Choice, Welfare, and Market Design:

An Empirical Investigation of Feeding America's Choice System

Sam M. Altmann*
Wadham College, University of Oxford

Click Here For The Most Recent Version

Abstract

Feeding America, an organisation responsible for feeding 130,000 Americans every day, distributes donated food among a network of participating food banks. Feeding America's allocation mechanism, the 'Choice System', uses first-price auctions to allow food banks to signal which types of food they need from Feeding America. This provides food banks a large degree of choice over the types of food they receive. This paper examines the welfare and distributional consequences of enabling this choice. I apply a dynamic auction model to Choice System bidding data, estimating the distribution of food banks' heterogeneous and time-varying needs. I then use these estimates to compare the Choice System to the previous allocation mechanism employed by Feeding America which gave food banks very limited choice. I estimate that the Choice System increased welfare by the equivalent of a 17.1% increase in the quantity of food being allocated.

*Department of Economics, University of Oxford, samuel.altmann@economics.ox.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 Ian Crawford, Howard Smith, Alex Teytelboym, Daniel Waldinger, Paul Klemperer, Francis DiTraglia, Shihang Hou, Ellen Lees, and Luke Milsom for useful comments and conversations, as well as seminar participants at EEA-ESEM Bocconi, The Annual Conference in Dynamic Structural Econometrics, The 23rd ACM conference on Economics and Computation, The Conference on Mechanism and Institution Design, The Network of Industrial Economists 2022, The University of Surrey, and The University of Oxford. Finally, thank you to Feeding America for giving me access to their data.

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 hospital boards 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 keep up with the ever-increasing demand for their services. Like many food bank networks 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 a truckload of food, and return to the back of the queue regardless of whether they accepted or rejected this load. This mechanism was unpopular among food banks as they 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 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. Feeding America's current allocation mechanism, the 'Choice System', consists of an auction market in which food banks are given an amount of virtual currency to bid on loads of donated food (Prendergast, 2017). This gives food banks a strong degree of control, and choice, over the food they receive.

In this paper I use a rich model of food bank bidding behaviour to investigate welfare under the Choice System, compared to alternative mechanisms that allow food banks varying degrees of choice. I develop an empirical strategy to estimate food banks' demand functions, applying a dynamic auction model 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,

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

cater to different numbers of clients, and receive different types of food from their local donors. I then use these estimates to evaluate equilibrium allocations under a number of alternative allocation mechanisms. Counterfactual simulations demonstrate that, relative to the Old System, the Choice System is extremely effective at achieving Feeding America's welfare and equity goals: Welfare is 17.1% higher under the Choice System than the Old System. This is roughly equivalent to an additional 50 tons of food allocated each day, enough to support an additional 22,300 people.

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 follows the empirical auction literature by employing the dynamic multi-object auction model of Altmann (2022), which builds on the models of Gentry et al. (2020) and Jofre-Bonet and Pesendorfer (2003). Descriptive evidence demonstrates the need for both the dynamic and multi-object framework: First, 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. Second, food banks certainly act as forward looking bidders, given that auctions happen so frequently. Conditional on winning a load, food banks are less likely to bid on similar loads on subsequent days. This suggests food banks treat loads as durable goods subject to storage costs, emphasising the need for an empirical model that accounts for the dynamic environment.

The value 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 capacity), and how it is used. It is further divided into 15 broad categories and 164 subcategories, each associated with distinct preference parameters. 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 food banks may irregularly receive donations from their local donors and irregularly give out food to clients.

The main estimation challenge is that I do not observe food banks' stocks. Current stocks are a key determinant of demand — if a food bank suddenly stops bidding on a particular type of food it might be because, unobserved by the econometrician, they recently received this from a local donor. A methodological contribution of this paper is to develop a procedure to estimate bidders' values in a dynamic multi-object auction environment when individual state variables (stocks) are unobserved. I overcome this problem using a Gibbs Sampling procedure, employing a data-augmentation step to draw the unobserved stocks from their conditional posterior distribution. The model is identified through observed variation in winnings which drives systematic variation in bidding behaviour. The change in the propensity to bid immediately after winning a lot identifies food banks' storage capacities: After a recent win, capacity constrained food banks will stop bidding on that type of food. Meanwhile, the length of time before food banks return to their average bidding propensity enables identification of the unobserved state transition process: If it takes them a long time to return to bidding on a particular type of food, this suggests they generally have access to that food from their local donors. To the best of my knowledge this is the first paper to estimate a model of this type.

I employ the three step estimation procedure introduced in Altmann (2022). In the first step I estimate equilibrium beliefs by estimating the conditional distribution of winning bids. I then invert food banks' first order conditions for optimal bidding, obtaining an inverse bidding system as in Guerre et al. (2000) and Gentry et al. (2020). 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 the pseudo-static payoff function. Therefore, in the third step I evaluate the estimated continuation value, before backing out the distribution of flow payoffs from the definition of the pseudo-static payoffs.

I find significant evidence of demand heterogeneity both across food banks and over time. I estimate large differences in access to local donors. For example, I find that food banks in urban areas generally have little access to fresh food, such as produce, so get much of it from Feeding America. This contrasts with more rural food banks which rarely need fresh food from Feeding America. This means it is important for food banks to be able to sort across types of food. Likewise, some food banks' local donations are estimated to be very variable over time, meaning it is important for them to be able to pick and choose different types of food as and when they are most needed. I estimate that day to day variation in stocks account for 45% of the daily variation in bidding behaviour. The model also suggests that food banks go through extended periods with high stocks, during which they very rarely place bids, and periods with low stocks, during which they bid very frequently. Therefore, over the long-run, I estimate that 72% of the variation in bidding behaviour can be attributed to unobserved variation in stocks.

Using the estimated model I consider equilibrium allocations under a number of alternative mechanisms that permit food banks varying degrees of choice. First, I consider the mechanism previously employed by Feeding America (The 'Old System'). This allows me to quantify and qualify the benefits of choice and the Choice System over lack of choice, complementing the evidence presented in Prendergast (2017) and Prendergast (2022). I find that welfare is 17.1% higher under the Choice System than under the Old System. The majority of this welfare gain is due to food banks having more control over their stocks, better tailoring their allocations to fit their most pressing needs first. This is as opposed to accepting sub-optimal food when they face significant storage costs; food that may be used more effectively by another food bank at that point in time. As a result, around 85% of food banks are estimated to be better off under the Choice System.

Feeding America's allocation problem is faced by numerous other food bank networks around the world, such as the European Federation of Food Banks (FEBA), and Food Bank Australia. Therefore I also consider mechanisms employed (often implicitly) by some of these other food bank networks.[2] A mechanism that offers food only to the nearest food bank, aiming to minimise transportation costs but allowing food banks even less choice than the Old System, achieves only 65% of the welfare under the Old System. This result arises because even under the Old System each load was offered to multiple food banks. Many food bank networks, including the Trussell Trust in the U.K., implicitly use this mechanism by not allocating food

[2]Other food bank networks often face somewhat different problems, in both scale and scope, to Feeding America. For example, transportation cost are known to be a larger factor in Australia. Likewise, FareShare (U.K.) face an allocation problem closer to an individual food bank allocating food to its associated food pantries. Therefore the results from this paper cannot be exactly applied to these other settings. That said, certain broader lessons are still valuable to these organisations. Future work, ideally using data from these other settings, is certainly needed.

centrally and instead linking food banks up with nearby donors. However, even if food is offered to every food bank (in order of distance), this is only marginally better than the Old System, but much worse in its distributional effects — food banks which happen to be well situated consume the most valuable food. This is because I estimate that transportation costs are not a large cost factor for most food banks.

The Choice System allocates food simultaneously in batches, rather than allocating food as donations arrive. Among other benefits, this 'batching' 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. The majority of other mechanisms employed for food allocation are sequential in nature.[3] I simulate an 'efficient' sequential mechanism and find that welfare is still around 12% lower than under the Choice System. This is because, while food is always allocated to the food bank that values it most, food banks are not always allocated the type of food they need the most. Then, when a donation does come along that they really want, they no longer have capacity to store it. This is essentially the same effect driving the poor welfare results for the Old System.

The paper proceeds as follows: Section 1.1 discusses how this paper contributes to the related literature. Section 2 describes the institutions and data being studied. Section 3 provides descriptive evidence that food banks have strongly heterogeneous preferences, as well as detailing several descriptive facts that must be taken into account in any empirical model. Section 4 outlines the empirical model of food bank bidding behaviour, making clear the assumptions necessary for identification. Section 5 describes the estimation procedure and parametric assumptions employed for the structural model, while Section 6 details the estimation results. Section 7 details the counterfactual mechanism considered before presenting the simulation results.

1.1 Related Literature

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

Empirical market design is a growing literature analysing preferences and allocations in centralised assignment markets, often employing techniques from Empirical Industrial Organisation. There is an extensive literature empirically analysing the al-

[3]Theoretical results suggest this is suboptimal (Akbarpour et al. (2020), Baccara et al. (2020)).

locative effects of centralised school choice and medical residency matching (see, for example, Agarwal and Somaini (2020) and Agarwal (2015)). However this literature typically employs static models, since students are only assigned a school once. The food allocation problem is both multi-object (many food banks are allocated many loads of food) and dynamic (food must be allocated repeatedly over time).

On the multi-object side Prendergast (2017) and Prendergast (2022) also estimate the welfare associated with Feeding America's transition to the Choice System. My structural approach is complementary to their descriptive and sufficient statistic approaches. However, my main results are relatively similar to Prendergasts'. This paper employs richer data that is disaggregated at the auction level and includes information on losing bids. By studying the exact timing of food banks' consumption, as well as losing bids, I gain a detailed understanding of how food banks make inter-temporal substitutions. This allows me to simulate alternative allocation mechanisms and consider additional counterfactuals of interest, for example investigating the benefits of batch versus sequential allocation. Budish and Cantillon (2012) study the course allocation problem, which also uses a system of virtual currency to allocated MBA courses to students. Similarly, Fox and Bajari (2013) use methods from the stable matching literature to measure the efficiency of the 1994 US spectrum auction. Preference heterogeneity is an important theme in these papers, even though their data is primarily cross-sectional. In contrast, in this paper I exploit the panel-dimension of the data, allowing unrestricted individual specific heterogeneity.

The empirical dynamic assignment literature often consists of evaluating waiting list design. Agarwal et al. (2020) and Agarwal et al. (2021) study the mechanisms used to offer deceased donor kidneys to transplant patients. Likewise Waldinger (2021) studies the allocation of public housing. Similar to this paper, they assess the value of giving agents choice over their allocations, considering the trade-offs between efficiency and other concerns of policy makers. Other work analysing dynamic multi-object allocation problems include Verdier and Reeling (2022) on hunting licenses, Gandhi (2019) on nursing homes, and Liu et al. (2019) on peer-to-peer ride sharing. This literature highlights the importance of heterogeneity in preferences and match values, but typically do not consider the role of heterogeneity over time, which is an important factor in the food allocation problem.

This paper applies the dynamic multi-object auction model of Altmann (2022), which combines the models of Jofre-Bonet and Pesendorfer (2003) and Gentry et al.

(2020). Unlike these papers reservation prices and endogenous entry are important in my application, with the average bidder only bidding on around 2% of auctions. I draw from Groeger (2014) and Balat (2013), who both introduce models of participation in dynamic first-price auctions. In the multi-object setting endogenous entry is structurally more complex due to the inherent combinatorial problem, and standard estimation procedures become computationally infeasible. The focus on a large auction market is similar to Backus and Lewis (2016) who introduce a framework for analysing dynamic bidding in a large single-unit second-price auctions, studying eBay's second hand camera market. Their framework has been employed a number of times, for example in Bodoh-Creed et al. (2021) and Hendricks and Sorensen (2015). To my knowledge, this is the first empirical auction paper to consider the role of time-varying unobserved heterogeneity.[4]

2 Institutional Background and Data

This section describes the Choice System and the Old System. Details of these mechanisms come from Prendergast (2017). Then in Section 2.2 I describe the data used in this paper.

2.1 Feeding America

Feeding America, formerly America's Second Harvest, began in 1976 as a collection of food banks that would solicit donations from local grocery stores and farms. As additional food banks joined their network it became necessary to co-ordinate resource sharing. In 2005, at the recommendation of a task force consisting of economists and food bank managers, they replaced the Old System with the Choice System.

Many of Feeding America's associated food banks operate as standard food pantries

[4] This paper also relates to the literature on the identification of dynamic models in the presence of unobserved states, building on Kasahara and Shimotsu (2009) Hu and Shum (2012), and Connault (2014). My identification argument is similar to that of Berry and Compiani (2020), using an instrument (observed winnings) to identify changes in the unobserved state. Unobserved stocks are also a key feature of the inventory models of Hendel and Nevo (2006) and Erdem et al. (2003) among others. The key distinction between my model and these examples is that agents also receive food from an external source - local donors. These local donations are expected to be a key driver of heterogeneous bidding behaviour across food banks and across time.

- directly giving out food to those in need. However, the majority of food banks act as food distributors; themselves responsible for storing and sending out food to *hundreds* of local food pantries.

2.1.1 The Old System

Under the Old System any truckload of food donated to Feeding America was offered to the head of a queue. The potential recipient had a few hours to accept or decline the load, 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 the loads. 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 the poverty in their local area relative to the national average. 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 offered in each truckload was essentially random, so that on average food banks received the same quantities of food per mouth. This would have been optimal if food banks all had the same preferences and capacities. In reality, different food banks needed different types of food at different times. Food banks use food from Feeding America to substitute for food they do not receive from their local donors. A food bank surrounded by farms is likely to have a weaker preference for fresh produce than a food bank in a city. Feeding America wanted to improve welfare by taking account of differing needs. They decided to use a market mechanism to give food banks power over the allocation they receive.

2.1.2 The Choice System

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

Food banks bid with a virtual currency called 'shares'. Other than storage and transportation costs, the only opportunity cost a food bank faces when bidding is that they will have fewer shares to bid on other lots. Feeding America can control which food banks have the most shares, ensuring that food banks with larger Goal Factors are allocated more shares and, consequently, receive more food (in the spirit of the Second Welfare Theorem). All spent shares are redistributed each night.[5] Food banks can save shares from one day to the next. Those with less than the median allocation of shares have access to interest-free credit, so that food banks can smooth their consumption over time. The money supply is set to ensure that prices remain constant (on average) over time, reacting to changes in the supply of food.

Food banks can bid negative amounts, down to a reserve price of -2000 shares. This incentivises food banks to accept undesirable loads, helping Feeding America maintain good relations with their donors by ensuring that every lot is graciously accepted. This had been a problem under the Old System - donors whose donations are refused are less likely to donate in future.[6] As all lots are eventually sold, donors now feel like their donations are always graciously accepted, and so they continue to donate. On average 21% of lots are sold at strictly negative prices, and 10% are sold at the reservation price.[7] Negative prices occur because food banks are capacity constrained. The marginal value of an additional load of food to a food bank with an already full warehouse is negative. The extra load will likely spoil and have to be thrown away, which creates a bad image.

The introduction of a market mechanism had the potential to disadvantage smaller food banks. Credit use, joint bidding and fairness committee mechanisms were introduced to alleviate this risk. Smaller food banks often choose to bid jointly, because they might not need a whole truckload of a food. A small number of food banks bid jointly for more than half their winning bids. Otherwise, food banks rarely place joint bids. For this reason I generally ignore the decision to bid jointly. Discussion of how I consider joint bidding is given in Appendix A.1.3. Feeding America also

[5] The redistribution creates a small positive externality. For every share spent, an individual food bank will receive around 1/210 of that share, which is negligible.

[6] This aspect of the Choice System contributed to the supply of donations to Feeding America increasing drastically since the introduction of the mechanism. Feeding America itself would often turn down donations under the Old System, fearing that no food bank would accept the load.

[7] While 22% of lots are not sold right away, most are sold the following day. Lots not sold right away are predominantly either multiple loads of fresh produce or large bottles of water. The numbers are skewed by 130 loads of 8 litre bottles of water that were sold over several months.

employ a fairness committee to enable food banks to raise any problems they have with the Choice System. So far there have been no complaints, and food banks have universally reported great satisfaction with the Choice System.

When multiple homogenous lots are auctioned simultaneously lots are allocated to the top bidders (who pay their bid) until the lots have been exhausted. These auctions resemble discriminatory first-price auctions, rather than simultaneous auctions. 7% of auctions fall into this category. The main text ignores these auctions, while estimation and analysis does not. Details of how the model and estimation procedure are extended to account for these auctions is given in Appendix E.

Feeding America allows food banks to sell the food they receive from local donors, making up 4.5% of lots. There are several distortions in this submarket: For equity reasons Feeding America taxes and redistributes shares earned, reducing incentives for foodbanks to sell. Similarly, foodbanks have always happily shared excess food with one another for free and selling one's excess is looked down on by the foodbanks.[8] In this paper I generally ignore food banks' decisions to sell food. This is because selling food is rare, particularly for the food banks most reliant on the Choice System. Incorporating the decision to sell adds too much complexity to the analysis.

2.2 The Data

Three sources of data were used for this paper. The main data is the Choice System dataset, which is not publicly available and was received directly from Feeding America. I also make use of an auxiliary data-set enabling the identification of the locations of 85% of the food banks. Lastly, I use information from Feeding America's on-line food poverty tracker tool to estimate Food banks' goal factors.[9] Detailed discussion of how I clean and categorise the data are included in Appendix A.

2.2.1 Choice System data

The Choice System dataset contains information on $26,617$ individual auctions run over the course of 44 months from January 2014, covering 165 food banks. The data included both winning and losing bids from each food bank, as well as information

[8]Reported in Planet Money, NPR (2015).

[9]This tool is accessible at https://map.feedingamerica.org

on the food composition and location of each lot.[10]

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 I further split food into 164 subcategories.[11] To capture storage costs I categorise food into five storage types: Dried, Tinned/Bottled, Refrigerated, Fresh, and Non-Food.[12] Many loads contain multiple types of food. I allow lots to contain up to four different items which I assume evenly make up the load, unless explicitly stated otherwise.

Figure 1 presents descriptive statistics on the lots being allocated, split by storage method. Several things are evident. First, that many lots are allocated simultaneously. Second, that lots come in very variable sizes. Third, only a small number of bidders bid on any given lot, and a large proportion of lots sell for negative prices - particularly Fresh produce and low quality beverages (included in the Tinned storage type). This suggests low demand for these types of food.[13]

2.2.2 Auxiliary Data

Food banks in the main Choice System data were anonymised. Using data from Feeding America's Food Bank Locator tool[14] I identified the locations for 85% of food banks, who together consumed just over 98% of all food on the Choice System. Figure 3 shows the approximate locations of foodbanks (black spots) and the origins of lots coming to auction, by storage type.

I did not receive access to recent Goal Factor figures. However, this data can be

[10]Importantly, I do not observe whether any given auction happened in the morning or afternoon. I assume that all auctions in a day happen at the same time. This presents a potential weakness of this analysis, however anecdotal evidence suggests that most food banks bid in only one auction round each day. This was suggested by Canice Prendergast, one of the original designers of the Choice System. If food banks are optimally choosing not to bid on any auction in a given round then the inaccuracy of my results will be minor. In future I will use simulations to consider the substantive impact on my estimates.

[11]See Appendix A for additional discussion of how food was categorised.

[12]Fresh food includes produce and baked goods, that generally have limited shelf-life. Refrigerated includes anything that needs to be stored in a fridge or freezer, such as meat and dairy. Tinned and Bottled food includes anything with a long shelf-life that is tinned or bottled, ranging from baked beans to bottled water. Dry food captures long shelf-life food such as cereal, pasta, or cookies. Non-Food includes anything not considered food, including cleaning and beauty products.

[13]Figure 1 shows that very few bidders are observed bidding in any given auction. This may suggest auctions are uncompetitive. In practice it is unlikely that food banks collude, given how this harms non-colluding food banks and that most food bank managers are extremely prosocial.

[14]Accessible at https://www.feedingamerica.org/find-your-local-foodbank

Figure 1: Descriptive Statistics, across lots

	Dried	Tinned	Fridge	Fresh	Non-food	Mixed	Total
Daily lots							
(mean)	9.19	5.3	4.32	10.56	2.52	4.6	29.36
(std)	5.95	3.81	2.92	5.79	2.06	3.22	13.32
Pounds per lot							
(mean, 000s)	22.5	34.3	28.3	40.1	20.4	27	28.8
(std, 000s)	9.7	8.3	10.1	3.9	12.2	10.6	11.3
Winning bid							
(mean)	2106	1085	2704	211	2967	2481	1802
(std)	5329	6414	6331	779	6436	5176	5375
No. bidders							
(mean)	2.95	2.7	2.54	1.22	3	2.78	2.59
(std)	3.14	3.5	3.17	0.64	3.26	3.06	3.04
% Allocated	93	83	91	71	91	96	88
% Negative prices	35	47	28	19	28	27	32

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

constructed using the locations of food banks, formulae given in Prendergast (2022), and information on local poverty and food insecurity rates from Feeding America's 'Hunger in America' on-line resource. See Appendix A for additional details on how I located food banks and constructed Goal Factors. 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 this section I do two things. First, in Section 3.1 I present descriptive evidence of heterogeneity. That is, heterogeneity in the food being allocated, and heterogeneity in needs across food banks and over time. This evidence highlights the value of choice for food banks. Second, in section 3.2 I investigate the key determinants of bidding, putting together several stylised facts motivating my model's key features.

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	380	13	44	166	442	844
Average Bid	3601	546	1202	2509	4067	6903
Lots Won	159	6	26	70	177	351
Average Payment	3803	485	1060	2476	4507	7414

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

3.1 Evidence of Heterogeneity

Under the Old System every food bank was, ex ante, offered the same allocation. If food banks have heterogeneous needs, unknown to the social planner, welfare might be increased by allowing food banks greater choice in their allocations. Therefore heterogeneity is a key determinant of the value of choice.

3.1.1 Differences Across Lots

There is a large degree of heterogeneity across different types of lots. Lots attract significantly different bids depending on their subcategory.[15] Figure 4 shows average winning bids across subcategories, controlling for the censoring caused by the reservation price. These averages are generally statistically different from one another, and a Likelihood Ratio test that subcategory coefficients are equal within a category is rejected at 1% significance level for all but the Pasta category. However there is still much variation in winning bids within subcategories: Variation in subcategories accounts for just 30% of the variation in winning bids. It is clear there is a great deal of heterogeneity between lots, and that these lots cannot be substituted one for one.

[15]Lots also sell for significantly different prices depending on the category. Goods such as cereal and pasta sell for much higher bids than fresh produce and beverages. These differences cannot only be explained by differences in supply: Both cereal and ready meals are in abundance, and sell for relatively high prices. Meanwhile, Health/Beauty and Baked goods are rare, and sell for relatively low prices. This suggests both demand and supply factors at work in determining the prices.

Figure 3: Locations of lots and (approximate) locations of food banks

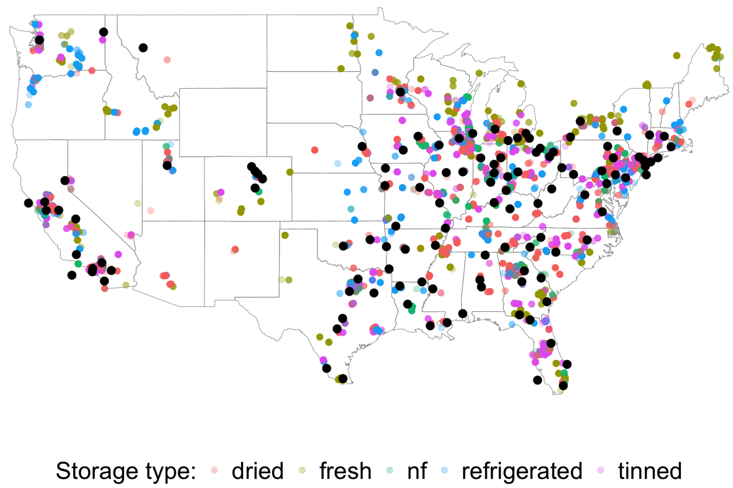

Note: Black spots give the approximate locations of food banks, jittered by an average of 200 miles. Excludes loads originating in Canada (2% of loads).

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. Therefore these five food banks are choosing to receive much cheaper food. This is likely because they rely on Feeding America for their staples, having fewer local donors than the other 122 food banks.

Figure 5 plots average bids and 95% confidence intervals across food banks and across different types of food. I focus on different types of food according to how they are stored. I use a Tobit specification to account for censoring caused by food banks only bidding on a small proportion of lots. Food banks are sorted according to their average bid for Dried food. That is, as one moves from left to right, the average 'Dried' bids increase monotonically. Several observations are clear. First, that there is evidence of systematic heterogeneity in average bids across food banks. Second, that there is evidence of systematic 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 Fresh food are higher than average

Figure 4: Heterogeneity in Lots

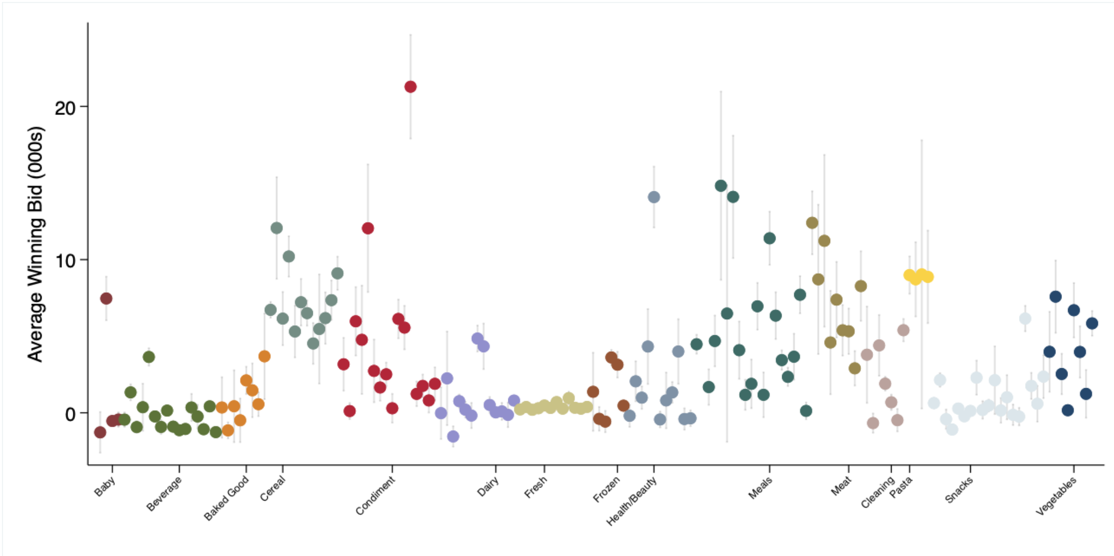

Note: Plots mean winning bids, and 95% confidence intervals, across subcategories, controlling for censoring and lot composition. Coefficients are ordered and coloured according to category.

bids on Dried food, but for other food banks this relationship is reversed. This demonstrates that bidding behaviour differs systematically across food banks.

3.1.3 Differences Over Time

To investigate the variation in bidding behaviour over time I run the same Tobit specification as above, considering how average bids vary from month to month. I focus on only those food banks who win at least 100 lots over my sample period, so that each food bank × type × month cell averages around 80 observations. I then consider the degree of variation in my estimated parameters. A likelihood ratio test that parameters are constant over time is rejected at 5% significance level for 96% of food banks. This is indicative of systematic heterogeneity in food banks' needs over time. Additional results are reported in Appendix B.

3.2 Stylised Facts

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

Figure 5: Heterogeneity Across Food Banks

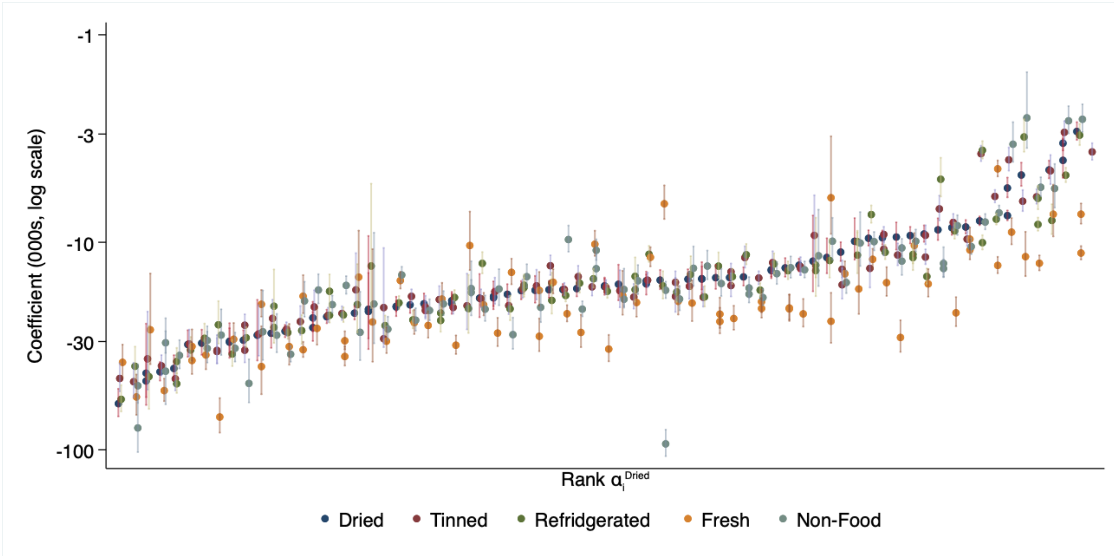

Note: The figure plots coefficients and 95% confidence intervals from a regression of Food bank × food storage type on bids, controlling for distance and censoring (non-bidding). Coefficients are ordered by the 'Dried' coefficients. Only includes results for food banks who placed at least 50 bids on a food type. Each food bank × type cell averages $3,450$ observations.

will become features of my model. In addition, I point to the importance of negative bidding, as well as both dynamic and static complementarities across lots.

3.2.1 Negative Bidding

Negative bidding is common: 27% of bids are negative. Furthermore, with a negative reservation price non-entry only happens when food banks have negative marginal valuations, when food banks must be paid to accept certain loads. This occurs in 98% of bidder × lot combinations. Negative valuations likely occur because of limited storage capacity, as emphasised in Prendergast (2017). They cannot throw away excess (non-expired) food as this sends a bad signal to donors. Therefore, I require a model that incorporates these storage costs and negative marginal valuations.

3.2.2 Dynamic Complementarities

Figure 6 panel (A) demonstrates that, conditional on winning food of a particular type at time 0, the probability of bidding on lots of the same type falls by around 25%

(1.5 pp) on subsequent days.[16] As food banks win more of a particular type of food, the less they are willing to pay for an additional lot from that type. Given that food banks are almost certainly forward looking, this finding highlights the need to model dynamics. Food banks treat these large loads of food like durable goods, working through their current stocks before returning to bidding on the Choice System.[17]

Figure 6: Evidence of Dynamic Complementarities

(A) (B)

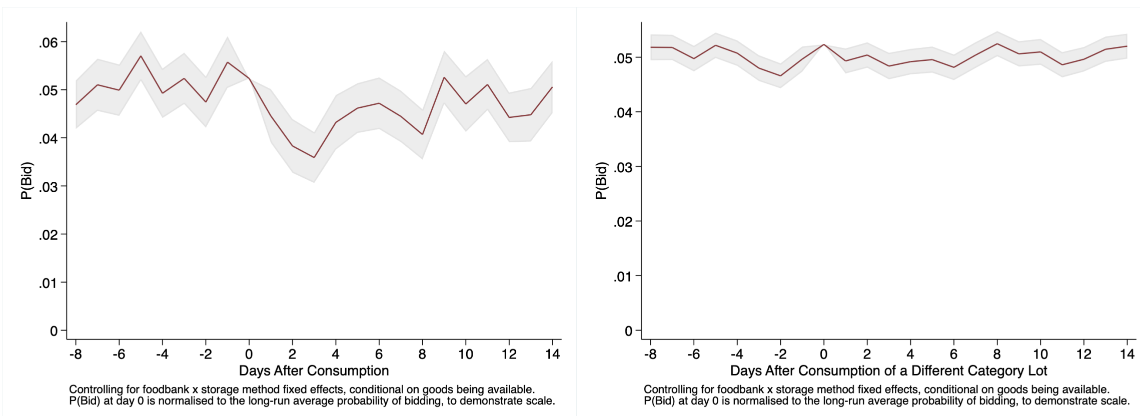

3.2.3 Static Complementarities

Figure 7 panel (A) demonstrates that, for a particular type of food, as the number of lots auctioned on a given day increases, food banks bid on a smaller proportion of lots. If pay-offs were additively separable we would see a horizontal line. This suggests that lots exhibit a negative complementarity (substitutes) within a storage type - they do not want to win more food than they can afford to store. I cannot treat auction

[16] This observation could be the result of transportation costs or budget constraints: Winning a lot at time 0 exhausted their budget, so the food bank must wait until they regain enough shares to begin bidding again. However this argument is inconsistent with the result from panel (B), that winning one type of food doesn't (meaningfully) impact the probability of bidding on a different type of food. The same argument holds for Figure 7 panel (B), given that the effect seen in Figure 7 panel (A) could also be caused by binding budget constraints.

[17] This alone does not necessarily require a dynamic model of forward looking agents. A dynamic model is only strictly required if the counterfactuals of interest sufficiently change the strategic environment, sufficiently altering the agents' continuation values. The counterfactual mechanisms I consider in section 7 are strategically very different from the Choice System, as food banks are generally less able to access the food they need than under the Choice System.

pay-offs as additively separable, and must instead take a multi-object approach, accounting for the simultaneous auction environment. Panel (B) demonstrates that we see a much weaker relationship as the total number of lots increases. This highlights the importance of treating different types of food as imperfect substitutes.

Figure 7: Evidence of Static Complementarities

(A) (B)

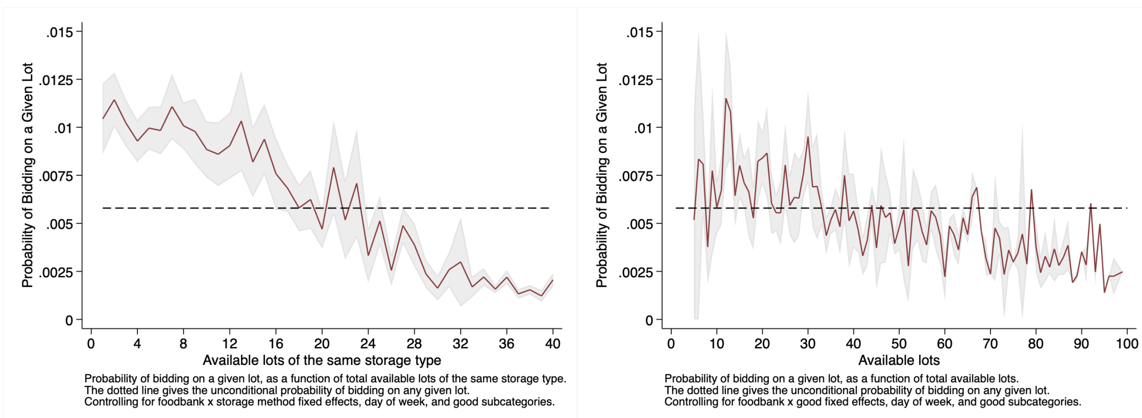

4 The Model

I now present the empirical model of food banks bidding in the Choice System. Section 4.1 introduces the market environment and the model primitives. Section 4.2 introduces the food banks' dynamic optimisation problem. Section 4.3 discusses the Markov Perfect Equilibrium and stationarity in this dynamic context. Finally, section 4.4 discusses Identification. Assumptions necessary for identification and feasibility of estimation are introduced as and when they are needed.

4.1 Market and Primitives

Each period t, over an infinite horizon, N food banks compete in up to L First-Price Sealed-Bid auctions. Food banks are denoted by i and lots are denoted by l. a is used to denote the combination outcome from a round of auctions. That is, which combination of lots food bank i won.

4.1.1 Auction Environment

Actions

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

Outcomes

Winners are announced simultaneously. Winners pay their bids, and every player observes the identities and bids of winners. Define player i's individual outcome vector $\mathbf{w}_{it}$ as the $L \times 1$ vector such that $w_{ilt} = 1$ if food bank i won lot l at time t, and zero otherwise. Ex-ante hypothetical outcomes are denoted by $\mathbf{w}^a_{it}$.

Lots and lot characteristics

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

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

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

4.1.2 Primitives

States

Food bank i begins the period in state $\mathbf{s}_{it} \in \mathbb{S}$. This represent the food bank's current stock of food. I primarily focus on their stocks from each storage method, so that the individual state has 5 dimensions.[18] This captures the dynamic costs of storing durable goods. If the outcome from period t is $\mathbf{w}_t^a$ they end the period in state $\mathbf{s}_{it}^a$. $\mathbf{s}_{it} = \mathbf{s}_{it}^a$ if and only if the player does not win a single lot. Writing $\mathbf{w}_{it}^T \mathbf{z}_t^g$ as i's winnings from period t I make the following assumptions about how stocks transition:

Assumption 2. *(i)* $\mathbf{s}_{it}$ *transitions according to the following process:*

$$\mathbf{s}_{it} = \mathbf{s}_{it-1} + \mathbf{x}_{it} + \mathbf{w}_{it-1}^T \mathbf{z}_{t-1}^g$$

(ii) $\mathbf{x}_{it} \sim F_i^x = N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ *is an exogenous innovation.*

I do not assume that individual states are observed. I assume that stocks are continuous.[19] Day-to-day variation in stocks is likely a major source of variation in bidding behaviour. Food banks supplement their stocks of one type of food they have not recently received from local donors with food from Feeding America. The random variable $\mathbf{x}_{it}$ is observed each morning before items are posted on the Choice System. It can be interpreted as the net daily change in food banks stocks - the food received from local donors, less the food given out to clients. Part *(ii)* of this assumption imposes that these changes are exogenous. Food banks don't turn down (or request additional) donations from their local donors.[20] Meanwhile, normality is reasonable for these large food distributors receiving many donations from many different sources, and sending out food to many different food pantries.[21]

---

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

[19] Difficulties in estimating dynamic models with continuous state variables is well known. However, continuous states are common in models of dynamic auctions. A previous version of this model discretised stocks. However the state had to be very finely discretised to capture all the possible combinatorial outcomes from a day's auctions.

[20] This simplification is likely to bias my results in favour of the Old System over the Choice System. If net donations were endogenous this would allow food banks to use their winnings to influence future net donations. Choice would be even more valuable. For example, if local donations were negatively correlated with previous winnings food banks could focus on only winning food from the Choice System they know they cannot get from local donors.

[21] This transition process incorporates two additional assumptions. First, food received from Feeding America, and food from local donors, are perfect substitutes. This is a necessary normalisation. Second, stocks do not degrade over time. This assumption was motivated by discussions with food bank volunteers. Most of the donated food, even fresh produce, have long shelf lives, so that any daily decay parameter is close to 1.

Pay-offs

Following Altmann (2022) and Gentry et al. (2020) I decompose the flow pay-off into a stochastic lot-specific component and a deterministic function of stocks:

Assumption 3. *(i) The flow pay-off from outcome a can be written as*

$$\mathbf{w}_{it}^{aT}\boldsymbol{\upsilon}_{it} + j(\mathbf{s}_{it}^{a})$$

(ii) The lot-specific pay-off $\boldsymbol{\upsilon}_{it}$ is a random variable with $\upsilon_{ilt} \sim F_i^{\upsilon} = N(\boldsymbol{\alpha}_i^T\mathbf{c}_{tl}, \sigma_l^2)$, known privately, observed before entry, and drawn independently across i and t.

(iii) The deterministic function $j : \mathbb{S}_i \to \mathbb{R}$ is finite, with $j(0)$ normalised to 0.

(iv) Pay-offs are quasi-linear in shares (virtual currency).

The flow payoff function j captures both the costs of storing food, and the utility from holding various types of food to be able to distribute them to clients. Part (*ii*) embeds two assumptions. Assuming the privately known $\boldsymbol{\upsilon}_{it}$ is conditionally independent across individuals imposes independent private values. Assuming conditional independence across time is a standard assumption in most dynamic models.[22] The assumption that j has finite range is predominantly for mathematical convenience, while the normalisation is required as only marginal pay-offs are identified. While I assume that pay-offs are quasi-linear in shares, as is standard in auction studies, I allow food banks to differ in their marginal value of wealth, given by $\lambda_i > 0$.[23]

I also assume players have temporally additively separable preferences, and make forward looking decisions with *annual* discount parameter $\beta = 0.99$, so that food banks are extremely patient. I assume F, j, $\mathbf{s}$, and β are common knowledge.

4.2 The Agent's Problem

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

[22] Note that the fact I don't observe the state variables is a violation of the conditional independence assumption. Assumption 2 is the weaker assumption required instead.

[23] Altmann (2022) shows that the quasi-linear model is observationally equivalent to a model with an inter-temporal budget constraint and constant marginal value of wealth. Constant marginal value of wealth is a reasonable assumption when food banks are sufficiently patient. This requires that day-to-day fluctuations in their budgets or stocks do not significantly impact expectations about how valuable accessing food from Feeding America will be in future.

a distribution of bids according to F_i, j_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 that they wins lot l, given their bid and entry decision, taking as given the strategies of other players. Denote $P_{ia}(\mathbf{b}, \mathbf{d}; \Lambda_{-i})$ player i's belief about the joint probability, conditional on $(\mathbf{b}, \mathbf{d}, \Lambda_{-i})$, that the outcome from the round of auctions is $\mathbf{w}_t^a$. These objects constitute food banks' beliefs about other players' equilibrium behaviour. In section 4.3 I make assumptions about these beliefs to make estimation feasible.[24]

4.2.2 Value Function

Assuming risk neutrality the bellman equation is given by:

$$W(\boldsymbol{v}, \mathbf{s}; j, \Lambda_{-i}) = \max_{\mathbf{b},\mathbf{d}} \{\Pi(\mathbf{b}, \mathbf{d}; \boldsymbol{v}, \mathbf{s}, j, \Lambda_{-i})\}$$

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

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

Continuation Value

The continuation value gives the expected pay-off from the start of the following period having ended the current period in state $\mathbf{s}^a$. This can be written as follows:

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

A further important object is the sum of the deterministic flow pay-off function and the discounted continuation value, denoted by $k(\mathbf{s}^a; \Lambda_{-i}) = j(\mathbf{s}_i^a) + \beta V(\mathbf{s}^a; \Lambda_{-i})$ and referred to as the 'Pseudo-Static' pay-off function. This is essentially the object one would estimate if one were to incorrectly assume bidders are myopic. Estimating this equilibrium object is key to estimating primitives in a dynamic multi-object

[24]Assumption 4, discussed shortly, ensures I only need to evaluate the 2^L probabilities of combinatorial outcomes for player i, not all N^L probabilities detailing which player won which lot.

model. The importance of this object arises because the value function (and hence the continuation value) can be written as functions of this pseudo-static pay-off:

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

4.3 Equilibrium

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

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

$$\textit{Optimality: } (\mathbf{b}_i^{\mathbf{\Lambda}^*}, \mathbf{d}_i^{\mathbf{\Lambda}^*}) = arg \ \max \left\{ \Pi(\mathbf{b}, \mathbf{d}; \boldsymbol{\upsilon}, \mathbf{s}, j, \Lambda^*_{-i}) \right\}$$

$$\textit{Consistency: } \Gamma_{il}(b_{il}, d_{il}; \Lambda^*_{-i}) = \mathbb{I}[d_{il} = 1] P(b_{il} > \max_{i' \neq i} \{b_{i'l}\} \, | \Lambda^*_{-i})$$

The optimality condition requires that agents maximise the net present value of pay-offs. The consistency condition requires that bidders' beliefs are consistent with the observed distribution of winning bids.[25] Symmetry requires that bidders with the same 'type', and the same beliefs, place the same bids. This allows us to write the equilibrium strategies as a function of the state: $\mathbf{\Lambda}^* = \mathbf{\Lambda}(\mathbf{s})$.

Altmann (2022) proved that, conditional on existence of a symmetric Pure Strategy Nash Equilibrium in the bidding game conditional on entry, such an equilibrium exists.[26] I make the following assumptions about equilibrium:

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

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

[25] This also requires bidders' beliefs about P are consistent with observed joint probabilities.

[26] Altmann (2022) also requires equilibrium pay-offs are continuous in $j + \beta V$. A full existence proof remains elusive. However this is not a practical problem. Numerous other papers studying sufficiently complex auction games are unable to guarantee neither existence nor uniqueness of equilibrium. This list includes, for example, Gentry et al. (2020) on simultaneous first-price auctions, Fox and Bajari (2013) on simultaneous ascending auctions, and Jofre-Bonet and Pesendorfer (2003) on dynamic single-object first-price auctions. The empirical strategy outlined in section 5 does not require existence of a MPE. Instead, it only requires that food banks have beliefs that are consistent with observed play.

(iii) $\forall$ *i and* $\mathbf{s}_i$ *the Hessian of the pseudo-static pay-off function k has full rank*

(iv) $\forall$ *i and joint outcome a* $P_{ia}(\mathbf{b}, \mathbf{d}|\mathbf{s}) = \prod_l \Gamma_{il}(b_{il}, d_{il}|\mathbf{s})^{w^a_{il}}(1 - \Gamma_{il}(b_{il}, d_{il}|\mathbf{s}))^{1-w^a_{il}}$

(v) $\forall$ *i, l,* b_{il} *and* d_{il} $\Gamma_{il}(b_{il}, d_{il}|\mathbf{s}) = \Gamma_l(b_{il}, d_{il}|\boldsymbol{\vartheta}(\{\mathbf{s}_i\}_N), \mathbf{s}_0)$

Part (*i*) is reasonably standard in studies of dynamic games, ensuring that the observed data is stationary. However, it embeds the stronger assumption that food banks' states are stationary. We expect $\mu_i < 0$; without access to Feeding America stocks will drift downwards over time. However food banks use the Choice System to supplement their stocks. When stocks get low, the food bank begins bidding to keep stocks up. This requires that, in equilibrium, food banks have enough control over their winnings to make this possible.[27]

Part (*ii*) of this assumption is required to ensure that standard first order conditions are necessary for optimality, so that primitives are point identified. I allow for the possibility of ties at the reservation price, which imply non-differentiability of Γ at R. Likewise, part (*iii*) is necessary for identification, as conditional on the function k it allows the first order conditions to be inverted for $\mathbf{s}_i$.

Part (*iv*) requires that, in equilibrium, food banks believe winning one lot is conditionally independent of winning any other lot. This essentially assumes that winning bids are conditionally independent across auctions, simplifying estimation considerably. In Appendix J.1 I test and present support for this simplification.

Part (*v*) is necessary for estimation to be feasible. 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 $\Gamma_{il}(.|\mathbf{s}_{t+1})$. As the number of bidders grows the probability of any individual and their state influencing prices falls to zero. This assumption ensures that equilibrium win probabilities Γ_i do not depend on the states of every player. Instead, they only depend on aggregate statistics of $\mathbf{s}$, using the aggregator $\boldsymbol{\vartheta}$ with known functional form.[28] For notational convenience I absorb $\boldsymbol{\vartheta}(\mathbf{s})$ into the

[27] Appendix C discusses the stationarity assumption in additional detail, demonstrating how we can test for stationarity. This assumption essentially requires that equilibrium winnings and net local donations are co-integrated, so that the *equilibrium* stock process is stationary. I demonstrate how we can test co-integration through equilibrium winnings without needing to observe stocks. Broadly, I find evidence of stationarity. Stationarity also requires that the distribution of net local donations is constant over my 3 year period. Feeding America's 'Hunger in America' resource shows that food bank usage and food insecurity remains stable over this period.

[28] This assumption presents a departure from both Jofre-Bonet and Pesendorfer (2003) and Gentry et al. (2020). It is similar to the large market Oblivious Equilibrium (Weintraub et al., 2008)

common state variable $\mathbf{s}_0$. This assumption also implies that we can write the value function, continuation value, and k, as functions of $\mathbf{s}_i$ and $\mathbf{s}_0$.

4.4 Identification

I now briefly discuss the identification of this model. A fully non-parametric proof goes beyond the scope of this paper. Altmann (2022) proves non-parametric identification of the model when all state variables are observed, also when individual states are not observed but the lot-specific value $\boldsymbol{\upsilon}$ is non-stochastic. In Appendix F I prove semi-parametric identification of the model, under the functional form assumption on k made in section 5.

Beliefs are identified trivially from the observed distribution of winning bids, conditional on lot characteristics. Parameters of the lot-specific pay-off distribution F^{υ} are identified from variation in lot characteristics, and how they are associated with differential bidding and entry decisions. The marginal utility of wealth parameters λ_i are identified by variation in the scale of bids across food banks, behaving in a similar manner to food bank specific lot-specific variances.

Permanent heterogeneity across food banks is identified from variation in bids and entry decisions across food banks. For example, we expect that food banks with fewer local donors or more clients (lower $\boldsymbol{\mu}_i$) to bid more frequently, while food banks with more variable local donations (larger $\boldsymbol{\Sigma}_i$) to bid more irregularly.[29] However this is not sufficient to identify the flow pay-off function j. Therefore, the model is additionally identified using two sources of variation in the data: First, variation in the set of lots being auctioned each day. Second, variation in observed winnings. Essentially, I make use of the same variation plotted in figures 7 and 6.

If the function j is very concave in $\mathbf{s}_i$, winning many lots simultaneously may be less desirable than smoothing out consumption over time. Therefore, when more lots are available in a given period, they will bid on a smaller proportion than if j was not so

and Moment-based Equilibrium (Ifrach and Weintraub, 2017), albeit in a game of incomplete information. Backus and Lewis (2016) employ a similar assumption in their dynamic auction framework. They argue that because there are many competitors it is unlikely that bidders follow the identities of which other bidders are likely to bid at any given time, and their states. It is unlikely that any given food bank keeps track of competitors' states. This assumption is tested on the empirical equilibrium winning probabilities in Appendix J.1.

[29] $\boldsymbol{\Sigma}_i$ is separately identified from variance parameters of F^{υ}, through time-series variation and auto-correlation, since stock shocks persist, while lot-specific shocks do not.

concave. This rests on Assumption 1 to ensure that the composition of lots available is conditionally exogenous. Likewise, if bidding generally stops altogether after a win, this suggests the food bank has very low capacity and j is again strongly concave. These two effects also allow us to tease out complementarities across different types of lots. For example, we can observe how bidding behaviour on milk is affected by cereal availability, or a recent cereal win.

The distribution of net donations is identified from the effect of winnings on bidding behaviour and how long these effects persist. If a food bank always returns to bidding quickly after a win this suggests they give out much of this type of food to clients, or receive little of this food from donors. Likewise, the degree of variation in this recovery time yields information about the variance $\mathbf{\Sigma}_i$. Importantly, these parameters are only identified conditional on the pseudo-static pay-off function k, since the state variable only influences bidding behaviour through k.[30]

5 Empirical Stategy

This section describes the estimation procedure used to estimate the model. Section 5.1 outlines the three step procedure, noting the relationship to the procedure of Jofre-Bonet and Pesendorfer (2003). Section 5.2 discusses parametrisation and estimation of beliefs, which are estimated using a likelihood procedure. Section 5.3 discusses the second estimation step, in which I simultaneously estimate the state transition process, the distribution of lot-specific values, and the pseudo-static pay-off function. In section 5.4 I detail how I disentangle the combinatorial flow pay-off j and the discounted continuation value from the pseudo-static pay-off. Full details of the estimation procedure are given in Appendix H.

5.1 The 3-Step Procedure

The standard approach to estimating dynamic auction games, from Jofre-Bonet and Pesendorfer (2003), relies on the ability to write the continuation value as a function

[30]This argument is similar to the identification argument of Berry and Compiani (2020). Identification of this model does not require randomisation nor strict exogeneity. Instead, conditional weak exogeneity is sufficient. The bidder's winnings in periods $t' < t$, as well as the set of available lots at time t, will affect their bidding behaviour at time t. I assume, plausibly, that there is no contemporaneous reverse causation; their bidding behaviour at time t does not impact the set of lots available at t, nor previous winnings.

of the distribution of bids *only*. This is not possible in the multi-object context because of an order problem: Bids are L dimensional, while values, and continuation values are 2^L dimensional. Full solution methods of Rust (1987) are computationally intractable in this setting - recursively evaluating the value function requires numerically maximising bids for each $\mathbf{s}, \boldsymbol{\upsilon}_i$.

Instead, Altmann (2022) demonstrates that we can write the continuation value as a function of the distribution of bids *and* pseudo-static pay-offs. If we know the pseudo-static pay-offs we can find the continuation value, which then allows us to back out the flow pay-off j from the definition of pseudo-static pay-offs: $k = j + \beta V$. To estimate the pseudo-static pay-off function we estimate the model as if we were estimating a static model, but allow pay-offs to depend on $\mathbf{s}_0$.

This procedure bears a strong relationship with the standard method of Jofre-Bonet and Pesendorfer (2003), and even Conditional Choice Probability (CCP) methods of Hotz and Miller (1993). Both methods involve first estimating non-primitive objects (CCPs and bid distributions), using these objects to back out the continuation value. Furthermore, this approach is numerically equivalent to a CCP approach in a dynamic discrete choice context.

5.2 Step 1. Beliefs

Assumption 4 ensures food banks form beliefs consistent with observed play. Therefore, we can estimate beliefs using the observed distribution of winning bids. Estimating beliefs in this way avoids the need to solve the model for equilibrium. This procedure is common in the empirical auction literature due to the extensive computational cost of finding equilibrium beliefs (Athey and Haile, 2007).

I make parametric assumptions about Γ to facilitate estimation. I assume winning bids follow a generalised extreme value distribution, censored at the reservation price:

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

Where the shape and scale parameters ξ and ζ are category specific. $\mathbf{c}_{lt}$ gives a vector of lot specific location shifters, such as the subcategory composition.

d_{lt} describes how the distribution varies with the state of the world, forming an index to be estimated. The index is a linear function of the quantity of food, by usage type, auctioned at t and also the quantity over the previous 30 days, up to

$t-1$. This is designed to capture competitive pressures on prices. If very little food has been auctioned over the previous month, one would expect a higher price. Estimating this demand index in the first stage allows us to use the estimated index in later estimation objects. In particular, when considering the transition process of common state variables, I can focus on just the transition process of d_{lt}.

The Fisher–Tippett–Gnedenko theorem establishes that the Generalised Extreme Value distribution is the limiting distribution of the maximum of independently distributed random variables. In an Independent Private Value framework, the winning bid is just the maximum of (conditionally) independent random variables. Therefore the GEV assumption is easily justified. Meanwhile the parametrisation is chosen to be suitably flexible, given the available data. Full details of how I estimate beliefs are included in Appendix H.1.[31]

5.3 Step 2. The Pseudo-Static Model

I now describe the second part of the estimation procedure, in which I jointly estimate $F^{\mathbf{x}}, F^{v}, k$ for each food bank. I begin by discussing the problems that must be overcome in this estimation step, before detailing the parametric assumptions made to enable estimation of each of the three sets of objects. The key functional form restriction is assuming that the pseudo-static pay-offs are quadratic in stocks, similar to the standard assumption of quadratic storage costs. Estimating the second-stage then requires estimating a censored Linear Gaussian State Space model.

The central estimation difficulty concerns the unobserved state and the unobserved bids. Bids are unobserved when a food bank chooses not to enter a particular auction, which occurs frequently. Ignoring these bids introduces the standard problems of censoring in econometrics. I estimate the model using a Gibbs Sampling procedure. I use data augmentation to iteratively sample both unobserved bids and unobserved states from their conditional posterior distributions, before updating my parameter estimates given the augmented data. Full details of the estimation procedure,

[31]This details additional covariates included in estimation. Also, how I estimate the probability of tieing at the reservation price. If no food bank bid on a lot, then a food bank would have won if they bid the reservation price. But, if some other food bank won at the reservation price, then the food bank would have tied had they bid the reservation price. Importantly, allowing for ties rationalises food banks choosing to bid just above the reservation price.

including assumptions on prior distributions, are given in Appendix H.2.[32][33]

5.3.1 Individual States

I estimate individual $\times$ storage type specific mean and variance parameters (μ_{ig}, Σ_{ig}) for the normally distributed net local donations $\mathbf{x}_{it}$. I make use of the prior information from Assumption 4, which requires the stock transition process is stationary. In Appendix C.4 I prove that stationarity requires:

$$\boldsymbol{\mu}_i = -E[\mathbf{w}_{it}^T \mathbf{z}_t^g] \qquad \boldsymbol{\Sigma}_i < 2Var[\mathbf{w}_{it}^T \mathbf{z}_t^g]$$

On average winnings must offset mean net donations and the variation in winnings reflects the variation in net donations. However I do not impose that either relationship holds exactly, and instead use them to build informative priors.[34] I have a standard initial state problem. The quadratic assumption I impose on k ensures the level of the state is not identified, so I normalise the initial state to zero.

5.3.2 Lot-Specific Pay-offs

I specify the mean of the lot specific pay-off υ_{ilt} as $\alpha_i distance_{ilt}$, so that the mean lot specific pay-off depends linearly on the distance between food bank i and lot l. The variance σ_l^2 is category combination specific. Assumption 3 imposes that the lot specific pay-offs are uncorrelated across t and i. To simplify estimation I also assume these variables are conditionally uncorrelated across lots l.

[32] Because heterogeneity is an important theme of the model I generally estimate separate parameters for each food bank. However, I do not always have enough identifying variation for each individual food bank. I use a Bayesian Hierarchical framework to flexibly introduce information pooling across bidders in my model. This approach is flexible enough to allow pooling for food banks that lack identifying variation, placing more weight on the hierarchical parameters, and allow separation for food banks that have a lot of identifying variation, placing more weight on the data.

[33] Due to computational requirements I focus my estimation on the 34 food banks that each won at least 150 lots (Type 1 food banks). These food banks consume 70% of the food from the Choice System. It is standard in empirical auction studies to estimate a main model and a model of 'fringe' bidders (For example Jofre-Bonet and Pesendorfer (2003) and Gentry et al. (2020)). I estimate a simpler (myopic) model for the remaining 88 food banks who won at least 30 lots and whose locations are known (Type 2 food banks). Details of this model and estimation is included in Appendix H.4. All counterfactual analysis uses the models from both sets of food banks.

[34] This is because the relationships only holds, averaged over time. Winnings $\mathbf{w}_{it}^T \mathbf{z}_t$ are both extremely 'lumpy' as well as auto-correlated, so that convergence to the true mean is slow.

5.3.3 Combinatorial Pay-offs

I fit a parametric form to the function $k(\mathbf{s}_i, \mathbf{s}_0)$. Within feasibility constraints, I choose a parametric function to reflect how food banks gain benefits from food according to how the food is used (according to it's subcategory) and how they face costs of storing the food (according to the storage method). I assume the following:

$$k(\mathbf{s}_i) = \Phi \mathbf{s}_i^h + \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g \tag{3}$$

Where Φ is an 1×164 row vector, and Ψ_i is a 5×5 dimensional matrix. The form of k as I have presented it above depends on both the stock of each storage type $\mathbf{s}_i^g$ and the stock of each subcategory $\mathbf{s}_i^h$. However, consider the marginal pseudo-static pay-off from winning lot l with characteristics $\mathbf{c}_l$:

$$k(\mathbf{s}_i + \mathbf{z}_l) - k(\mathbf{s}_i) = \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT} \Psi_i (\mathbf{z}_l^g + 2\mathbf{s}_i^g)$$

This does not depend on $\mathbf{s}_i^h$, because the hessian of the pseudo-static pay-off, with respect to $\mathbf{s}_i^h$, has rank 0. This means that $\mathbf{s}_i^h$ 'falls out' of the model, so that I can focus on $\mathbf{s}_i^g$ as the state variable.

In theory k should depend on $\mathbf{s}_0$, capturing how the continuation value depends on food banks' beliefs about future supply. If they believe many lots will be auctioned next period, prices are likely to be low in future, lowering the opportunity cost from *not* winning today. In practice I assume k is independent of $\mathbf{s}_0$ for two reasons. First, the supply of shares varies with supply to ensure prices remain approximately constant over time. Therefore we expect little variation in average prices over time. However, relative prices may vary with supply and this may impact the continuation value. As I show in Results section 6.1, the relationship between supply of different types of food and prices is not economically significant. Nonetheless, in Appendix J.2.1 I present results from an econometric specification that includes d_{ltg} (the demand index for food type g) as an input to k.

For computational tractability I impose that Φ is constant across i. Allowing it to vary introduces an unwieldy number of parameters to the model. This is a reasonable assumption as food banks likely gain the same benefit from different subcategories.

Three Equations

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

$$\mathbf{s}_{it}^g = \mathbf{s}_{it-1}^g + \mathbf{x}_{it} + \mathbf{w}_{it-1}^T \mathbf{z}_{t-1}^g \qquad \rightarrow \textbf{Transition Eq.}$$

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

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

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 $\mathbf{s}_{it}^g$. Therefore the model is a case of a Censored Linear Gaussian State-Space model.[35]

5.3.4 Estimation procedure

Unlike the non-censored case, the likelihood of the Censored Linear Gaussian State Space model is intractable. Instead, estimation is performed using a Gibbs Sampler, which consists of the following steps:[36]

1. Draw beliefs Γ from their posterior distribution using Metropolis Hastings

[35] This Observation Equation is endogenous - it contains a dependent variable on the right hand side, through b_{itm}. In general, b_{itm} may be correlated with υ_{itl} - when υ_{itl} is large, food bank i may prefer to win lot l instead of lot m (assuming negative complementarities), so lower their bid on lot m. In practice, however, simulations suggest the inconsistency caused by this endogeneity is very small, as $\Gamma_{im}(b_{itm})$ is generally very unresponsive to υ_{itl}, depending much more on $\upsilon_{itm}, \mathbf{z}_{itm}$ and even $\mathbf{z}_{itl}$. In Appendix J.2.2 I use the instrumental variable procedure of Altmann (2022), using $\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g$ as an instrument for $\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2\sum_{m \neq l} \Gamma_{im}(b_{itm}) \mathbf{z}_{tm}^g$.

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

2. Given Γ, the parameters of the pseudo-static model $\{k_i, F_i^{\upsilon}, F_i^{\mathbf{x}}\}_N$, and states $\{\mathbf{s}_{it}^g\}_{T,N}$, draw censored values of $\{y_{ilt}\}_{NTL}$ using the Censoring Equation

3. Given Γ, $\{k_i, F_i^{\upsilon}, F_i^{\mathbf{x}}\}_N$, and $\{y_{ilt}\}_{NTL}$, use the Carter-Kohn Algorithm to draw $\{\mathbf{s}_{it}^g\}_{T,N}$ using the Transition and Observation equations.

4. Given Γ, $\{y_{ilt}\}_{NTL}$ and $\{\mathbf{s}_{it}^g\}_{T,N}$, draw $\{k_i, F_i^{\upsilon}, F_i^{\mathbf{x}}\}_N$ from their posterior distributions using the Observation Equation.

5. Repeat

Additional details of the estimation procedure are given in Appendix H.2.

5.4 Step 3. The 'Dynamic' Game

At this point we have draws of beliefs, $\{k_i, F_i^{\upsilon}, F_i^{\mathbf{x}}\}_N$, and $\{\mathbf{s}_{it}^g\}_{T,N}$ from the posterior distribution. I now describe how I evaluate the continuation value $V(\mathbf{s}_i^g, \mathbf{s}_0)$. I make use of the following proposition:

Proposition 1. *The ex-ante Value Function can be expressed as:*

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

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

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

This proposition is proven in Appendix G. The identity $\pi(\mathbf{b}, \mathbf{d}|\mathbf{s}_i^g, \mathbf{s}_0)$ arises from substituting the first order conditions back into the maximand, writing the ex-ante value function as a function of bids and the pseudo-static pay-off function. The main proof then extends the key result from Arcidiacono and Miller (2011) to the continuous choice case. The sample counter-part to this object is then easily found. Full details of this procedure are given in Appendix H.3. I evaluate the ex-ante value function across a grid of states. I use a 20^5 grid evaluated evenly across points from the posterior sampled states.[37]

[37] Such a large grid is feasible in this context as I only need to perform the procedure once.

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

6 Estimation Results

This section presents the results from the three stages of estimation described in section 5. Only a subset of key results are reported in the text, focusing on the theme of heterogeneity and only presenting results for the 34 largest ('Type 1') food banks, who consume 70% of the food on the Choice System. Full results are reported in Appendix I, including Gelman-Rubin convergence statistics. When discussing statistical significance I focus on 95% credible intervals. I present several graphs plotting the individual parameters, and credible intervals, across food banks.[38]

6.1 First Stage Results

The key parameters estimated in the first stage are the shape, scale, and location parameters that describe the generalised extreme value distribution. The Shape parameters lie significantly within the interval $(-0.1, 0.5)$, with none of the parameters significantly below zero. The scale parameters are all estimated to be 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 (around 0.3 on average) 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 4 ($R^2 = 0.74$). The standard deviation of posterior means across subcat-

However, storing $34,000$ grids, one for each food bank $\times$ parameter draw, is not. I use a quadratic approximation of the ex-ante value function. In Appendix J.3.1 I evaluate the fit of this approximation by considering the R^2 of the approximation regression. 100% of these R^2s lie between 0.99 and 1. The fit is strong due to the quadratic term that appears in the ex-ante value function.

[38] Appendix I.6 discusses model fit, both in and out of sample. Broadly, the model fits the data well, matching average patterns of consumption across food banks and food types, as well as average propensities to place bids across food banks, categories of food, and months. The simulated distribution of bids conditional on entry does not fit the data as well, failing to match the observed long right tail of bids and over-estimating the mean and standard deviation of bids by a magnitude of 50%. However this is not a major problem, since for my counterfactual exercises it is food banks' allocations that matter, rather than their signal of preferences.

egories is 2400, while the mean posterior standard deviation is 800. This suggests much more variation across subcategories than uncertainty about subcategory posterior means. The previous 30 days supply of food is estimated to have a significant negative effect on prices for every type of good except Non-Food and Condiments. The coefficient on Meals is the largest, estimating that each additional increase in the previous 30 day supply by one thousand tons (approximately one hundred loads) decreases the winning bid by 350 shares. This magnitude, while statistically significant, is not economically significant (around 0.017 standard deviations), relative to the variation seen across different types of food through the subcategory parameters. The present day's supply of food is not estimated to have a significant effect on prices, however these estimates are noisy.

6.2 Second Stage Results

6.2.1 Unobserved State

Figure 8 plots the estimated μ_i and $\sqrt{\Sigma_i}$ parameters for each of the Type 1 food banks. 95% credible intervals are also plotted. Estimates are sorted according to the estimates for the Dried food type.

There are two key takeaways from these results: First, the extent and significance of the heterogeneity. The variation across food banks in the distribution of net donations of for Fresh food is particularly stark. It is also clear that the variation is not purely vertical: Some food banks have higher estimates for Dried than Tinned, while other food banks exhibit the opposite relationship. The second key takeaway concerns the differences in the scale between the two sets of graphs - the standard deviations are generally larger than the means, so we expect the unobserved state process to be noisy.[39]

6.2.2 Lot specific pay-off

The key lot specific parameters are the coefficient on distance between food bank i and lot l, the constant marginal value of shares λ_i, and the standard deviation of the

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

Figure 8: Estimated unobserved state parameters

(A) (B)

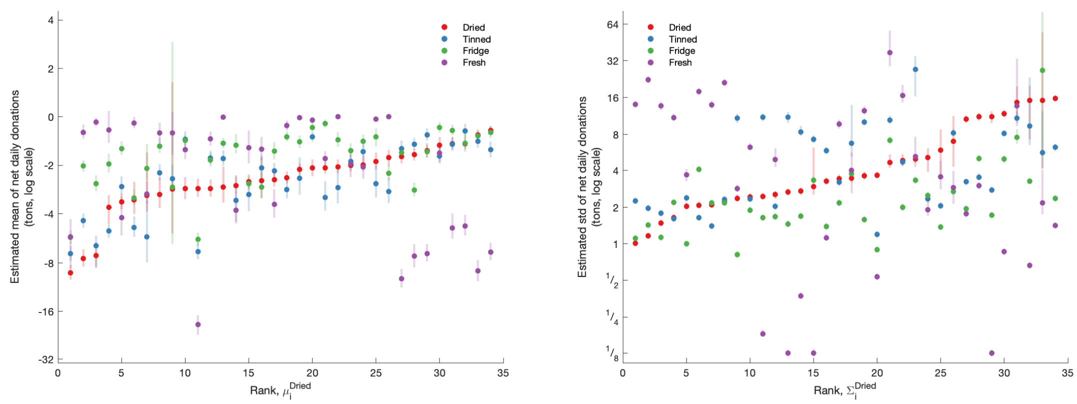

Note: The figure plots posterior means for the mean and standard deviations of net local donations, as well as 95% credible intervals. Results are sorted according to the estimates for the Dried storage type. The plot excludes Type 2 food banks, and the 'non-food' type, to improve graphability.

idiosyncratic lot specific shock υ_{ilt}.

Distance coefficients vary across food banks from a cost of 5 to 98 shares per km, with an average of 23. λ_i for the Type 1 food bank with median consumption is normalised to 1. The posterior means then vary from 0.5 to 5. I observe a negative relationship between the shadow price of shares and goal factor - food banks with a higher goal factor receive more shares. However this relationship is very weak, stressing the importance of unobserved food wealth. Standard deviation parameters are estimated to be large, with a mean posterior mean of 32,000. Such large standard deviations are needed to rationalise the small probability of bidding ($\approx 2\%$) with the relatively large variation in bids conditional on bidding.[40]

6.2.3 Combination pay-off

The estimated Φ parameters, associated with the marginal value of winning a pound of each subcategory, are strongly correlated with the first stage subcategory param-

[40] I decompose the residual variation into idiosyncratic variation in the lot specific value, and unobserved variation in the state. Condition on observing the previous period's state, 45% of this short-run variation is due to variation in the unobserved state, and 55% due to lot specific variation. When we consider long run variation, 72% of the unobserved variation is estimated to come from variation in the unobserved state.

eters ($R^2 = 0.82$). Panel A of Figure 9 plots food banks' willingness to pay for an additional 40,000 pounds from each storage type, evaluated when stocks are zero.[41] I estimate significant variation across food banks, and within food banks across storage types. The willingness to pays are generally (significantly) negative, as expected. I find broadly low figures for Dried food, which is driven by food banks generally bidding on fewer fresh items at a time, and bidding on fewer items after winning.

Figure 9: Estimates of Ψ_i and j_i

(A) (B)

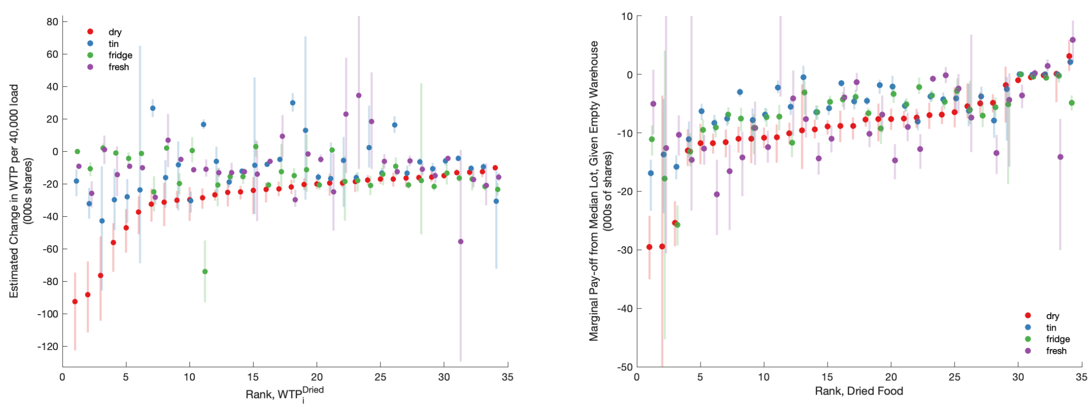

Note: Figures plot posterior mean equilibrium willingness to pay (A) and marginal flow pay-offs (B) for a 40,000 load for each storage type. Bars give the 95% credible intervals. Estimates are ordered according to the estimates for Dried loads. The plot excludes Type 2 food banks and estimates for non-food storage type. WTPs and marginal flow pay-offs are evaluated when stocks are zero.

6.3 Third Stage Results

Figure 9 panel (B) plots posterior means of the marginal flow-payoffs from receiving 40,000 pounds, evaluated when stocks are zero. Estimates are plotted for Type 1 food banks, sorted according to the estimate for the Dried storage type.

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

[41] Due to the normalisation made in section 5, stocks on March 1st, 2014 were normalised to zero, as the level of stocks is not identified.

empty warehouse. The results are broadly similar to those plotted in figure 9 panel (A) as one would expect. Because stocks are persistent, the continuation value from storing food tomorrow should be similar to the cost of storing food today. However, the absolute magnitudes in panel (B) are generally smaller than in panel (A). This is because the pseudo-static pay-offs account, not just for present storage costs, but also expected future storage costs and the future *Opportunity Cost* of storage.

To summarise, I estimate that food banks differ systematically in their net local donations, storage costs, transportation costs and marginal value of shares. This suggests that giving food banks choice over their allocations will be welfare improving.

7 Counterfactuals

Feeding America introduced the Choice System, replacing the Old System, in order to give food banks choice over the food they received. Feeding America also put significant resources into minimising the possible costs from the introduction of this system. They were worried that it might lead to an inequitable distribution of food, allowing smaller food banks to 'fall through the cracks'. Motivated by the introduction of this new system, this section investigates the welfare and distributional consequences of introducing the Choice System. I then consider a number of additional mechanisms used by other food bank networks around the world.

Section 7.1 briefly explains how I simulate equilibrium allocations under the Old System. Section 7.2 presents the results from this counterfactual exercise, as well as presenting descriptive analysis to understand what factors are driving these results. Section 7.3 introduces several additional allocation mechanisms and presents results from these additional counterfactual exercises. Additional details of how I simulate the counterfactual mechanisms and the Choice System are given in Appendix K.

7.1 The Old System

I model the Old System in continuous time. This is realistic since food banks could receive a call from Feeding America at any time. I also assume that food is given out to clients and received from local donors at random times during the day. Continuity of time ensures that the probability of a call from Feeding America, or local donors, occur simultaneously with probability almost surely zero. I assume food banks do not observe offers made to, nor decisions of, other food banks. They do not know their place in the queue; only their own Goal Factor, and when they were last offered a load. I assume they form beliefs about the rate they receive calls from Feeding America, and also the probability of being offered a load with characteristics $\mathbf{c}_l$, conditional on receiving a call. I assume these objects are independent of the time since their previous offer. In practice, given the frequency and irregularity with which food banks are offered food (on average, around 5 times per day with a standard deviation of 8) this simplification is unlikely to cause significant inaccuracy.

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 their beliefs, and that beliefs are consistent with the observed realisation of offer rates. Appendix K.1 details how equilibrium beliefs and equilibrium value functions are formed. Given beliefs I find each food banks' value function by numerically solving the Hamilton-Jacobi-Bellman differential equation. I then simulate the mechanism and update beliefs using observed offer rates, repeating until convergence.

7.2 Results

7.2.1 Welfare

My counterfactual simulations produce welfare measures in terms of consumer surplus, measured in shares. This has the benefit that consumer surplus is a cardinal measure, enabling inter-food bank comparisons. However the value of shares is difficult to interpret as they have no value outside the Choice System. Instead, similar to Agarwal et al. (2021), I report welfare as the equivalent increase in the supply of food that would have the same total value in shares.[42] This measure is valid under competitive equilibrium because the money supply adjusts to ensure prices are constant, given changes to the supply of food. Therefore, if consumer surplus under the Old System is double that under the Choice System, I liken this to double the nominal expenditure, which equates to double the supply of food.

[42]This is similar to how consumer surplus is typically measured in dollars, except that here I am measuring it in terms of how much food those dollars could purchase. I could also use distance travelled as a numeraire, reporting the equivalent reduction in the total distance travelled. Prendergast (2022) then measures welfare in dollar terms using estimates of trucking costs from the literature. Given my results that different food banks face very different transportation costs, as shown in Appendix I.2, this approach is unattractive.

Importantly, the 'level' of welfare is not identified because the levels of both stocks and flow payoffs $j(\mathbf{s}_i)$ are not identified. I use a random allocation as a benchmark counterfactual. This is a relevant benchmark since it can be considered the baseline worst case allocation mechanism. Results are reported on a scale of zero (food is allocated no better than random) to 292 tons (the daily average amount allocated under the Choice System).

I report both utilitarian welfare and a weighted sum using Goal Factors as priority weights. I also report descriptive measures of welfare. For example, the total amount of food allocated. This is an important measure given the political cost to Feeding America from being seen to waste food, or the indirect harm from donors being less likely to donate again in future.[43] The distance food must travel is another key metric. We expect food banks to sort on location as food banks choose nearby lots.

7.2.2 Importance of Choice

Figure 10 presents the posterior mean and 95% credible intervals for various measures under the Choice System and Old System. The first column gives the un-weighted sum of estimated welfare in equivalent tons of food. All welfare results are relative to the baseline random allocation. The mean welfare under the Choice System is mechanically equal to the average daily amount of food. While both Systems achieve significantly more welfare than a random allocation, the Choice System yields significantly higher welfare than the Old System. Welfare is on average 17.1% higher under the Choice System than the Old System, which is enough food to provide an additional 22,300 meals each day. When welfare is weighted according to Goal Factor, this figure increases to 22.9% higher under the Choice System. These results are extremely similar to those in Prendergast (2022), who finds that welfare is roughly 21% higher under the Choice System.

The third column shows that, under the Choice System, food banks sort into consuming closer lots, with around 6,000 km less transportation required each day. This is in spite of the result from the fourth column that around 22 additional tons of food is accepted each day under the Choice System.

[43]The welfare measures reported in this system do not account for endogeneity in the supply of food, with respect to the allocation mechanism employed. This relationship is unfortunately not identified, given that I only observe data from the Choice System. However, given that my results generally show more food being accepted under the Choice System, this simply means my results

Figure 10: Counterfactual Results

Mechanism	Welfare (unweighted)	Welfare (weighted)	Distance (000 km per day)	Allocated (tons per day)
Choice System	292 (276, 309)	745 (672, 815)	16 (14.6, 17.4)	271 (253, 284)
Old System	242 (203, 276)	576 (415, 706)	22.3 (22, 22.6)	249 (248, 251)

Note: This table displays posterior means and 95% credible intervals for various measures of welfare. Welfare is measured in food equivalent terms relative to a purely random allocation, and mean welfare of the Choice System is normalised to 292. Therefore welfare should be interpreted as pegged to this scale of 0 (as good as random) to 292 (as good as the Choice System).

When I decompose these welfare differences into the stock dependant component $j(\mathbf{s}_{it})$ and the lot specific component v_{ilt} (which contains transportation costs) we see that 81.5% of the welfare gains come from the stock dependant component. Reduced transport costs account for 6.53% of the gain. This is because transport costs are only estimated to be major cost for a small number of food banks. The additional food that is accepted under the Choice System only explains 1.72% additional welfare, as this food is typically lower quality. The remainder is attributed to food banks sorting into food with higher unobserved idiosyncratic payoffs.

A likely driving force behind these results is that under the Old System food banks accept food that does not meet their most pressing needs at that particular point in time. They accept food that might be more useful to a different food bank, and may prevent them from accepting food they value more in the near future.[44] This is evident for three reasons. First, most food banks are still offered enough food to prevent their stocks from trending downwards. Therefore it is not that food banks do not receive enough food. Second, stocks are more variable under the Choice System than the Old System. The average short run variance is around 25% higher under the Choice System, and the long variance is almost 100% higher. So, it is not about food banks smoothing their stocks. As evident from the main result, their stocks spend more time close to the optimum - the maximum of $j(\mathbf{s}_i)$. Therefore food banks are choosing to have more variable stocks, occasionally increasing stocks for food that

can be interpreted as lower bounds on the value of choice.

[44] Some food banks also accept food they would not have accepted under the Choice System just in case they need the food in future. Whereas under the Choice System they know they will be able to bid on this food when they need it, rather than having to wait and hope they will be offered it.

is particularly valuable. Finally, under the Choice System food banks were free to bid zero on the types of food they were offered in my counterfactuals, but chose not to. This revealed preference implies that, often, they only want certain types of food at certain times. Under the Old System they just have to wait to see what they are offered. Their stocks remain at consistent levels, suggesting they accept different types of food instead, accepting food that is just good enough.

7.2.3 Distributional Consequences

Figure 11 presents welfare results by food bank, plotting the difference between food bank specific welfare under the Choice System and the Old System. On average 85% of food banks are better off under the Choice System.

Figure 11: Individual Welfare

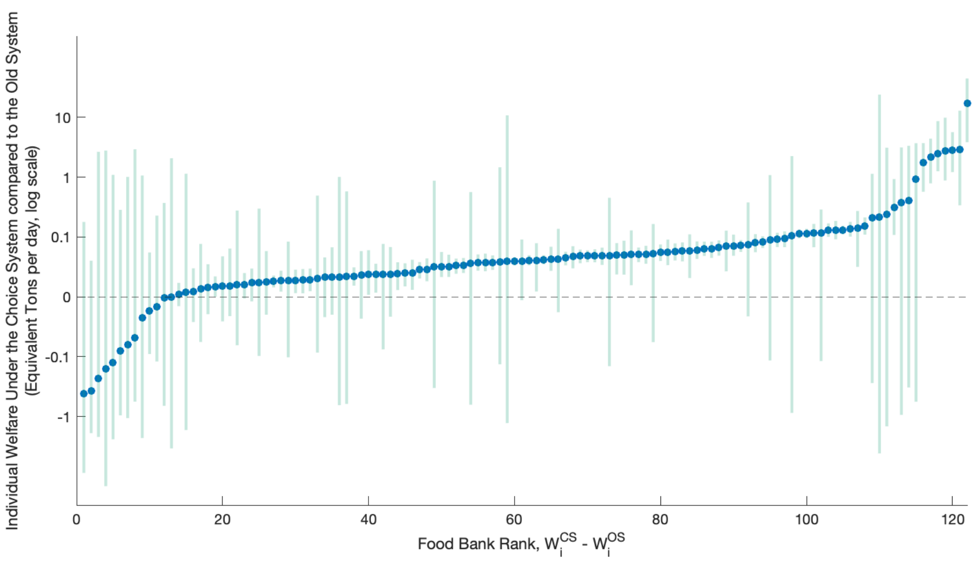

Note: Plots food bank specific welfare under the Choice System minus welfare under the Old System, ordered by the welfare difference, with 95% credible intervals across posterior draws. On average 85% of food banks are better off under the Choice System than the Old System.

Given the difference between the weighted and unweighted welfare estimates, it is unsurprising that there is a positive correlation between this welfare difference and Goal Factor ($\rho = 0.184$). I find negative correlation of -0.205 with estimated λ_is, the marginal value of wealth. This suggests that food banks who rely less on food from

Feeding America benefit more from choice, from being able to be picky.[45]

Relative welfare is negatively correlated with food banks' mean net donations $\boldsymbol{\mu}_i$ ($\rho \approx -0.2$). This is driven by food banks in the tail of the distribution of net donations, particular those whose stocks trend downwards under the Old System. These are the food banks that regularly bid and win food at negative prices under the Choice System. Welfare is also positively correlated with sampled $\boldsymbol{\Sigma}_i$ parameters, the variance of net donations, but only for non-food, dried, refrigerated, and fresh stocks ($\rho \approx 0.26$). This is sensible — food banks with more uncertain net donations benefit from being able to choose the food they receive from Feeding America. However I do not see these correlations for Tinned/Bottled food.

Additional analysis of the factors associated with food banks benefiting from choice would be valuable. However, these correlations should be interpreted with caution, as they are very dependent on the assumptions that underpin the structural model.

7.3 Additional Mechanisms

The allocation problem faced by Feeding America is faced by numerous other food bank organisations around the world, such as the European Foodbank Federation (FEBA), FareShare (U.K.), and Food bank Australia.[46] To my knowledge no other network employs an allocation mechanism that gives food banks nearly as much choice. Therefore an important question concerns how their mechanisms compare to the Choice System, and to what extent they might be able to benefit by giving food banks more choice over the types of food they receive.

Organisations that do give food banks choice typically allocate food sequentially, allocating them one at a time as they arrive.[47] This has a potential benefit over

[45] This relationship is small but worrying. Feeding America may not be setting budgets optimally - even a utilitarian social planner would equate marginal utility of wealth across food banks. I estimate a lot of variance in these parameters, shown in Figure 22 in Appendix I.2. However the λ_is are not well identified in my model, resting on the strong assumption that lot-specific payoffs v_{ilt} have the same variances across food banks. This is certainly an area for future work.

[46] It is worth recognising that these organisations often face different allocation problems to Feeding America. For example, transport costs are much more pertinent in Australia. Likewise, many of these organisations face a problem closer to the scale of individual food banks allocating food among food pantries. Nonetheless these results remain a useful starting point in analysing the efficacy of their mechanisms and proposed changes.

[47] Many food networks use modern technologies to make sequential allocation more feasible than it was for pre-Choice System Feeding America. Many organisations offer food using apps, and food banks make use of inventory management tools to quickly check the types of food they need.

the simultaneous allocation in the Choice System as food banks do not risk winning too many or too few loads. However, as Akbarpour et al. (2020) and others have highlighted, waiting until all the lots on a given day have arrived, then allocating simultaneously may yield better matches. In this setting, food banks may benefit from having information about everything being allocated that day. They do not risk accepting cornflakes when they really needed ready meals.

7.3.1 Mechanisms

I now briefly summarise the additional counterfactual mechanisms I consider. A detailed discussion of the mechanisms and how I solve for the equilibrium value functions is given in Appendix K.

Old System, all offers: Lots are offered to every food bank in the queue until the lot is accepted or it has been rejected by every food bank.

Closest: The lot is offered to the closest food bank, and no others. This mechanism is used implicitly by food networks who do not allocate food centrally, and instead link food banks up with additional local donors.

Closest, all offers: The lot is offered in order of distance from the lot.

Like: Each lot is offered to every food bank simultaneously. Food banks can "Like" the lot, or not. The lot is randomly (weighted by Goal Factor) assigned to one of the food banks who liked it (Walsh, 2015).

Efficient: The lot is assigned to the food bank with the highest marginal value (plus continuation value) for the lot. This assumes the social planner can observe food banks' needs, and is included for comparison purposes. This mechanism weakly dominates sequential first- and second-price auctions.

Aside from the "Closest mechanism", none of these mechanisms are used explicitly by other food bank networks. However, these give a good overview of some of these types of mechanisms. For example, the mechanism used by the Foodiverse platform in Ireland first offers lots to the nearest food bank, and then offers it to all the other food banks simultaneously as per the "Like" mechanism.

7.3.2 Results

Welfare estimates are presented in Figure 12. There are three key takeaways from this analysis. First, and most importantly, the Closest mechanism performs very poorly.

It achieves only 46.2% of the Welfare under the Choice System, and only around two thirds of the welfare under the Old System. This is indicative of the welfare benefits from centralised allocation. Food banks who are slightly further away from the donor are very often willing to pay the added transport costs.

Second, there are significant benefits from offering food to every food bank, as in the Like and 'all offers' mechanisms, all achieving around a 5% increase on the welfare from the standard Old System. This is very intuitive, since it ensures food is not turned away when there is still a food bank that needs it.

Finally, welfare under the 'Efficient' mechanism does not significantly exceed (nor even weakly dominate) welfare under the 'all offer' and 'Like' mechanisms. This is indicative of the limits of sequential allocation, since that is the main difference between this mechanism and the Choice System. Food is always allocated to the food bank who thinks they need it most in that moment. But it may not be the food most needed by the food bank at that time, and they may forfeit more needed food later in the day because storage is now more costly. This is very similar to the key reasons identified for the value of the Choice System over the Old System, that food banks were accepting sub-optimal food.

8 Conclusion

In this paper I examined the welfare and distributional consequences of Feeding America's implementation of the Choice System, and their decision to allow their food banks greater choice in what food they were allocated. I developed an empirical model of bidding on the Choice System to estimate food banks' demand functions. An important theme of this paper was the role of heterogeneity. I investigated heterogeneity across the types of food Feeding America allocates, as well as heterogeneity in food banks' needs - both across food banks and within food banks over time. This heterogeneity is important to understand how food banks' needs are determined by the types of food they have access to from their local donors.

I found that welfare was 17.1% higher under the Choice System than under the Old System. These results are driven by this heterogeneity, particularly heterogeneity over time. Choice allows food banks to focus their allocations on their most pressing needs, whereas under the Old System they might be offered food that was more useful to a different food bank at that particular time. I estimate that 85% of food banks

Figure 12: Counterfactual Results (2)

Mechanism	Welfare (unweighted)	Welfare (weighted)	Distance (000 km per day)	Allocated (tons per day)	% Better Off (under CS)
Choice	292	745	16	271	1
System	(276, 309)	(672, 815)	(14.6, 17.4)	(253, 284)	(1, 1)
Old	242	576	22.3	249	0.85
System	(203, 276)	(415, 706)	(22, 22.6)	(248, 251)	(0.803, 0.893)
Old System	263	649	19.4	264	0.799
All offers	(222, 297)	(477, 785)	(19, 20)	(263, 266)	(0.738, 0.852)
Closest	134	311	0.305	58	0.945
	(75.7, 183)	(43.9, 530)	(0.298, 0.311)	(57.4, 58.6)	(0.918, 0.967)
Closest	258	632	14	265	0.738
All offers	(218, 294)	(460, 772)	(13.7, 14.3)	(264, 267)	(0.676, 0.795)
Like	264	653	19.4	264	0.819
	(226, 296)	(489, 784)	(19, 20)	(263, 266)	(0.762, 0.877)
Efficient	266	661	19.2	246	0.718
	(228, 297)	(510, 783)	(18.7, 19.8)	(244, 249)	(0.656, 0.779)

Note: This table displays posterior means and 95% credible intervals for various measures of welfare. The final column gives the percentage of food banks who are estimated to be (weakly) better off under the Choice System than each alternative mechanism. A higher number is worse, except for the Choice System which has value of 1 by construction.

are better off under the Choice System. The largest benefits are seen by food banks with the fewest and most variable local donations, benefiting from the flexibility the Choice System permits. This study has important policy implications, both for Feeding America and other food bank networks around the world. I find that welfare under the Choice System significantly exceeds welfare under a number of alternative mechanisms. I found particularly poor welfare consequences of sending food only to the nearest food bank, and that mechanisms which allocate food sequentially as donations arrive are very limited in their efficacy. These finding highlight the importance of good market design.

Future work should consider the external validity of these results, and their applicability in other food bank settings. For example, applying the analysis to data from other food bank networks. Future work should also consider additional mechanisms, potentially building on the Choice System, for application in these other settings and even perhaps improving on the important work already done for Feeding America's allocation problem.

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. (2022). 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.

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

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

Baccara, M., Lee, S., and Yariv, L. (2020). Optimal dynamic matching. *Theoretical Economics*, 15(3):1221–1278.

Backus, M. and Lewis, G. (2016). Dynamic demand estimation in auction markets. Technical report, National Bureau of Economic Research.

Balat, J. (2013). Highway procurement and the stimulus package: Identification and estimation of dynamic auctions with unobserved heterogeneity. *Johns Hopkins University Mimeo*.

Berry, S. T. and Compiani, G. (2020). An instrumental variable approach to dynamic models. Technical report, National Bureau of Economic Research.

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

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

Budish, E. and Cantillon, E. (2012). The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, 102(5):2237–71.

Carter, C. K. and Kohn, R. (1994). On gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Connault, B. (2014). Hidden rust models. *Job Market Paper*.

Engle, R. F. and Granger, C. W. (1987). Co-integration and error correction: representation, estimation, and testing. *Econometrica: journal of the Econometric Society*, pages 251–276.

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.

Gandhi, A. (2019). Picking your patients: Selective admissions in the nursing home industry. *Available at SSRN 3613950*.

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

Gentry, M. L., Komarova, T., and Schiraldi, P. (2020). Preferences and performance in simultaneous first-price auctions: A structural analysis. *Available at SSRN 2514995*.

Groeger, J. R. (2014). A study of participation in dynamic auctions. *International Economic Review*, 55(4):1129–1154.

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

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

Hendricks, K. and Sorensen, A. (2015). The role of intermediaries in dynamic auction markets. *Working Paper*.

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

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

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

Kasahara, H. and Shimotsu, K. (2009). Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica*, 77(1):135–175.

Liu, T. X., Wan, Z., and Yang, C. (2019). The efficiency of a dynamic decentralized two-sided matching market. *Available at SSRN*, 3339394.

NPR (2015). Planet money, accessed: 2018-04-19. https://www.npr.org/templates/transcript/transcript.php?storyId=457408717.

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.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.

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

In this Appendix I present additional details on how I constructed the dataset used in my analysis. Appendix A.1 focuses on the Choice System data, received from Feeding America, outlining how I cleaned and categorised the data. This Appendix also details how I identified joint bidding and which lots were sold by food banks ('Maroon Pounds'). In Appendix A.2 I detail the auxiliary datasets used in my analysis, used to locate food banks and construct their Goal Factors.

A.1 Choice System data

The variables included in the data were as follows:

1. A unique auction identifier
2. Date of the auction
3. Details on all the goods included
4. The number of Pounds in the lot
5. The number of identical lots being auctioned
6. Bids placed on the lot
7. anonymised Foodbank ID that placed each bid
8. The winning Bid(s)
9. An indicator stating whether the Auction/Bid was cancelled
10. The geographic location of the lot

Food banks were anonymised and indexed from 1 to 165. I did not observe whether a bid was placed jointly, nor whether a load was sold by a food bank. I also do not observe whether an auction occurred in the morning or in the afternoon. Finally, I only observe an auction on a particular date if at least one bid is received on that lot. If an auction that consists of two identical loads is observed with just one bid on day t, and another observation with just one bid on day $t+2$ I assume that the auction also appeared on day $t+1$ with only one load available. I must assume that it is not the case that auction appeared on day $t-1$ but no one placed a bid.

A.1.1 data cleaning

The 1344 cancelled auctions and bids were removed from the data, with the assumption that bidding behaviour was not affected by cancellations.

There were various errors in the record data. Some errors could be corrected, such as misspelt names of products, while several had to be removed. Every load listed as being heavier than $97,000$ pounds (the maximum weight for a flat bed truck) was assumed to be a mistake, and fixed to $40,000$ pounds (the modal weight). Every load weighing less than 5000 pounds was also fixed to $40,000$. 6 auctions were removed from the data. These lots included items such as karaoke machines and a flat bed truck. These lots were removed under the assumption that the they fit outside the food banks' ordinary remits.

A.1.2 categorisation

Goods are classified into categories (mostly taken from Prendergast (2022)), subcategories,[48] Uses and Storage Method. Figure 13 panel (A) plots the categories as a proportion of the total amount of food auction. The plot excludes multiple identical auctions, which has the result of artificially reducing the proportions of Fresh Produce and Beverages down from 24% and 17% respectively. Figure 13 panel (B) plots subcategories as a word cloud, with more common subcategories larger.

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

Storage methods includes Shelf, Tinned, Refrigerated, Fresh, and Non-Food. The Non-food category is identical to the non-Food Use category. Shelf items can be stored on a shelf, are generally dried goods, and have extremely long shelf lives. They are generally light but bulky. Tinned food, which includes jars and bottles, have long shelf-lives and are generally compact and heavy. Refrigerated food must be stored in a fridge, but still expire reasonably quickly. Fresh food is food one wouldn't generally store in a fridge, and generally has only a limited shelf life. This includes both fresh produce and freshly baked goods such as bread. Any item that was additionally listed as 'Shelf Stable', such as UHT milk was put in the tinned storage category.

Figure 13: Composition of food allocated: Categories and Subcategories

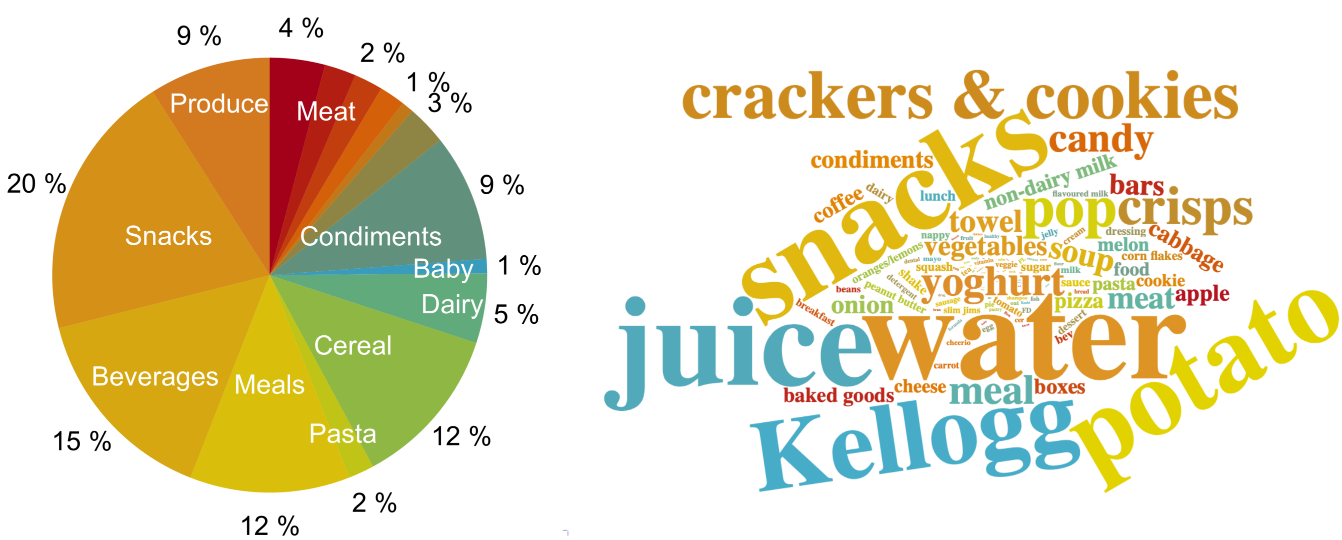

Code	Category	Subcategory	Use	Storage	Description

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

1	Baby		Non-food	N F	unspecified baby
2	Baby	diaper	Non-food	N F	nappies
3	Baby	food	Non-food	N F	Baby food
4	Baby	formula	Non-food	N F	Baby formula
5	Beverage		Snack	Tin	unspecified Bev
6	Beverage	capri	Snack	Tin	capri-sun
7	Beverage	coffee	Snack	Shelf	ground/instant
8	Beverage	dry	Snack	Shelf	chocolate/milk powder
9	Beverage	fj	Snack	Tin	Orange/Apple/Grape juice - high quality, "pure"
10	Beverage	gator	Snack	Tin	Sports drink (gatorade)
11	Beverage	ic	Snack	Tin	Iced/Alternate coffee
12	Beverage	juice	Snack	Tin	juices, lower quality, mixed, e.g. tropical punch, fruit shoot
13	Beverage	ka	Snack	Shelf	Kool-Aid
14	Beverage	pop	Snack	Tin	fizzy drinks, e.g. coke
15	Beverage	propel	Snack	Tin	Propel brand water/sports water
16	Beverage	pshake	Snack	Shelf	Protein shake/powder
17	Beverage	shake	Snack	Tin	Milk shakes
18	Beverage	tea	Snack	Shelf	Tea/Tea bags
19	Beverage	vf	Snack	Tin	V8 juices
20	Beverage	water	Snack	Tin	Bottled water
21	Baked Good		Snack	Shelf	unspecified BP
22	Baked Good	bread	Ingredient	Fresh	Bread
23	Baked Good	cake	Meal	Fresh	cake, cupcakes, muffins
24	Baked Good	dough	Ingredient	Fridge	cookie dough, bread dough, etc
25	Baked Good	flour	Ingredient	Shelf	flour, cake mix, bread mix
26	Baked Good	other	Snack	Shelf	miscellaneous BP
27	Baked Good	pastry	Snack	Fresh	croissants, waffles, pancakes etc
28	Baked Good	stuffing	Condiment	Shelf	Stuffing mix
29	Cereal		Meal	Shelf	unspecified cereal
30	Cereal	bran	Meal	Shelf	healthy bran cereal (fibre)
31	Cereal	cheerio	Meal	Shelf	Cheerios
32	Cereal	flake	Meal	Shelf	un-sweatened flakes (spK, corn etc)
33	Cereal	gran	Meal	Shelf	granola
34	Cereal	Kashi	Meal	Shelf	unspecified Kashi
35	Cereal	Kellogg	Meal	Shelf	unspecified Kellogg
36	Cereal	ns	Meal	Shelf	non-sugared cereal (rice-krispies)
37	Cereal	other	Meal	Fridge	miscellaneous non-dry cereal
38	Cereal	oat	Meal	Shelf	oats/grits/porridge
39	Cereal	PL	Meal	Shelf	unspecified Private Label (e.g. Post)
40	Cereal	sugar	Meal	Shelf	fruit loops, apple jacks etc
41	Condiment		Condiment	Tin	unspecified condiments
42	Condiment	dressing	Condiment	Tin	Salad dressings, glazes
43	Condiment	fruit	Condiment	Tin	Fruit sauces, preserves
44	Condiment	gravy	Condiment	Shelf	Gravy granules
45	Condiment	jelly	Condiment	Tin	Jam
46	Condiment	ketchup	Condiment	Tin	Ketchup
47	Condiment	mayo	Condiment	Tin	Mayonnaise
48	Condiment	mustard	Condiment	Tin	Mustard
49	Condiment	other	Condiment	Tin	Miscellaneous cond (e.g. frosting)
50	Condiment	oil	Ingredient	Tin	Cooking oils
51	Condiment	pasta	Condiment	Tin	Pasta sauces
52	Condiment	PB	Condiment	Tin	Peanut Butter
53	Condiment	pickle	Condiment	Tin	Pickled Gherkins
54	Condiment	salsa	Condiment	Tin	Salsa/Guacamole/dips
55	Condiment	sauce	Condiment	Tin	BBQ sauce, etc
56	Condiment	stock	Ingredient	Shelf	Stock (assumed cube form)
57	Dairy		Ingredient	Fridge	unspecified Dairy
58	Dairy	butter	Condiment	Fridge	Butter/Margarine/Spread
59	Dairy	cc	Condiment	Fridge	coffee-creamer, coffee-mate (assumed liquid)
60	Dairy	cheese	Ingredient	Fridge	mostly cottage/cream cheese
61	Dairy	cream	Condiment	Fridge	Mostly sour cream
62	Dairy	dessert	Meal	Fridge	cheese cake etc
63	Dairy	egg	Ingredient	Fresh	eggs
64	Dairy	ll	Ingredient	Tin	evaporated/preserved milk
65	Dairy	milk	Condiment	Fridge	milk
66	Dairy	milk-alt	Condiment	Fridge	non-Dairy milk
67	Dairy	milk-flav	Snack	Fridge	flavoured (chocolate) milk

68	Dairy	pie	Meal	Fridge	Sweet pies, e.g. Apple/custard
69	Dairy	yog	Snack	Fridge	Yoghurt
70	Fresh		Ingredient	Fresh	unspecified produce
71	Fresh	apple	Snack	Fresh	Apples
72	Fresh	cabbage	Ingredient	Fresh	Cabbages
73	Fresh	carrot	Ingredient	Fresh	Carrots
74	Fresh	citrus	Snack	Fresh	Citrus fruits
75	Fresh	corn	Ingredient	Fresh	Corn (maize)
76	Fresh	fruit	Snack	Fresh	unspec/misc fruit
77	Fresh	melon	Snack	Fresh	melons
78	Fresh	other	Ingredient	Fresh	Miscellaneous veg
79	Fresh	onion	Ingredient	Fresh	Onions or garlic
80	Fresh	potato	Ingredient	Fresh	Potatoes
81	Fresh	squash	Ingredient	Fresh	Squash/Pumpkin,Yams
82	Frozen		Ingredient	Fridge	unspecified/misc frozen
83	Frozen	bp	Snack	Fridge	Frozen Baked Goods, e.g. bread rolls
84	Frozen	dairy	Ingredient	Fridge	Frozen milk, butter, eggs
85	Frozen	meal	Meal	Fridge	Frozen meals/pies/pizza
86	Frozen	meat	Ingredient	Fridge	Frozen chickens etc
87	Frozen	veg	Ingredient	Fridge	peas, carrots etc
88	Health/Beauty		Non-Food	Non-Food	unspecified HBC
89	Health/Beauty	body	Non-Food	Non-Food	body creams/moisturiser
90	Health/Beauty	dental	Non-Food	Non-Food	dental hygiene
91	Health/Beauty	deod	Non-Food	Non-Food	deodorant
92	Health/Beauty	detergent	Non-Food	Non-Food	detergent powder/tablets
93	Health/Beauty	drug	Non-Food	Non-Food	medicines/ointments
94	Health/Beauty	nutri	Non-Food	Non-Food	vitamins / unspecified nutritional items (e.g. protein powder)
95	Health/Beauty	other	Non-Food	Non-Food	miscellaneous (e.g. razors)
96	Health/Beauty	shampoo	Non-Food	Non-Food	shampoo/conditioner
97	Health/Beauty	soap	Non-Food	Non-Food	hand/body soap
98	Health/Beauty	sun	Non-Food	Non-Food	sun-cream/block
99	Meal		Meal	Fridge	unspecified Meals
100	Meal	bert	Meal	Fridge	Bertolli ready meals
101	Meal	breakfast	Meal	Fridge	breakfast meals
102	Meal	broth	Ingredient	Tin	Broth - assumed carton stock
103	Meal	cb	Meal	Tin	Chef Boyardee ready meals
104	Meal	chang	Meal	Fridge	P.F. Chang ready meals
105	Meal	chilli	Meal	Tin	Tinned Chilli / meat 'n' beans
106	Meal	healthy	Meal	Fridge	Healthy/Nutritious ready meals (e.g. weight-watchers, fish)
107	Meal	lunch	Meal	Fridge	Lunchables (ready packed lunches)
108	Meal	mc	Meal	Shelf	Marie Callender ready meals
109	Meal	meat	Meal	Fridge	Meat based ready meals
110	Meal	other	Meal	Fridge	miscellaneous ready meals
111	Meal	pasta	Meal	Shelf	Pasta ready meals, mac n' cheese etc
112	Meal	pie	Meal	Fridge/Shelf	Savoury pies / pastries (often shelf stable)
113	Meal	pizza	Meal	Fridge	pizzas
114	Meal	sand	Meal	Fridge	sandwiches
115	Meal	side	Snack	Fridge	ready meal sides
116	Meal	soup	Meal	Tin	tinned soups
117	Meal	veggie	Meal	Fridge	vegetarian/vegan meals
118	Meat		Ingredient	Fridge	unspecified meat
119	Meat	bacon	Ingredient	Fridge	Bacon
120	Meat	beef	Ingredient	Tin	Mostly tinned savoury mince
121	Meat	burger	Ingredient	Fridge	various burger patties
122	Meat	chicken	Ingredient	Fridge	Chicken
123	Meat	fish	Ingredient	Fridge	Fish
124	Meat	lunch	Ingredient	Fridge	Deli/luncheon meat
125	Meat	other	Ingredient	Fridge	miscellaneous meats (e.g. pork)
126	Meat	sausage	Ingredient	Fridge	Mostly hot dog sausages
127	Non Food		Non-Food	Non-Food	unspecified non-food
128	Non Food	battery	Non-food	Non-food	batteries
129	Non Food	bleach	Non-food	Non-food	Bleach/solvent cleaning products
130	Non Food	box	Non-food	Non-food	banana boxes/crates
131	Non Food	other	Non-food	Non-food	e.g. clothes, bags, window cleaner, wipes
132	Non Food	salt	Non-food	Non-food	non-food salt
133	Non Food	towel	Non-food	Non-food	paper towels
134	Pasta		Meal	Shelf	Dried pasta
135	Pasta	ben	Ingredient	Shelf	Uncle Ben's rice
136	Pasta	other	Ingredient	Shelf	Miscellaneous pasta product (lasagna sheets etc)

137	Pasta	rice	Ingredient	Shelf	dried rice
138	Snack		Snack	Shelf	unspecified snack
139	Snack	bar	Snack	Shelf	snack/granola bars
140	Snack	bp	Snack	Shelf	baked snacks, e.g. butterfinger
141	Snack	candy	Snack	Shelf	candy/chocolate
142	Snack	chips	Snack	Shelf	crisps
143	Snack	cookies	Snack	Shelf	biscuits
144	Snack	crackers	Snack	Shelf	crackers
145	Snack	fruit	Snack	Shelf	rollups/cups
146	Snack	jelly	Snack	Shelf	jello (pre/unmixed)
147	Snack	kellogg	Snack	Shelf	unspecified Kellogg brand snacks
148	Snack	nuts	Snack	Shelf	nuts/trailmix
149	Snack	other	Snack	Shelf	miscellaneous snacks
150	Snack	pbar	Snack	Shelf	protein bars
151	Snack	pc	Snack	Shelf	pop-corn (mostly popped)
152	Snack	pretzel	Snack	Shelf	pretzels
153	Snack	pt	Snack	Shelf	pop-tarts
154	Snack	pud	Snack	Tin	tinned pudding
155	Snack	seed	Snack	Shelf	sunflower seeds
156	Snack	sj	Snack	Shelf	slim Jims, jerky, biltong
157	Vegetables		Ingredient	Tinned	unspecified non-fresh
158	Vegetables	beans	Ingredient	Tinned	baked beans
159	Vegetables	fruit	Ingredient	Tinned	canned fruit (escaloped apples etc)
160	Vegetables	fry	Ingredient	Fridge	chips/potato wedges/ fries
161	Vegetables	gbean	Ingredient	Tinned	beans (non-baked, mostly green)
162	Vegetables	other	Ingredient	Tinned	miscellaneous veg
163	Vegetables	potato	Ingredient	Fridge	ready to cook potatoes
164	Vegetables	tomato	Ingredient	Tinned	tinned tomatoes

A.1.3 Joint bidding

I did not receive information of joint bidding. However, in some circumstances joint bidding can be inferred. For example, when only one load is auctioned but multiple foodbanks are listed as winning. Likewise, I observe the amount paid by winners (separate to their bid): If two food banks jointly bid 50 shares each I observe that the "bid paid" was 100. I use these cases to identify common bidding coalitions. I then assume that whenever one of these coalitions appears to place a bid, that they are placing a joint bid. By this method I identify around 30 coalitions, and infer that 4.5% of bids are joint bids. This is slightly lower than the true value of 5% reported in Prendergast (2017). This is likely because I do not detect coalitions that never won together in the data.

I also risk classifying non-joint bids as joint bids when a coalition chooses not to bid jointly on occasions. This is only a problem if they did not win, and unlikely to lead to much inaccuracy if they do not win. A further problem is that I occasionally see multiple lots being auctioned, with more winners than lots (without a known coalition among these food banks). I am unable to infer which subset of bidders forms a coalition, and so am are forced to assume, incorrectly, that none of the bids are joint. This only happens a small fraction of the time, around 0.01%, so is unlikely to lead to much inaccuracy.

Joint bidding in the model

I do not consider the strategic considerations behind joint bidding. I do however consider how this impacts the inverse bid function and winnings. If a bid was joint between n people, I assume the pounds won are divided equally among the n bidders. As are the distance costs and the 'lot specific value'.[49] I also recognise how the food banks' beliefs about the probability they win given their joint bid is higher than either individual bid. Therefore I am able to recognise how the total (expected) surplus of the joint bid may exceed either individual surplus from placing a single bid equal to the joint bid - if storage costs are convex, sharing these load reduces the total cost incurred. When simulating the Choice System I am unable to simulate the joint bidding procedure.

A.1.4 Maroon Pounds

I do not observe which loads were sold by food banks ('Maroon Pounds'). However, after I had located food banks (discussed in Appendix A.2) considered whether any of the auction origin zipcodes matched the zipcodes of the food banks I had identified. Matched observations all had auction identifier codes that began with "ML" rather than "L" (followed by a string of numbers). Therefore, I focused on these auctions as Maroon Pounds, which make up 4.5% of unique auctions.

Maroon Pounds do not enter the current version of my model. Endogenising the decision to sell food adds too much complexity. However, the food banks responsible for consuming the most through the Choice System almost never sell food. As my results are predominantly driven by these food banks, ignoring Maroon Pounds is unlikely to lead to much inaccuracy. However, for the sake of posterity I will continue to describe how I match food banks to Maroon Loads.

One difficulty with matching food banks to maroon loads is that food banks move over time, often merging with other food bank organisations, so that the zipcode of a lot auctioned in 2014 may not match the zipcode I found for that food bank in 2019. Broadly speaking, I located food banks by finding the name of the city in which they are located, as well as their state. It is rare to have multiple food banks in the same. I therefore decided to match food banks to maroon loads under 3 conditions: First, if the zipcodes matched. Failing that, if they are located in the same city. Failing that, if they are within 20 miles of one another. I assume that the remaining Maroon Pounds are sold by the small food banks and food rescue organisations whose locations I cannot identify, or who are never observed bidding in my data.

A.2 Auxiliary data

I use five additional datasets in my analysis. Two data sets received from Prendergast (one of the original designers of the Choice System) containing losing bidders by auction for 2014, and also poverty figures by county. Third, food bank zipcode data from Feeding America's Food Bank Locator online tool.[50]. Fourth, Food bank catchment areas, defined at the County level, from Feeding America's 'Hunger in America' on-line resource.[51] Finally, Populations figures by county were then taken from the 2015 US census Small Area Income and Poverty Estimates (SAIPE).

These datasets were used to locate food banks and evaluate their Goal Factors.

[49]In principle I could split the lot according to the fraction of final expenditure, however joint bids in which one bidder bids an extremely small or zero amount, are not uncommon.

[50]Accessible at https://www.feedingamerica.org/find-your-local-foodbank

[51]Available here: https://map.feedingamerica.org

A.2.1 Locating food banks

Prendergast kindly sent me a dataset containing data on losing bidders by auction for 2014. Importantly, this data contained the nearby towns of the bidding food banks. That is, food banks were identified by the town they were located in. I was able to cross-reference this data with the Choice System data for 2014, merging by date and the origin of each lot.

For each anonymised ID I found the town that appeared in the largest proportion matched auctions. For each town I found the anonymised ID that appeared in the largest proportion of matched auctions (these two proportions need not be equal). If the two sets of pairings were identical, I listed the ID/town combination as matched, removed it from the pool of remaining IDs and towns, and continued the process until I was unable to remove any more matched pairs. This process allowed me to infer the nearby towns of all food banks who placed a bid in 2014. This allowed me to infer approximate locations for 85% of food banks, who together consumed just over 98% of all food on the Choice System. It was clear that my food bank ID numbers had been listed in alphabetical order from 1 to 165 before anonymisation, validating my location matches.[52]

Given knowledge of nearby towns, I then used Feeding America's Food Bank Locator online tool to find zip codes for these food banks. I was unable to find three food bank's locations in this way, as they listed town names which were nowhere near any of Feeding America's food banks. I kept their locations as unknown.

One of the most frequent bidders in the Choice System has a commonly occurring town name, with food banks listed in two of these towns. For these two candidate food banks I examined their annual financial statements from 2014 to find how much non-monetary donations they had received from Feeding America. One received an extremely large amount, while the other received a reasonably small amount. Because the food bank in question consumed an extremely large amount of food on the Choice System I reasoned it was most likely the food bank that received the larger non-monetary donations from Feeding America.

A.2.2 Distance

To find the distance between every lot $\times$ food bank combination I converted zipcodes into longitude/latitudes, then used the "distGeo" function from the R package "geosphere". This package finds the distance of the geodesic between any two points on the globe. In principle I could have found the shortest road distance using arcGIS software, as this would more accurately represent the transportation costs. However, this software is generally extremely computationally intensive. Given the large number of food bank $\times$ lot combinations (≈ 3.6 million) this option was not feasible.

A.2.3 Calculating Goal Factors

I did not receive recent Goal Factor figures. However, this data can be constructed using the locations of food banks, Goal Factor formulae given in **?**, and information on local poverty and

[52] Based on this alphetical order, and knowledge of all Feeding America's associated food banks, I was able to match an additional 3 food banks by visual inspection

food insecurity rates from Feeding America's 'Hunger in America' on-line resource. Under the Old System a food bank's Goal Factor was given by:

$$GF_i^{OS} = \frac{\text{Population}_i}{\text{Population}_{US}} + \frac{\text{Poverty}_i}{\text{Poverty}_{US}} \tag{5}$$

Where Population_i refers to the number of people living in food bank i's catchment area, and Poverty_i refers to the number of people living below the poverty line in food bank i's catchment area. Food bank catchment areas, defined at the County level, are given in Feeding America's 'Hunger in America' on-line resource. Populations figures were then taken from the 2015 US census Small Area Income and Poverty Estimates (SAIPE). Poverty rates, by county, are given in an additional dataset received from Prendergast, in turn received from Feeding America. Presumably these were the figures used to construct the Goal Factors to begin with.

Under the Choice System, the Goal Factor formula was updated to reflect that even individuals above the poverty line often use food banks. The new formula includes $\text{Poverty}'_i$, the number of people between the poverty line and 185% of the poverty line, as well as $\text{Population}'_i$, the number of people above 185% of the poverty line. These figures were included in the dataset I received. These figures are weighted according to empirical usage weights. The updated formula is given by:

$$GF_i^{CS} = \frac{0.73\,\text{Poverty}_i + 0.22\,\text{Poverty}'_i + 0.05\,\text{Population}'_i}{0.73\,\text{Poverty}_{US} + 0.22\,\text{Poverty}'_{US} + 0.05\,\text{Population}'_{US}} \tag{6}$$

I set the Goal Factors of food banks with unknown locations to the smallest known Goal Factor.[53] In this paper I only use the new Goal Factors.

For a small number of food banks their expenditure did not match up with their Goal Factors. That is, they spent *significantly* more shares than the amount they received (as implied by their Goal Factor). This is the case even when I take into account that food banks stop receiving new shares once they hit a 200,000 limit.

I calibrate Goal Factors and initial budgets to take this into account. I find the smallest absolute deviation from the Goal Factors implied by the formulae such that: 1. No food bank is ever in debt for longer than 30 days. 2. Food banks with above average Goal Factor are never in debt (as these food banks do not get access to credit). 3. Food banks' budgets cannot exceed the 200,000 share limit. 4. No food bank's budgets have a trend (positive or negative) of more than 100,000 shares over the period. 5. No food bank's budgets have a statistically significant (at the 5% level) trend. That is, we expect that their budgets should neither trend up nor downwards over time.

I perform this calibration as follows: Given proposed Goal Factors I find the initial share allocation that satisfies criteria 1-4. This is done by iteratively changing the initial allocation, simulating incomes (given observed expenditures) to find budgets, until the necessary initial allocation is converges. Then, in an outer loop, I find the Goal Factors that satisfy criteria 1,2, and 5. At each step

[53] Given that these food banks did not bid regularly enough for me to identify their location, presuming that these food banks also would not accept any loads they are offered in counter-factual simulations ensures that my estimates remain conservative. For some catchment areas Feeding America appears to have used slightly different cut-offs than the 100% and 185% lines.

I update Goal Factors by taking the average of the prior estimated Goal Factor and the implied new Goal Factor that satisfies the criteria. The process converges in around 100 iterations. For 95% of food banks Goal Factors change very little in relative terms. The largest change is seen by one food bank that consumes an extremely large quantity of food on the Choice System, but has an extremely low initial Goal Factor. However, inspection reveals that this food bank exists in a so-called 'food desert', meaning they likely have very little access to local donors, so must rely on Feeding America for the majority of their food.

To validate this approach I compare the distributions of calibrated initial allocations and Goal Factors to those used in Prendergast (2017), received from Prendergast, but that I was unable to link to my data. Importantly, these figures are around 5 years out of date relative to my data. The two sets of distributions are shown in 14. The distribution of intial budgets are relatively similar, as are the distributions of Goal Factors, with the exception that my estimated Goal Factors have a larger right tail. Importantly, however, the distribution of my estimated Goal Factors fits the observed distribution better than my initially calculated Goal Factors.

The Goal Factor was designed to ensure that a food bank with a 1% higher goal factor received 1% more food. Prendergast (2017) found that a 1% increase in Goal Factor was associated with a 0.45% increase in food won from the Choice System. I found that a 1% increase in estimated Goal Factor was associated with a 0.81% increase in consumption. Given that Prendergast's estimation was done on data with very different characteristics to mine, the inaccuracy of these estimated figures is unclear. The difference may be driven by the 15% of unknown Goal Factors in my data. If these food banks had relatively high goal factors this would drive the observed discrepancy, since we know these food banks choose not to consume much. Either way, the relationship between Goal Factor and consumption is not especially strong; the R^2 from a log-log regression is only 0.35. This weak correlation demonstrates the importance of food wealth in determining consumption behaviour. High Goal Factor food banks, who also happen to have many local donors, may not want to consume much food through the Choice System, weakening the correlation.

B Additional Descriptive Analysis

In this appendix I discuss additional descriptive analysis of the Choice System, building on the results in Section 3.1. I focus on establishing evidence of variation in bidding behaviour over time.

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

Figure 14: Distribution of Goal Factors and intial budgets

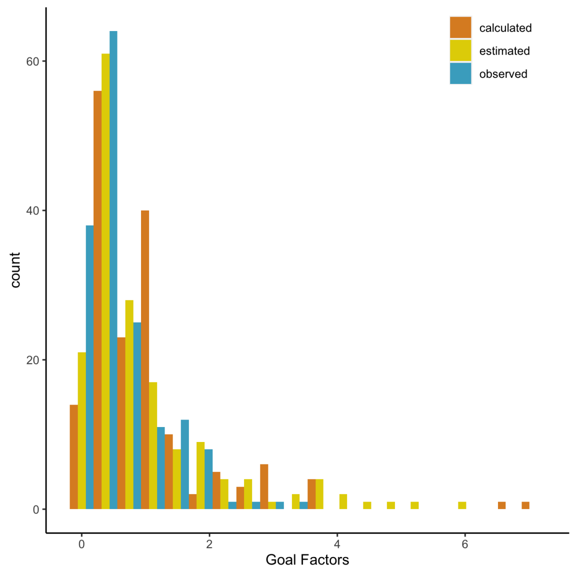

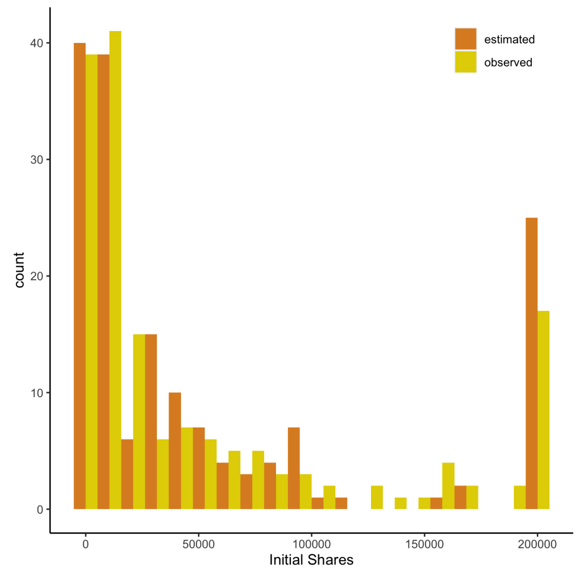

Note: These plots show histograms of Goal Factors and initial budgets across food banks. 'Observed' are the figures I received from Prendergast. 'Calculated' are the figures I calculated using poverty data and the formulae presented above. 'Estimated' are the figures I calibrated using the approach discussed above, that were then used in my counterfactual simulations.

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

This hypothesis test may be underpowered to reject a null hypothesis of constant average bids. It does not take into account that average bids likely don't shift neatly at the beginning of each month. It also does not take into account that large variation in bids within a month (which may cause failure to reject the null) are also indicative of variation in food banks' needs. If the within month variation is on a similar scale to the across month variation in average bids this reduces my power to reject the null hypothesis.

However, the test may be over powered if variation in factors other than food banks' needs is mistaken for variation in needs. For example, if the quality of food varies unobservably over time, this may cause systematic variation in bidding behaviour that should not be attributed to variation in food banks' needs. To account for this possibility I estimate a second restricted specification with food bank specific month fixed effects. These fixed effects will capture variation in bidding behaviour that is common across food types. Under this specification a rejection of the null is evidence of systematic variation over time in bidding behaviour *on specific types of food*. This specification is almost certainly underpowered. If food banks need more food of all types in certain months the fixed effects will also soak up this variation.

Figure 15 plots the likelihood ratio test statistic across food banks. The dotted lines gives the χ^2 critical values for tests at the 5% significance level. The red points give the baseline specification, while the blue points give the specification including month fixed effects.

Figure 15: Heterogeneity Across Time

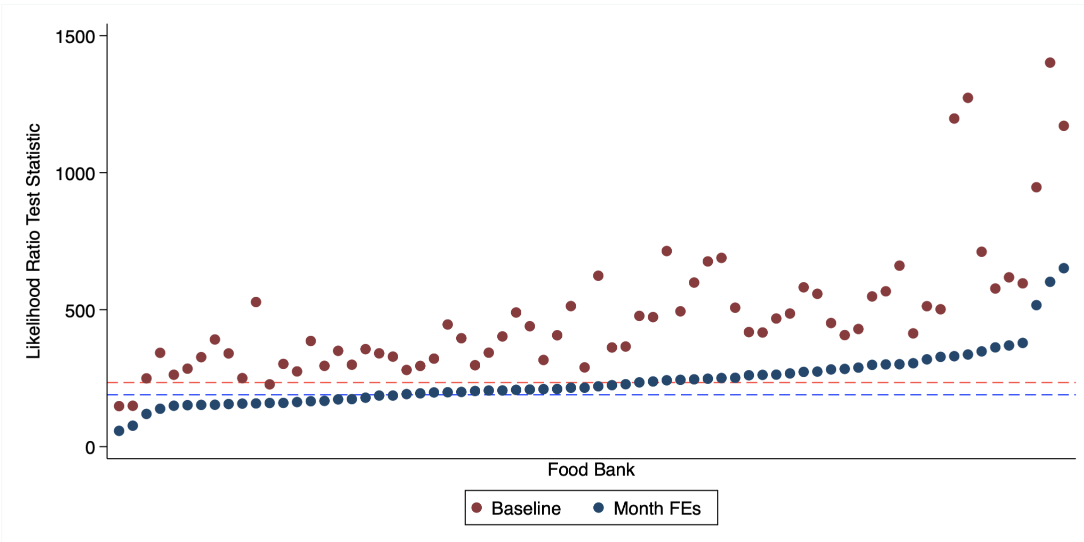

Note: This figure plots likelihood ratio test statistics for the hypothesis test that average bids for each type of food are constant over time, against the alternate hypothesis that bids vary by month. The estimated model controls for censoring, distance, and lot composition. The blue results also include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

I can reject the null hypothesis of constant average bids over time, at 5% significance level, for 96% of food banks in my baseline specification, and 70% of food banks for the month fixed effects specification. Therefore I have strong evidence that food banks' bidding behaviour, and hence their needs, vary significantly over time.

C Stationarity

In this appendix I present evidence that the equilibrium stock process is stationary - that the distribution of stocks remains constant over time. I focus on two types of stationarity: First, whether stocks trend over time. Second whether stocks follow a random walk.

In Appendix C.1 I present suggestive evidence that stocks neither trend upwards nor downwards over time, by testing for structural breaks in bidding behaviour. In Appendix C.2 I discuss how an additional assumption about how observed winnings reacts to changes in the stock allows me to test whether the equilibrium stock process follows a random walk. In Appendix C.4 I discuss how the results of this analysis gives us information about food banks' unobserved stock process, giving us natural priors for $\boldsymbol{\mu}_i$ and Σ_i.

C.1 Trend Stationarity

If stocks trend over time we expect that bidding behaviour should follow a similar pattern. Therefore we can investigate the existence of trends by looking for evidence of trends in bidding behaviour.

If stocks have a linear trend it is ex-ante unclear whether average bids will also have a linear trend. To allow for the possibility of non-linear trends in bidding behaviour I focus on testing for the existence of more general structural breaks in behaviour. I focus on average monthly bids by food bank and food type, using the estimated α_{igt} parameters estimated from the tobit specification in Section 3.1. If these parameters do exhibit a linear trend a generalised test for structural breaks should pick this up. I omit the Fresh storage type due to the structural break caused on day 553 of my sample when fresh food ceased to be allocated on the Choice System. Even if the Fresh stock process does not exhibit any structural break (as remains my hypothesis) I cannot estimate average bids after this break, given that no fresh food was allocated. In Appendix H.2 I give additional details of how I model this structural break.

I test for a structural break in the series $\{\alpha_{igt}\}_{t\in\{1\dots T\}}$. For each t I split the sample into a before and after group, then run a t-test on the equality of means. This is performed separately for each food bank $\times$ storage type combination, and I allow the variances to differ in the before and after periods. I then plot the distribution of estimated t-statistics. Under the null hypothesis of no structural breaks, these statistics are t-distributed with 40 degrees of freedom.[54] Therefore I can compare the resulting distribution of t-statistics from their distribution under the null hypothesis.[55] In figure 16 I plot the distribution of estimated t-statistics along side the distribution of these statistics under the null-hypothesis. Under the null we expect that 5% of test statistics will be above the critical values for a two-tailed test at 5% significance level (= 2.02). I find that 5.4% of test-statistics exceed the critical values. This gives some evidence that bidding behaviour does not exhibit trends or structural breaks, and so neither do stocks.

C.2 Cointegration

I now discuss how I can test whether the equilibrium stock process follows a random walk. I must maintain the assumption that stocks do not contain any sort of time trend, such as a linear trend. Fortunately the results presented in Appendix C.1 gives us evidence that equilibrium stocks are

[54] Although I have 44 months of data, I do not have all the data on the first and last month. So I only estimate 42 monthly means. I also adjust the t-statistics to account for sampling variation in the estimated α_{igt} using the law of total variance. Because I only have a finite number of months I assume normality of average monthly bids, resulting in this t-distribution.

[55] Using a Kolmogorv-Smirnov test I can reject the null hypothesis, at a 5% significance level, that the t-statistics come from this distribution. However, this is due to a lack of fit around the mode of the distribution, whereas the tails of the distribution fit much better.

Figure 16: Distribution of t-Test statistics

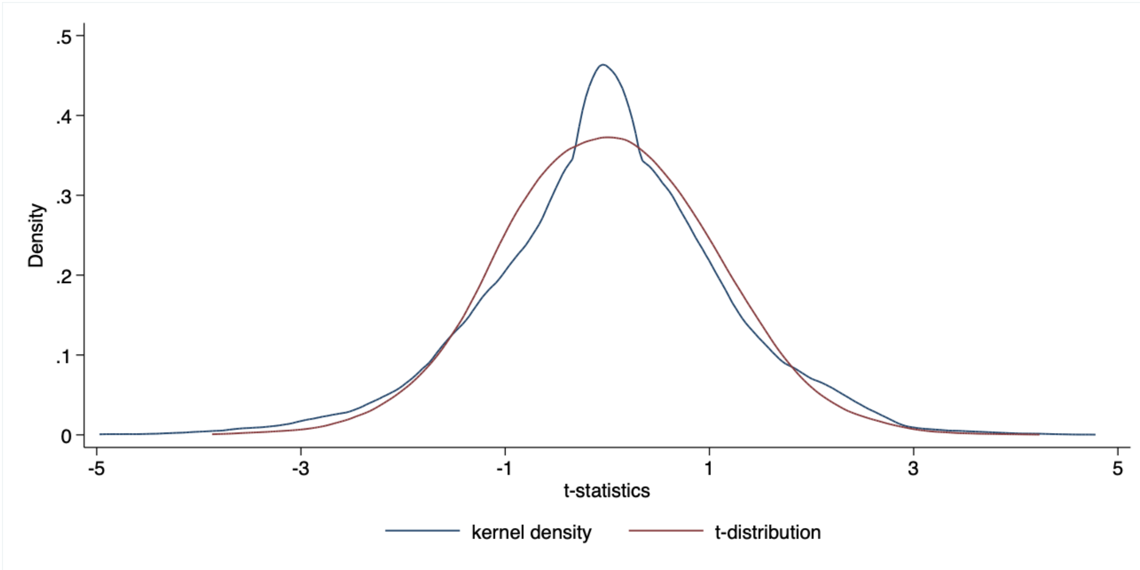

Note: This plot shows the distribution of t-test statistics, looking for structural breaks in average monthly bids. I focus on monthly bids between the 2nd and 43rd month. I also plot the standardised t-distribution with 40 degrees of freedom. I test for the presence breaks at each month between the 7th and 37th month. Tests are performed at the food bank × storage type level. Test-statistics are adjusted to account for sampling uncertainty in the estimates of average monthly bids.

unlikely to exhibit time trends.[56]

In the main text I focus on the morning state, before any auctions take place. However for the purposes of this appendix it is most convenient to focus on the evening state, $\mathbf{s}^e$, after the final auction has taken place. The evening stock transition process is therefore given by:

$$\mathbf{s}^e_{it} = \mathbf{s}^e_{it-1} + \mathbf{w}^T_{it}\mathbf{z}^g_t + \mathbf{x}_{it}$$

The only difference is that the superscript on winnings is not lagged, as it is for the morning process. Evidently $\mathbf{w}^T_{it}\mathbf{z}^g_t$ depends on net daily donations $\mathbf{x}_{it}$, which can be considered short term changes in the stock. It will also depend on the previous stock $\mathbf{s}^e_{it-1}$. This is where the similarity to a cointegration framework arises. Food banks likely have an ideal level of stock they would maintain - they do not want the warehouse to be too empty, nor too full. Therefore, as well as reacting to short term changes in net daily donations, winnings should also respond to how far off optimal the previous stock is.

[56]A linear trend is not identified from winnings data. Winnings are a measure of changes in the stock, so any linear trend is captured in the constant term. But, the constant term of average winnings also captures average net local donations $\boldsymbol{\mu}_i$. I am unable to disentangle these two objects using winnings data alone.

If I observed stocks as well as winnings I could easily test this relationship following the procedure of Engle and Granger (1987). Instead, I make the additional assumption that the equilibrium stock process is given by:

$$\mathbf{s}_{it}^{e} = \delta \mathbf{s}_{it-1}^{e} + \alpha \mathbf{x}_{it} + \boldsymbol{\varepsilon}_{it}$$

This assumption states, on average, the evening stock ends up as some fraction of the previous evening stock, plus some fraction of net local donations. These are the fractions that equilibrium winnings could not offset. I do not impose that this process is stationary. For example, it is possible that $\delta = I$, so that the process follows a random walk with drift. I discuss $\boldsymbol{\varepsilon}_{it}$ in detail shortly. I must assume there is not a linear trend in this process. For simplicity I focus on the case in which δ and α are diagonal matrices, essentially focusing on one component of stocks at a time.

This assumption on the equilibrium stock process is really an assumption on the equilibrium winnings process. Equating the two previous equations and rearranging yields something similar to a standard error correction process:

$$\mathbf{w}_{it}^{T} \mathbf{z}_{t}^{g} = (\alpha - I)\mathbf{x}_{it} + (\delta - I)\mathbf{s}_{it-1}^{e} + \boldsymbol{\varepsilon}_{it} \tag{7}$$

This states that, on average, winnings offset some fraction of that day's net donations, as well as some fraction of the previous stocks. The residual random variable $\boldsymbol{\varepsilon}$ captures idiosyncrasies that affect the food bank's winnings. For example, how much is actually actioned that period, how many rival active bidders there are, and other attributes of the lots. It likely exhibits correlation over time, is non-normally distributed, lumpy, and may have non-zero mean. Importantly, this variable is assumed independent of $\mathbf{x}_{it}$.

If $\delta = 0$ and $\alpha = 0$ then winnings perfectly offsets changes in the stock and daily donations. This means that equilibrium stocks only vary with $\boldsymbol{\varepsilon}_{it}$. Instead, if $\delta = I$, so that the stock process followed a random walk, this equation would state that winnings do not depend on the previous stock. This asserts that winnings only react to $\mathbf{x}_{it}$ and $\boldsymbol{\varepsilon}_{it}$. In this way, a random walk stock process suggests that the food bank only reacts to daily stock changes, treating past stock changes as a lost cause, not trying to offset past losses. This means allowing previous losses to propagate completely over time, creating the random walk. This gives us an intuitive way to test for stationarity - test whether winnings depend on previous stocks.

Substituting the equilibrium stock process into Equation 7 yields:

$$\mathbf{w}_{it}^{T} \mathbf{z}_{t}^{g} = (\alpha - I)\mathbf{x}_{it} + (\delta - I)\sum_{s=1}^{\infty} \delta^{s-1}[\alpha \mathbf{x}_{it-s} + \boldsymbol{\varepsilon}_{it-s}] + \boldsymbol{\varepsilon}_{it}$$

Now, suppose $\boldsymbol{\varepsilon}_{it}$ can be decomposed according to: $\boldsymbol{\varepsilon}_{it} = \gamma \mathbf{r}_{it} + \boldsymbol{\nu}_{it}$, where $\mathbf{r}_{it}$ is a vector of observables that impact the food bank i's winnings (all of which must be independent of $\mathbf{x}_{i}$). These can be considered non-stock factors that impact winnings. For example, I focus on the total supply in period t, by storage type, the number of other food banks who placed a bid on a lot of each storage type, as well as the minimum distance between food bank i and a lot of each storage type in period t. Importantly I require that $E[\mathbf{r}_{is} \otimes \boldsymbol{\nu}_{it}] = \mathbf{0}$. As above, I expect $\boldsymbol{\nu}_{it}$ to be lumpy, non-normal, and with possibly non-zero mean. We can then re-write the equilibrium winnings process as:

$$\mathbf{w}_{it}^T \mathbf{z}_t^g = \gamma \mathbf{r}_{it} + (\delta - I) \sum_{s=1}^{\infty} \delta^{s-1} \gamma \mathbf{r}_{it-s} + (\delta - I) \sum_{s=1}^{\infty} \delta^{s-1} [\alpha \mathbf{x}_{it-1} + \boldsymbol{\nu}_{it-s}] + (\alpha - I) \mathbf{x}_{it} + \boldsymbol{\nu}_{it}$$

Importantly, this is a regression equation that could, hypothetically, be consistently estimated. However, there are easier ways to consider a test of stationarity. Consider the simple null hypothesis that $\delta = I$, so that the equilibrium stock process is a random walk. Under this null hypothesis, the equilibrium winnings process does not respond to the previous period's evening state. Therefore, consider the following regression specification:

$$\mathbf{w}_{it}^T \mathbf{z}_t^g = \beta^0 + \beta^1 \mathbf{r}_{it} + \beta^2 \mathbf{r}_{it-1} + \vartheta_{it} \tag{8}$$

Under this null hypothesis $\beta^2 = \mathbf{0}$. This test is essentially an Anderson-Rubin test. The intuition is that we consider whether winnings respond to lagged non-stock factors. These lagged factors likely impact lagged winnings. If they are found to impact present winnings this suggests that current winnings responds to the lagged stock. If current winnings do not depend on the lagged stock, this suggest the food bank ignores their previous stocks, so the equilibrium stock process follows a random walk.

Results for this test are presented in Appendix C.3. As well as presenting aggregated results, assuming every food bank has the same δ, I consider disaggregated results. However in this case I have significantly less power. Broadly, however, I find strong evidence of stationarity.

C.3 Results

Figure 17 presents results from this regression. I run each regression separately for each storage method, focusing on food banks that won at least 100 lots. I also include a second lag of $\mathbf{r}_t$, as this allows us to interpret the coefficients on $\mathbf{r}_{t-1}$.[57] I include factors for all storage methods in every regression, but only present results for the matching storage types, since these are expected to be the most useful. I also include a dummy variable (and its lags) for whether no lots of a particular type were auctioned on a particular day.

There is evidence of stationarity, given that we can reject the null hypothesis $\delta = I$ at the 1% significance level ($p < 0.001$). Most of the coefficients have the expected signs. Non-lagged factors are almost all statistically significant, with winnings increasing in the amount of food allocated, decreasing in the minimum distance between the bidder and lots, and decreasing in the number of rival bidders bidding on lots of that type. Coefficients on the lagged factors generally have the opposite signs, as we expect. Lagged distance is never statistically significant, even though contemporaneous distance is, while lagged pounds and rival bidders are generally significant. Comparing coefficients across lagged and unlagged variables we can extract δ. Assume zero-off diagonal elements, then diagonal elements are estimated to be in the region of 0.8. This uses the formula $\hat{\delta} = 1 - \frac{\hat{\beta^1}}{\hat{\beta^2}}$.

[57] Conditional on $\mathbf{r}_{t-2}$ it is reasonable that $\mathbf{r}_{t-1}$ and $\mathbf{r}_t$ are uncorrelated with $\mathbf{r}_{t-s}$ for $s \geq 3$, the omitted variables.

I also consider results for individual food banks. I run a specification as above allowing for individual food bank coefficients. When considering winnings for storage type l I only include factors and lagged factors for storage type l. I can reject the null-hypothoses that $\delta = I$ for 25% of analysed food banks. These are predominantly the large food banks observed regularly bidding and winning on the Choice System. For many food banks the tests are under powered, with even 58% of the unlagged factors being statistically insignificant at 5% level of significance. This is most likely due to having only a small amount of variation in winnings, given that many food banks do not win very frequently. Interestingly Distance and Lagged Distance are much more likely to be statistically significant when I run the individual analysis.

Figure 17: Results: Stationarity

	Non-Food	Dried	Tinned/Bottled	Refrigerated	Fresh
Total Pounds	0.018	0.023	0.017	0.020	0.018
	(0.004)	(0.005)	(0.004)	(0.006)	(0.005)
Lagged Pounds	-0.001	-0.003	-0.001	-0.004	-0.001
	(0.001)	(0.001)	(0.001)	(0.003)	(0.001)
Minimum Distance	-0.856	-1.528	-0.884	-0.480	-0.288
	(0.285)	(0.450)	(0.264)	(0.751)	(0.107)
Lagged Distance	0.171	-0.024	0.187	0.116	0.030
	(0.118)	(0.160)	(0.107)	(0.320)	(0.074)
Active Bidders	-53.638	-71.039	-28.601	-593.755	-43.187
	(19.999)	(29.534)	(22.572)	(189.673)	(18.443)
Lagged Bidders	14.254	34.077	15.683	84.825	7.250
	(6.940)	(15.380)	(7.975)	(61.444)	(4.552)
FB Fixed Effects	✓	✓	✓	✓	✓
2nd Lags	✓	✓	✓	✓	✓

Note: Standard Errors clustered within food bank and period.
Regression for storage method l includes regressors for all other storage methods also.

C.4 Covariance Stationarity

I now demonstrate that this assumption about the equilibrium stock process yields two sets of intuitive information about parameters $\boldsymbol{\mu}_i$ and Σ_i.

C.4.1 μ_i

Take an expectation of the equilibrium winnings equation for:

$$E[\mathbf{w}_{it}^T \mathbf{z}_t^g] = (\alpha - I)\boldsymbol{\mu}_i + (\delta - I)\sum_{s=1}^{\infty} \delta^{s-1}[\alpha\boldsymbol{\mu}_i + E[\boldsymbol{\varepsilon}]] + E[\boldsymbol{\varepsilon}]$$

Recognise that $\sum_{s=1}^{\infty} \delta^{s-1} = (I - \delta)^{-1}$, so that we are left with:

$$E[\mathbf{w}_{it}^T \mathbf{z}_t^g] = -\boldsymbol{\mu}_i$$

Therefore, on average winnings offset net local donation. I use this constraint to build prior means for $\boldsymbol{\mu}$. I do not impose this relationship on account of the difficulty of efficiently estimating $E[\mathbf{w}_{it}^T\mathbf{z}_t^g]$ in the presence of auto-correlation, meaning my estimates are likely to be imprecise.

C.4.2 Σ_i

Take the variance of the equilibrium winnings equation, recognising that we assumed $\mathbf{x}$ is uncorrelated over time, and uncorrelated with $\boldsymbol{\varepsilon}$:

$$Var(\mathbf{w}_{it}^T\mathbf{z}_t^g) = (\alpha - I)\Sigma_i(\alpha - I) + (\delta - I)[\sum_{s=1}^{\infty}\delta^{s-1}\alpha\Sigma_i\alpha\delta^{s-1}](\delta - I) + Var(\boldsymbol{\varepsilon}_t + (\delta - I)[\sum_{s=1}^{\infty}\delta^{s-1}\boldsymbol{\varepsilon}_{t-1})$$

$Var(\boldsymbol{\varepsilon}_t + (\delta - I)[\sum_{s=1}^{\infty}\delta^{s-1}\boldsymbol{\varepsilon}_{t-1})$ is evidently positive definite. Meanwhile, the infinite geometric series does not have a simple form. To simplify matters, focus on the case where δ, α, and Σ_i are diagonal matrices. This is relevant since I impose that Σ_i is diagonal in the empirical model. In this case, applying the rule for infinite geometric series, with $\delta_l < 1$ I can write:

$$Var(\mathbf{w}_{it}^T\mathbf{z}_{lt}^g) = (\alpha_l - 1)^2\Sigma_i^{ll} + \frac{(\delta_l - 1)^2}{1 - \delta_l^2}\alpha^2\Sigma_i^{ll} + Var(\varepsilon_{lt} + (\delta_l - 1)[\sum_{s=1}^{\infty}\delta_l^{s-1}\varepsilon_{lt-1})$$

Since $\alpha \in [0, 1]$, the first part of this expression varies between Σ_i^{ll} and $[(1 - \delta_l)/2]\Sigma_i^{ll}$, when $\alpha = (1 + \delta_l)/2$. Therefore, since $Var(\varepsilon_{lt} + (\delta_l - 1)[\sum_{s=1}^{\infty}\delta_l^{s-1}\varepsilon_{lt-1}) > 0$, I can write:

$$Var(\mathbf{w}_{it}^T\mathbf{z}_{lt}^g) > \frac{1 - \delta_l}{2}\Sigma_i^{ll}$$

Therefore, I can bound Σ_i^{ll}, conditional on δ. I consider two benchmark cases, $\delta_l = 0$, so that, on average each period, winnings totally adapt to changes in previous stocks. This implies $\Sigma_i^{ll} < 2Var(\mathbf{w}_{it}^T\mathbf{z}_{lt}^g)$. I also consider $\delta = 49/50$, so that winnings do not strongly react to previous stocks. This implies $\Sigma_i^{ll} < 100Var(\mathbf{w}_{it}^T\mathbf{z}_{lt}^g)$.

I use the $\delta = 49/50$ case for a hard upper bound on Σ_i^{ll}, under the prior that $\delta < 49/50$. I use the $\delta = 0$ case for what is essentially the prior mean of Σ_i^{ll}, albeit with a very low prior weight, given by the degrees of freedom in the Normal-inverse-Gamma distribution. Full details of how I build these weakly informative priors is given in Appendix H.

D Inverse Bid System

In this Appendix I demonstrate that, in addition to the transition equation given in Assumption 2, a food bank's optimisation problem yields the Observation and Censoring equations given in Equation 4. I focus on the case for quadratic parametrisation of k. The general case is presented in (Altmann, 2022).

D.1 Set-up

Imposing the parametrisation given in section 5.3, and conditional on $\mathbf{d}_i^*$, the bidder's maximisation problem is given by:

$$\max_{\mathbf{b}}\left\{\sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(\upsilon_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{s}_i^{ah} + \mathbf{s}_i^{agT}\Psi \mathbf{s}_i^{ag}] \qquad s.t.\ b_l \geq R_l\right\}$$

D.2 Simplification

The maximand can be simplified for:

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

I now prove this. First, recognise that $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s}) = 1$, since we sum over mutually exclusive and exhaustive events. This allows me to write, for example, $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s})\mathbf{s}_i^{gT}\Psi \mathbf{s}_i^g = \mathbf{s}_i^{gT}\Psi \mathbf{s}_i^g$.

Second, recognise that $\mathbf{s}_i^a = \mathbf{s}_i + \mathbf{z}\mathbf{w}_i^a$, where $\mathbf{w}_i^a$ is the $L \times 1$ vector with entry l equal to 1 if i wins lot l in combinatorial outcome a and zero otherwise. The matrix z just gives the size and composition of lots. Exploiting $\sum_a P_a = 1$, re-write the maximand as:

$$\begin{aligned}
&\sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(\upsilon_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{s}_i^{ah} + \mathbf{s}_i^{agT}\Psi \mathbf{s}_i^{ag}] \\
&\quad = \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(\upsilon_l - b_l) + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{gT}\Psi \mathbf{s}_i^g + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{z}^h \mathbf{w}_i^a + \mathbf{s}_i^{agT}\Psi \mathbf{s}_i^{ag} - \mathbf{s}_i^{gT}\Psi \mathbf{s}_i^g] \\
&= \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(\upsilon_l - b_l) + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{gT}\Psi \mathbf{s}_i^{gT} + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})[\Phi \mathbf{z}^h \mathbf{w}_i^a + \mathbf{w}_i^{aT}\mathbf{z}^{gT}\Psi(\mathbf{z}^g \mathbf{w}_i^a + 2\mathbf{s}_i^g)]
\end{aligned}$$

Where the final line follows from quadraticness: $\mathbf{s}_i^{agT}\Psi \mathbf{s}_i^{ag} = (\mathbf{s}_i + \mathbf{z}\mathbf{w}_i^a)^T\Psi(\mathbf{s}_i + \mathbf{z}\mathbf{w}_i^a)$ and so $\mathbf{s}_i^{agT}\Psi \mathbf{s}_i^{ag} - \mathbf{s}_i^{gT}\Psi \mathbf{s}_i^g = \mathbf{w}_i^{aT}\mathbf{z}^{gT}\Psi(\mathbf{z}^g \mathbf{w}_i^a + 2\mathbf{s}_i^g)$.

Finally, recognise that $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s})\mathbf{s}_i^a = \mathbf{s}_i + \sum_l \Gamma_l(b_l, d_l; \mathbf{s})\mathbf{z}_l$. This arises because stocks are additive in winnings.[58] This ensures that:

$$\begin{aligned}
&= \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(\upsilon_l - b_l + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{gT}\Psi \mathbf{s}_i^g) + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{gT}\Psi \mathbf{s}_i^{gT} + \sum_a P_a(\mathbf{b}, \mathbf{d}^*; \mathbf{s})\mathbf{w}_i^{aT}\mathbf{z}^{gT}\Psi \mathbf{z}^g \mathbf{w}_i^a \\
&= \sum_l \Gamma_l(b_l, d_l^*; \mathbf{s})(\upsilon_l - b_l + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{gT}\Psi \mathbf{s}_i^g + \mathbf{z}_l^{gT}\Psi \mathbf{z}_l^g + \sum_{m\neq l}\Gamma_m(b_m, d_m^*; \mathbf{s})\mathbf{z}_l^{gT}\Psi \mathbf{z}_m^g) + \Phi \mathbf{s}_i^h + \mathbf{s}_i^{gT}\Psi \mathbf{s}_i^{gT}
\end{aligned}$$

Where the final line follows from two points: (1) that $(\mathbf{z}^g \mathbf{w}_i^a)^T\Psi(\mathbf{z}^g \mathbf{w}_i^a) = \sum_l \sum_m (w_{il}^a \mathbf{z}_l^g)^T\Psi(w_{im}^a \mathbf{z}_m^g)$, which arises from quadraticness. Because this object only depends on pairs of winnings, we can marginalise out the probability of receiving a particular pair, so that: $\sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s})(\mathbf{z}^g \mathbf{w}_i^a)^T\Psi \mathbf{z}^g \mathbf{w}_i^a =$

[58] This result can be seen by focusing on the expectation of $\mathbf{w}_i^{aT}\mathbf{z}_l$ for one particular l. As we sum across all the combinations in which they win lot l, the sum of these probabilities is just Γ_l, the marginal probability they win lot l.

$\sum_l \sum_m Prob(\text{win } l \text{ and } m|\mathbf{b},\mathbf{d};\mathbf{s})(\mathbf{z}_l^g)^T\Psi(\mathbf{z}_m^g)$. (2), imposing part (iv) of Assumption 4: $Prob(\text{win } l \text{ and } m|\mathbf{b},\mathbf{d};\mathbf{s}) = \Gamma_l(b_l,d_l;\mathbf{s})\Gamma_m(b_m,d_m;\mathbf{s})$ for $m \neq l$, and $\Gamma_l(b_l,d_l;\mathbf{s})$ otherwise.

D.3 First Order Conditions, conditional on entry

Written out in it's full simplified form, the lagrangian for this problem is given by:

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

Where Λ_l give the lagrangian multipliers. Necessary first order conditions are given by:

$$0 = \nabla_l\Gamma_l(b_l^*,d_l^*;\mathbf{s})(\upsilon_l - b_l^* + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi\mathbf{z}_l^g + 2\mathbf{z}_l^{gT}\Psi\mathbf{s}_i^g + 2\sum_{m\neq l}\Gamma_m(b_m^*,d_m^*;\mathbf{s})\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g) - \Gamma_l(b_l^*,d_l^*;\mathbf{s}) + \Lambda_l^*$$

Which rearranges for:

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

Let $y_l^* = b_l^* + \frac{\Gamma_l(b_l^*,d_l^*;\mathbf{s})}{\nabla_b\Gamma_l(b_l^*,d_l^*;\mathbf{s})} - \Lambda_l^*$. When we observe $b_l^* > R_l$, we can infer $\Lambda_l^* = 0$, so that $y_l^* = b_l^* + \frac{\Gamma_l(b_l^*,d_l^*;\mathbf{s})}{\nabla_b\Gamma_l(b_l^*,d_l^*;\mathbf{s})} = y_l$. In this case, the bidder is not constrained.

D.4 Reservation Price Bidding

When we observe $b_l^* = R_l$, the First Order Conditions break down, since as made clear in Section 5.2, beliefs are non-differentiable at the reservation price due to the non-negligible probability of ties. Therefore, consider the bidder's decision to bid at the reservation price, bidding vector b^*, compared to just above the reservation price at $R_l + 1$ playing vector b^+. Elements $m \neq l$ of these vectors will otherwise be equal. This implies that:

$$\sum_l \Gamma_l(b_l^*, d_l^*; \mathbf{s})(\upsilon_l - b_l^* + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{gT}\Psi\mathbf{s}_i^g + \sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s})\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)$$
$$\geq \sum_l \Gamma_l(b_l^+, d_l^*; \mathbf{s})(\upsilon_l - b_l^+ + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{gT}\Psi\mathbf{s}_i^g + \sum_m \Gamma_m(b_m^+, d_m^*; \mathbf{s})\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)$$

Therefore

$$\Gamma_l(R_l, d_l^*; \mathbf{s})(\upsilon_l - R_l + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s})\mathbf{z}_m^g))$$
$$\geq \Gamma_l(R_l+1, d_l^*; \mathbf{s})(\upsilon_l - R_l - 1 + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s})\mathbf{z}_m^g))$$

Therefore

$$[\Gamma_l(R_l+1, d_l^*; \mathbf{s}) - \Gamma_l(R_l, d_l^*; \mathbf{s})](\upsilon_l - R_l + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s})\mathbf{z}_m^g)) \leq \Gamma_l(R_l+1, d_l^*; \mathbf{s})$$

$$y_l = \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{gT}\Psi\mathbf{s}_i^g + \mathbf{z}_l^{gT}\Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s})\mathbf{z}_m^g) + \upsilon_l \leq R_l + \frac{\Gamma_l(R_l+1, d_l^*; \mathbf{s})}{\Gamma_l(R_l+1, d_l^*; \mathbf{s}) - \Gamma_l(R_l, d_l^*; \mathbf{s})} = y_l^*$$

Therefore, $y_l^* \geq y_l$

D.4.1 Bidding R_l vs Not Bidding

At the margin, the bidder must weakly prefer to enter and bid the reservation price, playing bidding/entry vector $\mathbf{b}^*, \mathbf{d}^*$, than to not enter at all, playing bidding/entry vector $\mathbf{b}^-, \mathbf{d}^-$. These vectors are identical apart from for lot l. This implies:

$$\sum_l \Gamma_l(b_l^*, d_l^*; \mathbf{s})(\upsilon_l - b_l^* + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{gT}\Psi\mathbf{s}_i^g + \sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s})\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)$$
$$\geq \sum_l \Gamma_l(b_l^-, d_l^-; \mathbf{s})(\upsilon_l - b_l^- + \Phi \mathbf{z}_l^h + 2\mathbf{z}_l^{gT}\Psi\mathbf{s}_i^g + \sum_m \Gamma_m(b_m^-, d_m^-; \mathbf{s})\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)$$

Therefore

$$\Gamma_l(R_l, d_l^*; \mathbf{s})(\upsilon_l - R_l + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi(\mathbf{z}_l^g + 2\mathbf{s}_i^g + 2\sum_m \Gamma_m(b_m^*, d_m^*; \mathbf{s})\mathbf{z}_m^g)) \geq 0$$

Therefore

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

Therefore $R_l \leq y_l \leq y_l^*$

D.5 Entry Decisions

If a food bank chooses not to enter the auction for lot l, then at the margin they must weakly prefer to not enter the auction, than to enter and bid the reservation price. This is just the complement of the previous inequality, allowing us to infer that $d_l^* = 0$ implies $y_l \leq R_l$.[59].

D.6 Monotonicity

I now prove that the inverse bid system is strictly monotonic for $b_l > R_l$. This broadly extends equivalent results from Altmann (2022). That is, I show that the Jacobian matrix of the inverse bid system, differentiated with respect to bids, is positive definite. I drop the dependence on t for notational convenience. The proof involves first finding the Hessian matrix for the bidder's maximisation problem. This matrix being negative definite is a necessary condition of optimising behaviour. I then differentiate the inverse bid system, before inserting the hessian matrix, and recognising that the resulting matrix must be positive definite.

Proof:

1. The Second order necessary conditions for maximising behaviour are given by the matrix with entry l, m:

$$\frac{\partial L(\mathbf{b}|\mathbf{d},\mathbf{s},\boldsymbol{\upsilon})}{\partial b_l \partial b_m} = \begin{cases} \nabla_l^2 \Gamma_l(b_l^*, d_l^*; \mathbf{s})(\upsilon_l - b_l^* + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi(\mathbf{z}_l^g + \mathbf{s}_i^g + 2\sum_{n\neq l}\Gamma_n(b_n^*, d_n^*; \mathbf{s})\mathbf{z}_n^g) \\ \qquad -2\nabla_l \Gamma_l(b_l^*, d_l^*; \mathbf{s}) & \text{if } l = m \\ 2\nabla_l \Gamma_l(b_l^*, d_l^*; \mathbf{s}) \nabla_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_l^{gT} \Psi \mathbf{z}_m^g & \text{if } l \neq m \end{cases}$$

   I focus on the region for which the constraint does not bind. Therefore a necessary condition for optimising behaviour is that this matrix is negative definite.

2. The first order conditions can be written as:

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

   Recognise that we can substitute this into the hessian for:

$$\nabla_{l,m}^2 L(\mathbf{b}|\mathbf{d},\mathbf{s},\boldsymbol{\upsilon}) = \begin{cases} \frac{\nabla_l^2 \Gamma_l(b_l^*, d_l^*; \mathbf{s})[\Gamma_l(b_l^*, d_l^*; \mathbf{s}) - \Lambda_l^*]}{\nabla_l \Gamma_l(b_l^*, d_l^*; \mathbf{s})} - 2\nabla_l \Gamma_l(b_l^*, d_l^*; \mathbf{s}) & \text{if } l = m \\ 2\nabla_l \Gamma_l(b_l^*, d_l^*; \mathbf{s}) \nabla_m \Gamma_m(b_m^*, d_m^*; \mathbf{s}) \mathbf{z}_l^{gT} \Psi \mathbf{z}_m^g & \text{if } l \neq m \end{cases}$$

3. The inverse bid system can be written as:

$$\xi_l(\mathbf{b}, \mathbf{d}; \mathbf{s}) = b_l + \frac{\Gamma_l(b_l, d_l; \mathbf{s})}{\nabla_l \Gamma_l(b_l, d_l; \mathbf{s})} - \Phi \mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi(\mathbf{z}_l^g + 2\mathbf{s}^g + 2\sum_{m\neq l}\Gamma_m(b_m, d_m)\mathbf{z}_m^g) \geq \upsilon_l$$

4. Differentiate this with respect to $\mathbf{b}$ for:

$$\frac{\partial \xi_l(\mathbf{b}, \mathbf{d}; \mathbf{s})}{\partial b_m} = \begin{cases} 2 - \frac{\Gamma_l(b_l, d_l; \mathbf{s}) \nabla_l^2 \Gamma_l(b_l, d_l; \mathbf{s})}{\nabla_l \Gamma_l(b_l, d_l; \mathbf{s})^2} \\ -2\nabla_m \Gamma_m(b_m, d_m) \mathbf{z}_l^{gT} \Psi \mathbf{z}_m^g \end{cases}$$

[59] Furthermore, negative definiteness of Ψ implies that if they prefer to not bid than enter and bid the reservation price, they also cannot prefer to bid strictly above the reservation price

5. Recognise that for $b_l > R_l$, so that $\Lambda_l = 0$, this is just the negative of the hessian divided by $\nabla_l \Gamma_l(b_l, d_l; \mathbf{s})$. Therefore, since $\nabla_l \Gamma_l(b_l, d_l; \mathbf{s})$ is strictly positive, the Jacobian of the inverse bid system must be positive definite.

□

E Discriminatory Auctions

In this Appendix I discuss how the discriminatory auctions of homogenous lots are take into account. In Appendix E.1 I explain the rules of the discriminatory auctions. In Appendix E.2 I derive the Inverse Bid System in the presence of discriminatory auctions, presenting a generalisation of Appendix D. In Appendix E.3 I discuss how I use the inverse bid system in estimation.

E.1 Framework

In a discriminatory auction of R homogenous lots, food banks place as many bids as they like, and lots are allocated to the R highest bidders. Bidders pay they bids. Lots are then allocated to the R highest bidders who pay their bids, if at least R bids were placed.

7% of unique auctions that occur contain more than one homogenous good, and are auctioned in discriminatory fashion. Of all the lots allocated, 21% of lots are auctioned with at least one additional identical load. The vast majority of lots sold at the reservation price are lots from discriminatory auctions with a large number of homogenous loads being auctioned in this manner (the remainder are fresh loads). Food banks recognise that the lowest winning bid is likely to be very low. Loads allocated in this manner are always homogenous and come from the same source.[60]

Figure 18 panel (A) shows a histogram of the number of loads included in each unique auction, conditional on at least two loads. As not every load may be sold on a particular date, Panel (B) shows the number of homogenous loads for each auction × date, conditional on at least two loads.

E.2 Adjusted Inverse Bid System

The only difference between simultaneous and discriminatory first-price auctions is that when the bidder wins on their rth lowest bid on lot l, they must also win on all higher bids. Write the rth bid on auction l as b_{lr}. Bids are ascending, such that $b_{lr} \leq b_{lr+1}$, up to $r = R_l$ the total number of loads contained in the lot. Food bank i's belief about the probability they win on their rth bid on lot l is given by $\Gamma_l(b_{lr})$, suppressing the entry decision and state variables for ease of notation. Their belief they win on bids r through to R, but no lower, is given by $\Gamma_l(b_{lr}) - \Gamma_l(b_{lr-1})$. This is the probability the lowest rival winning bid is between b_{lr} and b_{lr-1}.

Because lots are homogenous, I treat the lot specific value v_{il} as constant across the loads in auction l. Making use of the assumed parametrisation, the expected payoff is given by:

[60] However some homogenous lots, that are all donated by the same donors, are auctioned using simultaneous auctions. I do not analyse the determinants of this decision.

Figure 18: Distribution of Homogenous Loads

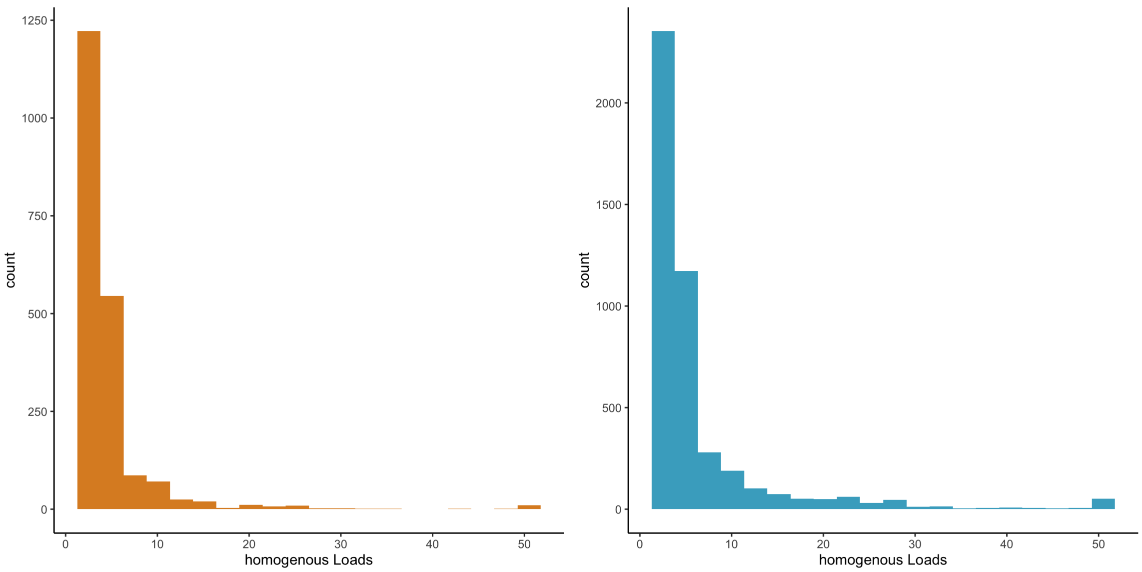

Note: These plots show histograms of the number of homogenous loads included in each auction, conditional on at least two homogenous loads. Panel (A) shows the number of loads for each unique auction. Panel (B) shows the number of loads for each date × auction, recognising that not every load is sold right away. More than 50 loads are grouped into the 50 category. This includes one auction with 181 loads, and 9 other auctions with between 60 and 80 Loads.

$$
\begin{aligned}
\pi(\mathbf{b},\mathbf{d}) = & \sum_{l}^{L}\sum_{r}^{R_l}\Gamma_l(b_{lr},d_{lr})v_l - b_{lr} + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g]) \\
& + \sum_{l}^{L}\sum_{r}^{R_l}\Gamma_l(b_{lr},d_{lr})\mathbf{z}_l^{gT}\Psi[(R_l-r)\mathbf{z}_l^g + \sum_{m\neq l}\sum_{n}^{R_m}\Gamma_m(b_{mn},d_{mn})\mathbf{z}_m^g] + \sum_{n=1}^{r-1}\Gamma_l(b_{ln},d_{ln})\mathbf{z}_l^{gT}\Psi\mathbf{z}_l^g
\end{aligned}
$$

The first line gives the lot-specific component of the pay-off, while the second line gives the combinatorial component - how the pay-off varies with winnings from other auctions. In the simultaneous only case this line is given by $\sum_l^L \Gamma_l(b_l,d_l)\mathbf{z}_l^{gT}\Psi[\sum_{m\neq l}\Gamma_m(b_m,d_m)\mathbf{z}_m^g]$ only. Recognise that if they win with bid b_{lr}, then they win the $(R_l - r)$ higher lots with certainty. This is why the $(R_l - r)\mathbf{z}_l^g$ term does not have a probability multiplier. Likewise, with probability $\Gamma_l(b_{ln},d_{ln})$ they also win on bid b_{lr} (for $n < r$). So, in expectation, they also gain this $\Gamma_l(b_{ln},d_{ln})\mathbf{z}_l^{gT}\Psi\mathbf{z}_l^g$ combinatorial term.

Conditional on $d_{lr} = 1$, first order conditions with respect to bid b_{lr} are given by:

$$\Gamma_l(b_{lr}, d_{lr}) = \nabla_{b_{lr}}\Gamma_l(b_{lr}, d_{lr})(\upsilon_l - b_{lr} + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g])$$
$$+ \nabla_{b_{lr}}\Gamma_l(b_{lr}, d_{lr})\mathbf{z}_l^{gT}\Psi[(R_l - r)\mathbf{z}_l^g + \sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_n^{r-1}\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^g]$$
$$+ \mathbf{z}^{gT}\Psi\nabla_{b_{lr}}\Gamma_l(b_{lr}, d_{lr})[\sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_{r+1}^{R_l}\mathbf{z}_l^g]$$

The inverse bid system is given by:

$$\xi_{lr}(\mathbf{b}; k) = b_{lr} + \frac{\Gamma_l(b_{lr}, d_{lr})}{\nabla_{b_{lr}}\Gamma_l(b_{lr}, d_{lr})} - \Phi\mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g]$$
$$- \mathbf{z}_l^{gT}\Psi[2(R_l - r)\mathbf{z}_l^g + 2\sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_n^{r-1}\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^g]$$

Next, we consider the decision to enter and bid the reservation price, rather than bid just above it. Importantly, if $b_{lr} = R + 1$ then all lower bids must be ether the same, at the reservation price, or not entered. Setting the utility at the reservation price greater than or equal to utility just above the reservation price we obtain the following upper bound on ξ_{lr}:

$$\xi_{lr}(\mathbf{b}; k) \leq b_{lr} + \frac{\Gamma_l(R+1, 1)}{\Gamma_l(R+1, 1) - \Gamma_l(R, 1)} - \Phi\mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g]$$
$$- \mathbf{z}_l^{gT}\Psi[2(R_l - r)\mathbf{z}_l^g + 2\sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_n^{r-1}\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^g]$$

Next, consider the decision to enter and bid the reservation price, versus not bidding at all. We obtain the following lower bound for ξ_{lr}:

$$\xi_{lr}(\mathbf{b}; k) \geq b_{lr} - \Phi\mathbf{z}_l^h - \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}_t^g + 2(R_l - r)\mathbf{z}_l^g + 2\sum_{m\neq l}\sum_n^{R_m}\Gamma_m(b_{mn}, d_{mn})\mathbf{z}_m^g + \sum_n^{r-1}\Gamma_l(b_{ln}, d_{ln})\mathbf{z}_l^g]$$

As in Appendix D, the decision not to enter then yields the same upper bound as this lower bound. There is a clear similarity between this system of equations and those presented previously.

E.3 Computation

E.3.1 First Stage

Food bank i wins lot l, load r, with the probability that their bid exceeds the lowest winning bid on that lot. Therefore it is their belief about the distribution of the lowest winning bid that matters. When fewer food banks place bids than loads are being auctioned, this is just the reservation price.

Therefore, when estimating beliefs I ignore all higher winning bids.

E.3.2 Second Stage

The difficulty in the second stage is that the lot specific idiosyncratic pay-off is assumed to be perfectly correlated within a discriminatory auction. Allowing these objects to vary, even if highly correlated, does not make sense given loads are perfectly homogenous. By definition of the homogenous lots, there can be no unobserved variation in lot characteristics across loads.

When performing the data augmentation step in the second stage of the procedure I treat discriminatory auctions correctly. That is, I correctly sample censored bids from their posterior distribution and sample states from their conditional posterior. However, in the Gibbs Sampling step I only use information from the highest bid (which may also be censored). If the idiosyncratic terms are perfectly correlated then I do not gain any additional information by using lower bids.[61] Finally, when considering model fit by simulating the Choice System, I do simulate these auctions (albeit limiting food banks to only be able to place up to 5 bids on a single auction). Therefore if my simplification does lead to inaccuracy or bias, this should become evident.

E.3.3 Third Stage

Finally, in the third step I treat these auctions properly in how I evaluate the ex-ante continuation value. The derivation presented in Appendix G extends easily to allow for the discriminatory case, and only requires summing over the probability of these combinatorial wins.

F Semi-parametric Identification

In this Appendix I prove that the model is semi-parametrically point-identified. I prove that it is point-identified under the parametric restriction on the pseudo-static payoff given by Equation 3. This restriction allows me to make use of the Inverse Bid System derived in Appendix D.

I make two additional simplifying assumption: I assume that reservation prices do not bind. Binding reservation prices ensure the first order conditions do not hold with equality. However, as discussed in Altmann (2022), reservation prices are not a first-order issue, and do not substantially alter the problem. In the same way that a censored regression model, which requires a Tobit or MAD specification, does not substantially alter the problem of identifying the regression coefficients. The key intuition garnered from this simplified approach extends to the case with reservation prices.

[61] In practice specification error will means I cannot rationalise the model with identical lot specific pay-offs. I could treat them as correlated, and estimate a correlation parameter, but in practice I will gain very little additional information from doing so. First, if they do not place bids at all in a particular auction I gain no additional information from considering additional bids. If they place one bid, but not a second, some information can be gained by considering why they didn't place this second bid. However, even given specification error we expect a very high degree of correlation, meaning I do not gain a large amount of information beyond that contained in the first bid.

I also assume that $\mathbf{z}_t^g$ has full rank, ensuring the generalised (left) inverse exists. This ensures that information about each dimension of the stock propagates through their bidding behaviour, so that we can separately identify the effects of each stock dimension. In practice this assumption generally holds, except in the case where one type of good is not auctioned one period. Strictly speaking, identification only requires 3 observations with full rank in a row.

While the argument only requires $T \geq 3$ for point identification, asymptotics require $T \rightarrow \inf$ because additional observations help us further pin down the identified objects objects. I treat Γ as trivially identified, because it is identified from observations of winning bids. I do not explicitly consider identification of the λ_i parameters, which give the marginal value of wealth for each bidder. These parameters are identified by comparing bidding behaviour across bidders, and requires the normalising assumption that the variance of lot-specific values is constant across bidders.

The identification argument proceeds as follows: In Appendix F.1 I detail sufficient assumptions model primitives for semi-parametric point identification. In Appendix F.2 I show that the Observation Equation and Transition Equations can be re-written as a 'Random-walk Linear State Space' model. In Appendix F.3 I show that under these assumptions the Random-walk Linear State Space model is identified, proving that k, F^x, and F^υ are point identified. Finally, in Appendix F.4 I prove that conditional on point identification of k, F^x, and F^υ, j is point identified.

F.1 Assumptions on the Pseudo-static Primitives

Identification of k, F^x, and F^υ requires assumptions 1 - 4, however I do not require that either $\boldsymbol{\upsilon}_{it}$ or $\mathbf{x}_{it}$ are normally distributed and strictly exogenous. Instead, I require the following (dropping the i subscript for notational clarity):

Assumption 5. *i)* $E[\boldsymbol{\upsilon}_s|\mathbf{z}_{lt}^g, \mathbf{z}_{lt}^h, d_{tl}] = 0 \;\; \forall l, m, t, s$

ii) $E[\mathbf{x}_s|\mathbf{z}_{lt}^g, \mathbf{z}_{lt}^h, d_{tl}] = 0 \;\; \forall l, m, t, s$

iii) $E[\boldsymbol{\upsilon}_s|\mathbf{w}_t^T \mathbf{z}_t^g] = 0 \;\; \forall t < s$

iiii) $E[\mathbf{x}_s|\mathbf{w}_t^T \mathbf{z}_t^g] = 0 \;\; \forall t < s$

v) $rank(E[(\mathbf{w}_t^T \mathbf{z}_t^g) \otimes (\mathbf{z}_t^g, \mathbf{z}_t^h, \mathbf{d}_t)]) = 5 \;\; \forall t$

Where d_{tl} gives the distance between the bidder and lot l. Essentially, I have de-meaned the lot-specific value υ_{lt}. The first two parts of the assumption ensures that the unobservables are mean independent of any of the lot-specific observables, a relatively standard assumption. Parts *iii)* and *iiii)* ensure that future unobservables are mean independent of past winnings. However, it does allow present and future winnings to depend on the un-observables. For example, the higher is υ_{ilt}, or the lower is $\mathbf{x}_{it}$ the more aggressively they bid, and the more likely they are to win the lot. Similarly, part *v)* actually necessitates this, essentially stating that winnings must be correlated with observables. This assumption is validated by Figure 6.

F.2 Random Walk Linear State Space form

In Appendix F.3 I demonstrate the identification of parameters in a Random Walk Linear State Space model. Therefore I first need to show that the model as presented above can be written in Random Walk Linear State Space form. I then show that the assumptions used in Appendix F.3 follow from the Assumption 5. I then briefly discuss the intuition underlying this identification argument, and how it maps to the intuition presented in Section 4.4.

A Random Walk Linear State Space model is written in the form: $\mathbf{y}_t = \Psi \mathbf{s}_t + \mathbf{x}_t\gamma + \boldsymbol{\upsilon}_t$. Importantly, $\mathbf{y}_t$ is affine in the unobserved state $\mathbf{s}_t$. The observation and transition equations of my model are given by:

$$b_{lt} + \frac{\Gamma_l(b_{lt})}{\nabla_b \Gamma_l(b_{lt})} = \Phi \mathbf{z}_{lt}^h + \alpha d_{lt} + \mathbf{z}_{lt}^{gT} \Psi(\mathbf{z}_{lt}^g + 2\mathbf{s}_t + 2\sum_{m \neq l} \Gamma_m(b_{mt})\mathbf{z}_{mt}) + \upsilon_{lt}$$

&

$$\mathbf{s}_t = \mathbf{s}_{t-1} + \mathbf{w}_{t-1} + \boldsymbol{\mu} + \mathbf{x}_t$$

Stack the observation equation over l, writing:

$$\begin{pmatrix} b_{1t} + \frac{\Gamma_1(b_{1t})}{\nabla_b \Gamma_1(b_{1t})} \\ \vdots \\ b_{Lt} + \frac{\Gamma_L(b_{Lt})}{\nabla_b \Gamma_L(b_{Lt})} \end{pmatrix} = \mathbf{z}_t^h \Phi^T + \mathbf{d}_t^T \alpha + \mathbf{z}_t^g \Psi \mathbf{s}_t + \begin{pmatrix} \mathbf{z}_{1t}^{gT} \Psi(\mathbf{z}_{1t}^g + 2\sum_{m \neq 1} \Gamma_m(b_{mt})\mathbf{z}_{mt}) \\ \vdots \\ \mathbf{z}_{Lt}^{gT} \Psi(\mathbf{z}_{Lt}^g + 2\sum_{m \neq L} \Gamma_m(b_{mt})\mathbf{z}_{mt}) \end{pmatrix} + \boldsymbol{\upsilon}_t$$

Unfortunately the $\mathbf{z}_{lt}^{gT} \Psi(\mathbf{z}_{lt}^g + 2\sum_{m \neq l} \Gamma_m(b_{mt})\mathbf{z}_{mt})$ terms do not have a simple vector form. However, we can write them in a 'linear-in-parameters' form.[62] Therefore, I write $\mathbf{z}_{lt}^{gT} \Psi(\mathbf{z}_{lt}^g + 2\sum_{m \neq l} \Gamma_m(b_{mt})\mathbf{z}_{mt}) = \mathbf{a}_{lt}\gamma$, where γ is a 25×1 vector with entry $\gamma_{5(p-1)+q} = \Psi_{p,q}$, and $\mathbf{a}_{lt}$ is a 1×25 row vector with entry $a_{5(p-1)+q,tl} = z_{p,tl}^g(z_{q,tl}^g + \sum_{m \neq l} \Gamma_m(b_{mt}) z_{q,tm}^g)$

Next, left multiply this system of equations by the generalised left inverse $(\mathbf{z}_t^g)^{-1}$:

$$(\mathbf{z}_t^g)^{-1}[\mathbf{b}_t + (\nabla_{\mathbf{b}} \Gamma(\mathbf{b}_t))^{-1}\Gamma(\mathbf{b})] = \Psi \mathbf{s}_t + (\mathbf{z}_t^g)^{-1}\mathbf{z}_t^h \Phi^T + (\mathbf{z}_t^g)^{-1}\mathbf{d}_t^T \alpha + (\mathbf{z}_t^g)^{-1}\mathbf{a}_t \gamma + (\mathbf{z}_t^g)^{-1}\boldsymbol{\upsilon}_t$$

This equation, coupled with the transition equation, constitute a Random Walk Linear State Space model. Essentially, Instead of considering an equation for each lot, I have reduced the dimensionality to focusing on only one equation for each dimension of the state variable. However, just like the original observation equation, this equation is also endogenous, $E[a_{n,t}\upsilon_{lt}] \neq 0$ due to possible correlation between $\Gamma_m(b_{itm})$ and υ_{lt}. However, we have an obvious instrument for $\mathbf{a}$, given by $\tilde{\mathbf{a}}_t$ where $\tilde{a}_{5(p-1)+q,tl} = z_{p,tl}^g z_{q,tl}^g$. The rank condition for this instrument is trivial. Now I show that the Assumption 6 follows from Assumption 5.

[62] This is just like how the regression $y_i = c + bx_i + ax_i^2 + \varepsilon_i$ can be written as $y_i = \begin{pmatrix} 1 & x_i \end{pmatrix} \begin{pmatrix} c & \frac{b}{2} \\ \frac{b}{2} & a \end{pmatrix} \begin{pmatrix} 1 \\ x_i \end{pmatrix} + \varepsilon_i$ or $y_i = \begin{pmatrix} 1 & x_i & x_i^2 \end{pmatrix} \begin{pmatrix} c \\ b \\ a \end{pmatrix} + \varepsilon_i$.

Proposition 2. *Let $\mathbf{q}_{nt}$ denote row n of $(\mathbf{z}_t^g)^{-1}$. Assumption 5 implies:*

i) $E[\mathbf{q}_{nt}(\mathbf{z}_t^h, \mathbf{d}_t^T)(\boldsymbol{\upsilon}_s^T)\mathbf{q}_{ms}^T] = 0 \ \ \forall n, m, t, s$

ii) $E[\mathbf{q}_{nt}(\mathbf{z}_t^h, \mathbf{d}_t^T)\mathbf{x}_s] = 0 \ \ \forall n, t, s$

iii) $E[\mathbf{w}_t^T \mathbf{z}_t(\boldsymbol{\upsilon}_s^T)\mathbf{q}_{ms}^T] = 0 \ \ \forall m, t < s$

iiii) $E[\mathbf{w}_t^T \mathbf{z}_t \mathbf{x}_s^T] = 0 \ \ \forall t < s$

v) $rank(E[(\mathbf{w}_t^T \mathbf{z}_t^g) \otimes (\mathbf{z}_t^g, \mathbf{z}_t^h, \mathbf{d}_t)]) = 5 \ \ \forall t$

Results *iiii)* and *v)* are as they were in Assumption 5. The remaining results follow by applying the Law of Iterated Expectations, before applying the relevant part of Assumption 5. The argument in Appendix F.3 requires that $(\mathbf{z}_{tl}^g, \mathbf{z}_{tl}^h, d_{tl})$ are not perfectly collinear, and vary over time.

The identification proof in the next section workers by first differencing both the augmented observation question, and minusing $\mathbf{s}_{t-1}$ from the transition equation. We then insert $\Delta\mathbf{s}_t$ into the observation equation, giving us an endogenous regression equation, with endogeneity arising through $\mathbf{a}_t$, but also through the correlation between $\mathbf{w}_{t-1}^T \mathbf{z}_{t-1}^g$ and $\boldsymbol{\upsilon}_{t-1}$.

Identification of Ψ is then driven by (the plausibly exogenous component of) variation in both $\mathbf{w}_{t-1}^T \mathbf{z}_{t-1}^g$ and $\mathbf{a}_t$, just as in a standard instrumental variable setting. Identification of $\boldsymbol{\mu}$ then arises, conditional on Ψ, from the average change in observed bids over time. Finally, Σ is identified by considering the covariance of bids over time, conditional on winnings, through a standard variance decomposition.

F.3 Identification of the RWLSS model

In this subsection I prove that the Random Walk State Space model, with state control variables and endogenous regressors, is point identified. The non-random walk case is a simple extension of this argument. The model can be written as follows:

$$\underbrace{\overbrace{\mathbf{y}_t}^{N\times 1} = \Psi \mathbf{s_t} + \overbrace{\mathbf{w}_t}^{exog}\beta + \overbrace{\mathbf{x}_t}^{endog}\gamma + \boldsymbol{\varepsilon}_t}_{\text{Observation}} \quad \& \quad \underbrace{\overbrace{\mathbf{s}_t}^{S\times 1} = \mathbf{s_{t-1}} + \overbrace{\mathbf{c}_t}^{control} + \overbrace{\boldsymbol{\mu}}^{constant} + \boldsymbol{\upsilon}_t}_{\text{Transition}} \quad \& \quad \underbrace{\overbrace{\mathbf{x}_t}^{M\times 1} = \Pi \overbrace{\mathbf{z}_t}^{Z\times 1} + \boldsymbol{\vartheta}_t}_{Instruments} \tag{9}$$

Without loss of generality I require that Ψ has full rank. $Var(\boldsymbol{\upsilon}_t) = \boldsymbol{\Sigma}$ is also a parameter of interest. Impose the following:

Assumption 6. *i)* $E[\mathbf{w}_t(\varepsilon_{ns}, \upsilon_{ms})] = 0, \ E[\mathbf{z}_t(\varepsilon_{ns}, \upsilon_{ms})] = 0 \ \ \forall n, m, t, s$

ii) $E[\mathbf{c}_t(\varepsilon_{ns}, \upsilon_{ms})] = 0 \ \ \forall n, m, t \leq s$

iii) $rank(\Pi) = M, \quad rank(E[\mathbf{c}_t(\mathbf{w}_{nt-1}, \mathbf{z}_t^T)]) = S \quad rank(E[\Delta\mathbf{w}_{nt}^T \Delta\mathbf{w}_{nt}]) = dim(\mathbf{w}_{nt})$

Importantly, this allows $E[\mathbf{c}_t(\varepsilon_{nt-1}, \upsilon_{mt-1})] \neq 0$.

Proposition 3. *Identification of the Random Walk State Space model*

Under Assumption 6, $\Psi, \beta, \gamma, \boldsymbol{\mu}$ are point identified.

Proof of this proposition simply involves an argument based on our moment conditions.

Proof:

1. Re-arrange the model equations for:

$$\mathbf{y}_t - \mathbf{w}_t\beta - \mathbf{x}_t\gamma - \boldsymbol{\varepsilon}_t = \Psi\mathbf{s}_t \qquad \& \qquad \Psi\mathbf{s}_t = \Psi\mathbf{s}_{t-1} + \Psi\mathbf{c}_t + \Psi\boldsymbol{\mu} + \Psi\boldsymbol{\upsilon}_t$$

2. Substitute the Observation equation into the transition equation, and re-arrange this for our estimating equation:

$$\Delta\mathbf{y}_t = \Psi\boldsymbol{\mu} + \Psi\mathbf{c}_t + \Delta\mathbf{w}_t\beta + \Delta\mathbf{x}_t\gamma + \Delta\boldsymbol{\varepsilon}_t + \Psi\boldsymbol{\upsilon}_t$$

   Write $\bar{\boldsymbol{\mu}} = \Psi\boldsymbol{\mu}$, where evidently if both Ψ and $\bar{\boldsymbol{\mu}}$ are point identified, then so is $\boldsymbol{\mu}$.

3. Focus on row n of this equation. Write Ψ_n as row n of the matrix Ψ. Exploiting that the transpose of a scalar is just itself, this equation can be re-written as: $\Delta y_{nt} = \bar{\mu}_n + \mathbf{c}_t^T\Psi_n^T + \Delta\mathbf{w}_{nt}\beta + \Delta\mathbf{x}_{nt}\gamma + \Delta\boldsymbol{\varepsilon}_{nt} + \Psi_n\boldsymbol{\upsilon}_t$

4. Identification then focuses on the following moment condition, which follows from Assumption 6: $E[(1, \mathbf{w}_{nt}, \Delta\mathbf{w}_{nt}, \Delta\mathbf{z}_t^T)^T(\Delta\boldsymbol{\varepsilon}_{nt} + \Psi_n\boldsymbol{\upsilon}_t)] = 0$. We essentially use $\mathbf{z}$ as an instrument for $\mathbf{x}$, and lagged $\mathbf{w}$ as an instrument for $\mathbf{c}$. Substitute in the estimating equation for:

$$E[(1, \mathbf{w}_{nt-1}, \Delta\mathbf{w}_{nt}, \Delta\mathbf{z}_{nt})^T(\Delta y_{nt} - \bar{\mu}_n - \mathbf{c}_t^T\Psi_n^T - \Delta\mathbf{w}_{nt}\beta - \Delta\mathbf{x}_{nt}\gamma)] = 0$$

5. Rearrange this equation for:

$$E[\begin{pmatrix} 1 \\ \mathbf{w}_{nt-1}^T \\ \Delta\mathbf{w}_{nt}^T \\ \Delta\mathbf{z}_t \end{pmatrix} \Delta y_{nt}] = E[\begin{pmatrix} 1 \\ \mathbf{w}_{nt-1}^T \\ \Delta\mathbf{w}_{nt}^T \\ \Delta\mathbf{z}_t \end{pmatrix} (1, \mathbf{c}_t^T, \Delta\mathbf{w}_{nt}, \Delta\mathbf{x}_{nt})] \begin{pmatrix} \bar{\mu}_n \\ \Psi_n^T \\ \beta \\ \gamma \end{pmatrix} = B \begin{pmatrix} \bar{\mu}_n \\ \Psi_n^T \\ \beta \\ \gamma \end{pmatrix}$$

   Part *iiii*) of Assumption 6 ensures the matrix B has full rank. Identification of the coefficient parameters then follows as, given B has full rank, there exists a unique set of coefficients for which this moment condition holds.

6. This argument holds for each of the N rows of $\mathbf{y}_t$.

□

To identify Σ we need some additional assumptions about the auto-correlation of the error terms. A simplest additional assumption as follows:

Assumption 7. *i*) $E[\boldsymbol{\varepsilon}_{nt}\varepsilon_{ns}] = 0 \;\; \forall n, t \neq s$
ii) $E[\boldsymbol{\upsilon}_{nt}\upsilon_{ns}] = 0 \;\; \forall n, t \neq s$
iii) $E[\boldsymbol{\upsilon}_{mt}\varepsilon_{ns}] = 0 \;\; \forall n, m, t, s$

Under this assumption, Σ is then also identified:

Proof:

1. As shown previously, the residual $\Delta\boldsymbol{\varepsilon}_{nt} + \Psi_n \boldsymbol{\upsilon}_t$ is point-identified.
2. The variance of this residual is given by $2var(\boldsymbol{\varepsilon}_{nt}) + \Psi_n \Sigma \Psi_n^T$
3. The lag-one auto-covariance is $-var(\boldsymbol{\varepsilon}_{nt})$
4. Therefore we can back out Σ from these two objects.

□

F.4 Identification of j, given $k, F^x, F^\upsilon, \Gamma$, and β

Proposition 4. *Conditional on k, F^x, F^υ, and Γ, being point identified, j is also point identified.*

This proposition demonstrates that identification of the value function, and hence continuation value, follows from identification of the primitives of the pseudo-static model. This is similar to the philosophy underlying the third stage estimation procedure discussed in the text.

Proof of this proposition involves first demonstrating that the value function is identified, which rests on the identification of k and Γ. I then show that the ex-ante value function is identified, which follows because F^υ is identified. Identification of the continuation value follows from the identification of F^x. Finally, given β, we can back j out from the definition of k. In practice, β is not identified, so must be fixed by the researcher.

Proof:

1. The Value Function can be written as:

$$W(\boldsymbol{\upsilon}, \mathbf{s}) = \max_{\mathbf{b},\mathbf{d}} \left\{ \sum_l \Gamma_l(b_l, d_l; \mathbf{s})(\upsilon_l - b_l) + \sum_a P_a(\mathbf{b}, \mathbf{d}; \mathbf{s}) k(\mathbf{s}^a) \right\}$$

   Given that both Γ and $k(\mathbf{s})$ are identified, so must be the value function. For any given $\boldsymbol{\upsilon}$ and $\mathbf{s}$ we could write down the maximand, and maximise it using numerical methods.

2. The ex-ante Value Function is given by $V^E(\mathbf{s}) = \int W(\boldsymbol{\upsilon}, \mathbf{s}) dF^\upsilon(\boldsymbol{\upsilon})$. Given that both F^υ and W are both point identified, so is the ex-ante value function.
3. The continuation value is given by $V(\mathbf{s}) = \int V^E(\mathbf{s}) dF^x(\mathbf{x})$. Given that both F^x and V^E are both point identified, so is the the continuation value.
4. The immediate pay-off function can be written as: $j(\mathbf{s}) = k(\mathbf{s}) - \beta V(\mathbf{s})$. Given identification of k and V, as well as the discount factor β, j is also point identified.

□

G Proof of Proposition 1.

In this Appendix I prove Proposition 1. The proof proceeds in three parts. In Appendix G.1 I prove the ex-ante value function result, essentially extending Arcidiacono and Miller (2011) to the continuous choice case. In Appendix G.2 I demonstrate that maximised expected payoffs $(\pi(\mathbf{b}, \mathbf{d}|\mathbf{s}^g, \mathbf{s}^0))$

takes this convenient analytic form in the case when reservation prices do not bind. In Appendix G.3 I extend this result to incorporate binding reservation prices. The only effect this has is that we must make a small adjustment when bids are observed *at* the reservation price.

The proposition to be proven, as stated in the main text, is as follows:

Proposition 1. *The ex-ante Value Function can be expressed as:*

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

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

$$\pi(\mathbf{b}, \mathbf{d}|\mathbf{s}_i^g, \mathbf{s}_0) = \sum_l \lambda_i \frac{\Gamma_l(b_{il}, d_{il}; \mathbf{s}_0)^2}{\nabla_b \Gamma_l(b_{il}, d_{il}; \mathbf{s}_0)} - \sum_{m \neq l} \Gamma_l(b_{il}, d_{il}; \mathbf{s}_0) \mathbf{z}_l^{gT} \Psi_i \mathbf{z}_m^g \Gamma_m(b_{im}, d_{im}; \mathbf{s}_0) + \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g$$

G.1 Expected Payoff, Given the State

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

$$E_{\boldsymbol{\upsilon}_t}[W(\boldsymbol{\upsilon}_t, \mathbf{s})|\mathbf{s}] = \frac{E_{\mathbf{b}_t}[q_t(\mathbf{s})\pi(\mathbf{b}_t|\mathbf{s})]}{E[q_t(\mathbf{s})]} \tag{10}$$

Where $q_t(\mathbf{s}) = f_{\mathbf{s}_t}(\mathbf{s}|\mathbb{O}_T)$ gives the posterior density, that the unobserved state is $\mathbf{s}$ at time t.

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

Proof:

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

$$\begin{aligned}
f_{\mathbf{b}_t}(\mathbf{b}|\mathbf{s}) &= \frac{f_{\mathbf{b}_t,\mathbf{s}_t}(\mathbf{b}, \mathbf{s})}{f_{\mathbf{s}_t}(\mathbf{s})} && \text{Bayes' rule} \\
&= \frac{E_{\mathbf{b}_t,\mathbf{s}_t}[\delta((\mathbf{b}_t, \mathbf{s}_t) - (\mathbf{b}, \mathbf{s}))]}{E_{\mathbf{s}_t}[\delta(\mathbf{s}_t - \mathbf{s})]} = \frac{E_{\mathbf{b}_t,\mathbf{s}_t}[\delta(\mathbf{b}_t - \mathbf{b})\delta(\mathbf{s}_t - \mathbf{s})]}{E_{\mathbf{s}_t}[\delta(\mathbf{s}_t - \mathbf{s})]} && \text{Definition of } \delta \\
&= \frac{E_{\mathbb{O}_T}[E_{\mathbf{b}_t,\mathbf{s}_t}[\delta(\mathbf{b}_t - \mathbf{b})\delta(\mathbf{s}_t - \mathbf{s})|\mathbb{O}_T]]}{E_{\mathbb{O}_T}[E_{\mathbf{s}_t}[\delta(\mathbf{s}_t - \mathbf{s})|\mathbb{O}_T]]} && \text{Law of Iterated Expectations} \\
&= \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t - \mathbf{b})E_{\mathbf{s}_t}[\delta(\mathbf{s}_t - \mathbf{s})|\mathbb{O}_T]]}{E_{\mathbb{O}_T}[E_{\mathbf{s}_t}[\delta(\mathbf{s}_t - \mathbf{s})|\mathbb{O}_T]]} && \text{as } \mathbf{b}_t \text{ is part of } \mathbb{O}_T \\
&= \frac{E_{\mathbb{O}_T}[\delta(\mathbf{b}_t - \mathbf{b})q_t(\mathbf{s})]}{E_{\mathbb{O}_T}[q_t(\mathbf{s})]} && \text{Definition of } q
\end{aligned} \tag{11}$$

2. Next recognise that we can write $E_{\boldsymbol{\upsilon}_t}[W(\boldsymbol{\upsilon}_t, \mathbf{s})|\mathbf{s}] = E_{\boldsymbol{\upsilon}_t}[\pi(\mathbf{b}(\boldsymbol{\upsilon}_t; \mathbf{s}), \mathbf{s})|\mathbf{s}]$ as $\mathbf{b}$ is set to

maximise the period payoff, given $\boldsymbol{v}_t$ and $\mathbf{s}$. Here, π is just some known function.

3. Applying a change of variables (the law of the unconscious statistician) ensures this equals $E_{\mathbf{b}_t}[\pi(\mathbf{b}_t, \mathbf{s})|\mathbf{s}]$. This requires that the mapping $\mathbf{b}_t = \mathbf{b}(\boldsymbol{v}_t; \mathbf{s})$ is monotonic (has a positive definite jacobian). This result is proven in Altmann (2022), and replicated in Appendix D.6.

4. Applying the result from step 1.:

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

5. Recognise that the denominator is not a function of the random variable $\mathbf{b}$, so we can pull it out of the integral. Then, move $\pi(\mathbf{b}, \mathbf{s})$ into the expectation for:

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

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

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

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

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

□

G.2 Maximised Payoff

The next part involves essentially applying Proposition 4 of Altmann (2022) to the quadratic pseudo-static pay-off. For notational simplicity I again drop the dependence on i and $\mathbf{s}_0$. I want to prove that when reservation prices do not bind, so that we can ignore the entry decision, we can write:

$$W(\boldsymbol{v}, \mathbf{s}) = \sum_l \lambda \frac{\Gamma_l(b_l(\boldsymbol{v}, \mathbf{s}))^2}{\nabla_b \Gamma_l(b_l(\boldsymbol{v}, \mathbf{s}))} - \sum_{m \neq l} \Gamma_l(b_l(\boldsymbol{v}, \mathbf{s})) \mathbf{z}_l^{gT} \Psi \mathbf{z}_m^g \Gamma_m(b_m(\boldsymbol{v}, \mathbf{s})) + \mathbf{s}^{gT} \Psi \mathbf{s}^g \tag{12}$$

Proof: 1. Writing $\mathbf{b}^* = \mathbf{b}(\boldsymbol{v}, \mathbf{s})$, and as discussed in section **??** the maximand and first order conditions are given by:

$$W(\boldsymbol{\upsilon},\mathbf{s})=\sum_l \Gamma_l(b_l)(\upsilon_l-\lambda b_l+\Phi\mathbf{z}_l^h+2\mathbf{z}_l^{gT}\Psi\mathbf{s}^g+\mathbf{z}_l^{gT}\Psi\mathbf{z}_l^g+\sum_{m\neq l}\Gamma_m(b_m)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)+\Phi\mathbf{s}^h+\mathbf{s}^{gT}\Psi\mathbf{s}^{gT}$$

$$0=\nabla_l\Gamma_l(b_l^*)(\upsilon_l-\lambda b_l^*+\Phi\mathbf{z}_l^h+\mathbf{z}_l^{gT}\Psi\mathbf{z}_l^g+2\mathbf{z}_l^{gT}\Psi\mathbf{s}^g+2\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)-\Gamma_l(b_l^*)\lambda \tag{13}$$

2. Divide the first order conditions by $\nabla_l\Gamma_l(b_l^*)$ (which by assumption is strictly positive), then multiply them by $\Gamma_l(b_l^*)$ and re-arrange for:

$$\begin{aligned}\Gamma_l(b_l^*)(\upsilon_l-\lambda b_l^*+\Phi\mathbf{z}_l^h+\mathbf{z}_l^{gT}\Psi\mathbf{z}_l^g+2\mathbf{z}_l^{gT}\Psi\mathbf{s}^g+\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)\\ =\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)}-\Gamma_l(b_l^*)\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)\end{aligned} \tag{14}$$

3. Summing this over l, we insert it back into the maximand for:

$$W(\boldsymbol{\upsilon},\mathbf{s})=\sum_l[\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)}-\Gamma_l(b_l^*)\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)]+\Phi\mathbf{s}^h+\mathbf{s}^{gT}\Psi\mathbf{s}^{gT}$$

4. $\mathbf{s}^h$ is not identified - it does not appear in the FOCs, so does not affect bidding behaviour. Since the $\Phi\mathbf{s}^h$ term is a level shift, without loss of generality, we can drop this term from the equation, essentially making the normalisation that $\mathbf{s}^h$ drops to zero at the end of each period (but only after food banks receive their pay-offs).

□

G.3 Maximised Payoff, given Reservation Prices

I now extent Proposition 1 to account for endogenous entry and binding reservation prices. As above, this essentially just applies the results from Altmann (2022) to the quadratic payoff case, then substitutes this in to the result proven in Appendix G.1 above. The proposition to be proven is as follows:

Proposition 1′. *The ex-ante Value Function can be expressed as:*

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

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

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

This function π is essentially the same as in Proposition 1, except that when summing over l we treat bids differently depending on whether they are at the reserve price or not. $E[\upsilon_l|b_{il} = R_l, b_{i,-l}]$ is evaluated using the bounds detailed in Appendix D and the formula for the first moment of the (doubly) truncated normal distribution.

Proof of this proposition proceeds in two parts. First, I show that the value function can be written in a form similar to the expression proven in Appendix G.2 - unlike in that Appendix, I cannot totally eliminate $\boldsymbol{\upsilon}$ from the expression. I then show that when we take an expectation over $\boldsymbol{\upsilon}$ we can still express this expectation as an integral over $(\mathbf{b},\mathbf{d})$ instead of $\boldsymbol{\upsilon}$. In Appendix G.1 I did this in step 3 of the proof, applying the law of the unconscious statistician. In this case I must first apply the law of iterated expectation to eliminate the $\boldsymbol{\upsilon}$ term left over in the value function. The remainder of the proof from Appendix G.1 can be applied.

G.3.1 Value Function

As before, write $(b_l^*, d_l^*) = (b_l(\boldsymbol{\upsilon},\mathbf{s}), d_l^*(\boldsymbol{\upsilon},\mathbf{s}))$ for the optimal bidding and entry decisions. I now show that the value function can be written as:

$$W(\boldsymbol{\upsilon},\mathbf{s}) = \Phi\mathbf{s}^h + \mathbf{s}^{gT}\Psi\mathbf{s}^{gT} + \sum_l \mathbb{I}[b_l^* > R_l]\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)} - \Gamma_l(b_l^*)\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)$$
$$+ \mathbb{I}[b_l^* = R_l]\Gamma_l(R_l,1)(\upsilon_l - \lambda R_l + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_m^g])$$

Proof:

1. Equation 12 gives the value function in parametric form. Recognise, since $\mathbb{I}[b_l^* > R_l] + \mathbb{I}[b_l^* = R_l] + \mathbb{I}[d_l^* = 0] = 1$ we can write:

$$W(\boldsymbol{\upsilon},\mathbf{s}) - \Phi\mathbf{s}^h - \mathbf{s}^{gT}\Psi\mathbf{s}^{gT} =$$
$$\sum_l \begin{cases} \mathbb{I}[b_l^* > R_l]\Gamma_l(b_l^*,d_l^*)(\upsilon_l - \lambda b_l + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_m^g]) \\ +\mathbb{I}[b_l^* = R_l]\Gamma_l(b_l^*,d_l^*)(\upsilon_l - \lambda b_l + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_m^g]) \\ +\mathbb{I}[d_l^* = 0]\Gamma_l(b_l^*,d_l^*)(\upsilon_l - \lambda b_l + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_m^g]) \end{cases} \tag{15}$$

2. By definition $\mathbb{I}[d_l^* = 0]\Gamma_l(b_l^*,d_l^*) = 0$, so the final row of equation 15 equals zero.

3. Next, since $\mathbb{I}[b_l^* = R_l]\Gamma_l(b_l^*, d_l^*) = \Gamma_l(R_l, 1)$, the second row of equation 15 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. Consider the first row of equation 15. As in Appendix D the FOCs for bid l, subject to the bid being above the reservation price and given $d_l = 1$, is given by:

$$0 = \nabla_l \Gamma_l(b_l^*, 1)(\upsilon_l - \lambda b_l^* + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m \neq l} \Gamma_m(b_m^*)\mathbf{z}_m^g]) - \Gamma_l(b_l^*)\lambda + \Lambda_l^*$$

Divide the FOCs by $\nabla_l \Gamma_l(b_l^*)$, multiply them by $\Gamma_l(b_l^*)$m and re-arrange for:

$$\begin{aligned}\Gamma_l(b_l^*)(\upsilon_l - \lambda b_l^* + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi \mathbf{z}_l^g + 2\mathbf{z}_l^{gT}\Psi \mathbf{s}^g + \sum_{m \neq l} \Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi \mathbf{z}_m^g) \\ = \lambda \frac{\Gamma_l(b_l^*)^2}{\nabla_l \Gamma_l(b_l^*)} - \Gamma_l(b_l^*)\sum_{m \neq l} \Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi \mathbf{z}_m^g) + \Lambda_l^* \quad (16)\end{aligned}$$

When $b_l^* > R_l$ the constraint does not bind, so $\mathbb{I}[b_l^* > R_l]\Lambda_l^* = 0$. Multiply both sides by $\mathbb{I}[b_l^* > R_l]$ to show that the first row on the right hand side of equation 15 equals $= \lambda \frac{\Gamma_l(b_l^*)^2}{\nabla_l \Gamma_l(b_l^*)} - \Gamma_l(b_l^*)\sum_{m \neq l} \Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi \mathbf{z}_m^g)$

□

G.3.2 ex-ante Value Function

I now show that $E_{\boldsymbol{\upsilon}}[W(\boldsymbol{\upsilon}, \mathbf{s})|\mathbf{s}]$ can be written as $E_{(\mathbf{b},\mathbf{d})}[\pi(\mathbf{b}, \mathbf{d}, \mathbf{s})|\mathbf{s}]$ for some function π. This ensures the proposition proven in Appendix G.1 applies in the binding reservation price case.

Proof: 1. Applying the result proven above, the ex-ante value function is given by:

$$\begin{aligned}E_{\boldsymbol{\upsilon}}[W(\boldsymbol{\upsilon}, \mathbf{s})] = \Phi \mathbf{s}^h + \mathbf{s}^{gT}\Psi \mathbf{s}^{gT} + \sum_l \\ E_{\boldsymbol{\upsilon}}[\mathbb{I}[b_l^* > R_l]\lambda \frac{\Gamma_l(b_l^*)^2}{\nabla_l \Gamma_l(b_l^*)} - \Gamma_l(b_l^*)\sum_{m \neq l} \Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi \mathbf{z}_m^g)] \\ + E_{\boldsymbol{\upsilon}}[\mathbb{I}[b_l^* = R_l]\Gamma_l(R_l, 1)(\upsilon_l - \lambda R_l + \Phi \mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m \neq l} \Gamma_m(b_m^*, d_m^*)\mathbf{z}_m^g)]]\end{aligned}$$

2. Next, we want to apply the the law of the unconscious statistician to the middle line, as we did in Appendix G.2. However, Even though this line conditions on $b_l > R_l$, so that we can apply this law for the bid on lot l, we cannot apply it to the whole term due to the potentially constrained bids that appear in the combinatorial term $\Gamma_l(b_l^*, d_l^*)\sum_{m \neq l} \mathbf{z}_l^{gT}\Psi \mathbf{z}_m^g \Gamma_m(b_m^*, d_m^*)$. Instead, apply Lemma G.1 proven below for the

requisite result:[63]

$$E_{\boldsymbol{\upsilon}}[\mathbb{I}[b_l^* > R_l]\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)} - \Gamma_l(b_l^*)\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)]$$
$$= E_{\mathbf{b}}[\mathbb{I}[b_l > R_l]\lambda\frac{\Gamma_l(b_l)^2}{\nabla_l\Gamma_l(b_l)} - \Gamma_l(b_l)\sum_{m\neq l}\Gamma_m(b_m)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)]$$

3. Now we focus on the bottom line. Apply the law of iterated expectations for:

$$E_{\boldsymbol{\upsilon}}[\mathbb{I}_{[b_l^*=R_l]}\Gamma_l(R_l,1)(\upsilon_l - \lambda R_l + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*, d_m^*)\mathbf{z}_m^g)]]$$
$$= E_{\mathbf{b}}[E_{\boldsymbol{\upsilon}}[\mathbb{I}_{[b_l^*=R_l]}\Gamma_l(R_l,1)(\upsilon_l - \lambda R_l + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*, d_m^*)\mathbf{z}_m^g)]|\mathbf{b}]]$$
$$= E_{\mathbf{b}}[\mathbb{I}_{[b_l=R_l]}E_{\boldsymbol{\upsilon}}[\Gamma_l(R_l,1)(\upsilon_l - \lambda R_l + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*, d_m^*)\mathbf{z}_m^g)]|b_l = R_l, \mathbf{b}_{-l}]]$$
$$= E_{\mathbf{b}}[\mathbb{I}_{[b_l=R_l]}\Gamma_l(R_l,1)(E_{\boldsymbol{\upsilon}}[\upsilon_l|b_l = R_l, \mathbf{b}_{-l}] - \lambda R_l + \Phi\mathbf{z}_l^h + \mathbf{z}_l^{gT}\Psi[\mathbf{z}_l^g + 2\mathbf{s}^g + \sum_{m\neq l}\Gamma_m(b_m^*, d_m^*)\mathbf{z}_m^g)]]$$

4. Substituting these two expressions back into the ex-ante value function yields the desired result.

□

Lemma G.1.

$$E_{\boldsymbol{\upsilon}}[\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)} - \Gamma_l(b_l^*)\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l]$$
$$= E_{\mathbf{b}}[\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)} - \Gamma_l(b_l^*)\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l > R_l]$$

Proof: 1. First we split the object into four components, each dealt with separately. This is performed using the linearity of the expectation operator and the law of iterated expectations. Write:

$$E_{\boldsymbol{\upsilon}}[\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)} - \Gamma_l(b_l^*)\sum_{m\neq l}\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l] =$$
$$= E_{\boldsymbol{\upsilon}}[\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)}|b_l^* > R_l]$$
$$-\sum_{m\neq l}\begin{cases} P[b_m^* > R_m]E_{\boldsymbol{\upsilon}}[\Gamma_l(b_l^*)\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l, b_m^* > R_m] \\ P[b_m^* = R_m]E_{\boldsymbol{\upsilon}}[\Gamma_l(b_l^*)\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l, b_m^* = R_m] \\ P[d_m^* = 0]E_{\boldsymbol{\upsilon}}[\Gamma_l(b_l^*)\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l, d_m^* = 0] \end{cases} \tag{17}$$

[63] Recognising that for random variables X and Y, $E[Y\mathbb{I}[X = x]] = P(X = x)E[Y|X = x]$.

Where the final equality holds by definition. We have essentially just applied the law of iterated expectation.

2. We can apply the law of the unconscious statistician for $E_{\boldsymbol{v}}[\lambda\frac{\Gamma_l(b_l^*)^2}{\nabla_l\Gamma_l(b_l^*)}|b_l^* > R_l] = E_{b_l}[\lambda\frac{\Gamma_l(b_l)^2}{\nabla_l\Gamma_l(b_l)}|b_l > R_l]$ as we condition on $b_l > R_l$, and we know that the function $b_l(\boldsymbol{v})$ is monotonic in this region.

3. Likewise

$$\begin{aligned}&E_{\boldsymbol{v}}[\Gamma_l(b_l^*)\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l, b_m^* > R_m]\\&\qquad= E_{b_l,b_m}[\Gamma_l(b_l)\Gamma_m(b_m)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l > R_l, b_m > R_m]\end{aligned}$$

as we condition on $b_l > R_l$ and $b_m > R_m$, and we know that the functions $b_l(\boldsymbol{v}), b_m(\boldsymbol{v})$ are monotonic in this region.

4. Considering the middle row:

$$\begin{aligned}&E_{\boldsymbol{v}}[\Gamma_l(b_l^*)\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l, b_m^* = R_m]\\&\qquad= \Gamma_m(R_m)E_{\boldsymbol{v}}[\Gamma_l(b_l^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l, b_m^* = R_m]\end{aligned}$$

Which is then just a function of b_l^*, which we know to be above the reservation price, so the law applies, for:

$$\begin{aligned}&E_{\boldsymbol{v}}[\Gamma_l(b_l^*)\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l, b_m^* = R_m]\\&\qquad= \Gamma_m(R_m)E_{b_l}[\Gamma_l(b_l)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l > R_l, b_m = R_m]\end{aligned}$$

5. Considering the final row, by the same logic as above we have

$$\begin{aligned}&E_{\boldsymbol{v}}[\Gamma_l(b_l^*)\Gamma_m(b_m^*)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l^* > R_l, d_m^* = 0]\\&\qquad= 0 = E_{b_l}[\Gamma_l(b_l)\Gamma_m(\emptyset, 0)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l > R_l, d_m = 0]\end{aligned}$$

6. Recognising that $P[b_m(\boldsymbol{v}) > R_m] = P[b_m > R_m]$, We have therefore shown that equation 17 is equal to:

$$\begin{aligned}&E_{b_l}[\lambda\frac{\Gamma_l(b_l)^2}{\nabla_l\Gamma_l(b_l)}|b_l > R_l]\\&\qquad-\sum_{m\neq l}\begin{cases}P[b_m > R_m]E_{b_l,b_m}[\Gamma_l(b_l)\Gamma_m(b_m)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l > R_l, b_m > R_m]\\P[b_m = R_m]E_{b_l,b_m}[\Gamma_l(b_l)\Gamma_m(b_m)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l > R_l, b_m = R_m]\\P[d_m = 0]E_{b_l,b_m}[\Gamma_l(b_l)\Gamma_m(b_m)\mathbf{z}_l^{gT}\Psi\mathbf{z}_m^g)|b_l > R_l, d_m = 0]\end{cases}\end{aligned}$$

7. Which, by the law of iterated expectations and linearity of the expectation operator,

is equal to

$$E_{\mathbf{b}}[\lambda \frac{\Gamma_l(b_l^*)^2}{\nabla_l \Gamma_l(b_l^*)} - \Gamma_l(b_l^*) \sum_{m \neq l} \Gamma_m(b_m^*) \mathbf{z}_l^{gT} \Psi \mathbf{z}_m^g) | b_l > R_l]$$

□

H Estimation Details

In this Appendix I give additional details of the estimation procedure outlined in Section 5. I outline my specification of priors, as well as computational details of how each step of the estimation procedure is performed. Appendix H.1 outlines details of the first estimation step, Appendix H.2 the second step, Appendix H.3 the third step, and finally Appendix H.4 details the model specification used for the Type 2 food banks.

H.1 Step 1.

In the first estimation step I estimate food banks beliefs about the probability they win a given lot given their bid. While I assume there is zero probability of ties above the reservation price, I allow for the possibility of ties at the reservation price. I begin by discussing how I conceptualise food banks' beliefs in Appendix H.1.1. I discuss ties in Appendix H.1.2. I then detail my parametrisation in Appendix H.1.3, before discussing how estimation is performed H.1.4.

H.1.1 Maximum Rival Bid

I do not explicitly parameterise bid distributions and use this to form food banks' beliefs about equilibrium win probabilities, as in Jofre-Bonet and Pesendorfer (2003) or Gentry et al. (2020). Instead I take an approach closer to that in Backus and Lewis (2016) and estimate the distribution of equilibrium winning bids. If the winning bid on auction l at time t was $\bar{b}_{lt}$, then food bank i knows they would have won the lot had they bid $b_{ilt} > \bar{b}_{lt}$. If no food bank placed a bid on lot l then food bank i knows they would have won if they had bid the reservation price. If a food bank won lot l at the reservation price, i knows they would have drawn had they bid the reservation price.

A food bank's ex-ante belief about the probability of winning given a bid is given by $P(\bar{b}_{lt} < b_{ilt} | \mathbf{s}_t)$, which requires the conditional cdf of the random variable $\bar{b}_{lt}$. However, this object is subject to censoring at the reservation price.

H.1.2 Ties

Ties are observed very rarely in the data, in around 0.02% of auctions, and all at the reservation price. Due to the continuity of bids, ties happen above the reservation price with probability zero.[64] However, because winning bids are observed more frequently at the reservation price, in around 20%

[64] In practice bids must be integer amounts, but because bids tend to range between −2000 and 4000 I treat this as continuous.

of auctions, food banks must consider the much larger chance of a tie if they bid the reservation price and no higher. Furthermore, food banks appear to recognise this, and often bid just above the reservation price. This means we get high density of winning bids just above the reservation price which food banks presumably also recognise, and so must be taken into account in the model.

The bidder wins lot l given bid b_{ilt} if $b_{ilt} > \bar{b}_{lt}$. If $b_{ilt} = \bar{b}_{lt}$ they win with probability 0.5.[65] 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). Therefore I introduce the latent random variable $\bar{b}^*_{lt}$, with cdf $G_l(b^*|\mathbf{s}_{0t})$, such that:

$$\bar{b}_{lt} = \begin{cases} \bar{b}^*_{lt} & \text{if } \ \bar{b}^*_{lt} > R_{lt} \\ R_l + \epsilon_{lt} & \text{if } \ \bar{b}^*_{lt} \in [R_l, \bar{R}_l) \\ R_l & \text{if } \ \bar{b}^*_{lt} \in [\bar{R}_l, \underline{R}_l) \\ \emptyset & \text{if } \ \bar{b}^*_{lt} \leq \underline{R}_l \end{cases}$$

Where $(\bar{R}_l, \underline{R}_l)$ are a category specific cutoff to be estimated. This is not dissimilar to cutoffs estimated in an ordered logit model, enabling me to capture the varying likelihood of winning bids at the reservation price across categories. This latent variable structure states that if the 'true', latent, winning bid $\bar{b}^*_{lt}$ is extremely low ($\leq \underline{R}_l$), then a competing food bank would win if it bid the reservation price. If it is somewhat higher $\bar{b}^*_{lt} \in [\bar{R}_l, \underline{R}_l)$ then the observed winning bid is just the reservation price - a competing food bank would draw if it bid the reservation price. The competing food bank may not value the lot enough to bid much above the reservation price, but may be willing to bid just one or two additional shares to ensure it doesn't risk a tie. Finally, if $\bar{b}^*_{lt}$ is just below the reservation price $\bar{b}^*_{lt} \in [R_l, \bar{R}_l)$ then the observed winning bid is actually just above the reservation price, where $\epsilon_{lt} \sim exponential(\alpha)$ and α is a parameter to be estimated. This means that a competing food bank must take into account the excess mass just above the reservation price - if it bids just one share above the reservation price it may lose out to equally strategic food banks.

This modelling approach is unusual, but enables the model to rationalise both the excess mass of winning bids at the reservation price, and also just above it. I assume that food banks do not internalise the probability of tieing at just one share above the reservation price (and likewise two, three, etc shares). Importantly, $\underline{R}_l$ is identified by the excess mass of winning bids at the reservation price (and how this varies across categories). $\bar{R}_l$ is identified from the excess mass just above the reservation price, and α is identified from how the excess mass diminishes as we move further from the reservation price.

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

$$P(i \text{ wins } l|b_{ilt}; \mathbf{s}_{0t}) = \Gamma_l(b_{ilt}|\mathbf{s}_{0t}) = \begin{cases} G_l(b_{ilt}|\mathbf{s}_{0t}) - f(b_{ilt}) & \text{if } \ b_{ilt} > R_{lt} \\ \frac{1}{2}G_l(\underline{R}^c|\mathbf{s}_{0t}) + \frac{1}{2}G_l(\bar{R}^c|\mathbf{s}_{0t}) & \text{if } \ b_{ilt} = R_{lt} \\ 0 & \text{otherwise} \end{cases} \tag{18}$$

[65] In other words I assume they tie with at most one other bidder, and the tie is broken with the flip of a coin. In practice, bidders believe that ties only occur with more than one bidder at the rate that ties are observed in the data, in 0.02% of auctions, which I deem as negligible.

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

H.1.3 Parameterisation

I normalise all bids by the reservation price, so that the estimated distributions can be considered the distribution of the difference between the winning bid and the reservation price. Therefore from here on, we can replace $\bar{b}^*_{lt}$ with $\bar{b}^*_{lt} - R_l$. Reservation prices are known to be -2000 for all lots except for fresh produce and Maroon lots which have $R_l = 0$.

As in Assumption 4 I assume that the distribution of $\bar{b}^*_{lt}$ is a function of $\boldsymbol{\vartheta}(\{\mathbf{s}_i\}_N)$, aggregate statistics of states only, so that it does not depend on the states of each individual food bank. In particular, I assume that it only depends on the previous 30 day supply of food from each storage type, as well as the supply of food allocated at time t from each storage type. This is intended to capture how prices vary with supply. This also ensures that food banks do not need to take into account exactly which food bank wins which combination of lots each period. Instead, i only needs to consider which lots they themselves win.

I assume that the latent random variable $\bar{b}^*_{lt}$, on lot l given common state $\mathbf{s}_{0t}$, follows a generalised extreme value distribution, with:

$$P(\bar{b}^*_{lt} \leq b|\mathbf{s}_{0t}) = exp(-t(\frac{b - \nu(\mathbf{s}_{0t})}{\zeta(\mathbf{s}_{0t})})) \quad \text{Where :} \quad t(x) = \begin{cases} (1 + \xi(\mathbf{s}_{0t})x)^{-\frac{1}{\xi(\mathbf{s}_{0t})}} & \text{if } \ \xi(\mathbf{s}_{0t}) \neq 0 \\ exp(-x) & \text{if } \ \xi(\mathbf{s}_{0t}) = 0 \end{cases}$$

The shape parameters ξ are category specific for categories with at least 500 loads, and the remainder are constrained to be equal to one another. Shape parameters are constrained to be > -1 to ensure bids are monotonic in values. This constraint does not bind.

The scale parameters ζ are also category specific. In addition, within a category if the subcategory is listed as "unspecified", "mixed" or "miscellaneous" these receive an additional fixed effect on their scale parameter. This is to allow me to capture additional variation due to uncertainty over the goods included in the lot. I constrain scale parameters to be strictly positive. Finally, I also allow the scale parameter for lot l to vary depending on whether the lot has also been auctioned in a previous period. If this is the case bidders gain information about rival bidders values for this lot, making it intuitive that the variance of rival bids is expected to decrease.

Each lot can contain up to four distinct categories, subcategories and storage types. Therefore, for both the shape and scale parameters, if the lot contains a mixture of categories, I use an average over the different categories / subcategories.

The location shifter ν varies with both lot specific covariates and the common state variable. I include subcategory fixed effects, as well as dummies for whether the lot includes free delivery, geographic restrictions, any unobserved notes about the lot contents, whether the lot is a "Maroon load" (category specific for categories with at least 50 maroon loads), which US region the lot originates in,[66] is shelf-stable, the number of distinct categories included in the lot, whether the lot

[66]Using the 8 US economic regions + Canada

has been auctioned previously, and the number of homogenous loads being auctioned simultaneously. It also varies with the log of the sum of the previous 30 day's supply of that type of food, up to $t-1$, by storage type, and the log of the sum of food of that storage type being auctioned simultaneously that day. As with the shape and scale parameters, if the lot includes multiple categories, subcategories, or storage types, I use the average location shifter.

The threshold cutoffs $\bar{R}_l$ and $\underline{R}_l$ are allowed to vary across categories, but only for categories in which there are at least 100 lots won at the reservation price. This includes Beverages, Cereal, Condiments, Fresh Produce, Meals, and Snacks. The remaining categories are grouped together. The exponential parameter α is constrained to be positive.

H.1.4 Computation

In the first stage I estimate 268 parameters, distributed over $26,000$ auctions, dropping the first 60 days to enable construction of the sum of the previous 30 days' supplies.

Parameters are initially estimated using maximum likelihood, using the implied distribution of $\bar{b}_{lt}$. I set initial shape parameters to 0, and initial scale parameters to the observed standard deviation of winning bids. I set all the location shifters to zero.

Having maximised the likelihood function, I then draw samples from the posterior distribution using the Metropolis Hastings algorithm. I use the inverse hessian from the maximum likelihood procedure for my proposal variance, and also adaptively tune this variance using the procedure of Atchadé and Rosenthal (2005) to ensure that on average 23.4% of proposed draws are accepted.

H.2 Step 2.

In the second estimation step I estimate the distribution of lot specific values, the distribution of net local donations, and the pseudo-static pay-off function. In Appendix H.2.1 I restate my parametrisation as in the main text. In Appendix H.2.2 I set out my assumptions on prior distributions. In Appendix H.2.3 I discuss the data augmentation algorithm, and in Appendix H.2.4 I discuss the gibbs sampling algorithm for drawing parameters from their conditional posteriors. In Appendix H.2.5 I discuss additional computational details. Note that this appendix makes heavy use of the results on the inverse bid system given in Appendix D.

H.2.1 Parametrisation

As stated in Section 5, the main parametric assumptions are as follows:

$$k(\mathbf{s}_i) = \Phi \mathbf{s}_i^h + \mathbf{s}_i^{gT} \Psi_i \mathbf{s}_i^g \qquad \upsilon_{ilt} \sim N(\alpha_i \text{distance}_{ilt}, \sigma_l^2) \qquad \mathbf{x}_{it} \sim N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$$

The 164 subcategory parameters of Φ are constant across food banks, and constrained to be positive. This is to allow us to interpret these subcategory weights as the benefit food banks receive from holding the various subcategories to give out to the clients. This constraint doe not bind. Ψ_i is a symmetric matrix that is allowed to vary across food banks, with 15 unique elements for each food bank (5 storage types and 10 interaction terms). I do not impose other constraints on this matrix, given that with probability 1 each draw of the matrix will have full rank.

α_i gives food bank i's cost, in shares, of transporting a lot an additional kilometre. These parameters are constrained to be negative (costs), and this does not bind. I do not allow this cost to vary with either the size of the lot, or the type of food. I allow the lot specific variance to vary depending on the combination of goods auctioned together in the lot. In particular, 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 variance parameter. I also include an 'other' variance parameter, which covers the remaining 5.5% of combinations. I parametrise the lot specific variances in this manner as it makes the problem of sampling from their posterior significantly easier. These parameters are constrained to be positive, and are not allowed to vary across food banks. The λ_i parameters, which capture the opportunity cost to food bank i of spending a share essentially capture variation in the variance of bids across food banks. These parameters are also constrained to be positive, and the parameter for the median consuming Type 1 food bank is constrained to 1.

Finally, for the distribution of net local donations I impose that Σ_i is a diagonal matrix. Informed by the analysis from Appendix C I impose that diagonal entries are strictly above $0.01Var[\mathbf{w}_{it}^T\mathbf{z}_t]$ and below $100Var[\mathbf{w}_{it}^T\mathbf{z}_t]$. For all parameters, if not otherwise constrained I impose an upper limit of e^{50} and a lower limit of $-e^{50}$.

H.2.2 Priors and Hierarchical Distributions

Write $\boldsymbol{\psi}_i$ as the vector of unique elements of the matrix Ψ_i. I assume these come from the hierarchical distribution, such that

$$\psi_i \sim N(\boldsymbol{\psi}, \Sigma^{\psi})$$

The hierarchical framework reduces the posterior variance of estimated parameters at a cost of bias, as estimated parameters are drawn together. Observations with a lot of identifying variation place little weight on the hierarchical parameters, whereas observations with little identifying variation place more weight on hierarchical parameters. Any bias caused by this framework causes parameters to be drawn together, so that my estimates will be biased in favour of the Old System rather than the Choice System. I assume weak inverse-Wishart priors for the hierarchical parameters $(\boldsymbol{\psi}, \Sigma^{\psi})$. The prior mean for $\boldsymbol{\psi}$ is -1 for diagonals and zero for off-diagonals of Ψ. The prior mean for Σ^{ψ} is set to be arbitrarily small. The two shape parameters are each set to 2.

I assume independent normal priors for Φ and α_i with means $\frac{1}{40}$ and 1 respectively, and prior variance 10000 to reflect my prior ignorance over these parameters. I assume weak inverse-gamma priors for the lot specific variance, with prior-mean set to the observed variance of bids, and shape parameter set to 2. For the lambda parameters I assume that λ_i^2 takes a prior gamma distribution with shape and rate parameters $T_{\lambda}^0 = 100$. This ensures λ_i^2 has prior mean of 1, and confidence about this prior mean equal to approximately 1/100th the weight placed on the data.

I assume normal-inverse-gamma priors for $\boldsymbol{\mu}_i$ and Σ_i, with prior means $\boldsymbol{\mu}_i^0 = \frac{1}{T}\sum_{t=1}^{T}\mathbf{w}_{it}^T\mathbf{z}_{it}$ and $\Sigma_i^0 = \frac{2}{T-1}\sum_{t=1}^{T}(\mathbf{w}_{it}^T\mathbf{z}_{it} - \frac{1}{T}\sum_{t=1}^{T}\mathbf{w}_{it}^T\mathbf{z}_{it})^2$ respectively. Given that in my estimation sample I have $T = 1075$, I set prior 'shape' parameters for these distributions to 107, essentially meaning that I place ten times as much weight on the data as I do on my priors.

H.2.3 Data Augmentation step

Given parameters and unobserved states I form the inverse bid system as in Appendix D. For observation ilt such that $b_{ilt} > R_{lt}$ the inverse bid system gives me a conditional observation of υ_{ilt}. For observations of bids at or below the reservation price I am only able to bound υ_{ilt}. I augment my data by drawing these conditional observations from their conditional posterior, the truncated normal distribution, using the sampling procedure of Botev (2017). I then revert these inverse bids into 'observations' y_{ilt}, essentially observations of bids from below the reservation price.

To sample the unobserved states from their posteriors I run a standard Kalman filter using the current draw of parameters and the current draw of censored observation. I begin the filter on day 61, as I do not have estimates of beliefs from before this point. I set the initial state to 0, essentially normalising the first set of stocks, with initial variance of zero.

I then run the Carter-Kohn algorithm (Carter and Kohn, 1994) to backwards sample the unobserved states from their conditional posterior. I only run it backwards until day 101, essentially discarding an extra 40 days of the filter. I do this to reduce the reliance on the initial state assumption. This is because even though the initial state is not identified, if there is significant bayesian shrinkage due to the hierarchical model, the initial state actually may be identified.

Finally, I also take into account the observed change in the supply of fresh produce that occurs on day 553 in my sample. Thereafter fresh produce stops being allocated through the Choice System and is instead allocated to food banks outside the system. Each food bank has two separate mean local donations for fresh food, one for before this period and one after this period. I fix the parameter for after period 553 to 0, so that on average food banks give out as much fresh food as they receive in net local donations. Anything else would lead stocks either to trend upwards or downwards indefinitely. Therefore I only estimate the expected net donation for fresh food using sampled states from before the break.

H.2.4 Gibbs Sample step

Given a sample of states $\{\mathbf{s}_{it}\}_{t\in\{1\ldots T\}}$ I back out a sample of net local donations by writing the transition equation as a function of $\mathbf{x}_{it}$. Write x_{itm} for the mth element of $\mathbf{x}_{it}$. The conditional posterior distribution is then normal-inverse-gamma, and given as:

$$(\mu_{im}, \Sigma_{imm})|\,\{x_{itm}\}_{t\in\{1\ldots T\}} \sim N - IG(A, B, C, D)$$

$$A = \frac{T^0\mu^0_{im} + T\bar{x}_{im}}{T^0 + T} \qquad B = T^0 + T \qquad C = \frac{T^0 + T}{2}$$

$$D = \frac{T^0\Sigma^0_{imm}}{2} + \frac{1}{2}\sum_{t=1}^{T}(x_{itm} - \bar{x}_{im})^2 + \frac{TT^0}{T + T^0}\frac{(\bar{x}_{im} - \mu^0_{im})^2}{2})$$

$$\bar{x}_{im} = \frac{1}{T}\sum_{t=1}^{T} x_{itm}$$

I then move on to the parameters of k and F^v. I focus on Ψ_i first, rewriting the observation equation (making use of both sampled states and censored observations) as:

$$Y_{ilt} = \lambda_i y_{ilt} - \Phi \mathbf{z}^h_{tl} - \alpha_i dist_{ilt} = \mathbf{z}^{gT}_{tl} \Psi_i(\mathbf{z}^g_{tl} + 2\mathbf{s}^g_{it} + 2\sum_{m \neq l} \Gamma_{im}(b_{itm}\mathbf{z}^g_{tm})) + \upsilon_{itl} - \alpha_i dist_{ilt} = \mathbf{X}_{itl}\boldsymbol{\psi}_i + \varepsilon_{itl}$$

Stacking Y_{itl} and $\mathbf{X}_{itl}$ over tl, the conditional posterior distribution of ψ_i is then multivariate normal, and given as:

$$\boldsymbol{\psi}_i|(\mathbf{Y}_i, \mathbf{X}_i), (\psi, \Sigma^{\psi}) \sim N(M, V)$$
$$M = V^{-1}(\Sigma^{-1\psi}\psi + \mathbf{X}_i^T\mathbf{Y}_i) \qquad \& \qquad V = (\Sigma^{-1\psi} + \mathbf{X}_i^T\mathbf{X}_i)^{-1}$$

The hierarchical parameters (ψ, Σ^{ψ}), which I sample after sampling Ψ_i, have normal-inverse Wishart distribution, with conditional posterior:

$$(\boldsymbol{\psi}, \Sigma^{\psi})|\,\{\boldsymbol{\psi}_i\}_{i \in \{1 \ldots N\}} \sim N - IW(A, B, C, D)$$
$$A = \frac{N\bar{\boldsymbol{\psi}} + \boldsymbol{\psi}^0}{N+1} \qquad B = N+1 \qquad C = N+1$$
$$D = \Sigma^{0\psi,-1} + \sum_i (\boldsymbol{\psi}_i - \bar{\boldsymbol{\psi}})(\boldsymbol{\psi}_i - \bar{\boldsymbol{\psi}})^T + \frac{N}{N+1}(\bar{\boldsymbol{\psi}} - \boldsymbol{\psi}^0)(\bar{\boldsymbol{\psi}} - \boldsymbol{\psi}^0)^T)$$
$$\bar{\boldsymbol{\psi}} = \frac{1}{N}\sum_i \boldsymbol{\psi}_i$$

I jointly sample the distance parameters α_i and subcategory weights Φ using standard bayesian regression, given normal priors, Ψ_i, sampled states and censored observations. I sample lot-specific variances just as in bayesian regression, given regression coefficients and $\{\lambda_i\}_{i \in \{1 \ldots N\}}$.

Finally, rewriting the observation equation as: $\lambda_i y_{ilt} = \mathbf{Z}_{itl}\delta + \varepsilon_{ilt}$, the conditional posterior pdf of λ_i is proportional to:

$$f(\lambda_i|\mathbf{Z}_i, \delta, \sigma) \qquad \propto \qquad (\frac{\lambda_i^{LT+2(T^0_\lambda - 1)}}{\prod_{lt}\sigma_l\sqrt{2\pi}})exp(-\frac{1}{2}(\lambda_i^2 \sum_{tl} \frac{y_{ilt}}{\sigma_{lt}}^2 - \lambda_i \sum_{tl} 2\frac{y_{ilt}\mathbf{Z}_{itl}\delta}{\sigma^2_{lt}} + 2T^0_\lambda))$$

I draw samples from this posterior distribution using a metropolis hastings step. I then divide the λ_is by that of the median food bank, ensuring the relevant normalisation.

H.2.5 Computation

I focus on data from only the highest 25 bids placed each day by each food bank. Even type 1 food banks rarely place more than between 5 and 10 bids each day - the 90th percentile food bank only bids on 4 lots each day. However on 50% of days with at least one auction there are more than 25 unique lots being auctioned simultaneously. In principle by ignoring that food banks also choose not to bid on any more than the first 25 lots I may bias my results towards food banks being willing to bid on a higher proportion of auctioned lots than in fact. However, the degree of this bias is unlikely to be large, since I am already taking into account that food banks only bid on maybe the first 10 lots, then choose not to bid on the next 15 lots. Furthermore, to show robustness to this assumption in Appendix J.2.4 I present results from considering 50 unique auctions each day. This assumption is useful in ensuring results converge relatively more quickly, since the higher the degree of censoring, the slower results are expected to converge.

The order of my data augmentation and Gibbs Sampling procedure is as it was presented in the main text. Every tenth iteration I draw a new sample of beliefs using five repetitions of Metropolis Hastings. At the very beginning of the procedure I run the data augmentation step 30 times without running the gibbs sampling step. This is to reduce the sensitivity to the initial draw of augmented data, in which it is assumed that states do not vary at all.

I run the full procedure for 300,000 iterations, and burnout the first 200,000 draws. I run 4 independent chains. For parameters with informative priors, initial points are drawn from the prior distribution. For parameters with diffuse priors I sample uniformly between 0 and 2 $\times$ the prior mean. I uniformly sample 250 points from each of the chains, so that I keep 1,000 parameter draws in total. I then use these parameters in evaluating both the third stage, and the Choice System simulations. I estimate around 1780 parameters across 1.1 million observations observations, 0.95 of which are censored (i.e. a bid is not placed).

H.3 Step 3.

In the third estimation 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.[67]

In Appendix H.3.1 I describe how I form the posterior probabilities that $\mathbf{s}_t = \mathbf{s}$, given by $q_t(\mathbf{s})$. In Appendix H.3.2 I outline how I evaluate the expression for the maximised expected pay-off. In Appendix H.3.3 I discuss how I pool information across food banks in this estimation stage. In Appendix H.3.4 I justify and detail the polynomial approximation used for the ex-ante value function. Finally, in Appendix H.3.5 I describe how I take the expectation of the ex-ante value function over states, yielding the continuation value, before backing out the combination flow payoff.

[67] One thing to note: The marginal welfare from consuming a lot with subcategory composition $\mathbf{z}_{tl}^h$ is just $\Phi\mathbf{z}_{tl}^h$, and does not depend on $\mathbf{s}_{it}^h$. I do not need to worry about the fact that Φ is a 'pseudo-static' object, not a present discounted sum of expected future flow payoffs from winning $\mathbf{z}_{tl}^h$. The reasons for this are simple - when flow payoffs are affine in $\mathbf{s}_{it}^h$, so is the pseudo-static payoff. Furthermore, bidding behaviour does not depend on $\mathbf{s}_{it}^h$, and so future bidding behaviour does not depend on $\mathbf{z}_{tl}^h$. Suppose the flow payoff is given by $\tilde{\Phi}\mathbf{s}_{it}^h + j(\mathbf{s}_{it}^g)$, and allow that $\mathbf{s}_{it}^h = \boldsymbol{\delta}\mathbf{s}_{it-1}^h + \mathbf{w}_{it-1}^T\mathbf{z}_{t-1}^h + \mathbf{x}_{it}^h$. In this case, the marginal welfare from winning a lot with subcategory composition $\mathbf{z}_{tl}^h$ (focusing only on the subcategory component) is just $\tilde{\Phi}\sum_{s=0}^{\infty}\beta^s\boldsymbol{\delta}^s\mathbf{z}_{tl}^h = \tilde{\Phi}(I-\beta\boldsymbol{\delta})^{-1}\mathbf{z}_{tl}^h$. Therefore $\Phi = \tilde{\Phi}(I-\beta\boldsymbol{\delta})^{-1}$.

H.3.1 Posterior Probabilities

Because states are continuous I must evaluate the continuation value over a finite set of states. For each food bank I form a 20^5 dimensional grid of states, so that each dimension of stocks is split into 20 evenly spaced points. For the minimum and maximum points I take the 2.5 and 97.5 percentiles of all their sampled states.

For each of these states I form the posterior probability density that at any given time this was the true state of their stocks, using my $1,000$ draws of states for each time period. I use an independent normal kernel, with Silverman's rule of thumb to calculate bandwidth h:

$$\hat{q}_{it}(\mathbf{s}) = \hat{p}(\mathbf{s}_{it} = \mathbf{s}_i|\text{data}) = \prod_{m=1}^{5} \frac{1}{1000} \sum_{r=1}^{1000} \frac{1}{h_m} \phi(\frac{s^r_{itm} - s_{im}}{h_m})$$

H.3.2 Maximised Payoff

I evaluate the maximised pay-off at each time period $\pi(\mathbf{b}_t, \mathbf{d}_t|\mathbf{s})$ using the reservation price adjusted formula discussed in Appendix G.3. This expression is evaluated once for each parameter draw, using the sample counterparts to the expectation operators given.

In principle I ought to take into account sampling variation in these finite sample expectations. However, given the large number of time periods we expect fairly little variation. One possibility is to use a bootstrap procedure when evaluating these averages to ensure that we introduce sampling variation alongside the variation in parameters from our draws. The difficulty is that this does not account for the correlations between sampled parameters and the sample expectations, so will overestimate posterior variances. This procedure is performed in Appendix J.3.3.

H.3.3 Information Pooling

I also pool information across food banks, in order to give me additional observations when evaluating the expectation. Similar to the bayesian hierarchical model, this is expected to bias my estimates in favour of the Old System, pushing food banks' flow pay-offs closer together. However, I use two adjustments to minimise this bias.

When constructing the ex-ante value function for food bank i given parameter draw θ_i, I find the probability density (using the same independent normal kernel as used above) that food bank j draws these parameters given their posterior distribution. The parameters I compare are the estimates for Ψ_i, α_i, and λ_i. I do not need to compare parameters for the net donation process, since the ex-ante value function is evaluated conditional on the state anyway. This density yields a weight for food bank j. I normalise the weights so they sum to 1. Then, when summing posterior probabilities across t I also multiply the probabilities by the associated food bank weight.

Finally, I also use a first-order adjustment to account for the fact that different food banks bids, and hence maximised payoffs, are determined by different parameters. Write:

$$\pi_i(\mathbf{b}_t, \mathbf{d}_t | \mathbf{s}_{it}; \theta_i) \approx \pi_j(\mathbf{b}_t, \mathbf{d}_t | \mathbf{s}_{it}; \theta_i) + \nabla_{\Psi,\alpha,\lambda} \pi_j(\mathbf{b}_t, \mathbf{d}_t | \mathbf{s}_{it}; \theta_i) \begin{pmatrix} \Psi_i - \Psi_j \\ \alpha_i - \alpha_j \\ \lambda_i - \alpha_j \end{pmatrix}$$

Where the derivative $\nabla_{\Psi,\alpha,\lambda} \pi_j(\mathbf{b}_t, \mathbf{d}_t | \mathbf{s}_i; \theta_i)$ comes from differentiating maximised expected payoff with respect to the parameters, employing the envelope theorem. This should reduce the bias caused by different food banks bidding subject to different parameters.

In Appendix J.3.2 I consider my results differ when I do not pool information across food banks.

H.3.4 Approximation

Having evaluated the ex-ante value function across the grid of states I fit a polynomial function of the states to the ex-ante value function. I include all interaction terms. The fit is performed using a standard weighted least squares procedure, weighting by the sum of posterior probabilities. This is to ensure that state observations that are more likely receive greater weight.

The main version uses a simple quadratic function. This is done primarily because my counterfactuals occasionally require extrapolation (given many of my counterfactual mechanisms do not allow food banks to maintain their balanced level of stocks). A quadratic polynomial has the appealing property that changes in the extrapolated values are constrained to be linear. The difficulty with higher-ordered polynomial (e.g. cubics, quintics etc) is that extrapolated values can be much further from interpolated values.

To validate the quadratic approximation I consider several measures of fit: The R^2 from the regression, as well as the results from considering higher order polynomials (up to order 6). These results are presented in Appendix J.3.1.

H.3.5 Continuation Value

Given the approximated ex-ante value function I evaluate the continuation value by taking an expectation of the polynomial function, given the distribution of $\mathbf{s}_{it+1}$ given $\mathbf{s}_{it}^a$. This is done using standard recursive formulae for the higher order moments of normally distributed random variables. With the continuation value in hand I evaluate the combinatorial flow-payoff using the definition of the estimated pseudo-static payoff function. All the above analysis is performed separately for time periods from both before and after period 553, when there is a structural break in the supply of fresh produce. I then average (weighting appropriately) the estimated flow payoffs from either side of the break. In future I will test whether continuation values and hence estimated combinatorial flow-payoffs are constant over the break.

H.4 Type 2 Food Banks

I now discuss the model of Type 2 food banks, the food banks who do not bid, nor win, regularly. This means I do not have significant identifying variation to allow estimation of their model parameters without a large degree of noise. Furthermore, many of these food banks never win certain types of food at all, meaning their parameters are not separately identified.

The Type 2 food banks consist of those food banks who win fewer than 200 lots over the sample period, and excludes the food banks who's locations are unknown or consume fewer than 30 lots over the period (which make up 2.5% of total consumption).

H.4.1 Differences to Type 1s

The key difference is that Type 2 food banks are assumed to be myopic bidders. That is, they are not forward looking. Aside from this assumption, I estimate the model using the same specification and estimation procedure as I used to estimate the pseudo-static payoff function for Type 1 food banks. I continue to recognise that I not observe food banks' stocks, and that variation in stocks is likely to be a key source of variation in bidding behaviour.

Due to the lack of variation in winnings and bids (an even more extreme degree of censoring), I assume that the combinatorial pay-off function for Type 2 food banks comes from the same hierarchical distribution as the pseudo-static payoff from Type 1 food banks. This means Type 2 food banks have the same Φ parameters, and their Ψ_i parameters are drawn from the same hierarchical distribution.[68] I also assume their lot specific values v_{itl} have the same variance as Type 1 food banks. I also assume they have the same beliefs as Type 1 food banks.

Therefore at each iteration of the estimation procedure I do the following: First, given the previous draw of parameters and unobserved states, draw censored observations from their conditional posterior. Second, given the previous draw of parameters and censored observations, use the Carter-Kohn algorithm to draw unobserved states from their conditional posterior distribution. Third, draw Φ, σ_l, and hierarchical parameters $(\boldsymbol{\psi}, \Sigma^{\psi})$ from the unconditional posterior distribution of Type 1 food banks. Fourth, use the Gibbs Sampling algorithm described above to draw Ψ_i, then α_i and λ_i parameters from their conditional posterior distribution. Finally, every 10th iteration, draw beliefs from their posterior distribution using Metropolis Hastings.

I use the same specification of priors as type 1 food banks and the same distributional/functional form assumptions. Again, I focus on just the first 25 unique auctions each period. I run 200,000 iterations, burning out the first 100,000 draws, and perform 4 independent chains. I keep 250 parameter draws from each chain, sampled uniformly, maintaining the correlations with the sampled parameters from Type 1 food banks. I estimate around 2400 parameters across 1.6 million observations, 1.45 million of which are censored.

[68] I previously considered imposing that they all have the same Ψ_i, however due to the large amount of data, and the large degree of censoring convergence is impractically slow.

H.4.2 Discussion

In reality even the Type 2 food banks are likely forward looking. And in principle, rather than interpret what I estimate as a static payoff, I could interpret it as another pseudo-static payoff function. Therefore I could apply the third stage estimation procedure. While this procedure is likely to produce fairly imprecise results, due to the lack of variation in bidding behaviour and imprecise estimates, I will consider this approach in a future robustness exercise.

However, the cost of misinterpreting their pseudo-static payoff function as a payoff function is potentially large. This is despite that Type 2 food banks only consume a relatively small amount of food under the Choice System, and even less under the Old System. This is because pseudo-static payoffs take into account expected future flow payoffs. Therefore when summing over time periods we essentially double count flow payoffs, skewing results towards these (typically lower priority) food banks. To alleviate this issue (and only in my final welfare calculations) I make the simplification that the exante value function for Type 2 food banks can be written as $K + \mathbf{s}_i^{gT}\Psi_i\mathbf{s}_i^g$. Comparing this to the ex-ante value function for Type 1 food banks, the quadratic term is just the pseudo-static payoff from winning no lots each period. This simplification therefore asserts that the ex-ante *marginal* value function, given here by K, is independent of $\mathbf{s}_i^g$. Given that type 2 food banks bid so infrequently this simplification is plausible. We can then write the flow payoff function as:

$$j(\mathbf{s}_{it}^g) = \mathbf{s}_{it}^{gT}\Psi_i\mathbf{s}_{it}^g - \beta(K + E[\mathbf{s}_{it+1}^{gT}\Psi_i\mathbf{s}_{it+1}^g|\mathbf{s}_{it}^g]) = (1-\beta)\mathbf{s}_{it}^{gT}\Psi_i\mathbf{s}_{it}^g - 2\beta\boldsymbol{\mu}_i^T\Psi_i\mathbf{s}_{it}^g - \beta[\boldsymbol{\mu}_i^T\Psi_i\boldsymbol{\mu}_i + Trace(\Psi_i\Sigma_j)]$$

Given that we normalise $j(0) = 0$, the constant term drops out. This process ensures that we do not double count flow payoffs. It is true that this simplification will impact my welfare calculations, but the effects are expected to be minor given that Type 2 food banks only consume a relatively small amount of food under the Choice System.

Even though I use this approach for welfare calculations, I continue to assume to assume myopia for counterfactual simulations. That is, their accept/reject decisions are based the estimated pseudo-static payoff under the Choice System, rather than the estimated flow payoffs and a counterfactual equilibrium continuation value. Therefore, my simulated counterfactual equilibrium will be invalid. Using the pseudo-static payoff essentially assume that the mechanism reverts to the Choice System in the following period. These food banks rarely bid, and rarely win, and are more likely than others to be picky. Therefore they are likely to reject more often than they should, as their value function incorrectly assumes they can be picky next period. In practice, for Type 1 food banks, I find that the discounted continuation value is not a large component of accept/reject decisions under the Old System. This is because the continuation value is extremely flat - much more so than the continuation value under the Choice System. Therefore this inaccuracy is likely to be minor.

I Additional Estimation Results

In this Appendix I report additional estimation results, adding to those in section 6. This includes tables and plots of parameter estimates, Gelman-Rubin Convergence tests, and model fit.

I.1 First Stage

I.1.1 Shape Parameters

Differ if more than 500 loads. Category specific shape parameters are given in column 1 of Figure 19. Only the 8 most common categories (with more than 500 loads auctioned over the period, excluding fresh produce) have category specific parameters, the remainder are constrained to be equal. The estimated shape parameters for mixed loads is 0.0538 (0.0383,0.0685). These parameters all lie within the interval (-0.05,0.1) with the exception of Condiments, Cereal, and Meals which have values exceeding 0.1, suggesting that winning bids on these types of food have larger right tails, likely due to subcategories, such as peanut butter, that attract extremely high bids. Only Dairy has a shape parameter that is estimated to be significantly below zero, meaning that winning bids on Dairy are bounded above. This is perhaps due to the storage requirements for Dairy products.

I.1.2 Scale Parameters

The standard deviation of the winning bid on lot l is given by $\frac{\sigma_l}{\xi_l}$. Column 2 of Figure 19 give the estimated scale parameters. These typically lie between 2000 and 5000, with the exception of fresh Produce, for which winning bids are typically clustered around the reservation price and have small standard deviation, and Pasta, which receives a small number of very high bids.

Column 2 of Figure 19 gives the additional scale fixed effect from a load being from the "other", "mixed" or "assorted" subcategories. If this value is negative it suggests these subcategories have a smaller standard deviation than other subcategories within the category. Most often, these parameters are not estimated to be significantly different from zero. I also estimate an additional parameters for when the load has already been auction previously, since food banks have additional information about previous bids on the load. This parameter is estimated at 3370 (3030,3710).

I.1.3 Location Parameters

Figure 20 plots the subcategory specific fixed effects. Results are strongly correlated with the results presented in figure 4 ($R^2 = 0.74$). Figure 21 plots the slope coefficients on the log of aggregate supply across the five "Use" types of food. Aggregate supply is given by both the previous 30 day aggregate supply and supply being auctioned that period. Coefficients are standardised, so that a one standard deviation increase in log 30 day supply of "Ingredients" decreases winning bids by 0.007 standard deviations. Estimated parameters are typically very small, and only significantly negative for the 30 day supply of Meals, Ingredients, and Snacks. I estimate a significantly positive and small (0.006 standard deviations) effect for contemporaneous supply of condiments, however this is difficult to interpret. It is possible that the small estimated coefficients are due to the extremely coarse food groupings I have used.[69] In practice it is unlikely that food banks keep track of the food supply from particularly detailed food groups. However, this remains a weakness of this analysis.

[69] however in a previous version of the model I estimated category specific slopes, and found similarly small effects.

Figure 19: Category Specific First Stage Parameters

Category	Shape	Scale	Scale (other)	Maroon	Loads	Threshold 1	Threshold 2
Baby	0.0886	3100	-1,200	-846	-10,500	-3,020	-175
	(0.0659,0.113)	(2790,3410)	(-1,970,-173)	(-1,520,-239)	(-17,700,-5,890)	(-3,230,-2,800)	(-216,-136)
Bev	0.0174	2170	-370	-756	-5,200	-3,470	-320
	(-0.00205,0.0388)	(2080,2280)	(-812,130)	(-2,470,835)	(-5,880,-4,590)	(-4,090,-2,850)	(-439,-203)
Baked	0.0886	2160	240	-846	-10,400	-3,020	-175
	(0.0659,0.113)	(1850,2460)	(-728,1370)	(-1,520,-239)	(-16,100,-6,040)	(-3,230,-2,800)	(-216,-136)
Cereal	0.133	4560	-520	-2,600	-2,750	-3,020	-175
	(0.0935,0.175)	(4320,4810)	(-886,-179)	(-4,710,-486)	(-4,700,-1,290)	(-3,230,-2,800)	(-216,-136)
Condiment	0.342	3660	-226	623	-9,140	-3,020	-175
	(0.277,0.417)	(3380,3910)	(-916,569)	(-1,070,2370)	(-11,800,-6,750)	(-3,230,-2,800)	(-216,-136)
Dairy	-0.0421	2340	-224	-846	-5,600	-2,680	-186
	(-0.0726,-0.00883)	(2220,2480)	(-945,817)	(-1,520,-239)	(-6,470,-4,660)	(-3,410,-1,940)	(-317,-64.5)
Fresh	0.0886	576	-3.88	-124	-2,940	-564	-474
	(0.0659,0.113)	(516,635)	(-73.8,63.4)	(-1,440,1230)	(-3,290,-2,600)	(-1,020,-122)	(-600,-356)
Frozen	0.0204	2560	1110	-1,390	-5,930	-3,020	-175
	(-0.0268,0.0762)	(2390,2780)	(-791,3930)	(-3,030,225)	(-9,350,-3,500)	(-3,230,-2,800)	(-216,-136)
H/B	0.0886	3570	-1,460	-846	-5,720	-3,020	-175
	(0.0659,0.113)	(3290,3900)	(-1,890,-977)	(-1,520,-239)	(-9,310,-2,590)	(-3,230,-2,800)	(-216,-136)
Meals	0.149	4020	-587	-846	-4,670	-3,020	-175
	(0.113,0.194)	(3830,4220)	(-1,090,-49.4)	(-1,520,-239)	(-5,970,-3,480)	(-3,230,-2,800)	(-216,-136)
Meat	0.0886	4360	1550	1240	-10,400	-3,020	-175
	(0.0659,0.113)	(3930,4840)	(577,2660)	(-1,260,3660)	(-15,300,-6,530)	(-3,230,-2,800)	(-216,-136)
Cleaning	0.173	2730	137	2110	-4,730	-3,020	-175
	(0.116,0.235)	(2520,2970)	(-393,692)	(426,3690)	(-6,190,-3,500)	(-3,230,-2,800)	(-216,-136)
Pasta	0.0886	5730	-612	-846	-6,790	-3,020	-175
	(0.0659,0.113)	(4360,7390)	(-2,600,1090)	(-1,520,-239)	(-11,900,-2,490)	(-3,230,-2,800)	(-216,-136)
Snack	0.0374	2220	-353	-767	-6,680	-3,830	-239
	(0.0211,0.057)	(2160,2310)	(-572,-119)	(-2,230,653)	(-7,980,-5,590)	(-4,490,-3,210)	(-331,-147)
Vegetables	0.0886	3540	-374	1230	-2,870	-3,020	-175
	(0.0659,0.113)	(3260,3850)	(-1,070,409)	(-761,3090)	(-5,150,-1,310)	(-3,230,-2,800)	(-216,-136)

Note: 95% Credible Intervals are given in parentheses.

Figure 19 column 4 displays the coefficients on maroon dummy variables. Parameters are constrained to be equal for categories with fewer than 50 maroon loads. Importantly, Maroon loads have a reservation price of zero, rather than -2000. The model I estimate focuses on difference from the reservation price, rather than raw winning bids. Therefore, with the exception of Fresh produce (which also has a reservation price of 0), I estimate that winning bids on maroon loads are systematically higher than winning bids on non-maroon loads when estimated parameters are significantly greater than -2000. On average Maroon loads attract winning bids around 1000 shares higher than non-maroon loads. Figure 19 column 5 shows the linear slope coefficient on the number of homogenous loads auctioned simultaneously. Estimates are all significantly below zero, so that when multiple loads attract lower bids than single loads. This is sensible, since there is less competition for each load.

For the 8 different economic regions (and Canada), only loads from the "MidEast" region attracts significantly different winning bids, but the effect is small at around 0.001 standard deviations. Previously auctioned loads attract significantly lower winning bids, with a point estimate of -11,700 (-12,400,-11,000), around a third of a standard deviation. Loads that contain different categories of food attract lower bids, but the magnitude is small (less than 0.001 standard deviations), and significantly different from zero only when the lot contains four distinct types of food. Loads with free delivery, additional notes, or additionally shelf stable products do not attract significantly different winning bids. Finally, loads with restrictions on where they can be sent, or how the food must be picked up attract significantly lower winning bids, with a point estimate of -3,490 (-4,360,-2,700).

Figure 20: Estimated Subcategory Fixed Effects

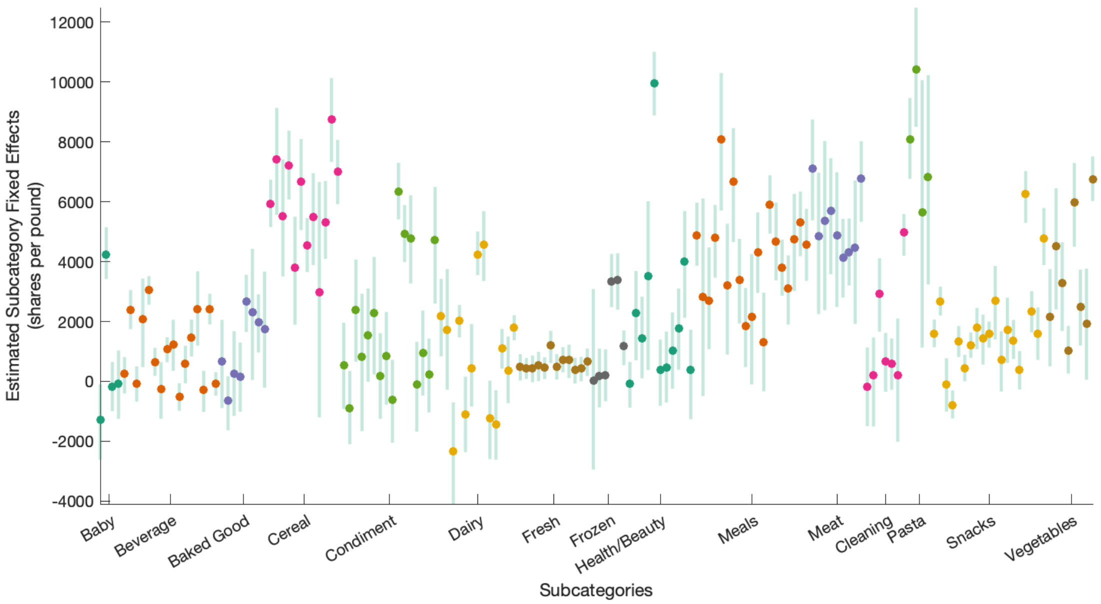

Note: Plot shows coefficients on location shifting subcategory specific dummy variables. Points give posterior means, and 95% Credible Intervals are given by the shaded lines. To interpret magnitudes in terms of standard deviations of winning bids, one must multiply by the associated shape parameter, and divide by the associated scale.

I.1.4 Threshold Parameters

Figure 19 column 6 gives the estimated threshold parameters $\underline{R}^c$. These are all estimated to be far from zero, indicative of the high likelihood of observing winning bids at the reservation price. Column 7 gives the other threshold parameters $\bar{R}^c$. To interpret the parameters in terms of the excess mass just above the reservation price, they must be divided by the standard deviations (around 30,000), so that on average there is around 0.5% more mass just above the reservation price than expected. The estimated exponential parameter is 0.395 (0.343,0.462), so that the model predominantly rationalises bids within 5 shares of the reservation price in this way.

I.2 Second Stage

I.2.1 Lot Specific Pay-off Parameters

Figure 22 panel (A) plots food banks' estimated transportation costs, measured in consumer surplus. Coefficients cannot be interpreted as willingness to pays as they are not divided by marginal value of wealth (λ_i). Coefficients are positive, suggesting transportation is costly, and we see significant differences across food banks. Figure 22 panel (B) plots the estimated log marginal value of wealth

Figure 21: Estimated Effect of Aggregate Supply on prices

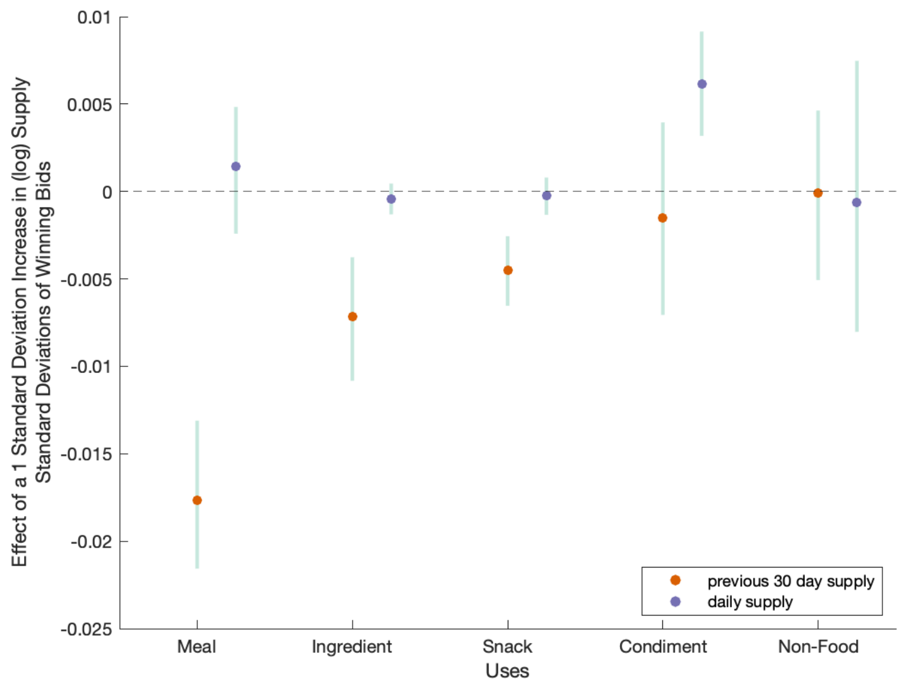

Note: Plot shows coefficients on aggregate supply, by food use for both daily supply and the previous month's supply. Points give posterior means, and 95% Credible Intervals are given by the shaded lines. In non-standardised terms at the mean of 1600 tons of meals (food that can be consumed as a meal in itself) per month, an increase in the previous 30 day supply of meals by 1000 tons, around 50 loads, decreases the expected winning bid by around 500 shares.

across food banks, $\hat{\lambda}_i$. Estimates above zero are relatively more budget constrained than the median food bank, and I estimate significant variation in these parameters. This suggests that shares are not allocated correctly, since to achieve efficiency the social planner would equate marginal values of wealth.[70]

Figure 23 panel (A) plots the estimated standard deviations of the lot specific idiosyncratic payoffs across category combinations. These tend to be between 10,000 and 30,000, which is significantly higher than the observed standard deviation of bids (around 3,000). They are large due in part due to the much larger variance of bids plus the markup term $b_l + \frac{\Gamma_l(b_l)}{\nabla_b \Gamma_l(b_l)}$, and also to rationalise the large degree of censoring. If only 2% of bids are observed, then the observed variation is only the variation in the far right hand tail.

[70]Note that identification of these parameters, which comes from variation in the variance of bids across food banks, strongly rests on the assumption that the lot specific payoff has the same variance across food banks. Therefore it is not possible to determine whether variation in these parameters is actually due to unmodelled variation in lot specific variances.

I.2.2 Combinatorial Pay-off Parameters

Figure 23 panel (B) plots the estimated Φ parameters, essentially how $k(\mathbf{s}_i, \mathbf{s}_0)$ vary with stocks of subcategories. Estimated parameters are strongly correlated with with the estimated first stage subcategory parameters plotted in figure 20 ($R^2 = 0.82$)

Figure 22: Estimated distance and marginal value of wealth parameters

(A) (B)

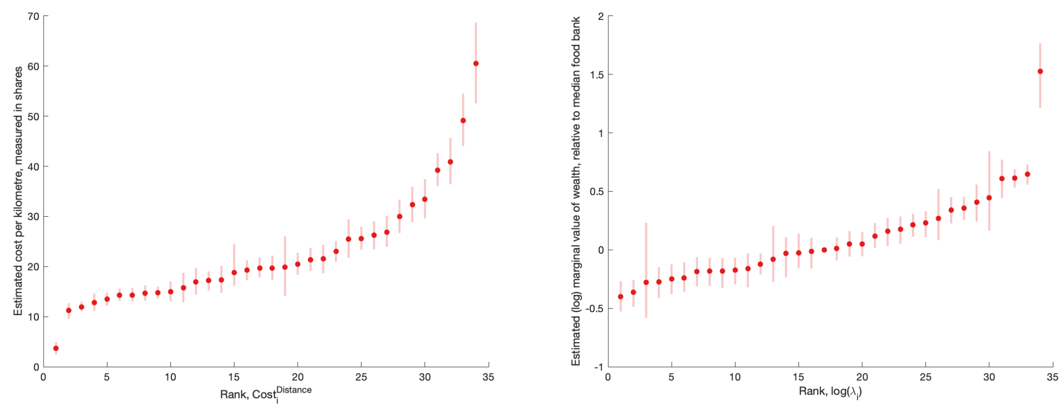

Note: Plot shows posterior means and 95% credible intervals of estimated distance coefficients (in km, panel (A)) and log marginal value of wealth (panel B). The Marginal value of wealth is normalised to 1 for the Type 1 food banks with median consumption.

I.3 Third Stage

I.4 Type 2s

I.4.1 Unobserved State

Figure 24 panel (A) presents estimated mean net donations for Type 2 food banks. Estimates are generally greater (less negative) than for Type 1 food banks. This is to be expected given these food banks typically win less food, suggesting they need less food in the first place. Likewise Figure 24 panel (B) presents estimated standard deviations of net donations for Type 2 food banks. Estimates are generally smaller than for Type 1 food banks. This is partly surprising, since one might expect that each period larger food banks give out a more predictable amount to their clients, and receive a more predictable amount from local donors. This is essentially a law of large numbers given Type 1 food banks are expected to give out and receive more food than Type 2 food banks. However, Type 2 food banks are not necessarily smaller than Type 1s, and many have larger Goal Factors. The fact they rely on the Choice System less is likely due to having many local donors, which may lead

Figure 23: Estimated lot specific standard deviations and subcategory parameters (Φ)

(A) (B)

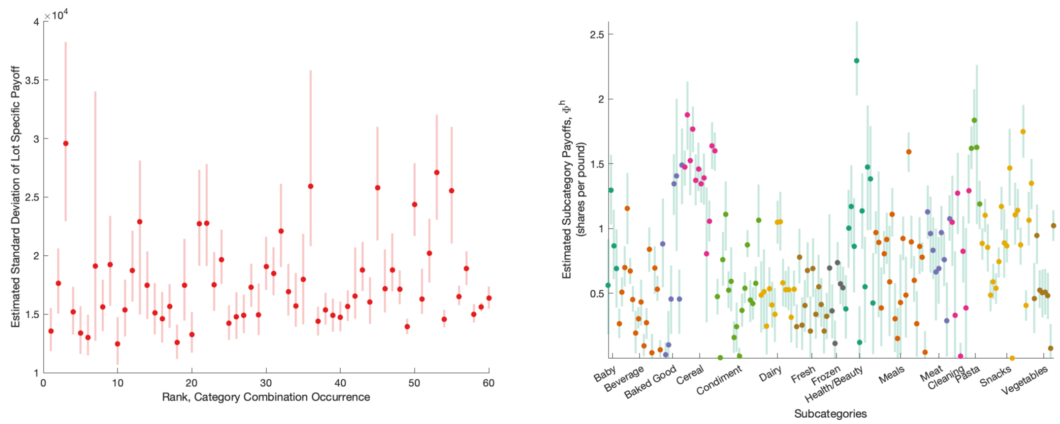

Note: Plot shows posterior means and 95% credible intervals of estimated standard deviations of υ_{ilt} (panel A), and of estimated Φ subcategory pay-off parameters (panel B). For panel (A) The x-axis shows different combinations of categories included in the same lot, as each of the 59 most common unique combinations receives their own parameter. More common combinations appear further to the right. The furthest right parameter corresponds to the remaining 423 observed unique combinations, which make up 7% of the data. For panel (B) The scale can be interpreted as consumer surplus, measured in shares. A coefficient of 1 can be interpreted as every addition pound of food increasing consumer surplus by one share.

to them receiving a more constant supply of local donations over time. It is also possible that these results are driven mechanically by my priors (as food banks who win more food may mechanically have a larger variance of their winnings), making it important that my priors about the variance are informative.

I.4.2 Lot Specific Payoff

Figure 25 panel (A) plots the distance costs for Type 2 food banks. Some of the estimated coefficients are in the ranges of the type 1 food banks, these are likely the large Type 2 food banks with access to many local donors. Mostly, Type 2 food banks have much higher distance costs. Figure 25 panel (B) plots the log of the estimated marginal value of wealth, which is still measured relative to the median Type 1 food bank. Type 2 food banks are estimated to generally have much lower estimates, suggesting that shares are not as valuable to Type 2 food banks as they are to Type 1 food banks. This is unsurprising given that Type 2 food banks are known to already need less food, either due to access to local donors, or as they have a smaller amount of poverty in their local area.

Figure 24: Estimated unobserved state parameters (Type 2)

(A) (B)

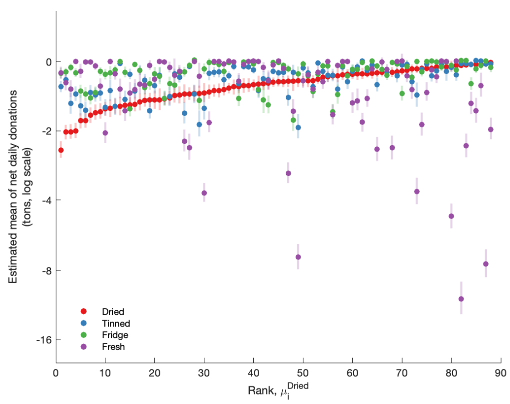

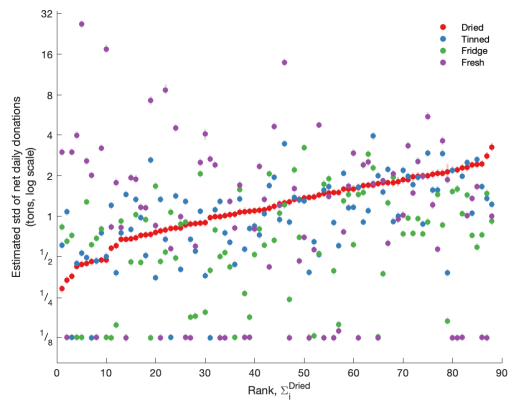

Note: The figure plots posterior means for the mean and standard deviations of net local donations, as well as 95% credible intervals. Results are sorted according to the estimates for the Dried storage type. The plot excludes the 'non-food' type, to improve graphability.

I.4.3 Combinatorial Payoff

Figure 26 plots the distribution of Type 2 food banks' willingness to pays for a single lot from each storage type, given stocks of zero. Estimates are most often negative, suggesting that storing food is costly. Estimates are generally higher (lower storage costs) for Type 2 food banks than Type 1 food banks. This is most likely because the state is normalised to zero for the first date in my sample period. The zero, normalised, state is likely higher for Type 1 food banks as they typically win more food. Therefore it is sensible that these food banks are more capacity constrained to begin with.

I.5 Diagnostics

This sub-Appendix reports Gelman-Rubin statistics, allowing us to assess model convergence. I follow the approach to constructing the test statistics laid out in Gelman et al. (1995). Figure 27 reports the results of this analysis. For each set of parameters I report the proportion of statistics below the recommended cutoffs of 1.2 and 1.1. I report results for both types of food banks, except in cases when the relevant parameters are the same for both types.

Broadly I have evidence of convergence, though not as strong as one might hope for the second stage parameters.

For the first stage parameters every parameter is found to converge except one of the 'other subcategory' scale parameters (likely due to a lack of observations). For the parameters of the second stage evidence of convergence is less strong, though only when we focus on the more stringent cutoffs. The lack of convergence is clustered within the four least regularly bidding Type 1 food banks, and the ten least regularly bidding Type 2 food banks. Each individual food bank has a

Figure 25: Estimated distance and marginal value of wealth parameters (Type 2)

(A) (B)

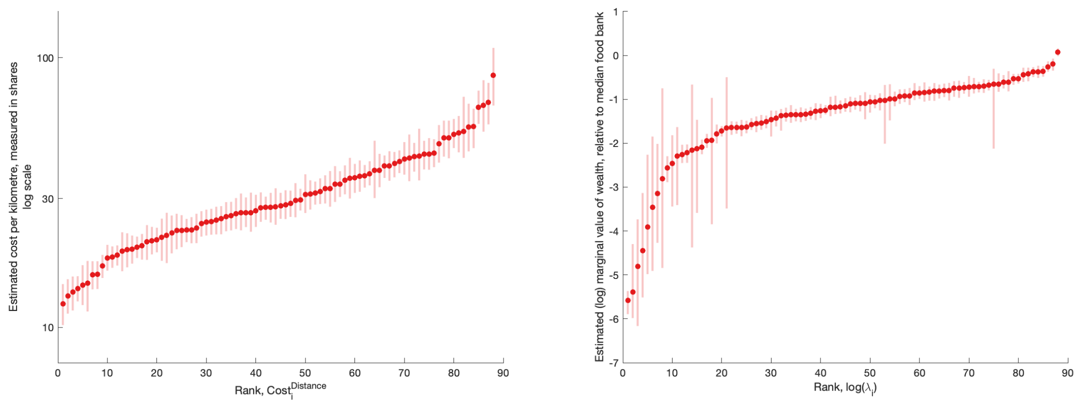

Note: Plot shows posterior means and 95% credible intervals of estimated distance coefficients (in km, panel (A)) and log marginal value of wealth (panel B). The Marginal value of wealth is normalised to 1 for the Type 1 food banks with median consumption.

relatively small effect on my counterfactuals (with the exception of the 5 largest food banks, all of whom converged). Therefore the effect on my main results of this non-convergence is likely to be very minor. For the sake of posterity, I now discuss this lack of convergence in more detail.

My sampler likely did not fully converge to the target distribution for certain parameters. Convergence was already expected to be slow due to the large degree of censoring, and the 300,000 iterations were likely insufficient to achieve full convergence for every parameter.[71]

I.6 Fit

I.6.1 First Stage

Figure 28 plots the estimated and empirical probability a food bank wins a lot given their bid, where bids are measured in distance from the reservation price. The estimation drops the first 60 days, and the final 150 days, and then randomly samples 95% of the remaining data. The other 5% is used as a validation dataset. Probabilities are plotted by taking an expectation over covariates. The discontinuity in probability occurring at zero is due to the non-negligible probability of ties at the reservation price.

[71]In future more iterations should be used. In particular, just as I presently only sample beliefs every 10th iteration due to computational cost, in future I will only run the Carter-Kohn algorithm (drawing stocks from their posterior) every 10th iteration. These Carter-Kohn steps were by far the most computationally intensive, taking around 90% of all the computation time. This will allow me to run the sampler for around five times as many iterations in the same space of time. This should improve convergence by allowing me to overcome the large degree of censoring.

Figure 26: Estimates of Ψ_i (Type 2)

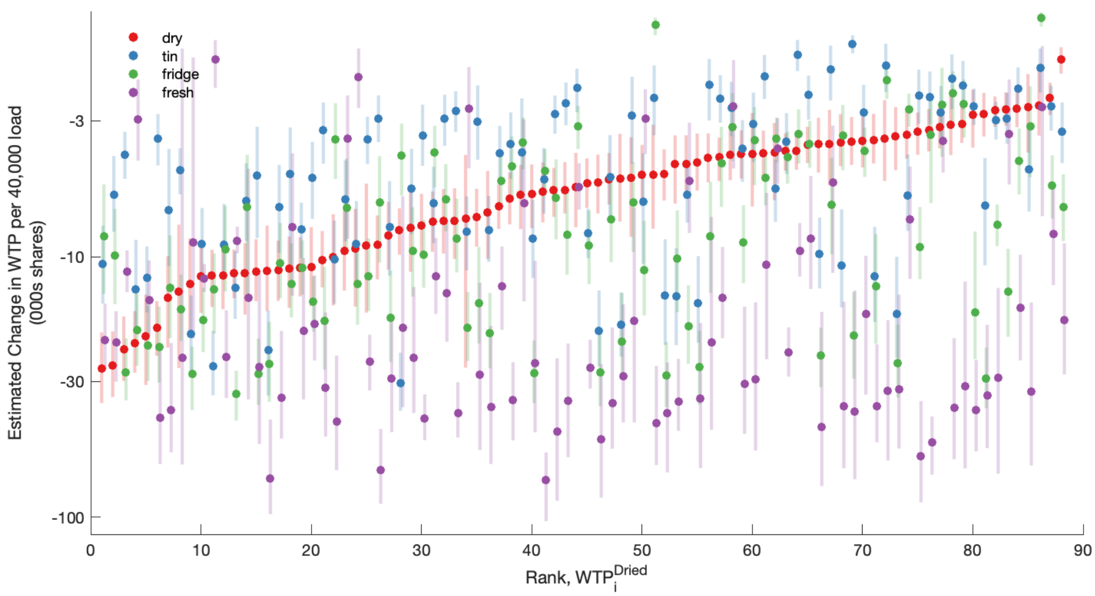

Note: Figure plots posterior mean equilibrium willingness to pay for a 40,000 load for each storage type. Bars give the 95% credible intervals. Estimates are ordered according to the estimates for Dried loads. The plot excludes estimates for non-food storage type. WTPs are evaluated when stocks are zero.

The model does a good job of matching the probability of winning at the reservation price, and just above it. The fit worsens around zero (typically 2000 shares above the reservation price) due to excess mass at zero, and hence excess cumulative probability above zero. The model is unable to rationalise food banks' bids being anchored around zero. However this inaccuracy is not large, even if it is statistically significant - the vertical distance between the two lines never exceeds 0.05.

I.6.2 Second Stage

Figure 29 presents several observed moments, comparing them to simulated moments. For the simulated moments I present the posterior mean moment, as well as the 2.5th and 97.5th percentiles. In sample moments are calculated on the training sample (in-sample), which cuts off the first 60 and final 150 days. The validation sample (out-of-sample) uses only the final 150 days.

Most of the moments I consider are self-explanatory, except for the number of lots won, relative to the observed number. For each simulation this takes the number of lots a food bank wins, dividing this by the number of lots they were observed winning. I then consider the average of this proportion across food banks. It mechanically equals one in the observed data. This moment is typically fit quite well by the model, with both Type 1 and Type 2 food banks winning similar amounts of food in y simulations compared to the true data. Type 2 food banks perhaps win too

Figure 27: Gelman-Rubin Convergence Statistics

	Type 1 food banks		Type 2 food banks	
Parameters	Prop < 1.1	Prop < 1.2	Prop < 1.1	Prop < 1.2
Γ	0.996	1	-	-
$\boldsymbol{\mu}_i$	0.947	0.976	1	1
$\boldsymbol{\Sigma}_i$	0.9	0.935	1	1
$\boldsymbol{\Phi}$	0.976	1	-	-
$\boldsymbol{\lambda}_i$	0.909	0.97	0.784	0.943
$\boldsymbol{\sigma}_l$	0.933	1	-	-
Distance	0.971	1	0.966	1
$\boldsymbol{\psi}_i$	0.902	0.975	0.92	0.986
$\boldsymbol{\psi}$	1	1	-	-
$\boldsymbol{\Sigma}^{\psi}$	1	1	-	-

many lots and Type 1 food banks perhaps slightly too few. This is likely due to the simulated over bidding of Type 2 food banks. This moment is important because it shows that the model correctly predicts food banks equilibrium allocations, even if it does not correctly predict bidding behaviour. A similar pattern is seen when we consider equilibrium expenditure (essentially weighting winnings by value) as well as food won by type of food.

Model fit, in terms of bidding behaviour, is poor for Type 2 food banks. They are predicted bidding more often than they actually do, and bidding too aggressively conditional on bidding. This is typically due to parameters failing to converge properly. In particular, the estimated marginal value of wealth for these food banks are far too low. Fortunately given that these food banks consume a relatively small proportion of total food, and this is predicted well by the model, these food banks have a small total impact on welfare, and so these inaccuracies will not majorly impact my counterfactuals.

The model fits much better for Type 1 food banks, though still estimates them bidding too aggressively, with average bids around 70% larger than observed. However, relative to the already large standard deviation of bids, this is not major. These inaccuracies seem to be caused by simulation error, as for each simulation I must numerically find optimum entry and bidding decisions, for which I use a simple hill-climbing heuristic that need not necessarily find a global optimum.

Figure 28: First Stage Fit, actual vs simulated

(A) (B)

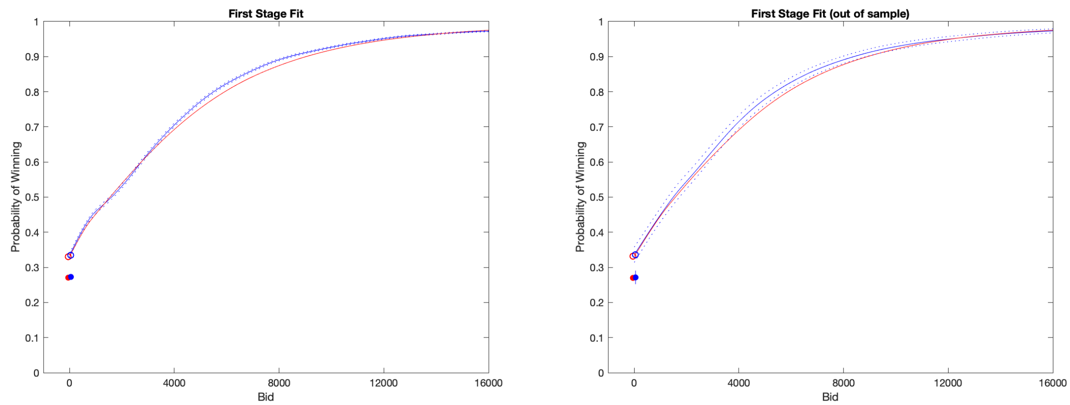

Note: probability of winning given bid, i.e. cdf of winning bids. discont due to ties. simulated values from estimated distribution vs empirical distribution. averaged over covariates.

J 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 J.1, J.2, and J.3 respectively.

J.1 First Stage

In this appendix I consider how robust is estimation to the specific assumptions I make about equilibrium beliefs. At present, I focus on the assumptions made on beliefs in Assumption 4. Specifically, that every food bank faces the same distribution of maximum rival bids, and that beliefs are not a function of individual food banks' states, but rather aggregate states. In Appendix J.1.3 I consider the assumption that winning bids are conditionally independent across auctions.

J.1.1 Food bank Specific Beliefs

Assumption 4 part (v) imposed that $\Gamma_i = \Gamma$, so that every food bank is assumed to face the same distribution of rival bids. This simplification allowed me to estimate Γ on the distribution of winning bids only. I can test this assumption, testing whether the distribution of food bank i's rival's highest bids is significantly different from the distribution of winning bids. In practice, this involves replacing food bank i's winning bids with the second highest bid in each of these auctions. This permits a simple hypothesis test for food bank i: We construct this alternative dataset and consider whether the estimated Γ_i from this dataset is significantly different from the estimated Γ constructed from winning bids only. This can be done using a simple Score Test.

Figure 29: Estimation Moments: Observed vs Simulated

Moment	FB Type	Observed	In-Sample Mean	p2.5	p97.5	Observed	Out-of-Sample Mean	p2.5	p97.5
Average No. bids	1	1.33	1.32	1.25	1.43	0.899	1.08	0.938	1.22
(per period)	2	0.252	0.596	0.554	0.633	0.146	0.525	0.476	0.57
P $\geq$ 1 bids	1	0.457	0.488	0.472	0.504	0.344	0.437	0.412	0.455
(per period)	2	0.157	0.363	0.345	0.38	0.0997	0.33	0.307	0.346
Mean bid	1	1640	2560	2260	2770	1010	2710	2350	3050
(given enter)	2	2590	6120	5600	6640	1590	6480	5790	7190
Std of bids	1	3890	4600	4420	4790	3020	4540	4290	4720
(given enter)	2	5290	8040	7380	8780	3260	7880	7080	8810
No. of lots won	1	1	0.829	0.562	1.16	1	0.887	0.429	1.53
relative to observed	2	1	1.54	0.979	2.32	1	2.62	1.04	6.02

Note: Moments are calculated for Type 1 and Type 2 food banks separately. 'Out of sample' refers to the final 150 days which are dropped from estimation. The final set of moments considers the number of lots won by each food bank in each simulation, and how this compares to the number of lots actually won. For each food bank and each simulation I consider the ratio of lots actually won and to the lots won in the simulation. Values closer to 1 are closer to the observed data. Values above 1 have food banks winning too many lots on average.

In figure 30 panel (A) I present the distribution of score test statistics across food banks. Under the null hypothesis these statistics take a χ^2 distribution with 268 degrees of freedom (the number of first stage parameters). None of these hypothesis tests can reject the null hypothesis at the 10% significance level.

J.1.2 Dependence on Aggregate Supply

Assumption 4 part (v) also requires that beliefs do not depend on any individual food banks' state, and instead only depend on aggregate statistics, such as the aggregate supply of various types of food. The argument is that equilibrium is sufficiently competitive that no individual food bank's behaviour is able to significantly shift the distribution of equilibrium winning bids. If this is the case, then no individual food bank's state will be able to significantly shift this distribution either.

To test this assumption I consider whether any food bank has an individually significant effect on the distribution of equilibrium winning bids. The results presented in Appendix J.1.1 act as evidence in favour of this hypothesis, essentially presenting the distribution of equilibrium winning bids when each food bank is 'removed' from the system in turn. To go further, I can also consider whether the distribution of equilibrium winning bids changes when data from food bank i *and* the auctions they won are removed from the system. If food bank i has a significant effect on the distribution of winning bids, we would expect that the distribution of winning bids is different when we drop all the data from food bank i. In figure 30 panel (B) I present the distribution of score test statistics across food banks. Under the null hypothesis these statistics take a χ^2 distribution with 268 degrees of freedom (the number of first stage parameters). None of these hypothesis tests can reject the null hypothesis at the 10% significance level.

Figure 30: Robustness: Stage 1

(A) (B)

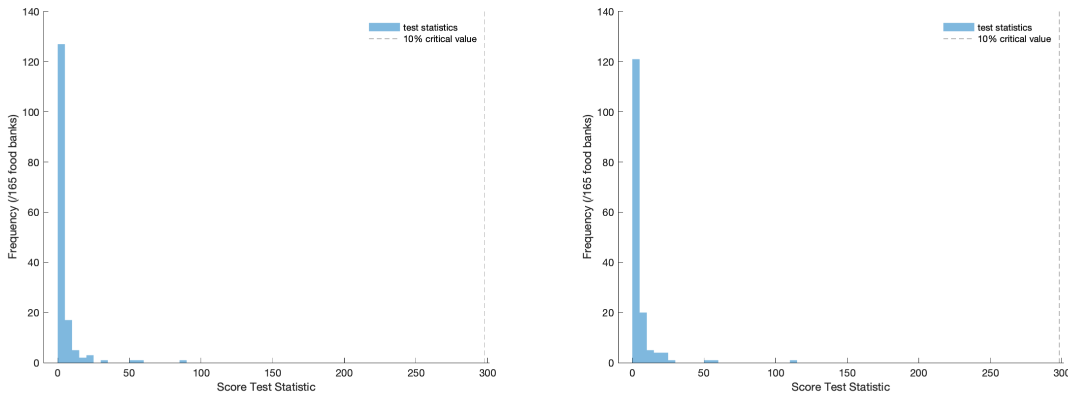

Note: The figure plots Score Test statistics from two robustness checks. Panel (A) relaxes the restriction that every food bank has the same equilibrium beliefs, while panel (B) tests whether any individual food bank's bidding behaviour has a significant effect on the distribution of winning bids.

J.1.3 Independence of Winning Bids

In this Appendix I investigate the assumption that winning bids within a period are conditionally independent across auctions. This assumption is necessary to ensure the joint probabilities of combinatorial outcomes $P(\mathbf{b}_t, \mathbf{d}_t | \mathbf{s}_t)$ can be written as products of the marginal distributions. Given the linear demand, or quadratic pseudo-static payoff parametrisation I make, I only require that winning bids are pairwise independent.

I investigate the validity of this assumption by investigating the degree of pairwise correlation in winning bids within a period. While a lack of correlation is not sufficient to infer independence, it at least suggests that one winning bid is not informative of another winning bid. This ensures that food banks' beliefs about joint probabilities of pairwise outcomes should be close to the product of the marginal win probabilities, meaning that errors from this misspecification are expected to be minor. Importantly, I only need to test for the presence of conditional correlation. Two winning bids are allowed to be conditionally correlated (conditional on covariates), as it is assume that food banks form beliefs conditional on the state. Correlation in winning bids is most likely to arise from the complementarity terms Ψ that also creates correlation in food banks' bids across different lots (given I assumed υ_{ilt} are uncorrelated across l). For example, since we expect lots to be substitutes a food bank's bids are likely to exhibit negative correlation (conditional on the state). Therefore we might also expect winning bids to exhibit negative correlation.

Writing $\bar{b}_{lt}$ for the winning bid on lot l at time t I investigate correlation between winning bids using the following regression specification:

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

$\mathbf{x}_{lt}$ give lot specific covariates, using the covariates included in the first stage of the estimation procedure such as subcategory fixed effects. $\mathbf{s}_t^0$ give time specific common state variables that do not vary across lots, just as in the first estimation stage. I include every pair of auctions (l, l') that occur simultaneously, giving me around $800,000$ observations (essentially including each pair twice, once on each side of the regression). Under the null hypothesis of independence β^1 should equal zero. However, this specification imposes that the relationship between every pair of winning bids is the same. One might expect negative correlation between substitutes and positive correlation between complements. To allow for differential correlations I also consider a specification that interacts $\bar{b}_{l't}$ with all three sets of covariates, allowing the correlation to depend on observable characteristics of the lots. I also consider a specification that includes triple interactions between $\bar{b}_{l't}$, $\mathbf{x}_{lt}$, and $\mathbf{x}_{mt}$. This predominantly involves including dummy variables for whether both lots come from the same subcategory, the same region, etc.

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

Results are presented in Figure 31. I can reject the null hypothesis of independence at the 1% significance level in all of my specifications. Therefore I have evidence that the independence assumption is invalid. However, as is evident from examining the R^2 values, the degree of dependence is extremely small. The covariates alone account for 41.2% of the variation in winning bids. Including the $\bar{b}_{l't}$ interactions is then only able to explain an additional 0.3% of the variation in winning bids. This suggests that winning bids are very close to being independent, even though we can reject independence. Therefore while I have found this independence assumption to be invalid, I have also found that it is likely to be a very good approximation to food banks' beliefs. This analysis does not presently account for censoring of winning bids at the reservation price, but will do in future.

One final point worth considering is how this simplification might impact my results. One difficulty with simultaneous auctions is that it makes it difficult for food banks to win precise numbers of loads. They might bid on two loads of cereal wanting precisely one, but there is a non-negligible risk they win both or neither. The more dependence there is between winning bids, the more information the food bank has about how they should bid, giving them even more control over which lots they win. In the cereal example, if winning bids are negatively correlated, the food bank knows they can place two middling bids and they will likely win precisely one lot. If they are positively correlated they know they should place one high and one low (if at all) bid. Therefore, this assumption is expected to bias my results against the Choice System.

Figure 31: Robustness: Independence of Winning Bids

Specification	Covariates	F test df	p-value	R^2
$\bar{b}_{l't}$		1	0	0.08734
Covariates only	✓			0.4121
$\bar{b}_{l't}$	✓	1	0	0.4124
$\bar{b}_{l't} \times (\mathbf{x}_{l't}, \mathbf{x}_{lt}, \mathbf{s}_{0t})$	✓	457	0	0.414
$\bar{b}_{l't} \times (\mathbf{x}_{l't}, \mathbf{x}_{lt}, \mathbf{s}_{0t}, [\mathbf{x}_{l't} \times \mathbf{x}_{lt}])$	✓	681	0	0.4153

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.

J.2 Second Stage

I consider four alternate model specifications for the second stage. These are designed to test the model's robustness to relaxing key simplifications made in the main model. In Appendix J.2.1 I allow the value function to depend on common state variables. In Appendix J.2.2 I account for endogeneity in the observation equation using an control function procedure. In Appendix J.2.3 I consider robustness to the assumption of normally distributed lot specific payoffs by allowing the lot specific idiosyncratic payoff to follow a normal-inverse-gamma distribution. In Appendix J.2.4 I consider how the simplification to only consider data from 25 auctions each period impacts my results, by estimating the model using 50 auctions from each period.

J.2.1 Incorporating the Common State

In general food banks' continuation values depend on the common state variables, which contain information about future prices. Common state variables are captured by the demand index estimated in the first stage, mapping common states onto parameters of the distribution of winning bids.

In order to allow continuation values to depend on common states, when estimating the pseudo-static payoff function k in the second estimation step, it is necessary to allow k to vary with the demand index. Importantly the index must be interacted with food bank specific state variables, so that the marginal payoff also depends on the index. Otherwise dependence on the index will not be identified from bidding behaviour alone. I introduce the demand index by specifying k as follows:

$$k(\mathbf{s}_i, \mathbf{s}^0) = \Phi(I + D^0)\mathbf{s}_i^h + \mathbf{s}_i^{gT}\Psi_i\mathbf{s}_i^g$$

Where D^0 is a diagonal matrix with entry $D^0_{hh} = \sum_u \delta_u d^{0u}\mathbb{I}[h \in u]$, where $\mathbb{I}[h \in u]$ is a dummy variable for whether subcategory h has usage type u. d^{0u} is the demand index for food from usage type u, and δ_u are parameters to be estimated. These parameters describe how strongly bidding behaviour changes given changes in aggregate supply. For example, when supply is high, and so d_t^{0u} is low, winning bids are expected to be low. If supply is also positively correlated over time, the opportunity cost from losing a lot today is low, as winning bids are also likely to be low in future.

Therefore bidding will be *even* less aggressive today, and so we expect $\delta_u > 0$. This specification is natural - if there is dependence on common states, we would expect to see evidence of it to show up in a linear term. I interact the index with the subcategory stock term, rather than the storage type term, because the subcategory term reflects a food banks' 'wants', while the storage term is intended to reflect the costs that the food bank must put up with. This is relevant because the index affects how easily the food bank can win the types of food it wants, on behalf of their clients.

Different sized food banks, with different budgets and storage capacities, are expected to respond differently to variation in common states. For example, a food bank that is not heavily reliant on the Choice System for their staples is unlikely to be responsive to common states. Therefore I allow the $\boldsymbol{\delta}$ parameters to vary across food banks, but again employ a bayesian hierarchical model to ensure a degree of shrinkage for food banks for whom identifying variation is scarce. I assume that $\boldsymbol{\delta}_i \sim N(\boldsymbol{\delta}, \Sigma^{\delta})$, where priors for $N(\boldsymbol{\delta}, \Sigma^{\delta})$ are weak normal-inverse-wishart. The parameters are identified using variation in the demand indices, which arise from variation in the common states, and seeing how this translates into variation in bidding behaviour.

In figure 32 panel (A) I plot estimates of δ_u across food banks. None are significant at the 5% significance level. This is predominantly caused by the lack of variation in the demand indices - as we saw in Figure 21 winning bids do not vary much with variation in the common state variables. This explains the extremely large credible intervals relative to the scale: if $\delta^g = 1$ this means that a one unit increase in d_t^g (associated with a one share expected increase in the winning bids) is associated with a $\Phi \mathbf{z}_{tl}^h$ unit increase in bids.

Therefore we have evidence that food banks' continuation values also do not vary with common state variables. Consequently, when evaluating maximised expected payoffs and the ex-ante value function in the third estimation step, we do not need to explicitly consider dependence on the common states, as this will not impact estimates of the flow payoffs j backed out in the final step.

J.2.2 Endogeneity of the Inverse Bid System

In this appendix I consider the endogeneity of the observation equation, caused by non-additivities across lots. In essence, I re-estimate the second stage of my estimation procedure using a control function approach. The observation equation is given by:

$$\lambda_i y_{ilt} = \Phi \mathbf{z}_{tl}^h + \mathbf{z}_{tl}^{gT} \Psi_i (\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2 \sum_{m \neq l} \Gamma_m(b_{itm}) \mathbf{z}_{tm}^g) + \upsilon_{ilt}$$

Given that this step is essentially estimated using a bayesian regression, estimation requires that the error term υ_{ilt} is independent of the regressors. In general there exists a dependency between υ_{ilt} and b_{itm} (for $m \neq l$), because optimum bids (and entry decisions) are a function of every lot specific payoff. That said, we have reason to think this dependency might be small, since typically $\Gamma_m(b_{itm})$ will depend much more strongly on things other than υ_{ilt}. However, allowing for this endogeneity is relatively easy. The endogenous regressor is $\mathbf{z}_{tl}^g(\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g + 2\sum_{m\neq l} \Gamma_m(b_{itm})\mathbf{z}_{tm}^g)^T$, and there exists an obvious instrument for this regressor: $\mathbf{z}_{tl}^g(\mathbf{z}_{tl}^g + 2\mathbf{s}_{it}^g)^T$. This is the same type of instrument used in Altmann (2022). What makes this instrumental variable procedure even easier is that our first stage is actually known, and given by the structure of the model.

Estimation is done using a control function approach. The basic idea is that we have a regression model along the following lines:

$$\begin{aligned} y_t &= \mathbf{x}_t^T\beta + u_t & & u_t|\mathbf{v}_t \sim N(\mathbf{v}_t\rho, \sigma^2) \\ & & \& & \\ \mathbf{x}_t &= \mathbf{z}_t + \mathbf{v}_t & & u_t|\mathbf{z}_t \sim N(0, \sigma^2) \end{aligned} \tag{19}$$

This is a standard case of endogeneity with an available instrument, except with a known first stage. In this setting β can be estimated consistently using the regression equation:

$$y_t = \mathbf{x}_t^T\beta + \mathbf{v}_t^T\rho + e_t$$

Because $y_t|\mathbf{x}_t, \mathbf{v}_t \sim N(\mathbf{x}_t^T\beta + \mathbf{v}_t^T\rho, \sigma^2)$. In my setting we specify the observation equation as we did previously, but include $\mathbf{z}_{tl}^g(2\sum_{m\neq l}\Gamma_m(b_{itm})\mathbf{z}_{tm}^g)^T$ as an additional regressor. The coefficient of this regressor is essentially an estimate of the endogeneity. I specify weak normal priors for $\boldsymbol{\rho}_i$.

In practice, the endogeneity is unlikely to be linear. However this remains a useful starting point for two reasons. First, even when the endogeneity is non-linear, so that $E[u_t|\mathbf{v}_t] = f(\mathbf{v}_t)$, the posterior distribution for β, marginalised over ρ, remains correct. This is because, conditional on $\mathbf{v}_t$, $\mathbf{x}_t$ is independent of any non-linear functions of $\mathbf{v}_t$ that remain in the error term. Second, if there are non-linearities and the effect of endogeneity on $\mathbf{x}_t$ is large, the linear term should still pick up evidence of endogeneity.

Estimation remains as it was in the main text, except that before sampling ψ_i parameters, I form the conditional joint posterior distribution for (ψ_i, ρ_i), before marginalising over ρ_i, and sampling ψ_i from the marginal posterior for ψ_i.

In Figure 32 panel (B) I present posterior means of ρ_i across food banks. Estimates are presented on the same scale as the results from Figure 9. Only 16% of estimates are individually significant at the 5% significance level. Furthermore, the estimated magnitudes of the bias are small - between 2% and 8% of the relevant entries of Ψ_i.

J.2.3 Normal-Inverse-Gamma Idiosyncratic Payoff

In this appendix I relax the assumption that the lot specific idiosyncratic terms υ_{ilt} are normally distributed. Instead, I allow for the possibility that they take a normal inverse-gamma distribution:

$$\upsilon_{ilt} \sim N(0, \sigma_l^2 U_{ilt}) \qquad \text{where} \qquad U_{ilt} \sim IG(\alpha, \alpha)$$

This distribution has heavier tails than the normal distribution. The distribution can be interpreted as taking into account unobserved variation in lot specific attributes that affect the variance of the payoff. For example, I do not take into account different varieties of apples. The quality of some types of apples may be significantly more variable than other.

The procedure described shortly can also be extended to net donations $\mathbf{x}$. However the assumption that net donations are normally distributed is significantly more reasonable, given that these

Figure 32: Robustness: Stage 2 (1)

(A) (B)

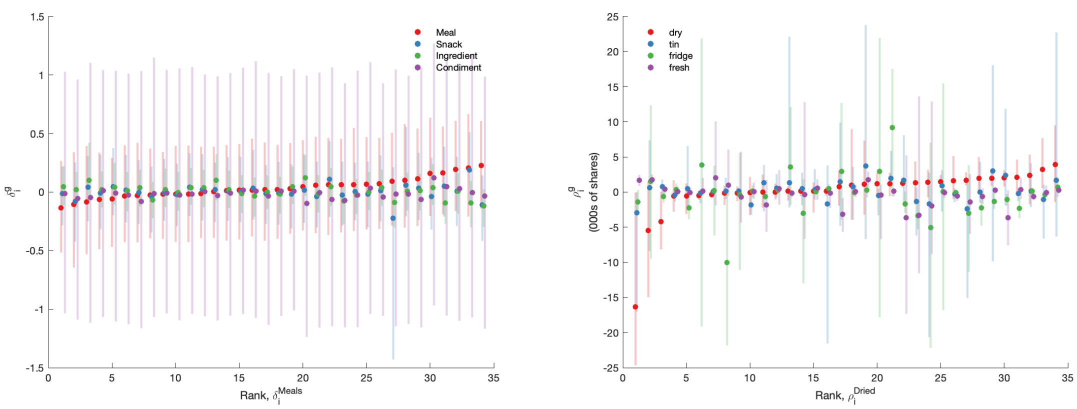

Note: The figure plots deviation statistics from two robustness checks. Panel (A) plots posterior means and 95% credible intervals of the coefficients on the demand indices across food usage types. To interpret the scale of the coefficients, if $\delta^g = 1$ this means that a one unit increase in d_t^g (associated with a one share expected increase in the winning bids) is associated with a $\Phi \mathbf{z}_{tl}^h$ unit increase in bids. We expect δs to be small and positive. Panel (B) plots posterior means and 95% credible intervals of the ρ parameters, which can be interpreted as estimates of the bias in the Ψ_i parameters caused by endogeneity of the observation equation. The estimated ρs are presented on the same scale as estimates of Ψ presented in Figure 9.

net donations are the sum of many local donations and many loads sent out to food pantries.

Sampling υ_{ilt}s from their censored distributions and sampling $\mathbf{s}_{it}^g$ using the Carter-Kohn algorithm both rest strongly on the normal distribution assumption. In particular, the posterior distributions of $\mathbf{s}_{it}^g$ is intractable when υ_{ilt} is non-normal. However, conditional on $\{U_{ilt}\}_{ILT}$, we have normality again. Therefore I use an additional data augmentation step in which I sample $\{U_{ilt}\}_{ITL}$ conditional on $\{\upsilon_{ilt}\}_{ITL}$, α, and σ_l^2. Given known difficulties associated with estimating the shape parameters of these types of distributions I fix $\alpha = 5$, ensuring the first four moments of the distribution exist. The σ_l parameters are just a rescaling of those presented in the main text.

This data augmentation step is performed using the following conditional posterior:[72]

$$U_{ilt} | \upsilon_{ilt}, \sigma_l^2, \alpha \sim \textit{scaled-inv-}\chi^2(2\alpha + 1, \frac{2\alpha + \frac{\upsilon_{ilt}^2}{\sigma_l^2}}{2\alpha + 1})$$

In Figure 33 I present the results of Wald tests from different groups of parameters, considering how this alteration to the model changes the estimated model parameters. I can reject that the

[72]This result comes from the fact that the marginal distribution of the normal inverse-gamma distribution for υ_{ilt} only is a t-distribution with 2α degrees of freedom. I then use the standard result that t-distributions can be written as a scale mixture of normals, for which the conditional posterior is readily available.

parameters have the same posterior means for an overall Wald test. However it is useful to see where the main differences are coming from. I find that we can only reject the null hypothesis that posterior means are equal for the lot specific variance, and marginal value of wealth parameters.[73] Estimated variance parameters are on average lower than those form the baseline specification, and λ_i parameters higher. This is because this specification does not need an excessively large variance in order to rationalise the heavy right tail of bids. In future I will investigate how these differences lead to different welfare effects from my simulations.

Figure 33: Robustness tests, differences in posterior means

Alternate	Statistic	Parameters						
Model		Ψ_i	Φ	λ_i	σ_l	Distance	$\boldsymbol{\mu}_i$	$\boldsymbol{\Sigma}_i$
NIG υ_{ilt}	χ^2	570	10.2	85.2	97.1	4.24	1.91	37.7
	p-val	0.0344	1	1.7e-06	0.00172	1	1	1
50 auctions	χ^2	2490	31.8	51.9	314	8.52	0.761	19
	p-val	0	1	0.0192	0	1	1	1

Note: This table presents test statistics and p-values from Wald tests for differences in posterior means across several alternate model specifications. Tests are performed separately across groups of parameters. The test statistic has an asymptotic chi^2 distribution with degrees of freedom given by the number of parameters of that type.

J.2.4 Accounting for Additional Non-Entered Lots

In the main text I estimated the model using only 25 unique auctions held each day. Although no food bank was every observed placing more than 25 bids, on around 50% of days there were more than 25 unique auctions. As many as 87 auctions unique auctions were observed being held simultaneously in my data. This simplification risks introducing bias as it does not recognise food banks' decisions not to bid on these additional lots. This bias is similar to the possible bias in a standard tobit model from simply dropping half the censored observations. However, given that food banks rarely place more than 10 bids each period, so that I am already taking into account their decision not to bid on 15 lots, these additional observations are unlikely to yield much additional information. The simplification was made to speed up the convergence of my Gibbs Sampler, as the large degree of censoring will typically harm this.

To investigate robustness to this decision I estimate the model using data on 50 auctions held each day. To get around convergence problems I begin estimation from the final iteration(s) of the

[73] I can also reject the hypothesis for the Ψ_i parameters at the 5% significance level, however this appears to come from a small number of food banks, and preliminary investigation suggests it may relate to non-convergence of their parameters.

main model. I only observe days with more than 50 auctions on 5% of days. If I find that my results are generally robust to doubling the number of auctions considered each day, it is unlikely that including the remaining auctions will change the results either.

In Figure 33 I present the results of Wald tests from different groups of parameters, considering how this alteration to the model changes the estimated model parameters. I can reject that the parameters have the same posterior means for an overall Wald test. I can also reject that posterior means are equal for the pseudo-static payoff function parameters Ψ_i, marginal values of wealth λ_i, and the lot specific variances. I need a larger variance to rationalise the lower probability of bidding on any given auction. For the Ψ_i parameters, as expected, a number of them failed to converge properly. This warrants additional investigation in future.

J.3 Third Stage

J.3.1 Quadratic Approximation

After evaluating the ex-ante value function across a 3.2×10^6 grid of states, I take a quadratic approximation across this grid, weighted by each states' estimated relevance. This is necessary because I must evaluate the ex-ante value function for each food bank and each draw, and I am unable to save all 34×1000 grids. A legitimate concern is whether this approximation is accurate. Figure 34 panel (A) presents a histogram of the R^2s from forming this approximation. 100% of these value lie between 0.99 and 1. The fit is strong because of the quadratic term which appears in equation 12.

As an additional robustness test I consider whether fitting higher order polynomials improves the fit significantly. I consider up to a 4th order polynomial. Figure 34 panel (B) plots the results from this analysis. None of these test statistics exceeds 2, far below the critical values. This is evidence that including higher order polynomials does not yield better fit than using a quadratic approximation.

J.3.2 No Information Pooling

As discussed in Appendix H.3.3 I pool information across food banks when evaluating the ex-ante value functions. This may introduce bias, by drawing food banks' estimated flow payoffs together. I already use a first-order correction for this bias, however as an additional robust test I consider how estimated parameters vary when I do not use this information pooling. Figure 35 panel (A) displays the results from this analysis, presenting estimates of the marginal flow payoff across food banks just as in Figure 9. My estimates remain in line with the previous results.

J.3.3 Sampling Variation in Means

As discussed in Appendix H.3.2 I do not take into account sampling variation in my finite sample evaluations of the expectation terms in Proposition 1. Therefore I likely underestimate the posterior variance of the flow payoffs. In this appendix I consider how results change when I employ a

Figure 34: Robustness: Stage 1

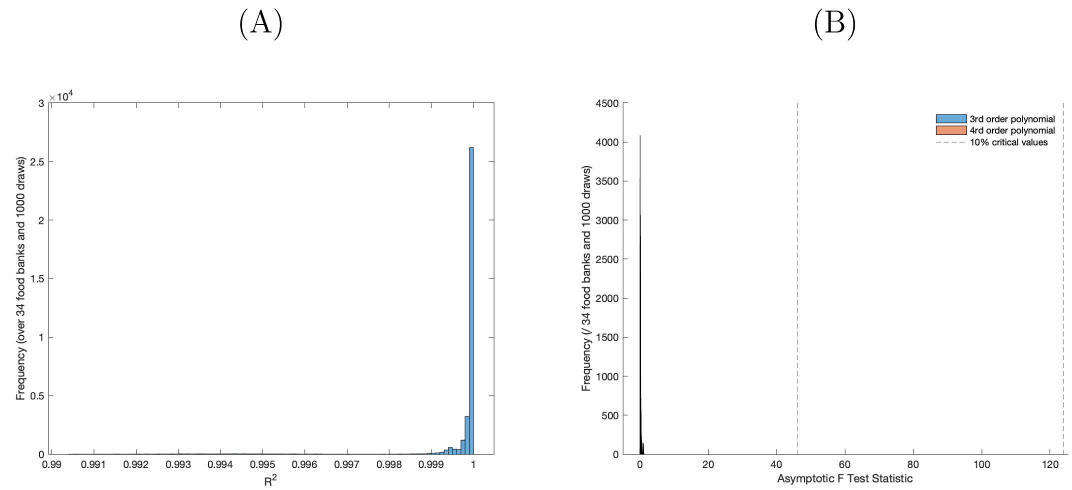

Note: The figure considers the accuracy of my quadratic approximation. Panel (A) presents R^2 statistics from the least squares quadratic fit, while panel (B) considers third and fourth order polynomials, presenting F statistics for whether the additional parameters are significant at the 10% level.

bootstrap resampling procedure on the estimated expectations. When estimating these means, for each draw from my posterior distribution of second stage parameters I randomly draw (with replacement) the time periods used to evaluate the means. This procedure should overestimate the posterior variance, as it does not account for covariance between the sampled parameters and the sampled time periods. Figure 35 panel (B) displays the results from this analysis, presenting estimates of the marginal flow payoff across food banks just as in Figure 9. Credible intervals become somewhat larger, but not by much, and the plot remains similar. This is unsurprising given the long panel, and that I am pooling information across food banks.

K Simulation Details

In this Appendix I describe how the counterfactual simulations are performed. Appendix K.1 focuses on the Old System, detailing how I numerically solve for the equilibrium value function. Appendix K.2 outlines how I simulate the Choice System. Appendices K.3 - K.6 detail how I simulate equilibrium allocations under the remaining counterfactuals. The additional counterfactuals all use the same basic continuous time set up as the Old System. They only differ in the offer and acceptance processes. Due to computational constraints equilibria are evaluated only at the posterior means of my parameter draws. This simplification is unlikely to have a major impact on my results as equilibrium accept/reject decisions are much more strongly determined by the flow payoffs than the continuation values.

For all the counterfactuals there is a risk that stocks trend downwards for some food banks, and

Figure 35: Robustness: Stage 1

(A)

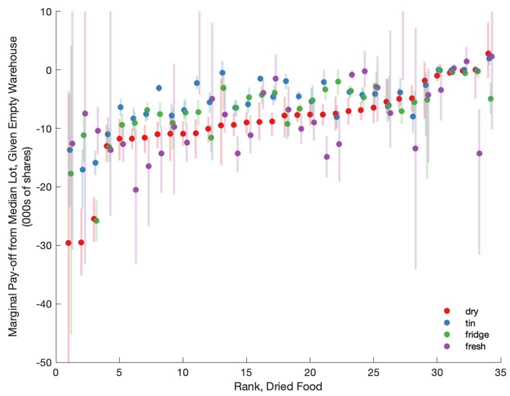

(B)

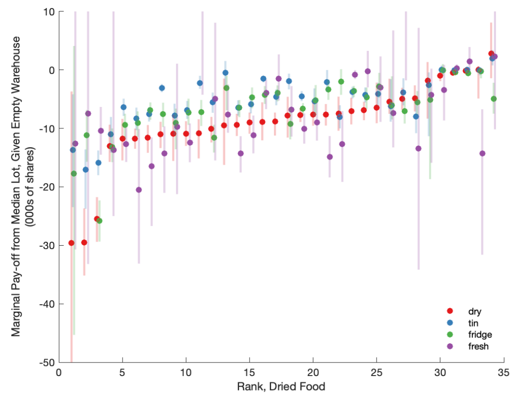

Note: The figure plots two sets of estimates for the marginal flow payoff, just as in Figure 9 panel (B). Here, panel (A) estimates the marginal flow payoff without pooling information on bidding behaviour across food banks, while panel (B) estimates this object taking into account sampling variation in the finite sample evaluations of the expectation terms in Proposition 1.

may trend upwards for others (particularly in the random allocation). If this case I would have to extrapolate estimated flow-payoffs into regions of the state space that were never visited under the Choice System. This is predominantly a problem under the random allocation, the limited offer Old System, and the limited offer Closest mechanism. To alleviate these concerns I make two simplifications. First, if stocks exceed the highest sampled stock for a given food bank for a particular type of food, the food bank turns down all subsequent offers for that type of food until stocks return to levels within the sampled space. Likewise if stocks stray below the minimum sampled stock, all additional loads are unconditionally accepted. Second, whenever stocks exceed the maximum sampled stock (or stray below the minimum) by more than one standard deviation of sampled stocks, I do not extrapolate the flow payoff to that state. Instead, I fix the flow payoff to the minimum flow payoff from across the sampled states. When I do not use these simplifications, payoffs are significantly lower under these three counterfactual mechanisms.

K.1 Old System

In this Appendix I detail how the simulations of the Old System are performed. I use the same procedure for both the Old System with only 10 offers, and the Old System with food offered to every food bank. I treat time as continuous, and each day is of length 1. This means I assume local donations and offers of food from Feeding America are received continuously during the day. However, to ensure that results are easily comparable across the Choice System and Old System simulations, when evaluating welfare I treat local donations as only altering stocks at the end of the day. Likewise, that flow payoffs only accrue at the end of the period. In evaluating the equilibrium value function, however, I treat both these objects as continuous.

K.1.1 Set Up

Arrivals

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

Priorities

Food is offered to whichever food bank is at the head of a queue. A food bank's position in the queue is given by their rank in a priority ordering. The priority ordering at time t is given by the difference between the total amount of the food bank has been offered up to t, and that food bank's target amount at t. Food bank i's target amount at time t is given by $\frac{GF_i}{\sum_j GF_j} \times$ The total amount of food allocated up to time t.

Because the ordering is a function of the amount of food offered to food bank i, not the amount of food actually allocated to them, once a food bank is offered a load, their new priority is independent of whether they accept or reject the load.

I set initial priorities equal to long-run average priorities, with minor perturbation to ensures initial priorities differ across simulation draws. I also drop results from the first 100 days in my sample period, reducing the dependence of my results on the initial priorities.

Net Local Donations

Estimated local donations arrive at discrete intervals, however I must translate this into continuous time. I assume that local donations arrive at exogenous Poisson rate $\mathbf{q}_i$, with one element for each of the storage methods. Denote Net Local Donations at time t from storage method l by $\tilde{x}_{il}(t)$. This is non-zero with rate q_{il}. Conditional on being non-zero, this donation X_{ilt} is drawn from distribution F^X. The assumption I make on F is discussed shortly.

Daily local donations of type l are given by $\int_0^1 \tilde{x}_{il}(t)dt$. Given the assumptions of our model, this is expected to be normally distributed with mean μ_{il} and variance Σ_{il}. This is a sum of i.i.d. random variables, ensuring that the choice of F^X should not matter if q_{il} is sufficiently high so that I can apply the central limit theorem. In general, both F^X and q are not jointly identified from discrete data. However, the requirement that total net daily local donations has mean $\boldsymbol{\mu}_i$ and variance Σ_i, plus a functional form assumption on F^X, can be used to pin down F^X and $\mathbf{q}_i$.

I assume that $X_{ilt} \in \left\{\underline{X}_{il}, \bar{X}_{il}\right\}$ with probabilities $1 - r_{il}$ and r_{il} respectively. Conditional on $\underline{X}_{il}, \bar{X}_{il}$ (which I discuss shortly), I set q_{il} and r_{il} to ensure that the mean and variance of daily donations equal $\boldsymbol{\mu}_i$ and Σ_i.[74]

Payoffs and States

If food bank i accepts lot l they receive lot specific flow-payoff υ_{ilt}. If they are in state $\mathbf{s}_i$, then they also receive combination specific flow-payoff $\Phi\mathbf{z}^h_{lt} + j(\mathbf{s}_i + \mathbf{z}^g_{lt})$. If they reject the lot, they only receive flow payoff $j(\mathbf{s}_i)$. Importantly, this means that at every continuous moment in time (with density zero), the food bank receives flow payoff $j(\mathbf{s}_i)$.

I discretise the individual state space in the same way as done in Section 5.4, using a grid formed of 20 evenly spaced points from each dimension of the state. This means that accepting a lot, or receiving a local donation, can only move the state in a finite number of ways, which I now discuss.

From state $\mathbf{s}_i$ if they accept lot l their stocks would increase by $\mathbf{z}^g_l$. Therefore, for each lot $\times$ state combination I find the nearest discretised state that minimises the euclidean distance to $\mathbf{s}_i + \mathbf{z}^g_l$. This allows me to define the $20^5 \times 20^5$ matrix Z_l containing a single 1 in each row (corresponding to a particular state) in the column that corresponds to this closest state from accepting lot l.

I do a similar thing for the local donations. I set $\bar{X}_{il} = -\underline{X}_{il}$ equal to the distance between my grid points, so that with rate q_{il} the food bank moves up or down a grid point, with probabilities r_{il} and $1 - r_{il}$ respectively. This allows me to define the transition matrix Q_l containing two non-zero values (r_{il} and $1 - r_{il}$) in each row, in the columns corresponding to the states one discrete notch above and below (along dimension l) of each state.

K.1.2 Equilibrium

The Agent's Problem

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

Beliefs

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

[74] I also impose that $r_{il} \in [0, 1]$. Average net daily donations are then given by $q_{il}(\underline{X}_{il} + r_{il}(\bar{X}_{il} - \underline{X}_{il}))$. The variance is given by $q^2_{il} r_{il}(1 - r_{il})(\underline{X}_{il} - \bar{X}_{il})^2$. Under the additional restriction that $\bar{X}_{il} = -\underline{X}_{il}$, it can be shown that $q_{il} = \sqrt{\frac{\Sigma_{il} + \mu^2_{il}}{-\bar{X}_{il}\underline{X}_{il}}}$ and $r_{il} = \frac{\mu_{il}/q_{il} - \underline{X}_{il}}{\bar{X}_{il} - \underline{X}_{il}}$

Equilibrium

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 their beliefs about p and f^c, and that their beliefs about p and f^c are consistent with the observed realisation of the rates at which Feeding America offers them loads.[75] As I have assumed a stationary equilibrium, I require that p and f are conditionally independent of t.

Because equilibrium value functions must be calculated over a large state space I make a number of simplifying assumptions about equilibrium beliefs. I assume food banks do not observe other food banks' stocks, nor when loads are offered to other food banks (hence aggregate supply is also unobserved). They only observe when Feeding America offers them a load. I therefore assume the only objects used to form their beliefs are $\mathbf{s}_i$, their own (relative) Goal Factor, and the time since they were last offered a load τ. I assume that f^c, the distribution of lot characteristics, depends only on GF_i. I assume that the offer rate p also depends only on GF_i. In principle I could allow p to depend on τ, however for simplicity I assume it does not.[76] I will consider this dependence as a robustness check in future. I also assume food banks beliefs do not change conditional on the previous history of offers. That is, food banks do not infer from frequent offers that offers will be more frequent in future. I will allow for this dependence in a future robustness check.

I assume parametric forms for both these objects. Broadly, for f_i^c, I split lots into the same 60 discrete category combinations used for the lot specific variances σ_l, detailed in Appendix H.2. Therefore f_i^c can be interpreted as conditional probabilities. Then, conditional on the category combination, I assume food banks believe that, in equilibrium, the distance between the lot and a given food bank is normally distributed with some mean and variance. I also assume that, conditional on category combination, food banks believe $\Phi\mathbf{z}^h$ is also normally distributed. I then assume a multinomial logit form for the probabilities of each category combination being offered, with probabilities allowed to vary with GF_i. Finally, I assume that p_i is (logit) linear in GF_i.

K.1.3 The Optimal Control Problem

Under the assumptions outlined above, we can write the value function as a function of $\mathbf{s}_i$ and τ, and GF_i. It does not depend on t due to the stationarity assumption. For numerical convenience I

[75] One problem with assuming a stationary equilibrium is that food banks' stocks may not be stationary under a counterfactual. If $\mu_i < 0$ then all else equal their stocks trend down over time. I previously made the assumption that food banks are able to consume enough food from the Choice System to ensure their unobserved stock process remains stationary. Priors suggest the Old System is worse at allowing food banks to access the food they need. Therefore there is no guarantee their unobserved states will be stationary. However, this does not rule out the possibility of a stationary equilibrium. If stocks quickly trend down to the point where food banks accept any load they are offered, this equilibrium is still stationary. It is therefore important that I drop the first 100 days in my analysis (and equilibrium belief update), to ensure food banks are in this equilibrium.

[76] This is unlikely to be a problem, since, for most food banks, loads are offered to them so frequently it is unlikely they will ever have to wait particularly long before receiving another offer. I cannot allow f^c to depend on τ for computational reasons.

absorb GF_i into the individual specific value function. Therefore write the value function as $V_i(\tau, \mathbf{s}_i)$.

Accept/Reject decision

Food bank i, that is offered load l, accepts the load if $\upsilon_{il} + \Phi\mathbf{z}_l^h + V_i(0, \mathbf{s}_i + \mathbf{z}_l^g) \geq V_i(0, \mathbf{s}_i)$. Regardless of whether they accept or reject the load, τ resets to 0.

HJB Equation

The Hamilton-Jacobi-Bellman differential equation is given by:

$$(\rho + p_i + \sum_{l=1}^{5} q_{il})V_i(\tau, \mathbf{s}_i) = p_i \int_{\mathbf{c}_l} \max\left\{\upsilon_l + \Phi\mathbf{z}_l^h + V_i(0, \mathbf{s}_i + \mathbf{z}_l^g), V_i(0, \mathbf{s}_i)\right\} dF_i^c(\upsilon_l, \mathbf{z}_l^h, \mathbf{z}_l^g)$$
$$+ \sum_{l=1}^{5} q_{il} \int V_i(\tau, \mathbf{s}_i + X) dF^X(X_l) + j(\mathbf{s}_i) + \frac{\partial V_i(\tau, \mathbf{s}_i)}{\partial \tau} \quad (20)$$

Where ρ gives the discount rate ($= (1-\beta)/\beta$). To solve this differential equation, write V_i in vector form, stacking over all the possible individual states $\mathbf{s}_i$ (i.e. our discretised states). Also discretise the category combinations across c. The equation can then be written as:

$$(\rho + p_i + \sum_l q_{il})\mathbf{V}_i(\tau) = p_i \mathbf{H}_i(\tau) + \sum_l q_{il} Q_l \mathbf{V}_i(\tau) + \mathbf{j} + \nabla_\tau \mathbf{V}_i(\tau)$$
$$\text{Where} \quad H_i(\tau, \mathbf{s}_i) = \sum_c f_c^i E[\max\left\{\upsilon_l + \Phi\mathbf{z}_l^h + Z_c^g \mathbf{V}_i(0), V_i(0, \mathbf{s}_i)\right\} |c, \mathbf{s}_i] \quad (21)$$

Where Z_c^g gives the transition matrix defined by the pounds from a load of category combination c, and Q_l gives the transition matrix formed from the net local donations. The expectation is taken over $\upsilon_l + \Phi\mathbf{z}_l^h$. This vector differential equation does not have an analytic solution. However, recognising that $\mathbf{H}_i(\tau) = \mathbf{H}_i(0)$ it is clear that there exists a solution for $\mathbf{V}$ which is independent of τ, for which we can solve using numerical methods.

Numerical Solution

For a given $\mathbf{V}_i^k$ and beliefs (p^k, f^{ck}) I compute $\mathbf{H}_i^k$, then evaluate $\mathbf{V}_i^{k+1} = ([\rho + p_i]I + \sum_l q_{il}[I - Q_l])^{-1}(\mathbf{j} + p_i^k H_i^k)$, repeating until the magnitude of the normal vector $|\mathbf{V}_i^{k+1} - \mathbf{V}_i^k|$ is less than 1. I use these successive approximations, and switch to Newton-Kantorovich algorithm as in Rust (1987) when progress slows. Inverting the matrix Q_l is not feasible due to its size. However multiplying by Q_l is trivial given its sparsity. Therefore in evaluating this matrix inverse, and the inverse used in Newton-Kantorovich, I use the Neumann formula for matrix inversion. This procedure generally converges in around 100 iterations.

I then simulate the Old System using these value functions, before updating beliefs. I update p_i^{k+1} by running a Poisson regression on the number of offers each food bank receives each day, conditional on goal factor, and dropping the first 100 days. I then update f^{ck+1} by estimating a multinomial logit model on the category combination that composes each offer. I repeat this process until the rates and estimated probabilities change by a total less than 10^{-4}. Beliefs converge extremely quickly, generally around 4 iterations, as value functions are relatively insensitive to beliefs.

K.2 Choice System

I now detail how I simulate the Choice System. I simulate the mechanism as laid out in section 2. I use these simulations both as a comparison for my counterfactuals and to assess model fit.

K.2.1 Basics

I simulate the system once for each of the 1000 posterior parameter draws. For each of these draws I use the associated draw of net donations (given by the unobserved stocks less their observed winnings). The set of objects being allocated each period is taken as given. As 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. Instead, estimated beliefs Γ can be used as equilibrium beliefs in my simulations, and the estimated pseudo-static payoff function $k(\mathbf{s}_i, \mathbf{s}_0)$ can be used in place of flow payoffs plus the equilibrium discounted continuation value. This approach would not be valid if I wanted to consider changes to the Choice System, such as changes in food banks' budgets.

I treat maroon pounds as exogenously determined, so I continue to not model food banks decisions to sell their local donations. I also treat joint bidding as exogenous - if a bid is placed jointly I have each food bank optimally set their bid *taking as given* the other player's bid. Food banks continue to split winnings evenly. This is a major simplification, but one that I would not expect to significantly impact the results, particularly given that joint bidding makes up a small fraction of bids. I treat discriminatory auctions of multiple loads correctly, though only allow food banks to place up to 5 bids on each set of these auctions. This is done using the adjustment to payoffs and first order conditions discussed in Appendix E.

The central problem in estimating the Choice System concerns the bidding function, as this involves a complex combinatorial problem of deciding which combination of lots to bid on. It is made more complicated by the possibility of multiple local optima.

K.2.2 The Bidding Function

I describe how I find optimal bids by first discussing how I optimise bids conditional on an entry decision $\mathbf{d}_{it}$, before discussing how I find the optimal entry decisions. I then discuss how I validate any optimised bids. The key simplification I make is assuming that Ψ_i is negative definite. This would imply that the payoff function is concave in bids and entry decisions, allowing me to exploit standard results from convex optimisation. If Ψ_i is indefinite, the problem of finding an optimum is NP-hard. In practice my sampled Ψ_is for Type 1 food banks are not always negative definite, but they are in the vast majority of cases. Even when they are not negative definite, most often they are fairly 'close' to negative definite, in that the largest (positive) eigenvalue is orders of magnitude smaller than the smallest (most negative) eigenvalue.

Conditional on an entry decision $\mathbf{d}$ I use standard interior point methods to numerically maximise payoffs subject to reservation prices. I begin the maximisation process at $b^0_{ilt} = R_l + 1$. In principal there may exist multiple local optima when Ψ_i is non-negative definite. However even for simulated Ψ_i matrices that exhibited large complementarities I was unable to find evidence of multiple optima. This is likely on account of my quadratic assumption, ensuring that for any $\mathbf{b}_{-l}$ (any vector of bids excluding b_l) payoffs are quasi-concave in b_l.

Considering every permutation of entry decisions is not feasible. Instead I find initial optimum entry decisions $\mathbf{d}^*$ using a hill climbing procedure. Under this procedure there are three options for each l, either $d_l = 0$, $d_l = 1$ and $b_l = R_l$, or $d_l = 1$ and $b_l = R_l + 1$. I begin with every $d_l = 0$ and run a hill climb until reaching a local optimum. If Ψ_i is negative definite this is guaranteed to be the global optimum, and any optimal bids found after this procedure are also guaranteed to be optimal. In simulations I found numerous occasions in which there were multiple local optima, but only in cases when the simulated Ψ_i exhibited sufficiently strong complementarities. Sampled (non-negative definite) Ψ_i do occasionally admit multiple optima, particularly when the number of desirable lots is large.

Finally, I use the First Order Conditions to 'check' my optimum. For $b_{ilt} > R_l$ I check that the partial derivatives are less than 10^{-5}. For $d_{ilt} = 0$ I check that the food bank does not strictly prefer $d_{ilt} = 1$, and for $b_{ilt} = R_l$ I check the food bank does not strictly prefer either $b_{ilt} > R_l$ nor $d_{ilt} = 0$. If Ψ_i is negative definite then solutions found by hill climb followed by numerical optimisation are guaranteed to satisfy these conditions. If either of these conditions fails (and if so, it is always one of the latter two conditions) I repeat the hill climb from this point, and repeat the process until I find a solution that does satisfy these conditions. By construction, each time I repeat this process the expected payoff increases, ensuring that this process terminates in a finite number of iterations. This occurs in a limited number of cases for Type 1 food banks only.

Two final things are worth mentioning. First, the problem of multiple optima is unlikely to be significantly impacting my results as it is only a problem for certain draws of Ψ_i. Even then the model fit is typically fairly good (at least for Type 1 food banks, for whom multiple optima is a problem). Second, food bank managers likely do not solve the full combinatorial problem, and instead likely use heuristics. It is also not impossible that they also get stuck at local optima. A hill climbing heuristic is possibly even more sophisticated than they might use in practice (as it can require many iterations to find a solution). Therefore this algorithm could be considered a reasonable approximation to their behaviour.

K.3 Random Allocation

I now detail how I perform the random allocation. This mechanism is fundamentally the same as the Old System in which food is offered to every food bank, using the same queueing system. The only difference is that food banks are not given a choice to reject the lot. The only time a food bank is not offered a lot is if their stocks are above the maximum of the stocks sampled under the Choice

System. This is in order to prevent having to make large extrapolations, and keep the resulting welfare comparable to that under the Choice System and other counterfactuals.

K.4 Closest Mechanism

The closest mechanism offers food to the nearest food bank first and, in this case of the 'all offers' version, then works down food banks in order of distance. Strategically it is very similar to the Old System, except that offers (and characteristics conditional on an offer) will be much more food bank specific, rather than determined by Goal Factor. It is also much more likely that these objects do not depend on the time since the previous offer.

I follow the continuous time modelling approach used for the Old System, so that the Hamilton-Bellman-Jacobi equation remains fundamentally the same. Food banks form beliefs about the rate p_i at which they receive offers of food. Conditional on an offer, the load has characteristics $\mathbf{c}$ with probability f_i^c. As for the Old System I group food into the 60 category combinations used for the lot specific variances. Conditional on a category $\times$ food bank combination, I assume the distance between the food bank and the lot, as well as $\Phi\mathbf{z}_l^h$, is normally distributed.

As with the Old System I numerically solve the Hamilton-Bellman-Jacobi equation. For the 'single offer' Closest mechanism I can directly estimate p_i, f_i^c, and the means and variances of the normally distributed lot characteristics by considering the set of lots for which they are the closest food bank. For the 'all offer' version, given initial beliefs, I must repeatedly evaluate the value function and simulate the system until beliefs about these objects converge.

K.5 Like Mechanism

Details of the Like mechanism come from Walsh (2015). Under the Like mechanism each load is offered to every food bank simultaneously. The load is then randomly assigned, with some probability, among the food banks that 'Liked' it. This assignment probability is given by a food banks' Goal Factor, divided by the sum of Goal Factors of the food banks which 'Liked' the load.

Once more, I model this allocation problem in continuous time. I assume food banks form beliefs about the probability that food is offered p, and the characteristics of food being offered f_i^c.[77] Neither of these objects depend on the actions of other food banks. I also assume they form beliefs about the probability π_i^c of winning any given lot conditional on 'Liking' it and characteristics c. I assume these equilibrium probabilities does not depend on time, nor on other aspects of the state. While food banks have more information than under the Old System, I assume they still do not observe which food bank wins the food.[78]

[77] The i subscript here is just to recognise that food banks face different distributions of distances.

[78] It is plausible that these beliefs should depend on aggregate states, just as win probabilities did in the Choice System model. However given the small effects I found in the first stage estimation it is unlikely there would be any economically meaningful changes with the state.

Under this set up, the Hamilton-Bellman-Jacobi differential equation is given by:

$$(\rho+p+\sum_{l=1}^{5} q_{il})V_i(\tau,\mathbf{s}_i) = p\sum_c f^c\pi_i^c(\tau)E[\max\left\{\upsilon_l + \Phi\mathbf{z}_l^h + V_i(\tau,\mathbf{s}_i+\mathbf{z}_c^g) - V_i(\tau,\mathbf{s}_i), 0\right\}|c,\mathbf{s}_i] + V_i(\tau,\mathbf{s}_i)$$
$$+\sum_{l=1}^{5} q_{il}\int V_i(\tau,\mathbf{s}_i+X)dF^X(X_l) + j(\mathbf{s}_i) + \frac{\partial V_i(\tau,\mathbf{s}_i)}{\partial\tau} \quad (22)$$

As in the main text I will assume a symmetric Markov Perfect Equilibrium, so that V and π_i^c are independent of τ. I solve for equilibrium just as I did for the Old System.

K.6 Efficient Sequential Mechanism

Under the Efficient Sequential mechanism each load is allocated to the food bank with the highest value, or discarded if no food bank has a weakly positive marginal value. Value includes both flow payoffs and the continuation value. Strategically this mechanism is very similar to the Old System, except by construction food banks will always accept any load they are offered.

To evaluate the equilibrium value function I again assume food banks form beliefs about the rate p of food being donated to Feeding America, and the probabilities of loads coming from each category combination f^c. They then believe they have the highest marginal value for that lot with probability Γ. I assume this probability function takes the same generalised extreme value form as in the specification of beliefs for the Choice System. This object will be a function of their marginal value from winning the lot. As above, I again assume that food banks also form beliefs about the distribution of distances and $\Phi\mathbf{z}_l^h$ conditional on food from category combination c, which again I treat as normally distributed. The Hamilton-Bellman-Jacobi equation is given by:

$$(\rho+p+\sum_{l=1}^{5} q_{il})V_i(\tau,\mathbf{s}_i) = pH(\tau,\mathbf{s}_i) + \sum_{l=1}^{5} q_{il}\int V_i(\tau,\mathbf{s}_i+X)dF^X(X_l) + j(\mathbf{s}_i) + \frac{\partial V_i(\tau,\mathbf{s}_i)}{\partial\tau}$$
$$\text{Where} \quad H(\tau,\mathbf{s}_i) = V_i(\tau,\mathbf{s}_i) + \sum_c f^c E[\Gamma(B(\upsilon_l+\Phi\mathbf{z}_l^h,\tau,\mathbf{s}_i))B(\upsilon_l+\Phi\mathbf{z}_l^h,\tau,\mathbf{s}_i)|c,\mathbf{s}_i]$$
$$B(\upsilon_l+\Phi\mathbf{z}_l^h,\tau,\mathbf{s}_i) = \upsilon_l + \Phi\mathbf{z}_l^h + V_i(\tau,\mathbf{s}_i+\mathbf{z}_c^g) - V_i(\tau,\mathbf{s}_i) \quad (23)$$

Once more I assume a symmetric Markov Perfect Equilibrium, and fit the same empirical specification to Γ estimating food bank specific separate shape, scale, and location parameters for each category combination. Unlike in the previous mechanisms this conditional expectation does not have a closed form solution, so I find it through simulation. I then numerically solve this equation as I did for the mechanisms as described previously. I simulate the system and estimate the function Γ just as in the main text, repeating this procedure until beliefs converge.