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 increasing the supply of donated food by around a third. 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, NYU, 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. Using a structural model allows me to build on previous evidence in Prendergast (2022) using detailed equilibrium analysis of welfare under counterfactual mechanisms. I develop an 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 numbers of

---

[1] I use the term 'needs' to capture the determinants of food banks' demand functions, incorporating both what they have a preference for, on behalf of their clients, as well as what they have capacity to store in their warehouse.

clients, and receive different types of food at different times from local donors. I then use these estimates to evaluate equilibrium allocations under 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 36% higher than under the Old System. Furthermore, despite fears that a market mechanism could lead to a more unequal distribution of food, on average 89% 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 (2023). 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 observed 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 Sampler, employing data-augmentation 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 adapt the three step estimation procedure introduced in Altmann (2023) to allow for unobserved states. 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 flow payoffs from 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.

Using the estimated model I then simulate equilibrium allocations under the Old

---

[2] How bidding behaviour varies with the number 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 identifies 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.

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 36% higher under the Choice System than under the Old System. This is roughly equivalent to an additional 84 tons of food allocated each day, enough to support an additional 34,000 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 89% of food banks are significantly better off under the Choice System.

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. 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, 24% of this improvement comes from the signalling effect, while 55% comes from the batching effect. This is an important result because other food relief organisations have been hesitant to adopt the Choice System, deeming a fake money mechanism too complex and too risky. However, in this setting, the artificial currency is predominantly used to facilitate signalling. So, I then simulate a signal free mechanism that offers food out in batches, and find that it is able to capture 62% of the welfare benefits of transitioning from the Old System to the Choice System.

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 30% 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. There is an extensive literature empirically analysing the allocative effects of centralised school choice and medical residency matching (see, for example, Agarwal and Somaini (2020) and Agarwal (2015)). 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 the relative importance of various design features. Furthermore, in estimating a structural model I am able to explicitly estimate food banks' demand functions and so model counterfactual behaviour under the Old System. I do not need to assume that food banks would have essentially accepted all the food they were offered, biasing results in favour of the Choice System. Meanwhile, by employing a dynamic multi-object auction model, capturing more intricacies of the true allocation process, I am able to consider additional counterfactual mechanisms that appear equivalent using Prendergast's consumption based approach: For example, despite both being 'market mechanisms', simultaneous auctions and sequential auctions yield very different welfare implications due to the batching and exposure effects.

Relative to static allocation settings, the literature on empirical dynamic allocation systems remains relatively nascent. Papers studying other dynamic 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 (2023), which combines the models of Jofre-Bonet and Pesendorfer (2003) and Gentry et al. (2023). 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. I extend this literature by allowing bidder's state variables to be unobserved. In empirical auction studies it is standard to impute bidders' state variables, such as backlogs of contracts or stocks. However this is known to introduce non-classical measurement error, biasing estimates towards values being unresponsive to state variables (Raisingh, 2020).

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 states.

1.2 Overview

Section 2 describes the institutions and data being studied. Section 3 provides suggestive evidence of heterogeneity and presents 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 presents the estimation results. Section 7 details the counterfactual mechanisms and the simulation results.

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 retailers and manufacturers. Food banks stored this food until it was collected by local food pantries, who were then responsible for distributing it among those in need. Food banks are well distributed across the United States, but differ in both the number of food pantries in their catchment areas as well as local food insecurity rates. This then drives heterogeneity in the demand for their services. Similarly, because of regional variation in food production, food banks also face regional variation in the types of food they can solicit from local donors. So, as the food bank network expanded, Feeding America began directly soliciting donations they could then allocate centrally, helping ensure that food banks all had some access to food from further afield than just their local donors. See Appendix B for additional information on the geography of Feeding America's food banks.

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. As highlighted in Prendergast (2022), 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 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, so, receive more food. Shares are redistributed each night, and 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, meaning they can actually receive shares when they win. This incentivises food banks to accept undesirable loads, helping Feeding America maintain good relations with their donors by ensuring food is rarely turned away. 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

---

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

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.

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]

---

[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 discriminatory auctions distinct from the simultaneous auctions — the only difference is that one cannot lose a unit for a high bid yet win for a low bid. Meanwhile, food sold by food banks only makes up 4.5% of lots. The food banks who rely most on Feeding America (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. Consequently I ignore the decision to sell food.

[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. That said, the unobserved auction timing really only impacts the model by introducing measurement error into estimating equation 4: When bidding on lot l the food bank has to account for the probability they also win lot m. If the two lots were auctioned at different times, then either this probability has already been realised or was certainly zero if the lot was not yet available. The instrumental variable procedure I discuss in Appendix I.2.2 shows evidence of negligible bias.

[7] Refrigerated includes anything stored in a fridge or freezer, such as meat and dairy. Tinned

Figure 1: Descriptive Statistics, across lots

	Dried	Tinned	Fridge	Non-food	Total
	mean (sd)				
Daily lots	10.74 (8.35)	10.83 (10.8)	5.11 (3.76)	3.3 (3.21)	32.74
Tons per lot	11.25 (4.85)	17.15 (4.15)	14 (5.1)	10.2 (6.1)	13.25 (5.45)
Winning bid	2106 (5329)	1085 (6414)	2688 (6176)	2967 (6436)	2134 (5818)
No. bidders	2.95 (3.14)	2.7 (3.5)	2.5 (3.12)	3 (3.26)	2.8 (3.21)
% Allocated	93	83	91	91	91
% Negative prices	35	47	29	28	34

Note: Total includes mixed loads (14% of lots). Winning bids includes the reserve if no bids are received. 'Allocated' gives the percentage of lots receiving at least one bid. Negative prices include loads allocated for 0 shares.

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]

---

and Bottled food includes products with long shelf-lives that are tinned or bottled, including baked beans and bottled water. Dry food captures long shelf-life food such as cereal and cookies. Non-food items includes non-edible items, predominantly cleaning products and baby food. See Appendix A for additional details.

[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 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 evaluated across food banks. The mean Goal Factor is normalised to 1. Population and Poverty figures refer to the people in a food bank's catchment area.

2.2.2 Auxiliary Data

I construct food banks' Goal Factors using the formulae in Prendergast (2022), publicly 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 3: Heterogeneity in Lots

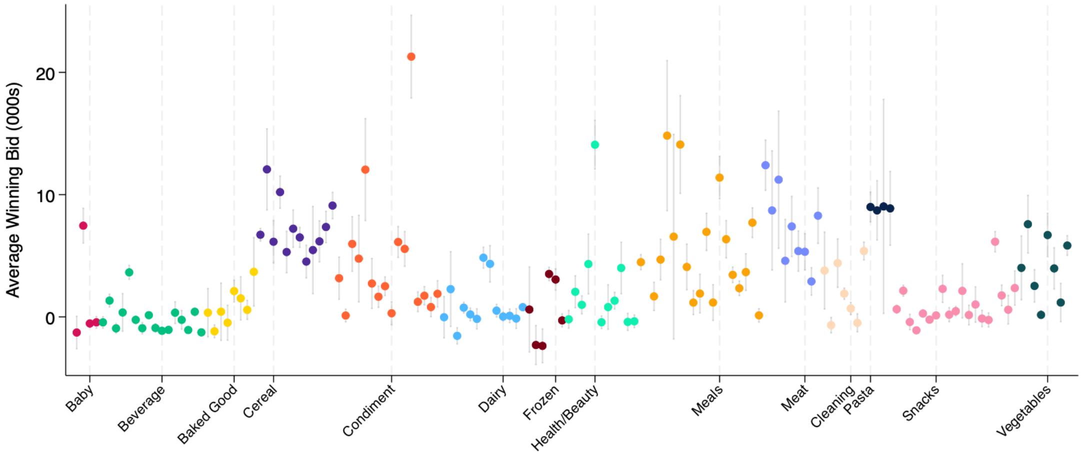

Note: Plots average winning bids across subcategories (subcategory fixed effects), accounting for censoring and composition. Estimates coloured according to category, with 95% confidence intervals.

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. Therefore heterogeneity is a key determinant of the value of choice. To investigate heterogeneity I demonstrate how bidding behaviour varies over types of food, over food banks, and within food banks over time. I consider several reduced form Tobit regressions, controlling for lot composition, distance, and the censoring caused by bidders' entry decisions.

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 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. The five food banks choose 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.

Next, I consider how bidding behaviour varies across food banks and across different types of food according to how they are stored. I use a Tobit regression of food bank × storage type fixed effects against bids, accounting for 'censoring' when bidders choose not to bid. Results are plotted in Figure 4. 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, but allow average bids to vary across months. I only include food banks who win at least 100 lots over my sample period. I then investigate variation in average bids over time. A likelihood ratio test that parameters are constant is rejected at 5% significance level for 98% 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.

Figure 4: Heterogeneity Across Food Banks

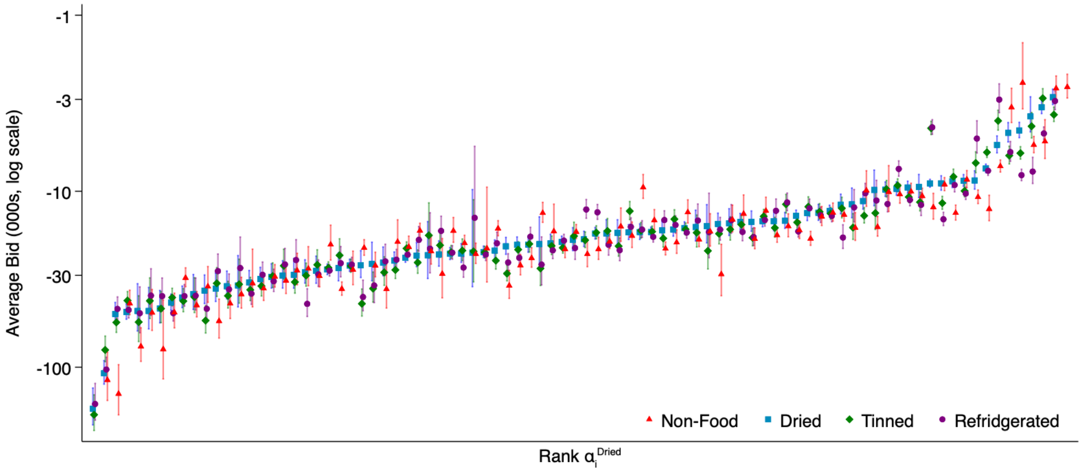

Note: Plots coefficients and 95% robust confidence intervals from a Tobit regression of food bank × storage type dummies on bids, accounting for distance and censoring/non-bidding. Food banks are ordered by estimates for Dried food (blue points monotonically increase from left to right). Includes food banks who won at least 50 lots.

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. This implies negative values in 98% of bidder × lot combinations. Negative valuations likely occur because of limited storage capacity or storage costs, as emphasised in Prendergast (2017). They cannot throw away excess food as this sends a bad signal. These storage costs include the opportunity cost of volunteer time and effort from repackaging and properly storing the food. Transportation costs also factor into these negative valuations.

3.2.2 Dynamic Complementarities

Figure 5 panel (A) demonstrates that, conditional on winning food of a particular storage type at time 0, the probability of bidding on lots of the same type falls by 35% on subsequent days. Therefore, as food banks win more of a particular type of food, the less they are willing to pay for an additional lot of that same storage type. The blue series shows that after winning one type of food we see only a 3% drop in bidding probability on different types, suggesting that this observation is not driven

Figure 5: Evidence of Dynamic and Static Complementarities

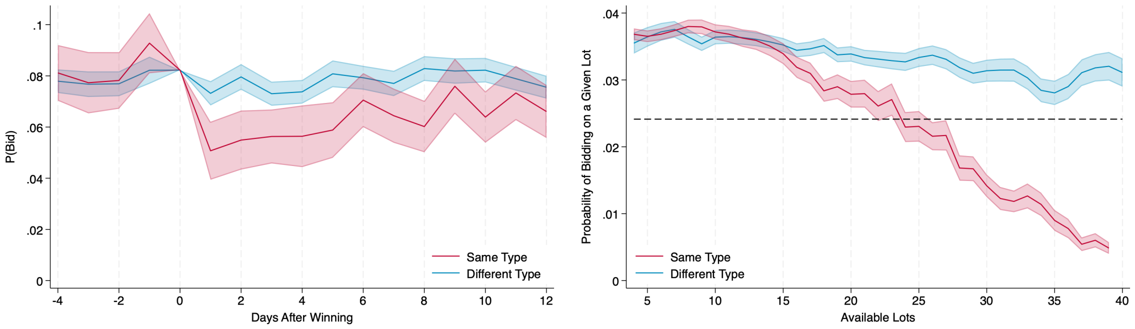
(A) (B)

Note: Panel (A) plots the probability of bidding on a particular type of food, conditional on winning a load of that same storage type at time zero. The blue series plots the impact of winning a different type. The probability at zero is set at the average for scale. Panel (B) plots the probability of bidding on a given lot, conditional on the number of lots available from the same storage type (versus different types, in blue). The dotted line is the unconditional bidding probability. Both plots include food bank × storage dummies, and condition on the type being available.

by transportation costs or budget constraints.

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. In Appendix B I repeat this analysis using alternate categories, splitting food by usage, and find that this relationship only occurs when food is split by storage type. This lends support to the importance of storage costs in this environment.

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 those lots. In particular, a 1 *sd* increase in lots decreases the probability of bidding on any given lot by 18% relative to the unconditional mean. Furthermore, they can always bid zero or negative, meaning this is not explained by binding budget constraints. If payoffs were additively separable we would see a horizontal line, even when payoffs are correlated. Instead, 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 payoffs as additively separable, and must instead take a multi-object approach, accounting for the simultaneous auction environment. The blue series demonstrates a significantly weaker relationship when considering how bidding varies with the total number of lots.

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. I assume there are no entry costs, since food banks simply have to log in to the platform, view the available lots, and decide how to bid. Each player bids an L dimensional vector, 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 $d_{itl} = 1$. Winners are announced simultaneously. Winners pay their bids, and every player observes the identities and bids of winners.

Lots are characterised by a vector of characteristics $\mathbf{c}_{tl}$, consisting of the location, size, categories (c), subcategories (h), and storage method (g) of the lot. The quantity of food in each lot by 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 stocks of each category increases by $\mathbf{z}^c_{tl}$. For notational convenience I absorb these variables into the common state variable $\mathbf{s}_{0t}$. I then assume that supply is Markovian:

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

I can allow supply to depend on outcomes from the previous round of auctions. However, this is not observed in practice, partly because donors do not observe prices or bids. They only observe whether their donation was successfully allocated, which virtually all donations are, either that day or the day after. However, Prendergast (2017) reports that after implementing the Choice System food banks started accepting more food, and that this increased donations. So, this assumption may not hold in the long run. Unfortunately, because I only have data from the Choice System, I cannot allow for supply endogeneity with respect to the allocation mechanism.

4.1.2 Primitives: States and Transitions

Food bank i begins the period in state $\mathbf{s}_{it} \in \mathbb{S}$, representing their current stocks of food. I primarily focus on their stocks from each storage method, so that the individual state has 4 dimensions.[10] This captures the dynamic costs of holding storable goods. Importantly, stocks are not observed by the econometrician. If the combinatorial 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. $\mathbf{s}_{it}$ is observed each morning before items are posted on the Choice System, and depends on the stocks and winnings from the previous period. 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) $E[\mathbf{s}_{it}] = 0$

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. Day-to-day variation in stocks is expected to be 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. 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

---

[10] I also use their stock of each subcategory h to capture food banks' preferences over how the food is used. However, I will assume payoffs are affine in subcategory stock (not subject to diminishing returns as they always have people to feed), so the level of these stocks are neither identified nor welfare relevant up to normalisation.

in stocks are identified. Any true long run average of stocks is not separately identified from π_i, so is set to zero. ‘Negative’ stocks should be interpreted as just being below this long run average.

4.1.3 Primitives: Payoffs

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

Assumption 3. *(i) If food bank i ends with stocks $\mathbf{s}_{it}$, they receive payoff $\pi_i(\mathbf{s}_{it})$. The deterministic function $\pi_i : \mathbb{S}_i \to \mathbb{R}$ is bounded, with $\pi_i(0)$ 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_i captures both the costs of storing food, and the utility from holding various types of food to be able to distribute them to clients. The assumption that π has finite range is predominantly for mathematical convenience, while the normalisation is required as only marginal payoffs are identified.

Part (*ii*) embeds two assumptions. Assuming the privately known $\boldsymbol{\upsilon}_{it}$ is conditionally independent across individuals imposes independent private values and no unobserved auction level heterogeneity. This is reasonable since my data contains the same information shown to food banks. Assuming conditional independence across time is standard in most dynamic models. The mean independence assumption is required to separately identify F_i^{υ} and π_i, where the mean of $\boldsymbol{\upsilon_{it}}$ is treated as an element of π_i.

Quasi-linearity is standard in empirical auction studies. However, with an intertemporal budget constraint, quasi-linearity can be thought of as imposing that the marginal value of wealth, λ_{it}, is constant over time. This means that the value of an extra share today is the same as an extra share far in the future; food banks are sufficiently patient as to recognise that day-to-day fluctuations in budgets or stocks do not significantly impact expectations about how valuable access to the Choice System will be over the long run. This argument is well known in the consumption literature and in Appendix C I formalise this argument, proving that if the equilibrium process exhibits weak dependence ($\mathbf{s}_t$ and $\mathbf{s}_{t+\tau}$ are independent for sufficiently large τ), then quasi-linearity is observationally equivalent to the budget constraint. In support of this argument I also demonstrate that bidding behaviour does not generally depend on budgets, with even credit constrained food banks rarely going near the constraint. Finally, I allow the marginal value of wealth λ_i to vary across food banks, capturing that some food banks rely on the Choice System more than others.

I assume players have temporally additively separable payoffs, and make forward looking decisions with discount parameter $\beta = 1$, so that food banks are extremely patient.[11] I assume risk neutrality, and common knowledge of F, π, β, and $\mathbf{s}$.[12]

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, 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.

---

[11] This was motivated by conversations with food bank managers, and because the interest rate is zero. $\beta = 1$ means the value function, an infinite sum, does not converge. Instead, we need convergence in the difference between the value function at any two states, requiring the same weak dependence as Assumption 3 part *iii*).

[12] Assuming food banks observe, and track, each others stocks is a strong assumption, particularly given the large number of food banks. Likewise, whether they would keep track of each others' bids and winnings. In section 5.1 I assume, and test, that the market is sufficiently large and competitive that variation in any individual food banks' stocks do not shift the equilibrium distribution of winning bids, similar to an Oblivious Equilibrium (Weintraub et al., 2008). Consequently, individual stocks will not matter, and food banks only need to keep track of aggregate statistics of the food supply, meaning the estimated model is robust to this assumption.

4.2.2 Value Function and Continuation Value

The Bellman equation is given by: $W_i(\boldsymbol{\upsilon}, \mathbf{s}; \pi, \Lambda_{-i}) = \max_{\mathbf{b},\mathbf{d}} \left\{ \bar{W}_i(\mathbf{b}, \mathbf{d}; \boldsymbol{\upsilon}, \mathbf{s}, \pi, \Lambda_{-i}) \right\}$, Where $\bar{W}_i(\mathbf{b}, \mathbf{d}; \boldsymbol{\upsilon}, \mathbf{s}, \pi, \Lambda_{-i}) =$

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

The continuation value gives the expected payoff from ending the period in state $\mathbf{s}^a$: $V_i(\mathbf{s}^a; \Lambda_{-i}) = \int_{\tilde{\mathbf{s}}} \int_{\tilde{\boldsymbol{\upsilon}}} W_i(\tilde{\boldsymbol{\upsilon}}, \tilde{\mathbf{s}}; \pi, \Lambda_{-i}) dF_i^{\boldsymbol{\upsilon}}(\tilde{\boldsymbol{\upsilon}}|\tilde{\mathbf{s}}) dF^{\mathbf{s}}(\tilde{\mathbf{s}}|\mathbf{s}^a)$.. A further important object is the sum of the deterministic payoff function and the discounted continuation value, denoted $\kappa_i(\mathbf{s}; \Lambda_{-i}) = \pi_i(\mathbf{s}_i) + \beta V_i(\mathbf{s}; \Lambda_{-i})$. Define κ as the 'Pseudo-Payoff' function. This is essentially what we estimate if we incorrectly assume myopic bidding. Estimating κ 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:

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

4.3 Equilibrium

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

Definition 4.1. : *A symmetric MPE consists of strategies $\mathbf{\Lambda}^*$ and beliefs $(\mathbf{\Gamma}^*, \mathbf{P}^*)$ such that for all i and any $(\boldsymbol{\upsilon}, \pi, \mathbf{s})$: (Optimality) Λ_i^* maximises the present value of payoffs given beliefs, (Consistency) Beliefs are consistent with $\mathbf{\Lambda}_{-i}^*$, (Markovian) $\mathbf{\Lambda}^*$ only depends on t through $\mathbf{s}_t$, (Symmetry) Agents with the same 'type' and beliefs place the same bids.*

This allows us to write equilibrium beliefs as a function of the state: $\Gamma_{il}(\mathbf{b}, \mathbf{d}; \Lambda_{-i}^*(\mathbf{s})) = \Gamma_{il}(\mathbf{b}, \mathbf{d}; \mathbf{s})$. Altmann (2023) proved that, conditional on existence of a symmetric Pure Strategy Bayesian Nash Equilibrium in the bidding game conditional on entry, such an equilibrium exists.[13] I make the following assumptions about equilibrium:

Assumption 4. *(i) The data are generated by a strategy profile* $\mathbf{\Lambda}^*$*, 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}*.*

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' stock processes are stationary.[14] This assumption does not impose the equilibrium is unique, instead that the same equilibrium is played throughout the data period. 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.

4.4 Identification

Altmann (2023) proves that π_i and F_i^v 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_i, as well as the transition process for the unobserved variable $F_i^{\mathbf{s}}$, for each food bank. 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

---

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

[14] I require that the distribution of local donations and food sent to clients is constant over 2014-2017. Feeding America's 'Hunger in America' resource shows that food bank usage and food insecurity is stable over this period.

stock process, from which Altmann (2023)'s results apply. I require three 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 $\mathbf{w}_{it}$ *and any real function* g*:* $E[g(\mathbf{s}_{it})|\mathbf{w}_{it}, \mathbf{s}_{it+1}] = 0$ *implies* $g(\mathbf{s}_{it}) = 0$*.*

(iii) For any $\mathbf{s}_{0t+1}$ *there exists a pair* $(\mathbf{w}_t, \mathbf{s}_{0t})$ *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 are standard 'completeness' conditions, the non-parametric analogue of rank conditions required in linear 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, given winnings $\mathbf{w}_t$, the distribution of stocks at $t+1$ can be inverted for the distribution of stocks at t. This will hold if, for example, $\mathbf{s}_{t+1}$ is monotonically increasing in $\mathbf{s}_t$, which seems reasonable. Part *iii*) 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 *iii*) 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 D. The proof is omitted for brevity as it requires formal definition of linear operators and their spectral decomposition. The proof builds on 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, as well as weaker normalisations because of how 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 they will frequently bid 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 only bid on a few of them. They do not want to win too many loads or face excessive storage costs. This is precisely the behaviour seen in practice in Figure 5 panel (B): When there are 10 lots of dried food available food banks bid on a given lot with 3.5% probability, placing on average 0.35 bids per round. But, when there are 35 lots available they only bid with around 1% probability, still only placing 0.35 bids each round. We then 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 relationship is 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 this load out quickly and immediately need more. On the other hand, if they take several days to give out this cereal then they may stop bidding on cereal and similar dried foods while their stocks are high, to avoid increased storage costs. Meanwhile, 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 more cereal out to clients than usual, and request less from donors. Their stocks quickly return to normal levels, and so bidding behaviour also bounces back quickly. Whereas, if they have only limited control then they are unable to shift the extra cereal any faster than usual, keep receiving the same 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 extends Altmann (2023) to allow for unobserved stocks. Section 5.1 discusses parametrisation and estimation of beliefs. Section 5.2 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 5.3 I discuss disentangling the flow payoffs $\pi_i s$ and the discounted continuation value from the pseudo-payoffs κ_i. See Appendix G for additional details.[15]

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

Instead, Altmann (2023) 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-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: $\kappa = \pi + \beta V$.[16] To estimate the pseudo-payoff function we begin by essentially estimating a misspecified static model. While Altmann (2023) proposed a non-parametric procedure, this is inapplicable when stocks are unobserved and so I take a parametric approach.

---

[15] Due to computational requirements I focus my analysis on the 90 food banks that each won at least 50 lots, consuming 94% of the food from the Choice System. Because heterogeneity is an important theme 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.

[16] In single-object environments this procedure reduces to Jofre-Bonet and Pesendorfer (2003), while for discrete choice problems it is numerically equivalent to Hotz and Miller (1993). This is because, when policy functions are invertible there is a one-to-one relationship between the policy function and the pseudo-payoff function.

5.1 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.[17] This assumption is tested on the empirical equilibrium winning probabilities in Appendix I.

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 I.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. The winning bid is just the maximum of conditionally independent variables, so I assume winning bids follow a generalised extreme value distribution censored at the reservation price:

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

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

---

[17] This assumption is similar to the large market Oblivious Equilibrium (Weintraub et al., 2008). 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 food banks tracks their competitors' states, or their bids and winnings.

the competitive pressures on prices: If little food has been auctioned recently then on average food banks' stocks will be low, and we expect higher prices. See Appendix G.1 for additional details of how I estimate beliefs.

5.2 Step 2. The Pseudo-Static Model

I now describe the second estimation step, in which I jointly estimate $F_i^{\mathbf{s}}, F_i^v$, and κ_i. The central difficulty concerns the unobserved state and the bids that are unobserved when food banks do not enter certain auctions. I employ a Gibbs Sampler: I use data augmentation to iteratively sample unobserved bids and unobserved states from their conditional posterior distributions, before updating parameters given the augmented data. Full details of the procedure, including assumptions on prior distributions, are given in Appendix G.2.

5.2.1 Stock Process

Stock transitions follow $\mathbf{s}_{it+1}^g = \mathbf{s}_{it}^g + \mathbf{w}_{it}^g + \mathbf{x}_{it+1}$, where $\mathbf{x}_{it+1}$ gives the daily local donations minus food taken by clients. I assume $\mathbf{x}_{it+1}|\mathbf{s}_{it}^g \sim N(\boldsymbol{\delta}_i[\mathbf{s}_{it}^g + \mathbf{w}_{it}^g] + \boldsymbol{\mu}_i, \Sigma_i)$. This process incorporates only simple feedback from food banks' stocks to net donations, as this process is difficult to identify. $\boldsymbol{\delta}_i$ controls the responsiveness of their net donations to previous stocks, while $\boldsymbol{\mu}_i$ controls the unconditional average. I impose that $\boldsymbol{\delta}_i$ is diagonal, with entries $\delta_{ig} \in [-1, 0]$ to ensure stationarity. When $\delta_{ig} = 0$ 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.[18] This stock process is therefore essentially an Auto Regressive process, with coefficient $(I + \boldsymbol{\delta}_i)$, drift $\boldsymbol{\mu}_i$, and noise $\tilde{\mathbf{x}}_{it+1} \sim N(0, \Sigma_i)$.

Unlike other storable goods models I assume stocks do not exogenously decay over time.[19] Seasonality is indirectly captured through variation in stocks over time, as

---

[18] 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 to larger welfare benefits of choice as food banks can focus their consumption on the types of food they cannot so easily procure from local donors.

[19] I cannot separately identify decay parameters from δ. However, most of the donated food, even fresh produce, have long shelf lives, meaning daily decay parameters are close to 1 anyway.

net donations may be lower in winter when demand for food banks is high.[20] 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 × storage type specific feedback, mean and variance parameters $(\delta_{ig}, \mu_{ig}, \Sigma_{ig})$.

5.2.2 Lot-Specific Payoffs

I assume $v_{ilt} \sim N(\alpha_i distance_{ilt}, \sigma_l^2)$, 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 assume they are also conditionally uncorrelated across lots.

5.2.3 Combinatorial Payoffs

To make estimation feasible I fit a quadratic form to the pseudo-payoff function $\kappa_i(\mathbf{s}_i, \mathbf{s}_0)$, ensuring that inverse-demand is affine in the unobserved stocks. This is similar to standard assumptions of quadratic storage costs, or opportunity cost of storage. Non-linear demand would require multidimensional particle filters, exponentially increasing computation requirements. Within these constraints, I choose a functional form to reflect how food banks benefit from food based on how it is used (depending on the subcategory) and how they face costs of storing it (depending on the storage method). I impose the following:

$$\kappa_i(\mathbf{s}_i, \mathbf{s}_0) = \Phi_i \mathbf{s}_i^h - \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g \tag{3}$$

Where $\mathbf{s}_i^g$ gives the food bank's stock of each storage type and $\mathbf{s}_i^h$ their stock of each subcategory. Φ_i is an 1×152 row vector and Ψ_i is a 4×4 dimensional matrix that I constrain positive, imposing gross substitutes. This form imposes that the benefit of holding food to give out to clients, which comes through $\mathbf{s}_i^h$, does not exhibit decreas-

---

[20] 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.

ing 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. κ_i should depend on $\mathbf{s}_0$, however in the main specification I impose that it only depends on their own stocks $\mathbf{s}_i$.[21]

5.2.4 Estimation algorithm

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 E. This gives us the following three equation model:

$$\mathbf{s}_{it}^g = (I + \boldsymbol{\delta}_i)[\mathbf{s}_{it-1}^g + \mathbf{w}_{it-1}^g] + \boldsymbol{\mu}_i + \tilde{\mathbf{x}}_{it} \qquad \rightarrow \textbf{Transition Eq.}$$

$$\lambda_i y_{ilt} = \Phi_i \mathbf{z}_{tl}^h - \mathbf{z}_{tl}^{gT} \Psi_i (\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2\sum_{m \neq l} \Gamma_m(b_{itm}) \mathbf{z}_{tm}^g) + \upsilon_{ilt} \qquad \rightarrow \textbf{Observation Eq.}$$

$$y_{itl}^* = \begin{cases} b_{itl} + \frac{\Gamma_l(b_{itl})}{\nabla_b \Gamma_l(b_{itl})} & if \;\; b_{itl} > R \\ R + \frac{\Gamma_l(R+1)}{\Gamma_l(R+1) - \Gamma_l(R)} & if \;\; d_{itl} = 1, b_{itl} = R \\ R & if \;\; d_{itl} = 0 \end{cases} \qquad \rightarrow \textbf{Censoring Eq.} \tag{4}$$

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 $\mathbf{s}_{it}^g$.[22] The likelihood for this model intractable. Instead, estimation is performed using a Gibbs Sampler, which consists of repeating the following steps:[23]

---

[21] κ_i should capture how bidding depends on beliefs about future prices and supply. However, the money supply varies with food supply to keep prices roughly constant. That said, relative prices may vary, so in section 6.1, I show that the relationship between supply of different types and their prices is not economically significant. Nonetheless, in Appendix I.2 I include the demand indices ϑ_{ltg} as inputs to κ, finding these have no effect.

[22] 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 m. Results in Altmann (2023) 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 I.2 I use an instrumental variable procedure to demonstrate evidence of negligible bias.

[23] The Bernstein-von-Mises Theorem ensures the posterior mean of this sampler is asymptotically equivalent to the maximum likelihood estimator. Recognise how this procedure employs the identification argument presented in 4.4. In step 3. I use variation in winnings and the effect on bids

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

5.3 Step 3. The 'Dynamic' Game

Next, given draws of beliefs, $\{\kappa_i, F_i^v, F_i^{\mathbf{s}}\}_N$, and $\{\mathbf{s}_{it}^g\}_{T,N}$ from their posterior distribution, I use the following proposition to evaluate the continuation value $V_i(\mathbf{s}_i^g, \mathbf{s}_0)$:

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

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

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

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

I prove this proposition in Appendix F. The identity $\tilde{W}_i(\mathbf{b}, \mathbf{d}|\mathbf{s}_i, \mathbf{s}_0)$ follows from substituting the Inverse Bid System into the maximand, writing the value function as a function of bids and pseudo-payoffs. The main proof extends results from Arcidiacono and Miller (2011) to continuous choices. The sample counter-part to this object is then easily found. See Appendix G.3 for full details. I evaluate the value function across a 30^4 grid of stocks, evaluated evenly across points from the posterior sampled states, before fitting a quadratic approximation. In Appendix I.3 I evaluate this approximation using the R^2s of the approximation regressions: 100% of these R^2s lie between 0.9 and 1.

---

to infer changes in stocks, pinning down $F_i^{\mathbf{s}}$. In step 4. variation in $\mathbf{z}_t$, and how this impacts bids, pins down κ.

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

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 H. 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 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 has a significant negative effect on prices for most types of food. While statistically significant, these results are not economically significant: Each increase in the previous 30 day supply by one hundred loads ($\approx 1sd$) decreases the winning bid by only 0.017 standard deviations. Given the money supply varies with the food supply to ensure prices remain stable over time, this is unsurprising.

6.2 Second Stage Results

Results are presented in Figure 6. I do not report parameters separately for every food bank, instead I present three measures: A) The posterior means, averaged across food banks, B) The posterior standard deviations, averaged across food banks, and

C) The standard deviation across food banks of the posterior means. Measure C highlights the heterogeneity in the parameters across food banks, while measure B highlights the extent of the sampling variation. In general, the degree of heterogeneity across food banks is large.

Estimates of feedback parameters $\boldsymbol{\delta}$ are heavily skewed — While the mean posterior means are around -0.25, around half are estimated to be above -0.1. This suggests most food banks have little control over their net donations, and that their stocks exhibit a large degree of persistence. While precisely estimated, only 18% of the $\boldsymbol{\mu}_i$ parameters are significantly different from zero, predominantly for food banks who are observed almost always bidding on the same types of food.[24] Estimates of $\sqrt{\Sigma_i}$ are large, heterogenous, and also skewed. 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^v and $F^{\mathbf{s}}$, around 93% of the unexplained variation in bidding behaviour is from variation in unobserved stocks.

The average transport cost is 39.7 shares per mile, and varies significantly across food banks from 5 to 120. These are higher than Prendergast (2022)'s figure of around 0.16 shares per mile. However the figures are not directly comparable: Prendergast takes this figure as the coefficient from a regression of distance on the observed winning bid. When I perform this exercise I get 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.

Estimates of Ψ_i and Φ_i are presented in Figure 6, averaging Φ_i across subcategories with the same uses. I estimate significant heterogeneity in 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 (± 2,400). This varies from -71,500 to -14,800 across different 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 demand elasticities for food given the specific setting of large food banks using fake money to bid on food, rather than

---

[24] I also correlate estimates with observable characteristics of food banks, such as population density and local agricultural rents. This analysis is omitted as I do not find any particularly striking results.

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 from local donors.

6.3 Third Stage Results

Parameters from the quadratic approximation of the flow payoff function π are displayed in Figure 6. These storage cost estimates can be interpreted as follows: For the average food bank with the long run average stocks, receiving a ton of Dry food yields a payoff equivalent to receiving an extra 180 shares. Positive marginal payoffs suggests food banks benefit from not having an empty warehouse. However, every additional ton of Dry food increases storage costs by 53.8. These numbers are smaller in magnitude than estimates of willingness to pay and pseudo-payoffs, which are forward looking objects.

Storage costs for Dry food are on average larger than other types of food due to space constraints — dried food such as cereal or crackers tend to be light but bulky. There is again evidence of a lot of heterogeneity in these storage costs across food banks, with the standard deviation in posterior means across food banks similar in magnitude to the mean posterior mean. Once more, we see the distribution is heavily skewed, with a small number of food banks facing very large storage costs for certain types of food. We can also reject that different types of food are perfectly substitutable in terms of how they are stored.

6.4 Model Fit

Figure 7 displays true and simulated 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 per day. I slightly under predict the probability of bidding on any given lot and the average distance each load travels. Figure 8 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

Figure 6: Parameter Estimates

Parameters										
Stock Process			Storage Type							
			Dry		Tin		Fridge		Non-Food	
δ_{ig}			-0.353	0.0798 0.326	-0.267	0.0716 0.3	-0.27	0.0814 0.282	-0.19	0.0535 0.266
μ_{ig}	A	B C	-0.0945	1.17 2.35	-0.0415	1.55 4.19	0.114	0.768 3.33	1.44	1.1 4.53
$\sqrt{\Sigma}_{ig}$			6.54	1.15 6.12	7.1	1.55 8.69	5.1	1.13 5.91	6.82	1.36 6.59
Pseudo-Payoff			Usage Type							
	Meal		Ingredient		Snack		Condiment		Non-Food	
Φ_i	2420	312 1070	2280	305 987	2350	366 1130	2050	341 1130	1880	414 977
			Storage Type							
			Dry		Tin		Fridge		Non-Food	
Ψ_i	Dry		-176	12.2 68.8	-120	10.4 46.5	-164	11.4 56.4	-107	12.5 51
	Tin		-		-89.9	9.66 39.1	-108	9.38 43	-74.4	10.3 36.9
	Fridge		-		-		-139	11.9 44.2	-101	11.7 40
	Non-Food		-		-		-		-141	14.1 49.3
Flow-Payoff	Constant		180	124 271	140	88.9 225	222	122 317	319	121 479
π_i	Dry		-53.8	10.2 43	-38.4	7.47 29.8	-54.6	9.3 36.2	-30.6	6.21 23
	Tin		-		-21.3	5.56 22.9	-31.3	6.26 24.2	-18.3	4.36 17.2
	Fridge		-		-		-34.3	9.5 29	-27.6	6.4 21.8
	Non-Food		-		-		-		-26.6	6.21 32.5

Note: For each parameter the table presents 3 measures: $A = E_i[E_r[\theta_{ir}]]$, the posterior mean (over draws r) averaged across food banks, $B = E_i[\sqrt{V_r[\theta_{ir}]}]$, the posterior standard deviation for the parameters averaged across food banks, and $C = \sqrt{V_i[E_r[\theta_{ir}]]}$, the standard deviation across food banks of the posterior means.

Figure 7: Model Fit

Measure	Mean				Std			
	True	Mean	q0.025	q0.975	True	Mean	q0.025	q0.975
Average Bid	1950	1860	1720	2040	4290	4350	4270	4470
Pr. Bid	0.0275	0.0239	0.0232	0.0252	0.164	0.153	0.151	0.157
No. Bids	0.676	0.683	0.664	0.721	1.66	1.33	1.3	1.36
Dist. Travelled	503	467	459	478	397	367	360	379

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

true data. In Appendix H I present Gelman-Rubin convergence statistics. Generally, the data converged well.[25]

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.

---

[25] 100% of first stage parameters are below the standard threshold of 1.1, while 91% of second stage parameters are below 1.1 and 94% are below the more lenient 1.2 threshold. The model did not converge for 5 food banks, which I removed due to implausible parameter estimates. These only consume 1.5% of the food on the system.

Figure 8: Model Fit

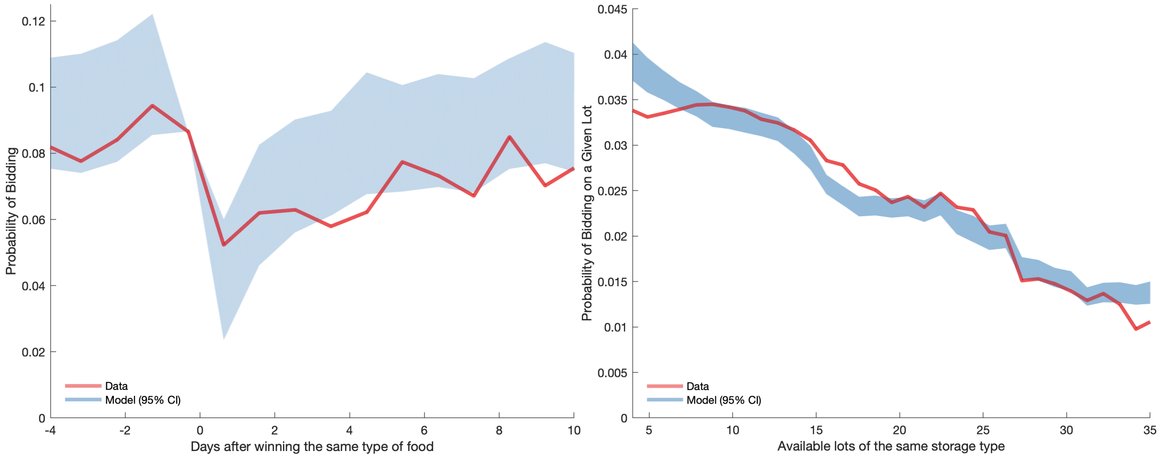

(A) (B)

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

7.1 The Old System

As discussed in Section 2, under the Old System each load was offered to food banks sequentially along a queue. Their position in the queue was determined by their 'Goal Factor' and how recently they had been offered food. Food banks could receive an offer from Feeding America at any point in time, or they could receive a shipment of food from local donors, or send a shipment out to clients.

With this in mind I model the Old System in continuous time, ensuring that multiple of these events occurs simultaneously with probability zero. Because food banks did not observe offers made to anyone else, nor their place in the queue, I assume they form beliefs over the poisson rate they are offered food by Feeding America as a function of the time since their last offer τ: $p_i(\tau)$.[26] Lots and lot characteristics are drawn from the empirical distribution. Denote i's Value Function under the Old System as $\tilde{V}_i(\mathbf{s}_i; \mathbf{s}_0, \tau)$. Upon being offered lot l, food banks accept if and only if $v_{ilt} + \tilde{V}_i(\mathbf{s}_i^g + \mathbf{z}_l^g, \mathbf{s}_0, 0) \geq \tilde{V}_i(\mathbf{s}_i^g, \mathbf{s}_0, 0)$.

---

[26] This may also depend on common state variables $\mathbf{s}_0$, such as the recent supply of lots. However, unlike under the Choice System, food banks do not observe how much food has been donated recently, so can only base inference on τ. Furthermore, equilibrium offer rates are so high that the value function is insensitive to τ, so I impose that p_i is independent of τ in the end anyway.

The model primitives π and $F^{\mathbf{s}}$ estimated above are defined for discrete intervals, and I require continuous time analogues to solve for the food banks' Value Functions. I calibrate the continuous time flow payoff $\tilde{\pi}$ such that food banks receive the same flow payoff as in the main model if stocks remain constant across a period. Likewise, I assume net donations arrive at constant rate q_i, then calibrate $\tilde{F}^{\mathbf{s}}$ to ensure that the distribution of net donations across a period is the same as in the main model. Finally, define $\rho = \frac{1-\beta}{\beta}$ as the discount rate. $\tilde{V}_i$ is then the solution to the Hamilton-Jacobi-Bellman equation:

$$\begin{aligned}(\rho + p_i(\tau) + q_i)\tilde{V}_i(\mathbf{s}_i, \mathbf{s}_0, \tau) =& \tilde{\pi}_i(\mathbf{s}_i) + \frac{\partial \tilde{V}_i(\mathbf{s}_i, \mathbf{s}_0, \tau)}{\partial \tau} + q_i \int \tilde{V}_i(\mathbf{s}_i', \mathbf{s}_0, \tau) d\tilde{F}^{\mathbf{s}}(\mathbf{s}_i'|\mathbf{s}_i) \\ &+ p_i(\tau) E_{\upsilon_{il}, \mathbf{z}_l} \left[\max \left\{ \upsilon_{il} + \tilde{V}_i(\mathbf{s}_i + \mathbf{z}_l, \mathbf{s}_0, 0), \tilde{V}_i(\mathbf{s}_i, \mathbf{s}_0, 0) \right\} \right]\end{aligned}$$

An equilibrium consists of Value Functions that satisfy the optimal control problem given beliefs p_i, as well as beliefs consistent with observed offer rates. I compute equilibrium by numerically finding the Value Function, holding beliefs fixed. I then simulate the Old System using these Value Functions and update the observed offer rates.[27] I repeat this process until beliefs converge. See Appendix J for additional details.

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 with the food supply to ensure prices are constant: If the food supply 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,

---

[27] 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.

Figure 9: Welfare

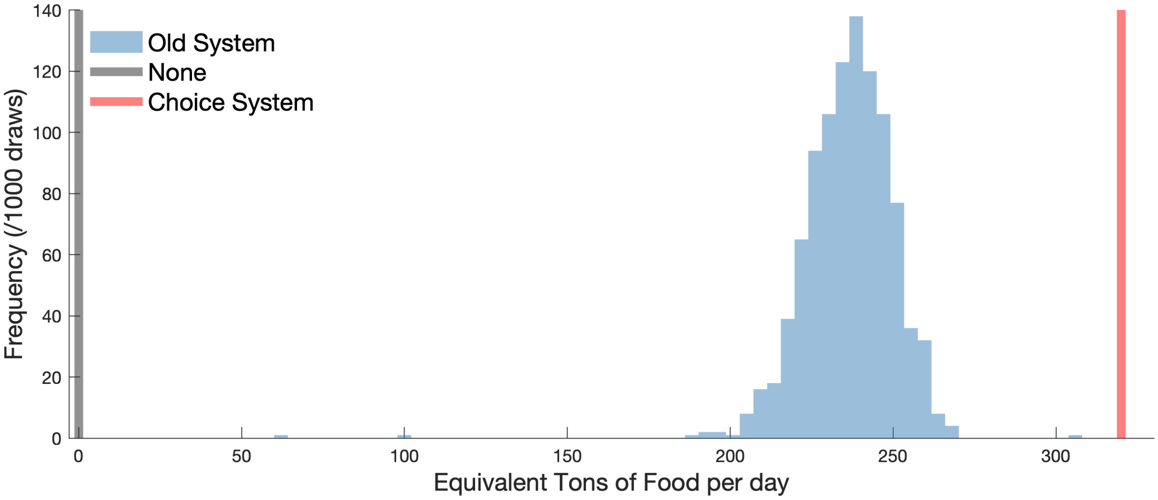

Note: This plot shows the posterior distribution of welfare under each mechanism, evaluated over 1000 draws from the posterior distribution of parameters. Welfare is measured relative to the Choice System and no allocation mechanism. On average, welfare is 36% higher under the Choice System.

which equates to double the supply of food.[28]

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.[29]

7.2 Results

Figure 9 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. Welfare is on average 36% 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 84 tons per day, which is enough to provide an additional 34,000 meals. These figures are statistically significant at the 0.1% level. When welfare is weighted by Goal Factor, this figure increases to 39% 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, in particular heterogeneity within food banks over time, it is unsurprising I find a larger effect.

Additional outcome measures are reported in Figure 11. Only 224 tons of food are successfully allocated each day, compared to 300 tons under the Choice System, so wastage is considerably reduced. Furthermore, food banks sort into consuming closer lots — the average ton of food travels 80 miles under the Old System, compared to 60 miles under the Choice System. Total storage costs are only 27% higher under the Choice System, despite 34% 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

Feeding America put significant resources into minimising possible negative distributional consequences of the Choice System. Figure 10 presents welfare results by

---

[28] This will actually be a lower bound on the value of the Choice System: Because payoffs are concave (storage costs are convex), in order to double consumer surplus it requires more than doubling the food supply.

[29] 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: Individual Welfare

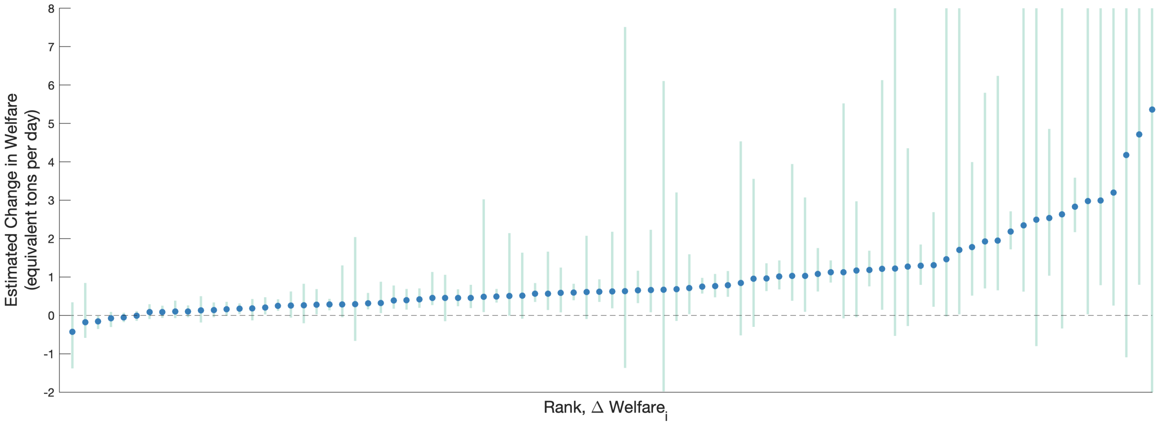

Note: This plot shows the difference in food bank specific welfare under the two mechanisms, ordered by the welfare difference and including 95% credible intervals. On average 89% of food banks are better off under the Choice System. 5 food banks have large credible intervals due to long left tails under the Old System.

food bank, plotting the difference between food bank specific welfare under the two systems. On average 89% of food banks are better off under the Choice System.

The 11% of Food Banks who are occasionally 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 accept essentially any food rejected by other food banks, benefiting from the bad market design. Otherwise, we see only small correlation between Goal Factor and welfare gain ($\rho = 0.197$). We do see that the food banks serving known food deserts are all significantly better off under the Choice System.[30]

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.

---

[30] We do not see any strong correlations for factors such as population density or deprivation indices. We do see some correlation ($\rho \in [0.08, 0.15]$) between these welfare differences and state transition parameters δ and Σ. Those with more uncertain net donations benefit more from choice. However, this is driven by two food banks with particularly high Σs observed performing very well under the Choice System.

First, I simulate two sets of mechanisms used in practice. I investigate the importance of allowing food banks to turn down food, as well as the impact of only offering food to the closest food bank, aiming to minimise transport costs.[31]

Next, I examine four of the most central differences between the Old System and the Choice System. *Offers*: Food banks can bid on any load, whereas under the Old System each load was only offered to around 10 food banks. *Signals*: Food banks can signal the intensity of their preferences. *Batches*: Food is allocated in batches, so that food banks have better information about all the food being offered that day. *Exposure*: Because food is allocated simultaneously, and food banks cannot place combination bids, they risk winning too many or too few loads. I simulate three counterfactual mechanisms aimed at teasing apart the importance of these different features. I then propose one final mechanism aimed at exploiting the most important features. Details of how I compute equilibria under these mechanisms are given in Appendix J.[32]

7.4.1 Random Allocation

Many food relief organisations pressure food banks to accept any food offered, attempting to minimise food waste. To examine the importance of allowing food banks to turn down food I consider an extreme mechanism: Randomly allocating food in proportion to Goal Factor. I only allow food banks to reject an offer if it is further from the food bank than 90% of lots, so that transport costs are in the top decile. In Figure 11 I show that welfare is equivalent to allocating just 64 tons (20%) of food each day under the Choice System. Allowing food banks to turn down food is extremely important. Furthermore, this is precisely the counterfactual exercise used in Prendergast (2022) to simulate the Old System, highlighting the importance of using a structural model to simulate optimal equilibrium decisions or risk underestimating

---

[31] 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.

[32] I also considered optimising budget allocations by accounting for different access to local donors, not only local poverty rates. I do this by equalising the marginal values of wealth $\lambda_i = \lambda$ across food banks. I ran these counterfactuals for both Sequential and Combinatorial Auctions. In both cases the welfare improvement was only marginal and insignificant, suggesting we would not expect to see significant improvement under the Choice System. This arises because, as I discuss below, the benefits of signalling are not the key driver of welfare.

counterfactual welfare.

7.4.2 Closest Offers

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. This mechanism is likely to be effective if transport costs are a suitably important margin. I also consider a version where food is offered sequentially in order of distance.

Offering food only to the nearest food bank is equivalent to distributing only 95 tons of food each day under the Choice System, a 70% 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 that extra distance. When food is offered to every food bank welfare increases to 258 tons per day, extremely similar to the Sequential Offer mechanism considered shortly, despite a 55% reduction in transportation costs. This is because it is always the same food banks being offered the same types of food.

7.4.3 Sequential Offers

Under the Old System loads were only offered to around 10 food banks. While infeasible due to time constraints, removing this constraint highlights the importance of ensuring food is offered to everyone, as under the Choice System. This mechanism is strategically equivalent to the 'Like' mechanism of Walsh (2015) used by Food Bank Local.

Welfare under Sequential Offers is 7.3% higher than under the Old System, equivalent to 15 additional tons donated each day. This suggests around 20% of the benefit from transitioning to the Choice System can be attributed to offering food to every food bank. The effect is partly mechanical: Donations that would have been returned instead go to someone further down the queue. However, food banks can also afford to be more picky in equilibrium, knowing they will not have to wait as long until they are next offered food.

7.4.4 Sequential Auctions

As we move from Sequential Offers to Sequential Auctions we can evaluate the importance of allowing food banks to signal the intensity of their preferences. This ensures food always goes to the food bank who values it most in that moment, rather than being offered out essentially at random. Results are presented in Figure 11. Moving from Sequential Offers to Sequential Auctions increases welfare by the equivalent of distributing an extra 21 tons of food per day, with a credible interval of (4.4,35). The welfare benefit from signalling increases to 32 tons per day $(10, 51)$ when welfare is weighted by Goal Factor.

Either way, this finding suggests that the signalling effect is only small, and is responsible for around 24% of the benefit from transition from the Old System to the Choice System. 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 because storage is now more costly. This result is evident from the data: The average lot receives only 3 bids. The relative importance of making sure the lot goes to the food bank with the highest value, out of only 3 bidders, is marginal.

7.4.5 Combinatorial Auctions

Comparing Sequential Auctions to a Combinatorial Auction, such as a daily combinatorial clock, allows us to evaluate the impact of batching. By allocating food in batches, rather than sequentially, food banks have information about all the food being allocated that day and so we expect better matches (Akbarpour et al., 2020). This is important because the majority of other food relief organisations distribute food sequentially. Meanwhile, unlike in simultaneous auctions in which food banks risk winning too much or too little food, a combinatorial auction eliminates the exposure effect by allowing food banks to express substitutability between lots (Gentry et al., 2023). The sequential mechanisms considered so far are also not subject to this exposure effect.

The gross substitutes assumption simplifies simulation of these auctions. Results are presented in Figure 11. Moving from Sequential to Combinatorial Auctions, welfare increases to the equivalent of 355 tons of food being allocated each day under the Choice System, 81 tons (30%) greater than Sequential Auctions, and 35 tons

(11%) greater than the Choice System. This implies 55% of the benefit of moving from the Old to the Choice System is driven by the shift to simultaneous allocation, comprised of a 96% batching benefit offset by a 41% exposure cost. Once more, this result is clear from the data: Even though only a few bidders bid on any given lot, 91% of lots still get allocated immediately. This means food banks are bidding on different lots, suggesting that horizontal heterogeneity is more important than vertical heterogeneity.

7.4.6 Simultaneous Offers

The mechanisms considered above are not feasible in practice, partly due to the burden of combinatorial bidding or needing to bid/respond to offers repeatedly throughout the day. Furthermore, the use of fake money as a signalling device has been cited as a barrier to adoption of the Choice System by other food relief organisations, believing these elements to be too risky. However, I have established that it is the batching element of the Choice System, rather than signalling, that is key in this environment. I now propose a Simultaneous Offer mechanism designed to exploit this feature.

As lots arrive they are added to a menu, which is then offered to the food bank at the head of a queue. They can take a lot from the menu, or not, and return to the back of the queue either way. The menu is offered down the queue until a lot is taken, ensuring that the menu is always around the same size, consisting of around 20 lots. This has the benefit that food banks choose from multiple lots at a time, but because food banks make their choices sequentially, it is not subject to the exposure effect. Details of how I simulate this mechanism are given in Appendix J. In brief, I assume food banks form beliefs over the expected maximum payoff from the menu, using an 'inclusive value' approach, and simulate the mechanism until beliefs match the observed distribution of payoffs.

I find that welfare under Simultaneous Offers is equivalent to distributing 289 tons of food under the Choice System, 22% greater than under the Old System and 62% of the way to the Choice System. This result is important as it demonstrates that even if food relief organisations are hesistant to fully adopt the Choice System, or similar systems involving fake money, they can still get most of the benefits using a relatively simpler mechanism. Lastly, this finding also demonstrates that while signalling is not

Figure 11: Additional Mechanisms

Mechanism	Welfare	Welfare	Distance	Allocated	% Better Off
	(unweighted)	(weighted)	(000 miles per day)	(tons per day)	(compared to CS)
Choice System	320	320	18	300	1
	(320, 320)	(320, 320)	(17.9, 18.1)	(299, 301)	(1, 1)
Old System	236	231	17.8	224	0.115
	(209, 260)	(208, 251)	(17, 18.5)	(218, 230)	(0.071, 0.165)
Closest Offers	95	82.3	0.502	57.4	0.054
Single Offer	(67.8, 124)	(56.2, 114)	(0.474, 0.529)	(53.2, 61.7)	(0.024, 0.094)
Closest Offers	258	254	10.3	305	0.316
All Offers	(222, 288)	(214, 285)	(9.91, 10.7)	(303, 307)	(0.247, 0.388)
Random Offers	63.6	6.02	28.2	320	0.001
	(6, 108)	(-58.4, 65.4)	(27.9, 28.5)	(320,320)	(0, 0.012)
Sequential Offers	254	252	22.8	299	0.242
	(226, 275)	(231, 271)	(22.3, 23.4)	(297, 302)	(0.176, 0.306)
Sequential Auctions	274	290	19.2	291	0.444
	(243, 298)	(260, 315)	(18.5, 19.8)	(286, 295)	(0.388, 0.506)
Combinatorial Auctions	355	365	19	281	0.804
	(347, 365)	(355, 377)	(18.6, 19.3)	(277, 284)	(0.741, 0.865)
Simultaneous Offers	289	294	22	302	0.369
	(257, 321)	(251, 325)	(19.6, 24)	(294, 311)	(0.247, 0.518)

Note: This table displays posterior means and 95% credible intervals for various measures. The final column gives the proportion weakly better off under each alternative mechanism compared to the Choice System.

so valuable on its own, signalling and batching appear to be complementary.

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 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, a key determinant of bidding behaviour. I proved the 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 36%, driven by the estimated scale of heterogeneity in the types of food needed by different food banks at different points in time. On average 89% of food banks are better off under the Choice System. Among other things, these results are driven by 'batching', that food is allocated in batches, whereas under the Old System food was allocated sequentially. This is an important finding as most other food relief organisations around the world allocate 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 the importance of using fake money would be extremely valuable.

References

Agarwal, N. (2015). An empirical model of the medical match. *American Economic Review*, 105(7):1939–78.

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.

Agarwal, N. and Somaini, P. (2020). Revealed preference analysis of school choice models. *Annual Review of Economics*, 12:471–501.

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. (2023). Identification and estimation of a dynamic multi-object auction model, available here.

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

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.

Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory*, 16(2):252–292.

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.

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.

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.

Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. *Journal of political economy*, 86(6):971–987.

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.

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.

Raisingh, D. (2020). The effect of pre-announcements on participation and bidding in dynamic auctions. Technical report, working paper.

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.[33] Categories come from Prendergast (2022), combining several small categories (such as pasta and rice). I form subcategories according to the most common product names, ensuring at least 30 lots per subcategory. Subcategories are more granular the more observations there are. E.g. for cereal this includes brands, whereas all cheese is together.

As well as the four Storage Methods outlined in the main text, food is also categorised by how it is used, including Meals, Ingredients, Condiments, Snacks, and Non-Food. Meals can be eaten on their own as part of a reasonably healthy diet for breakfast, lunch, or dinner. Multiple Ingredients can be mixed together to form a meal. Condiments are 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.

To find the distance between every lot × food bank combination I convert zipcodes into longitude/latitudes, then found the geodesic between these zipcodes.

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.

---

[33] 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.

Figure 12: Additional Descriptives

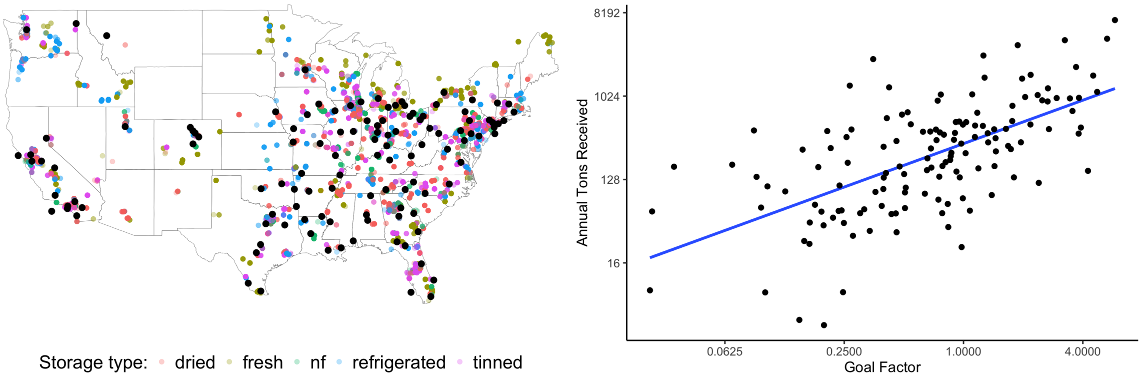
(A)

(B)

Note: Panel (A) shows approximate locations of food banks and donations, by storage type. Panel (B) plots annual consumption by food bank against Goal Factor, essentially the number of mouths they have to feed.

B Additional Descriptive Analysis

B.1 Geography of Food Banks and Donations

Figure 12 Panel (A) plots locations of food banks and donated lots by storage type. Locations are jittered by 200 miles to preserve anonymity. Locations of food and food banks are evidently correlated. Eastern states and California often have multiple food banks, while the West and Midwest usually have a single large food bank serving the entire state. Figure 12 Panel (B) plots annual winnings against Goal Factor, a combined measure of poverty and population in their catchment areas that proxies ‘mouths to feed’. While there is a clear positive relationship, evidently many other factors also determine their consumption.

B.2 Importance of Storage Costs

Storage costs are a key part of the model, forming the dynamic linkage. The decision to focus on storage costs was partly motivated by conversations with food bankers, and was also highlighted as key in Prendergast (2022). Figure 13 Panel (B) presents evidence that storage costs are the empirically relevant feature. We can replicate Figure 5 using food won by category, or by use type, instead of storage type, and see a similar relationship, suggesting that these dynamics may be driven by some other diminishing returns to payoffs. That said, this finding could arise because storage type and category/use type are correlated. Instead, Figure 13 examines the probability of bidding on a given lot, conditional on both whether they previously won food of the same storage type and whether they previously won food of the same use type. We see that Storage type remains the relevant feature.

B.3 Heterogeneity Over Time

I now present additional evidence of systematic variation in bidding behaviour over time. I use a fixed effects Tobit specification, investigating how food bank i's bid on food of storage type g varies across months m, writing α_{igm} for these fixed effects. 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.

This test may be overpowered if, for example, variation in factors other than food banks' needs, such as unobserved variation in quality, is mistaken for variation in preferences. Therefore, I also estimate a restricted specification with food bank specific month fixed effects. These capture variation in bidding common across food types. Under this null specification, a rejection 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 13 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 blue points give the baseline specification, while the red points give the month fixed effects specification. I reject the null hypothesis at 5% significance level, for 98% of food banks in my baseline specification, and 85% of food banks for the second specification. I have evidence that the types of food each food bank wants from Feeding America varies significantly over time.

Figure 13: Additional Descriptives

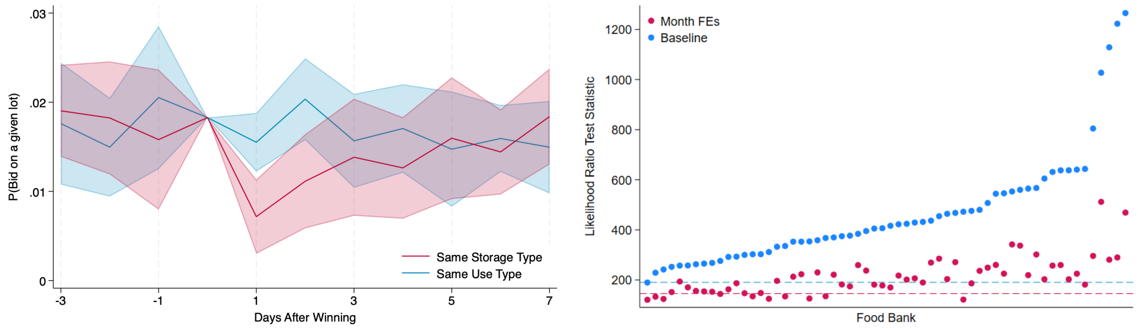

(A)

(B)

Note: Panel (A) plots the probability of bidding on a particular type of food, conditional on whether they won a load of the same storage type (in red) and conditional on whether they won a load of the same use type (in blue) at time zero. Panel (B) plots test statistics for tests that average bids for each type of food are constant over time. Estimation includes food banks who win at least 100 lots over the period, controlling for the distance.

C Budget Constraints

I now prove that the quasi-linear model is observationally equivalent to a model in which food banks bid subject to an inter-temporal budget constraint in a weakly dependent environment. This argument is well known in the consumption literature (e.g. Bewley (1977) and Hall (1978)). The idea is that the marginal value of an extra unit of wealth is the same today as far in the future, since any wealth shocks or changes in the state do not matter in the long run. Writing X_t as the state of the world at t, weak dependence means that the equilibrium distribution of food banks' states far in the future are independent of their states at t: $\lim_{\tau \to \infty} F_{X_{t+\tau}|X_t} = F_{X_{t+\tau}}$.

Suppose food banks face the standard inter-temporal budget constraint, so that in each period $E_t \sum_{s=0}^{\infty} \frac{\sum_l b_{ilt+s} - y_{it+s}}{(1+r)^s} = A_{it}$, Where A_t gives their savings at time t, y_t their 'income', and E_t is the expectation operator, conditional on information available at t. Suppressing i subscripts, a food bank's Lagrangian is then: $L(\mathbf{b}_t; X_t) = \sum_l \Gamma_l(b_{lt})\upsilon_{lt} + \sum_a P_a(\mathbf{b}_t)\kappa(X_t^a) - \lambda_t[\sum_l \Gamma_l(b_{lt})b_{lt} - y_t - A_t + E_t \sum_{s=1}^{\infty} \frac{\sum_l b_{ilt+s} - y_{it+s}}{(1+r)^s}]$. Write the optimised lagrangian multiplier, the marginal value of wealth, as $\lambda_t^* = \lambda(X_t)$.

At the optimum, $\lambda_t^* = E_t[\lambda_{t+\tau}^*]$ for all $\tau > 0$ (Hall, 1978). Otherwise a food bank could increase expected expenditure in one period, decrease it in another, increasing their expected payoff. Next, consider how $\lambda^*_{t+\tau}$, and hence λ^*_t, vary with X_t: $\lambda^*_{t+\tau}(= \lambda(X_{t+\tau}))$ only depends on X_t through the dependence in the conditional distribution, so that $\lambda(X_t) = E_t[\lambda^*_{t+\tau}] = \int_{X_{t+\tau}} \lambda(X_{t+\tau}) dF(X_{t+\tau}|X_t)$. Next, impose weak dependence and recognise that continuity of expected equilibrium payoffs in X_t, which is necessary for equilibrium existence, ensures continuity of $\lambda(X_t)$, allowing us to take the limit within the integral. Therefore, for $\tau \to \infty$, $E_t[\lambda^*_{t+\tau}]$ and hence λ^*_t are invariant to X_t and constant over time.

Now, consider the first order conditions of the Lagrangian: $\lambda^*[\frac{\partial \Gamma_l(b^*_{lt})}{\partial b_{lt}} b^*_{lt} + \Gamma_l(b^*_{lt})] = \frac{\partial \Gamma_l(b^*_{lt})}{\partial b_{lt}} \upsilon_{lt} + \sum_a \frac{\partial P_a(\mathbf{b}^*_t)}{\partial b_{lt}} \kappa(X^a_t)$. Recognise that these are the same optimality conditions as for the quasilinear model, derived in Altmann (2023) and Appendix E. Therefore the empirical implications of these models are the same, implying observational equivalence when λ^* is constant over time.

This argument is dependent on bidders not facing credit constraints, and credit is only available to food banks with below median Goal Factor. However, above-average Goal Factor food banks by definition have larger budgets and so the credit constraint rarely binds: On average they go below 1,500 shares 5% of the time, so they may still have constant marginal value of wealth. We can investigate this empirically: If the marginal value of wealth is constant then bids should not depend on A_t, similar to consumption smoothing. In Figure 14 panel (A) I estimate a Tobit model regressing food banks' assets on their bids, and plot the standardised coefficients in blue (red for the credit unconstrained). The relationship is positive and significant for three quarters of credit constrained food banks. However, we see that coefficients are small: A one standard deviation increase in budget leads to, on average, a 0.04 standard deviation change in bids. This is particularly small relative to the impact of distance, plotted in orange, or the type of food.

D 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, before introducing some additional terminology required for the proof. Then, in D.1 I present the two step proof. Altmann (2023) showed that conditional on having identified *i*) the conditional equilibrium distribution of bids $f_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t}$ and *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 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), following their spectral decomposition techniques for linear operators. The distinction between our arguments consists of the introduction of an additional intermediate variable (winnings) that acts as an observed shifter of the unobserved process, weakening the necessary identifying assumptions.[34] It also yields clear intuition behind identification of the latent state: Variation in this observed shifter 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. Lastly, note that this argument holds for multiple strategic agents and endogenous supply, 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 increase 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.[35]

D.0.1 Linear Operators and Spectral Decomposition

A linear operator $L_{z,x}$ is a map from the $L^{|x|}$ space of functions of x to the $L^{|z|}$ functions of z, such that for function $g : \mathbb{R}^{|z|} \to \mathbb{R}^{|x|} : (L_{z,x}g)z = \int f_{z,x}(z,x)g(x)dx$. Likewise, define the diagonal operator $D_{z,x}$ as: $(D_{z,x}g)z = f_{z,x}(z,x)g(x)$. Injectivity and surjectivity of these mappings are tied to completeness. $E[g(\mathbf{x})|\mathbf{z}] = 0$ for all $\mathbf{z}$ implies $g(\mathbf{x}) = 0$ for all $\mathbf{x}$ if and only if the mapping $L_{z,x}$ is injective and $f(\mathbf{z}) > 0$, so the left inverse of $L_{z,x}$ exists. Variation in $\mathbf{z}$ yields enough variation in $\mathbf{x}$ to pin down functions of $\mathbf{x}$.

---

[34] 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. This framework encapsulates multivariate latent states, as the signals (bids and winnings) are both also multivariate.

[35] The assumption nests two implicit assumptions: Reservation prices do not bind, and $\mathbf{z}_t^g$ has full rank. Binding reservation prices mean first order conditions do not hold with equality. However, as in Altmann (2023), reservation prices are not first-order, not substantially altering the identification problem. The rank condition on $\mathbf{z}_t^g$ requires something of each type is auctioned each period so that bids are informative of every type of stocks.

The proof below requires spectral decomposition of certain linear operators. As in Hu and Shum (2012), I require the eigenvalues of this decomposition are unique. This requires that for any pair $(\bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t})$ satisfying assumption 5 part *ii*) and any $\mathbf{s}_t$, that the density $f_{\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t}(\bar{\mathbf{w}}_t|\bar{\mathbf{s}}_{0t}, \mathbf{s}_t) = \int f_{\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{b}_t}(\bar{\mathbf{w}}_t|\bar{\mathbf{s}}_{0t}, \mathbf{b}_t) f_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t}(\mathbf{b}_t|\bar{\mathbf{s}}_{0t}, \mathbf{s}_t) d\mathbf{b}_t$ 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: $\Lambda(\mathbf{w}_t, \mathbf{s}_{0t}, \bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}, \mathbf{s}_t) = \frac{f_{\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t}(\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t) f_{\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t}(\bar{\mathbf{w}}_t|\bar{\mathbf{s}}_{0t},\mathbf{s}_t)}{f_{\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t}(\bar{\mathbf{w}}_t|\mathbf{s}_{0t},\mathbf{s}_t) f_{\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t}(\mathbf{w}_t|\bar{\mathbf{s}}_{0t},\mathbf{s}_t)}$. 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 5, 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.

D.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.

Lemma D.1. *The density* $f_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}}$ *completely determines* $f_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t}$.

I broadly follow the proof of Hu and Shum (2012) Lemma 3, using $(\mathbf{b}_{t+1}, \mathbf{w}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1})$ in place of $(V_{t+1}, w_t, w_{t-1}, V_{t-2})$, so 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:

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

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 D.0.1 above, ensures $D_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t}$ is invertible for any $(\mathbf{w}_t, \mathbf{s}_{0t})$.

4. From Hu and Schennach (2008) Lemma 1, assumption 5 *ii*) ensures 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, from equation 5, we consider two distinct $(\mathbf{w}_t, \bar{\mathbf{w}}_t) \in \bar{\mathbb{W}}$ for:

$$\begin{aligned} &L_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} L^{-1}_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\bar{\mathbf{w}}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} \\ &\qquad = L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} D_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t} D^{-1}_{\mathbf{s}_{0t+1},\bar{\mathbf{w}}_t|\mathbf{s}_{0t},\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\bar{\mathbf{w}}_t,\mathbf{s}_t}. \end{aligned} \tag{6}$$

5. Next, as in Hu and Shum (2012) equation 44, using equation 6 for the four different permutations of $\mathbf{w}_t, \bar{\mathbf{w}}_t, \mathbf{s}_{0t}, \bar{\mathbf{s}}_{0t} \in (\bar{\mathbb{W}}, \bar{\mathbb{S}}_0)$, i.e. variation in winnings and choice sets:

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

   Where the diagonal operator:[36] $(D_{\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\bar{\mathbf{w}}_t,\bar{\mathbf{s}}_{0t},\mathbf{s}_t} h)(\mathbf{s}_t) = \Lambda(\mathbf{w}_t, \mathbf{s}_{0t}, \bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}, \mathbf{s}_t) h(\mathbf{s}_t)$.

6. Equation 7 states the left hand side has an eigen decomposition given by the right hand side. The discussion in D.0.1 ensures eigenvalues are bounded, so these operators are also bounded. Theorem XV.4.3.5 from Dunford and Schwartz (1971) then ensures uniqueness of the decomposition. Our assumptions on Λ varying with $\mathbf{s}_t$ ensures that eigenvalues for different $\mathbf{s}_t$ are distinct (for some $(\mathbf{w}_t, \mathbf{s}_{0t}) \neq (\bar{\mathbf{w}}_t, \bar{\mathbf{s}}_{0t}) \in (\bar{\mathbb{W}}, \bar{\mathbb{S}}_0)$).

7. 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*). This requires that for any vector $\mathbf{x}$:

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

   Because g is invertible, the only function that satisfies perfect substitutes

---

[36] 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 for endogenous supply ($f(\mathbf{s}_{0t+1}|\mathbf{s}_{0t}, \mathbf{w}_t, \mathbf{s}_t) = f(\mathbf{s}_{0t+1}|\mathbf{s}_{0t}, \mathbf{w}_t)$) if this observed density is > 0 everywhere.

is $g(\mathbf{s}) = \mathbf{s} + \bar{\mathbf{s}}$, so that stocks are identified up to location. Imposing long run zero mean ensures $0 = E[\mathbf{s}] = E[g(\mathbf{s})] = E[\mathbf{s} + \bar{\mathbf{s}}] = E[\mathbf{s}] + \bar{\mathbf{s}}$ which holds only for $\bar{\mathbf{s}} = 0$.

8. Therefore, the density $f_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t}(\mathbf{b}_{t+1}|\mathbf{s}_{0t+1}, \mathbf{w}_t, \mathbf{s}_t)$ 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 identification for all $\mathbf{w}_t$.

□

Lemma D.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. $f_{\mathbf{b}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} = \int_{\mathbf{s}_t} f_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t}(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t) f_{\mathbf{s}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}}(\mathbf{s}_t, \mathbf{s}_{0t}, \mathbf{w}_{t-1}) d\mathbf{s}_t$. 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. Take sequential left inverses of equation 5 for:
   $L_{\mathbf{s}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} = D^{-1}_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} L_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}}$.

3. Substitute into the above for:
   $L_{\mathbf{b}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}} = L_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t} D^{-1}_{\mathbf{s}_{0t+1},\mathbf{w}_t|\mathbf{s}_{0t},\mathbf{s}_t} L^{-1}_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} L_{\mathbf{b}_{t+1},\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_{0t},\mathbf{w}_{t-1}}$.

4. 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 $f_{\mathbf{b}_t|\mathbf{s}_{0t},\mathbf{s}_t}(\mathbf{b}_t|\mathbf{s}_{0t}, \mathbf{s}_t)$ is identified.

5. $f_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} = \int_{\mathbf{s}_{t+1}} f_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{s}_{t+1}} f_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t} d\mathbf{s}_{t+1}$. In operator notation:
   $L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{w}_t,\mathbf{s}_t} = L_{\mathbf{b}_{t+1}|\mathbf{s}_{0t+1},\mathbf{s}_{t+1}} L_{\mathbf{s}_{t+1}|\mathbf{w}_t,\mathbf{s}_t}$.

6. 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.

□

E Inverse Bid System

In this Appendix I show that the food bank's optimisation problem yields the Observation and Censoring equations given in text. For the most part, I simplify the result from Altmann (2023) for a quadratic parametrisation of κ_i.

Parameterising κ_i, and given entry decision $\mathbf{d}_i^*$, the maximisation problem is given by: $\max_{\mathbf{b}} \Big\{ \sum_l \Gamma_l(b_l, d_l^*)(\upsilon_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*)[\Phi \mathbf{s}_i^{ah} - \mathbf{s}_i^{agT} \Psi \mathbf{s}_i^{ag}] \quad s.t.\ b_l \geq R_l \Big\}$. The maximand can be simplified, and so write the lagrangian as: $L(\mathbf{b}|\mathbf{d}^*, \boldsymbol{\upsilon}, \mathbf{s}) = -\mathbf{s}_i^{gT}\Psi\mathbf{s}_i^{gT} + \sum_l \Gamma(b_l, d_l^*)(\upsilon_l - b_l + \Phi\mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_i^g + \sum\limits_{m\neq l}\Gamma(b_m, d_m^*)\mathbf{z}_m^g]) + \Phi\mathbf{s}_i^h - \sum_l \Lambda_l(R_l - b_l)$. Λ_l are lagrangian multipliers. This simplification just exploits the quadratic κ, employing $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s})\mathbf{s}_i^{gT}\Psi\mathbf{s}_i^g = \mathbf{s}_i^{gT}\Psi\mathbf{s}_i^g$ and $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s})\mathbf{s}_i^a = \mathbf{s}_i + \sum_l \Gamma_l(b_l, d_l; \mathbf{s})\mathbf{z}_l$. Assumption 4 ensures $\nabla_{b_l}\Gamma_l > 0$ for $b_l > R_l$. Taking First Order Conditions and rearranging:

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

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 $y_l^* = b_l^* + \frac{\Gamma_l(b_l^*, d_l^*; \mathbf{s})}{\nabla_b\Gamma_l(b_l^*, d_l^*; \mathbf{s})} = y_l$.

Γ_l is non-differentiable at R_l, due to the risk of ties 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 (2023) 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:

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

This implies, 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 reverse this inequality (they prefer not to enter than enter at the reserve), and so $y_l \leq R_l$.

F 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. The modified proposition is given as:

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

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

$q_t(\mathbf{s}_i^g)$ gives the posterior probability that $\mathbf{s}_{it}^g = \mathbf{s}_i^g$ and $E[W_i(\boldsymbol{\upsilon}_{it}, \mathbf{s}_i, \mathbf{s}_0)|\mathbf{b}_{it}, \mathbf{d}_{it}, \mathbf{s}_i, \mathbf{s}_0] =$

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

The proof consists of two parts, first proving equality 8, and then equality 9.

F.1 Proof of Equality 8

To simplify notation, I drop i subscripts and dependence on the observe state $\mathbf{s}^0$. This is trivially introduced by multiplying objects by $\mathbb{I}[\mathbf{s}_t^0 = \mathbf{s}^0]$. I also drop dependence on the discrete actions $\mathbf{d}$, trivially introduced by multiplying objects by $\mathbb{I}[\mathbf{d}_t = \mathbf{d}]$ and summing over actions, as in the discrete choice case. In the below, $\delta()$ gives Dirac's delta function.

Proof:

1. First, I prove $f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s}) = \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t - \mathbf{b})|q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}$. Applying Bayes rule for $f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s}) = \frac{f_{\mathbf{b}_t,\mathbf{s}_t}(\mathbf{b},\mathbf{s})}{f_{\mathbf{s}_t}(\mathbf{s})}$. Focusing on the numerator, employ the definition of the Delta Function for $f_{\mathbf{b}_t,\mathbf{s}_t}(\mathbf{b}, \mathbf{s}) = E_{\mathbf{b}_t,\mathbf{s}_t}[\delta(\mathbf{b}_t - \mathbf{b})\delta(\mathbf{s}_t - \mathbf{s})]$. Next, apply Iterated Expectations for $= E_{\mathbb{O}_T}[E_{\mathbf{b}_t,\mathbf{s}_t}[\delta(\mathbf{b}_t - \mathbf{b})\delta(\mathbf{s}_t - \mathbf{s})|\mathbb{O}_T]]$. Now, recognise that $\mathbf{b}_t$ is a part of the observation set $\mathbb{O}_T$ for $= E_{\mathbb{O}_T}[\delta(\mathbf{b}_t - \mathbf{b})E_{\mathbf{s}_t}[\delta(\mathbf{s}_t - \mathbf{s})|\mathbb{O}_T]]$. Finally, the definition of the Delta Function and the posterior probabilities q_t yields $= E_{\mathbb{O}_T}[\delta(\mathbf{b}_t - \mathbf{b})|q_t(\mathbf{s})]$. Applying the same logic to the denominator gives the result.
2. Apply iterated expectations to the left side of equality 8 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 convenience write $\tilde{W}(\mathbf{b}_t, \mathbf{s}) = E_{\boldsymbol{\upsilon}_t}[W(\boldsymbol{\upsilon}_t, \mathbf{s})|\mathbf{b}_t, \mathbf{s}]$.
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 $\delta(.)$ the expectation is zero for $\mathbf{b} \neq \mathbf{b}_t$, so I can replace $\tilde{W}(\mathbf{b}, \mathbf{s})$ with $\tilde{W}(\mathbf{b}_t, \mathbf{s})$. Then, move 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:
   $= \frac{E_{\mathbb{O}_T}[\int_{\mathbf{b}} \delta(\mathbf{b}_t - \mathbf{b}) d\mathbf{b} \; \tilde{W}(\mathbf{b}_t,\mathbf{s}) q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} = \frac{E_{\mathbb{O}_T}[\tilde{W}(\mathbf{b}_t,\mathbf{s}) q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]}.$

□

F.2 Proof of equality 9

I now prove equality 9. If reservation prices do not bind then $\mathbf{b}$ directly pins down $\boldsymbol{\upsilon}$. As in Appendix E, with binding reserve prices $\boldsymbol{\upsilon}$ is determined up to a convex set. So we must integrate over this set. Use $\mathbf{b}^*$ and $\mathbf{d}^*$ to denote optimised bids / entry, so $\mathbf{b}_i^* = \mathbf{b}(\boldsymbol{\upsilon}_i, \mathbf{s}; \kappa)$ 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: $W(\boldsymbol{\upsilon}, \mathbf{s}) = \Phi \mathbf{s}^h - \mathbf{s}^{gT} \Psi \mathbf{s}^{gT} +$

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

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 the second row 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 E 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, integrate $W(\boldsymbol{\upsilon}, \mathbf{s})$ over $\boldsymbol{\upsilon}$, given $\mathbf{b}_{it}, \mathbf{d}_{it}, \mathbf{s}_i, \mathbf{s}_0$. The first row only depends on $\boldsymbol{\upsilon}$ through $\mathbf{b}^*, \mathbf{d}^*$, so we essentially have a change of variables.

The second row is affine in v_l, and otherwise only depends on $\boldsymbol{v}$ through $\mathbf{b}^*, \mathbf{d}^*$. Therefore, we take a simple conditional expectation of v_l given $b_{ilt} = R_l, \mathbf{b}_{i-lt}$. This yields equality 9.

$\square$

G Additional Estimation Details

G.1 Step 1.

First, I estimate beliefs about the probability of winning lot l given bid b_{itl}. While I assume there is zero probability of ties above reserve, I allow for ties at the reservation price. While this only occurs in 0.02% of auctions, 15% of winning bids are at the reserve price, so bidders must consider the non-zero probability of tieing were they to bid reserve. Food banks evidently recognise this and regularly bid between 1 to 50 shares above -2000.

The bidder wins lot l given bid b_{ilt} if $b_{ilt} > \bar{b}_{lt}$, which is 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 estimate, similar to ordered logit thresholds. 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 reserve, where $\epsilon_{lt} \sim \exp(\alpha)$ and α is to be estimated. So, competing food banks consider the excess mass just above the reservation price. This modelling approach is unusual, but helps rationalise the excess mass of bids at, and just above, reserve. I assume bidders do not internalise the probability of tieing one share above reserve. These thresholds are identified by the excess mass of winning bids at and just above the reserve. We then write down food banks beliefs Γ using the implied cdf of the observed maximum rival bid $\bar{b}_{lt}$.

G.1.1 Parameterisation and Computation

I normalise winning bids by the reservation price and use $\bar{b}^{*}_{lt} - R_l$. Lots contain up to four distinct subcategories, 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, and include additional scale fixed effects if the lot has been unsuccessfully auctioned previously or is "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.

The demand index $\boldsymbol{\vartheta}_{lt} = \mathbf{s}_{0t}^{T}\boldsymbol{\vartheta}$ is linear in the log quantity of food allocated at t and the log quantity allocated over the previous 30 days. I form the index for each of the 5 'Use' types, resulting in 10 parameters to estimate, dropping the first 60 days to construct the previous 30 days' supply. I draw samples using adaptively tuned Metropolis Hastings.

G.2 Step 2.

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

Pseudo-payoffs are parameterised as $\kappa_i(\mathbf{s}_i, \mathbf{s}_0) = \Phi_i \mathbf{s}_i^h - \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g$. Estimating Φ_i separately by food bank introduces too many parameters. I exploit that 152 subcategories (h) are nested within the 5 'use' 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 h has use u. So, I estimate the 152 parameters of $\bar{\Phi}_h$ and 5×90 use parameters $\tilde{\Phi}_{iu}$.

The standard deviation of the lot specific idiosyncratic value σ_l vary depending on the combination of goods in the lot. To simplify 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, and an 'other' parameter for the remaining 5.5% of combinations.

G.2.1 Priors and Hierarchical Distributions

Write $\boldsymbol{\psi}_i$ as the stacked vector of elements of Ψ_i and $\tilde{\Phi}_i$, which are drawn from hierarchical distribution $N(\boldsymbol{\psi}, \Sigma^{\psi})$.[37] The mean of $\tilde{\Phi}_i$ is not separately identified from $\bar{\Phi}$, so I use strong priors of 1 for corresponding elements of $\boldsymbol{\psi}$. I constrain Σ^{ψ} to be diagonal to simplify sampling (discussed shortly). Otherwise, I use weak normal-inverse-gamma priors.

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. Prior means of μ_{ig} and Σ_{ig} are set to the mean and variance of $-w_{igt}$ and the prior mean of δ_{ig} is 0.[38] I set prior shape parameters to ensure 100 times more weight is put on the data than on the priors.

G.2.2 Computation

I focus on data from food banks' 25 highest bids each day.[39] I begin the Carter-Kohn algorithm from day 61 using the long-run mean and (10×)variance, then discard the first 40 days sampled. I impose elements of Ψ and Φ are positive using truncated priors, for identification requirements and because the sampler often diverged other-

---

[37] The hierarchical framework reduces variance at a cost of bias. Food banks with little identifying variation put lots of weight on the hierarchical parameters, whereas bidders with lots of variation put very little weight. The bias causes parameters to be drawn together (too little heterogeneity), biasing results in favour of the Old System.

[38] On average winnings should off-set shortfalls in net donations. $\delta_{ig} = 0$ biases results towards the Old System: If net donations are endogenous ($\delta < 0$) food banks can influence future donations. This makes control over $\mathbf{w}_{it}$ even more valuable as they can focus on winning food they have little influence over from local donors.

[39] This considerably speeds up computation and convergence, and should not significantly impact results. Even Type 1 food banks only place more than 15 bids 1% of the time. However, there are more than 25 lots on 25% of auction days, and ignoring that bidders do not bid on these extra lots will bias results towards food banks bidding too frequently. However, this bias should be small: if they place maybe 5 bids, then I already account for not bidding on lots 6-25. Relatively little additional information is conveyed by them also not bidding on lots 26+.

wise. The $\bar{\Phi}^h$ constraint binds for 5 unpopular subcategories with limited shelf lives (bread and milk). I impose $\boldsymbol{\delta}$ and Σ diagonal, for stationarity and because separately identifying off-diagonals from the off-diagonals of Ψ proved difficult.

Priors are generally conjugate, enabling sampling from conditional posteriors. The exceptions are the hierarchical parameters (ψ, Σ^{ψ}), due to the constraints on Ψ, where I use Metropolis Hastings. I sample beliefs every 5 iterations, and run the sampler for one million iterations, burning out the first half. Initial points are drawn from the priors. I run two independent chains and uniformly sample 500 draws from each chain for later steps.

G.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.[40]

Because stocks are continuous I evaluate the continuation value over a uniform grid of 30 points for each of the 4 dimensions, taken from the 2.5 and 97.5 percentiles of sampled states. For each t I form the (approximate) posterior distribution of stocks $q_t(\mathbf{s}_i^g)$ using my draws of $\mathbf{s}_{it}^g$ and a gaussian kernel. I evaluate the maximised payoff at t, $\tilde{W}(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s})$, using finite sample approximations of the formulae in Appendix F for each parameter draw.[41]

I fit a polynomial function of the states to the ex-ante value function, including interaction terms. I use least squares, weighted 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 in Appendix I.3.

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

---

[40] The marginal welfare from gaining $\mathbf{z}_{tl}^h$ is $\Phi_i \mathbf{z}_{tl}^h$, and does not depend on $\mathbf{s}_{it}^h$, even though Φ is not a structural parameter. The marginal pseudo-payoff gives the expected benefit going forward of these extra stocks, and is independent of any other stocks because pseudo-payoffs, and hence also payoffs, are affine in $\mathbf{s}_{it}^h$.

[41] As I have large T I do not account for sampling variation in the finite sample expectation. I previously used a bootstrap, which gave similar posterior means but ignores dependence within draws so overestimates variances.

Figure 14: Additional Parameter Estimates

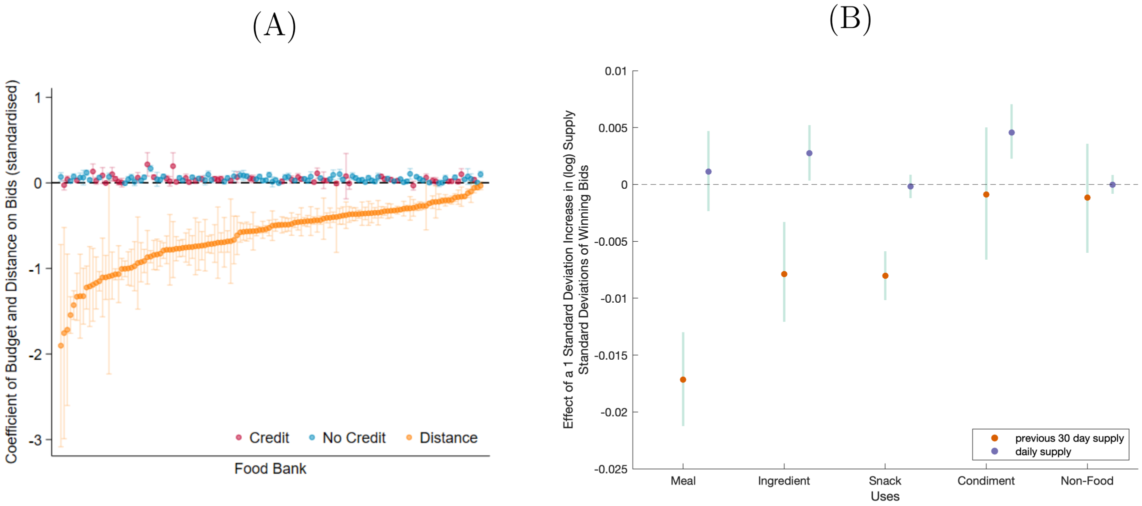

Note: Panel (A) plots coefficients of budget against bids across food banks, compared to the coefficients of distance against bids (orange). Excludes estimates that did not converge (all with credit access). Variables are standardised, includes food bank × month fixed effects. Panel (B) displays posterior means and 95% Credible Intervals of $\boldsymbol{\vartheta}$, the first stage coefficients on aggregate supply, for both daily and the previous month's log supply, split by Use type. These parameters form the demand index $\boldsymbol{\vartheta}_{lt}$.

H Additional Estimation Results

I now report additional estimation results. Full results available on request. Figure 14 panel (B) displays estimated first stage demand parameters, forming the index $\boldsymbol{\vartheta}$.

Figure 15 reports Gelman-Rubin statistics, displaying the proportion of statistics below the recommended cutoffs of 1.2 and 1.1. I report results for all 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. Convergence failure was generally due to multiple modes, specific to food bank × storage type combinations for those who rarely bid on certain types of food (suggesting an identification problem). Because the model still fits the data well, I do not worry about the lack of convergence. Furthermore, these Statistics assume the target distribution is approximately normal, and may fail to detect convergence if the target distribution is sufficiently skewed. I also consider convergence of the first and second moments only, and find proportions increase by around 5 percentage points.

Figure 15: Gelman-Rubin Convergence Statistics

Parameter		Γ	$\boldsymbol{\delta}_i$	$\boldsymbol{\mu}_i$	$\boldsymbol{\Sigma}_i$	$\bar{\boldsymbol{\Phi}}$	$\tilde{\boldsymbol{\Phi}}_i$	λ_i	σ_l	α_i	$\boldsymbol{\Psi}_i$
Type 1	Prop < 1.1	1	0.857	0.884	0.777	0.98	0.971	1	0.983	0.964	0.886
	Prop < 1.2	1	0.902	0.893	0.839	0.993	0.993	1	1	1	0.918
Type 2	Prop < 1.1	-	0.774	0.832	0.824	-	0.948	1	-	0.976	0.807
	Prop < 1.2	-	0.826	0.859	0.859	-	0.969	1	-	1	0.849

I 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 I.1, I.2, and I.3 respectively.

I.1 First Stage

I.1.1 Food bank Specific Beliefs

I implicitly imposed that food banks all have the same beliefs. This permits estimating Γ on the distribution of winning bids only. I can test whether the distribution of i's rivals' highest bids is different from the distribution of winning bids using a score test. In figure 16 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 tests can reject the null at 10% significance level.

I.1.2 Dependence on Aggregate Supply

For feasibility I require that Γ does not depend on any individual food banks' state. If equilibrium is sufficiently competitive then no individual's behaviour can significantly shift the distribution of winning bids. If so, variation in an individual food bank's state will not shift this distribution either. I consider whether the distribution of winning bids changes when data from food bank i *and* the auctions they won are removed from the data. If i has a significant effect on the distribution of winning bids, we would expect that the distribution is different when we drop all their data.

In figure 16 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 reject the null at 10% significance level.

Figure 16: Robustness: Stage 1

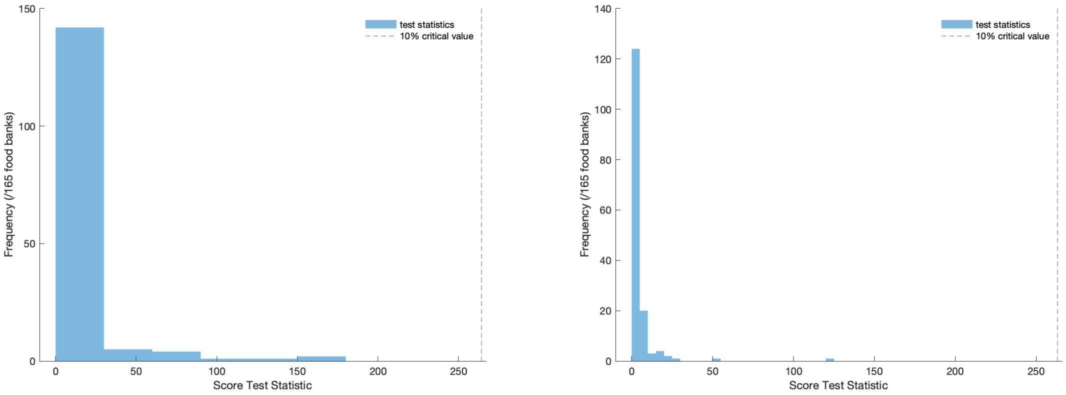

(A) (B)

I.1.3 Independence of Winning Bids

Next, I investigate whether winning bids are independent within a period, conditional on covariates. This ensures that $P(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s}_t)$ are products of marginal distributions. Because κ_i is quadratic I only require that winning bids are pairwise conditionally independent.

Writing $\bar{b}_{lt}$ for the winning bid on lot l, I investigate dependence by regressing $\bar{b}_{lt}$ on $\bar{b}_{l't}$ and include all first stage covariates $\mathbf{x}_{lt}$, $\mathbf{x}_{l't}$, and common state variables $\mathbf{s}_t^0$. I include every pair of auctions (l, l') that occur simultaneously ($\approx 8e5$ observations). Under the null of independence the coefficient on $\bar{b}_{l't}$ is zero. The degree of correlation may depend on lot characteristics, so I also interact $\bar{b}_{l't}$ with $(\mathbf{x}_{lt}, \mathbf{x}_{l't}, \mathbf{s}_t^0)$. I also consider the triple interaction between $\bar{b}_{l't}$, $\mathbf{x}_{lt}$, and $\mathbf{x}_{l't}$. 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 17. I can reject the null at 1% significance in all specifications. However, examining the R^2s we see the degree of dependence is small. Covariates alone account for 42.76% of the variation in winning bids. Including the $\bar{b}_{l't}$ interactions only explains an extra 0.3% of variation. Despite rejecting independence, this suggest winning bids are close to independent, so a reasonable approximation for food banks' beliefs.

Figure 17: 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.

I.2 Second Stage

I.2.1 Incorporating the Common State

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

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

$\boldsymbol{\vartheta}(\mathbf{s}^0) \cdot \mathbf{s}_i^u$ gives the elementwise product of the demand indices and the food banks' stocks by use type. Υ_i is a 1×5 vector of i specific coefficients. If the dependence is strong, we should detect this using a linear approximation. I interact the index with stocks so that it impacts marginal pseudo-payoffs. 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 future.

In Figure 18 panel (A) I plot estimates of Υ_i. I focus on just the Type 1 food banks — if any, it is these large food banks who will be dynamically strategic. 27.8% of parameters are significant at the 5% significance level. However, parameters are generally very small. Only two of these have economically significant magnitudes, with a $1sd$ deviation difference in the demand indices leading to more than a 100 share difference in willingness to pay.

I.2.2 Endogeneity of the Inverse Bid System

The Observation Equation is technically endogenous, and may bias estimates. In addition, unobserved auction timing introduces measurement error in $\Gamma_m(b_{itm})$. However, we expect the degree of bias to be small. To investigate the bias, as in Altmann (2023), I use $\mathbf{z}_{tl}^g(\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g)^T$ to instrument for $\mathbf{z}_{tl}^g(\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2\sum_{m\neq l}\Gamma_m(b_{itm})\mathbf{z}_{tm}^g)^T$. This instrument is clearly valid and relevant: We instrument using variation in bids caused by variation in the characteristics of that lot only, excluding endogenous variation through $\Gamma_m(b_{itm})$.

I do not run a full bayesian instrumental variable procedure, which requires normality assumptions on the endogeneity. This is difficult to justify as $\Gamma_m(b_{itm})$ is observed is not normal. Instead, for each draw of unobserved stocks I run frequentist 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. 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 as usual. I repeat this exercise for each of the 1,000 draws of augmented data (producing something like a credible interval). The diagonal coefficients on these control functions are plotted for the 28 Type 1 food banks in Figure 18 panel (B). Results are normalised by corresponding components of Ψ_i highlighting their relative magnitude. Estimates are clustered close to 0, suggesting that any endogeneity is extremely small. All but 1 have absolute magnitude of the endogeneity less than 10%.

Figure 18: Robustness: Stage 2

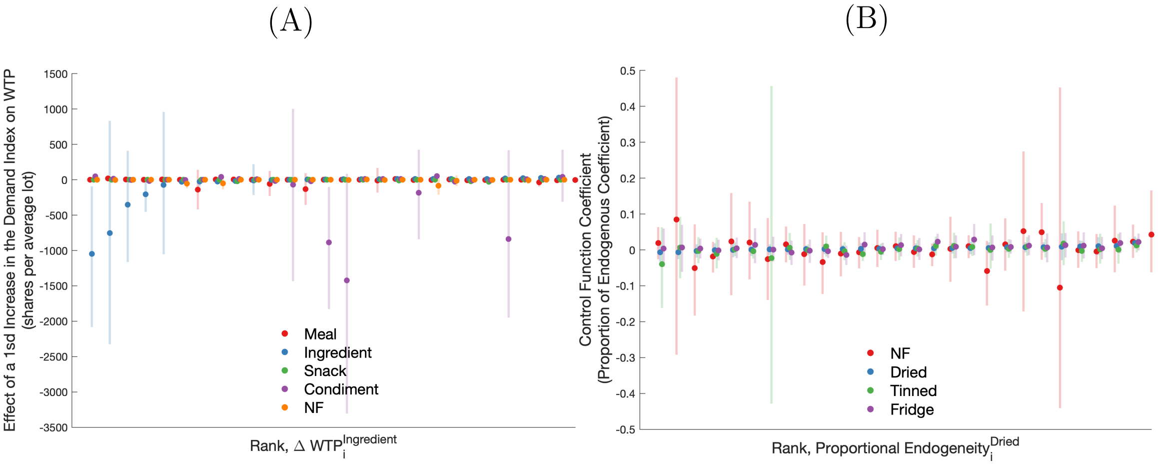

I.3 Third Stage

The quadratic approximation made in Step 3. is technically incompatible with the parametric assumptions made in Step 2. However, this approximation fits well. I consider the R^2s from 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 in equation 9.

J 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. I do not find equilibrium for all 1000 parameter draws, as this requires too much computation. Instead, I use a random sample of 30 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 draws.

J.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 be invalid if I wanted to consider changes to the Choice System. The central problem then concerns the bidding function, which involves a complex combinatorial problem of deciding the combination of lots to enter. 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 positive, imposing gross substitutes, this algorithm is guaranteed to find the global optimum.

J.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.

J.2.1 Set Up

Food with characteristics $\mathbf{x}_l$ is donated to Feeding America at some exogenous rate. This rate, and the distribution 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 Poisson rate q_i. Conditional on arriving, the net donation $\tilde{\mathbf{x}}_{it}$ is normally distributed: $N(\tilde{\boldsymbol{\delta}}_i \mathbf{s}_{it} + \tilde{\boldsymbol{\mu}}_i, \tilde{\Sigma}_i)$. The parameters $\tilde{\boldsymbol{\delta}}_i, \tilde{\boldsymbol{\mu}}_i, \tilde{\Sigma}_i$ are calibrated so that integrating over the day's net donations, the resulting distribution has the same mean and standard deviation as those 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 $\upsilon_{ilt} + \tilde{\pi}_i(\mathbf{s}_{it} + \mathbf{z}_l)$, or $\tilde{\pi}_i(\mathbf{s}_{it})$ if they reject. The payoff function $\tilde{\pi}_i$ is calibrated so that integrating over a day, if stocks do not change, the payoff is the same as 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.

J.2.2 Equilibrium

Write the agent's value function as $V_i(t, \mathbf{s}_i, \mathbf{s}_0)$, 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 are 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 accept if $\upsilon_{ilt} + V_i(t, \mathbf{s}_i + \mathbf{z}^g_{ilt}, \mathbf{s}_0) \geq V_i(t, \mathbf{s}_i, \mathbf{s}_0)$. The agent believes that Feeding America will offer them a load at Poisson rate $p_i(t, \mathbf{s}_0)$. 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 $f^c_i(\upsilon_i, \mathbf{z}^g, \mathbf{z}^h; t, \mathbf{s}_0)$.

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 offers made to others, nor their stocks. I assume the only objects used to form their beliefs are $\mathbf{s}_i$, their own Goal Factor, and the time since their last offer τ. The offer rates p_i and the distribution of offered lot characteristics f^c_i are food bank specific. For simplicity I impose p does not depend on τ. This is because I found that in equilibrium offers occur so frequently (3-4 per day) that Value Functions are insensitive to τ, so imposed this independence. I also assume beliefs do not change based on the previous history of offers: Food banks do not infer from frequent offers how offers will increase in future. For f^c_i, I split lots into the same 60 discrete category combinations used for the lot specific variances σ_l, detailed in Appendix G.2. Then, conditional on the category combination, I assume food banks believe that, in equilibrium, the distance between them and the lot k 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.

J.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:

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

The solution must be independent of τ. Holding constant beliefs (p_i^k, f^{ck}) I solve the equation numerically using successive approximations, switching to a dampened Newton-Kantorovich algorithm as progress slows. I then simulate the Old System 28 times using these value functions, before updating beliefs, (p_i^k, f_i^{ck}), using a frequency estimator. I repeat this process until I cannot reject that beliefs change at 10% statistical significance.

In the simulations, I force food banks to reject lots if accepting would push their stocks outside the grid. Likewise, to avoid too much extrapolation in evaluating welfare, if the state drops too far below the lowest state 'observed' under the Choice System, I replace flow payoffs with the 2.5% flow payoff, essentially bounding payoffs below. This predominantly binds in the "no Feeding America", Random, and Closest counterfactuals.

J.2.4 Other Mechanisms

The Sequential Offer Mechanism is simulated the same as above, but allowing every load of food to be offered to every food bank. Likewise, the two Closest mechanisms simply involve offering food by order of distance from the lots' origin.

For Sequential Auctions I exploit that, in equilibrium, lots are always 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:

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

p^c gives the arrival rate of lots with characteristics c, and $\tilde{\Gamma}_c(\upsilon_{il} + V_i(0, \mathbf{s}_i + \mathbf{z}_l) - V_i(0, \mathbf{s}_i)|c)$ 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 impose 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.

For Combinatorial Auctions, I impose that value functions are quadratic, and exploit that gross substitutes ensures period outcomes are efficient and greedy solvable. This allows me to evaluate payoffs from a particular round of auctions, given value functions from a previous iteration. I repeat this 300 times per period (using new υ_{ilt} draws), and then can evaluate an updated value function. I repeat this procedure until value functions converge. I drop food banks that did not converge for a particular parameter draw.

Simultaneous Offers has food banks form beliefs over the rate they receive calls from Feeding America, but differs from the Old System in their beliefs about expected payoffs from receiving a call. Upon receiving a call, denote a food bank's maximum payoff from the menu (or nothing) as the solution to: $x = \underset{l \cup \emptyset}{max}\{\upsilon_{ilt} + V_i(\mathbf{s}_i + \mathbf{z}_{lt}, x, 0)\}$. I treat x as an inclusive value, so that the expected payoff from receiving a call is $E[x|\mathbf{s}_i, x,\tau]$, depending on $\mathbf{s}_i$, time since last call, and the previous menu's inclusive value. Assuming $E[x|\mathbf{s}_i, x,\tau]$ is quadratic in $(\mathbf{s}_i, x,e^{\tau})$ implies V has an analytic solution. To find equilibrium beliefs I simulate the mechanism, storing the xs, before fitting the conditional expectation using NLLS, repeating the simulations until these beliefs no longer change statistically significantly.