types of food they have space for. Food banks are less picky under the Old System, with food banks around 3 times more likely to accept any given load. They do not know when they can next access the types of food they really want, so accept food even if it does not precisely meet their needs. Food that might better meet the needs of a different food bank that just happened to be lower in the queue.

7.3 Distributional Consequences

Figure 11: Individual Welfare

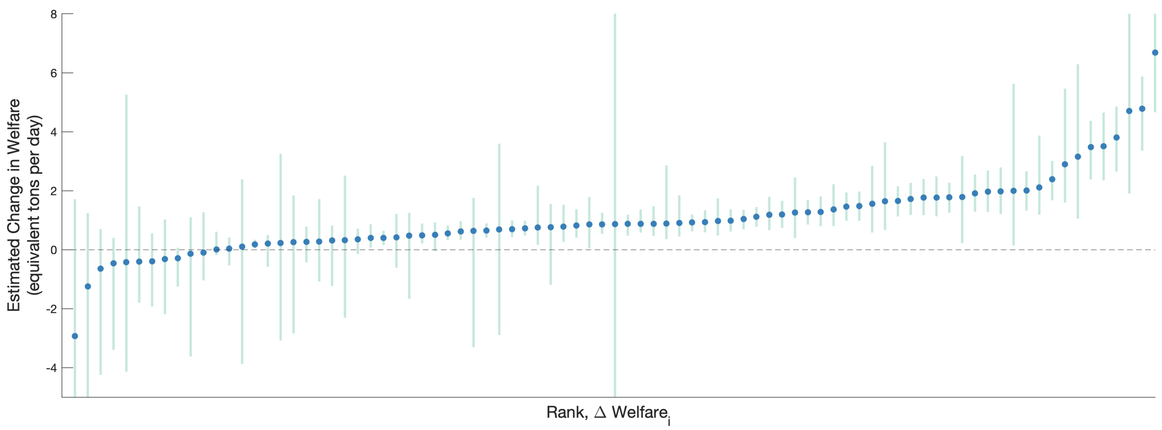

Note: This plot shows food bank specific welfare under the Choice System minus welfare under the Old System, ordered by the welfare difference, with 95% credible intervals across posterior draws. On average 87.9% of food banks are better off under the Choice System than the Old System. 3 food banks have large credible intervals due to their Old System welfare exhibiting long right tails, while another food bank exhibits a long left tail.

Feeding America put significant resources into minimising possible negative distributional consequences of the Choice System. Figure 11 presents welfare results by food bank, plotting the difference between food bank specific welfare under the two systems. On average 87.9% of food banks are better off under the Choice System.

The $\approx$ 12% of Food Banks who are often worse off under the Choice System consume more and higher quality food under the Old System, and tend to be food banks with lower than average Goal Factor. They appear to accept essentially any food rejected by other food banks, benefiting from the bad market design. Otherwise, we see very little correlation between a food bank's welfare difference and their Goal Factor ($\rho = 0.059$).[33] Likewise, we do not see any strong correlations for factors such as population density or deprivation indices. We do see that the food banks which serve known food deserts are all significantly better off under the Choice System.

[33]We see some small positive correlations ($\rho \in [0.08, 0.15]$) between these welfare differences and the estimated state transition parameters δ and Σ. Food banks with more uncertain net donations benefit from being able to choose the food they receive from Feeding America. However these are predominantly driven by one or two food banks with, for example, particularly high Σs who are observed perform particularly well under the Choice System.

7.4 Key Features

I now use additional simulations, varying the different features of these mechanisms, to understand what is driving these welfare gains. This is important for identifying key features we might take to other food allocation settings. I consider four sets of additional mechanisms, some of which are also used in practice by other food relief organisations.[34] Additional computational details are given in Appendix I.

7.4.1 Sequential Offers

Under the Choice System any food bank can bid on any load of food, while under the Old System each load of food could only be offered to around ten food banks. To understand how important it is to allow every food bank access to every load of food I consider the impact of removing the ten food bank constraint. This 'Sequential Offer' mechanism is strategically and outcome equivalent to the 'Like' mechanism proposed by Walsh (2015) and used by Food Bank Local.

Figure 12 plots welfare under 'Sequential Offers' alongside welfare under the Old System and the Choice System. Welfare under Sequential Offers is around 8% higher than under the Old System, equivalent to an additional 17 tons per day. Welfare is still lower than under the Choice System, suggesting that around 16% of the welfare gain from transitioning from the Old System to the Choice System can be attributed to being able to offer food to every food bank. Interestingly, three tons more food is successfully allocated under Sequential Offers than even under the Choice System, corroborating that food banks become less picky under this mechanism.

7.4.2 Sequential Auctions

By simulating sequential (rather than simultaneous) auctions, I consider the relative importance of two key features of the Choice System: That food is allocated in batches, and that food banks can signal the intensity of their preferences.

The value of allowing food banks to (credibly) signal their preferences is evident, ensuring food is more likely to go to the food bank that values it most. Meanwhile,

[34] It is worth recognising that other food relief organisations often face different allocation problems to Feeding America. Nonetheless, these results remain a useful starting point. It would also be interesting to consider directly tweaking features of the Choice System. However, even numerically solving for the new MPE remains intractable. This is an important area for future research.

Figure 12: Additional Mechanisms (1)

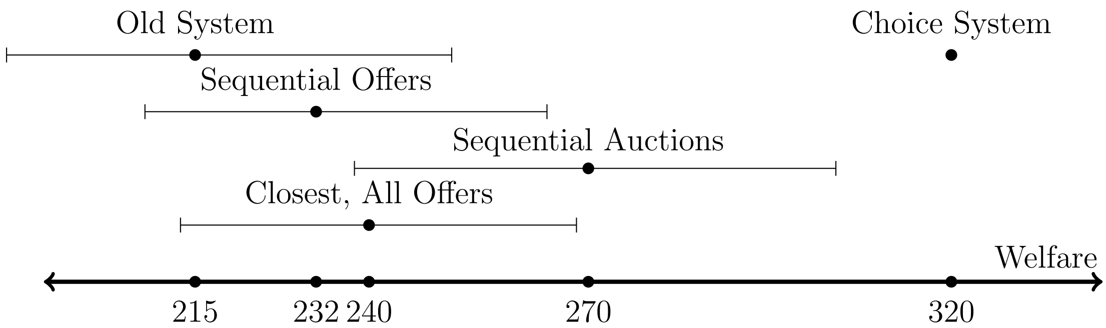

Note: This figure shows estimated welfare, in equivalent tons per day, under the Choice System, Old System, Sequential Offers (Old System with all offers), Closest — All Offers (sequential offers by distance), and Sequential Auctions. Results for the Random and Single Offer Closest mechanisms are excluded for graphability. 95% (marginal) Credible Intervals are plotted.

by allocating food in batches food banks have better information about all the food being allocated that day and so we expect better matches (Akbarpour et al., 2020). Like most other food relief organisations, under the Old System food was allocated sequentially. Each load was distributed before the next arrived. This has a potential benefit that food banks do not risk winning too many or too few loads, alleviating the exposure problem of simultaneous auctions highlighted in Gentry et al. (2023).

Considering sequential auctions teases apart the relative importance of these two features. Comparing Sequential Offers to Sequential Auctions pins down the value of allocating food to the food bank that values it most in that moment, as opposed to essentially offering the food out at random. Then, comparing Sequential Auctions to the Choice System, pins down the importance of allocating food using simultaneous versus sequential auctions.[35]

Results of this decomposition are displayed in Figure 12. Moving from Sequential Offers to Sequential Auctions increases welfare by the equivalent of distributing an extra 38 $(4, 66)$ tons of food per day. As we go from sequential to simultaneous auctions, welfare increases by the equivalent of an additional 50 $(16, 83)$ tons per day, so the batching effect dominates the exposure effect. Therefore, around 36% of the welfare of moving from the Old System to the Choice System can be attributed to the cardinal signal, while 48% can be attributed to the batching effect.

[35]It would be interesting to consider other simultaneous mechanisms, such as a VCG mechanism or simultaneous allocation without the cardinal signals. However simulating equilibrium under such mechanisms in a dynamic setting remains computationally intractable and is left for future work.

Under Sequential Auctions food is always allocated to the food bank who values it most *in that moment*, but it may not be the food that *they* most need at that time. Then, they may forfeit more needed food later in the day because storage is now more costly. These results are identified by the fact that even though only a small number of bidders tend to bid on any given lot, 90% of lots still get allocated immediately. Food banks are bidding on different lots, suggesting that horizontal heterogeneity is more important than vertical heterogeneity. Therefore, it is more important that food banks have information on all the different types and locations of lots than making sure that food always goes to the food bank that values it most.

7.4.3 Closest Offers

Figure 13: Additional Mechanisms (2)

Mechanism	Welfare (unweighted)	Welfare (weighted)	Distance (000 miles per day)	Allocated (tons per day)	% Better Off (compared to CS)
Choice	320	320	18	300	1
System	(320, 320)	(320, 320)	(17.9, 18)	(299, 301)	(1, 1)
Old System	230	215	18.3	227	0.121
	(202, 261)	(188, 251)	(17.4, 19.3)	(220, 234)	(0.071, 0.176)
Sequential	249	232	23.6	303	0.162
Offers	(219, 278)	(208, 264)	(22.9, 24.4)	(301, 306)	(0.106, 0.224)
Closest	81.8	63.6	0.45	61.8	0.066
Single Offer	(54.7, 110)	(42.6, 86.1)	(0.416, 0.487)	(55.6, 68.7)	(0.035, 0.106)
Closest	254	240	9.74	307	0.22
All Offers	(219, 284)	(213, 268)	(9.22, 10.3)	(306, 309)	(0.165, 0.271)
Random	-355	-383	45.7	320	0.001
	(-531, -150)	(-543, -171)	(45.2, 46.2)	(320,320)	(0, 0.012)
Sequential	286	270	20.4	295	0.249
Auctions	(249, 324)	(237, 304)	(19, 22.1)	(291, 300)	(0.188, 0.318)

Note: This table displays posterior means and 95% credible intervals for various measures of welfare. The final column gives the proportion of food banks who are estimated to be (weakly) better off under each alternative mechanism compared to under the Choice System.

Transportation costs are an important determinant of food bank's decision making. To examine the relative importance of transport costs I now consider a mechanism in which food is offered to only the nearest food bank, aiming to minimise transport costs. I also consider a version where food is offered sequentially in order of distance.

These are an important set of mechanisms that are used in practice, implictly or explicitly, by several food relief organisations. Many organisation, including the Trussell Trust (U.K.) and Second Bite (Australia), do not allocate food centrally. Instead, they link their partner food banks up with nearby donors, which is equivalent to offering food only to the closest food bank.

Offering food only to the nearest food bank is equivalent to distributing only 63 tons of food each day under the Choice System, an 80% reduction. This is driven by a comparable drop in the actual quantity of food successfully allocated: There is almost always a slightly more distant food bank willing to pay to transport the lot that extra distance. When food is offered to every food bank welfare increases to roughly the same as under Sequential Offers (240 tons), despite a 60% reduction in transportation costs. This is because it is always the same food banks being offered the same types of food. Transportation costs are not so important a factor as to design a mechanism entirely focused around their minimisation.

7.4.4 Random Allocation

Many food relief organisations pressure food banks to accept all the food they are offered, attempting to minimise food waste. A central question concerns how important it is to allow food banks to turn down food they do not have room for.

I consider an extreme mechanism: Randomly allocating food in proportion to a food bank's Goal Factor. Food banks are only able to reject the load if its stocks for that type of food exceed the highest stocks ever 'observed' under the Choice System. In Figure 13 I show that both weighted and unweighted welfare are significantly lower than *even* allocating no food at all. In 99.5% of the simulations, every food bank would be better off if Feeding America did not exist. Allowing food banks to turn down food is extremely important.

8 Conclusion

The efficient and equitable allocation of food to food banks is of first-order importance for the welfare of many of America's most vulnerable. In this paper I examined the welfare and distributional consequences of giving food banks choice over the types of food they receive, studying Feeding America's Choice System. I developed an empirical model of food banks bidding for food on the Choice System. The central challenge was that I do not observe food banks' stocks, as I do not observe the food taken by clients nor the food received from local donors — the key determinants of bidding behaviour. Nonetheless, I proved this model is identified from standard data, and proposed a Gibbs Sampling procedure to estimate the model primitives.

I then used counterfactual simulations to compare equilibrium allocations under the Choice System to their Old System, which gave food banks very limited choice. The transition to the Choice System increased welfare by 32.7%, driven by the estimated scale and scope of heterogeneity in the types of food needed by different food banks at different points in time. On average 87.9% of food banks are better off under the Choice System. Among other features, these results appear to be driven by 'batching', that food is allocated simultaneously in batches, whereas under the Old System food was allocated sequentially. This is a particularly important finding as most other food relief organisations around the world still allocated food sequentially.

How applicable my results are for other food relief organisations remains an open question. Future work should apply this type of analysis to data from other food bank networks. I also explored only a limited space of counterfactual mechanisms. Further analysis of mechanism that can exploit these batching effects would be extremely valuable, as well as further analysis of the importance of using fake money.

References

Agarwal, N., Ashlagi, I., Rees, M. A., Somaini, P., and Waldinger, D. (2021). Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys. *Econometrica*, 89(1):37–76.

Agarwal, N., Hodgson, C., and Somaini, P. (2020). Choices and outcomes in assignment mechanisms: The allocation of deceased donor kidneys. Technical report, National Bureau of Economic Research.

Akbarpour, M., Li, S., and Gharan, S. O. (2020). Thickness and information in dynamic matching markets. *Journal of Political Economy*, 128(3):783–815.

Altmann, S. (2022). Identification and estimation of a dynamic multi-object auction model, available here.

Altmann, S. M. (2024). *Choice and auction design in the allocation of food to food banks*. PhD thesis, University of Oxford.

Arcidiacono, P. and Miller, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867.

Atchadé, Y. F. and Rosenthal, J. S. (2005). On adaptive markov chain monte carlo algorithms. *Bernoulli*, 11(5):815–828.

Athey, S. and Haile, P. A. (2007). Nonparametric approaches to auctions. *Handbook of econometrics*, 6:3847–3965.

Backus, M. and Lewis, G. (2024). Dynamic demand estimation in auction markets. *Review of Economic Studies*, page rdae031.

Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

Berry, S. T. and Haile, P. A. (2014). Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797.

Bodoh-Creed, A. L., Boehnke, J., and Hickman, B. (2021). How efficient are decentralized auction platforms? *The Review of Economic Studies*, 88(1):91–125.

Botev, Z. I. (2017). The normal law under linear restrictions: simulation and estimation via minimax tilting. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(1):125–148.

Dunford, N. and Schwartz, J. (1971). *Linear Operators*, volume 3. Wiley.

Erdem, T., Imai, S., and Keane, M. P. (2003). Brand and quantity choice dynamics under price uncertainty. *Quantitative Marketing and economics*, 1(1):5–64.

Fox, J. T. and Bajari, P. (2013). Measuring the efficiency of an fcc spectrum auction. *American Economic Journal: Microeconomics*, 5(1):100–146.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (1995). *Bayesian data analysis*. Chapman and Hall/CRC.

Gentry, M., Komarova, T., and Schiraldi, P. (2023). Preferences and performance in simultaneous first-price auctions: A structural analysis. *The Review of Economic Studies*, 90(2):852–878.

Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574.

Hendel, I. and Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673.

Hotz, V. J. and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529.

Hu, Y. and Schennach, S. M. (2008). Instrumental variable treatment of nonclassical measurement error models. *Econometrica*, 76(1):195–216.

Hu, Y. and Shum, M. (2012). Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics*, 171(1):32–44.

Ifrach, B. and Weintraub, G. Y. (2017). A framework for dynamic oligopoly in concentrated industries. *The Review of Economic Studies*, 84(3):1106–1150.

Jofre-Bonet, M. and Pesendorfer, M. (2003). Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489.

Newey, W. K. and Powell, J. L. (2003). Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578.

Prendergast, C. (2017). How food banks use markets to feed the poor. *Journal of Economic Perspectives*, 31(4):145–62.

Prendergast, C. (2022). The allocation of food to food banks. *Journal of Political Economy*, 130(8):000–000.

Verdier, V. and Reeling, C. (2022). Welfare effects of dynamic matching: An empirical analysis. *The Review of Economic Studies*, 89(2):1008–1037.

Waldinger, D. (2021). Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation. *American Economic Review*, 111(8):2660–96.

Walsh, T. (2015). Challenges in resource and cost allocation. In *AAAI*, pages 4073–4077.

Weintraub, G. Y., Benkard, C. L., and Van Roy, B. (2008). Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411.

Appendix

A Data

I now present additional details on how I construct the dataset used in my analysis. Specifically, how I categorised the data and construct Goal Factors.[36]

Categories come from Prendergast (2022), combining several small similar categories (such as pasta and rice). Figure 14 panel (A) plots the categories as a proportion of the total amount of food auction. The plot excludes multiple identical auctions, which has the result of artificially reducing the proportions of Fresh Produce and Beverages down from 24% and 17% respectively. Within categories I form subcategories according to the most common product names.[37] Figure 14 panel (B) plots subcategories as a word cloud, with more common subcategories larger.

Food is categorised by how it is used, including Meals, Ingredients, Condiments, Snacks, and Non-Food. Meals are items that could be eaten on its own as part of a reasonably healthy diet for either breakfast, lunch, or dinner. Multiple Ingredients can be mixed together to form a meal. Condiments can be added to a meal to enhance it. Snacks can be eaten on their own, though not necessarily part of a meal. Snacks includes beverages. Non-food items are inedible items, such as cleaning products. This also includes formula and baby food.

Storage methods includes Dried, Tinned, Refrigerated, and Non-Food. The Non-food category is identical to the non-Food Use category. Dried items can be stored on a shelf, have extremely long shelf lives, and are generally light but bulky. Tinned food, which includes jars and bottles, have long shelf-lives and are generally compact and heavy. Refrigerated food must be stored in a fridge, but still expire reasonably quickly. Any item that was additionally listed as 'Shelf Stable', such as UHT milk was put in the tinned storage category.

To find the distance between every lot $\times$ food bank combination I convert zipcodes into longitude/latitudes, then found the geodesic between these zipcodes using the "distGeo" function from the R package "geosphere". In principle I could have found

[36] I do not directly observe joint bids or food sold by food banks. Because they are not core features of the model, how I identify this data is detailed in the replication package.

[37] This is performed to ensure at least 30 lots per subcategory. Subcategories are more granular the more observations there are. E.g. for cereal and beverages this includes brands, whereas all cheese is lumped together.

Figure 14: Composition of food allocated: Categories and Subcategories

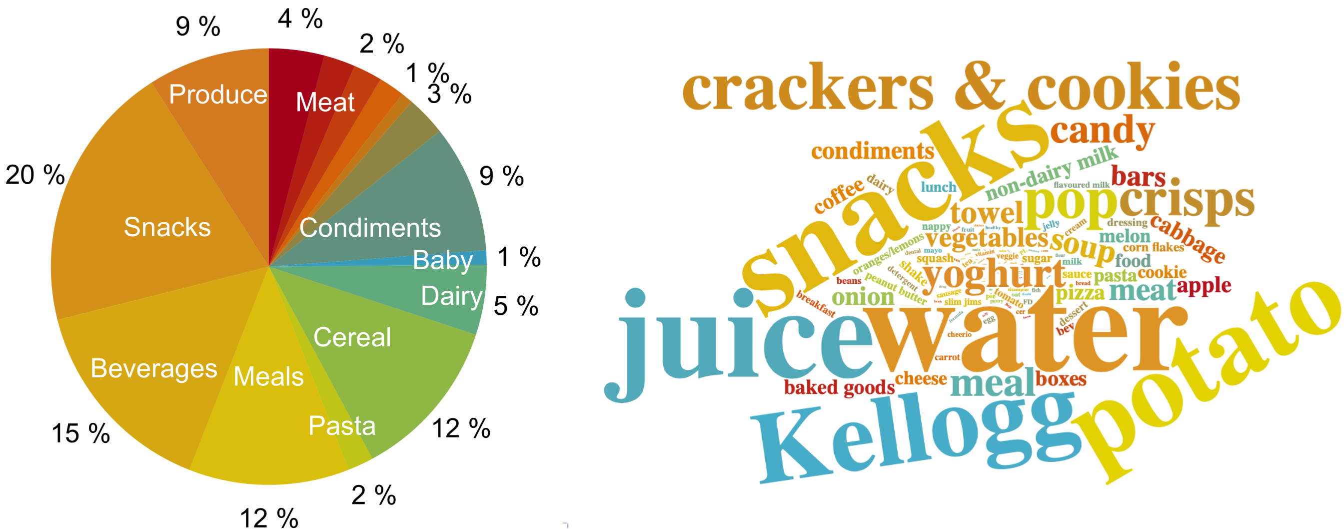

the shortest road distance using arcGIS software, which would more accurately represent the transportation costs. However, this proved too computationally intensive because of the large number of food bank $\times$ lot combinations ($\approx$ 3.6 million).

I imputed Goal Factor figures using the (updated) formulae in Prendergast (2022). For a small number of food banks their expenditure did not match up with their Goal Factors. The largest deviation occurs for a food bank in a known food desert, suggesting that food banks likely contact the 'fairness committee' to request a higher Goal Factor. I account for this deviation by calibrating food banks' Goal Factors and their budget at the beginning of my data period, minimising the distance from my calculated Goal Factor, while maintaining several known constraints (such as shares cannot exceed 200,000). The resulting distribution of calibrated Goal Factors matches the known distribution of Goal Factors (which I have, but cannot link). Details are available in the replication package.

B Additional Descriptive Analysis

In this appendix I present additional evidence of systematic variation in bidding behaviour over time. I use a fixed effects Tobit specification, given in equation B below. I investigate how each food bank i's bid on food of storage type g varies across months m, writing α_{igm} for these average bids. I estimate the model only on food banks who win at least 100 lots over the period. I also control for the distance between the food bank and the lot. I drop the first and last months due to incomplete data. Each $i \times g \times m$ cell averages 80 observations. I also use a restricted model with average bids α_{ig} fixed over time, the same model used to create Figure 4. The hypothesis test of interest is whether $\alpha_{igm} = \alpha_{ig}$ for all m.

$$b_{itl} = \alpha_{igm} + \beta_i distance_{itl} + \varepsilon_{itl} \qquad b^*_{itl} = \begin{cases} b_{itl} & \text{if} \quad b_{itl} \geq R_h \\ R_h & \text{if} \quad \text{Otheriwse} \end{cases} \qquad \varepsilon_{itl} \sim N(0, \sigma_{ih})$$

This hypothesis test may be underpowered, since food banks' needs also vary within a month and this may be of the same scale as the across month variation. The test will be overpowered if variation in factors other than food banks' needs is mistaken for variation in needs. For example, if the quality of food varies unobservably over time, this may cause systematic variation in bidding behaviour that should not be attributed to variation in needs. To account for this I estimate a second restricted specification with food bank specific month fixed effects.[38]

Figure 15 plots the likelihood ratio test statistic across food banks. The dotted lines gives the χ^2 critical values for tests at the 5% significance level. The red points give the baseline specification, while the blue points give the month fixed effects specification. I reject the null hypothesis at 5% significance level, for 96% of food banks in my baseline specification, and 70% of food banks for the month fixed effects specification. Therefore I have strong evidence that food banks' bidding behaviour, and hence their needs, vary significantly over time.

C Non-parametric Identification

I now prove Proposition 1, that the model is non-parametrically point identified. I begin by discussing in slightly more detail the necessary assumptions underlying the argument, before introducing some additional terminology required for the proof. Then, in C.1 I present the two step proof. Altmann (2022) showed that conditional on having identified i) the conditional equilibrium distribution of bids $f(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t)$ and

[38] These fixed effects capture variation in bidding that is common across food types. Under this specification a rejection of the null is evidence of systematic variation over time in bidding behaviour *on specific types of food*. This specification is almost certainly underpowered. If food banks need more food of all types in certain months the fixed effects also soak up this variation.

Figure 15: Heterogeneity Across Time

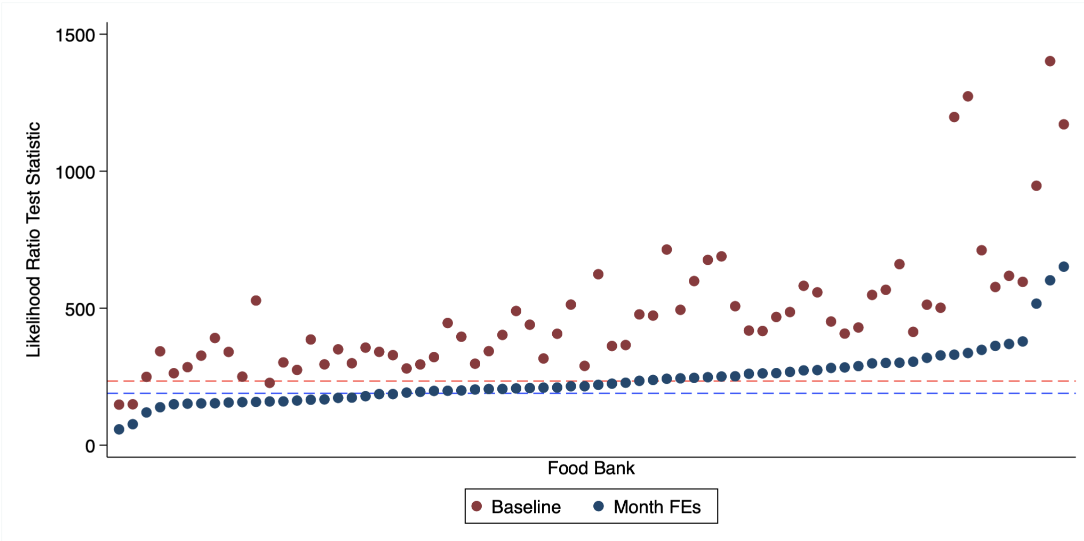

Note: This figure plots likelihood ratio test statistics for the hypothesis test that average bids for each type of food are constant over time, against the alternate hypothesis that bids vary by month. The estimated model controls for censoring, distance, and lot composition. The blue results also include month fixed effects. Under this null hypothesis the test statistic takes a χ^2 distribution with 200 (red) or 160 (blue) degrees of freedom. Critical values for tests at the 5% significance levels are plotted as horizontal lines.

ii) the transition process $f(\mathbf{s}_{t+1}|\mathbf{w}_t, \mathbf{s}_t)$, that variation in the state $(\mathbf{s}_{0t}, \mathbf{s}_t)$ is sufficient for non-parametric identification of both π, and F^v. Therefore, it suffices to prove that both $f(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t)$ and $f(\mathbf{s}_{t+1}|\mathbf{w}_t, \mathbf{s}_t)$ are identified.

The argument builds on the key result in Hu and Shum (2012), closely following their spectral decomposition techniques for linear operators. The key distinction between our arguments consists of the introduction of an additional intermediate variable (winnings) that acts as an observed shifter of the unobserved process. This requires that conditional on $\mathbf{s}_t$ and $\mathbf{s}_{0t}$, that both bids and winnings are independent of past winnings. This reduces the assumptions required for identification.[39] It also yields clear intuition behind identification of the latent state: Variation in this observed shifter of the unobserved state pins down the relationship between bids and the unobserved state. Then variation in bids over time, holding constant the observed shifter, enables identification of the state transition process. The final distinction

[39]Hu and Shum (2012)'s version of assumption 5 requires that conditional on observations at t *and* $t-1$, variation at $t-2$ is sufficient to pin down functions of observables at $t+1$. Meanwhile, assumption 2 i) means we do not have to normalise the latent state up to monotone transformation.

is that this framework encapsulates multivariate latent states, as the signals of the latent state (bids and winnings) are both also multivariate. Lastly, note that this identification argument works in the case of multiple strategic agents, not only the single agent or 'large market' setting as considered in this paper. $\mathbf{b}_{it+1}$ is allowed to depend on $\mathbf{s}_{jt+1}$, so that we essentially see how i's behaviour varies after j wins a lot, as their stocks have now increased in an observable way.

Like all assumptions about completeness, Assumption 5 is strong and requires an unrealistic amount of variation in $\mathbf{w}_{t-1}$ and $\mathbf{b}_{t+1}$. Nonetheless it is still important for understanding the conditions for non-parametric identification, demonstrating that identification is not driven by (potentially stronger) parametric assumptions.[40]

C.0.1 Linear Operators and Spectral Decomposition

A 'linear operator' $L_{x,y}$ is a map from the $L^{|y|}$ space of functions of y to the $L^{|x|}$ space of functions of x, such that for function $g : \mathbb{R}^{|x|} \to \mathbb{R}^{|y|}$: $(L_{x,y}g)x = \int f(x,y)g(y)dy$. Likewise, define the diagonal operator $D_{x,y}$ as follows: $(D_{x,y}g)x = f(x,y)g(y)$.

Linear operators are close to infinite dimensional counterparts to matrices, with similar properties. Injectivity and surjectivity of these mappings is inherently tied to completeness. $E[g(\mathbf{y})|\mathbf{x}] = 0$ for all $\mathbf{x}$ implies $g(\mathbf{y}) = 0$ for all $\mathbf{y}$ if and only if the mapping $L_{x,y}$ is injective and $f(\mathbf{x}) > 0$, so the left inverse of $L_{x,y}$ exists. That is, variation in $\mathbf{x}$ yields enough conditional variation in $\mathbf{y}$ to allow us to pin down functions of $\mathbf{y}$. So assumption 2 iii) and 5 i) equivalently assume the existence of left inverses of the operators $L_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$, $L_{b_t|\mathbf{s}_{0t},\mathbf{s}_t}$, while 5 part ii) ensures the right invertibility of $L_{b_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}}$. See Hu and Schennach (2008) for additional discussion.

The proof below relies on taking a spectral decomposition of certain linear operators, essentially the linear operator equivalent of eigenvalue decomposition. Just as in Hu and Shum (2012), I require the eigenvalues of this decomposition are unique. This

[40] The assumption nests two additional implicit assumptions: That reservation prices do not bind, and that $\mathbf{z}_t^g$ has full rank. Binding reservation prices mean first order conditions do not hold with equality. However, as discussed in Altmann (2022), reservation prices are not a first-order issue, not substantially altering the identification problem. In the way that a censored regression model, which requires a Tobit or MAD specification, does not substantially alter the regression identification problem. The key intuition garnered from this simplified approach extends to the case with reservation prices. Meanwhile the rank condition on $\mathbf{z}_t^g$, the size and composition of lots auctioned each day, simply ensures that at least one of each food type is auctioned each period, so that bids in every period are informative of every dimension of stocks. As it turns out, identification actually only requires this rank condition for 3 consecutive periods, which holds.

requires that for any, for any pair $(\bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t})$ satisfying assumption 5 part *ii*) and any $\mathbf{s}_t$, that the density $f(\bar{\mathbf{w}}_t|\bar{\mathbf{s}}_{0t}, \mathbf{s}_t) = \int f(\bar{\mathbf{w}}_t|\bar{\mathbf{s}}_{0t}, \mathbf{b}_t) f(\mathbf{b}_t|\bar{\mathbf{s}}_{0t}, \mathbf{s}_t) d\mathbf{b}_t$ is strictly positive and bounded above. This is already guaranteed by the model setup. I also require that for any tuple $(\mathbf{w}_t, \mathbf{s}_{0t}, \bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t})$ and for any $\bar{\mathbf{s}_t} \neq \mathbf{s}_t$ satisfying this assumption, that $\Lambda(\mathbf{w}_t, \mathbf{s}_{0t}, \bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}, \mathbf{s}_t) \neq \Lambda(\mathbf{w}_t, \mathbf{s}_{0t}, \bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}, \bar{\mathbf{s}}_t)$, where:

$$\Lambda(\mathbf{w}_t, \mathbf{s}_{0t}, \bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}, \mathbf{s}_t) = \frac{f(\mathbf{w}_t|\mathbf{s}_{0t}, \mathbf{s}_t) f(\bar{\mathbf{w}}_t|\bar{\mathbf{s}}_{0t}, \mathbf{s}_t)}{f(\bar{\mathbf{w}}_t|\mathbf{s}_{0t}, \mathbf{s}_t) f(\mathbf{w}_t|\bar{\mathbf{s}}_{0t}, \mathbf{s}_t)}.$$

That is, variation in the unobserved stocks yields variation in the relative conditional win probabilities (integrating over bids) for these pairs of winnings and available lots. In practice, this result follows from Assumption 4 part *iii*), which ensures that (conditional on $\boldsymbol{\upsilon}$), bids are monotonic in stocks. However proof of this proposition is tedious, so this should instead be considered an auxiliary assumption.

C.1 Proof of Proposition 1

The proof follows the argument of Hu and Shum (2012). First I show that the conditional density $f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t)$ is completely determined by the observed joint density $f(\mathbf{b}_{t+1}, \mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1})$. Then I show that $f(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t)$ and $f(\mathbf{s}_t|\mathbf{w}_{t-1}, \mathbf{s}_{t-1})$ are point identified given $f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t)$ and other observed joint densities.

C.1.1 Identification of $f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t)$ by spectral decomposition

Lemma C.1. $f(\mathbf{b}_{t+1}, \mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1})$ *completely determines* $f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t)$.

The argument broadly follows the proof of Lemma 3 from Hu and Shum (2012), using $\mathbf{b}_{t+1}$ in place of their V_{t+1}, $\mathbf{w}_{t-1}$ in place of V_{t-2}, $\mathbf{w}_t$ for w_t, and $\mathbf{s}_{0t}$ for w_{t-1}. Consequently I do not elaborate the proof in excessive detail.

Proof:

1. From our exclusion restrictions we can write: $f(\mathbf{b}_{t+1}, \mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1}) = \int_{\mathbf{s}_t} f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t) f(\mathbf{s}_{0t+1}, \mathbf{w}_t|\mathbf{s}_t, \mathbf{s}_{0t}) f(\mathbf{s}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1}) d\mathbf{s}_t$
2. In Operator notation, for fixed $(\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_{0t})$, this can be written as:

$$L_{\mathbf{b}_{t+1}, \mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1}} = L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t} D_{\mathbf{s}_{0t+1}, \mathbf{w}_t|\mathbf{s}_{0t}, \mathbf{s}_t} L_{\mathbf{s}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1}} \quad (5)$$

3. Assumptions 5 part *i*) and 2 part *iii*) ensure that for any $\mathbf{s}_{0t+1}, \mathbf{w}_t$, $\mathbf{b}_{t+1}$ is also complete for $\mathbf{s}_t$, so that the left inverse of $L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t}$ exists.

Likewise, $f(\mathbf{s}_{0t+1}, \mathbf{w}_t|\mathbf{s}_{0t}, \mathbf{s}_t) > 0$, which follows from the discussion in C.0.1 above, ensures $D_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t}$ is invertible. Therefore, we can write:

$$L_{\mathbf{s}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} = D^{-1}_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} L_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} \tag{6}$$

4. As proven in Hu and Schennach (2008), assumption 5 *ii*) ensures that for any $\mathbf{s}_{0t+1}$ there exists a neighbourhood near this fixed $\mathbf{w}_t, \mathbf{s}_{0t}$, $(\bar{\mathbb{W}}, \bar{\mathbb{S}}_0)$, such that for all $(\bar{\mathbf{s}}_{0t}, \bar{\mathbf{w}}_t) \in (\bar{\mathbb{W}}, \bar{\mathbb{S}}_0)$, $L_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\bar{\mathbf{w}}_t,\bar{\mathbf{s}}_{0t},\mathbf{w}_{t-1}}$ is right invertible. Therefore, using equations 6 and 5 we can write:

$$\begin{aligned} L_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} & L^{-1}_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\bar{\mathbf{w}}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} \\ & \times L_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\bar{\mathbf{w}}_t,\bar{\mathbf{s}}_{0t},\mathbf{w}_{t-1}} L^{-1}_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\bar{\mathbf{s}}_{0t},\mathbf{w}_{t-1}} \\ & = L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} D_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t} D^{-1}_{\mathbf{s}_{0t+1},\bar{\mathbf{w}}_t|\mathbf{s}_{0t},\mathbf{s}_t} \\ & \times D_{\mathbf{s}_{0t+1},\bar{\mathbf{w}}_t|\bar{\mathbf{s}}_{0t},\mathbf{s}_t} D^{-1}_{\mathbf{s}_{0t+1},\mathbf{w}_t|\bar{\mathbf{s}}_{0t},\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} \\ & = L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} D_{\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\bar{\mathbf{w}}_t,\bar{\mathbf{s}}_{0t},\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t}. \end{aligned} \tag{7}$$

Where the diagonal operator:[41]
$(D_{\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\bar{\mathbf{w}}_t,\bar{\mathbf{s}}_{0t},\mathbf{s}_t} h)(\mathbf{s}_t) = \Lambda(\mathbf{w}_t, \mathbf{s}_{0t}, \bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}, \mathbf{s}_t)h(\mathbf{s}_t)$.

5. Equation 7 states the left hand side has an eigenvalue-eigenfunction decomposition given by the right hand side. The discussion in C.0.1 above ensures the eigenvalues are bounded, so these operators on are similarly bounded. Therefore we can apply Theorem XV.4.3.5 from Dunford and Schwartz (1971), ensuring uniqueness of the decomposition.[42] Our assumptions on Λ varying with $\mathbf{s}_t$ ensures that eigenvalues for different values of $\mathbf{s}_t$ are distinct (for some $(\mathbf{w}_t, \mathbf{s}_{0t}) \neq (\bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}) \in (\bar{\mathbb{W}}, \bar{\mathbb{S}}_0)$).

6. Eigenvalues and eigenfunctions are unique up to invertible transformations of $\mathbf{s}_t$. Let $g : \mathbb{R}^{|s|} \to \mathbb{R}^{|s|}$ denote any invertible function of stocks, so that $\mathbf{s} = g(\tilde{\mathbf{s}})$. Consider the set of g that satisfy assumption 2 part *i*).

[41] By assumption 1 the $f(\mathbf{s}_{0t+1}|\mathbf{s}_{0t}, \mathbf{w}_t, \mathbf{s}_t) = f(\mathbf{s}_{0t+1}|\mathbf{s}_{0t})$ terms cancel out. We can easily allow $f(\mathbf{s}_{0t+1}|\mathbf{s}_{0t}, \mathbf{w}_t, \mathbf{s}_t) = f(\mathbf{s}_{0t+1}|\mathbf{s}_{0t}, \mathbf{w}_t)$ if this (observed) density is > 0 everywhere.

[42] Linear operators are bounded by their largest eignevalue. This theorem then ensures uniqueness of the eigenfunctions up to a scalar multiple. The requirement that the eigenfunctions are proper densities that integrate to one pins down their scale.

This requires that for any vector $\mathbf{x}$:

$$\begin{aligned} f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t + \mathbf{s}_t) =& f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, (\mathbf{w}_t - \mathbf{x}) + (\mathbf{s}_t + \mathbf{x})) \\ =& f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, (\mathbf{w}_t - \mathbf{x}) + g(\tilde{\mathbf{s}}_t + \mathbf{x})) \end{aligned}$$

This is the crux of the perfect substitutes assumption. Because g is invertible, the only function that satisfies perfect substitutes is $g(\mathbf{s}) = \mathbf{s} + \boldsymbol{\mu}$, so that stocks are identified up to location. Then, imposing that stocks have long run mean of zero ensures we can write $0 = E[\mathbf{s}] = E[g(\mathbf{s})] = E[\mathbf{s} + \boldsymbol{\mu}] = E[\mathbf{s}] + \boldsymbol{\mu}$ which holds only for $\boldsymbol{\mu} = 0$.

7. Therefore, the density $f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t)$ is point identified for $\mathbf{w}_t$ satisfying assumption 5 *ii*). Assumption 2 *i*) implies the density can be written as $f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t + \mathbf{s}_t)$, ensuring it is identified for all $\mathbf{w}_t$.

□

C.1.2 Identification of $f(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t)$ and $f(\mathbf{s}_t|\mathbf{w}_{t-1}, \mathbf{s}_{t-1})$ given $f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t)$

Lemma C.2. *If $f_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t}$ identified, then so is $f(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t)$ and $f_{\mathbf{s}_t|\mathbf{w}_{t-1},\mathbf{s}_{t-1}}$.*

Proof:

1. $f(\mathbf{b}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1}) = \int_{\mathbf{s}_t} f(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t) f(\mathbf{s}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1}) d\mathbf{s}_t$. In operator notation, $L_{\mathbf{b}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} = L_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t} L_{\mathbf{s}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}}$.
2. Substituting in $L_{\mathbf{s}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}}$ from equation 6:

$$L_{\mathbf{b}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} = L_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t} D^{-1}_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} L_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}}.$$

3. Taking sequential right inverses yields:

$$L_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t} = L_{\mathbf{b}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} L^{-1}_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} D_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t}.$$

   Therefore $f(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t)$ is identified.
4. $f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t) = \int_{\mathbf{s}_{t+1}} f(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{s}_{t+1}) f(\mathbf{s}_{t+1}|\mathbf{w}_t, \mathbf{s}_t) d\mathbf{s}_{t+1}$. In operator notation, $L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} = L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{s}_{t+1}} L_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$.
5. From assumption 5 *i*) the left inverse of $L_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t}$ exists, and so $L_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} = L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{s}_{t+1}} L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t}$, therefore $f(\mathbf{s}_{t+1}|\mathbf{w}_t, \mathbf{s}_t)$ is identified also.

□

D Inverse Bid System

In this Appendix I demonstrate that, in addition to the transition equation given in Assumption 2, a food bank's optimisation problem yields the Observation and Censoring equations given in text. For the most part, I simply the results presented in Altmann (2022) for the quadratic parametrisation of κ.

Parameterising κ, and given entry decision $\mathbf{d}_i^*$, the maximisation problem is given by: $\max_{\mathbf{b}} \Big\{ \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(\upsilon_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{s}_i^{ah} - \mathbf{s}_i^{agT} \Psi \mathbf{s}_i^{ag}] \qquad s.t.\ b_l \geq R_l \Big\}$. The maximand can be simplified, allowing us to write the lagrangian as:

$$L(\mathbf{b}|\mathbf{d}^*, \boldsymbol{\upsilon}, \mathbf{s}) = \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(\upsilon_l - b_l + \Phi \mathbf{z}_l^h - \mathbf{z}_l^{gT} \Psi [\mathbf{z}_l^g + 2\mathbf{s}_i^g + \sum_{m \neq l} \Gamma_m(b_m, d_m^*; \mathbf{s}) \mathbf{z}_m^g]) + \Phi \mathbf{s}_i^h - \mathbf{s}_i^{gT} \Psi \mathbf{s}_i^{gT} - \sum_l \Lambda_l (R_l - b_l).$$

Λ_l are lagrangian multipliers.[43] Taking First Order Conditions and rearranging yields:

$$b_l^* + \frac{\Gamma_l(b_l^*, d_l^*; \mathbf{s})}{\nabla_b \Gamma_l(b_l^*, d_l^*; \mathbf{s})} - \Lambda_l^* = \Phi \mathbf{z}_l^h - \mathbf{z}_l^{gT} \Psi (\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2 \sum_{m \neq l} \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_l^g) + \upsilon_l = y_l.$$

Λ_l^*, and hence y_l, is unobserved. Let $y_l^* = b_l^* + \frac{\Gamma_l(b_l^*, d_l^*; \mathbf{s})}{\nabla_b \Gamma_l(b_l^*, d_l^*; \mathbf{s})}$ be what we observe. When $b_l^* > R_l$, we infer $\Lambda_l^* = 0$, so that $y_l^* = b_l^* + \frac{\Gamma_l(b_l^*, d_l^*; \mathbf{s})}{\nabla_b \Gamma_l(b_l^*, d_l^*; \mathbf{s})} = y_l$.

Because of the non-zero probability of tieing at the reserve price Γ_l is non-differentiable at R_l. However, because they prefer to bid the reserve price (and risk tieing), rather than bidding 1 share above reserve, assuming exogenous tie breaking Altmann (2022) demonstrates that, for l such that $b_l^* = R_l$, we have:

$$y_l \leq R_l + \frac{\Gamma_l(R_l + 1, d_l^*; \mathbf{s})}{\Gamma_l(R_l + 1, d_l^*; \mathbf{s}) - \Gamma_l(R_l, d_l^*; \mathbf{s})} \qquad (= y_l^*).$$

They also demonstrate that because the bidder is observed entering, they must prefer to enter and bid the reserve, than not enter at all. This yields the inequality:

$$(y_l =) \qquad \Phi \mathbf{z}_l^h - \mathbf{z}_l^{gT} \Psi_i (\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2 \sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_m^g) + \upsilon_l \geq R_l.$$

[43] Derivation of the simplification just exploits the quadraticness of κ, and involves tedious algebra. It employs $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s}) \mathbf{s}_i^{gT} \Psi \mathbf{s}_i^g = \mathbf{s}_i^{gT} \Psi \mathbf{s}_i^g$ and $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s}) \mathbf{s}_i^a = \mathbf{s}_i + \sum_l \Gamma_l(b_l, d_l; \mathbf{s}) \mathbf{z}_l$. Note that Assumption 4 ensures $\nabla_{b_l} \Gamma_l > 0$ for $b_l > R_l$.

This implies that, for l such that $b_l^* = R_l$, $R_l \leq y_l \leq y*_l$. Finally, for l such that $d_l^* = 0$, we can reverse this inequality (they prefer not to enter than to enter at the reserve), and hence $y_l \leq R_l$.

E Proof of Proposition 2.

In this Appendix I prove Proposition 2. However, I prove a modified version of the proposition that accounts for binding reserve prices, excluded from the text for ease of exposition. The modified proposition is given as:

Proposition 2′. *The ex-ante Value Function can be expressed as:*

$$E[W(\boldsymbol{\upsilon}_{it}, \mathbf{s}_i, \mathbf{s}_0)|\mathbf{s}_i, \mathbf{s}_0] = \frac{E[q_t(\mathbf{s}_i^g)E[W(\boldsymbol{\upsilon}_{it}, \mathbf{s}_i, \mathbf{s}_0)|\mathbf{b}_{it}, \mathbf{d}_{it}, \mathbf{s}_i, \mathbf{s}_0]|\mathbf{s}_0]}{E[q_t(\mathbf{s}_i^g)|\mathbf{s}_0]} \tag{8}$$

Where $q_t(\mathbf{s}_i^g)$ gives the posterior probability that $\mathbf{s}_{it}^g = \mathbf{s}_i^g$ and

$$\begin{aligned} E[W(\boldsymbol{\upsilon}_{it}, \mathbf{s}_i, \mathbf{s}_0)|\mathbf{b}_{it}, \mathbf{d}_{it}, \mathbf{s}_i, \mathbf{s}_0] &= -\mathbf{s}_i^{gT}\Psi_i\mathbf{s}_i^g \\ &+ \sum_l \begin{cases} \mathbb{I}[b_l > R_l]\left(\lambda\frac{\Gamma_l(b_l,d_l)^2}{\nabla_b\Gamma_l(b_l,d_l)} + \sum_{m\neq l}\Gamma_l(b_l,d_l)\mathbf{z}_l^{gT}\Psi_i\mathbf{z}_m^g\Gamma_m(b_m,d_m)\right) \\ \mathbb{I}[b_l = R_l]\Gamma_l(R_l,1)\begin{pmatrix} E[\upsilon_l|b_l = R_l, \mathbf{b}_{-l}] - \lambda R_l + \Phi\mathbf{z}_l^h \\ -\mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m,d_m)\mathbf{z}_m^g] \end{pmatrix} \end{cases} \end{aligned} \tag{9}$$

The proof consists of two parts. First, I prove equality 8, extending results from Arcidiacono and Miller (2011) to the continuous choice case. Then I prove equality 9, applying a result from Altmann (2022) for quadratic payoffs. The conditional expectation $E[\upsilon_l|b_l = R_l, \mathbf{b}_{-l}]$ is just a truncated expectation, and given the gaussian assumptions has a simple analytic expression using the bounds derived in D.

E.1 Proof of equality 8

To simplify notation, I drop i subscripts and dependence on the observe state $\mathbf{s}^0$. This can be trivially introduced by multiplying objects by $\mathbb{I}[\mathbf{s}_t^0 = \mathbf{s}^0]$. I also drop dependence on the discrete actions $\mathbf{d}$, also trivially introduced by multiplying objects by $\mathbb{I}[\mathbf{d}_t = \mathbf{d}]$ and summing over possible actions, as in the discrete choice case.[44]

[44]The proof uses the Dirac Delta function, defined for continuous random variable $\mathbf{B}$ with density $f_{\mathbf{B}}$ such that $E_{\mathbf{B}}[\delta(\mathbf{B} - \mathbf{b})] = f_{\mathbf{B}}(\mathbf{b})$ and with the property that $\int_{\mathbf{B}} \delta(\mathbf{B} - \mathbf{b})d\mathbf{B} = 1$. I also use that $\delta((\mathbf{B}, \mathbf{S}) - (\mathbf{b}, \mathbf{s})) = \delta(\mathbf{B} - \mathbf{b})\delta(\mathbf{S} - \mathbf{s})$.

Proof: 1. First, I prove that $f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s}) = \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})|q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$:

$$
\begin{aligned}
f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s}) =& \frac{f_{\mathbf{b}_t,\mathbf{s}_t}(\mathbf{b},\mathbf{s})}{f_{\mathbf{s}_t}(\mathbf{s})} && \text{Bayes' rule}\\
=&= \frac{E_{\mathbf{b}_t,\mathbf{s}_t}[\delta(\mathbf{b}_t-\mathbf{b})\delta(\mathbf{s}_t-\mathbf{s})]}{E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})]} && \text{Definition of Dirac } \delta \text{ Function}\\
=& \frac{E_{\mathbb{O}_T}[E_{\mathbf{b}_t,\mathbf{s}_t}[\delta(\mathbf{b}_t-\mathbf{b})\delta(\mathbf{s}_t-\mathbf{s})|\mathbb{O}_T]]}{E_{\mathbb{O}_T}[E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})|\mathbb{O}_T]]} && \text{Iterated Expectations}\\
=& \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})|\mathbb{O}_T]]}{E_{\mathbb{O}_T}[E_{\mathbf{s}_t}[\delta(\mathbf{s}_t-\mathbf{s})|\mathbb{O}_T]]} && \text{as } \mathbf{b}_t \text{ is part of } \mathbb{O}_T\\
=& \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} && \text{Definition of } q
\end{aligned}
$$

2. Apply iterated expectations for $E_{\boldsymbol{\upsilon}_t}[W(\boldsymbol{\upsilon}_t,\mathbf{s})|\mathbf{s}] = E_{\mathbf{b}_t}[E_{\boldsymbol{\upsilon}_t}[W(\boldsymbol{\upsilon}_t,\mathbf{s})|\mathbf{b}_t,\mathbf{s}]|\mathbf{s}]$. For notational convenience, let $\tilde{W}(\mathbf{b}_t,\mathbf{s}) = E_{\boldsymbol{\upsilon}_t}[W(\boldsymbol{\upsilon}_t,\mathbf{s})|\mathbf{b}_t,\mathbf{s}]$.

3. Applying the result from step 1.:

$$E_{\mathbf{b}_t}[\tilde{W}(\mathbf{b}_t,\mathbf{s})|\mathbf{s}] = \int_{\mathbf{b}} \tilde{W}(\mathbf{b},\mathbf{s}) f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s}) d\mathbf{b} = \int_{\mathbf{b}} \tilde{W}(\mathbf{b},\mathbf{s}) \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} d\mathbf{b}$$

4. The denominator is not a function of the random variable $\mathbf{b}$, so pull it from the integral. Then, move $\tilde{W}(\mathbf{b},\mathbf{s})$ into the expectation for:

$$= \frac{\int_{\mathbf{b}} \tilde{W}(\mathbf{b},\mathbf{s}) E_{\mathbb{O}_T}[\delta(\mathbf{b}_t-\mathbf{b})q_t(\mathbf{s})] d\mathbf{b}}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} = \frac{\int_{\mathbf{b}} E_{\mathbb{O}_T}[\tilde{W}(\mathbf{b},\mathbf{s})\delta(\mathbf{b}_t-\mathbf{b})q_t(\mathbf{s})] d\mathbf{b}}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$$

5. From the definition of the delta function the expectation equals zero for $\mathbf{b} \neq \mathbf{b}_t$, so that I can replace $\tilde{W}(\mathbf{b},\mathbf{s})$ with $\tilde{W}(\mathbf{b}_t,\mathbf{s})$. Then, swap the order of integration, moving the integral into the expectation for:

$$= \frac{\int_{\mathbf{b}} E_{\mathbb{O}_T}[\tilde{W}(\mathbf{b}_t,\mathbf{s})\delta(\mathbf{b}_t-\mathbf{b})q_t(\mathbf{s})] d\mathbf{b}}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} = \frac{E_{\mathbb{O}_T}[\int_{\mathbf{b}} \tilde{W}(\mathbf{b}_t,\mathbf{s})\delta(\mathbf{b}_t-\mathbf{b})q_t(\mathbf{s}) d\mathbf{b}]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$$

6. Within the expectation, $\mathbf{b}_t$ and $\mathbf{s}$ are constant, so pull $\tilde{W}(\mathbf{b}_t,\mathbf{s})q_t(\mathbf{s})$ out of the integral, before applying the definition of the delta function:

$$= \frac{E_{\mathbb{O}_T}[\int_{\mathbf{b}} \delta(\mathbf{b}_t-\mathbf{b}) d\mathbf{b}\, \tilde{W}(\mathbf{b}_t,\mathbf{s})q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} = \frac{E_{\mathbb{O}_T}[\tilde{W}(\mathbf{b}_t,\mathbf{s})q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$$

$\square$

E.2 Proof of equality 9

I now prove that the conditional expectation $E_{\boldsymbol{\upsilon}_i}[W(\boldsymbol{\upsilon}_i, \mathbf{s}_i, \mathbf{s}_0)|\mathbf{b}_{it}, \mathbf{d}_{it}, \mathbf{s}_i, \mathbf{s}_0]$ has the convenient form above. If reserve prices do not bind then, as in Altmann (2022), $\mathbf{b}$ directly pins down $\boldsymbol{\upsilon}$. As shown in appendix D, with binding reserve prices $\boldsymbol{\upsilon}$ is only determined up to a convex set. Therefore, we must take an expectation over this set. I use $\mathbf{b}^*$ and $\mathbf{d}^*$ to denote optimised bids / entry, i.e. $\mathbf{b}_i^* = \mathbf{b}(\boldsymbol{\upsilon}_i, \mathbf{s}; \kappa)$ is a function of $\boldsymbol{\upsilon}$. Trivially, $E_{\boldsymbol{\upsilon}_i}[h(\mathbf{b}^*)|\mathbf{b}_{it}] = h(\mathbf{b}_{it})$ for any function h.

Proof:

1. $\mathbb{I}[b_l^* > R_l] + \mathbb{I}[b_l^* = R_l] + \mathbb{I}[d_l^* = 0] = 1$, so we can write the (parametrised) value function as: $W(\boldsymbol{\upsilon}, \mathbf{s}) - \Phi \mathbf{s}^h + \mathbf{s}^{gT}\Psi\mathbf{s}^{gT} =$

$$\sum_l \begin{cases} \mathbb{I}[b_l^* > R_l]\Gamma_l(b_l^*, d_l^*)(\upsilon_l - \lambda b_l + \Phi \mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_m^g]) \\ +\mathbb{I}[b_l^* = R_l]\Gamma_l(b_l^*, d_l^*)(\upsilon_l - \lambda b_l + \Phi \mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_m^g]) \\ +\mathbb{I}[d_l^* = 0]\Gamma_l(b_l^*, d_l^*)(\upsilon_l - \lambda b_l + \Phi \mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_m^g]) \end{cases} \tag{10}$$

2. By definition $\mathbb{I}[d_l^* = 0]\Gamma_l(b_l^*, d_l^*) = 0$, so the final row equals zero.
3. Next, $\mathbb{I}[b_l^* = R_l]\Gamma_l(b_l^*, d_l^*) = \Gamma_l(R_l, 1)$, so second row of equation 10 equals: $\mathbb{I}[b_l^* = R_l]\Gamma_l(R_l, 1)(\upsilon_l - \lambda R_l + \Phi \mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_m^g])$
4. Reserve prices do not bind for the first row, so the FOCs hold with equality, meaning we can substitute in the inverse bid system $\xi_l(\mathbf{b}, \mathbf{d})$ derived in Appendix D in place of υ_l, giving: $\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l \Gamma_l(b_l^*)} + \Gamma_l(b_l^*)\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)$
5. Finally, take the expectation of $W(\boldsymbol{\upsilon}, \mathbf{s})$ over $\boldsymbol{\upsilon}$, conditional on $\mathbf{b}_{it}, \mathbf{d}_{it}, \mathbf{s}_i, \mathbf{s}_0$. The first row only depends on $\boldsymbol{\upsilon}$ through $\mathbf{b}^*, \mathbf{d}^*$, so we essentially just get a change of variables. The second row is affine in υ_l, and otherwise only depends on $\boldsymbol{\upsilon}$ through $\mathbf{b}^*, \mathbf{d}^*$. Therefore, we take a simple conditional expectation of υ_l given $b_{ilt} = R_l, \mathbf{b}_{i-lt}$. This yields equality 9.

$\square$

F Additional Estimation Details

In this Appendix I detail the estimation procedure outlined in Section 5, including specification of priors and the sampling algorithm.

F.1 Step 1.

First, I estimate food banks beliefs about the probability of winning lot l given bid b_{itl}. While I assume there is zero probability of ties above the reservation price, I allow for ties at the reservation price. This occurs occurs in 0.02% of auctions. However, 15% of winning bids are at the reserve price, so food banks must consider the non-zero probability of tieing if they bid exactly the reserve. Food banks recognise this and regularly bid just above -2000, which causes high density of winning bids 1 to 50 shares above the reserve, so must also be accounted for in the model.

The bidder wins lot l given bid b_{ilt} if $b_{ilt} > \bar{b}_{lt}$ (the highest rival bid). If $b_{ilt} = \bar{b}_{lt}$ they win with probability 0.5. Like i's bids, $\bar{b}_{lt}$ is censored both at R_l (when the maximum rival bid equals the reservation price) and below it (when no rivals place bids). I introduce the latent random variable $\bar{b}^*_{lt}$, with cdf $G_l(b^*|\mathbf{s}_{0t})$, such that:

$$\bar{b}_{lt} = \begin{cases} \emptyset & \text{if } \bar{b}^*_{lt} \leq \underline{R}_l & \leftarrow \text{No rivals enter} \\ R_l & \text{if } \bar{b}^*_{lt} \in [\bar{R}_l, \underline{R}_l) & \leftarrow \text{Rival bids } R_l \\ R_l + \epsilon_{lt} & \text{if } \bar{b}^*_{lt} \in [R_l, \bar{R}_l) & \leftarrow \text{Rival bids just above } R_l \\ \bar{b}^*_{lt} & \text{if } \bar{b}^*_{lt} > R_{lt} & \leftarrow \text{Rival bids } > R_l \end{cases}$$

$(\bar{R}_l, \underline{R}_l)$ are category specific cutoffs to be estimated, similar to the cutoffs estimated in ordered logit models. This latent variable structure states that when $\bar{b}^*_{lt} \leq \underline{R}_l$ (or $\bar{b}^*_{lt} \in [\bar{R}_l, \underline{R}_l)$), i would win (or tie) if they bid reserve. Meanwhile, if $\bar{b}^*_{lt} \in [R_l, \bar{R}_l)$ then the observed winning bid is actually just above the reservation price, where $\epsilon_{lt} \sim \mathit{exponential}(\alpha)$ and α is a parameter to be estimated. So, a competing food bank must take into account the excess mass just above the reservation price.[45]

[45]This modelling approach is unusual, but enables the model to rationalise the excess mass of winning bids at, and just above, reserve. I assume food banks do not internalise the probability of tieing at just one share above the reservation price (and likewise two, three, etc). The cutoffs are identified by the excess mass of winning bids at/just above the reservation price.

Given the distribution of $\bar{b}^*_{lt}$, and implied distribution of $\bar{b}_{lt}$, i's beliefs are:

$$P(i \text{ wins } l|b_{ilt};\mathbf{s}_{0t}) = \Gamma_l(b_{ilt}|\mathbf{s}_{0t}) = \begin{cases} G_l(b_{ilt}|\mathbf{s}_{0t}) - f(b_{ilt}|\mathbf{s}_{0t}) & \text{if } b_{ilt} > R_{lt} \\ \frac{1}{2}G_l(\underline{R}^c|\mathbf{s}_{0t}) + \frac{1}{2}G_l(\bar{R}^c|\mathbf{s}_{0t}) & \text{if } b_{ilt} = R_{lt} \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

Where $f(b_{ilt}|\mathbf{s}_{0t}) = [G_l(R_l|\mathbf{s}_{0t}) - G_l(\bar{R}^c|\mathbf{s}_{0t})]e^{-\alpha b_{ilt}}$ captures the probability that i loses out to a food bank bidding just above the reservation price. This probability features a discontinuity at the reserve price.

F.1.1 Parameterisation and Computation

I normalise winning bids by the reservation price, estimating the distribution of $\bar{b}^*_{lt} - R_l$. Lots contain up to four distinct categories, subcategories and storage types, reflected in the shape, scale and location parameters. The shape parameters ξ are category specific for categories with at least 500 loads. The scale parameters ζ are all category specific. I include additional scale fixed effects if the lot has been unsuccessfully auctioned previously, and for lots with subcategories listed as “mixed”. The location shifter ν includes the common state variables, subcategory fixed effects, and dummies for several observables such as whether the lot is sold by a food bank. The threshold cutoffs $\bar{R}_l$ and $\underline{R}_l$ vary across categories for which at least 100 lots were won at the reservation price. The remaining categories are grouped together. The exponential parameter α is constrained positive.

I drop the first 60 days to construct the previous 30 days' supply. I maximise the posterior likelihood, then sampled from the posterior distribution using Metropolis Hastings. I use the MLE inverse hessian for the proposal variance which I adaptive tune as per Atchadé and Rosenthal (2005).

F.2 Step 2.

Within the food banks I use for estimation I further split food banks into Type 1, the 28 food banks who win more than 200 loads over my sample (65% of total consumption), and 62 Type 2 who win more than 50 loads (29%). Parameters that are common across food banks, including hierarchical parameters, are estimated on data from Type 1 food banks only, for whom I have sufficient identifying variation.

F.2.1 Parametrisation

The pseudo-payoff function is parameterised as $\kappa(\mathbf{s}_i) = \Phi_i \mathbf{s}_i^h - \mathbf{s}_i^{gT}\Psi_i\mathbf{s}_i$. I do not allow all elements of Φ_i to vary for each food bank, as this introduces too many parameters to estimate. Instead I exploit that 152 subcategories are nested within the 5 different 'usage' types u, so that $\Phi_i \mathbf{s}_i^h = \sum_h \sum_u \tilde{\Phi}_{iu}\bar{\Phi}_h k_{hu} s_{ih}^h$ where $k_{hu} = 1$ if subcategory h has usage u, and zero otherwise. So, I estimate 152 subcategory parameters $\bar{\Phi}_h$ and also 5 parameters $\tilde{\Phi}_{iu}$ for the usage types for each i, the mean of which (over i) is not identified so normalised to 1.

The standard deviation of the lot specific idiosyncratic value σ_l vary depending on the combination of goods auctioned together in the lot. To simplify posterior sampling I find the 60 most common category combinations (e.g. $\frac{2}{3}$ dairy $\frac{1}{3}$ cereal), and associate each combination with a unique parameter, as well as an 'other' parameter for the remaining 5.5% of combinations.

F.2.2 Priors and Hierarchical Distributions

Write $\boldsymbol{\psi}_i$ as the stacked vector of unique elements of the matrix Ψ_i and food bank specific components of Φ_i ($\tilde{\Phi}_i$), which are drawn from hierarchical distribution $N(\boldsymbol{\psi}, \Sigma^{\psi})$. The hierarchical framework reduces the posterior variance of estimated parameters at a cost of bias, as estimated parameters are drawn together. Observations with a lot of identifying variation place little weight on the hierarchical parameters, whereas observations with little identifying variation place more weight on hierarchical parameters. Any bias caused by this framework causes parameters to be drawn together, so that my estimates will be biased in favour of the Old System rather than the Choice System. For the parameters of $\boldsymbol{\psi}$ corresponding to the mean of $\tilde{\Phi}_{iu}$ I use strong priors of 1, as the mean of these parameters is not separately identified from $\bar{\Phi}$. I constrain Σ^{ψ} to be a diagonal matrix to simplify sampling (discussed shortly). Otherwise, I assume weak normal-inverse-gamma priors for these objects.

For both the distance coefficients and subcategory weights $\bar{\Phi}$ I use weak normal priors. For σ_l I use weak inverse gamma priors, while for λ_i I use gamma priors, placing around 100 times more weight on the data than on the prior of mean of 1. For the parameters of the transition process $(\delta_{ig}, \mu_{ig}, \Sigma_{ig})$, because the requirements for identification are strong, I use informative normal-inverse-gamma priors. The means of μ_{ig} and Σ_{ig} are taken as the mean and variance of $-w_{igt}$ and the prior mean of δ_{ig} is set to 0.[46] I set the prior shape parameters, which influence prior variances, to ensure around 100 times more weight is put on the data than on the priors.

F.2.3 Computation

I focus on data from only the highest 25 bids placed each day by each food bank.[47] I begin the Carter-Kohn algorithm from the 61st day in my sample. To reduce dependence on the intial state, which I set to the long-run mean, I set the initial state variance to 10 times the long-run variance and also discard the first 40 days sampled. All truncated normal draws are performed using Botev (2017)'s procedure.

I constrain elements of Ψ and Φ to be positive using truncated priors, for identification requirements and also because the sampler often diverged otherwise. The $\bar{\Phi}^h$ constraint binds for just 5 apparently unpopular subcategories with limited shelf lives, including fresh bread and milk. I constrain $\boldsymbol{\delta}$ and Σ to be diagonal, as separately identifying off-diagonals from the off-diagonal components of Ψ proved difficult. I also impose diagonal elements of $\boldsymbol{\delta}$ to be within $[-1, 0]$, to ensure the process remained stationary.

As all my priors are conjugate I sample directly from the conditional posterior distributions. The only exception is for the parameters of the hierarchical distribution ψ, Σ^{ψ}, due to the constraints on Ψ, for which I using adaptively tuned Metropolis Hastings. I sample beliefs only every fifth iteration, and run the full procedure for one million iterations, burning out the first half. Initial points are drawn from the prior distributions. I run two independent chains, then uniformly sample 500 points from each chain, keeping 1,000 parameter draws in total, which I use in later steps.

[46] This is because we expect that on average winnings more or less off-set shortfalls in net donations. $\delta_{ig} = 0$ ensures my results are biased in favour of the Old System. If net donations are endogenous ($\delta < 0$) this allows food banks to use their winnings to influence future net donations. Choice is then even more valuable. When local donations are negatively correlated with previous winnings food banks can focus on only winning food from the Choice System they know they cannot get from local donors.

[47] This assumption considerably speeds up both computation and convergence, and is not expected to significantly impact results. Even Type 1 food banks place more than 15 bids only 1% of the time. However on 25% of auction days more than 25 unique (non-homogenous) lots are auctioned simultaneously. By ignoring that food banks also choose not to bid on any more than the first 25 lots I bias my results towards food banks being willing to bid on too high a proportion of lots. However, the degree of this bias should be small. I already take into account that food banks only bid on maybe 5 lots, then choose not to bid on the other 20, so that there is relatively little extra information conveyed by them also choosing not to bid on lots 26+.

F.3 Step 3.

In the third step I evaluate the continuation value as a function of observed bids and the pseudo-static pay-off, before backing out the combination flow pay-off.[48]

Because states are continuous I must evaluate the continuation value over a finite set of states. For each food bank I form a 30^4 dimensional grid of state, with 30 points per dimension from the 2.5 and 97.5 percentiles of their sampled states. For each t I form the (approximate) posterior distribution of stocks $q_t(\mathbf{s}_i^g)$ using my $1,000$ draws of $\mathbf{s}_{it}^g$. I use an independent normal kernel and Silverman's rule of thumb. I evaluate the maximised pay-off at each time period $\tilde{W}(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s})$ using the formulae in Appendix E (with finite sample approximations) for each parameter draw.[49]

I fit a polynomial function of the states to the ex-ante value function, including interaction terms. I use a standard least squares procedure, weighting by the sum of posterior probabilities. The main version uses a quadratic. This is primarily because my counterfactuals occasionally require extrapolation. Higher order polynomials typically lead to extrapolated values much further from the interpolated values. I validate this approximation by considering the R^2s from these regressions, detailed in Appendix H.3.

Given the approximated ex-ante value function I evaluate the continuation value by taking an expectation of the polynomial function, given the distribution of $\mathbf{s}_{it+1}$ given $\mathbf{s}_{it}^a$. I then back out $\hat{\pi}$ using the pseudo-payoffs $\hat{\kappa}$ and the continuation value.

G Additional Estimation Results

I now report additional estimation results not reported in 6. This includes tables and plots of parameter estimates, Gelman-Rubin Convergence tests, and model fit.

The key category specific first stage parameters are given in Figure 16. Meanwhile,

[48] A small note: The marginal welfare from consuming a lot with subcategory composition $\mathbf{z}_{tl}^h$ is just $\Phi_i \mathbf{z}_{tl}^h$, and does not depend on $\mathbf{s}_{it}^h$. It does not matter that Φ is a pseudo-payoff object, not a structural parameter. This is because, if flow-payoffs are affine in $\mathbf{s}_{it}^h$, so is the pseudo-payoff. The marginal pseudo-payoff gives the expected benefit going forward of these extra stocks. This does not interact with any other stocks due to affinity. Therefore, in counterfactuals the expected discounted benefit of receiving an extra $\mathbf{z}_{tl}^h$ is the same as under the Choice System, and independent of $\mathbf{s}_{it}^h$.

[49] I should account for sampling variation in these finite sample expectations. The large number of time periods means we expect fairly little variation. I previously considered a bootstrap procedure, however this does not account for correlations between sample expectations and the sample parameters, so over estimates the posterior variances.

the subcategory fixed effects and demand parameters are given in Figure 17.

Figure 16: Category Specific First Stage Parameters

Category	Shape	Scale	Scale (other)	Maroon	Loads	Threshold 1	Threshold 2
Baby	0.0856	3090	-379	-963	-9,400	-3,270	-199
	(0.0514,0.121)	(2840,3390)	(-502,-234)	(-1,620,-301)	(-14,800,-5,550)	(-3,520,-3,020)	(-240,-164)
Bev	0.0175	2180	-379	-808	-5,200	-3,590	-307
	(-0.00124,0.0396)	(2080,2290)	(-502,-234)	(-2,360,800)	(-5,790,-4,610)	(-4,270,-2,940)	(-427,-206)
Baked	0.0856	2270	-379	-963	-11,000	-3,270	-199
	(0.0514,0.121)	(1980,2600)	(-502,-234)	(-1,620,-301)	(-20,400,-6,380)	(-3,520,-3,020)	(-240,-164)
Cereal	0.128	4540	-379	-2,340	-2,410	-3,270	-199
	(0.0857,0.174)	(4340,4740)	(-502,-234)	(-4,570,-91.2)	(-4,530,-816)	(-3,520,-3,020)	(-240,-164)
Condiment	0.296	3880	-379	425	-7,280	-3,270	-199
	(0.235,0.361)	(3640,4110)	(-502,-234)	(-1,460,2260)	(-9,830,-4,990)	(-3,520,-3,020)	(-240,-164)
Dairy	-0.0352	2430	-379	-963	-5,980	-3,270	-199
	(-0.0685,6.35e-05)	(2300,2570)	(-502,-234)	(-1,620,-301)	(-7,090,-5,040)	(-3,520,-3,020)	(-240,-164)
Frozen	0.0195	2630	-379	-1,110	-6,380	-3,270	-199
	(-0.0271,0.0708)	(2440,2850)	(-502,-234)	(-2,740,567)	(-9,900,-3,890)	(-3,520,-3,020)	(-240,-164)
H/B	0.0856	3460	-379	-963	-7,070	-3,270	-199
	(0.0514,0.121)	(3190,3770)	(-502,-234)	(-1,620,-301)	(-10,500,-4,270)	(-3,520,-3,020)	(-240,-164)
Meals	0.141	4070	-379	-963	-4,060	-3,270	-199
	(0.0999,0.181)	(3890,4260)	(-502,-234)	(-1,620,-301)	(-5,400,-2,910)	(-3,520,-3,020)	(-240,-164)
Meat	0.0856	5180	-379	1110	-12,500	-3,270	-199
	(0.0514,0.121)	(4760,5570)	(-502,-234)	(-1,560,3840)	(-19,600,-7,510)	(-3,520,-3,020)	(-240,-164)
Cleaning	0.174	2900	-379	1790	-4,830	-3,270	-199
	(0.12,0.232)	(2690,3110)	(-502,-234)	(120,3450)	(-6,150,-3,580)	(-3,520,-3,020)	(-240,-164)
Nutri	0.0856	2280	-379	-963	0	-3,270	-199
	(0.0514,0.121)	(1770,2910)	(-502,-234)	(-1,620,-301)	(0,0)	(-3,520,-3,020)	(-240,-164)
Pasta	0.0856	5360	-379	-963	-8,040	-3,270	-199
	(0.0514,0.121)	(4670,6080)	(-502,-234)	(-1,620,-301)	(-13,800,-3,950)	(-3,520,-3,020)	(-240,-164)
Snack	0.0391	2230	-379	-776	-6,540	-3,950	-222
	(0.0203,0.0602)	(2150,2310)	(-502,-234)	(-2,310,711)	(-7,780,-5,540)	(-4,650,-3,230)	(-316,-135)
Vegetables	0.0856	3620	-379	1100	-2,890	-3,270	-199
	(0.0514,0.121)	(3330,3950)	(-502,-234)	(-824,3090)	(-4,790,-1,400)	(-3,520,-3,020)	(-240,-164)

Note: 95% Credible Intervals in parentheses. Several parameters are constrained equal, due to lack of observations.

For the second stage parameters, plots include the net donation means $\boldsymbol{\mu}_i$ and feedback parameters $\boldsymbol{\delta}_i$ in Figure 18, transport costs and the marginal value of wealth parameters λ_i in Figure 19, the standard deviation of the lot specific idiosyncratic values σ_l and subcategory weights Φ^h in Figure 20, and the determinants of the pseudo-payoff function Φ_i and Ψ_i in Figure 21.

Figure 22 reports Gelman-Rubin statistics, as discussed in Gelman et al. (1995), displaying the proportion of statistics below the recommended cutoffs of 1.2 and 1.1. I report results for all the food banks, and separately for the 'Type 1' food banks specifically, as these are the most important for the analysis. Broadly I have evidence of convergence. 5 'Type 2' food banks had implausible parameter estimates and were dropped. Convergence failure was generally due to multiple modes, and was specific to food bank $\times$ storage type combinations for those who rarely bid on (or never won) certain types of food. This suggests a problem of non-identification. Because the model still fits the data well, I do not worry about the lack of convergence.[50]

[50]The Gelman-Rubin statistics presented assume the target distribution is approximately normal, and may fail to detect convergence if the target distribution is sufficiently skewed. I also consider the convergence statistics for the first and second moments, rather than the full distribution, and find proportions increase by around 5 percentage points.

Figure 17: First Stage Parameter Estimates

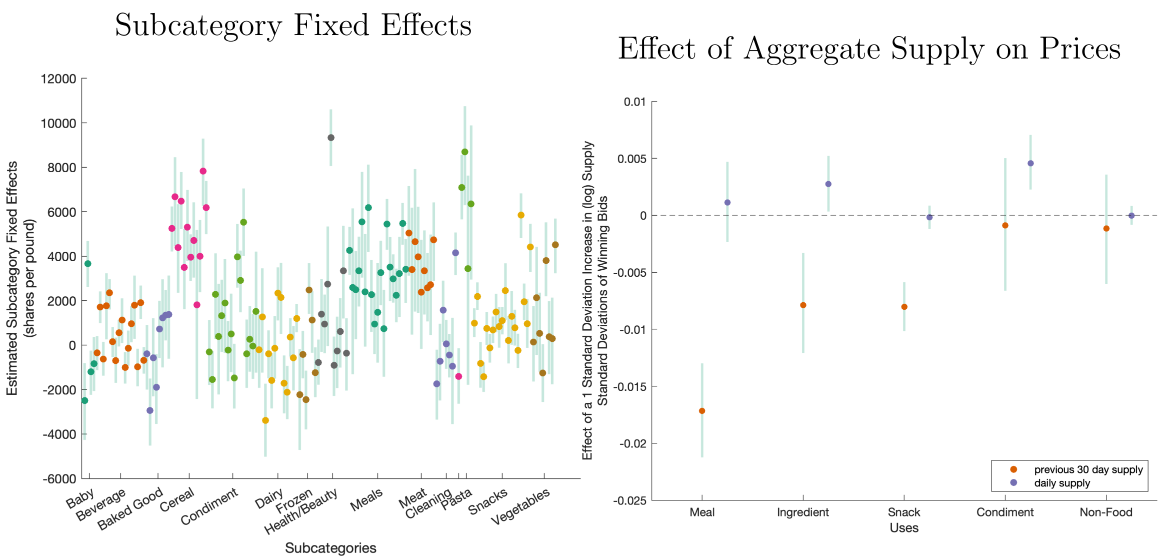

Note: Panel (B) shows coefficients on aggregate supply, by use for both daily and the previous month's supply. Points give posterior means, and 95% Credible Intervals are given by the shaded lines. In non-standardised terms at the mean of 1600 tons of meals (food that can be consumed as a meal in itself) per month, an increase in the previous 30 day supply of meals by 1000 tons, around 100 loads, decreases the expected winning bid by 350 shares.

Figure 23 plots the estimated and empirical probability a food bank wins a lot given their bid, where bids are measured in distance from the reservation price. In estimation I drop the first 60, and the final 150, days and then randomly samples 95% of the remaining data. The 5% is used as a validation dataset. While the parametric specification generally fits the data well, it cannot rationalise food banks' bids being anchored around zero. However this inaccuracy is not large, even if statistically significant — the vertical distance between the two lines never exceeds 0.05.

H Robustness

This Appendix investigates how robust my results are to certain key assumptions and simplifications made in the main text. Robustness exercises are split across the three stages of my estimation procedure in Appendices H.1, H.2, and H.3 respectively.

Figure 18: Estimated mean net donations $\boldsymbol{\mu}_i$ and donation feedback parameters $\boldsymbol{\delta}_i$

(A) (B)

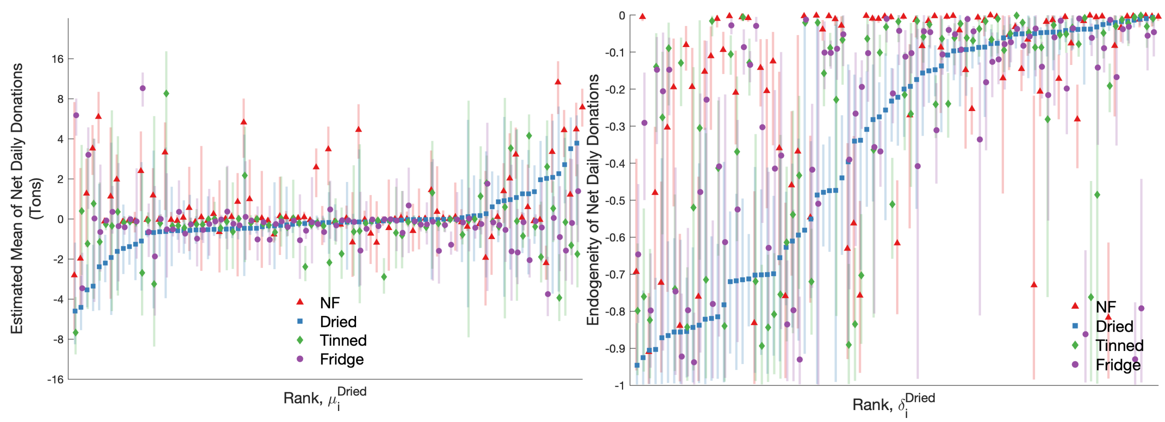

Figure 19: Estimated distance and marginal value of wealth (λ_i) parameters

(A) (B)

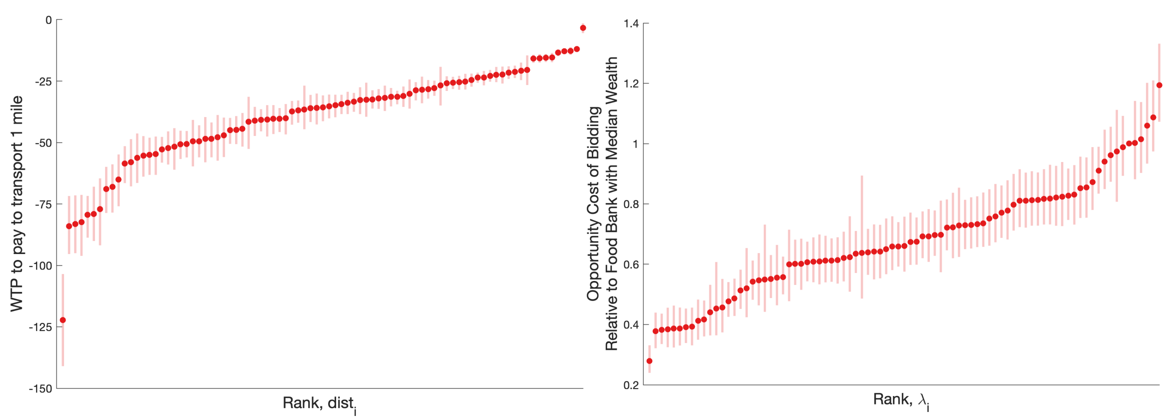

H.1 First Stage

H.1.1 Food bank Specific Beliefs

I implicitly imposed that $\Gamma_i = \Gamma$: Every food bank faces the same distribution of rival bids. This permits estimating Γ on the distribution of winning bids only. This is testable by testing whether the distribution of food bank i's rivals' highest bids is significantly different from the distribution of winning bids using a score test. In figure 24 panel (A) I present the distribution of test statistics across food banks. Under the null hypothesis these statistics take a χ^2 distribution with 235 degrees of freedom (the number of first stage parameters). None of these hypothesis tests can

Figure 20: Lot specific standard deviations and subcategory weights (Φ^h)

(A) (B)

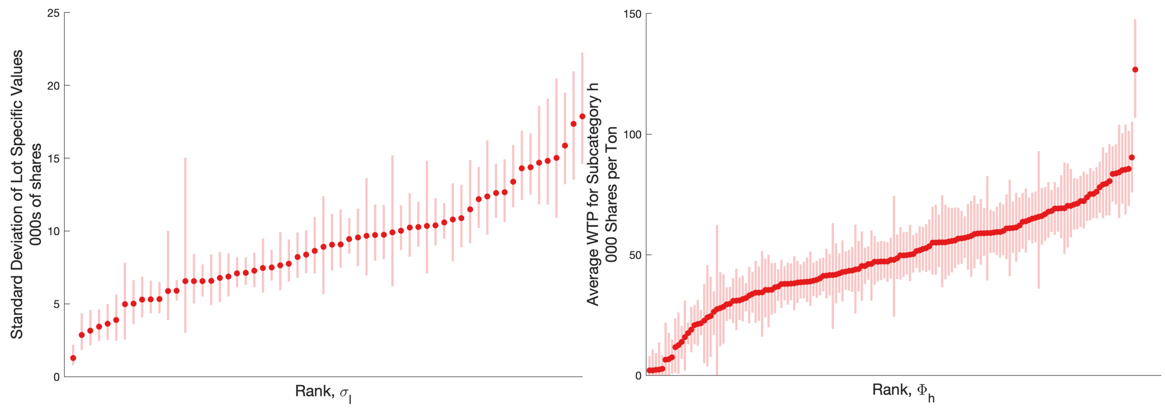

Figure 21: Estimated Pseudo-Payoff parameters Φ_i and Ψ_i.

(A) (B)

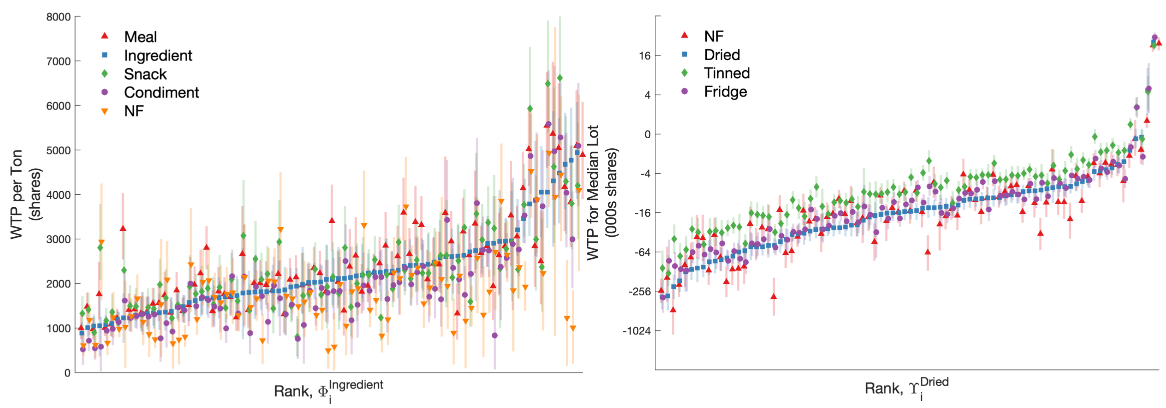

reject the null at 10% significance level.

H.1.2 Dependence on Aggregate Supply

For estimation to be feasible I require that beliefs do not depend on any individual food banks' state. If equilibrium is sufficiently competitive then no individual food bank's behaviour will be able to significantly shift the distribution of winning bids. If so, then variation in an individual food bank's state will not shift this distribution either. The results presented in Appendix H.1.1 support this hypothesis. I also consider whether the distribution of equilibrium winning bids changes when data from food bank i *and* the auctions they won are removed from the data. If i has

Figure 22: Gelman-Rubin Convergence Statistics

	Type 1 food banks		all food banks	
Parameters	Prop < 1.1	Prop < 1.2	Prop < 1.1	Prop < 1.2
Γ	1	1	-	-
$\boldsymbol{\delta}_i$	0.857	0.902	0.774	0.826
$\boldsymbol{\mu}_i$	0.884	0.893	0.832	0.859
$\boldsymbol{\Sigma}_i$	0.777	0.839	0.824	0.859
$\bar{\boldsymbol{\Phi}}$	0.98	0.993	0.98	0.993
$\tilde{\boldsymbol{\Phi}}_i$	0.971	0.993	0.948	0.969
$\boldsymbol{\lambda}_i$	1	1	1	1
$\boldsymbol{\sigma}_l$	0.983	1	0.983	1
Distance	0.964	1	0.976	1
$\boldsymbol{\Psi}_i$	0.886	0.918	0.807	0.849
$\boldsymbol{\psi}$	0.867	0.933	0.867	0.933
$\boldsymbol{\Sigma}^{\psi}$	0.867	0.933	0.867	0.933

a significant effect on the distribution of winning bids, we would expect that the distribution is different when we drop all the data from i. In figure 24 panel (B) I present the distribution of score test statistics across food banks. Under the null these statistics take a χ^2 distribution with 235 degrees of freedom. None of these tests can reject the null hypothesis at 10% significance level.

H.1.3 Independence of Winning Bids

Next, I investigate whether winning bids within a period are conditionally independent across auctions. This assumption ensures the joint probabilities of combinatorial outcomes $P(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s}_t)$ can be written as products of the marginal distributions. Given my quadratic pseudo-payoff parametrisation I only require that winning bids are pairwise independent conditional on observed covariates. Write $\bar{b}_{lt}$ for the winning bid. I investigate dependence with the following specification:

$$\bar{b}_{lt} = \beta^1 \bar{b}_{l't} + \beta^2 \mathbf{x}_{lt} + \beta^3 \mathbf{x}_{l't} + \beta^4 \mathbf{s}_t^0 + \varepsilon_{lt}$$

$\mathbf{x}_{lt}$ give lot specific covariates, using the covariates included in the first estimation step. $\mathbf{s}_t^0$ give time specific common state variables that do not vary across lots. I include every pair of auctions (l, l') that occur simultaneously ($\approx 800{,}000$ observations). Under the null of independence β^1 should equal zero. The degree of correlation may depend on the lot characteristics, so I also consider two relaxed specifications. First

Figure 23: First Stage Fit, actual vs simulated

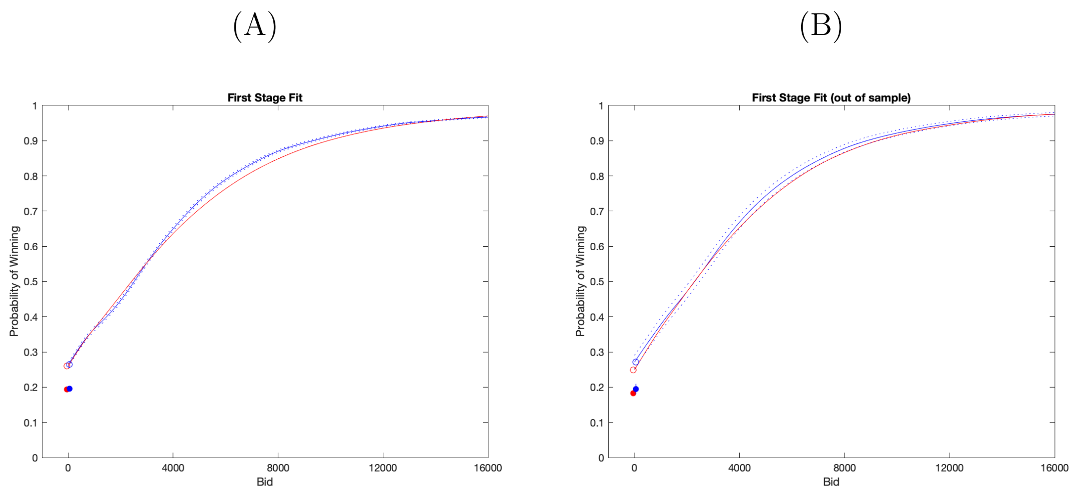

Note: probability of winning given bid, i.e. cdf of winning bids. discont due to ties. simulated values from estimated distribution vs empirical distribution. averaged over covariates.

I allow the correlation to vary with covariates by interacting $\bar{b}_{l't}$ with all three sets of covariates. Second, I allow for triple interactions between $\bar{b}_{l't}$, $\mathbf{x}_{lt}$, and $\mathbf{x}_{mt}$.

I consider significance of the $\bar{b}_{l't}$ coefficients using asymptotic F-tests. However, it is also important to consider how much variation in $\bar{b}_{lt}$ $\bar{b}_{l't}$ can explain. If $\bar{b}_{l't}$ has very little explanatory power, then the extent of the dependence is minor. Any departure from independence is unlikely to cause much inaccuracy in my results, since the true joint probabilities are close to the product of the marginal probabilities.

Results are presented in Figure 25. I can reject the null hypothesis of independence at the 1% significance level in all specifications: The independence assumption is invalid. However, as is evident from examining the R^2 values, the degree of dependence is extremely small. The covariates alone account for 42.76% of the variation in winning bids. Including the $\bar{b}_{l't}$ interactions is then only able to explain an additional 0.3% of the variation in winning bids. This suggests that winning bids are very close to being independent, even though we can reject independence. Therefore while I have found this independence assumption to be invalid, I have also found that it is likely to be a very good approximation to food banks' beliefs.

Figure 24: Robustness: Stage 1

(A) (B)

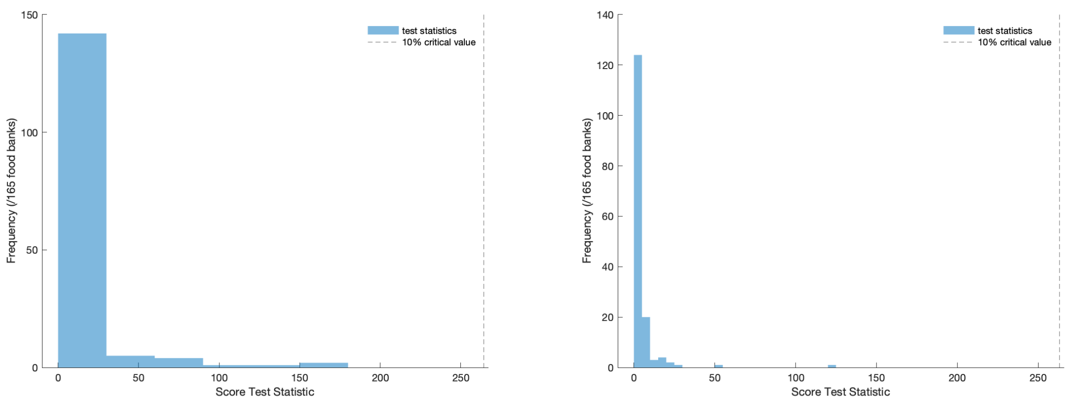

Note: The figure plots Score Test statistics from two robustness checks. Panel (A) relaxes the restriction that every food bank has the same equilibrium beliefs, while panel (B) tests whether any individual food bank's bidding behaviour has a significant effect on the distribution of winning bids.

Figure 25: Robustness: Independence of Winning Bids

Specification	Covariates	F test df	p-value	R^2
$\bar{b}_{l't}$		1	0	0.1097
Covariates only	✓			0.4276
$\bar{b}_{l't}$	✓	1	6.49e-66	0.4279
$\bar{b}_{l't} \times (\mathbf{x}_{l't}, \mathbf{x}_{lt}, \mathbf{s}_{0t})$	✓	424	1.85e-211	0.43
$\bar{b}_{l't} \times (\mathbf{x}_{l't}, \mathbf{x}_{lt}, \mathbf{s}_{0t}, [\mathbf{x}_{l't} \times \mathbf{x}_{lt}])$	✓	632	0	0.4313

Note: The F test degrees of freedom and p-value refer to the hypothesis tests that all coefficients on $\bar{b}_{l't}$ are equal to zero, where the degrees of freedom gives the number of coefficients being considered.

H.2 Second Stage

H.2.1 Incorporating the Common State

The pseudo-payoff function κ ought to depend on the common state variables $\mathbf{s}_{0t}$, even though we have several reasons to expect that this relationship will be weak. I consider a parsimonious specification allowing κ to depend linearly on the demand indices estimated in the first step, as the common states will only enter the continuation value, and hence κ, through beliefs. I estimate the following:

$$\kappa(\mathbf{s}_i, \mathbf{s}^0) = \Phi_i \mathbf{s}_i^h + \Upsilon_i[\boldsymbol{\vartheta}(\mathbf{s}^0) \cdot \mathbf{s}_i^u] - \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g$$

Where $\boldsymbol{\vartheta}(\mathbf{s}^0) \cdot \mathbf{s}_i^u$ gives the elementwise product of the demand indices and the food banks' stocks by usage type. Υ_i is a 1×5 vector of i specific coefficients. This simple specification can be considered a linear approximation of any true dependence. If the dependence is particularly strong, we would expect this to be picked up in a linear relationship. I interact the index with food banks' usage stocks to ensure it impacts their marginal pseudo-payoff. I focus on stocks by usage type as the index affects how easily the food bank can win the types of food it wants, on behalf of their clients.

In Figure 26 panel (A) I plot estimates of Υ_i. I focus on just the 28 Type 1 food banks, as we expect it is the larger food banks who (if any) will be dynamically strategic. 27.8% of parameters are significant at the 5% significance level. However, parameters are generally very small. Only two of the significant parameters have economically significant magnitudes, with a $1sd$ deviation difference in the demand indices leading to more than a 100 difference in willingness to pay.

H.2.2 Endogeneity of the Inverse Bid System

In Step 4. of the Gibbs Sampler, I run a bayesian regression on the Observation Equation. However, this equation is endogenous, featuring $\Gamma_m(b_{itm})$ (generally a function of υ_{itl}) as a regressor. As discussed in Altmann (2022) we do not expect this endogeneity to be large: $\Gamma_m(b_{itm})$ depends much more strongly on υ_{itl} and other variables. Nonetheless, I now investigate this suggestion.

As in Altmann (2022) we use $\mathbf{z}_{tl}^g(\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g)^T$ as an instrument for $\mathbf{z}_{tl}^g(\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2\sum_{m \neq l} \Gamma_m(b_{itm})\mathbf{z}_{tm}^g)^T$. It is clear how this instrument is both valid and relevant. We instrument using variation in bids caused by variation in the characteristics of that lot only, excluding the effect of other lots.

I do not run a full bayesian instrumental variable procedure, which would require augmenting my sampler to include instruments and require normality assumptions on the endogeneity. This is difficult to justify given we observe $\Gamma_m(b_{itm})$ and know it is far from normal. Instead, for each draw of my unobserved states I run (non-bayesian) IV and look for evidence of endogeneity. This heuristic procedure means we cannot perform valid inference, but if the degree of endogeneity were severe it should shed light on its existence. In particular, I treat the augmented data from steps 2 and 3 of the sampler as 'known', then use the instrument to form a control function and proceed in the standard frequentist way. I repeat this exercise for each of the 1,000 draws of augmented data (giving me something similar to a credible interval). The diagonal coefficients on these control functions (corresponding to the endogeneity of the diagonal elements of Ψ_i) are plotted for the 28 Type 1 food banks in Figure 26 panel (B). Results are normalised by the corresponding components of Ψ_i to highlight the magnitude of the relative endogeneity. The results are clustered close to 0, suggesting that if there is endogeneity it is extremely small. All but 1 have absolute magnitude of the endogeneity less than 10%.

Figure 26: Robustness: Stage 2

(A) (B)

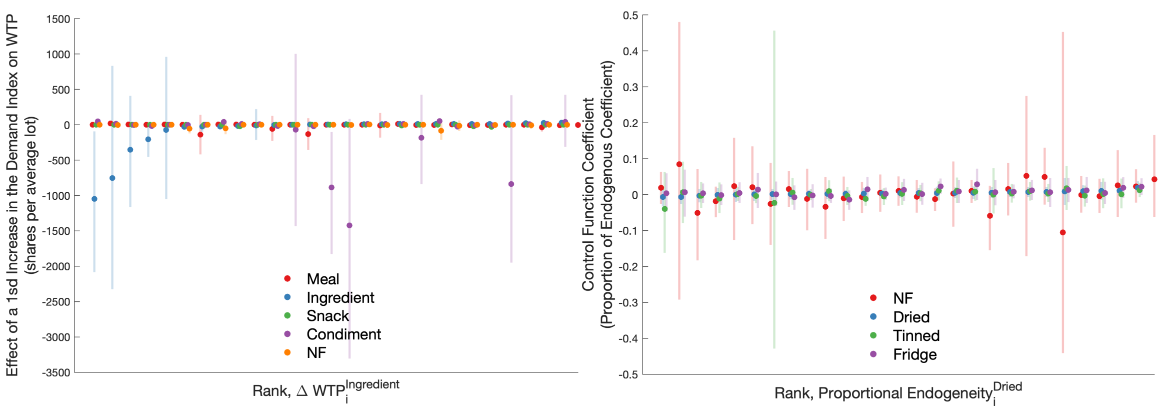

Note: Panel (A) plots posterior means and 95% credible intervals of the demand index coefficients. Estimates give the effect on Willingness to Pay of a $1sd$ increase in the demand index for that usage type. Panel (B) plots means and 95% (approximate) credible intervals of the control function parameters, presented as proportions of Ψ.

H.3 Third Stage

The quadratic approximation made in the third estimation step is technically incompatible with the parametric assumptions made in the second step. However, this approximation fits well. I consider the R^2s from forming this least squares approximation. 90% of these value lie between 0.99 and 1, with the lowest at 0.91. The fit is strong because of the quadratic term which appears in equation 9.

I Simulation Details

In this Appendix I describe my simulations: First, the Choice System, then the Old System. The other mechanisms only involve minor alterations of the Old System.

In the simulation for parameter draw r I use the corresponding sampled $\mathbf{x}_{it}^r$ and $\boldsymbol{\upsilon}_{it}^r$, thereby maintaining correlations between these and the model parameters. This ensures results are slightly more robust to model misspecification. Meanwhile, to assess model fit I draw the lot-specific values $\boldsymbol{\upsilon}_{it}$ from their estimated posterior distributions, else the simulated bids are trivially the same as observed bids.[51]

I.1 Choice System

Because I observe and estimate my model on equilibrium bidding data under the Choice System, I do not need to solve for equilibrium beliefs or continuation values. I can instead use the estimated beliefs and pseudo-payoffs. This approach would not be valid if I wanted to consider changes to the Choice System, such as changes in food banks' budgets. The central problem then concerns the bidding function, as this involves a complex combinatorial problem of deciding which combination of lots to bid on. I use a greedy algorithm for this purpose. Beginning with no auctions entered, I iteratively add the lot that yields the highest marginal improvement in payoffs, then re-optimise bids. Because I impose Ψ_i is strictly negative, imposing gross substitutes, this algorithm is guaranteed to find the global optimum.

I.2 Old System

I treat time as continuous, and each day is of length 1. Therefore local donations and offers of food from Feeding America are received continuously during the day. To ensure results are easily comparable across the Choice System and Old System simulations, when evaluating welfare I treat local donations and flow payoffs as only accruing at the end of the day. To evaluate the equilibrium value function I treat both these objects as continuous.

[51] I do not find equilibrium for all 1000 posterior draws of the parameters, as this would require too much computation power. Instead, I use a random sample of 30 parameter draws. I then average value functions over these 30 draws, dropping any that did not converge ($\approx 2\%$). I find relatively little variation in equilibrium objects across parameter draws.

I.2.1 Set Up

Food is donated to Feeding America at some exogenous rate. Conditional on arriving, the load has various characteristics. The rate and probability of these characteristics are taken from the empirical distribution. What matters in the agent's problem is their belief about the rate at which they are offered food, and the probabilities of characteristics they are offered. Food banks' positions in the queue are determined exactly as described in Prendergast (2022).

Net donations arrive at exogenous Poisson rate q_i. Conditional on arriving, the net donation $\tilde{\mathbf{x}}_{it}$ is normally distributed: $N(\tilde{\boldsymbol{\delta}}\mathbf{s}_{it} + \tilde{\boldsymbol{\mu}}_i, \tilde{\Sigma}_i)$. The parameters $\tilde{\boldsymbol{\delta}}, \tilde{\boldsymbol{\mu}}_i, \tilde{\Sigma}_i$ are set so that when we integrate over the day's net donations, the resulting distribution has the same mean and standard deviation as estimated. Finally, I normalise q_i to 1, as this is not identified from my discrete data.

The lot specific payoff is the same as in the text. The deterministic flow-payoff is accrued at every instant, so that if i accepts load l at t they receive $\upsilon_{ilt} + \tilde{\pi}(\mathbf{s}_{it} + \mathbf{z}_l)$, or $\tilde{\pi}(\mathbf{s}_{it})$ if they reject. The payoff function $\tilde{\pi}$ is set so that integrating over a day, if stocks do not change, they receive the same payoff as in the discrete time case.

I discretise the individual stock space using a grid formed of eleven evenly spaced points along each dimension. Points range from one interquartile range below the 2.5% percentile of sampled stocks, up to one range above the 97.5% percentile. I use flexible Bernstein polynomials basis functions to interpolate the Value Function. These polynomials are convenient as I can easily form expectations of the value function after receiving normally distributed net donations, and use simple matrix multiplication to evaluate expectations of being offered lots with various characteristics.

I.2.2 Equilibrium

Write the agent's value function as $V(t, \mathbf{s}_i, \mathbf{s}_0)$. This gives their presented discounted value from state $(\mathbf{s}_i, \mathbf{s}_0)$ at time t. I augment the common state to include the newly defined priorities and Goal Factors. If the food bank is offered a load at t they must be at the head of the queue, and so have the highest priority. If they are offered load l characterised by $(\upsilon_{ilt}, \mathbf{z}^g_{lt}, \mathbf{z}^h_{lt})$, they will accept if $\upsilon_{ilt} + V(t, \mathbf{s}_i + \mathbf{z}^g_{ilt}, \mathbf{s}_0) \geq V(t, \mathbf{s}_i, \mathbf{s}_0)$.

The agent believes that Feeding America will offer them a load at Poisson rate $p_i(t, \mathbf{s}_0)$. In principle this should depend on the state of every food bank, including i, however I will assume that food banks do not observe each others' states. The agent then believes that, conditional on receiving an offer, the load will have characteristics $(\upsilon_i, \mathbf{z}^g, \mathbf{z}^h)$ with probability density $f_i^c(\upsilon_i, \mathbf{z}^g, \mathbf{z}^h; t, \mathbf{s}_0)$.

I assume a Markov Perfect Equilibrium in symmetric strategies, as defined in section 4. This requires that food banks make optimal accept/reject decisions given beliefs about p and f^c, and that their beliefs are consistent with the observed realisation of the rates at which Feeding America offers them loads. As I have assumed a stationary equilibrium, I require that p and f are conditionally independent of t.

I assume food banks do not observe others' stocks, nor when loads are offered to any one else (hence aggregate supply is also unobserved). I therefore assume the only objects used to form their beliefs are $\mathbf{s}_i$, their own (relative) Goal Factor, and the time since their last offer τ. The offer rates p_i and the distribution of offered lot characteristics f_i^c are food bank specific. In principle I should allow p to depend on τ, however for simplicity I assume it does not.[52] I also assume food banks beliefs do not change conditional on the previous history of offers. That is, food banks do not infer from frequent offers that offers will be more frequent in future. For f_i^c, I split lots into the same 60 discrete category combinations used for the lot specific variances σ_l, detailed in Appendix F.2. Therefore f_i^c can be interpreted as conditional probabilities. Then, conditional on the category combination, I assume food banks believe that, in equilibrium, the distance between the lot and a given food bank is normally distributed with some mean and variance. I also assume that, conditional on category combination, food banks believe $\mathbf{z}^h$ is also normally distributed.

I.2.3 The Optimal Control Problem

Under the assumptions outlined above, we can write the value function as $V_i(\tau, \mathbf{s}_i)$. Food bank i, that is offered load l, accepts the load if $\upsilon_{il} + V_i(0, \mathbf{s}_i + \mathbf{z}_l^g) \geq V_i(0, \mathbf{s}_i)$. The Hamilton-Jacobi-Bellman differential equation is given by:

$$\begin{aligned}(\rho + p_i + q_i)V_i(\tau, \mathbf{s}_i) &= q_i E_X[V_i(\tau, \mathbf{s}_i + X)|\mathbf{s}_i] + \tilde{\pi}(\mathbf{s}_i) + \frac{\partial V_i(\tau, \mathbf{s}_i)}{\partial \tau} \\ &\quad + p_i \sum_c E_{\upsilon,z}[\max\{\upsilon_{il} + V_i(0, \mathbf{s}_i + \mathbf{z}_l), V_i(0, \mathbf{s}_i)\}\,|c, \mathbf{s}_i, \tau] f_i^c \quad (12)\end{aligned}$$

[52]This is unlikely to be a problem, since, for most food banks, offers occur so frequently it is unlikely they ever have to wait particularly long before receiving the next offer.

Where ρ gives the discount rate ($= (1-\beta)/\beta$). To solve this differential equation recognise that the solution must be independent of τ. The equation can then be solved using numerical methods,

For given $\mathbf{V}_i^k$ and beliefs (p_i^k, f^{ck}) I evaluate Equation 12, using Bernstein Polynomials for interpolation. I use successive approximations and switch to a dampened Newton-Kantorovich algorithm when progress slows. I then simulate the Old System using these value functions, before updating beliefs, (p_i^k, f_i^{ck}), using a frequency estimator. I repeat this process until the rates and estimated probabilities change by a total less than $10^{-4}\times$ the euclidean norm of (p_i^k, f_i^{ck}).

I.2.4 Other Mechanisms

The Sequential Offer Mechanism is done exactly the same as above, but allowing every load of food to be offered to every food bank. The two Closest mechanisms are exactly the same, but offering food by order of distance from the lots' origin. The random allocation does not require solving for equilibrium.

For the Sequential Auction mechanism I exploit that, in equilibrium, lots will always be allocated to the food bank with the highest marginal benefit (if this is positive). Therefore, I simulate the probability that a given food bank has the highest marginal benefit (essentially, highest bid), given their value. In this case, the Hamilton-Bellman-Jacobi equation can be written as:

$$
\begin{aligned}
&(\rho + p_i + q_i)V_i(\tau, \mathbf{s}_i) = q_i E_X[V_i(\tau, \mathbf{s}_i + X)|\mathbf{s}_i] + \tilde{\pi}(\mathbf{s}_i) + \frac{\partial V_i(\tau, \mathbf{s}_i)}{\partial \tau} \\
&+\sum_c p^c E_{\upsilon,z}[\tilde{\Gamma}_c(\upsilon_{il} + V_i(0, \mathbf{s}_i + \mathbf{z}_l) - V_i(0, \mathbf{s}_i)|c) \max\{\upsilon_{il} + V_i(0, \mathbf{s}_i + \mathbf{z}_l), V_i(0, \mathbf{s}_i)\}\,|c, \mathbf{s}_i, \tau]
\end{aligned} \tag{13}
$$

Where p^c gives the rate of the arrival of donations with characteristics c, and $\tilde{\Gamma}_c(\upsilon_{il} + V_i(0, \mathbf{s}_i + \mathbf{z}_l) - V_i(0, \mathbf{s}_i)|c)$ gives food bank i's belief about the probability they have the highest value. $\tilde{\Gamma}$ is very similar to the Γ marginal win probabilities from the main model, so I assume the same Generalised Extreme Value parameterisation, with a distinct set of parameters for each set of characteristics. Numerically solving for equilibrium then follows as for the Old System.