Estimating the Demand for Used Tractors: Do Farmers Prefer Older Tractors?

James Keeler*

February 19, 2026

Abstract

There is anecdotal evidence that farmers are choosing older, less software dependent tractors to address a lack of repairability in newer tractors. Are farmers really willing to trade the productivity benefits of modern technologies and amenities for cheaper, timely, and more accessible repairs? Does a farmer's income or salience of repair costs affect their willingness to make this trade-off? I address these questions by estimating a randomized coefficients discrete choice demand model using over 97,000 observations of used tractor auction listings. I find that higher income farmers are more likely to buy older tractors, but the salience of repair costs has no effect on preferences for older tractors. I also report own-price and cross-price demand elasticities for used tractors.

*Department of Agricultural and Resource Economics, University of California, Davis

1 Introduction

If new tractors are prohibitively expensive to repair, farmers may look to the secondary market to meet their equipment needs. As the right-to-repair movement has grown in agriculture, there is anecdotal evidence that farmers are choosing older, less software dependent tractors to meet their needs (Gault, 2020). Investigating this response in the secondary market has the potential to illuminate how farmers value repairability and the factors that influence their individual preferences and relevant choices.

The right-to-repair movement has also generated interest in the development of policies directed towards improving equipment repairability on the farm. Right-to-repair policies include mandating fair access to repair manuals and replacement parts, eliminating software restricting repairs (parts pairing), and incentivizing repair friendly designs. Such policies may reduce repair costs, but they may also raise the prices of both new and used tractors if manufacturers' costs of improving repairability are passed on to farmers. These policy impacts are likely to vary across tractor characteristics, brands, and farmer demographics. Thus, forecasting and evaluating the impacts of right-to-repair policies requires in part understanding how a diverse population of farmers and operations responds to changes in the price and availability of used tractors across the space of tractor characteristics. Policy impacts could extend beyond farmer welfare. For example, improvements to repairability for new tractors with low emission engines could lead farmers to substitute away from older tractors with high emission engines.

Addressing right-to-repair policy impacts in the agricultural industry and beyond can be accomplished by estimating a demand function for farm tractors that accounts for changes in repairability. Such a function could then be used to generate predictions of farmer welfare impacts, as well as substitution effects across tractor brands and characteristics. However, compared to markets for commonplace commodities or differentiated products like food, electronics, or cars, the market for farm tractors is small and observations of new tractor sales are not readily available. Repairability is also a relatively imprecise characteristic that cannot be observed or measured directly, compounding the difficulty of estimating a relevant demand function. In this paper I attempt to address this challenge by using data from used tractor auctions to estimate market-level demand as a function of various proxies for repairability.

The random coefficients discrete choice model of demand introduced by Berry et al. (1993) (henceforth BLP) is a prevailing tool for estimating the demand for differentiated products. In addition to providing demand elasticity estimates, a major benefit of the BLP method is that it produces estimates of taste preferences for individual product characteristics. In the context of used tractors, this method permits analysis of preferences for characteristics like the age of a used tractor which may serve as observable proxies for repairability. The BLP method accounts for observed and unobserved heterogeneity in preferences for characteristics, further extending the analysis to study how preferences for tractor repairability vary with demographics like farm income or exposure to shocks in repair prices. The BLP method also provides a structural framework for studying and interpreting policy impacts and counterfactuals. For example, if right-to-repair policy reduces repair costs by 10% using BLP we could identify the subsequent change in per dollar willingness to pay for any observed tractor characteristic.

The BLP method has been applied to many goods; two of the more well known studies have examined the demand for cars and ready-to-eat cereals (Berry et al., 1993; Nevo, 2001). This paper follows Nevo (2000)'s long standing guide for implementing, estimating, and interpreting random coefficients discrete choice models for demand estimation. Conlon and Gortmaker (2020) is a more recent and expanded guide to BLP estimation that is the companion to the their PyBLP package, which this paper employs to estimate the BLP model. Berry and Haile (2021) is a excellent guide for designing and applying BLP models. Other papers specifically applied to auction data are also relevant here (Backus and Lewis, 2024). As far as the author is aware, no other studies have used the BLP method to estimate demand for agricultural capital or inputs.

Many influential economic studies over the past several decades have examined demand and supply elasticities in agricultural markets. Griliches (1959) is one, if not the first, study to estimate a demand elasticity for tractors and other agricultural inputs. Nerlove (1956) and Nerlove and Addison (1958) are also classic studies of elasticity estimation for agricultural goods and supply. Roberts and Schlenker (2013) applies more modern elasticity estimation methods with a more specific agricultural policy application.

This article makes two primary contributions to the farm input and capital demand literature: (1) estimation of own and cross price elasticities and farmer tastes for used tractors across characteristics including brand, horsepower, and era; (2) incorporation of farm income and repair cost shock observations to study their effects on farmer price sensitivity and preferences.

I find that while market shares for older used tractors are higher, this does not translate to farmers preferring them more than newer tractors, all else equal. Farmers generally prefer newer and higher horsepower John Deere tractors to older, lower horsepower, and/or Case IH tractors. However, I also find these results are sensitive to market definition. Using a market definition that is too large or small to adequately characterize the farmer's choice set threatens the internal and external validity of randomized coefficients discrete choice logit models. I discuss alternative market definitions that should be explored in future work.

2 Model

The model follows Nevo (2000) and Berry and Haile (2021) as applied to the market for used tractors. Farmers select either a single used tractor from a combination of brand, horsepower, and era characteristics, or choose an outside option. The outside options include foregoing purchase entirely, purchasing a new tractor, or purchasing a used tractor that is not represented in the combinations of characteristics. Their are two brands (John Deere, Case IH), three horsepower categories (HP < 40, $40 \leq$ HP ≤ 99, HP ≥ 100), and two era categories (1970-1999, 2000 - 2025), composing 12 unique combinations of characteristic groups. I refer to each combination as a single product or inside option.

Let $j = 0, \ldots 12$ index the finite discrete set of options with $j = 0$ representing the outside option. Each farmer i has conditional indirect utility u_{ij} for option j and their objective is to select the option that maximizes their utility. The BLP model uses the random utility specification

$$u_{ijt} = -\alpha_{it} p_{jt} + x_{jt} \beta_{it} + \xi_{jt} + \epsilon_{ijt} \tag{1}$$

for $j > 0$, and

$$u_{i0t} = \epsilon_{i0t}. \tag{2}$$

The index t denotes a market which is a specific combination of time period and geographic area, e.g. used tractors available at auction in 2023 in the US Heartland. p_{jt} is the price of option j in market t. x_{jt} is a K-dimensional row vector of characteristics for option j.[1] ξ_{jt} is an unobserved demand shock assumed to be correlated with price and ϵ_{ijt} is an i.i.d error term drawn from a standard type-1 extreme value distribution.

The terms α_{it} and β_{it} are the random coefficients that represent individual farmer preferences and price sensitivity. The β_{it} term is typically further decomposed into observed and unobserved demographics to model the distribution of farmer preferences. For observed demographics I include measures of farm income and exposure to a repair cost shock. As in Nevo (2000), unobserved demographics are i.i.d drawn from a multivariate standard normal distribution. Given the assumed distribution of ϵ_{ijt}, market share for option j in market t takes the form of a canonical multinomial logit choice probability.

[1] $K \geq 12$ since I define options according to the 12 unique combinations of characteristics. K will exceed the number of unique combinations if additional characteristics are included, e.g. nonlinear terms

3 Data

A panel of over 97,000 in-person and online auction results for used tractors from June 1, 2021 to June 1, 2024 was obtained from TractorHouse.com. Listings are typically updated by the listing owner with their final sale price after an auction ends. In practice, these auction results are intended to be used as references for pricing of similar used tractors in future auctions. There are 1,494 unique auctioneers or dealers reporting results, 26 of which provided over 500 results. The majority of listings are created by companies that primarily provide auctioneer services, but some come directly from tractor dealerships or individual farms[2]. Dealerships that obtain used tractors from trade-ins likely contract with these major auctioneers to reduce their used inventory. Estate sales are also a common source of tractors on auctioneer lots. All listings include price, listing title, model, manufacturer, horsepower category, machine location, and general description. Additional information like model year, engine hours, condition, serial number, and add-ons are reported at the discretion of the listing owner.

To fill in some of the sparse listing information, particularly model year, I obtain additional characteristic information by tractor model from TractorData.com. I match the tractor model reported in the listing with the model production period reported by TractorData and use the last year the model was in production to generate three era categories: pre-1970, 1970-1999, 2000-2025.[3] These eras were chosen to separate primarily mechanical tractors from tractors that are more dependent on electronics and software with potentially limited repairability. Pre-1970 models were omitted because they are more likely to be purchased for their value as a collectible and less likely to be used for commercial farm operations.

I also assign each used tractor to horsepower and brand categories. The horsepower categories are taken directly from TractorHouse's listings format: HP < 40, $40 \leq$ HP ≤ 99,

[2]Only the name of the listing company or owner is provided, so it's not feasible to generate statistics about listing owners. Yet, it is feasible to use these names to generate auctioneer fixed effects.

[3]For example, if a tractor model was produced from 1964-1972 I use the last year (1972) to assign the model to the 1970-1999 era

HP ≥ 100. To reduce the number of products and the subsequent dimensions of the elasticity matrix, I restrict the sample to used tractors under the two largest brands: John Deere and Case IH. Consolidation in the agricultural equipment industry over the last half-century complicates comparisons across brands and time as brands available in 1970-1999 have been absorbed into their modern counterparts. I account for this by assigning defunct or subsidiary brands to their modern owner or lineage. For example, International Harvester and New Holland tractors are assigned to the Case IH modern brand. I omit sales of all other tractor brands that are not related to John Deere or Case IH. I also omit Case IH tractors with horsepower < 40 as there are relatively few observations of them in the sample. I generate an annual price for each combination of brand, era, and horsepower category using the median price of tractors sold in the category in a year. I define years starting on June 1, so the 2024 year spans June 1 2023 to May 31 2024.

For demographics, I obtain county-level farm-related income receipts from the 2022 USDA Census of Agriculture. I also obtain county-level supply and repair expenses (excluding lubricants) from the 2017 and 2022 Census of Agriculture. After adjusting for inflation I use their ratio as a repair cost shock.[4] I include the repair shock demographic to evaluate to what extent the salience of repair costs affect tastes for old era tractors.

Market definition is one of the most critical and sensitive components of a BLP model. Market definition will determine how the shares of farmer choices are computed and which observations are included in each market. The share of each farmer choice is the number of tractors sold (in a brand/hp/era category) in a market in a given year divided by the total size of the market. The total size is not the total tractors sold in the market observed because we must account for farmers that choose the outside good. Instead, I use the total county-level inventories of tractors greater than five years old in a market from the 2022 Census of Agriculture. However, farmers often own multiple tractors and may not replace them for decades, so total inventories overestimate the number of farmers that are considering

[4] I use the ratio so I can log transform it later.

buying a used tractor. To address this I follow Nevo (2000) and use a scale factor to reduce the inventories to a level that better represents the expected market shares for used tractor purchases and preserves a healthy share of the outside good. I choose a scale factor of 0.3 to target a conservative market share for used tractor sales in a year is 3% to 5%.[5] Thus, I compute the empirical share for option j in a market t as

$$s_{jt} = \frac{\text{\# of tractors sold for brand/HP/era } j}{0.3 \times \text{total inventory of tractors } \geq \text{ 5yrs}}. \tag{3}$$

For the spatial component of market definition, I consider two delineations: (1) USDA Land Resource Regions (LRRs), and (2) 2020 USDA commuting zones (CZs) provided by Fowler (2024). LRRs separate the US into approximate broad agricultural market regions. For example, the Central Feed Grains and Livestock Region covers much of the corn-soybean rotation acreage from Eastern Nebraska to Ohio. This delineation groups farmers and tractors from similar agricultural operations and product markets. LRRs are not based on county borders. To aggregate county-level demographics I assign counties to LRRs according to which LRR covers the largest area in a county.[6]

A weakness of using LRRs for market definition is that farmers are unlikely to travel the breadth of the region just to purchase a used tractor, so an LRR based market does not represent a practical set of used tractor options. As an alternative I consider commuting zones. Commuting zones are significantly smaller than LRRs and are designed to delineate local economics by grouping rural counties based on their nearest urban economic center. Defining markets according to commuting zones may better represent the options available to farmers at nearby local auctions. However, a weakness of commuting zones is that they may cut used tractor markets too finely. A farmer on the edge of two commuting zone may look at auctions in both, or the supply of used tractors at particularly large auctions may come from multiple nearby commuting zones. By providing results for competing market

[5]The 0.3 and 3% target is just a coincidence.

[6]Each LRR is a group of USDA Land Resource Units whose borders are generally defined by state soil maps.

definitions, I intend to show how the strengths and weaknesses of each definition affect the validity and sensitivity of the demand estimation results.

For both the LRR and commuting zone market definitions I limit the sample to counties and tractors sold in the Midwest and Eastern conterminous US where the majority of tractor sales are observed. There are 13 LRRs and 438 commuting zones observed over 3 years. Hence, there are 39 and 1,314 markets using the LRR and commuting zone market definitions, respectively. Figure 1 provides a map of all tractor sales, LRRs, and commuting zones included in the sample.[7]

[7] Note that some of the LRRs split up clusters of tractor sales. This is another reason why it is important to consider different market share definitions

Figure 1: Map of Tractor Sales, LRRs, and Commuting Zones

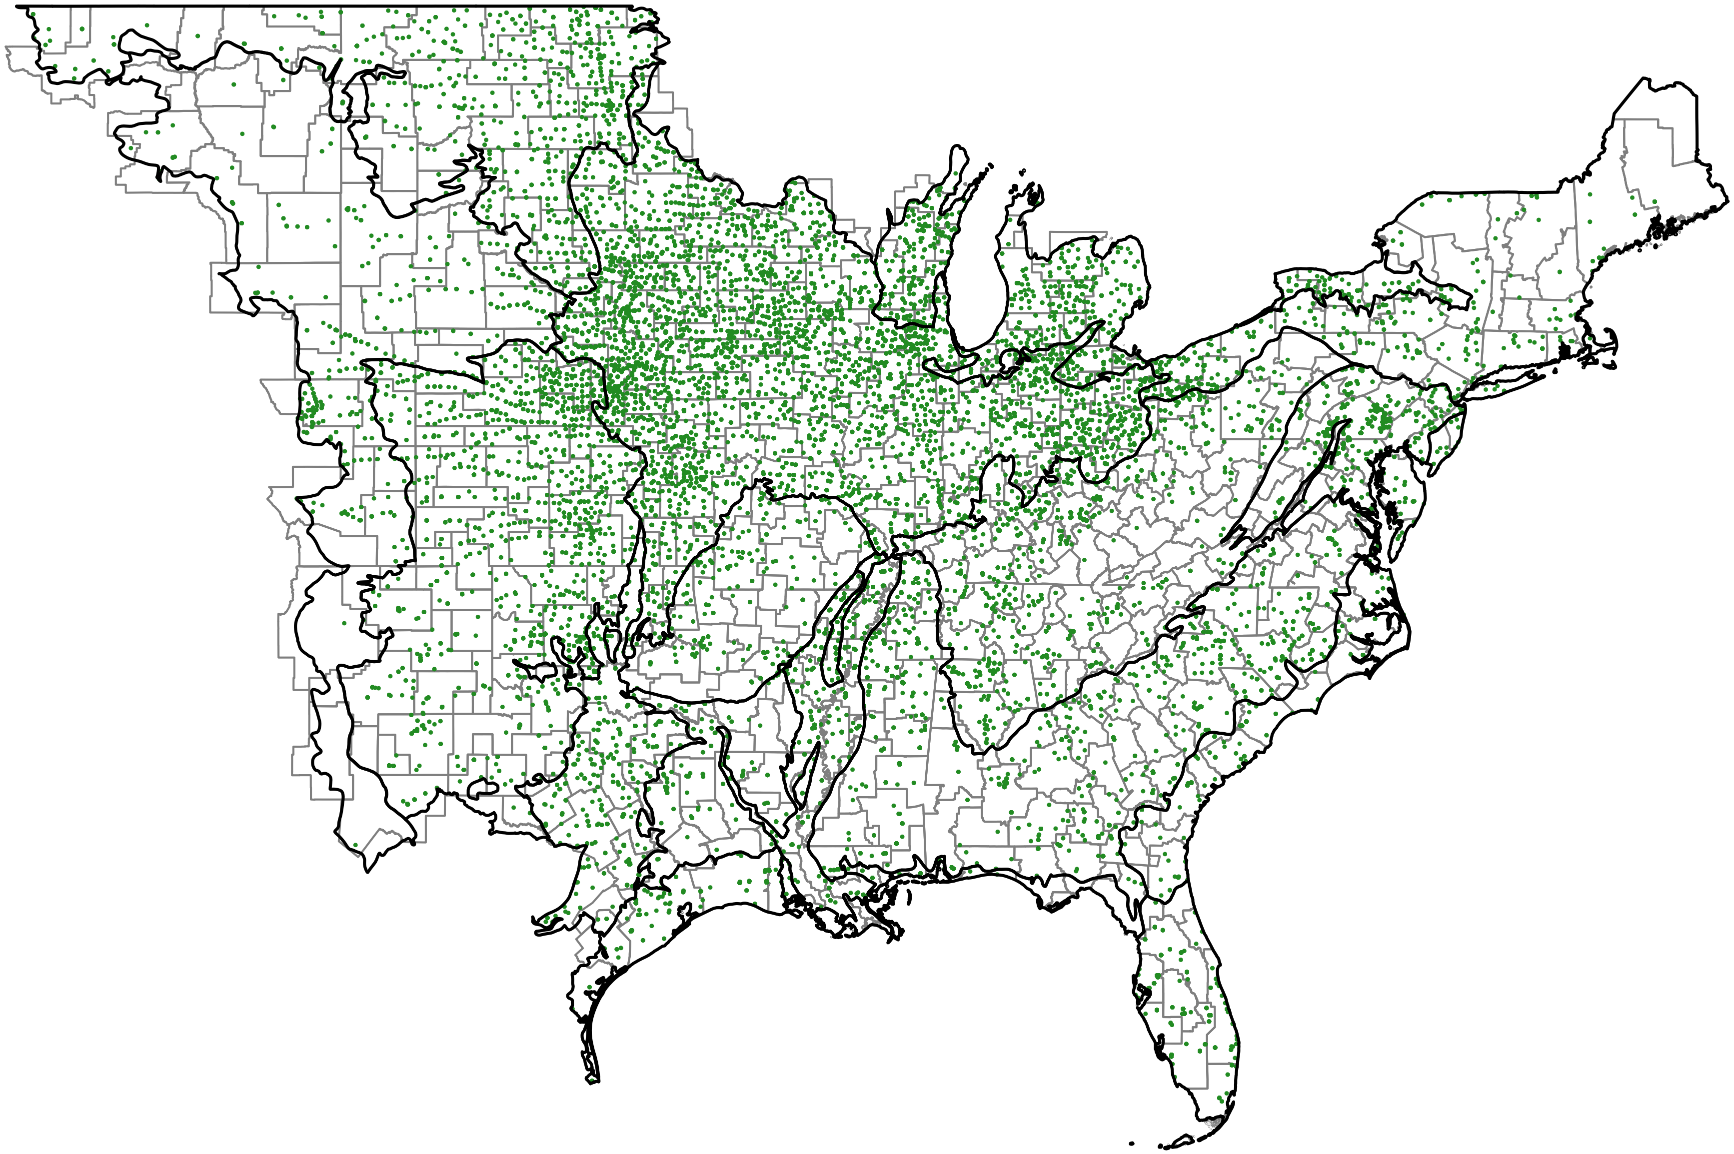

3.1 Descriptive Statistics

Tables 1 and 2 report medians and upper and lower quantiles for market share and price by tractor categories. John Deere tractors are universally sold at a premium and 1970-1999 era tractors sell for less than half the price of comparable 2000-2024 tractors. However, the 1970-1999 era tractors generally have higher market shares across all categories. Note how the some upper and lower quantiles vary considerably between the LRR and commuting zone tables. This is likely a result of the larger within market sample size for the LRR market definition.

Table 1: Market Share and Price Descriptive Statistics - LRR

			Market Share			Price		
Brand	HP	Era	0.15	Median	0.85	0.15	Median	0.85
CASE IH	100 HP or Greater	1970-1999	0.09	0.33	0.53	$8,220	$12,050	$14,337
CASE IH	100 HP or Greater	2000-2025	0.06	0.14	0.25	$34,222	$46,200	$58,298
CASE IH	40 HP to 99 HP	1970-1999	0.16	0.27	0.57	$4,000	$4,900	$5,651
CASE IH	40 HP to 99 HP	2000-2025	0.03	0.06	0.18	$11,016	$16,112	$25,232
CASE IH	Less than 40 HP	1970-1999	0.04	0.08	0.17	$2,100	$2,636	$4,567
JOHN DEERE	100 HP or Greater	1970-1999	0.21	0.50	0.91	$15,568	$17,975	$21,925
JOHN DEERE	100 HP or Greater	2000-2025	0.21	0.41	1.21	$37,690	$56,944	$75,600
JOHN DEERE	40 HP to 99 HP	1970-1999	0.14	0.21	0.48	$6,275	$7,348	$8,544
JOHN DEERE	40 HP to 99 HP	2000-2025	0.10	0.22	0.87	$16,458	$20,588	$26,388
JOHN DEERE	Less than 40 HP	1970-1999	0.03	0.06	0.17	$3,100	$4,502	$5,778
JOHN DEERE	Less than 40 HP	2000-2025	0.07	0.12	0.37	$8,509	$11,393	$14,242

Table 2: Market Share and Price Descriptive Statistics - Commuting Zones

			Market Share			Price		
Brand	HP	Era	0.15	Median	0.85	0.15	Median	0.85
CASE IH	100 HP or Greater	1970-1999	0.03	0.10	0.32	$6,500	$12,682	$24,055
CASE IH	100 HP or Greater	2000-2025	0.02	0.06	0.17	$30,000	$52,262	$95,166
CASE IH	40 HP to 99 HP	1970-1999	0.02	0.07	0.27	$3,150	$4,950	$8,000
CASE IH	40 HP to 99 HP	2000-2025	0.01	0.03	0.10	$7,212	$17,032	$31,541
CASE IH	Less than 40 HP	1970-1999	0.01	0.04	0.13	$1,350	$2,625	$4,761
JOHN DEERE	100 HP or Greater	1970-1999	0.04	0.15	0.51	$12,125	$19,159	$27,500
JOHN DEERE	100 HP or Greater	2000-2025	0.03	0.11	0.38	$33,515	$63,250	$109,460
JOHN DEERE	40 HP to 99 HP	1970-1999	0.02	0.07	0.22	$4,750	$7,500	$11,000
JOHN DEERE	40 HP to 99 HP	2000-2025	0.02	0.06	0.26	$12,859	$20,820	$31,279
JOHN DEERE	Less than 40 HP	1970-1999	0.01	0.04	0.09	$2,695	$4,410	$7,426
JOHN DEERE	Less than 40 HP	2000-2025	0.02	0.05	0.16	$6,972	$11,418	$16,380

Table 3 reports county-level descriptive statistics for farm income receipts, repair cost shocks, and the inventory of tractors at least 5 yrs old. All values are reported as per operation averages within a county. Farm income receipts are lower than other typical measures of gross cash farm income because they only include income from selling agricultural commodities and payments from government programs.[8] Most farms in the census use more than one tractor and have seen an increase in supplies and repairs from 2017 to 2022. As the gap between the 15% quantile and median for farm income receipts is significantly smaller than the gap between the median and 85% quantile, the distribution of farm income receipts is right-tailed.

[8]Farm income receipts are the total of Section 34 "Income from Farm-related sources" in the census. It is reported in USDA Quick Stats under the data item INCOME, FARM-RELATED - RECEIPTS, MEASURED IN $ / OPERATION

Table 3: County-level Per Operation 2022 Descriptive Statistics

	0.15	Median	0.85
Farm Income Receipts	\$10,620	\$23,724	\$49,734
Repair Cost Shock	-0.08	0.14	0.35
Trac Inv. $\geq$ 5 yrs old	1.81	2.38	3.25

4 Estimation

I specify the BLP model for both used tractor market definitions as follows. For linear terms I include the market price, a binary variable for John Deere branding, two categorical variables for $40 \leq \text{HP} \leq 99$ and $\text{HP} \geq 100$, and another categorical variable for 1970-1999 era tractors. I treat market price and era as random coefficients, each interacted with farm income receipts and only era interacted with the repair cost shock. All noncategorical variables are log transformed. Finally, all models include a linear constant term and its interaction with farm income. I use the PyBLP python package to perform the GMM estimation (Conlon and Gortmaker, 2020). Monte Carlo integration with 500 simulated individuals is used to compute the market share integrals. Individual farmer demographic data is generated by taking 1000 samples from the set of county-level farm income and repair costs shocks for each market with equal weights across individuals. To solve the nonlinear unconstrained GMM optimization problem, I use the BFGS algorithm with a tolerance of 0.0001 and all initial values set to 1.

A significant challenge for any demand estimation model is the endogeneity of prices. BLP models generally address this with instrumental variables for price. Two of the most common types of instruments used in BLP models are characteristic-based and Hausman instruments. Because I define options for used tractors according to groups of observed characteristics, the characteristic-based approach to instrumenting for price used by Berry et al. (1993) is not feasible. Instead, I follow Nevo (2001) by considering Hausman instruments (Hausman et al., 1994).[9] Hausman instruments assume product valuations are independent across markets, but not within markets. This implies that product prices in neighboring markets in space and/or time can be used as valid instrumental variables. The Hausman independence assumption will not hold if auction prices, demand, and/or supply are coordinated across markets. For example, if an auctioneer manages auctions across multiple markets then their marketing practices could generate correlation in tractor demand and supply across markets[10]

I use two separate specifications of Hausman instruments for the LRR and CZ market definitions. For each LRR I use prices from the other 12 LRRs as instruments. For example, the price of 2000-2025 era, HP $\geq$ 100, John Deere in 2024 for LRR M is matched with the price of the same era, horsepower, and brand category in 2024 for LRR F. For each commuting zone I use prices from the 25 nearest commuting zones as instruments in the same manner as the LRR instruments. First stage results for the LRR and commuting zone specifications are reported in Tables 4 and 5, respectively.

[9] Some alternative instruments I am exploring are buyer or seller auction premiums, national retail tractor inventories, and agricultural commodity prices.

[10] This could be addressed by including auctioneer fixed effects which are observed in the TractorHouse data.

Table 4: IV First Stage - LRR Model

	Model 1
$40 \leq$ HP ≤ 99	-0.008 (0.068)
HP ≥ 100	0.029 (0.147)
Era 1970-1999	0.073 (0.119)
Inst. 1	0.014 (0.111)
Inst. 2	-0.138 (0.076)
Inst. 3	0.297 (0.137)*
Inst. 4	0.090 (0.131)
Inst. 5	0.814 (0.167)***
Inst. 6	-0.558 (0.105)***
Inst. 7	0.485 (0.129)***
Inst. 8	0.049 (0.095)
Inst. 9	0.297 (0.116)*
Inst. 10	-0.423 (0.180)*
Inst. 11	0.057 (0.097)
Inst. 12	-0.115 (0.094)
R^2	0.926
Adj. R^2	0.924
Num. obs.	462
F statistic	373.993

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 5: IV First Stage - CZ Model

	Model 1
$40 \leq$ HP ≤ 99	0.040 (0.021)
HP ≥ 100	0.120 (0.039)**
Era 1970-1999	-0.097 (0.030)**
Inst. 1	0.074 (0.011)***
Inst. 2	0.050 (0.011)***
Inst. 3	0.049 (0.011)***
Inst. 4	0.046 (0.011)***
Inst. 5	0.058 (0.011)***
Inst. 6	0.041 (0.011)***
Inst. 7	0.041 (0.011)***
Inst. 8	0.053 (0.011)***
Inst. 9	0.025 (0.012)*
Inst. 10	0.039 (0.011)***
Inst. 11	0.038 (0.011)***
Inst. 12	0.030 (0.011)**
Inst. 13	0.044 (0.011)***
Inst. 14	0.029 (0.011)*
Inst. 15	0.020 (0.011)
Inst. 16	0.030 (0.011)**
Inst. 17	0.035 (0.011)**
Inst. 18	0.025 (0.011)*
Inst. 19	0.022 (0.011)
Inst. 20	0.026 (0.011)*
Inst. 21	0.044 (0.011)***
Inst. 22	0.031 (0.011)**
Inst. 23	0.014 (0.011)
Inst. 24	0.037 (0.011)***
Inst. 25	0.009 (0.011)
R^2	0.756
Adj. R^2	0.755
Num. obs.	8201
F statistic	903.051

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

For the LRR specification only a few instruments are individually significant, primarily the LRRs with the majority of used tractor sales. Yet, the f-statistic exceeds 10, which I take as evidence that the LRR instruments have a sufficiently strong first stage. The commuting zone specification also exhibits a strong first-stage with a larger sample size.

5 Results

Tables 6 and 7 report estimates for the LRR and commuting zone specifications, respectively. Note that since this is a random coefficients model, we have a distribution of estimates for each coefficient estimate, not just a single estimate. Hence, the tables report the mean and standard deviations of each coefficient. The estimated standard deviations are not statistically significant for the LRR and commuting zone specifications, which is evidence that estimated coefficients do not vary considerably with unobserved demographics. Since we are estimating the parameters of an indirectly utility function, the coefficients on individual variables are interpreted in units of "utils". A better approach to interpreting the results for specific characteristics is to look at estimated elasticities across characteristics and markets.

Used tractors follow the law of demand. The strength of the relationship varies with market definition and farm income. Both the LRR and commuting zone estimates for the price coefficient are negative. The LRR estimate suggests that a 1% increase in the used tractor price reduces the median (by farm income receipts) farmer's utility by approximately $21.47 - 1.88 \times \log(23,724) = 13.26$ "utils". The reductions for the 15% and 85% quantile farmer are -13.9 and -12.64, respectively. This is evidence the higher income farmers are less sensitive to prices in deciding to purchase a used tractor. The CZ price estimates are considerably smaller with a coefficient for the median farmer of only -4.48. A possible explanation for the gap in estimates between the LRR and commuting zone models is that computed market shares for the commuting zones are smaller so any changes in quantity demanded due to changes in price will be smaller. This is evident from Tables 1 and 2. The commuting zone estimates are also less statistically significant which may be due to the smaller sample size within commuting zones and greater variance in used tractor pricing at the commuting zone level.

Farmers are willing to pay more for John Deere, higher horsepower, and new tractors. The LRR estimates for the linear characteristics are all statistically significant. Farmers prefer John Deere to Case IH tractors regardless of era and horsepower. Used tractors with more than 100 horsepower exhibit the greatest preference. In contrast, farmers are willing to pay considerably less for older era tractors. The 1970-1999 coefficient for the median farmer is -14.18 in the LRR model. A weakness of the aggregated tractor data is that it omits engine hours which likely confounds estimated era effects as older tractors will have more engine hours on average. Farms may be willing to pay more for older less software-dependent tractors, but only for those with relatively few engine hours or refurbished engines. Coefficients for the commuting zone model depict similar preferences for brand, horsepower, and era characteristics, but the estimates are again weaker and less statistically significant.

Farmers with higher farm income are willing to pay more for older era tractors than those with lower income. The interaction between 1970-1999 era and farm income receipts is positive and statistically significant in both the LRR and commuting zone models. This is surprising considering that higher income farmers are more likely to be able to afford newer tractors with more modern amenities and precision technology. My hypothesis is that farmers with higher incomes may be willing to pay more only for older tractors with few engine hours which also sell for a considerably higher price. The relationship between may be concave instead of linear as assumed in the reported specification. Another possible explanation is that higher income farmer have more tractors per operation on average. They may purchase older era tractors to fill a less necessary equipment role or replace cheaper farm labor. Older era tractors may be more of a want than a need. Lower income farmers may prefer a newer used tractor because that single tractor needs to cover a greater variety of tasks on the farm.

There is no evidence that exposure to a positive repair cost shock affects willingness to pay for older era tractors. The interaction between the 1970-1999 era variable and the repair cost shock is not statistically significant in either the LRR or commuting zone models. This may be evidence that salience of equipment repair costs is not a factor in used tractor purchasing decisions, but could also be a result of the quinquennial Census of Agriculture not adequately capturing the specific cost of tractor repairs or the opportunity cost of tractor downtime.

Table 6: BLP Estimates - LRR Market Definition

Variable	Means (β's)	Standard Deviations (σ's)	Interactions with Demographic Variables: Farm Income Receipts	Interactions with Demographic Variables: Supply & Repair Expense Shock
log(Price)	-21.47	0.24	1.88	
	(6.89)	(4.77)	(0.67)	
Constant	207.10	2.55	-19.44	
	(69.58)	(3.87)	(7.03)	
John Deere	1.53			
	(0.43)			
$40 \leq$ HP ≤ 99	2.55			
	(0.69)			
HP ≥ 100	5.04			
	(1.59)			
Era: 1970 - 1999	-23.15	1.52	2.05	-2.74
	(9.61)	(4.23)	(0.87)	(20.16)
Hausman Test stat	31.02			
LM Test Stat	0.01			

Standard errors are in parentheses.

Table 7: BLP Estimates - Commuting Zone Market Definition

Variable	Means (β's)	Standard Deviations (σ's)	Interactions with Demographic Variables: Farm Income Receipts	Interactions with Demographic Variables: Supply & Repair Expense Shock
log(Price)	-7.63	0.03	0.72	
	(4.02)	(0.32)	(0.39)	
Constant	76.78	0.50	-8.14	
	(37.43)	(3.93)	(3.67)	
John Deere	0.3			
	(0.16)			
$40 \leq$ HP ≤ 99	0.54			
	(0.26)			
HP ≥ 100	1.13			
	(0.62)			
Era: 1970 - 1999	-31.38	0.47	3.07	-0.58
	(12.46)	(3.96)	(1.11)	(7.29)
Hausman Test stat	15.11			
LM Test Stat	0.01			
Standard errors are in parentheses.				

5.1 Elasticity Estimates

Tables 8 and 9 report matricies of elasticity estimates for the LRR and commuting zone models, respectively. The own-price elasticities are negative and the cross-price elasticities are positive for the LRR model, as expected for used tractors with many substitutes. In contrast, the commuting zone elasticities are inconsistent and imprecise, again likely do to the relatively small market shares for each tractor category. For this reason I focus on the LRR elasticity estimates in the analysis.

Note that demand elasticities for individual products are typically larger than elasticities for entire industries because of the availability of substitutes. For example, insulin is a classic example of a product with a highly inelastic demand. But if there are competing insulin brands or products and the price of one brand increases, then its demand may decrease as consumers switch to the competitor, implying that demand for the individual insulin product is still somewhat elastic. The estimated aggregate elasticity of demand for used tractors is a relatively inelastic -0.19 according to the LRR model. This is close to Griliches (1959)'s short-run elasticity estimate of -0.25 for tractors, but far from his long-run estimate of -1.5. We would expect the demand for all tractors to be inelastic as they are essential for modern commercial farming. However, we might expect demand for used tractors to be more elastic given the availability of new tractors as substitutes.

Used Case IH tractors exhibit generally more elastic own-price elasticities than competing John Deere tractors across all horsepower and era categories. According to the cross-price elasticity estimates, increases in Case IH prices also have a greater affect on the demand for John Deere tractors than increases in John Deere prices have on the demand for Case IH tractors. This is evidence of greater substitution to John Deere tractors, which is consistent with the estimated farmer tastes for tractor brand.

Higher horsepower tractors exhibit smaller own-price elasticities across all brands and eras. Cross-price elasticities do not exhibit any noticeable substitution patterns across horsepower categories. This can be explained by each tractor horsepower filling distinct roles on the farm. Large tractors with more than 100 horsepower are more likely used for bigger tasks like row cropping while smaller horsepower tractors are better for smaller tasks like baling small batches of hay for livestock.

Older era tractors also exhibit smaller own-price elasticities than newer tractors, though this difference is not as substantial as the differences for brand or horsepower categories. The tractor with the overall lowest own-price elasticity is a John Deere with over 100 horsepower from 1970-1999. In general, price increases in newer tractors have a larger affect on demand for older era tractors than vice versa.

Table 8: Elasticity Estimates - LRR Market Definition

			CASE IH					DEERE					
			HP $\geq$ 100		$40 \leq$ HP ≤ 99		HP <40	HP $\geq$ 100		$40 \leq$ HP ≤ 99		HP <40	
			1970-1999	2000-2024	1970-1999	2000-2024	1970-1999	1970-1999	2000-2024	1970-1999	2000-2024	1970-1999	2000-2024
CASE IH	HP $\geq$ 100	1970-1999	-17.87	0.23	1.28	0.21	0.37	1.53	0.80	1.05	0.77	0.29	0.41
		2000-2024	0.50	-18.22	0.51	0.39	0.13	0.69	1.72	0.45	1.47	0.11	0.73
	$40 \leq$ HP ≤ 99	1970-1999	1.28	0.24	-23.24	0.36	0.84	1.43	0.78	1.436	1.15	0.62	0.866
		2000-2024	0.62	0.59	1.17	-27.37	0.57	0.71	1.85	0.73	2.63	0.30	1.83
	HP <40	1970-1999	1.22	0.211	2.96	0.55	-27.11	1.2	0.63	1.60	1.43	0.82	1.26
DEERE	HP $\geq$ 100	1970-1990	0.96	0.20	0.77	0.13	0.18	-13.13	0.71	0.76	0.54	0.16	0.23
		2000-2024	0.45	0.47	0.46	0.33	0.11	0.64	-15.48	0.41	1.31	0.10	0.61
	$40 \leq$ HP ≤ 99	1970-1999	1.24	0.25	1.66	0.28	0.53	1.57	0.87	-20.85	0.96	0.39	0.56
		2000-2024	0.62	0.61	0.92	0.69	0.33	0.76	1.00	0.66	-23.58	0.22	1.47
	HP <40	1970-1999	1.26	0.23	2.52	0.40	1.00	1.37	0.74	1.46	1.23	-25.67	1.03
		2000-2024	0.63	0.56	1.40	0.90	0.60	0.66	1.74	0.77	2.77	0.43	-28.11

Table 9: Elasticity Estimates - CZ Market Definition

			CASE IH					DEERE					
			HP ≥ 100		40 ≤ HP ≤ 99		HP < 40	HP ≥ 100		40 ≤ HP ≤ 99		HP < 40	
			1970-1999	2000-2024	1970-1999	2000-2024	1970-1999	1970-1999	2000-2024	1970-1999	2000-2024	1970-1999	2000-2024
CASE IH	HP ≥ 100	1970-1999	2.13	-0.01	-0.03	-0.04	-0.05	-0.07	-0.01	-0.02	0	0	0
		2000-2024	-0.41	-3.22	0	0	0	0	0.03	0.01	0.04	0.01	0.04
	40 ≤ HP ≤ 99	1970-1999	-0.09	0	-0.03	-0.01	-0.02	-0.03	0	-0.01	0.01	0	0.01
		2000-2024	-0.39	-0.67	-1.72	-4.07	0.01	0.01	0.06	0.02	0.09	0.02	0.09
	HP < 40	1970-1999	-0.03	-0.04	-0.22	-0.39	-0.81	-0.02	0.01	0	0.02	0	0.01
DEERE	HP ≥ 100	1970-1990	0.08	0.37	0.64	0.85	0.67	0.57	-0.02	-0.03	0	-0.01	0
		2000-2024	-0.25	-0.25	-0.42	-0.37	-0.61	-1.05	-1.84	0.01	0.04	0.01	0.04
	40 ≤ HP ≤ 99	1970-1999	-0.02	0.03	0.05	0.03	0.09	0.24	0.17	0	0	0	0.01
		2000-2024	-0.2	-0.23	-0.21	-0.49	-0.61	-0.73	-1.1	-1.8	-3.45	0.02	0.07
	HP < 40	1970-1999	-0.04	-0.02	0	0.02	0.07	-0.03	-0.05	-0.09	-0.16	-0.7	0.01
		2000-2024	-0.05	-0.11	-0.05	-0.14	-0.3	-0.47	-0.42	-1.01	-1.41	-2.88	-6.991

6 Discussion and Future Work

While the BLP parameter estimates for the used tractor market are consistent with descriptive statistics and expectations from my own priors, they alone are not directly informative for evaluating prospective policies, like the right to repair, for agricultural equipment markets. Instead, as BLP is a structural model of decision making, these estimates are most useful when applied to predict farmer choices and adaptation strategies under hypothetical market scenarios. For example, we may hypothesize that right-to-repair policies will increase the price of new tractors as manufacturer costs of increasing repairability are passed on to farmers. Alternatively, we could consider including a new tractor among farmer choices. The BLP model allows us to predict how market shares across all included tractor options would adjust under these scenarios. Unfortunately, to address these counterfactuals directly, observations of new tractor market shares and prices are required.

Though the model presented here does not include new tractors as a choice for farmers, it does provide an initial framework for evaluating such a scenario. Observations of new tractor sales are not readily available in the public domain. However, if I assume that market shares for similar used tractors are representative of market shares for new tractors, then using new tractor prices from TractorData I could add new tractors to the model. Further, I could use annual changes in national retail inventories for tractors across horsepower categories reported by the Association of Equipment Manufactures to capture variation in new tractor market shares over time.

In addition to including new tractors, the results of this initial model illuminate where it can be refined to improve eventual counterfactual analysis. First, it is evident that LRRs and commuting zones are imperfect delineations for tractor markets. LRRs are likely to large and overexaggerate the set of used tractor options. Commuting zones are almost certainly to small and muddy the spatial variation and identification of farmer choices. A potential goldilocks alternative delineation is USDA's Major Land Resource Areas (MLRAs) which are subsets of LRRs and are a more granular grouping of agricultural land uses. Using MRLAs would also refine the set of Hausman instruments by allowing for a larger or more selective set of external market prices.

Second, engine hours and auctioneer fixed effects may confound estimated preferences for tractor age and its interaction with demographics. Auctioneer information is available for all listings and including their fixed effects would account for possible selection bias in the set of tractors available at auctions. While engine hours are not available for all auction listing observations, there is a critical mass of listings with engine hours to generate estimates. An alternative specification would include engine hours as either a continuous variable or set of categorical variables to identify farmer preferences for hours. Engine hours should play an important role in used tractor purchase decisions because of their relationship to the marginal revenue product and marginal cost of tractor use. This relationship is unlikely to be linear as tractors with few hours have a greater risk of being lemons and tractors with many hours may be costlier to keep running with each additional engine hour. Including engine hours would also present the opportunity to estimate farmer perceptions of the marginal revenue product of an engine hour for a used tractor.

Finally, a cleaner definition of observed repairability and repair costs is necessary to evaluate potential right-to-repair policy questions and counterfactuals. The repair cost shock demographic in the initial model added no statistically significant information, nor identify the impact of repair costs on farmer preferences. An alternative approach to identifying repair relevant characteristics may be to mine listing descriptions and classify tractors in terms of required repairs. Then estimating willingness to pay for tractors needing repairs across tractor characteristics and farmer demographics would better illuminate preferences for repairability and generate predictions relevant for right-to-repair policy.

References

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

Berry, S. T. and Haile, P. A. (2021). Foundations of demand estimation. In *Handbook of industrial organization*, volume 4, pages 1–62. Elsevier.

Berry, S. T., Levinsohn, J. A., and Pakes, A. (1993). Automobile prices in market equilibrium: Part i and ii.

Conlon, C. and Gortmaker, J. (2020). Best practices for differentiated products demand estimation with PyBLP. *The RAND Journal of Economics*, 51(4):1108–1161.

Fowler, C. S. (2024). New commuting zone delineation for the us based on 2020 data. *Scientific Data*, 11(1):975.

Gault, M. (2020). Farmers Are Buying 40-Year-Old Tractors Because They're Actually Repairable. *VICE*.

Griliches, Z. (1959). The demand for inputs in agriculture and a derived supply elasticity. *Journal of Farm Economics*, 41(2):309–322.

Hausman, J., Leonard, G., and Zona, J. D. (1994). Competitive analysis with differenciated products. *Annales d'Economie et de Statistique*, pages 159–180.

Nerlove, M. (1956). Estimates of the elasticities of supply of selected agricultural commodities. *American Journal of Agricultural Economics*, 38(2):496–509.

Nerlove, M. and Addison, W. (1958). Statistical estimation of long-run elasticities of supply and demand. *Journal of Farm Economics*, 40(4):861–880.

Nevo, A. (2000). A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of economics & management strategy*, 9(4):513–548.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Roberts, M. J. and Schlenker, W. (2013). Identifying supply and demand elasticities of agricultural commodities: Implications for the us ethanol mandate. *American Economic Review*, 103(6):2265–2295.