

NBER WORKING PAPER SERIES

TRADE-INS AND TRANSACTION COSTS IN THE MARKET FOR USED BUSINESS JETS

Charles Hodgson

Working Paper 30497

<http://www.nber.org/papers/w30497>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

September 2022

The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Charles Hodgson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Trade-ins and Transaction Costs in the Market for Used Business Jets
Charles Hodgson
NBER Working Paper No. 30497
September 2022
JEL No. L13,L14,L20,L93

ABSTRACT

Manufacturers of durable goods can encourage consumers facing transaction costs to upgrade by accepting used units as trade-ins. These “buyback schemes” increase demand for new units, but increase the supply of used units if trade-ins are resold. In this paper, I investigate the equilibrium effects of buyback schemes in the market for business jets. I find that buyback increases demand for new units by 37% at fixed prices. However, in equilibrium this increase in sales is diminished by 38% due to substitution away from new jets among first time buyers. Because of this cannibalization, offering buyback is a dominant strategy for only 3 of the 6 major firms, with 3 firms offering buyback as a best response to other firms' policies. I show that equilibrium buyback policies can change under counterfactual market structures: a simulated merger leads to a reduction in consumer welfare, 70% of which is due to a change in buyback policy.

Charles Hodgson
Department of Economics
Yale University
37 Hillhouse Avenue
New Haven, CT 06511
and NBER
cbchodgson@gmail.com

1 Introduction

In durable goods industries, used units are often traded in decentralized secondary markets. Durable goods typically depreciate over time, resulting in gains from trade when consumers with a high willingness to pay to sell depreciated goods to consumers with lower willingness to pay. In a model of a used goods market with frictionless trade (e.g. Rust (1986)), goods are never held for more than one period due to depreciation.¹ In reality, used durables are held for extended periods of time. This behavior is typically rationalized by the presence of transaction costs. A consumer may prefer to hold a new good to the used good she currently owns but choose not to upgrade if the cost of executing the exchange is too high.

In such a market, manufacturers of new goods may have an incentive to reduce transaction costs and thereby encourage consumers to upgrade to new goods more frequently. One way of doing this is to offer the buyer the opportunity to trade in their used unit when upgrading. If consumers hold one unit at a time and it is costly to sell used units on the secondary market, then the opportunity to sell a used unit back to the manufacturer allows the consumer to avoid some of these transaction costs. This type of manufacturer policy, which I refer to as *a buyback scheme*, is used in numerous durable goods industries: manufacturers of cars, airplanes, and cell phones, for example, all provide trade-in incentives that encourage owners to sell their used units back to the manufacturer or dealer when upgrading.

Manufacturer buyback increases demand for new units by increasing the frequency with which consumers upgrade, and by encouraging upgrading consumers to substitute from buying used units to buying new units. However, if manufacturers resell the used units they receive as trade-ins, then manufacturer buyback also increases the supply of used units. In equilibrium, this will lower the price of used units and may cause customers, in particular first time buyers who cannot benefit from buyback, to substitute *away* from new units, cannibalizing the gains in manufacturer revenue from upgrading consumers.

The firm's decision to offer buyback depends on the extent to which the benefits from directing trade towards their own new units outweighs the costs of increasing the supply of used goods traded in the secondary market². Additionally, in an oligopolistic market, a firm's optimal buyback policy depends on the policies of its competitors. By accepting trade-ins when its competitors do not, a manufacturer can encourage upgrading consumers to substitute away from its competitors' products. Alternatively, it may be that the cannibalization effect of buyback is sufficiently large that

¹This is true in a simple model of demand. In particular, suppose utility is quasi-linear in price. If the flow utility from new and used goods are (u^N, u^U) and the prices are (p^N, p^U) . Suppose a consumer prefers a new to a used good, $u^N - p^N > u^U - p^U$. If this consumer held a used good then, absent transaction costs, they would choose to sell it and upgrade to a new good because $u^N - p^N + p^U > u^U$. In more complex settings, for example when prices or the set of available products are changing over time, this revealed preference argument may not apply.

²Chen et al. (2013) refer to these two effects on manufacturer revenue as the "allocative effect" and the "substitution effect".

offering buyback is only optimal for a firm when its competitors also participate in the secondary market. In equilibrium, manufacturers might offer buyback because doing so is a best response to other manufacturers' policies. Indeed, all manufacturers might offer buyback in equilibrium even though they would all have higher profits if they jointly agreed not to accept trade-ins. The equilibrium buyback policies in a particular market therefore depend on market structure and the extent of demand substitution between different manufacturers' products, as well as the extent of substitution between used and new units.

In this paper, I focus on a particular industry in which buyback schemes are common, the market for business jets. Business jets are long-lived durable goods produced by a small number of manufacturers with an active secondary market. For jet owners, selling a used jet involves significant transaction costs. The market for any particular jet model is thin, and finding a buyer typically requires paying for the services of an aircraft broker. These transaction costs may give manufacturers an incentive to buy back used jets from upgrading consumers. Indeed all major manufacturers participate in the secondary market by accepting same-brand used units as trade-ins and reselling them. Using data on all transactions in the new and used business jet market between 1961 and 2000, I estimate a model of jet demand which measures the size of transaction costs and the reduction in transaction costs that can be attributed to manufacturer buyback. I use the estimated model to explain the ubiquity of manufacturer buyback by simulating market equilibrium and computing manufacturer revenue under different combinations of buyback policies.

I measure the average transaction cost paid by upgrading consumers to be \$1.8 million, or approximately 27% of the average jet price. I find that manufacturer buyback schemes eliminate between 6.7% and 11.1% of these transaction costs. At fixed prices, removing buyback from all manufacturers decreases the number of new jets bought as upgrades over the 20 year period by 442.4, or 27%. To evaluate the extent to which this increase in demand for upgrades is cannibalized by substitution away from new units among first time buyers, I simulate a counterfactual equilibrium in which I allow used and new jet prices to adjust. The increase in used jet supply due to buyback reduces the average price of used units by 3.5% relative to the no-buyback equilibrium prices. The resulting reduction in quantity demanded for new jets among first time buyers is 38% of the increase in quantity demanded among upgraders. The size of this effect depends on the substitution between used and new jets among first time buyers and upgrading consumers. I show that repeating this exercise under the assumption of no heterogeneity in consumer preferences reduces this measure of revenue cannibalization to 2%.

To investigate the firm's decision of whether to offer buyback, I compute threshold per-unit buyback cost ranges under which each firm's buyback policy is a dominant strategy, and (higher) cost ranges under which operating buyback is a best response to other firms' policies. If the per-unit cost of buyback to the firm is equal to the reduction in transaction costs faced by consumers, then for three of the six major manufacturers, operating buyback is a dominant strategy. For the three

other manufacturers, Cessna, IAI, and Raytheon, operating buyback is a best response to other firm’s policies, but each firm would be better off if no firms offered buyback. In the case of Cessna, I show that this difference in the incentive to engage with the secondary market is due to the close substitutability of Cessna’s new jets with used jets.

Finally, I show how equilibrium buyback policies can change under counterfactual market structures. I simulate a merger of Bombardier’s small Jet business with Cessna, creating a new firm with a dominant position in the small jet sector, and show that the merged firm will choose not to offer buyback in equilibrium. Cessna’s and Bombardier’s small jets are close substitutes, and the merged firm therefore suffers a substantially smaller loss in profits due to cross-manufacturer substitution when it removes buyback than the unmerged firms. The merger simulation allows an evaluation of the importance of equilibrium buyback policies to consumer welfare. If buyback remained in place, the merger would reduce consumer welfare by \$600 million due to higher prices. The removal of buyback by the merged firm leads to a further reduction in consumer welfare of \$1.4 billion. 70% of the total loss in consumer welfare comes from the removal of buyback in equilibrium, rather than higher prices.

Together, these results show that durable goods manufacturers’ engagement with the secondary market depends sensitively on both demand substitution patterns and the market structure. Realistic changes in market structure due to mergers, entry, or exit, can lead to changes in equilibrium buyback policies with significant impacts of consumer welfare. In particular, an analysis of the simulated merger that did not take into account changes in equilibrium buyback would underestimate the effect of the merger on consumer welfare by 70%. This suggests that changes in firms’ engagement with secondary markets can be of first order importance to merger analysis in durable goods markets.

1.1 Related Literature

This paper contributes to an existing, largely theoretical, literature on the interaction of manufacturers with secondary markets. Fudenberg and Tirole (1998) show that it can be optimal for the manufacturer to offer upgraders a lower price than first time buyers, and to buy back and destroy used units in order to maintain high resale prices. Rao et al. (2009) motivate the role of trade-ins in a durable goods industry as a solution to the “lemons problem”. In their model trade-in incentives encourage consumers who own high-quality used goods to upgrade rather than hold, thus increasing the average quality of used goods on the secondary market. Unlike this paper, both of these studies assume a frictionless secondary market in which trade-in incentives have no effect on the *supply* of used units.

Closet in spirit to this paper are Hendel and Lizzeri (1999) and Chen et al. (2013). Hendel and Lizzeri (1999) identify the manufacturer’s key tradeoff in allowing trade in a secondary market:

although used units are a substitute for new units, a liquid used market allows consumers who prefer new units to upgrade more frequently. In their model, a monopolist would not want to close the secondary market entirely. They speculate that this result rationalizes the existence of manufacturer policies that facilitate trade in secondary markets, including manufacturers buying back and reselling used goods.³

Chen et al. (2013) calibrate a dynamic model of demand for new and used goods to the market for cars to quantify the effects of closing the secondary market on manufacturer revenue. They show that whether or not opening the secondary market increases manufacturer profits depends on the heterogeneity of consumer preferences and the depreciation rate of the good. The current paper advances this literature by measuring the effect of actual manufacturer policies that increase liquidity in the used market on manufacturer revenue in equilibrium. Crucially, this paper seeks to explain individual manufacturers' observed *participation* in the secondary market in equilibrium, rather than quantifying the effect of shutting down the secondary market altogether. In particular, this paper's analysis of observed buyback policies as best responses by heterogeneous firms that could change under counterfactual market structures is new to the literature.⁴

In terms of empirical methodology, this paper builds on Schiraldi (2011). Schiraldi examines the effects of scrappage policies in the Italian used car market using a dynamic demand model in which cars depreciate and there are transaction costs which prevent owners from upgrading immediately. This paper extends Schiraldi's methodology in several ways. I allow holders of different jets face different transaction costs because of the presence of heterogeneous buyback schemes across manufacturers. To identify the heterogeneity in preferences induced by buyback, I combine aggregate market shares with transaction-level "micro-moments" along the lines of Petrin (2002). These micro-moments also allow me to estimate rich preference heterogeneity at the consumer level, and keep track of the evolution of the distribution of preferences among holders of different jets over time, similar to other dynamic demand papers including Gowrisankaran and Rysman (2012), Carranza (2010), and Esteban and Shum (2008).

I use the estimated model to perform counterfactual equilibrium simulations. I combine a model of Nash-Bertrand pricing of new jets by manufacturers with equilibrium conditions in used jet markets. Equilibrium requires all used goods markets to clear, and for each consumer type to have consistent beliefs about the inclusive value of holding each available jet model. The method for computing equilibrium discussed in this paper, which relies on inclusive value sufficiency, is an alternative to Gillingham et al.'s (2021) recent work on the full solution approach to computing equilibrium in durable goods markets with heterogeneous consumers. This approach is also related

³Hendel and Lizzeri (1999) highlight certified pre-owned cars and IBM's resale of used typewriters and computer equipment as examples of this type of policy.

⁴There is a related literature on firm behavior in markets with switching costs. For example, Chen and Sacks (2021) study firm reimbursement of switching costs, a policy that is common among wireless carriers. This policy is similar to buyback, but typically applies to consumers switching firms, rather than upgrading with the same firm, and (in the examples cited by Chen and Sacks) does not involve resale in a secondary market.

to the literature on calibrated models of equilibrium in used durable goods markets (Srolyarov (2002). Chen et al. (2013), Gavazza et al. (2014)).

This is one of the few papers to study the business jet market. Gilligan (2004) uses FAA airworthiness directives to measure uncertainty about jet quality, and finds evidence for adverse selection in the used jet market. Gavazza (2016) emphasizes the significant search frictions in the market for second hand business jets, and calibrates a model of search and bargaining in a used goods market to aggregate data on business jet transactions.

The rest of the paper proceeds as follows. In Section 2 I provide an overview of the market for business jets and outline the data. Section 3 describes the features of this data that allow me to measure the effects of manufacturer buyback on demand. Section 4 presents a model of demand for new and used jets. Section 5 describes the estimation and identification of the model, and results are described in Sections 6 and 7. Section 8 concludes.

2 Data and Setting

2.1 The Market for Business Jets

The market for used business jets is typical of a durable goods industry with active trade in used goods as well as prolonged holding. Between 1961 and 2000, the six leading business jet manufacturers - Cessna, Bombardier, Dassault, Gulfstream, IAI, and Raytheon, - sold 10,938 new jets. Over the same period, there were 40,845 sales of these manufacturers' jets on the used market. Jets are long lived and can have many owners over their lifetime - the average 1971 Cessna Citation 1, for example, had 9.67 owners between 1971 and 2000. A typical owner holds a jet for between three and four years.

There are significant costs to selling a jet on the used market. Unlike in the market for used cars, aircraft dealers (or "brokers") do not always buy used jets outright. Jet brokers are closer to real estate agents - they advertise jets and facilitate transactions, and either charge a fixed fee or take a share of the sale price in commission. Arranging a sale is complicated, and even if the seller does not use a broker, there are substantial taxes and legal fees. In addition, the small number of potential buyers and sellers for a particular model of jet means that jet markets are "thin", and there are substantial search and matching costs. These costs are highlighted by Gavazza (2016) who models the market for used business jets as an asset market with search frictions.

Manufacturer buyback policies allow used jet owners to avoid paying the transaction costs associated with selling their jet, as long as they replace it with a new jet from that manufacturer. This kind of assistance can take two forms. In one, the manufacturer will accept the used unit as a trade-in and literally "buy it back" from the owner who is upgrading to a new jet. In the other, the

manufacturer will instead act as a broker for the upgrading consumer, and facilitate the sale of the used jet to a new owner without formally taking ownership of the jet itself. Under either policy, the used jet owner is able to upgrade to a new jet without paying brokerage fees and avoiding some share of the search costs involved in finding a buyer and completing a sale. As discussed further below, the fact that not all transactions facilitated by these policies involve a transfer of ownership of the used jet back to the manufacturer is important because it limits what can be observed in ownership data.

Manufacturers understand these policies as a means of stimulating demand for new jets from upgrading consumers. In the words of one salesperson who I talked to, buyback is a “necessary evil” that manufacturers use to convince jet owners to upgrade. Used jets that are bought back by manufacturers are almost always resold since the price a jet will earn on the used market usually exceeds scrap value. Holding used jets for an extended period is costly, and the stock of used jets held by a manufacturer is frequently discussed as an important measure of firm health in industry reports and the press.⁵

An important feature of buyback in this industry is that trade-ins are overwhelmingly own-brand. That is, Bombardier typically only buys back used Bombardier jets, etc. The salesperson I talked to explained that cross-brand buybacks were a rare exception that might be allowed when dealing with “important customers”. Industry participants rationalize this feature of buyback using two arguments. First, it serves as a means of strengthening product differentiation - the value of a Cessna jet includes the option of trading in that jet for a new Cessna in future. Second, manufacturers can more easily maintain, upgrade, and market their own jets. For instance, a 1982 advertisement for Learjet in the Wall Street Journal asked “What’s the next-best thing to a factory-new Learjet? A used Learjet from that same factory.” The “own brand” feature of jet buyback programs will play an important role in the empirical analysis discussed below.

2.2 Data

The analysis uses a data set constructed from FAA registration records which record all transactions of new and used business jets registered in the United States from 1961 to 2000. An observation in the data is a change to registration record, which could be the manufacture of a jet, the sale of a jet, the retirement of a jet, etc. The data includes the date of the activity, the identity of the owner and operator, the manufacturer, model, and serial number of the jet. This data allows me

⁵For example, a 1984 article in Canada’s Globe and Mail claimed that Canadair Ltd. (Bombardier) was “renewing efforts to sell its inventory of used Challenger business jets” by upgrading them with new features before putting them on the market. Similarly, a 1995 article in Canada’s National Post described Bombardier’s decision to “write down the value of approximately 65 used business jets it received on trade-in.” More recently, Bombardier’s 2021 Annual Report predicts a positive outlook for the business based partly on “low pre-owned inventory levels”. The cost to the manufacturer of holding jets for extended periods will not be explicitly considered in this paper. I will assume that jets can be immediately resold by manufacturers on the used market at the prevailing price.

Table 1: Market Share by Manufacturer

Manufacturer	New Market Share 1961 - 2000			Used Market	
	Small	Medium	Large	Resale Ratio	Annual Used Sales
Bombardier	32%	8%	33%	23.3%	309.6
Cessna	52%	11%	0%	26.6%	392.2
Dassault	4%	22%	14%	26.6%	165.8
Gulfstream	0%	0%	54%	19.7%	88.8
IAI	0%	14%	0%	29.5%	100.0
Raytheon	10%	26%	0%	25.2%	168.6

Notes: Columns 1-3 record the market share of the top 6 manufacturer in new jet sales between 1961 and 2000 in each jet category. Column 4 records the average resale ratio between 1961 and 2000 - the share of existing units that are resold in a given year. Column 5 records the average number of used units resold in a year between 1961 and 2000.

to track jets across owners over time from manufacture to retirement, and to track owners as they buy, hold, and sell jets.

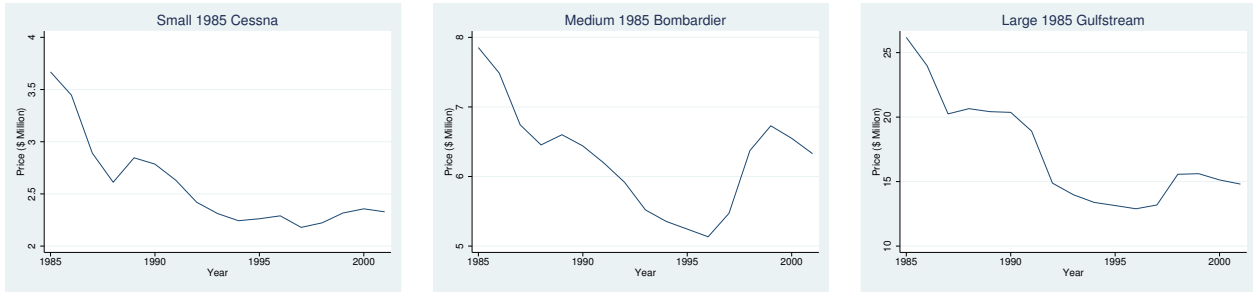
Business jets are typically marketed as belonging to one of several size classes: light, super-light, medium, medium-heavy, or heavy. For this paper, I aggregate these into three categories - small (comprising light and super-light), medium, and large (comprising medium-heavy and heavy). These categories are roughly defined by engine size, range, and capacity, as illustrated by Appendix Table A.6.⁶ Table 1 records manufacturer market shares of new jet sales for the six major manufacturers in each of the three market segments.⁷ Note that the small jet market is dominated by Cessna, the large jet market is dominated by Gulfstream, and that together, the six firms listed make up over 81% of each of the three segments. Table 1 also records the average number of used market transactions per year for each manufacturer’s jets, as well as the average annual resale rate, which is the number of used transactions divided by the stock of used aircraft for each manufacturer expressed as a percentage. Resale rates are between 19% and 30%, which is on an order similar to those recorded by Schiraldi (2011) for used cars. The resale rates indicate that there is an active market for used jets, but that jets are typically held for several years before being resold.

I supplement the registration data with prices from the *2001 Blue Book of Aircraft Values*, published by Penton Information Services. The Blue Book contains quarterly prices for new and used

⁶I obtain jet characteristics from Frawley’s (2003) *International Directory of Civil Aircraft 2003/2004*. For each jet model I record the jet’s maximum range (in km), total engine power (in kN), and maximum takeoff weight (in kg).

⁷Note that Bombardier acquired Learjet in 1990. Here, and for the rest of this paper, I record Bombardier and Learjet as the same manufacturer for the full sample (not only after 1990). Bombardier and Learjet never competed in the same market category - all small and medium jets produced by “Bombardier” are Learjet models, and Learjet never produced a large jet. Raytheon manufactures the Hawker jet series. These models were originally produced by Hawker-Siddeley until 1977, when the company was merged into British Aerospace. In 1993 the business was acquired by Raytheon. Since none of these companies ever separately competed in the market for business jets, I classify them as one manufacturer, “Raytheon”.

Figure 1: Used Jet Prices



Notes: Prices in 2009 \$. Prices for a manufacturer-size-year are averages over all model variants. Missing price data is filled in using the procedure described in the text.

jets, broken down by model and model-year. For example, an observation could be the price of a 1970 Gulfstream II in 1985 Q1. These prices are comparable to blue book prices in used car markets. They are guideline prices that should reflect the expected price for a given jet at a given time. They are not averages of actual transaction prices - in many of the quarters where a price is recorded, no aircraft of that type were actually sold. These price data were used by Gilligan (2004), and similar blue book prices have been used in comparable studies of the used car market (Schiraldi (2011), Porter and Sattler (1999)).

The raw price data series are incomplete - for example, there is no data on the price of Large Gulfstream jets manufactured in 1980 before 1985. Among all (j, t) pairs in the raw data, where j is a model (such as a Large 1980 Gulfstream) available in year t , 17.6% of prices are missing. This missing data mostly comprises older jets and earlier years in the sample - only 13.3% of observed *purchases* are of a jet with a missing price. In order to estimate the model described in Section 4, I need prices for every model that is available in every year of the sample. To fill in the missing prices, I run regressions of log price on US GDP, jet age, year fixed effects, and a time trend separately for each manufacturer-segment. I then use fitted values from these regressions to fill in the missing price observations.

Appendix Table A.6 records summary statistics about jet prices and characteristics. There is significant heterogeneity in prices across jets and over time. The average large jet is over 5 times more expensive when new than the average small jet. Prices for used medium jets drop by 5% on average in the first year, and then by additional 27% over the next four years. Figure 1 illustrates these patterns for three 1985 models. Note that prices appear to reflect differences in jet size, the age of the model, and aggregate demand shocks.

The FAA classifies jet owners into several types: dealers, manufacturers, finance companies, corporations, private owners, government, and air transport. Table 2 records the means and standard deviations of holding times, fleet sizes, and transaction probabilities by owner category.

Corporations, private owners, and air transportation companies have similar holding and purchase

Table 2: Purchase and Holding Behavior by Owner Type

	Corp.	Private	Air Transport	Manuf./Dealer	Finance	Gov.
Holding Time (Months)	47.319 (47.296)	40.323 (41.675)	47.816 (46.049)	11.597 (21.071)	40.524 (40.204)	147.594 (107.900)
Fleet Size	1.282 (0.981)	1.067 (0.302)	1.644 (1.922)	2.652 (5.106)	1.732 (3.939)	4.630 (16.419)
Purchases per Year	0.300 (0.630)	0.280 (0.540)	0.383 (0.980)	1.609 (4.189)	0.491 (2.305)	0.312 (1.737)
Sales per Year	0.248 (0.608)	0.236 (0.518)	0.317 (0.877)	2.179 (8.299)	0.415 (1.919)	0.127 (0.609)
Owner Count	14110	1518	2197	2120	2021	358
In Estimation Sample?	Yes	Yes	Yes	No	No	No

Notes: Table records means and standard deviations (in parentheses). Holding time observations are jet-owner pairs. The value of each observation is the number of months that pair is observed. Fleet size observations are owner-month pairs. The value for each observation is the number of jets owned by that owner in that month. Purchases and sales per year are owner-year observations.

patterns. On average, owners in these categories hold jets for between 3 and 4 years before selling them. The average monthly fleet size for owners in these categories is between one and two, with the modal owner holding one jet at a time. Purchase and sale rates are close to each other, suggesting fleet sizes are relatively stable, and these transaction rates are consistent with average holding times of between three and four years.

Dealers and manufacturers hold larger fleets, and hold individual jets for less than a year on average. This is consistent with their role in the market as retailers and intermediaries who sell new units and facilitate the trade of used units between consumers. Owners classified as finance companies are typically the legal owners of larger fleets of aircraft operated under lease or credit arrangements that do not involve a transfer of the title to the operator. Finally, government agencies (such as the Air Force) hold large fleets of aircraft for longer periods of time, over 12 years on average.

2.3 Estimation Sample

To use this ownership data to estimate consumer demand, I first define the relevant market and decide which owners should be counted as consumers and included in the estimation sample. Naturally, I exclude manufacturers and dealers from the set of consumers. I also exclude finance corporations because they are typically not the operators of the jets they own, and could be con-

sidered as a operating in a separate market for aircraft leases.⁸ The remaining owner types are all jet operators, and can be thought of as the “consumers” in this industry. I exclude government agencies because they make purchases through contract tendering procedures, hold very large fleets, and are probably not represented by the demand model developed in Section 4. The main estimation sample therefore includes corporations, individuals, and air transport companies as the relevant consumers.⁹

This reduces the number of owners in the sample from 22,324 to 17,825. I define a time period in the data as one calendar year. For this sample of owners, the first observed jet purchase is in 1961. The remainder of the analysis will therefore focus on the period 1961-2000. As described in Table 2, owners may hold more than one aircraft at a time. To estimate a discrete choice model of jet demand, I construct a panel in which each owner holds a *single* jet for each year. I follow the *first jet owned* by each owner and its *successors*. When I observe multiple jets held simultaneously, I split the owner into two, and the panel records a new jet owner entering and purchasing the second jet. This results in a panel of 22,793 owners. The algorithm used to construct this panel is described in detail Appendix Section A.2. The mean owner is in the sample for 5.3 years and makes 0.23 upgrade purchases. The data used for estimation includes 121,635 owner-year observations.

Application of this panel to the model discussed below implicitly assumes that the utility obtained from one jet does not depend on whether the owner holds another jet - there are no complementarities or portfolio effects in holding multiple jets. For example, this assumption rules out efficiency gains from owning multiple aircraft of the same brand rather than multiple aircraft of different brands. Appendix Table A.7 records the share of owner-years for which multiple jets are purchased for different sets of owners. Corporations, which make up 79% of owners in the estimation sample, make purchases in 27% of owner-year observations but purchase multiple jets in less than 2% of owner-year observations. The low rate of multiple jet purchases suggests that corporations do not regularly purchase “bundles” of jets, consistent with the assumption of no portfolio effects in demand.

I aggregate the available choices to the manufacturer-segment-model year level. Model year refers to the year the model was manufactured. For example, an owner making a choice in 1985 could choose to buy a large 1972 Gulfstream or a medium 1980 Cessna, both of which would be used, or a medium 1985 Bombardier, which would be new. I also collapse all manufacturers other than the top 6 into a composite “other” category. Many of these model categories contain multiple jet model variants. For example there are several variant models in the medium 1980 Cessna category.

⁸Leased aircraft may be registered under the name of the leasing organization or under the name of the operator depending on the form of the lease. In particular, if leases include a purchase option meeting certain criteria, the FAA considers this legally equivalent to ownership, and jets held under such leases are recorded as owned by the lessee in the data. Jets held under (usually shorter term) leases that do not meet these criteria are recorded as owned by the leasing company. In excluding finance companies from my analysis, I am implicitly defining the relevant market as jet “ownership” as defined by the FAA.

⁹In Appendix A.4, I report results using a sample that excludes air transport companies.

I map price data to model categories by averaging over the “true” model variants in that category. The raw price data is quarterly. Prices for a given choice in a given year are the average of all jet models in that category over all quarters in the year.¹⁰

3 Empirical Strategy

In this section I discuss the variation in the data that I use to measure the impact of manufacturer buyback schemes on new jet demand and used jet supply. I first present statistics on observed buyback transactions to illustrate the role of buyback in the business jet industry. These suggest that consumers are able to trade in a used jet with a manufacturer only when they upgrade to a new jet of the same brand. I then show that the availability of buyback appears to increase demand for new jets relative to used jets. In particular, demand for new jets is higher among upgrade buyers than first time buyers, and is higher among same-brand upgrades than among different-brand upgrades. It is this variation in demand that drives the identification of the model discussed in Section 4.

3.1 Buyback Patterns

Recall three features of manufacturer buyback in the business jet industry: manufacturers offer buyback incentives when a consumer upgrades from a used jet to a new jet, manufacturers typically restrict buyback to own-brand used jets, and manufacturers may facilitate upgrades either by purchasing the used jet or by acting as a broker. The ownership data provides some direct evidence for the first two features. The third feature limits what can be observed in the data.

To measure buyback in the ownership data, I first define an upgrade as the sale of a used jet by a consumer followed by the purchase of another (used or new) jet within 12 months. Next, I manually identify from the list of dealers and finance companies those that appear to be manufacturer owned.¹¹ For example, the largest dealer in the data is “Bombardier Aerospace Corporation.” I then identify transfers of ownership of used jets from consumers to manufacturers, manufacturer finance companies, or manufacturer owned dealers. These “observed buyback” events include all manufacturer-facilitated upgrades in which ownership of the used jet is actually transferred to the manufacturer. Importantly, observed buybacks do not include manufacturer-brokered upgrades in which ownership is transferred between consumers.

¹⁰Note that aggregation to the yearly level poses a problem for the definition of “new” jets when jets manufactured near the end of one year are sold early in the next calendar year. To deal with this, I record all year $t + 1$ sales of new jets manufactured in year t as occurring in year t . Moving these sales back in time by one year affects around 24% of new jet sales.

¹¹I researched the ownership of all dealers whose name contained certain keywords (variants of manufacturer and model names), as well as the top 100 dealers. This procedure is imperfect because many dealers have similar names, and many companies in the data are now defunct and difficult to track down.

Table 3: Buyback Patterns

Manufacturer	Share of Observed Buybacks		Share of Potential Buybacks
	Upgrades to New	Own Brand	Sold to Manufacturer
Bombardier	84%	87%	39%
Cessna	92%	93%	37%
Dassault	92%	80%	21%
Gulfstream	82%	82%	43%
IAI	80%	71%	6%
Raytheon	81%	87%	27%

Notes: The first column records used-new upgrades as a share of all upgrades in which a jet is sold from a consumer back to a manufacturer. I define an upgrade here as the sale of one jet followed by the purchase of another within 12 months. The second column records the share of jets bought by manufacturers that are of their own brand. The third column records the share of potential buybacks in which the used unit is sold to the manufacturer. potential buybacks are defined here as upgrades from used to new units of the same brand.

The first column of Table 3 records the share of new jet purchases among all observed buyback upgrades. For each of the major manufacturers, between 80% and 92% of upgrades in which the consumer's jet is sold back to the manufacturer result in the sale of a new jet, rather than a used jet. These statistics confirm that manufacturers buy back used jets from owners who wish to upgrade to new units, and it is uncommon for a buyback-facilitated upgrade to involve the sale of a used jet to the upgrading consumer. That sales to manufacturers rarely take place as part of used-used upgrades provides some assurance that manufacturers use buyback primarily to drive the sale of new jets and are not acting as general used jet dealers.

The data also indicates that manufacturers largely buy back used jets of their own brand. The second column of Table 3 records own brand jets bought back as a share of all jets bought back for each of the six major manufacturers. The shares are over 80% for five of the six major manufacturers, and as high as 93% for Cessna. That is, 93% of jets sold to Cessna are used Cessna jets. Consistent with my discussions with industry participants, this suggests that manufacturers might require trade-ins to be of their own brand or offer more favorable terms to owners trading in an own brand jet.

The third column of Table 3 records the number of *potential* buyback transactions in which the used jet is sold to the manufacturer. Based on the statistics recorded in the first two columns, I define a potential buyback as an upgrade from a used jet to a new jet of the same brand. Table 3 indicates that in around 30% of these upgrades from used to new jets, the used jet is bought by a manufacturer. The figure ranges from 6% to 43% across manufacturers. Owners are *observed* to take advantage of buyback in a large share of used-new upgrades.

These statistics provide direct evidence of the important features of buyback, but are unlikely to provide reliable measurement of *all* buyback events in which the manufacturer assumes some of the transaction costs faced by the seller. I therefore use the patterns recorded in Table 3 to

Table 4: Used and New Upgrade Shares by Same Brand

	(A) All Owners			(B) Observed Buyback Users			(C) Excluding Buyback Users		
	Used	New	Diff	Used	New	Diff	Used	New	Diff
Same Brand	0.629	0.371	0.258	0.360	0.640	-0.280	0.732	0.268	0.464
Diff. Brand	0.746	0.254	0.492	0.516	0.484	0.032	0.811	0.189	0.622
Diff	-0.117	0.117	-0.234	-0.156	0.156	-0.312	-0.079	0.079	-0.158
SE	(0.014)	(0.014)	(0.028)	(0.030)	(0.030)	(0.060)	(0.014)	(0.014)	(0.028)

Notes: Analysis performed on estimation sample. Panel A records the share of new and used jets purchased among all first time and upgrade purchases made by all consumers. Panel B repeats the calculations in Panel A including only jet owners who are observed to ever sell a jet to a manufacturer. Panel C repeats the calculation excluding jet owners that are observed to sell a jet back to a manufacturer. Standard errors of the difference in means for each column on parentheses.

motivate an identifying assumption that allows the effect of buyback on demand to be estimated from ownership data without relying on observed buyback counts.

3.2 Buyback as a Demand Shifter

To estimate the effect of buyback on demand, I assume that buyback is always available to consumers who upgrade from a used jet to a new jet of the same brand. The size of the transaction costs that are avoided by trading in a used jet rather than using an independent broker can then be identified by comparing same-brand upgrades to different-brand upgrades. Under this assumption we would expect, all else equal, the market share for new jets to be higher for same-brand upgraders, who benefit from buyback, than for different brand upgraders, who do not.

Table 4 shows that this pattern holds in the data. Panel A shows the market shares for new and used jets among same brand and different brand upgrade purchases. Upgraders who buy the same brand of jet as they sell are 11.7 percentage points more likely to buy a new jet than upgraders who change brands. This statistic is consistent with the availability of manufacturer buyback increasing demand for new jets among upgraders.

This pattern could be driven by systematic correlations in preferences rather than buyback. In particular, the same pattern would obtain if consumers with strong brand loyalty also prefer new jets. To provide evidence that buyback seems to be driving a large part of these patterns, panels B and C of Table 4 repeat the exercise on two subsamples of consumers. Panel B shows that among consumers who are *observed* to sell a jet back to the manufacturer in the data the difference in new jet share between same brand and different brand upgrades is significantly higher, at 15.6 percentage points. In Panel C, which excludes observed buyback users, this difference is significantly smaller, although still significantly different from zero. Since the measure of buyback transactions is imperfect it is not surprising that there is a difference in the new jet share even when observed buyback users are excluded.

Differences in the share of new jets bought between same-brand and different-brand upgrades are the key variation that I use to identify the effect of buyback on demand. To obtain estimates of the transaction costs that buyback schemes circumvent, I estimate a structural model of new and used jet demand and holding behavior. In this model, the differences in the relative market shares for new and used jets between same brand and different brand upgrades reported in Table 4 will be attributed to differences in transaction costs between buyback eligible and non-eligible purchases.

An advantage of exploiting this variation is that it provides an estimate of the effect on demand of *all* manufacturer programs that encourage same brand used-new upgrades. In particular, the differences in Table 4 include the effect of manufacture-brokered upgrades that are not directly observed in the data. By conditioning these differences on different manufacturers, it is also possible to allow for heterogeneity in the effect of buyback.¹² The downside of this approach is that it cannot provide estimates of the relative value to consumers of different *types* of buyback program (e.g. trade-ins vs. manufacturer brokered sales). An alternative approach to identification would use differences in buyback policies across manufacturers and time to evaluate the impact of buyback schemes on demand. This is difficult because systematic documentation of these policies is not available and there are no obvious natural experiments that can be used to examine how changes in buyback schemes affected demand. Furthermore, any buyback incentives given to cross-brand upgrades will not be captured.¹³

4 Demand Model

4.1 Model Description

In this section, I present a model of new and used jet demand which incorporates the decision of which jet to buy for first time buyers, and the decisions of whether to hold, sell, or upgrade for jet owners.¹⁴ The model is adapted from existing work on demand for durable goods (Gowrisankaran and Rysman, 2012; Schiraldi, 2011).

I assume that consumers hold at most one jet in any given period. During each year, t , the set of existing new and used jet models is J_t . As discussed in Section 2.3, a model $j \in J_t$ is defined by its year, manufacturer, and segment (e.g. “a medium 1980 Cessna”). p_{jt} is the price of jet t .

¹²The main specification of the demand model allows for heterogeneity across manufacturers. In Appendix A.3 I examine heterogeneity across time periods, jet types, and consumer types. The basic patterns recorded in Table 4 hold throughout the years 1961-2000.

¹³As discussed above, industry participants have told me that cross-brand buyback is uncommon, and this is supported by Table 3. In terms of the model developed below, any cross brand buyback will be subsumed into the estimated value of transaction costs, τ .

¹⁴Note that although I will describe the model in the language of consumer choice and utility maximization, the consumers being studied are corporations. For this reason, I refer to consumers by “they” instead of “he or she”.

Consumer i 's flow utility from holding jet j in period t is

$$u_{jjt}^i = \gamma_{ijt} + \epsilon_{ijt}, \quad (1)$$

where γ_{ijt} is the individual-specific mean flow utility for jet j in year t , and ϵ_{ijt} is an iid type 1 extreme value shock to preferences. Consumer i 's flow utility from upgrading from jet type k to jet type j is

$$\begin{aligned} u_{kjt}^i &= \gamma_{ijt} + \tilde{\gamma}_{kj} + (p_{kt} - p_{jt}) \alpha_i^p - \tau_{ikj} + \epsilon_{ijt} \\ u_{k0t}^i &= p_{kt} \alpha_i^p - \tau_{ik0} + \epsilon_{i0t}. \end{aligned} \quad (2)$$

I indicate the outside option of “no jet” in equation 2 with $k = 0$ or $j = 0$. When upgrading from a jet of type k to a jet of type j , consumers pay the difference in prices, $p_{kt} - p_{jt}$, and an individual-specific transaction cost, τ_{ikj} . Consumers have individual-specific price sensitivities, α_i^p . When selling jet k and exiting the market, consumers receive the market price for jet k , p_{kt} , and pay the transaction cost τ_{ik0} . The mean flow utility of holding no jet is normalized to 0. $\tilde{\gamma}_{kj}$ captures any additional utility from upgrading beyond the flow utility of ownership and transaction costs.

Note that transaction cost τ_{ikj} can depend on the model being sold, k , and the model being purchased, j . This allows for the effect of buyback programs that can differentially affect transaction costs across different pairs of models. I assume that $\tau_{ijj} = 0$ and $\tau_{i0j} = 0$. That is, buyers who do not sell a jet - those who choose to hold their current jet or first time buyers - do not pay any transaction costs.

Consumers are forward looking. A consumer who holds jet model k at the beginning of period t has a value function given by Bellman equation 3, where Ω_{it} is consumer i 's information set.

$$V_i(k, \Omega_{it}, \epsilon_{it}) = \max \left\{ \max_{j \in J_t} \{ u_{kj}^i + \epsilon_{ijt} + E[V_i(j, \Omega_{it+1}, \epsilon_{it+1}) | \Omega_{it}] \}, u_{k0}^i + \epsilon_{i0t} \right\} \quad (3)$$

The expectation in equation 3 is taken over the vector of taste shocks ϵ_{it+1} , which is iid over time, and Ω_{it+1} which is assumed to evolve according to a Markov process, $P(\Omega_{it+1} | \Omega_{it})$. This state variable includes all relevant information about the state of the market, including available models, prices, and the holdings of all consumers. Consistent with the data construction described in Section 2, I assume that when consumers choose no jet they exit the market for good.

The probability that a consumer i who holds model k at date t upgrades to model j is thus given by

$$P_{kjt}^i = \frac{\exp(u_{kj}^i + E[V_i(j, \Omega_{it+1}, \epsilon_{it+1}) | \Omega_{it}])}{\exp(u_{k0}^i) + \sum_{l \in J_t \cup 0} \exp(u_{kl}^i + E[V_i(l, \Omega_{it+1}, \epsilon_{it+1}) | \Omega_{it}])}, \quad (4)$$

with analogous expressions for the probability of holding model k and exiting the market.

4.2 Inclusive Value Sufficiency

The consumer faces a dynamic discrete choice problem with a high dimensional state variable, Ω_{it} . Solving for the consumer choice probabilities in equation 4 requires specifying consumer expectations about the evolution of the state variable, Ω_{it} , and solving the Bellman equation for each point in the state space. Since the state variable includes the price and characteristics of all available aircraft and any market characteristics which may influence pricing in future periods, for instance the distribution of jet holdings among all consumers, solving this dynamic problem is impractical. To simplify the problem, I adopt a version of the inclusive value sufficiency assumption used by Hendel and Nevo (2006), Gowrisankaran and Rysman (2012) and Schiraldi (2011). I define the inclusive value of holding jet j at time t for consumer i as,

$$\delta_{ikt} = \log \left(\sum_{j \in J_t} \exp(u_{kj}^i + E[V_i(j, \Omega_{it+1}, \epsilon_{it+1}) | \Omega_{it}]) + \exp(u_{k0}^i) \right) \quad (5)$$

Notice that $\delta_{ijt} = E_\epsilon [V_i(j_{it}, \Omega_{it}, \epsilon_{it})]$, so the Bellman equation can be rewritten using iterated expectations as,

$$V_i(k, \Omega_{it}, \epsilon_{it}) = \max_{k \in J_t \cup 0} \{u_{j(i,t)k}^i + \epsilon_{ikt} + E[\delta_{ikt+1} | \Omega_{it}]\}. \quad (6)$$

Where the expectation is over δ_{ikt+1} conditional on the current state, Ω_{it} . That is, the consumer's optimal choice at date t depends only on date t flow utilities and the expected value of δ_{ikt+1} for all available jets $k \in J_t$. This form of the Bellman equation makes it clear that the dynamic problem can be simplified by imposing a restriction of consumer beliefs about the evolution of δ_{ikt} .

- Assumption 1: Each consumer i believes that δ_{ijt} evolves according to a first order Markov process $G_i(\delta_{ijt+1} | \delta_{ijt+1})$. In particular, $G_i(\delta_{ikt+1} | \Omega_{it}) = G_i(\delta_{ijt+1} | \delta_{ijt})$.

This assumption implies that consumers are boundedly rational because they do not condition on all available information when making predictions about future δ_{ijt+1} . In particular, It is clear that different states Ω_{it} could induce the same values of δ_{jt} for a product j , but lead to different distributions of δ_{ijt+1} . However, this form of inclusive value sufficiency makes solving the consumer's problem computationally tractable, and is flexible enough to capture the dynamic incentives that are induced by transaction costs that vary across products. To see this, note that under Assumption 1 I can rewrite the expression for the inclusive value (equations 5) as

$$\delta_{ikt} = \log \left(\sum \exp(\tilde{u}_{kj}^i - \tau_{kj} + E[\delta_{ijt+1} | \delta_{ijt}]) + \exp(u_{k0}^i) \right). \quad (7)$$

Where \tilde{u}_{kj}^i is flow utility net of transaction costs. Suppose that a buyback program eliminates transaction costs for holders of jet A but not for an otherwise identical jet B , so $\tau_{iBj} > \tau_{iAj} = 0$. Because upgrading from A to the consumer's most preferred jet is cheaper than upgrading from B , the inclusive value of holding jet A will be higher, $\delta_{iAt} > \delta_{iBt}$. Note however, that the inclusive value not only includes the cost of upgrading this period, but also implicitly accounts for future optimizing behavior through the terms $E[\delta_{ijt+1}|\delta_{ijt}]$. The value of the lower transaction costs from holding jet A enter δ_{iBt} through the term $E[\delta_{iAt+1}|\delta_{iAt}]$, which is the expected inclusive value of holding jet A in the next period. That is, δ_{iBt} incorporates the benefit of lower transaction costs from holding jet A along future paths where the consumer chooses to upgrade from B to A .¹⁵

Consumer choice probabilities (equation 4) can also be rewritten as

$$P_{kjt}^i = \frac{\exp(u_{kj}^i + E[\delta_{ijt+1}|\delta_{ijt}])}{\exp(u_{k0}^i) + \sum_{l \in J_i \cup 0} \exp(u_{kl}^i + E[\delta_{ilt+1}|\delta_{ilt}])}. \quad (8)$$

Given consumer beliefs, $G_i(\delta_{ijt+1}|\delta_{ijt+1})$, and flow utilities, u_{kj}^i , equations 7 can be solved for δ_{ijt} , and equation 8 can be used to recover choice probabilities. For this to constitute a solution to the consumer's problem, it must be that consumers beliefs $G_i(\delta_{ijt+1}|\delta_{ijt+1})$ are rational. I specify $G_i(\delta_{ijt+1}|\delta_{ijt+1})$ as a first order autoregressive process,

$$\delta_{ijt+1} = \rho_{1i} + \rho_{2i}\delta_{ijt} + \eta_{ijt}, \quad (9)$$

where $E(\eta_{ijt}|\delta_{ijt}) = 0$ and ρ_{1i} and ρ_{2i} are consumer specific incidental parameters. For a given vector of inclusive values, $\boldsymbol{\delta}_i$, the regression equation 9 yields beliefs $G_i(\delta_{ijt+1}|\delta_{ijt+1})$ defined by a vector of parameters $\boldsymbol{\rho}_i$. The solution to consumer i 's problem is therefore a fixed point of the vectors $(\boldsymbol{\delta}_i, \boldsymbol{\rho}_i)$ in the two equations 7 and 9.

4.3 Econometric Specification

In the main specification, consumer i 's mean utility for jet j at date t , γ_{ijt} , and the additional utility from upgrading from jet k to jet j , $\tilde{\gamma}_{kj}$, are defined as

$$\begin{aligned} \gamma_{ijt} &= \nu_i^0 + \gamma_j + \gamma_t + \nu_i^{m(j)} + new_{jt}\alpha^{new} + \xi_{jt} \\ \tilde{\gamma}_{kj} &= new_{jt}\alpha_{upgrade}^{new} + 1(m(j) = m(k))\alpha^{sb}. \end{aligned} \quad (10)$$

¹⁵The form of inclusive value sufficiency assumed here is different than in Schiraldi (2011). Schiraldi assumes that the transaction cost does not depend on the good currently held, so the inclusive value of upgrading does not have a j subscript. This allows him to write the consumer's problem as a static decision that depends only on the flow utility of the good currently held and the inclusive value of upgrading. In my setting, this is not possible because of the dependence of τ_{ikj} on manufacturer-specific buyback programs.

The flow utility from holding a jet, γ_{ijt} , depends on, γ_j , a jet fixed effect that captures average jet quality, γ_t , a year fixed effect, an indicator for whether the jet is new, and consumer specific random coefficients ν_i^0 and $\nu_i^{m(j)}$. ν_i^0 is a consumer specific intercept that captures heterogeneous preferences over the inside good, and $\nu_i^{m(j)}$ is a consumer specific preference for manufacturer $m(j)$. ξ_{jt} is jet-year level unobservable quality that captures (for example) jet specific deterioration in quality.

The additional utility from upgrading from jet k to jet j , $\tilde{\gamma}_{kj}$, depends on $new_{jt}\alpha_{upgrade}^{new}$, which allows upgraders to have a different preference for new jets than first time buyers, and $1(m(j) = m(k))$ which is an indicator for whether the manufacturer of jet j is the same as the manufacturer of jet k . The coefficient α^{sb} therefore captures consumer inertia in brand choice.¹⁶

I specify transaction costs as:

$$\tau_{ikj} = (\tau - new_{jt}1(m(j) = m(k))b_{m(j)}) 1(j \neq 0) + \tau^{exit}1(j = 0) + \nu_i^\tau \quad (11)$$

The transaction cost of upgrading from k to j is composed of two terms: a uniform transaction cost parameter τ that applies to all upgrades, and a buyback parameter $b_{m(j)}$ that applies only when the consumer who upgrades from k to j can take advantage of a buyback scheme. The main specification assumes that a consumer can trade in her jet to a manufacturer when upgrading from a used jet to a new jet of the same brand. b_m is therefore the coefficient on an interaction of a manufacturer fixed effect, an indicator for the purchased jet being new, new_{jt} , and an indicator for the jet purchase, j , having the same manufacturer as the jet sold, k . Thus, utility is shifted by b_m when a used jet of brand m is sold and a new jet of the same brand is purchased. When a consumer sells their jet and exits the market they pay a transaction cost τ^{exit} . Finally, there is an additive individual-specific term, ν_i^τ that allows for individual-specific heterogeneity in the level of transaction costs. The assumptions about the structure of buyback policies - that a consumer can take advantage of a manufacturer's buyback policy when they upgrade to a new jet from a used jet of the same brand, and that the effect of buyback policies on demand are different for different manufacturers - are based on the descriptive patterns on buyback use discussed in Section 3.¹⁷

I specify the individual specific price coefficient, α_i^p as

¹⁶I allow for both individual-specific brand preferences through γ_{ij} and consumer inertia in brand choice through the coefficient α^{sb} in equation 10 (see Keane, 1997 and Dubé et al., 2010). I argue in Section 5 below that these sources of persistent brand choice are separately identified.

¹⁷Note that this model cannot rationalize the imperfect takeup of buyback schemes observed in the data. In the model, buyback programs increase the utility of certain eligible choices, and consumers are not able to choose whether or not to make use of a buyback scheme. The buyback parameters b_m should therefore be interpreted as shifts to mean utility that explain differences in the level of demand for new jets between buyers who do not benefit from buyback and buyers who do benefit from buyback. Estimating a more detailed model in which consumers are able to choose whether or not to make use of buybacks would require more reliable data on observed buybacks and transaction prices.

$$\log(\alpha_i^p) = \alpha^p + \nu_i^p. \quad (12)$$

The individual-specific utility parameters, $\boldsymbol{\nu}_i = (\nu_i^p, \nu_i^\tau, \nu_i^0, \nu_i^1, \dots, \nu_i^7)$ (there are 7 manufacturers) are distributed joint normal, $\boldsymbol{\nu}_i \sim N(0, \Sigma)$. Σ has diagonal elements $(\sigma_p, \sigma_\tau, \sigma_0, \sigma_m, \dots, \sigma_m)$ and has all off-diagonal elements equal to zero except $Cov(\nu_i^p, \nu_i^\tau) = \sigma_{p\tau}$. I restrict the variances of the six brand preference parameters to be the same to help with identification and computation. I allow ν_i^p and ν_i^τ to be correlated because the transaction cost includes both explicit costs such as broker fees and taxes as well as implicit costs such as search cost or the cost of adverse selection, so consumer sensitivity to price is likely correlated with sensitivity to transaction costs, although not perfectly so.

Preferences are drawn i.i.d. across consumers when they enter the market. However, selection into the market will mean that preferences will be distributed differently among *holders* of different jet types in different years. The assumption of heterogeneous preferences rationalizes the fact that jets of all vintages are traded in the data. If all consumers had the same willingness to pay for quality, then only one type of jet would be demanded (up to the presence of ϵ_{ijt}), and there would be no gains from trade in the secondary market. As discussed further in Sections 6 and 7 below, the extent to which preferences are heterogeneous across consumers is an important determinant of the substitution patterns between new and used jets and the net effect of buyback policies on firm profits.

Finally, I follow Schiraldi (2011) in assuming that product level unobservables evolve according to a first order autoregressive process,

$$\xi_{jt+1} = \lambda \xi_{jt} + \omega_{jt}, \quad (13)$$

where ω_{jt} is mean 0 and independent of ξ_{jt} and λ is a parameter to be estimated. The other parameters to be estimated are the mean utility parameters $(\alpha^p, \alpha^{new}, \alpha_{upgrade}^{new}, \alpha^{sb}, \gamma_j, \gamma_t)$, the covariance parameters of the random coefficients $(\sigma_p, \sigma_\tau, \sigma_0, \sigma_m, \dots, \sigma_m, \sigma_{p\tau})$, the transaction cost parameters τ and τ^{exit} , and the buyback parameters b_m for each manufacturer m . Denote the vector of parameters by θ .

5 Estimation and Identification

5.1 Estimation Procedure

Estimation is based on the Generalized Method of Moments (GMM) procedure of Berry, Levinsohn, and Pakes (1995) nesting a fixed point procedure that solves the consumer's dynamic problem,

similar to Rust (1987). This procedure is close to that applied by Schiraldi (2011). The following paragraphs provide an overview of estimation with details provided below in Sections 5.1.1-5.1.3.

The inner loop that solves the consumer’s dynamic problem proceeds as follows. Starting with a candidate parameter vector, θ , a vector of product-year unobservables, ξ , and initial values for the incidental parameters (ρ_{1i}, ρ_{2i}) in 9, I use equation 7 to solve for the vector of inclusive values δ . I then estimate the regression equation 9 to recover new values of the parameters (ρ_{1i}, ρ_{2i}) . I repeat this procedure until I achieve convergence in δ . I perform this procedure for 1000 consumer types drawn from the distribution of the random coefficients ν_i .

The outer loop finds the vector of product-year unobservables, ξ , that rationalize observed *aggregate* market shares. I apply the BLP contraction mapping to log market shares, and for each new candidate value of ξ , I rerun the inner loop that solves for δ . I then use equation 13 to obtain product-year innovations in the unobservable, ω_{jt} , and form moments by interacting ω_{jt} with instruments.

Note that in addition to the standard preference heterogeneity across consumers induced by the “random coefficients”, there is heterogeneity in preferences across consumers that hold different jets induced by transaction costs, buyback, and inertia in brand choice. In order to identify this preference heterogeneity I augment the estimation procedure with “micro-moments” as in Petrin (2002).

5.1.1 BLP-Style Moments

Let M_{kt} be the number of consumers that hold jet k at the beginning of year t , and s_{kjt} be the share of those consumers who upgrade to jet j in year t . M_{0t} is the number of first time buyers that arrive in the market at date t and s_{0jt} are purchase shares among first-time buyers.¹⁸ Define the aggregate market share for jet j in year t as

$$s_{jt} = \frac{\sum_{k \in J_{t-1} \cup 0} M_{kt} s_{kjt}}{\sum_{k \in J_{t-1} \cup 0} M_{kt}}. \quad (14)$$

To construct moments, I match these observed market shares with model-implied market shares

$$\hat{s}_{jt}(\theta, \xi) = \frac{\sum_{k \in J_{t-1} \cup 0} M_{kt} \hat{s}_{kjt}(\theta, \xi)}{\sum_{k \in J_{t-1} \cup 0} M_{kt}}, \quad (15)$$

¹⁸As is common in demand estimation studies, I do not directly observe the “outside good” share since M_{0t} , which includes consumers who chose not to purchase a jet, is not observed in the data. In estimation, I set M_{0t} equal to the number of consumers who purchased jets for the first time in years t to $t + 5$. Under this assumption, the market is growing over time for reasons unrelated to the jet characteristics included in the model. In appendix A.4 I consider estimation of a model with an alternative definition of M_{0t} that uses data on counts of public companies.

where $\hat{s}_{kjt}(\theta, \boldsymbol{\xi}, \boldsymbol{\delta})$ is the model-implied share of jet k holders who choose to upgrade to jet j , which depends on the parameters, θ , and the vector of product-year level unobservables, $\boldsymbol{\xi}$. This is given by

$$\hat{s}_{kjt}(\theta, \boldsymbol{\xi}) = \int P_{kjt}^i(\theta, \boldsymbol{\xi}) dF_{kt}(\boldsymbol{\nu}_i | \theta, \boldsymbol{\xi}), \quad (16)$$

where $P_{kt}^i(\theta, \boldsymbol{\xi})$ are obtained by first solving the consumer's dynamic problem (the "inner loop" described above), obtaining inclusive values $\boldsymbol{\delta}$, and applying by equation 8. $F_{kt}(\boldsymbol{\nu}_i | \theta, \boldsymbol{\xi})$ is the joint distribution of random coefficients for holders of jet k in period t .

Recall that consumer preferences are drawn i.i.d. when consumers *first enter the market*. The distribution $F_{kt}(\boldsymbol{\nu}_i | \theta, \boldsymbol{\xi})$ for $k \neq 0$ differs for each k because it depends on the selection into ownership of consumers of different types. In particular, the distribution $F_{kt}(\boldsymbol{\nu}_i | \theta, \boldsymbol{\xi})$ for holders of jet k in year t depends on the probability that consumers with different values of $\boldsymbol{\nu}_i$ choose to purchase and hold jet k in all previous years.

Computation of $\hat{s}_{kjt}(\theta, \boldsymbol{\xi})$ therefore proceeds by sequential simulation, starting with the first year of the sample, $t = 1961$. There are no jet holders in the sample before 1961, so $M_{k1961} = 0$ for all $k \neq 0$.¹⁹ For each year, I draw new consumers from the unconditional distribution of $\boldsymbol{\nu}_i$, with the number of new entrants proportional to M_{0t} .²⁰ I solve the dynamic problem of all consumers and simulate the choices of new entrants and existing jet holders in each year t . Year t choices determine the distributions of consumer types in year $t + 1$, $F_{kt+1}(\boldsymbol{\nu}_i | \theta, \boldsymbol{\xi})$. Given these simulated distributions, I can compute $\hat{s}_{kjt}(\theta, \boldsymbol{\xi})$ and $\hat{s}_{jt}(\theta, \boldsymbol{\xi})$ from equations 15 and 16.

For a given candidate parameter vector θ I find a vector of product level unobservables $\boldsymbol{\xi}$ such that $s_{jt} = \hat{s}_{jt}(\theta, \boldsymbol{\xi})$ by iterating according to Berry, Levinsohn, and Pakes (1995) contraction mapping procedure. I then obtain innovations to the unobservables, $\omega_{jt}(\theta)$ by applying equation 13 and use instruments Z_{jt} to construct empirical analogues of the moments conditions

$$E[\omega_{jt}(\theta)Z_{jt}] = 0.$$

Instruments include all product characteristics included in the utility specification except for price. To instrument for p_{jt} , I use counts of the number of consumers in the sample that hold close substitutes to jet type j at the beginning of year t and lagged prices. To measure close substitutes, I use the number of jets of the same size (small, medium, large) of the same age as jet j and the number of jets of the same size that are one year older than jet j .

It is clear that the more jets of type j held by owners, the higher the quantity of jet j supplied

¹⁹The choice of 1961 as a starting year means that I do not have to pick initial distributions of preferences among jet holders, under the assumption that the distribution of $\boldsymbol{\nu}_i$ is the same among consumers arriving to the market for the entire sample period.

²⁰To reduce computational time, I draw consumers from a set of 1000 types drawn from the unconditional distribution of $\boldsymbol{\nu}_i$.

on the used market at a given price level. Thus the number of model j jets and close substitutes held is a supply shifter that is correlated with p_{jt} , but is uncorrelated with ω_{jt} since the number of jets held at date t is determined at date $t - 1$. I use ω_{jt} rather than ξ_{jt} to construct these moments to account for the possibility of autocorrelation in ξ_{jt} , which would invalidate moments based on interactions of ξ_{jt} and holdings of jets in period $t - 1$. The innovations ω_{jt} can be thought of as shocks to product quality that are unanticipated at date $t - 1$. For instance, news about the reliability of used aircraft or maintenance requirements. In Appendix Table A.8, I present a diagnostic “first stage” regression which shows that the instruments predict prices conditional on jet and year fixed effects. In what follows, I call this vector of moments $G(\theta)$.

5.1.2 Micro-Moments

Note that the BLP moments are constructed using aggregate market shares s_{jt} not market shares conditioned on current holdings s_{kjt} . An attempt to match market shares conditional on current holdings s_{kjt} would run into the “zero market shares problem” (see for example Gandhi, Lu, and Shi, 2019; Quan and Williams, 2018), since the share of holders of jet k that upgrade to jet j in a particular year t is frequently 0. However, aggregate market shares do not capture preference heterogeneity across holders of different jets, including the effect of buyback on demand for new jets. To identify this heterogeneity, I add a set of “micro-moments” in the spirit of Petrin (2002).

The micro moments are computed using averages across consumers from the estimation sample. Denote the jet owned by consumer i at the beginning of period t as j_{it} . Let $j_{it} = 0$ if i does not own a jet at date t . Define indicators for whether a consumer upgraded at date t and whether a consumer made their first purchase at date t as

$$\begin{aligned} upgrade_{it} &= I(j_{it+1} \neq j_{it} \ \& \ j_{it+1} \neq 0 \ \& \ j_{it} \neq 0) \\ first_{it} &= I(j_{it+1} \neq 0 \ \& \ j_{it} = 0). \end{aligned}$$

Table 5 lists the included micro-moments along with the parameters that are most closely related to each moment. The relationship between moments and parameters is discussed further in section 5.2 below.

To illustrate the implementation of these moments, consider moment 3 in Table 5. I compute the empirical probability of upgrading to a new jet conditional on upgrading at all as

$$\hat{h}_3 = \frac{\frac{1}{N} \sum_{i=1}^N \frac{1}{T} \sum_{t=1}^T I(age(j_{it+1}) \leq 1) upgrade_{it}}{\frac{1}{N} \sum_{i=1}^N \frac{1}{T} \sum_{t=1}^T upgrade_{it}}. \quad (17)$$

As in Petrin (2002), the conditional probability defined by equation 17 is written as a ratio of averages over N observations corresponding to the N owners in the sample. In this case, the ratio

Table 5: Micro-Moments

	Moment	Related Parameters	
1	Upgrade conditional on holding	$P(\text{upgrade}_{it} j_{it} \neq 0)$	τ
2	Exit conditional on holding	$P(j_{it+1} = 0 j_{it} \neq 0)$	τ^{exit}
3	New jet conditional on upgrade	$P(\text{age}(j_{it+1}) \leq 1 \text{upgrade}_{it} = 1)$	$\alpha_{upgrade}^{new}$
4	New jet conditional on first purchase	$P(\text{age}(j_{it+1}) \leq 1 \text{first}_{it} = 1)$	α^{new}
5	Same brand conditional on upgrade	$P(m(j_{it+1}) = m(j_t) \text{upgrade}_{it} = 1)$	α^{sb}
6-12	Difference in new jet share between same-brand upgrades and brand switchers	$P(\text{age}(j_{it+1}) \leq 1 m = m(j_{it+1}) = m(j_t) \& \text{upgrade}_{it} = 1)$ $-P(\text{age}(j_{it+1}) \leq 1 m = m(j_{it+1}) \neq m(j_t) \& \text{upgrade}_{it} = 1)$	b_m, σ_m
13-15	Expected purchase price conditional on upgrading from jets of different ages	$E(\text{price}(j_{it+1}) \text{age}(j_{it} < 5) \& \text{upgrade}_{it} = 1)$ $E(\text{price}(j_{it+1}) \text{age}(5 \leq j_{it} < 15) \& \text{upgrade}_{it} = 1)$ $E(\text{price}(j_{it+1}) \text{age}(15 \leq j_{it}) \& \text{upgrade}_{it} = 1)$	$\sigma_p, \sigma_{p\tau}$
16-18	Exit conditional on holding jets of different ages	$P(j_{it+1} = 0 \text{age}(j_{it} < 5))$ $P(j_{it+1} = 0 \text{age}(5 \leq j_{it} < 15))$ $P(j_{it+1} = 0 \text{age}(15 \leq j_{it}))$	σ_0
19	Upgrade conditional on past upgrade	$P(\text{upgrade}_{it} j_{it} \neq 0 \& \max_{\tilde{t} < t} \{\text{upgrade}_{i\tilde{t}}\} = 1)$	σ_τ
20	Brand held before conditional on upgrade	$P(m(j_{it+1}) \in \{m(j_{i\tilde{t}})\}_{\tilde{t} < t} \text{upgrade}_{it} = 1)$	σ_m

is the number of upgrades to new jets across all owners and years in the sample to the number of upgrades to any jet. The model-implied analogue is given by

$$h_3(\theta) = \frac{\sum_{t=1}^T \sum_{k \in J_{t-1}} M_{kt} \sum_{j \in J_t^N} \hat{s}_{kjt}(\theta, \boldsymbol{\xi}(\theta))}{\sum_{t=1}^T \sum_{k \in J_{t-1}} M_{kt} \sum_{j \in J_t} \hat{s}_{kjt}(\theta, \boldsymbol{\xi}(\theta))}, \quad (18)$$

where J_t^N is the set of *new* jets available in year t . The estimation procedure looks for parameters that minimize the difference $H_3(\theta) = \hat{h}_3 - h_3(\theta)$. Other moments are constructed analogously. The vector of micro-moments is $H(\theta)$ with i th entry $H_i(\theta)$.

5.1.3 Objective Function

I use the two-step optimally weighted GMM estimator introduced by Hansen (1982). The first step objective function is

$$\tilde{\theta} = \arg \min_{\theta} \begin{bmatrix} G(\theta) \\ H(\theta) \end{bmatrix} [G(\theta)', H(\theta)']'. \quad (19)$$

Using consistent estimates $\tilde{\theta}$, I then construct a weighting matrix $\Omega(\tilde{\theta})$ which is an estimate of the asymptotic covariance matrix of $\begin{bmatrix} G(\tilde{\theta}) \\ H(\tilde{\theta}) \end{bmatrix}'$. As in Petrin (2002), $\Omega(\tilde{\theta})$ is block diagonal since $G(\theta)$ are averages over a sample of product-years and $H(\theta)$ are functions of averages over a sample of consumers.²¹ The covariance of $G(\theta)$ is straightforward to compute, the covariance of $H(\theta)$ is obtained by re-computing the micro-moments for 200 bootstrap samples of consumers. The final estimates are then given by:

$$\theta^* = \arg \min_{\theta} \begin{bmatrix} G(\theta) \\ H(\theta) \end{bmatrix} \Omega(\tilde{\theta}) [G(\theta)', H(\theta)']'. \quad (20)$$

Standard errors are obtained using the usual GMM formula as in Petrin (2002).

5.2 Identification

The identification of the key parameters relies on the assumption that manufacturers only accept trade-ins of their own brands. This assumption allows preference for newness to be different in the first time and replacement markets, and to be identified separately from the effect of buyback.

²¹In adopting this weighting matrix, I am following convention in assuming that the market shares used to construct the moments $G(\tilde{\theta})$ are observed without error. This assumption is commonplace in the demand estimation literature following Berry, Levinsohn and Pakes (1995). In reality, the market share is constructed from a finite sample of actual purchases. Since the sample of consumer data used to construct the market share is the same as that used to construct the micro-moments, relaxing this assumption would yield an optimal weighting matrix that is not block diagonal. Note that the choice of weighting matrix only affects the efficiency of the parameter estimates, not their consistency.

As recorded in panel A of Table 4, conditional on purchase, a replacement buyer is more likely to buy a new jet than a first time buyer. In this model, this is explained by the parameters $\alpha_{upgrade}^n$ and b_m , both of which shift the utility of new jets for replacement buyers. Similarly, the tendency of replacement buyers to buy jets of the same brand (manufacturer) as those they sell is captured by α^{sb} , individual-specific brand preferences, ν_i^m , and b_m . b_m is separately identified by the interaction of these two effects - the extent to which *same brand upgraders* are more likely to purchase a new jet than *brand switchers*. Alternatively, b_m can be thought of as being identified by the extent to which the share of same brand purchases among all replacements of *used jets with new jets* is greater than the share of same brand purchases among all replacements of *used jets with used jets*. The identification is similar to the classic difference in differences approach - after preference for the same brand and preference for newness are controlled for, any additional effect of the interaction - new jets of the same brand - is identified with the effect of buyback schemes.

In particular, the patterns identified in panel A of Table 4 that identify b_m enter directly into the GMM objective function through micro-moments 6-12 as detailed in Table 5. α^{sb} and $\alpha_{upgrade}^{new}$ are identified by the probability of upgrading to a jet of the same brand, and a new jet respectively. These probabilities enter through micro-moments 3 and 5. Brand inertia, α^{sb} , is separately identified from the variance of individual-specific brand preferences, σ_m , by micro moment 20, which is the probability of upgrading to a jet of a previously held brand (not necessarily the same brand as the currently held jet).

The transaction cost parameter, τ , is identified by the frequency with which owners in the sample upgrade. This is captured by micro-moment 1. If $\tau = 0$, then owners would upgrade frequently as their jets age and provide less utility. τ therefore rationalizes the average holding time observed in the data of around 4 years.

The mean coefficients on price and other jet characteristics are identified by the correlation between market shares and instruments, as usual in BLP-style estimation. The heterogeneity in preferences, captured by the covariance parameters $(\sigma_0, \sigma_p, \sigma_\tau, \sigma_{p\tau})$ are identified by micro moments 13-19. Moments 13-15 measure the expected purchase price for upgraders that hold of jets of different ages. If there was no heterogeneity in preferences, then choice probabilities would be identical for holders of different jets, and there would be no relationship between the age of the jet currently held and the price of the upgrade. Heterogeneity means that consumers with low values of α_i^p are more likely to hold older jets and more likely to upgrade to cheaper jets. These moments also help identify the covariance parameter, $\sigma_{p\tau}$. If $\sigma_{p\tau}$ is negative, then consumers who are price sensitive also rarely upgrade, and purchase prices among upgraders will be higher. Moments 16-18 measure the probability of exiting the market among holders of different ages. Following a similar logic, owners with low values of α_i^0 should be more likely to hold older, cheaper jets, and more likely to exit the market conditional on holding. Moment 19 measures the probability of upgrading conditional on having upgraded in the past. If $\sigma_\tau > 0$, then the set of consumers who

have upgraded at least once should be selected to have lower transaction costs, and thus be more likely to upgrade again.

Additional variation that helps identify these distributional parameters comes from the instrument that records the supply of used jets from the previous period. As supply increases, the increase in the number of first time buyers that choose to purchase a jet depends on the distribution of α_i^0 - if there is no preference heterogeneity then the inside market share will increase according to the logit formula, while greater heterogeneity will diminish this market expansion effect.

6 Results

6.1 Estimated Parameters

The main parameter estimates are presented in Table 6. Appendix Table A.10 records the fit of the micro-moments from Table 5 at the estimated parameters.

The median coefficient on price, which is the marginal utility of \$100,000, is $\alpha_i^p = 1.517$. The individual-specific coefficient on price which enters the utility function is distributed according to $\alpha_i^p \sim \log N(\alpha^p, \sigma^p)$. This is the distribution of price coefficients from which consumers' preferences are drawn when they enter the market. Appendix figure A.2 illustrates how the estimated distribution of price coefficients changes *conditional on jet holdings*, illustrating selection into jet ownership of less price-sensitive consumers. The mean price parameter among jet holders is $E(\alpha_i^p | hold) = 0.466$, and the mean price parameter conditional on upgrading is $E(\alpha_i^p | upgrade) = 0.300$. There is also significant heterogeneity in the value of the inside good and preferences for manufacturers, indicated by the estimated values of σ_0 and σ_m .

The covariance parameter $\sigma_{p\tau}$ is significant and negative, indicating that consumers who are more price sensitive are also more likely to face higher idiosyncratic transaction costs ν_i^τ (equation 11). Idiosyncratic transaction costs are not perfectly correlated with price sensitivity, suggesting that transaction costs include non-pecuniary costs such as the time cost of finding a buyer. However, the standard deviation of ν_i^τ , σ_τ , is small relative to the average transaction cost, indicating that there is limited cross-consumer heterogeneity in transaction costs. This finding can be understood by examining the micro-moments recorded in Appendix Table A.10. Consistent with the intuition in Section 5.2 above, the probability of upgrading *conditional* on having upgraded previously is close to the unconditional probability of upgrading.

The estimated median transaction cost for upgrades at the expected price parameter among jet holders, $\frac{\tau}{E(\alpha_i^p | hold)}$, is approximately \$1.8 million or 27% of the average jet price of \$6.7 million. The estimated median transaction costs faced by consumers exiting the market, $\frac{\tau^{exit}}{E(\alpha_i^p | hold)}$, are lower at \$812 thousand. By way of comparison, Schiraldi (2011) finds that transaction costs (defined

Table 6: Parameter Estimates

Parameter	Estimate	SE	Parameter	Estimate	SE	Buyback Parameter b_m	Estimate	SE
α^p	-1.517	0.069	σ_p	1.785	0.187	Bombardier	0.915	0.247
α^{new}	2.892	0.114	σ_τ	0.391	0.119	Cessna	0.648	0.104
$\alpha_{upgrade}^{new}$	1.743	0.104	$\sigma_{p\tau}$	-0.973	0.099	Dassault	0.664	0.127
α^{sb}	0.548	0.185	σ_0	1.895	0.115	Gulfstream	0.828	0.124
τ	8.225	0.160	σ_m	0.237	0.060	IAI	0.907	0.101
τ^{exit}	3.785	0.163				Raytheon	0.550	0.213
						Other	0.558	0.199

Notes: Table reports estimated parameters and standard errors for the demand model. Prices are in hundreds of thousands of 2009 \$.

in a similar manner) in the market for new and used cars are between 10% and 80% of the sale price. Note that conditional on upgrading, realized transaction costs are on lower than this because consumers with low values of the idiosyncratic component of transaction costs, ν_i^τ , are more likely to upgrade. Similarly, consumers only upgrade when they receive a sufficiently high draw of the logit shocks ϵ_{ijt} , which could be interpreted to include shocks to transaction costs.

As expected based on the descriptives recorded in Table 4, the manufacturer-specific buyback parameters b_m are positive and statistically significant for all manufacturers. As discussed in Section 5.2, this reflects the fact that the difference in probability between same-brand upgrades and different-brand upgrades is higher for new jets than for used jets. The estimates indicate that the impact of buyback schemes on demand is equivalent to a reduction of transaction costs of between 6.7% for Dassault and 11.1% for Bombardier. In dollar terms the value of buyback is about 3% of the average jet price. This is roughly consistent with buyback eliminating the direct cost of brokerage. Quotes for various private aircraft broker fees services range from 0.5% to 10% of the sale price. The remainder of transaction costs that are not eliminated buyback could include sales taxes, legal fees, and more general non-pecuniary costs such as the cost of inconvenience and delays. Differences in b_m across manufacturers could reflect differences in the costs associated with upgrading different types of jets and differences in the generosity of buyback policies across manufacturers.

The estimated model includes jet and time fixed effects, and therefore the effect of jet characteristics on utility are not immediately apparent from the estimated parameters. To illustrate the effect of jet characteristics and age on utility, I regress jet mean utility, $\int \gamma_{ijt} dF(\boldsymbol{\nu}_i)$, on jet characteristics.

The regression coefficients recorded in the first two columns of Appendix Table A.9 imply that

the loss of utility from a jet aging one year, at the mean price parameter among jet holders, is equivalent to an increase in price of between \$10 and \$20 thousand. Compare this to the difference in utility between a new and used jet, $\frac{\alpha^{new}}{E(\alpha_i^{hold})}$, which is equivalent to a change in price of \$118 thousand. The drop in quality once a jet becomes used is an order of magnitude larger than the annual depreciation in quality thereafter. This suggests that the jet market might exhibit the ‘lemons’ effect of adverse selection on the used market, as suggested by Gilligan (2004).

Utility is increasing in jet power and range, with an increase in power of 1 kN equivalent to a reduction in price of around \$6000 and an increase in range of 100 km equivalent to a reduction in price of around \$13000. Utility is decreasing in maximum weight which could reflect factors such as fuel consumption.

6.2 The Effect of Buyback on Demand and Supply

Manufacturer buyback increases the demand for new jets by encouraging owners to upgrade to a new unit instead of holding their current used unit, upgrading to another used unit, or selling their jet and not upgrading. Buyback also increases the supply of used jets, since units that are bought back are resold by manufacturers. In this section, I measure the direct effect of buyback schemes on demand and supply *at fixed prices* by comparing a simulation of the demand model at the estimated parameter values to a counterfactual simulation under which no manufacturers offer buyback.

To benchmark counterfactual simulations, I simulate market outcomes at the estimated parameters and product unobservables, holding fixed prices and market sizes at the observed levels. Starting with the first year of the sample, I draw preferences from the estimated distribution for M_{0t} first time buyers. I then simulate choices by all first time buyers and jet holders and proceed to the next period, keeping track of preferences and holdings. I repeat this simulation under a “no buyback” scenario by setting $b_m = 0 \forall m$ and holding all other parameters, prices, market sizes, and consumer beliefs fixed. Note that this is not an equilibrium counterfactual, as I do not allow prices or consumer beliefs to adjust. Rather, the comparison between these two simulations measures *direct effect* on demand for new and used jets of removing buyback for all firms.

The first three columns of Table 7 illustrate how buyback shifts the demand for new and used jets at fixed prices. Buyback increases the demand for new units from upgrading consumers over the sample period by 442.4 units or about 37% of the no-buyback demand. There are three margins of substitution that enter this increase in demand for new jets among jet holders: substitution away from upgrades to used jets, substitution away from exiting the market, and substitution away from holding (not upgrading). The second and third rows of Table 7 indicate that of the 442.4 additional new jets demanded, 29% comes from consumers substituting from used jets, 2.5% comes

Table 7: Demand Simulations

	Baseline Parameters			No Heterogeneity		
	No Buyback	Buyback	Δ	No Buyback	Buyback	Δ
(1) Upgrades to New	1193.0	1635.4	442.4	1042.2	1476.5	434.3
(2) Upgrades to Used	3185.1	3056.4	-128.7	3359.5	3325.5	-34
(3) Exits	16700.0	16689.0	-11.0	15459.0	15434.0	-25
Used Jet Supply to First Time Buyers = (1) + (3)	17893.0	18324.4	431.4	16501.2	16910.5	409.3

Notes: The first three columns record statistics computed under simulations of the demand model at fixed prices at the estimated parameters. The first column records simulations with $b_m = 0$ for all manufacturers. the second column records simulations with b_m set to the estimated values. The third column records the difference between the first and second columns. The fourth, fifth and sixth columns record equivalent numbers for demand simulations under the non-heterogeneity parameters recorded in Appendix Table A.11. All columns are averages over 100 simulations. All figures are totals for the period 1961-2000.

from substitution from market exits, and the remaining 68.5% comes from substitution from jet holding.

These substitution patters affect the extent to which buyback shifts the supply of used units to first time buyers. The supply of used jets to first time buyers can be decomposed as follows,

$$SupplyUsed = UpgradesToNew + Exits. \quad (21)$$

In particular, substitution from holding to upgrading to a new jet increases the number of used jets supplied by one, and substituting from holding to exiting the market increases the number of used jets supplied by one. On the other hand, substituting from holding to upgrading to a used jet does not affect the net supply of used jets to first time buyers, and thus the number of upgrades to used does not enter equation 21.

The implication of this decomposition is illustrated by the fourth line of Table 7. The net change in the supply of jets to first time buyers is the sum of the positive change in the number of upgrades to new and the negative change in the number of exits. At the estimated parameters and observed prices, supply of used jets to first time buyers increases by 97.5% of the increase in new jet demand, since most of the increase in demand comes from substitution form holding. If more of the increase in new jet demand came from substitution from exiting the market, the effect on used jet supply would be lower.

To illustrate how alternative demand substitution patterns change the effects of buyback, the second set of columns in Table 7 record analogous demand simulations using parameters estimated

under the restriction of no-heterogeneity. That is, $\Sigma = 0$.²² Without preference heterogeneity, only 8% of the increase in demand for new units comes from substitution from upgrades to used units, compared to 29% in the baseline simulation. The high degree of substitution between upgrades of to used and new jets in the baseline simulation is driven by preference heterogeneity. For example, consumers with low idiosyncratic transaction costs, ν_i^T , are likely to be on the margin between upgrading to a used unit and upgrading to a new unit, and thus buyback should draw disproportionately more consumers away from upgrading to a used unit.

Similarly, preference heterogeneity means that consumers who exit the market are likely to have low draws of ν_i^0 , and should therefore be unlikely to substitute towards new units. Indeed, removing heterogeneity increases the share of substitution from exits from 2.5% to 6% at the baseline parameters, and reduces the increase in used jet supply from 97.5% to 94% of the increase in demand for new jets. This comparison illustrates the importance of measuring preference heterogeneity in determining the effect of buyback on the supply of used units.

7 Buyback in Equilibrium

In equilibrium, the effect of buyback on the supply of used units will lower the price of used jets and induce substitution of first time buyers away from new jets, eating into the increase in manufacturer revenue from increased sales to upgrading consumers. Manufacturer revenue cannot be *decreased* by introducing buyback, but it can be completely cannibalized by the effect on the used market, leading to profit loss if buyback is costly.

To see this, notice that each additional new unit purchased because of buyback is either an additional upgrade (an increase in used jet supply) or a substitution to new from used (a reduction in demand for used jets). Either way, the net supply of used jets to the first time buyer market increases by one. The number of additional new units sold to upgrading customers must therefore equal the number of additional used units purchased by first time buyers. The extent of cannibalization depends on the substitution patterns among first time buyers. Each additional used unit sold to first time buyers must represent either a substitution from a new jet or a substitution from no purchase. If all of the increase in used jet demand comes from substitution from new jets, then the increase in new jet sales in the upgrade market is entirely cannibalized by the reduction in demand in the first time buyer market (subject to the mix of models/prices being the same). I formalize this argument in Appendix A.3.

In this section, I run *equilibrium* simulations to measure the effect of buyback on revenue *net* of this cannibalization effect.

²²Parameters estimated under this restriction are recorded in Appendix Table A.11.

7.1 Computing Equilibrium

To measure the effect of buyback on manufacturer revenue in equilibrium, I simulate the adjustment of new and used jet prices to the changes in demand and supply discussed in the previous section. Define total demand for jet j at the estimated parameters as

$$D_{jt} = \sum_{k \in J_{t-1} \cup 0} M_{kt} s_{kjt}(\hat{\theta}, \hat{\xi}).$$

Recall that consumers are forward looking, and that under assumption 1 and equation 9, consumer i 's beliefs about product-specific continuation values are defined by the inclusive values δ_i and the incidental parameters ρ_i . At the estimated parameters $(\hat{\theta}, \hat{\xi})$, the joint distribution of (ρ_i, δ_i) and consumer preferences are implicitly defined by the solution to a fixed point problem conditional on all prices and other product characteristics. To compute demand under counterfactual prices it is therefore necessary to solve for new fixed point values of (ρ_i, δ_i) so that consumers' beliefs are consistent with the flow utilities that obtain under the new price vector. To make explicit the dependence of demand on prices p and beliefs (ρ_i, δ_i) , I write demand as

$$D_{jt}(p; (\rho, \delta)) = \sum_{k \in J_{t-1} \cup 0} M_{kt} s_{kjt}(\hat{\theta}, \hat{\xi}; p, (\rho, \delta)).$$

Using this expression, I define an equilibrium in prices and consumer beliefs.

- Assumption 2: Equilibrium is defined by a vector of prices p and a distribution of beliefs (ρ_i, δ_i) such that:

- Let J_{ft} be the set of new jets offered by firm f in year t .²³ New jet prices maximize firms' static profits conditional on consumer beliefs,

$$\Pi_{ft} = \sum_{j \in J_{ft}} D_{jt}(p; (\rho, \delta))(p_{jt} - c_{jt}). \quad (22)$$

- Used jet markets clear.²⁴

$$ED(p; (\rho, \delta)) \equiv D_{jt}(p; (\rho, \delta)) - M_{jt} = 0 \quad \forall j, t : new_{jt} = 0. \quad (23)$$

²³As mentioned above, I treat age 0 jets as "new". I therefore assume firms sell only year t models in year t , subject to the adjustment described in footnote 10.

²⁴I assume that all used jet markets clear within one period (1 year). I do not explicitly model the role of independent jet dealers who mediate a substantial share transactions on the secondary market, as shown in Appendix Figure A.1. Table 2 shows that these jet dealers hold jets for an average of 11.5 months, suggesting that the assumption of markets clearing in one year may not be unreasonable. As discussed in the text, τ includes the cost of using a dealer including the dealer's markdown on the price of used jets. Since τ is held fixed across counterfactuals in Section 7, I am implicitly assuming that dealer markdowns are held fixed.

- Consumer beliefs $(\boldsymbol{\rho}_i, \boldsymbol{\delta}_i)$ are given by the solution to the fixed point defined by equations 7 and 9.

The first part of this definition implies the familiar first order condition for new jet prices,

$$p_{jt} = c_{jt} - \frac{D_{jt}(p; (\boldsymbol{\rho}, \boldsymbol{\delta})) + \sum_{k \in J_{ft}, k \neq j} \frac{\partial D_{kt}(p; (\boldsymbol{\rho}, \boldsymbol{\delta}))}{\partial p_{jt}} (p_{kt} - c_{kt})}{\frac{\partial D_{jt}(p; (\boldsymbol{\rho}, \boldsymbol{\delta}))}{\partial p_{jt}}}. \quad (24)$$

I evaluate the markup term on the right hand side of equation 24 at the estimated parameters for each jet-year pair, and use the observed prices to back out marginal costs, c_{jt} . The distribution of implied markups, defined as $\frac{p_j - c_j}{c_j}$, is illustrated in Figure 2. The average marginal cost is \$102 million and the average markup is 12.6%. The estimated markups are broadly consistent with figures from manufacturers’ financial reports.²⁵ In columns 3-4 of Appendix Table A.9 I present regressions of estimated marginal costs on jet characteristics. Marginal costs are significantly positively correlated with range, power, maximum weight, and mean utility, γ_{ijt} .

This definition of equilibrium assumes that firms set prices to maximize static profit, holding fixed jet characteristics.²⁶ Because of the dynamics of consumer demand, a forward looking firm might be able to increase profit beyond that achieved by a static profit maximizer. As a robustness test, I also estimate marginal costs for forward looking firms with perfect foresight about future demand and market structure (that is, firms internalize the effect of a price change on future profit). Appendix figure A.3 plots marginal cost estimates from the static pricing model against marginal cost estimates from the forward looking model. The two sets of estimates are highly correlated across products and the static marginal cost estimates do not appear to be systematically higher or lower than the forward looking marginal costs, suggesting that dynamic concerns are second order in the firm’s pricing problem.

Given the implied marginal costs, a parameter vector θ (not necessarily equal to the estimated parameters) and unobservables $\hat{\xi}$, I solve for the equilibrium as follows.


First, I simulate consumer demand to construct market shares according to equation 15. For new jets, I then compute the prices implied by equation 24. As usual (see Nevo, 2001), iteratively substituting price vectors into these equations yields equilibrium new jet prices, holding used jet prices and consumer beliefs fixed.

Second, I find a vector of used jet prices such that all used markets clear. That is, the excess of the number of units demanded (where demand includes the decision of holders of jet j not to upgrade)

²⁵Actual gross profit margins for 2021 for the holding companies of the five largest manufacturers are reported in Appendix Table A.12. This data is not available for IAI, which is a privately held firm owned by the Israeli government.

²⁶Simplified supply side models are common in the literature on dynamic demand. For instance, neither Gowrisankaran and Rysman (2012) nor Schiraldi (2011) consider supply side responses in counterfactuals. Lee (2013) adopts a “fixed markup” rule on the supply side.

Figure 2: Implied Markups



Notes: This figure is a histogram of the implied markups $\frac{p_j - c_j}{c_j}$, where p_j is the observed price and c_j is the implied marginal cost, among all new jets available in the estimation sample, where jets are defined by a manufacturer-year-segment. The figure drops the 1.1% of estimated markups that are greater than 0.3. The marginal costs are backed out from the manufacturers' first order conditions as described in the text.

over the existing stock of jet j held at the beginning of year t is zero. To solve for equilibrium used jet price, fixing new jet prices, I increase the prices of jets for which $ED_{jt}(p, \theta) > 0$ and reduce the prices of jets for which $ED_{jt}(p, \theta) < 0$ until $ED_{jt}(p, \theta) = 0$ for all used jets.²⁷

Finally, I solve for equilibrium consumer beliefs (ρ_i, δ_i) for each simulated consumer using the fixed point discussed above. I alternate between solving for new jet prices, used jet prices, and beliefs until convergence is achieved.

7.2 The Effect of Buyback in Equilibrium

To illustrate the effect of buyback on new jet sales, manufacturer profit, and consumer surplus in equilibrium, I simulate equilibrium at the estimated parameters and in a no-buyback counterfactual in which $b_m = 0 \forall m$. In particular, I calculate how much of the increase in demand for new jets from additional upgrades is cannibalized by the substitution away from new jets by first time buyers that results from lower used jet prices. The first three columns of Table 8 record total quantities sold, average prices, total manufacturer profit, and consumer surplus in the buyback

²⁷Since the estimation sample includes a subset of all owners (see Section 2.3) I have to account for used jets flowing in and out of the sample. I achieve this by subtracting from the excess demand given by equation 23 by the net supply of jets into the sample from outside owners. In all counterfactual simulations I keep these quantities fixed. The absolute value of this “outside supply” is 7.1% of M_{jt} on average across models and years. The total net outside supply in each year is on average 0.7% of $\sum_j M_{jt}$. Part of this flow of jets in and out of the sample is due to dealers holding inventory for more than one year.

Table 8: Equilibrium Simulations

	Baseline Parameters			No Heterogeneity		
	No Buyback	Buyback	Δ	No Buyback	Buyback	Δ
Upgrades to New	1304.8	1635.4	330.6	1053.5	1476.5	423.0
Used Jet Supply to First Time Buyers	17834.8	18324.4	489.6	16396.5	16910.5	514.0
New Sales to First time Buyers	6131.6	6004.9	-126.7	6099.7	6090.0	-9.7
Average Used Jet Price (\$ Million)	5.85	5.64	-0.21	56.4	56.3	-0.1
Manufacturer Profit (\$ Billion)	8.0	9.4	1.4	0.5	0.5	0.0
Consumer Surplus (\$ Billion)	215.2	220.3	5.1	9.2	9.4	0.2

Notes: The first three columns record statistics computed under equilibrium simulations at the estimated parameters. The first column records simulations with $b_m = 0$ for all manufacturers. The second column records simulations with b_m set to the estimated values. The third column records the difference between the first and second columns. The fourth, fifth and sixth columns record equivalent numbers for equilibrium simulations under the non-heterogeneity parameters recorded in Appendix Table A.11. All columns are averages over 100 simulations. All figures are totals for the period 1961-2000.

and no-buyback simulations at the estimated parameters and *equilibrium* prices.

Moving from the no-buyback equilibrium to the buyback equilibrium increases the number of new jets purchased as upgrades by 330.6, or 25%. As discussed in Section 6 above, this increase in upgrades to new jets translates into an increase in the supply of used jets. Moving from the no-buyback to buyback equilibrium increases the supply of used jets to first time buyers by 489.6 units. Notice that the increase in supply of used jets is higher than the increase in the number of upgrades to new jets. Per equation 21, this implies that the number of market exits also increases in equilibrium when buyback is introduced. Note that the numbers here differ from those recorded in Table 7 because of equilibrium price effects. In particular, the *equilibrium* increase in upgrades to new is smaller than in Table 7 since the introduction of buyback raises new jet prices and lowers used jet prices.

The increase in used jet supply lowers the average used jet price by 3.6% from \$5.85 million in the no buyback equilibrium to \$5.64 million in the buyback equilibrium. This fall in price induces substitution towards used jets among first time buyers. This cannibalization effect means that 126.7 fewer new jets are purchased by first time buyers in the buyback equilibrium. However, the net effect of buyback on new jet sales is still positive, and buyback increases the total profits of all manufacturers in equilibrium by \$1.4 billion, or 17.5%.

The size of this cannibalization effect depends crucially on the substitution patterns of first time buyers and therefore the degree of heterogeneity in consumer preferences. In the second set of three columns in Table 8, I record equilibrium quantities under buyback and no-buyback simulations

using parameters estimated under the restriction of no-heterogeneity. Under these parameters, introducing buyback increases the supply of used jets to first time buyers by 514 units, but only reduces the number of new jets purchased by first time buyers by 9.7 units. Intuitively, the greater the degree of heterogeneity in consumer preferences, the less likely consumers are to be on the margin between the inside and outside goods, meaning that more of the substitution towards used jets comes from consumers who otherwise would have purchased new jets.²⁸

Finally, note that buyback also increases equilibrium consumer surplus by 2.3%, from \$215.2 billion to \$220.3 billion. This change in consumer surplus is the result of several effects. Consumer welfare is directly increased by lower transaction costs and indirectly increased by allowing consumers to upgrade to their preferred model more frequently. Price changes have offsetting effects, with lower used prices increasing welfare and higher new prices reducing welfare. Note that the magnitude of consumer surplus means that the average surplus generated for one consumer in one year is around \$2.6 million.²⁹

7.3 Optimality of Buyback

Although these results show that buyback increases manufacturer profit, they do not account for any costs incurred by firms from operating buyback programs. Indeed, the notion that the firm is relieving the consumer of some of the transaction costs associated with selling a used jet suggests that the firm itself is taking on these costs. When is it optimal for firms to offer buyback, and how does competition affect this incentive?

The observation that all manufacturers actually operate buyback does not imply that all manufacturers are better off accepting trade-ins than they would be if no manufacturers operated buyback. Consider the game in which manufacturers simultaneously decide whether or not to operate buyback schemes. It may be that this game is a prisoner’s dilemma in which all manufacturers would be better off if they jointly agreed not to accept trade-ins, but offering buyback is a best response to other firms’ policies. On the other hand it could be that operating buyback is a dominant strategy, and that firms would choose to do so even without competitive pressure. Which of these equilibria prevails depends on the costs associated with operating buyback.

To compute ranges of costs for which offering buyback is a dominant strategy for each firm, and ranges of costs for which it is a best response to other firms’ policies but not a dominant strategy, I simulate additional counterfactual equilibria. For each manufacturer, m , I simulate a “unilateral deviation” price equilibrium in which $b_m = 0$ and $b_n = \hat{b}_n \forall n \neq m$. Firm m ’s profit under these unilateral deviations is lower than under the no-buyback equilibrium because it is less likely that

²⁸Note also that the no-heterogeneity parameters generate implausibly small levels of profit and consumer surplus.

²⁹The fact that consumer surplus is substantially larger than firm profits suggests that market power is limited by competition between manufacturers and between manufacturers and the secondary market.

Table 9: Buyback Cost Ranges

Manufacturer	Dominant Strategy Costs		Prisoner's Dilemma Costs		$\frac{b_m}{E(\alpha_i^p hold)}$	$\frac{b_m}{E(\alpha_i^p upgrade)}$
	Min	Max	Min	Max		
Bombardier	0	443.3	443.3	640.0	196.3	305.1
Cessna	0	61.6	61.6	241.7	139.1	216.1
Dassault	0	443.3	443.3	780.8	142.5	221.4
Gulfstream	0	1106.2	1106.2	1458.0	177.7	276.1
IAI	0	86.6	86.6	461.6	194.6	302.5
Raytheon	0	78.0	78.0	355.8	118.0	186.1
Other	0	82.7	82.7	167.3	119.7	183.4
Cessna +	0	110.5	110.5	205.3	139.1	216.1

Notes: Table records computed per-unit buyback costs in thousands of dollars. Max dominant strategy costs are computed by dividing the difference between the buyback and no-buyback equilibrium profit by the number of buyback-eligible upgrades in the buyback equilibrium. Max prisoner's dilemma costs are computed by dividing the difference between buyback and unilateral deviation equilibrium profits by the number of buyback-eligible upgrades in the buyback equilibrium. Profits used in the calculations are recorded in Appendix Table A.13. Cessna + records costs from simulations in which Bombardier's small jets are merged with Cessna. The fifth and sixth columns record ratios of the estimated buyback parameters, b_m , to the expected marginal utility of \$1000 for jet holders and jet upgraders. All figures are totals for the period 1961-2000.

upgraders from other brands with switch brand to m . Appendix Table A.13 records manufacturer specific profit under buyback, no-buyback, and unilateral deviation.

For each manufacturer m , I compute ranges of "buyback costs" under which operating buyback is a dominant strategy - net profit for firm m when all firms offer buyback is greater than profit when no firms offer buyback - and under which it is a best response to other firms - net profit for firm m when all firms offer buyback is greater than unilateral deviation profit. I assume buyback costs to be "per unit" - if a firm offers buyback, they incur a cost for every upgrade that is eligible for buyback (every same brand used-new upgrade).

Table 9 records the computed cost ranges. The first two columns record the range of per-unit buyback costs under which offering buyback is a dominant strategy. The third and fourth columns record the range of per-unit costs under which firms are better off under no-buyback, but offering buyback is a best response to other firms' policies. The fact that all firms offer buyback means that buyback costs are bounded above by the values in the fourth column.

For comparison, the final two columns records the estimated values of $\frac{b_m}{E(\alpha_i^p|hold)}$ and $\frac{b_m}{E(\alpha_i^p|upgrade)}$, the per-unit reduction in transaction costs from buyback evaluated at the average price coefficient among jet holders and upgraders. If buyback was a one to one transfer of transaction costs from consumers to manufacturers, then the cost of buyback to the firm would be equal to the reduction in transaction costs faced by the consumer. The computed cost ranges are consistent with this hypothesis. For three of the manufacturers, Bombardier, Dassault, and Gulfstream, $\frac{b_m}{E(\alpha_i^p|hold)}$ and $\frac{b_m}{E(\alpha_i^p|upgrade)}$ are within the range of dominant strategy costs. For the remaining

firms, the reduction in transaction costs falls in the range of costs for which offering buyback is not a dominant strategy but is a best response to other firms' policies. Of course, the computed cost ranges are also consistent with the cost of buyback to manufacturers being less than the reduction in transaction costs, for example because of efficiency savings from centralizing the exchange of used jets. The per-unit buyback cost firms face will also be reduced by any profit (and increased by any loss) made on the resale of used jets if the resale price is different from the buyback price. Anecdotal evidence (Globe and Mail, 1994; National Post, 1994) suggests that manufacturers often make losses on refurbishment and resale.

Why is buyback is not a dominant strategy for all firms? For instance, Cessna and Dassault have similar values of b_m , but buyback is a dominant strategy for Dassault and not Cessna. Cessna does not benefit from buyback to the same extent as other firms because of the position of its products in the market and the substitution patterns induced by the estimated distribution of consumer preferences. As recorded in Table 1, Cessna is the dominant firm in the small jet market. Cessna's jets are smaller and cheaper than those of other manufacturers, and consumers that select into purchasing Cessna Jets are therefore different from consumers who select into other manufacturers.

This is illustrated by Appendix Figure A.4, which records the distribution of α_i^p among jet holders in 2000, conditional on the brand of jet held. The median value of α_i^p for holders of Cessna jets is 0.213, while the median value for all other manufacturers is between 0.123. Cessna holders are significantly more price sensitive than holders of other jets. First time buyers of Cessna jets are therefore more likely to be on the margin between buying used and new jets than first time buyers of other brands. Indeed, Cessna sells 60.5 fewer new jets to first time buyers in the buyback scenario compared to the no-buyback scenario. This is 59% of the increase in new Cessna jets sold to upgraders. The equivalent cannibalization rate for all other firms is 29%. Furthermore, the selection of buyers into Cessna ownership means that they are unlikely to switch to another, more expensive, brand. In the no-buyback scenario, Cessna receives 174.4 buyback-eligible upgrades, compared to between 7.7 and 156.1 for the other manufacturers. The combination of serving a large number of infra-marginal upgraders and the high degree of new-used substitution in the first time buyer market makes buyback less profitable for Cessna than for other manufacturers.

7.4 Buyback and Market Structure

The fact that substitution patterns and market structure matter for the firm's optimal buyback strategy suggests that *changes* in market structure might change firms' buyback policy. To examine this possibility, I consider a counterfactual acquisition of Bombardier's small jets (marketed as Learjet) by Cessna. Bombardier is the second largest manufacturer of small jets, with a 32% market share, and 84% of all new small jets sold in the data are either Cessna or Bombardier models. I run equilibrium simulations in which the merged firm, which I call "Cessna +" exists

throughout the time period covered by the data. I apply Cessna’s buyback parameter, b_{Cessna} to the merged firm and keep all other parameters and product unobservables as in the baseline estimates.³⁰

The final row of Table 9 records the dominant strategy and prisoner’s dilemma cost ranges for the merged firm, computed as described above. Notice that $\frac{b_m}{E(\alpha_i^p|upgrade)}$ is now *above* the maximum prisoner’s dilemma cost of buyback. This means that, under the assumption that the per-unit cost of buyback is equal to $\frac{b_m}{E(\alpha_i^p|upgrade)}$, it is not a best response for the merged firm to offer buyback. Unlike Cessna in the main simulations, Cessna + can increase profit net of the cost of buyback by setting $b_{Cessna+} = 0$. To understand this, recall that the maximum prisoner’s dilemma cost is higher than the maximum dominant strategy cost because a unilateral deviation from the equilibrium in which all firms offer buyback induces more substitution of demand towards other firms. In the merger counterfactual, Cessna + now owns Bombardier’s small jets, which are the second most popular small jets on the market, and therefore likely a close substitute to Cessna’s own jets. Indeed, in the non-merged case, Bombardier’s profit increases more than all other firms in equilibrium when Cessna does not offer buyback (by \$10 million compared to a maximum of \$5 million for the other firms). By acquiring its closest competitor, Cessna + now internalizes substitution to Bombardier’s small jets, reducing losses from substitution to other-firm jets when the merged firm removes buyback.

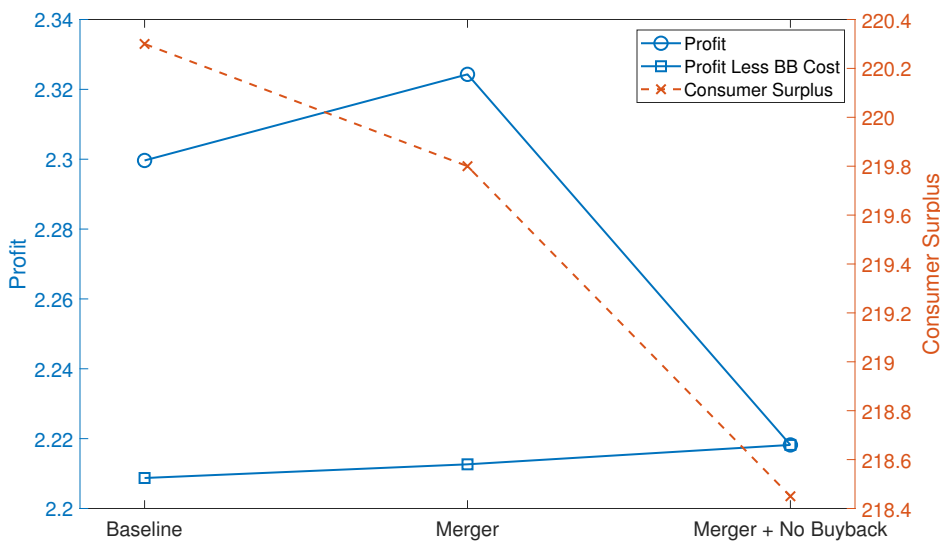
In Figure 3, the solid blue lines illustrate the profit of the merged firm (or the components of the merged firm) before the merger, after the merger with buyback, and after the merger without buyback. Moving from the left to the right, the merger increases profit (before the cost of buyback) by about \$30 million, as the merged firm is able to raise the prices of Cessna and Bombardier small jets and now benefits from more upgrading across its models. The merger also increases profit when the cost of buyback is taken into account. Notice that the increase in profit is rather small in percentage terms. The merged firm still faces competition from the other manufacturers and the secondary market. Because the merged manufacturer sells mostly lower-cost small jets, used units are likely to be close substitutes and discipline how much the firm can raise prices.

Setting $b_m = 0$ for the merged firm then reduces gross profit by about \$100 million, but increases profit net of the cost of buyback by about \$10 million. Although the merged manufacturer sells fewer new jets when $b_m = 0$, the loss in profit from jet sales is outweighed by the reduction in buyback cost.

The dashed red line in Figure 3 records the effect of the merger on consumer surplus. Moving from left to right, the merger with buyback reduces consumer surplus by around \$600 million. Consumers are worse off because of slightly higher new jet prices. This effect may be tempered by the change in ownership structure allowing consumers to take advantage of buyback when

³⁰In particular, this means that the utility of small Bombardier jets includes the firm-specific preference $\nu_i^{Bombardier}$ and not ν_i^{Cessna} .

Figure 3: Simulated Effects of Merger



Notes: Solid lines report profit and correspond to the left y-axis. Dashed line reports consumer surplus and corresponds to the right y-axis. The blue circles record profit for new sales of Cessna and small Bombardier jets under three scenarios: the baseline, a merger on Cessna and Bombardier’s small jets in which $b_m = b_{Cessna}$ for the merged firm, and a merger in which $b_m = 0$ for the merged firm. The Blue squares record profit less the cost of buyback, where the cost of buyback is taken to be $\frac{b_{Cessna}}{E(\alpha_i^p | Upgrade)}$. The red crosses record consumer surplus. All results are averages over 100 simulations. All figures are totals for the period 1961-2000.

upgrading from a small Cessna to a small Bombardier jet. Setting $b_m = 0$ for the merged firm reduces consumer surplus by a further \$1.4 billion.

This merger simulation suggests that if Cessna had acquired Bombardier’s small jets, the merged firm would have found it optimal to remove the buyback incentives offered to consumers. This removal of buyback would reduce consumer surplus significantly relative to the effect of the merger alone - ignoring the change in equilibrium buyback policy would understate the effect on consumer surplus by 70%.

These conclusions rest on the assumed cost of buyback to the firm. The merged firm would only find it optimal not to offer buyback at a per-unit buyback cost between \$205.3 thousand and \$241.7 thousand, a range which includes the benefit of buyback to the consumer, $\frac{b_m}{E(\alpha_i^p | upgrade)}$. Furthermore, the binary comparison between buyback and no-buyback precludes intermediate policies, for example changing the terms of the buyback policy to cover only certain models. Despite these caveats, the merger simulation illustrates that changes in market structure can change firms’ profit maximizing buyback policies, and that ignoring the effect of market structure on buyback can lead to misleading conclusions about consumer welfare. This highlights the importance of accounting for firms’ incentives to engage with the secondary market when evaluating the effects of entry, exit, and mergers in durable goods markets.

8 Conclusion

When manufacturers of durable goods engage with secondary markets they face a trade off between encouraging consumers to upgrade to new units and facilitating trade in used units. Previous studies have shown, theoretically and using models calibrated to industry data, that manufacturers may have an incentive to increase the liquidity of secondary markets, even though this can lead to substitution of first time buyers away from new goods. One way that manufacturers do this in reality is by buying back and reselling used units from upgrading consumers. In this paper, I estimate the effect of these manufacturer buyback policies on demand and supply in the market for business jets and separately quantify the increase in demand for upgrades due to buyback and the equilibrium decrease in new jet sales among first time buyers. I then show how offering buyback can arise as an equilibrium strategy, and how this depends on demand substitution patterns and market structure.

I estimate a dynamic demand model in which consumers enter the market with heterogeneous preferences and hold jets over time. I estimate the model parameters by matching aggregate market shares and micro-moments in a GMM framework. Relying on assumptions about the structure of buyback policies allows me to estimate the size of transaction costs and the effect of buyback on demand without observing exogenous variation in policies over time or across manufacturers. In particular, matching the annual upgrade probability allows me to identify the size of transaction costs, and matching the difference in the probability of new jet purchases between same brand and different brand upgrades allows me to estimate the effect of buyback for each brand.

I simulate equilibrium in the new and used jet markets to illustrate how the effects of buyback on demand and supply interact to affect manufacturer profit in equilibrium. The direct effect buyback on the demand for upgrades to new units increases revenue. This is counteracted by the equilibrium effect on used jet prices, which encourages substitution towards used jets. I find that this equilibrium cannibalization effect reduces the increase in the number of new jet sales from buyback by 38%.

Buyback increases equilibrium profit for the industry as a whole by 17.5% (gross of any cost of buyback). However, this masks substantial heterogeneity across manufacturers. To illustrate this, I compute threshold “buyback cost” levels under which manufacturers are better off when they all offer buyback relative to the no-buyback counterfactual, and under which offering buyback is a best response to other firms’ policies. If the cost of buyback to manufacturers is equal to the reduction in transaction costs, then three of the top six firms are in the dominant strategy region. The remaining firms only offer buyback as a best response to other firms’ policies. Whether or not buyback is a dominant strategy for a firm depends on the substitutability of its new units with used units among first time buyers.

Finally, I illustrate how changes in market structure can affect equilibrium buyback by simulating

a merger of Bombardier's small jets into Cessna. In equilibrium, it is no longer a best response for the merged firm to offer buyback. In particular, the merged firm now controls most of the small jet market, so that when it unilaterally deviates from the buyback equilibrium, there is limited substitution to other manufacturers. Of the total effect of the merger on consumer surplus, 70% is due to the removal of buyback by the merged firm. Ignoring the effect of market structure on buyback would result in substantial underestimation of the effect of the merger.

This paper provides a framework for thinking about the equilibrium effects of buybacks in a durable goods industry with an active secondary market and transaction costs. Engagement in the secondary market can take other forms, for example explicit price discrimination between first time buyers and upgraders, which the framework presented in this paper can be adapted to analyze. Generally, firms' choice of profit maximizing secondary market policies will depend on market structure, cross-brand substitution patterns, and the cannibalization effect which depends on substitution between new and used units. Changes in market structure can change equilibrium secondary market policies, with potentially large effects on consumer welfare. Indeed, the results in this paper suggest that changes in secondary market engagement as the result of a merger can have first order effects on consumer welfare. This has direct implications for antitrust analysis of durable goods markets. The typical focus of such analysis on price effects alone may substantially understate consumer welfare losses from mergers when secondary markets are important.

References

- Berry, S., J. Levinsohn, and A. Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841-890
- Carranza, J. E. 2010. "Product innovation in market equilibrium: the case of digital cameras." *International Journal of Industrial Organization* 28: 604-618
- Chen, J., and M. Sacks. 2021. "Reimbursing Consumers' Switching Costs in Network Industries." *Working Paper*
- Chen, J., S. Esteban, and M. Shum. 2013. "When Do Secondary Markets Harm Firms?" *American Economic Review* 103 (7): 2911-2934
- Dubé, J. P., G. J. Hitsch, and P. E. Rossi. 2010. "State Dependence and Alternative Explanations for Consumer Inertia." *RAND Journal of Economics* 41 (3): 417-445
- Esteban, S, and M. Shum. 2008. "Durable-goods oligopoly with secondary markets: the case of automobiles." *RAND Journal of Economics* 38 (2): 332-354
- Fudenberg, D., and J. Tirole. 1998. "Upgrades, Trade-ins, and Buybacks." *RAND Journal of Economics* 29 (2): 235-258
- Gandhi, A., Z. Lu, and X. Shi, 2019. 2019. "Estimating Demand for Differentiated Products with Zeroes in Market Share Data." *Working Paper*
- Gavazza, A. 2016. "An Empirical Equilibrium Model of a Decentralized Asset Market." *Econometrica* (Forthcoming)
- Gavazza, A., A. Lizzeri, and N. Roketskiy. 2014. "A Quantitative Analysis of the Used Car Market." *American Economic Review* 104 (11): 3668-3700
- Gilligan, T. W. 2004. "Lemons and Leases in the Used Business Aircraft Market." *Journal of Political Economy* 112 (5): 1157-1180
- Gillingham, K., F. Iskhakov, A. Munk-Nielsen, J. Rust, and B. Schjerning. 2021. "Equilibrium Trade in Automobiles." *Working Paper*
- Gowrisankaran, G., and M. Rysman. 2012. "Dynamics of consumer Demand for New Durable Goods." *Journal of Political Economy* 120 (6): 1173-1219
- Hendel, I., and A. Lizzeri. 1999. "Interfering with Secondary Markets." *RAND Journal of Economics* 30 (1): 1-27
- Keane, M. "Modeling heterogeneity and State Dependence in Consumer Choice Behavior." *Journal of Business Economics & Statistics* 15 (3): 310-327
- Lee, R. 2013. "Vertical Integration and Exclusivity in Platform and Two-Sided Markets." *American Economic Review* 103 (7): 2960-3000

- Nevo, A. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica* 69 (2): 307-342
- Petrin, A. 2002. "Quantifying the Benefits of New Products: The Case of the Minivan." *Journal of Political Economy* 110 (4): 705-729
- Porter, R. H., and P. Sattler. 1999. "Patterns of Trade in the Market for Used Durables: Theory and Evidence." *NBER Working Paper*
- Quan, T. W., and K. R. Williams, 2018. "Product variety, across-market demand heterogeneity, and the value of online retail." *RAND Journal of Economics* 49 (4): 877-913
- Rao, R. S., O. Narasimhan, and G. John. 2009. "Understanding the Role of Trade-ins in Durable Goods Markets: Theory and Evidence." *Marketing Science* 28 (5): 950-967
- Rust, J. P. 1986. "When Is It Optimal to Kill Off the Market for Used Durable Goods?" *Econometrica* 54 (1): 65-86
- Schiraldi, P. 2011. "Automobile Replacement: A Dynamic Structural Approach." *RAND Journal of Economics* 103 (7): 2911-2934
- Stolyarov, D. 2002. "Turnover of used Durables in a Stationary Equilibrium: Are Older Goods Traded More?" *Journal of Political Economy* 110 (6): 1390-1413
- Train, K. 2003. "Discrete Choice Methods with Simulation." New York: Cambridge University Press.

A Appendix

A.1 Jet Holding Algorithm

The model assumes that jet owners can only hold one jet at a time. In the data, 21.4% of owners (excluding manufacturers and dealers) own more than one jet at some point. To deal with this, I construct a mapping of jet owners to *single* jets for each year by following the *first jet owned* by each owner and its *successors*. when a jet owner holds multiple jets at the same time I split that owner into two owners etc. The algorithm used to construct this panel is as follows.

For each owner, I record the set of jets owned in December of each year. I assign the jet owned in the first year I observe an owner to that owner for that year. If the owner has two jets in December of the first year, I assign the jet that was purchased first (i.e. earlier in the year) to that owner for that year. I then look at the second year for that owner. If the owner still owns the jet I assigned to them in the first year, I assign this jet to them in the second year, regardless of any other jets they might own. If the owner no longer owns the jet assigned in the first year, and has acquired some other jet in the second year, I record the owner as having upgraded to the new jet in the second year. For cases where more than one jet is purchased, I assign the jet that was purchased earlier in the year. I repeat this procedure for subsequent years. If I observe the previously held jet being sold and no new jet being purchased, I record the owner as exiting the market. I then repeat the entire procedure, starting at the first year the owner is observed in the data for the jets that were not assigned during the first iteration. The second iteration generates a second sequence of jets and is recorded as a second owner in the final data. I repeat this until all the jet-years in the data have been assigned to an owner. If there is a gap in ownership - say an owner exits the market for a period of time and then reenters, I record these two stints as two separate owners. This procedure generates a panel of jet owners observed once a year, holding at most one jet each year.

A.2 The Effect of Buyback on Equilibrium Revenue: Theory

To see why buyback can only increase revenue, consider the following single firm example. There are two types of goods: used and new. Let p_U be the price of a used unit and p_N be the price of a new units. There are M_1 jet holders who own used units. Jet holders can choose to upgrade to a new unit, hold their used unit, or sell their used unit and exit the market. Demand for new units from jet holders is $D_1^N(p_U, p_N, b)$, demand for holding is $D_1^U(p_U, p_N, b)$, and demand for exiting the market is $D_1^{exit}(p_U, p_N, b)$ where $M_1 = D_1^N + D_1^U + D_1^{exit}$, $\frac{\partial D_1^N}{\partial p_N} < 0$, $\frac{\partial D_1^N}{\partial p_U} > 0$, $\frac{\partial D_1^U}{\partial p_N} > 0$, $\frac{\partial D_1^U}{\partial p_U} < 0$, $\frac{\partial D_1^{exit}}{\partial p_N} > 0$, and $\frac{\partial D_1^{exit}}{\partial p_U} > 0$. The variable b indicates buyback, where $\frac{\partial D_1^N}{\partial b} > 0$, $\frac{\partial D_1^U}{\partial b} < 0$ and $\frac{\partial D_1^{exit}}{\partial b} < 0$.

There are M_0 potential first time consumers who do not own used units. Demand for new units from first time customers is $D_0^N(p_U, p_N)$, demand for used units is $D_0^U(p_U, p_N)$, and demand for not buying (exiting the market) is $D_0^{exit}(p_U, p_N)$ where $M_0 = D_0^N + D_0^U + D_0^{exit}$. Partial derivatives with respect to price have the same sign as for jet holders.

New jets, N , are supplied by the manufacturer according to supply function $S_N(p_N)$, where $\frac{\partial S_N(p_N)}{\partial p_N} > 0$. The supply of used jets is given by the sum of upgrades to new and exits, $S_U(p_N, p_U, b) = D_1^N(p_U, p_N, b) + D_1^{exit}(p_U, p_N, b)$. In equilibrium, equilibrium quantities sold are given by,

$$Q_N = S_N(p_N) = D_1^N(p_U, p_N, b) + D_0^N(p_U, p_N) \quad (25)$$

$$Q_U = S_U(p_N, p_U, b) = D_0^U(p_U, p_N). \quad (26)$$

Proposition 1. *Let equilibrium manufacturer revenue be $R(b) = Q_N p_N$. $\frac{\partial R(b)}{\partial b} \geq 0$.*

Proof. Suppose $\frac{\partial R(b)}{\partial b} < 0$. this implies that there is some $b' > b$ such that $R(b') < R(b)$. Let primed variables represent equilibrium under b' and unprimed variables represent equilibrium under b .

Since $S_N(p_N)$ does not depend on b , it must be that $p'_N < p_N$ and $Q'_N < Q_N$. By equation 25 there are three possible cases:

1. $D_1^N(p'_U, p'_N, b') < D_1^N(p_U, p_N, b)$ and $D_0^N(p'_U, p'_N) \geq D_0^N(p_U, p_N)$
2. $D_1^N(p'_U, p'_N, b') \geq D_1^N(p_U, p_N, b)$ and $D_0^N(p'_U, p'_N) < D_0^N(p_U, p_N)$
3. $D_1^N(p'_U, p'_N, b') < D_1^N(p_U, p_N, b)$ and $D_0^N(p'_U, p'_N) < D_0^N(p_U, p_N)$.

I take these cases one at a time.

Case 1: Since $p'_N < p_N$ and $b' > b$, for $D_1^N(p'_U, p'_N, b') < D_1^N(p_U, p_N, b)$ to hold it must be that $p'_U < p_U$. This means that $D_1^{exit}(p'_U, p'_N, b') < D_1^{exit}(p_U, p_N, b)$ and $D_0^{exit}(p'_U, p'_N) < D_0^{exit}(p_U, p_N)$.

Differencing equation 26 and substituting in $D_0^U = M_0 - D_0^N - D_0^{exit}$, we have,

$$\begin{aligned} & -D_0^N(p_U, p_N) + D_0^N(p'_U, p'_N) - D_0^{exit}(p_U, p_N) + D_0^{exit}(p'_U, p'_N) \\ & = D_1^N(p_U, p_N, b) - D_1^N(p'_U, p'_N, b') + D_1^{exit}(p_U, p_N, b) - D_1^{exit}(p'_U, p'_N, b') \end{aligned}$$

This implies $D_0^N(p'_U, p'_N) - D_0^N(p_U, p_N) > D_1^N(p_U, p_N, b) - D_1^N(p'_U, p'_N, b')$, which implies $Q'_N > Q_N$, a contradiction.

Case 2: Since $p'_N < p_N$, for $D_0^N(p'_U, p'_N) < D_0^N(p_U, p_N)$ to hold it must be that $p'_U < p_U$. Following the same steps as above we can obtain $D_0^N(p_U, p_N) - D_0^N(p'_U, p'_N) < D_1^N(p'_U, p'_N, b') - D_1^N(p_U, p_N, b)$, which implies $Q'_N > Q_N$, a contradiction.

Case 3: Again, it must be the case that $p'_U < p_U$, so $D_0^{exit}(p'_U, p'_N) < D_0^{exit}(p_U, p_N)$. Since $D_0^U = M_0 - D_0^N - D_0^{exit}$ it must be that $D_0^U(p'_U, p'_N) > D_0^U(p_U, p_N)$. However, it is also the case

that $D_1^{exit}(p'_U, p'_N, b') < D_1^{exit}(p_U, p_N, b)$, and therefore $S_U(p'_U, p'_N, b') < S_U(p_N, p_U, b)$, contradicting equation 26.

In each case, I derive a contradiction, so it must be that $\frac{\partial R(b)}{\partial b} \geq 0$. □

The intuition for the proof is that for revenue to fall, it must be that total quantity demanded for the new good is lower and the price of the new good is lower when buyback is increased. This can only be the case if the price of the used good falls sufficiently for consumers to substitute away from the new good, but if both prices fall then consumers also substitute away from exiting the market. This means that the quantity supplied of used goods is lower - jet holders are less likely to upgrade or exit - but quantity demanded for used goods is higher - first time buyers are also less likely to buy new goods or exit/buy nothing. Since there is a fixed number of used jets, this cannot be an equilibrium.

A.3 Buyback Heterogeneity

A significant assumption of the model presented in Section 4 is that the buyback parameters, b_m , are constant across consumers, jets, and time within a manufacturer. This assumption of a constant buyback policy within firms makes the counterfactual analysis in which I shut down buyback for one or more firms tractable by reducing the space of buyback policies (buyback is either “on” or “off” for the entire time period covered by the data). In principle, it would be possible to relax this assumption and estimate richer heterogeneity in buyback parameters. This would require adding further micro moments to the GMM objective function. For example, to vary b_m over time would require analogues of moments 6-12 in Table 17 conditioned on different time periods. Because the dataset is relatively small, conditioning these moments further would lead to imprecision in the micro-moments, and heterogeneity in b_m would likely not be estimated precisely (for example, there are only 608 upgrades from used Dassault jets, 63 of which are to new jets of the same brand).

To test whether this restriction on b_m is reasonable, I present descriptive statistics that examine the extent of apparent heterogeneity in buyback across three dimensions: jet owners, time, and jet models.

Figure A.1 reports statistics on upgrades by consumer type. I split all consumers in the estimation sample into terciles of maximum fleet size. Recall that the estimation data, in which each consumer holds one jet at a time, is constructed from raw data in which some firms hold multiple jets. This Figure examines whether upgrade behavior differs across consumers that hold fleets of different sizes. The left panel records the share of buyback-eligible upgrades (upgrades from a used unit to a new unit of the same brand) in which the used unit is: (1) sold to an independent dealer, (2) sold back to the manufacturer, or (3) transferred directly to another consumer (“back to back”).

Recall that both (2) and (3) could be manufacturer facilitated “buyback” transactions. I find that the share of transactions in each of these three categories is broadly similar across consumers with different fleet sizes. Consumers with large fleets (> 7 jets) are somewhat more likely to sell jets directly rather than using a dealer. Although these statistics do not enter the estimation of the model, if it is easier for some companies (e.g. those with large fleets) to sell directly on the secondary market, this should show up in upgrade frequency and be captured by heterogeneity in ν_i^T .

The right panel shows the same statistics for all other (non-buyback eligible) upgrades. Notice that for all three fleet size categories the number of sales back to manufacturers is very small, and the share of sales to independent dealers is larger than in the left panel. Comparing the two panels, the upgrade transactions of consumers of all fleet sizes respond similarly to buyback eligibility.

Table A.1 records analogues of the statistics in Table 3 for all manufacturers in different time periods. For each of the three time periods reported, the majority of observed buybacks involve upgrades from used to new jets and are of the manufacturer’s own brand. However, the share of used-new upgrades are higher later in the sample period. Similarly, the share of potential buybacks that are observed to be sold back to the manufacturer is lower earlier in the sample, rising from 13% before 1980 to 39% after 1990. These patterns raise the possibility that buyback policies changed over time.

To investigate this possibility further, Table A.2 breaks down changes in observed buyback over time by manufacturer. Each column records coefficients from a separate regression. Observations are all buyback-eligible upgrades for the manufacturer indicated by the column header. The dependent variable is an indicator for whether the used unit in the upgrade was observed to be sold back to the manufacturer. Regressors are time dummies for each decade in the data. Significant coefficient estimates would suggest that, for a specific manufacturer, upgrades are more likely to be bought back in that period of time. Note that there is at least one omitted category for each regression, and some regressions omit two categories because there are no observations in the first decade of the sample.

Although the coefficients on 1991-2000 are positive for all manufacturers, only one (Cessna) is significant at the 5% level. Notice that the number of observations in each regression is quite small. This illustrates the difficulty in precisely estimating heterogeneity in buyback at the manufacturer-decade level. Furthermore, the statistics reported in Tables A.1 and A.2 examine *observed* sales to manufacturers only. As discussed in Section 3.1, these statistics may not capture all manufacturer-facilitated upgrades that I want to classify as “buyback”, and this variation is not used in estimation of the model.

To directly examine heterogeneity in the differences in upgrade probability used to identify buyback in the model, I replicate panel A of Table 4 for different time periods, and record these in Table A.3. Recall that b_m is identified by the difference in new jet (or used jet) share between same

brand and other brand upgrades. In Table 4 I report a significant difference of 11.7 percentage points. Analogous numbers for the four decades of the sample reported in Table A.3 are: 7.8, 7.7, 10.7, and 12.1 percentage points. These estimates are statistically significant at at least the 10% level for 1971 onward. The point estimates show a slight upwards trend over time, but it is not possible to reject the hypothesis that any two of these estimates are equal.

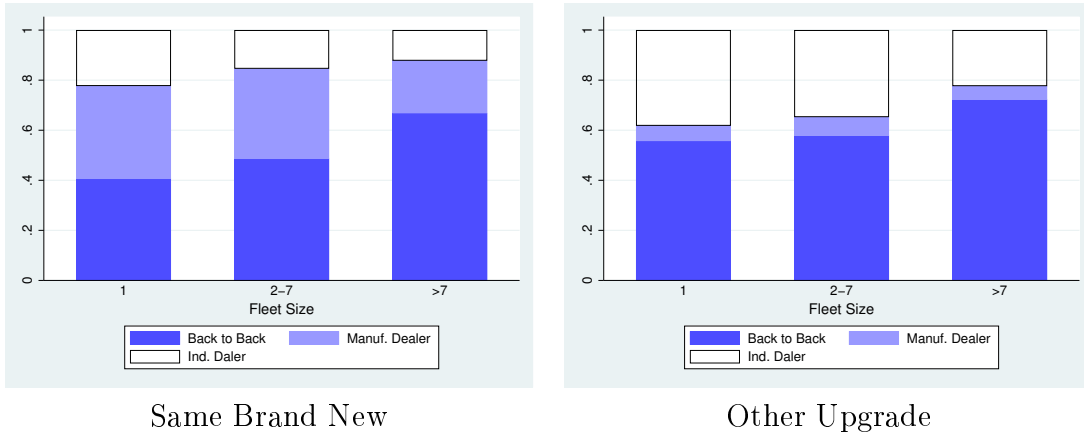
As with Table A.2, there is some evidence of changes in the effects of buyback over time, but they are difficult to measure precisely because of small sample size. Note that these descriptive exercises do not take into account the changing composition of consumer types over the time, which is accounted for in the model and might drive some of these patterns (if for example early adopters have low ν_i^T). Because of this, I view my estimates of b_m in the main specification as average effects of buyback on utility over the entire sample period.

Finally, I examine heterogeneity in observed buyback by characteristics of jet traded in and purchased. In principle, the optimal buyback policy need not be identical for all jets, as I assume it to be in the model. To test for this type of heterogeneity, Table A.4 reports coefficients from regressions of an indicator for observed buyback (sale of the used unit back to the manufacturer) on jet characteristics. The sample is all buyback-eligible upgrades. The first three columns report coefficients from a single regression, arranged over three columns for clarity. They report the coefficient on indicators for the interaction of the size of the jet sold and the size of the jet purchased. Buyback is observed significantly more frequently for upgrades from small and medium jets to new medium jets. The pattern of coefficients suggests that “down-upgrades” from larger to smaller jets are less likely to use buyback, and “up-upgrades” from smaller to larger jets are more likely to use buyback, consistent with heterogeneous buyback policies being used to encourage consumers to “size up”. However, most of the coefficients are not statistically significant.

In the fourth column I report coefficients from a regression of observed buyback frequency on jet price and model age. Model age is measured as the number of years since the characteristics (engine power, range, etc.) of the jets in a manufacturer-size category have changed. That is, how long has it been since a manufacturer’s production model has been updated. I find that, although there is no effect of price, observed buybacks are more common for older jet models. The final three columns include both sets of regressors and report similar results, though with slightly reduced significance.

The results reported in Table A.4 point to some heterogeneity in buyback policy across models within the same manufacturer (all regressions include manufacturer and year fixed effects). It may be that firms increase the generosity of buyback programs as their jet models age and become less competitive with other firms’ new models. Solving jointly for the firm’s dynamically optimal buyback and model upgrade policy is an interesting challenge that is outside the scope of this paper.

Figure A.1: Upgrade Type Heterogeneity by Fleet Size



Notes: The left panel records the share of upgrades from used units to new units of the same brand in which the used unit is sold to an independent dealer, manufacturer dealer, or another consumer (“back to back”). Each bar presents these statistics for consumers whose maximum number of jets held simultaneously falls in the “Fleet Size” bin denoted on the x-axis. The right panel repeats this exercise for upgrades that are no same-brand used-new.

Table A.1: Buyback Patterns for Different Sub-Periods

Manufacturer	Years	Share of Buybacks		Share of Potential Buybacks
		Upgrades to New	Own Brand	Sold to Manufacturer
All	1961-1980	67%	82%	14%
All	1981-1990	76%	91%	19%
All	1991-2000	90%	87%	39%

Notes: Statistics identical to those recorded in Table 3, recorded separately for transactions in three time periods. “All” manufacturers includes only the top si manufacturers as defined in the text.

Table A.2: Firm-Specific Buyback Heterogeneity by Period

Dependent Variable: Manufacturer:	Used Jet Bought Back by Manufacturer						
	Bombardier	Cessna	Dassault	Gulfstream	IAI	Raytheon	Other
1961-1970
1971-1980	0.063 (0.340)	.	0.321 (0.410)	.	.	0.000 (0.324)	0.125 (0.120)
1981-1990	0.121 (0.330)	0.243 (0.177)	0.093 (0.407)	0.387 (0.299)	0.000 (0.186)	0.214 (0.283)	0.000 (0.132)
1991-2000	0.513 (0.323)	0.407** (0.162)	0.288 (0.407)	0.480 (0.294)	0.167 (0.186)	0.291 (0.259)	0.000 (0.270)
<i>N</i>	168	235	135	84	15	156	32

Notes: Each column records estimated coefficients from a regression of an indicator for observed buyback on decade indicators. The sample for each regression is all used-new upgrades of the same brand for the brand indicated in the column. The dependent variable is equal to 1 when the used unit is sold back to the manufacturer. Standard errors are in parentheses. * indicates significance at the 10% level. ** indicates significance at the 5% level. *** indicates significance at the 1% level.

Table A.3: Time Heterogeneity in Upgrade Shares by New and Same Brand

	(A) 1961-1970				(B) 1971-1980			
	Used	New	Diff	<i>N</i>	Used	New	Diff	<i>N</i>
Same Brand	0.243	0.757	-0.514	37	0.436	0.564	-0.128	326
Different Brand	0.321	0.679	-0.358	28	0.513	0.487	0.026	271
Diff	-0.078	0.078	-0.156		-0.077	0.077	-0.154	
SE	(0.113)	(0.113)	(0.226)		(0.041)	(0.041)	(0.082)	

	(C) 1981-1990				(D) 1991-2000			
	Used	New	Diff	<i>N</i>	Used	New	Diff	<i>N</i>
Same Brand	0.667	0.332	0.335	721	0.672	0.328	0.344	1188
Different Brand	0.774	0.226	0.548	796	0.793	0.207	0.586	1153
Diff	-0.107	0.107	-0.214		-0.121	0.121	-0.242	
SE	(0.023)	(0.023)	(0.046)		(0.018)	(0.018)	(0.036)	

Notes: Analysis performed on estimation sample. Each panel records the share of new and used jets purchased among all first time and upgrade purchases made by all consumers. Each panel subsets the sample to transactions in the indicated date range. Standard errors of the difference in means for each column on parentheses.

Table A.4: Buyback Heterogeneity by Jet Types

Upgrading From:	Dependent Variable: Observed Buyback					
	Upgrading to			Upgrading to		
	Small	Medium	Large	Small	Medium	Large
Small	0.039 (0.078)	0.270*** (0.092)	0.013 (0.156)	0.039 (0.112)	0.244** (0.121)	0.193 (0.197)
Medium	-0.006 (0.088)	0.197** (0.084)	0.109 (0.111)	-0.014 (0.113)	0.201* (0.108)	0.132 (0.139)
Large	-0.009 (0.153)	-0.064 (0.114)		-0.215 (0.219)	-0.017 (0.188)	
Price				-0.003 (0.003)	-0.000 (0.004)	
Model Age				0.006*** (0.002)	0.006*** (0.002)	
<i>N</i>	825			746	746	

Notes: Table reports coefficients from a regressions of an indicator for observed buyback on jet characteristics. The sample includes all used-new same brand upgrades. The first three columns all correspond to the same regression, with estimated coefficients on interactions of sold and purchased jet size are arranged in a 3x3 matrix. Similarly, the last three columns correspond to one regression with coefficients arranged in a 3x3 matrix. All regressions include year and manufacturer fixed effects. Standard errors are in parentheses. * indicates significance at the 10% level. ** indicates significance at the 5% level. *** indicates significance at the 1% level.

A.4 Alternative Specifications

In this section I report results from two alternative specifications.

First, I estimate the main model under an alternative market size assumption. In the main specification, I use the number of consumers who buy jets in the next five years as the size of the first time buyer market in year t , M_{0t} . This choice has the problem of being a function of future market outcomes. As an alternative, I estimate a model in which M_{0t} is set using counts of active public companies at each year, t . In particular, I use data from the Center for Research on Security prices to generate a count of the number of publicly listed firms each year. This alternative approach has the downside of not counting potential private owners.

The overall magnitudes of the baseline and alternative market sizes are fairly similar. The average value of M_{0t} in the baseline specification is 3125, and the average alternative market size is 2519. Column two of Table A.11 records estimated parameters using the alternative market sizes. They are broadly similar to the main estimates. The first three columns of Table A.5 repeat the simulations reported in Table 7 using the alternative market share estimates. The simulated demand patterns are similar.

Second, I estimate the main model using a different sample of consumers. As discussed in Section 2.3, the estimation sample contains private owners, corporations, and air transport companies. As I illustrate in Table 2, these three groups of owners have similar holding and upgrading patterns. However, one might suspect that air transport companies have different preferences than private users and should be treated separately. To test the sensitivity of the main results, I estimate the baseline model dropping air transport companies from the set of consumers. This reduces the number of consumers in the estimation sample by 15%.

The third column of Table A.11 records the estimated parameters for this specification. Parameter estimates are close to the main specification. One notable difference is that the buyback parameters, b_m , are smaller in the alternative specification (although for only one firm, Bombardier, is the difference statistically significant). In the second set of three columns in Table A.5 I present demand simulations of this specification with and without buyback. The demand patterns are similar to those in the main specifications, with absolute values somewhat lower, consistent with a smaller consumer sample.

Table A.5: Demand Simulation of Alternative Specifications

	Alternative Market Size			Excluding Air Transport		
	No Buyback	Buyback	Δ	No Buyback	Buyback	Δ
(1) Upgrades to New	1273.9	1602.5	328.6	1180.4	1441.9	261.5
(2) Upgrades to Used	3280.3	3148.3	-132	2115.8	2019.3	-96.5
(3) Exits	17334.0	17327.0	-17	13156.0	13146.0	-10
Used Jet Supply to First Time Buyers = (1) + (3)	18607.9	18929.5	311.6	14336.4	14587.9	251.5

Notes: Table reports statistic analogous to those in Table 7 for alternative specifications. The first three columns report outcomes simulated using parameters estimated under an alternative market size assumption. In particular, I use the count of all publicly listed US firms as the number of potential first time buyers, M_{ot} , each year. The second set of three columns report simulated outcomes using parameters estimated using a reduced sample that drops air transportation companies from the set of consumers. All columns are averages over 100 simulations. All figures are totals for the period 1961-2000.

A.5 Additional Tables and Figures

Table A.6: Aircraft Characteristics

Segment Model Years		All		Small	Medium	Large
		< 1990	≥ 1990			
Range	Mean	2351.7	2882.3	1779.5	2401.2	3898.7
	SD	(812.8)	(978.4)	(272.0)	(578.6)	(461.8)
Power	Mean	21.5	25.7	13.3	18.7	46.0
	SD	(12.8)	(14.6)	(1.7)	(3.3)	(10.4)
Max. Weight	Mean	13123.2	14844.1	7204.5	12319.3	26156.0
	SD	(7289.8)	(8292.2)	(984.5)	(2356.7)	(6024.8)
New Price (\$ Millions)	Mean	8.987	17.896	5.168	10.379	27.053
	SD	(6.931)	(10.894)	(1.871)	(5.592)	(7.787)
1 Year Used Price (\$ Millions)	Mean	8.270	17.037	4.955	9.873	26.177
	SD	(6.490)	(10.652)	(1.711)	(5.440)	(7.960)
5 Year Used Price (\$ Millions)	Mean	5.523	12.786	3.558	7.197	20.813
	SD	(4.711)	(8.930)	(1.344)	(4.417)	(7.380)
N		218	109	104	152	71

Notes: An observation is a manufacturer-segment-year. Year refers to the year of manufacture. Characteristics are averaged over models within each manufacturer-segment-year (for example if there are multiple large 1990 Bombardier jets, their characteristics are averaged and treated as one observation). Prices are in millions of year 2009 \$. Columns 1 and 2 record the mean and standard deviation of model characteristics for models manufactured before and after 1990. Columns 3 to 6 record the mean and standard deviation of model characteristics by market segment (jet size).

Table A.7: Multiple Jet Ownership

	Share of Owner-Years		
	Corporations	Individuals	Air Transport
Multiple Jets Held	16.25%	5.64%	27.12%
At Least One Jet Purchased	26.71%	26.94%	28.81%
Multiple Jets Purchased	1.66%	0.57%	4.23%
Share of Owner Years	79.00%	6.82%	14.18%

Notes: Figures calculated using estimation sample before owners with multiple jets are split into separate owners, as described in the text.

Table A.8: First Stage and 2SLS

Dependent Variable:	First Stage				OLS	2SLS
	Price				$\log(s_{jt}) - \log(s_{0t})$	
Price					-0.023*** (0.002)	-0.118*** (0.008)
Substitute Jets (1)	-0.021*** (0.001)					
Substitute Jets (2)		-0.027*** (0.002)				
Lagged Price			-0.004 (0.004)	0.042*** (0.005)		
1st Stage F-Stat						166.214
Jet Model and Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: An observation is a model-year. Substitute Jets (1) is the number of currently held jets in the same category of the same age. Substitute Jets (2) is the number of currently held jets in the same category one year older. Regressions all contain controls for jet characteristics, GDP growth, and manufacturer dummies. The final column is a 2SLS regression using all three instruments. Standard errors are in parentheses. * indicates significance at the 10% level. ** indicates significance at the 5% level. *** indicates significance at the 1% level.

Table A.9: Mean Utility and Cost Regressions

Dependent Variable:	Mean Utility		Marginal Cost	
Range	0.622*** (0.027)		29.011*** (4.274)	8.678** (3.985)
Power	0.030*** (0.006)		-0.335 (0.505)	6.866*** (0.714)
Max Weight	-0.109*** (0.009)		4.250*** (1.018)	2.434** (1.121)
Age	-0.058*** (0.001)	-0.083*** (0.001)		
Mean Utility			25.913*** (3.773)	8.830*** (3.026)
New Jets Only	No	No	Yes	Yes
Year FE	No	Yes	Yes	Yes
Jet FE	Yes	No	No	No
Manuf-Segment FE	No	Yes	No	Yes
N	5874	5874	339	339


Notes: The first two columns record coefficients from regressions of the estimated values of mean utility, $\int \gamma_{ijt} dF(\nu_i)$ on jet characteristics. Observations are jet-years, and the sample includes all new and used jets in the data. The third and fourth columns record coefficients from regressions of estimated marginal costs, c_j , on jet characteristics. The sample includes all new jets in the data. Standard errors are in parentheses. * indicates significance at the 10% level. ** indicates significance at the 5% level. *** indicates significance at the 1% level.

Table A.10: Micro-Moments

Moment	Data	Model	Moment	Data	Model
1	0.043	0.041	13-15	9.675	11.766
				9.347	12.034
2	0.133	0.132		8.087	8.767
3	0.301	0.325	16-18	0.099	0.120
				0.131	0.113
4	0.269	0.2693		0.156	0.167
5	0.503	0.460	19	0.942	0.914
6-12	0.070	0.072	20	0.529	0.410
	0.042	0.053			
	0.066	0.090			
	0.190	0.155			
	0.022	0.050			
	0.130	0.059			
	0.168	0.074			

Notes: Table records the values of the micro-moments used in estimation as defined in Table 5, both in the data, and as implied by the model at the estimated parameters.

Figure A.2: Distribution of Price Coefficient



Notes: Figure records lognormal distributions fitted to the empirical distribution of α_i^P among all consumers, all consumers who purchase a jet, and all consumers whose first purchase is a new jet in model simulations at the estimated parameters.

Table A.11: Parameter Estimates: Alternative Specifications

Parameter	No Heterogeneity	Alternative Market Size	Excluding Air Transport
α^p	-1.566 (0.201)	-1.887 (0.112)	-1.425 (0.151)
α^{new}	2.809 (0.229)	2.551 (0.090)	2.000 (0.065)
$\alpha_{upgrade}^{new}$	2.676 (0.122)	0.968 (0.064)	1.094 (0.047)
α^{sb}	1.372 (0.039)	0.259 (0.087)	0.259 (0.075)
τ	8.880 (0.104)	8.486 (0.176)	8.354 (0.085)
τ^{exit}	8.559 (0.127)	2.486 (0.072)	4.030 (0.106)
σ_p		1.392 (0.152)	1.558 (0.121)
σ_τ		0.802 (0.109)	0.368 (0.076)
$\sigma_{p\tau}$		-0.311 (0.101)	-0.963 (0.040)
σ_0		1.965 (0.129)	1.739 (0.093)
σ_m		0.609 (0.099)	0.345 (0.043)
Buyback Parameter b_m	No Heterogeneity	Alternative Market Size	Excluding Air Transport
Bombardier	0.730 (0.475)	0.830 (0.143)	0.334 (0.061)
Cessna	0.597 (0.304)	0.644 (0.124)	0.574 (0.113)
Dassault	0.785 (0.604)	0.419 (0.168)	0.652 (0.135)
Gulfstream	0.423 (0.171)	0.290 (0.210)	0.585 (0.128)
IAI	0.390 (0.198)	0.726 (0.085)	0.660 (0.095)
Raytheon	0.582 (0.277)	0.521 (0.121)	0.502 (0.180)
Other	0.705 (0.171)	0.389 (0.080)	0.415 (0.140)

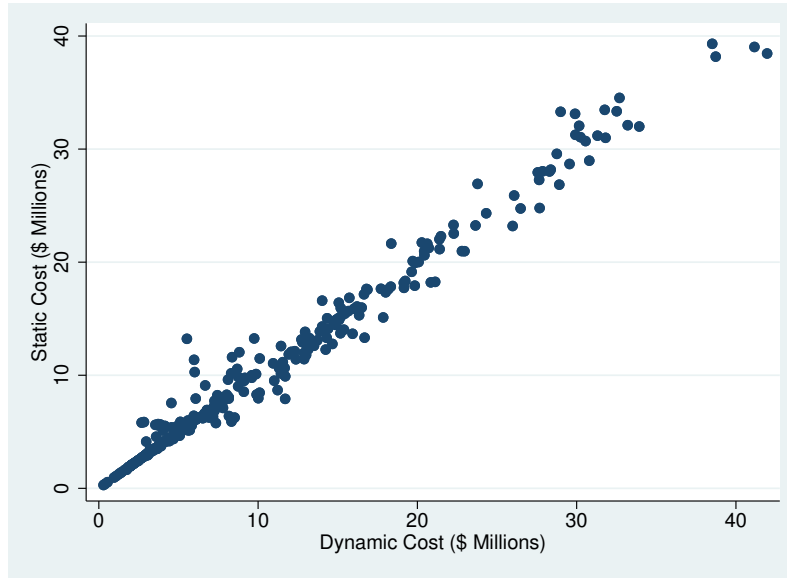
Notes: The first column reports estimated parameters and standard errors for the demand model under the restriction of no preference heterogeneity. The second column reports estimated parameters and standard errors estimated under an alternative market size assumption. In particular, I use the count of all publicly listed US firms as the number of potential first time buyers, M_{ot} , each year. The third column reports estimated parameters and standard errors estimated using a reduced sample that drops air transportation companies from the set of consumers. Prices are in hundreds of thousands of 2009 dollars.

Table A.12: Gross Profit Margins from Financial Reports

Manufacturer	Holding Company	Gross Profit Margin $\frac{p-c}{c}$
Bombardier	Bombardier	9.91%
Cessna	Textron	21.84%
Dassault	Dassault	13.20%
Gulfstream	General Dynamics	16.66%
Raytheon	Raytheon	19.42%

Notes: Figures from annual reports summarized by the Wall Street Journal. Note that profit margins are across all businesses and products, not only business jets.

Figure A.3: Marginal Cost Comparison




Notes: Figure plots marginal costs estimated using a model of static pricing against marginal costs using a model of dynamic pricing in which firms are forward looking and have perfect foresight.

Table A.13: Firm Profit

Manufacturer	BB units	No BB Profit	BB Profit	Unilateral Deviation Profit
Bombardier	295.31	1.78	1.91	1.72
Cessna	277.59	1.57	1.59	1.52
Dassault	147.03	1.40	1.47	1.36
Gulfstream	146.36	1.96	2.12	1.91
IAI	16.29	0.21	0.21	0.20
Raytheon	100.12	0.90	0.91	0.87
Other	22.48	0.21	0.21	0.21

Notes: Table records manufacturer profits in billions of 2009 \$ for various equilibrium simulations. No BB profit is computed using simulations which set b_m to 0 for all manufacturers and otherwise use the estimated parameters. BB profit is computed using simulations at the estimated parameters. Unilateral deviation profit is computed using simulations which set b_m to zero for the firm in question and set b_m to the estimated values for all other firms. BB units records the number of same-brand used-new upgrades for each firm in the buyback simulation.

Figure A.4: Distribution of Jet Holders in 2000



Notes: Figure records lognormal distributions fitted to the empirical distribution of α_i^p among holders of each jet brand in equilibrium simulations at the year 2000, using the estimated parameters.