Trade-ins and Transaction Costs in the Market for Used Business Jets*

Charles Hodgson†

October 15, 2019

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 manufacturer revenue by 7.2% at fixed prices. However, in equilibrium this revenue gain is diminished by 43% due to substitution away from new jets among first time buyers. I show how the size of this cannibalization effect depends on preference heterogeneity.

Keywords: durable goods, transaction costs, buyback, trade-ins, secondary markets

JEL Classification: L13, L14, L20, L93

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 frictionless market, consumers who prefer new goods would upgrade as frequently as possible. 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. In reality, used durables are held for extended periods of time. This behavior is typically rationalized by the presence of transaction costs. A consumer

*I would like to thank Liran Einav, Timothy Bresnahan, Matthew Gentzkow, and participants at the Stanford IO workshop for helpful comments and discussion. I am especially grateful to Liran Einav and Thomas Gilligan for sharing their data.

†Department of Economics, Yale University, 37 Hillhouse Avenue, New Haven, CT 06511 Email: charles.hodgson@yale.edu

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 market[1].

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. All major manufacturers participate in the secondary market by accepting used units as trade-ins and reselling them. Using data on all transactions in the new and used business jet market between 1961 and 1999, I estimate a model of jet demand which I use to measure the size of transaction costs and the reduction in transaction costs that can be attributed to manufacturer buyback. I then simulate demand without buyback schemes to measure the effect of buyback on the demand for new jets and the supply of used jets.

I measure the average transaction cost paid by upgrading consumers to be $734 thousand, or approximately 11% of the average jet price. I find that manufacturer buyback schemes eliminate between 4.4% and 10.2% 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 389, or 21.6%. Aggregated across manufacturers, the increase in demand for upgrades due to buyback increases revenue by 7% if prices are held fixed. To evaluate the extent to which this gain in revenue is cannibalized by substitution away from new units among first time buyers, I

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

simulate counterfactual equilibria in which I allow used jet prices to adjust. The increase in used jet supply due to buyback reduces the average price of used units by 1.4% relative to the no-buyback equilibrium prices. The resulting reduction in quantity demanded for new jets among first time buyers is 31% of the increase in quantity demanded among upgraders. In equilibrium, this cannibalization effect diminishes the gain in revenue due to buyback by 43%. 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 to 7.6%.

In an oligopolistic market, buyback also has competitive effects. By accepting trade-ins when its competitors do not, a manufacturer can encourage upgrading consumers to substitute away from its competitors' products. 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. 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 four of the five major manufacturers, operating buyback is a dominant strategy. For the fifth manufacturer, Cessna, operating buyback is a best response to other firm's policies, but the firm's profit would be higher if no firms offered buyback. 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 ask how reducing the level of transactions costs in the market changes the incentive for firms to operate buyback. I run simulations where transactions costs are reduced by 1% and 5%, and show that there is a non-monotonic effect on the incentive to operate buyback.

This paper contributes to an existing, largely theoretical, literature on the interaction of manufacturers with secondary markets. Fudenberg and Tirole (1998) model a durable goods monopolist who sells successive generations of a product to a stock of consumers who may be able to trade in a secondary market. They 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.

Closer in spirit to this paper, 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[2]. 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.

In terms of empirical methodology, this paper is closely related to Schiraldi (2011). Schiraldi examines the effects of scrappage policies in the Italian used car market using a model in which cars depreciate and there are transaction costs which prevent owners from upgrading immediately. I adapt Schiraldi's structural approach to identifying transaction costs. My methodology is different in that I do not use a dynamic model and I combine aggregate market shares with transaction-level "mirco-moments" along the lines of Petrin (2002) to help identify key parameters that control the heterogeneity of preferences among holders of different jets. The model allows for heterogeneous preferences, and keeps track of the evolution of the distribution of preferences among holders of different jets over time, similar to other durable goods demand papers including Gowrisankaran and Rysman (2012), Carranza (2010), and Esteban and Shum (2008). Also relevant to this paper's methodology is the literature on on models of equilibrium in used durable goods markets. Stolyarov (2002), Chen et al. (2013), and Gavazza et al. (2014) calibrate equilibrium models to match facts about the rate of resale in used car markets and simulate equilibira under different parametric assumptions. In this paper, although I do not estimate a full equilibrium model, I perform counterfactual equilibrium simulations using the parameters recovered from the demand estimation in order to estimate the effects of buyback on new and used unit prices in equilibrium.

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]Hendel 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.

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 five leading business jet manufacturers - Cessna, Bombardier, Raytheon, Dassault, and Gulfstream - sold 10,295 new jets. Over the same period, there were 40,295 sales of these manufacturers' jets on the used market. A typical owner holds a jet for between three and four years. 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.

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 typically 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. While it is possible for a jet sale to be arranged without a broker, this is rare. 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 by selling it directly to a manufacturer, as long as they replace it with a new jet from that manufacturer. Manufacturer buyback was commonplace in the industry until the 2008 recession. 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," and described several of the recent trade-ins received by Learjet. A similar 1983 Wall Street Journal advertisement for Cessna boasted of "liberal trade-in allowances." More recently, manufacturers have moved towards offering free brokerage for owners who upgrade to new jets instead of accepting trade-ins. 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.

Used jets that are bought back by manufacturers are almost always resold. Jets are long lived, and the price a jet will earn on the used market typically exceeds scrap value. Manufacturers hold unsold used jets for extended periods and try to resell the units they receive as trade-ins even when demand is low. For 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."[3]

2.2 Data

The analysis uses a data set which records all transactions of new and used business jets in the United States from 1957 to 2000. The data was constructed from FAA registration records for all business jet aircraft first registered before 2001. 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 to track jets across owners over time from manufacture to retirement, and to track owners as they buy, hold, and sell jets.

Owners are classified into one of nine types: dealers, manufacturers, corporations, financial institutions, airlines, air taxis, cargo, government, and private owners. Corporations are by far the largest owner category (exuding dealers and manufacturers), comprising 62% of owner-jet pairs. The top panel of Table 1 records the average observed holding lengths (in months) for dealers, corporations and other owners. Corporate owners hold jets for an average of 46.5 months (about 3.9 years) before reselling them. Dealers hold jets for about one year on average, significantly shorter than any other owner type. The bottom panel records the average fleet size for these owner categories. Corporations hold between one and two jets at a given time and dealers hold more than two on average. This finding is consistent with industry practice. Most dealers or brokers typically arrange direct transactions between owners, and occasionally buy jets for short periods in order to facilitate "back-to-back" transactions. However, some dealers acquire used jets even when a buyer is not lined up, similar to used car dealers, and other companies classified as dealers in this data may also lease 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, medium, and large. These categories are roughly defined by engine size, range, and capacity, as illustrated by Table 2. Table 3 records manufacturer market shares of new jet sales for the five major manufactures in each of the three market segments.[4] Note that the small jet market is dominated by Cessna, the large jet market is dominated by Gulfstream, and that together, the five firms listed make up over 83% of each of the three segments. Table 3 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

[3]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.

[4]Note 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.

Table 1: Holding Time and Fleet Size by Owner Type

		Broker	Corporation	Other
Holding Time	Mean	12.576	46.539	49.492
	SD	18.230	45.840	58.310
	N	11804	17461	10832
Fleet Size	Mean	2.144	1.282	1.967
	SD	3.668	0.981	6.586
	N	91130	849650	412031

Notes: Holding time observations are jet-owner pairs. The value of each observation is the number of months that pair is observed. The sample includes all jet-owner pairs except those whose last observation is in December 2000 (the last period of the panel) since the observed holding times for these pairs are truncated. Fleet size observations are owner-month pairs. The value for each observation is the number of jets owned by that owner in that month. The sample includes all owner-month pairs. The sample for the bottom two rows includes all owner-years.

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 in 2001 by Penton Information Services. The Blue Book contains quarterly prices for new and used 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 last three rows of Table 2 record the mean and standard deviation of the prices of jets of different sizes and production years when they are new, and the used prices 1 and 5 years after manufacture. I express all prices in year 2000 dollars. Prices for used medium jets drop by 5% on average in the first year, and then by additional 27% over the next four years.

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 data. In order to estimate the model of demand described in Section 3, I fill in this missing data using a procedure described the Appendix Section A.1.

Finally, I obtain jet characteristics from Frawley's (2003) *International Directory of Civil Aircraft*

Table 2: Aircraft Characteristics

Segment		All		Small	Medium	Large
Model Years		< 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 2000 dollars. Columns 1 and 2 record the mean and standard deviation of model characteristics for models manufacturers before and after 1990. Columns 3 to 6 record the mean and standard deviation of model characteristics by market segment (jet size).

Table 3: Market Share by Manufacturer

	New Market Share 1961 - 2000			Used Market	
Manufacturer	Small	Medium	Large	Resale Ratio	Annual Used Sales
Cessna	52%	11%	0%	26.6%	392.2
Bombardier	32%	8%	33%	23.3%	309.6
Raytheon	10%	26%	0%	25.5%	168.6
Dassault	4%	22%	14%	26.6%	165.8
Gulfstream	0%	0%	54%	19.7%	88.8

Notes: Columns 1-3 record the market share of the top 5 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.

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). These characteristics are described in Table 2 for aircraft of different sizes and production years.

2.3 Estimation Sample

For the estimation of the model presented in Section 4, I exclude from the sample owners classified as manufacturers, dealers, financial institutions, and government agencies. This reduces the number of owners in the sample from 20,258 to 17,816.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 1, 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 aricraft of the same brand rather than multiple aircraft of different brands. Appendix Table A.1 records the share of owner-years for which multiple jets are purchased for different sets of owners. Corporations, which are included in the estimation sample, make purchases in 25% of owner-year observations but purchase multiple jets in only 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. The equivalent numbers for dealers, who are excluded from the estimation sample, are 52% and 20%.[5]

I aggregate the available choices to the manufacturer-segment-model year level. For example, an owner making a choice in 1985 could choose to buy a large 1972 Gulfstream or a medium 1980 Cessna. I also collapse all manufacturers other than the top 5 into a composite "Other" category. Many of these model categories contain multiple jet models. For example there are several variant

[5] I exclude dealers, government agencies, and financial institutions from the estimation sample because these buyers hold multiple jets with greater frequency, are likely not well represented by the demand model developed in Section 4, and may face different prices. For example, Large financial institutions lease aircraft, and likely face significant efficiency gains in holding a fleet of a single brand due to maintenance costs. Jet purchases by the Air Force are made though contract tendering procedures.

models in the medium 1980 Cessna category. I map price and jet characteristic data to model categories by averaging over the 'true' models 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.

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

The raw data provides direct evidence that buyback schemes are an important feature of the used jet industry. Although precise details about when and whether each firm operated buyback schemes and the terms of these programs are not available, it is possible to use the data to study buyback schemes directly by identifying sales of jets from corporations back to manufacturers or to manufacturer-affiliated dealers. I manually identify from the list of dealers those that appear to be manufacturer affiliated, based on their name. For example, the largest dealer in the data is "Bombardier Aerospace Corporation."

Table 4 records statistics on the used jets bought by manufacturers and manufacturer affiliated dealers from corporations. The first column records the share of jets bought back by each manufacturer that are part of an upgrade to a new jet by the corporation the sells the used jet. For these statistics, I define an upgrade as the sale of a used jet by a corporation followed by the purchase of another jet within 12 months. The share of buybacks that are part of an upgrade to new is between 79% and 93% for each of the major manufacturers. These statistics confirm that manufacturers buy back used jets from owners who wish to upgrade to new units. he fact that sales to manufacturers rarely take place as part of used-used upgrades provides some assurance that manufacturers are not acting as general used jet dealers.

The data also indicates that manufacturers buy back used jets of their own brand almost exclusively. The second column of Table 4 records own brand jets bought back as a share of all jets

Table 4: Buyback Patterns

Manufacturer	Share of Buybacks Upgrades to New	Share of Buybacks Own Brand	Share of Potential Buybacks Sold to Manufacturer
Bombardier	86%	86%	46%
Cessna	93%	93%	42%
Dassault	91%	80%	26%
Gulfstream	79%	82%	44%
Raytheon	82%	87%	56%

Notes: The first column records that share of all jets sold back to a manufacturer that are part of a used-new upgrade. 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 are upgrades from used to new units of the same brand.

bought back for each of the five major manufacturers. The shares are over 80% for each manufacturer, and as high as 93% for Cessna. That is, 93% of jets sold to Cessna are used Cessna jets. 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. Indeed, if manufacturers had no own-brand preference in the jets they accept as trade ins, we would expect the share of own-brand buybacks recorded in Table 4 to be closer to 59%, which is the share of same brand upgrades among all used-new upgrades. As I discuss below, the fact that 59% of all used-new upgrades are "same brand" suggests persistent heterogeneity in brand preference. In the demand model estimated in Section 4 I allow for surch persistence in brand preference.

The third column of Table 4 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 4 indicates that in around 40% of these upgrades from used to new jets, the used jet is bought by a manufacturer. The figure ranges from 26% to 56% across manufacturers. Owners therefore take advantage of buyback in a significant share of observed used-new upgrades. Appendix Table A.2 repeats these exercises for different time periods, demonstrating that this pattern holds throughout the years 1961-2000.

3.2 Buyback as a Demand Shifter

The ubiquity of buyback across manufacturers and time makes it difficult to evaluate the impact of buyback schemes on demand and manufacturer revenue. The data suggests that there is variation in the terms of these schemes across manufacturers, but documentation of the policies is not available. There are therefore no obvious natural experiments that can be used to examine how changes in buyback schemes over time or differences across manufacturers drive differences in jet demand. In addition, observed sales back to manufacturers may understate the extent of manufacturer involvement in facilitating used-new upgrades - for instance, a manufacturer who acts as a broker will not necessarily be recorded as the owner of that jet in the data.

In order to estimate the effect of buyback on demand, I make the assumption 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 the market for upgrades to the market for first time buyers, and by comparing same brand upgrades to different brand upgrades. First time jet buyers do not have to pay the transaction costs associated with selling a jet, and therefore do not benefit from buyback schemes. Furthermore, the patterns illustrated in Table 4 suggest that it is plausible to assume manufacturers only accept trade-ins of their own brands of jets. Under this assumption we would expect, all else equal, the market share for new jets to be higher among upgrade buyers than among first time buyers, and higher among same brand upgraders than among different brand upgraders.

Table 5 shows that these patterns hold in the data. Panel A shows the market shares for new and used jets among first time and upgrade purchases. An upgrade purchase is defined as a purchase by an owner who sold a jet less than 12 months earlier. Used jets have a higher share among both sets of buyers, but the difference between the used and new share is 9 percentage points higher in the first time buyers market. As expected, conditional on making a purchase, replacement buyers are more likely to buy a new jet. Panel C 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 14.1 percentage points more likely to buy a new jet than upgraders who change brands.

These differences in the share of new jets bought are the key variation that I use to identify the effect of buyback. Of course, there could be systematic differences in preferences between, for example, first time and replacement buyers that could produce the same patterns. To demonstrate that buyback seems to be driving a large part of these differences, panels B and D of Table 5 repeat the exercises in panels A and C, but exclude all owners who *ever* sell a jet back to a manufacturer. Thus, the comparison in panel B is between first time and replacement buyers, where the buyers do not benefit from buyback in either market. The difference in market shares between the two markets is 2 percentage points, much smaller than in panel A, and not statistically different from zero. Similarly, the comparison in panel C is between same brand and different brand upgraders who never trade in a used jet. In this case the difference in new jet market shares is reduced from 14.1 to 4.1 percentage points, and is no longer statistically significant.

These statistics provide suggestive evidence that the availability of manufacturer buyback increases demand for new jets among upgraders. 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. The differences in the relative market shares for new and used jets between same brand and

Table 5: Upgrade Shares by New and Same Brand

	(A) All Owners			(B) Excluding Buyback Users		
	Used	New	Diff	Used	New	Diff
First Jet	0.702	0.298	0.404	0.679	0.321	0.358
Upgrade	0.657	0.343	0.314	0.689	0.311	0.378
Diff	0.045	-0.045	0.090	-0.010	0.010	-0.020
SE	(0.008)	(0.008)		(0.009)	(0.009)	

	(C) All Owners			(D) Excluding Buyback Users		
	Used	New	Diff	Used	New	Diff
Same Brand	0.545	0.455	0.090	0.677	0.323	0.354
Different Brand	0.686	0.314	0.372	0.718	0.282	0.436
Diff	-0.141	0.141	-0.282	-0.041	0.041	-0.082
SE	(0.020)	(0.020)		(0.024)	(0.024)	

Notes: Panel A records the share of new and used jets purchased among all first time and replacement purchases made by all owners excluding manufacturers and dealers. Owners who only ever own one jet are also excluded from the calculation. A replacement purchase is defined as a purchase that occurs less than 12 months after a sale. Panel B repeats the calculations in Panel A excluding all jet owners who ever sell a jet to a manufacturer. Panel C records the share of new jets and used jets purchased among same brand upgrades and different brand upgrades made by all owners excluding manufacturers and dealers. Panel D repeats the calculations in Panel C excluding all jet owners who ever sell a jet to a manufacturer. Standard errors reported are the standard error of the difference in means for each column.The sample used includes all purchases between 1961 and 2000.

different brand upgrades reported in Table 5 drive the identification of this model. In particular, these differences in demand will be attributed to differences in transaction costs between buyback eligible and non-eligible purchases. The model enables me to run counterfactual simulations that measure the effect of buyback schemes on the demand for new jets and the supply of used jets.

4 Demand Model

4.1 Model Description

In this section, I present a model of 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.[6] 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"), and described by its observable characteristics, x_{jt} and its price, p_{jt}. Let X_{kjt} be a matrix that depends on x_{jt}, x_{kt}, p_{jt} and p_{kt}.

[6] Note 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

$$u^i_{jjt} \quad = \quad X_{jjt}\alpha_i + \xi_{jt} + \epsilon_{ijt}, \tag{1}$$

where ξ_{jt} is unobserved product quality that is allowed to vary at the product-year level, and ϵ_{ijt} is an i.i.d. type 1 extreme value shock to preferences. α_i is a vector of individual-specific coefficients. Consumer i's flow utility from upgrading from jet type k to jet type j is

$$u^i_{kjt} \quad = \quad X_{kjt}\alpha_i - \tau_{kj} + \xi_{jt} + \epsilon_{ijt}, \tag{2}$$

where τ_{kj} is an unobserved transaction cost that is incurred when upgrading from a jet of type k to a jet of type j. I indicate the outside option of "no jet" in equation 2 with $k = 0$ or $j = 0$. For instance, u^i_{0jt} is the utility from purchasing jet j for a consumer who does not hold a jet, and u^i_{k0} is the utility from selling jet of type k and exiting the market. I assume that $\tau_{jj} = 0$ and $\tau_{0j} = 0$. That is, buyers who do not sell an existing jet in order to upgrade or exit the market are assumed not to pay any transaction costs and therefore not to benefit from any buyback schemes. The model yields the familiar logit choice probabilities,

$$P^i_{kjt} = \frac{exp(X_{kjt}\alpha_i - \tau_{kj} + \xi_{jt})}{\sum_{l \in J_t \cup 0} exp(X_{klt}\alpha_i - \tau_{kl} + \xi_{lt})}. \tag{3}$$

The part of utility that depends on jet characteristics, $X_{kjt}\alpha_i$, is defined as follows,

$$\begin{aligned}
X_{jjt}\alpha_i =& \alpha_i^0 + x_{jt}\alpha \\
X_{0jt}\alpha_i =& \alpha_i^0 + x_{jt}\alpha - p_{jt}\alpha_i^p \\
X_{kjt}\alpha_i =& \alpha_i^0 + x_{jt}\alpha + (p_{kt} - p_{jt})\,\alpha_i^p + new_{jt}\alpha^{new}_{upgrade} + i = 1(m(j) = m(k))\alpha^{sb} - \tau_{kj} \\
X_{k0t}\alpha_i =& p_{kt}\alpha_i^p - \tau_{k0}.
\end{aligned} \tag{4}$$

The flow utility from holding a jet, $X_{jjt}\alpha_i$, depends only on the characteristics of jet j. The utility obtained from purchasing jet j for a consumer that does not hold a jet, $X_{0jt}\alpha_i$, depends on characteristics x_{jt} and the price p_{jt}.

The flow utility from upgrading from k to j, $X_{kjt}\alpha_i$, also depends on $new_{jt}\alpha^{new}_{upgrade}$, 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 persistent brand preferences or consumer inertia in brand choice.

Finally, the utility from selling jet k and exiting the market, $X_{k0t}\alpha_i$, depends on the price of jet k and a transaction cost, τ_{k0}. The outside option of no jet yields a normalized utility of 0. Notice that I allow for two dimensions of preference heterogeneity - an individual specific coefficient on price, α_i^p, and an individual specific coefficient on the inside good, α_i^0.

As I discuss further in Section 4.3, I make the strong assumption that consumers are not forward looking and make decisions each period based on the flow utilities without considering how this decision affects their expected utility in future periods. However, I assume they are able to borrow at an interest rate r_t and costlessly refinance held jets each period. The relevant price is therefore the *one period rental rate.* I set $p_{jt} = \tilde{p}_{jt} - \delta_t \tilde{p}_{jt+1}$ where $\delta_t = \frac{1}{1+r_t}$ and $\tilde{p}_{jt}$ is the 'sticker price' of jet j at time t. For example, if jet j is new at date t, then $\tilde{p}_{jt}$ is the price of the new jet and $\tilde{p}_{jt+1}$ is the used price after one year. When deciding whether to purchase a jet, consumers compare the instantaneous utility they obtain from owning the jet this period to the cost of owning the jet for one period, which is given by $(1 + r_t)\, \tilde{p}_{jt} - \tilde{p}_{jt+1}$, discounted at rate δ_t. Note that this formula also assumes that consumers know the period $t + 1$ price of a jet at period t. I refer to this rental rate as the jet's "price" from here on.

4.2 Econometric Specification

In the main specification, the vector x_{jt} contains an indicator for each manufacturer, the jet's maximum range (in km), total engine power (in kN), maximum takeoff weight (in kg), the jet's age (in years), and an indicator for whether the jet is new, n_{jt}. To capture the effects of the business cycle on jet demand, I also include annual GDP growth in x_{jt}. Prices, p_{jt}, are scaled to hundreds of thousands of 2009 Dollars.

I specify transaction costs as:

$$\tau_{kj} = \left(new_{jt}1(m(j) = m(k))b_{m(j)} - \tau\right) 1(j \neq 0) - \tau^{exit}1(j = 0). \tag{5}$$

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}.

I assume distributions on that the individual specific price coefficient, α_i^p, and intercept, α_i^0. I assume both parameters are is drawn iid across owners when they enter the market, before they make their first purchase decision, according to

$$
\begin{aligned}
\alpha_i^p &\sim logN(\alpha^p, \sigma^p) \\
\alpha_i^0 &= \alpha^0 + \nu_i \delta \\
\nu_i &\sim Bernoulli(\pi).
\end{aligned} \tag{6}
$$

That is, the coefficient on price is distributed across consumers according to a log normal distribution, and the inside good coefficient is equal to α^0 with probability $(1-\pi)$ and $\alpha^0 + \delta$ with probability π. Note that although preferences are drawn i.i.d. across consumers when they enter the market, 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 heterogreneous 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.

The parameters to be estimated are the mean utility parameters $(\alpha, \alpha_{upgrade}^{new}, \alpha^{sb})$, the distributional parameters of the random coefficients $(\alpha^p, \sigma^p, \alpha^0, \delta, \pi)$, the transaction costs τ and τ^{exit}, and the buyback parameters b_m for each manufacturer m. Denote the vector of parameters to be estimated by θ.

4.3 Discussion

The purpose of this model is to measure the size of transaction costs and the amount by which these costs are reduced by buyback. The key parameters to be estimated are therefore τ and b_m. The transaction cost, τ, 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. Buyback schemes are designed to reduce these transaction costs for consumers who upgrade to new jets. By offering to buy a potential customer's old jet, the manufacturer eliminates the cost of hiring a broker, advertising, and the time cost of waiting to complete the sale of the old jet. This reduction in transaction costs is measured by b_m, and is allowed to be different for each manufacturer. 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.

Note that this model cannot rationalize imperfect takeup of buyback schemes. As recorded in

Table 4, only about 40% of all upgrades to new jets by corporate owners involve buyback. 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.

The imperfect takeup of buyback observed in the data suggests that the benefit of buyback is heterogeneous across the population of consumers, and not universally positive. While some consumers may benefit from the reduction in transaction costs, others may be able to obtain a better price from using a broker or selling directly to another consumer. The benefit likely depends on the resources an owner has to market their used jet independently, and the ability of the owner to wait to obtain a high sale price. Estimating a model with this level of heterogeneity in consumer preferences and in prices would require more reliable data on observed buybacks and transaction level price data.

Finally, note that this model does not capture dynamic incentives in consumer behavior. For example, consumers might take into account differences in buyback schemes across manufacturers when making their first time purchases if they are forward looking and anticipate the future cost of upgrading. Both first time and replacement buyers obtain this future benefit of buyback, but only replacement buyers receive the instantaneous reduction in transaction costs. The forward-looking benefit of buyback will therefore be absorbed into the manufacturer fixed effects, and the parameters b_m can still be interpreted as measuring the reduction in transaction costs induced by buyback.

5 Estimation and Identification

5.1 Estimation Procedure

Estimation is based on the nested fixed-point Generalized Method of Moments (GMM) procedure of Berry, Levinsohn, and Pakes (1993). In particular, for each candidate parameter vector, θ, I find values of the product-year level unobservables, ξ_{jt}, that rationalize observed *aggregate* market shares and form moments by interacting ξ_{jt} with instruments. Note that in addition to the standard preference heterogeneity across consumers induced by the "random coefficient" α_i^p on price and α_i^0 on the inside good, 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.[7] Define the aggregate market share for jet j in year t as

$$s_{jt} = \frac{\sum_{k \in J_{t-1} \cup 0} M_{kt} s_{kjt}}{\sum_{k \in J_{t-1} \cup 0} M_{kt}}. \tag{7}$$

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

$$\hat{s}_{jt}(\theta, \xi) = \frac{\sum_{k \in J_{t-1} \cup 0} M_{kt} \hat{s}_{kjt}(\theta, \xi)}{\sum_{k \in J_{t-1} \cup 0} M_{kt}}, \tag{8}$$

where $\hat{s}_{kjt}(\theta, \xi)$ 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, ξ. This is given by

$$\hat{s}_{kjt}(\theta, \xi) = \int P^i_{kjt}(\theta, \xi) dF_{kt}(\alpha^p_i, \alpha^0_i | \theta, \xi), \tag{9}$$

where $P^i_{kt}(\theta, \xi)$ are given by equation 3, and $F_{kt}(\alpha^p_i, \alpha^0_i | \theta, \xi)$ is the joint distribution of random coefficients for holders of jet k in period t.

Recall from equation 6 that consumer preferences are dawn i.i.d. when consumers *first enter the market*. The distribution $F_{kt}(\alpha^p_i, \alpha^0_i | \theta, \xi)$ 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 $F_{kt}(\alpha^p_i, \alpha^0_i | \theta, \xi)$ for holders of jet k in year t depends on the probability that consumers with different values of α^p_i and α^0_i choose to purchase and hold jet k in all previous years.

Computation of $\hat{s}_{kjt}(\theta, \xi)$ therefore proceeds sequentially, 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 $k \neq 0$.[8] The model-implied aggregate market share for jet j is therefore equal to the share of first time buyers who purchase jet j, $\hat{s}_{j1961}(\theta, \xi) = \hat{s}_{0j1961}(\theta, \xi)$, which is computed according to equation 9, where the distribution $F_{01961}(\alpha^p_i, \alpha^0_i | \theta, \xi)$ is given by equation6.

For subsequent years, $\hat{s}_{kjt}(\theta, \xi)$ depends on the distribution of preferences to pay among jet holders.

[7] As 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.

[8] Note that the 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 α^p_i and α^0_i are the same among consumers arriving to the market for the entire sample period.

For jet k in year t, this distribution can be computed with Bayes' rule as a function of the year $t-1$ distributions and choice probabilities as follows,

$$f_{jt}(\alpha_i^p, \alpha_i^0 | \theta, \xi) = \frac{1}{\sum_{k \in J_{t-1} \cup 0} M_{kt}} \sum_{k \in J_{t-1} \cup 0} M_{kt-1} \frac{P_{kjt}^i(\theta, \xi) f_{kt-1}(\alpha_i^p | \theta, \xi)}{\hat{s}_{kjt}(\theta, \xi)}. \tag{10}$$

To ease the computational burden in practice, I replace the continuous distribution of price coefficients with a discrete approximation with support points at quantiles of the log normal distribution specified in equation 6.[9] The updating described by equation 10 is then performed a finite number of times corresponding to the number of points in the support of the joint distribution of (α_i^p, α_i^0) for each (k, t).

For a given candidate parameter vector θ I find a vector of product level unobservables ξ such that $s_{jt} = \hat{s}_{jt}(\theta, \xi)$ by iterating according to Berry, Levinsohn, and Pakes (1993) contraction mapping procedure. I then use instruments Z_{jt} to construct empirical analogues of the moments conditions

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

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 jets of type j at the beginning of year t and the number of consumers who hold close substitutes at the beginning of year t. I use two counts for close substitutes - the number of jets of the same size (small, medium, large) of the same age as jet j, the number of jets of the same size that are one year older than jet j, and the number of models available. It is clear that the more jets of type j held by owners, the higher the quantity of jet j supplied 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 ξ_{jt} since the number of jets held at date t is determined at date $t-1$. In Appendix Table A.3, I present a diagnostic "first stage" regression which shows that the jet holdings instruments predict prices conditional on other jet characteristics. 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

[9] In particular, for a given (α^p, σ^p) I use a distribution with 11 points of support given by $\omega_\ell = exp(\alpha^p + \ell\sigma^p)$ for $\ell \in \{-2.5, -2, -1.5...2, 2.5\}$, where $P(\alpha_i^p = \omega_\ell) = \frac{\phi(\ell)}{\sum_k \phi(k)}$.

Table 6: Micro-Moments

		Moment	Related Parameters
1	Upgrade conditional on holding	$P\left(upgrade_{it} \mid j_{it} \neq 0\right)$	τ
2	Exit conditional on holding	$P\left(j_{it+1} = 0 \mid j_{it} \neq 0\right)$	τ^{exit}
3	New jet conditional on upgrade	$P\left(age(j_{it+1}) \leq 1 \mid upgrade_{it} = 1\right)$	$\alpha^{new}_{upgrade}$
4	New jet conditional on first purchase	$P\left(age(j_{it+1}) \leq 1 \mid first_{it} = 1\right)$	α^{new}
5	Same brand conditional on upgrade	$P\left(m(j_{it+1}) = m(j_t) \mid upgrade_{it} = 1\right)$	α^{sb}
6-11	Difference in new jet share between same-brand upgrades and brand switchers	$P\left(age(j_{it+1}) \leq 1 \mid m = m(j_{it+1}) = m(j_t)\,\&\, upgrade_{it} = 1\right)$ $-P\left(age(j_{it+1}) \leq 1 \mid m = m(j_{it+1}) \neq m(j_t)\,\&\, upgrade_{it} = 1\right)$	b_m
12-14	Expected purchase price conditional on upgrading from jets of different ages	$E(price(j_{it+1}) \mid age(j_{it} < 5)\,\&\, upgrade_{it} = 1)$ $E(price(j_{it+1}) \mid age(5 \leq j_{it} < 15)\,\&\, upgrade_{it} = 1)$ $E(price(j_{it+1}) \mid age(15 \leq j_{it})\,\&\, upgrade_{it} = 1)$	σ^p
15-17	Exit conditional on holding jets of different ages	$P\left(j_{it+1} = 0 \mid age(j_{it} < 5)\right)$ $P\left(j_{it+1} = 0 \mid age(5 \leq j_{it} < 15)\right)$ $P\left(j_{it+1} = 0 \mid age(15 \leq j_{it})\right)$	δ, π

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

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

Table X 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 X. I compute the empirical probability of upgrading to a new jet conditional on upgrading at all as

$$\hat{h}_3 = \frac{\frac{1}{N}\sum_{i=1}^{N}\frac{1}{T}\sum_{t=1}^{T} I(age(j_{it+1}) \leq 1)upgrade_{it}}{\frac{1}{N}\sum_{i=1}^{N}\frac{1}{T}\sum_{t=1}^{T} upgrade_{it}}. \tag{11}$$

As in Petrin (2002), the conditional probability defined by equation 11 is written as a ratio of averages over N observations corresponding to the N owners in the sample. In this case, the ratio 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

$$h_3(\theta) = \frac{\sum_{t=1}^{T}\sum_{k\in J_{t-1}} M_{kt}\sum_{j\in J_t^N} \hat{s}_{kjt}(\theta,\xi(\theta))}{\sum_{t=1}^{T}\sum_{k\in J_{t-1}} M_{kt}\sum_{j\in J_t} \hat{s}_{kjt}(\theta,\xi(\theta))}, \tag{12}$$

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 ith 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

$$\tilde{\theta} = \arg\min_{\theta} \begin{bmatrix} G(\theta) \\ H(\theta) \end{bmatrix} [G(\theta)', H(\theta)']. \tag{13}$$

Using consistent estimates $\tilde{\theta}$, I then construct a weighting matrix $\Omega(\tilde{\theta})$ which is an estimate of the asymptotic covariance matrix of $\left[G(\tilde{\theta}), H(\tilde{\theta})\right]'$. 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. 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:

$$\theta^* = \arg\min_{\theta} \begin{bmatrix} G(\theta) \\ H(\theta) \end{bmatrix} \Omega(\tilde{\theta}) [G(\theta)', H(\theta)']. \tag{14}$$

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. As recorded in panel A of Table 5, 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 $\alpha^n_{upgrade}$ 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 both α^{sb} 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 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 panels C and D of Table 5 that identify b_m enter directly into the GMM objective function through mirco-moments 6-11 as detailed in Table 6. α^{sb} and $\alpha^n_{upgrade}$ 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. Note that α^{sb} could measure either inertia in brand preference or unspecified persistent heterogeneity in brand preferences (see Keane, 1997 and Dubé et al., 2010).

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 each year 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 variance parameter on the price coefficient σ^p and the parameters δ and π which determine the distribution of the coefficient on the inside good, α^p_i, are identified by micro moments 12-14 and 15-18. Moments 12-14 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^p_i are more likely to hold older jets and more likely to upgrade to cheaper jets. Moments 15-18 measure the probability of exiting the market among holders of different ages. Following a similar logic, owners with low values of α^0_i should be more likely to hold older, cheaper jets, and more likely to exit the market conditional on holding. Additional variation that helps identify these distributional parameters comes from the instrument that records the number of products available. As more products become available, 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 7. Table A.4 in the Appendix records the fit of the micro-moments from Table 6 at the estimated parameters.

The median coefficient on price, which is the marginal utility of \$100,000, is $\alpha_i^p = 3.813$. The individual-specific coefficient on price which enters the utility function is distributed according to $\alpha_i^p \sim logN(\alpha^p, \sigma^p)$. This is the distribution of price coefficients from which consumers' preferences are drawn when they enter the market. Figure 1 illustrates how the estimated distribution of price coefficients changes *conditional on jet holdings.* Panel A shows that consumers who hold jets are on average less price sensitive than consumers who do not. The mean price parameter among jet holders is $E(\alpha_i^p|hold) = 1.168$. Panel B shows that consumers who purchase new jets are on average less price sensitive than consumers who purchase used jets. Similarly α_i^0, the intercept term in the utility function, is drawn from a binomial distribution when consumers enter the market. The expected value of α_i^0 is -1.588 among first time buyers and -4.640 among entering consumers who do not purchase a jet.

The baseline model parameter estimates imply that the loss of utility from a jet aging one year, at the mean price parameter among jet holders, $\frac{\alpha^{age}}{E(\alpha_i^p|hold)}$, is equivalent to an increase in price of \$17.6 thousand. The difference in utility between a new and used jet is equivalent to a change in price of \$123.2 thousand. Note that the drop in quality once a jet becomes used is around 7 times 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 \$7000 and an increase in range of 10 km equivalent to a reduction in price of \$2000, at the mean price parameter among jet holders. Surprisingly, utility is decreasing in maximum weight, though the effect is economically negligible. This could reflect other factors associated with increased weight beyond capacity such as fuel consumption. Finally, note that the utility from holding a jet is higher during periods of higher GDP growth. Empirically, companies are less likely to purchase business jets during recessions.

The estimated transaction cost for upgrades at the expected price parameter among jet holders, $\frac{\tau}{E(\alpha_i^p|hold)}$, is approximately \$734 thousand or 11% of the average jet 'sticker price' of \$6.7 million.

Table 7: Parameter Estimates

Parameter	Estimate	SE	Parameter	Estimate	SE	Buyback Parameter b_m		
α^p	3.813	0.603	τ	8.580	0.032	Bombardier	0.378	0.141
σ^p	1.300	0.106	τ^{exit}	5.188	0.168	Cessna	0.423	0.135
α^{age}	-0.205	0.004	α^{sb}	1.467	0.048	Dassault	0.463	0.183
α^n	1.497	0.103	$\alpha^n_{upgrade}$	1.439	0.117	Gulfstream	0.666	0.208
α^{range}	0.023	0.001	α^0	-5.683	0.359	Raytheon	0.872	0.166
α^{power}	0.072	0.014	δ	6.361	0.199	Other	0.453	0.192
α^{weight}	-0.000	0.000	π	0.247	0.014			
$\alpha^{GDPgrowth}$	11.650	1.307						

Notes: Table reports estimated parameters and standard errors for the demand model. Prices are in hundreds of thousands of 2009 dollars. Age is measured in years, range is measured in km, weight is measured in kg, and power is measured in kN. GDP growth is % change from the previous year.

The estimated transaction costs faced by consumers exiting the market, $\frac{\tau^{exit}}{E(\alpha_i^p|hold)}$, are lower at \$444 thousand. By way of comparison , Schiraldi (2011) finds that transaction costs (defined in a similar manner) in the market for new and used cars are usually between 10% and 80% of the sale price. Note that owners upgrade only when they obtain a high draw of ϵ_{ijt} for some jet j which they do not currently own. If we interpret ϵ_{ijt} to include idiosyncratic transaction costs, then τ is an upper bound on the size of the transaction costs actually paid when owners upgrade.

Note that the estimated size of τ does not depend on the definition of a time period in the data. If the data was aggregated to five year periods, then the frequency of upgrades per period would be five times higher than in the annual data used for estimation. The estimate of τ would therefore be smaller under a coarser aggregation. However, the relevant prices would be five year rental rates, and therefore the estimated average price parameter would be approximately five times smaller. The estimated value of τ converted into dollar units would therefore be approximately the same under different levels of temporal aggregation.

As expected based on the descriptive recorded in Table 5, the manufacturer-specific buyback parameters b_m are positive and statistically significant. 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 4.4% for Bombardier and 10.2% for Raytheon. Differences in b_m across manufacturers could reflect differences in the costs

Figure 1: Distribution of Price Coefficient

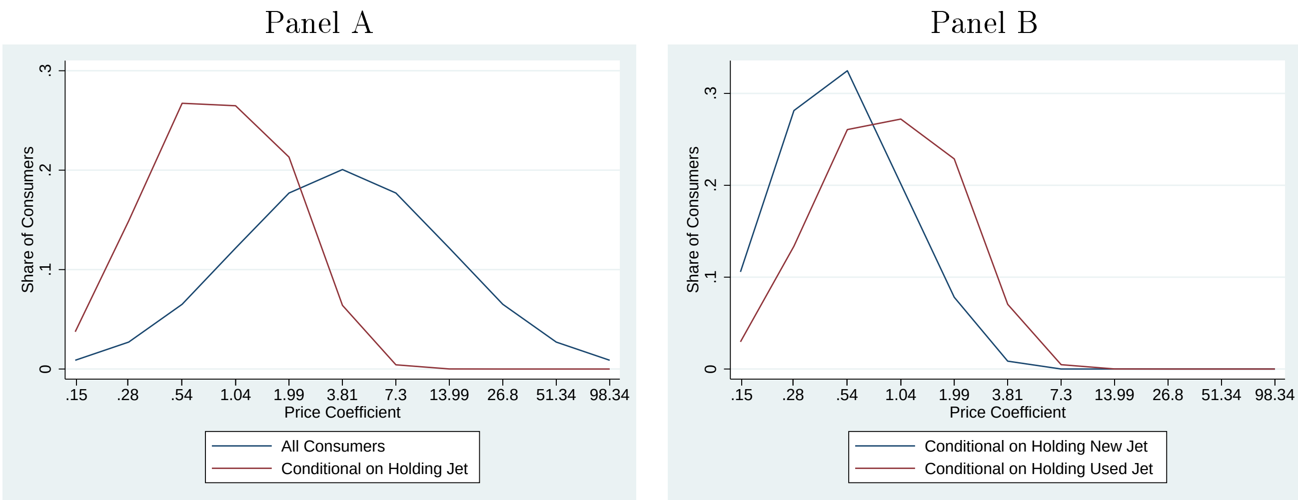

Notes: Panel A records the distribution of α_i^p among all consumers and all consumers who hold a jet in model simulations at the estimated parameters. An observation is a consumer-year, so consumers who hold jets for longer are more heavily weighted in the distribution. Panel B records equivalent distributions for consumers who hold new and used jets. Note that the distribution of α_i^p is discrete with 11 points of support, as discussed in footnote 9.

associated with upgrading different types of jets - one might expect the reduction in transaction costs for larger jets to be higher since the market for these jets is thin and it may take longer to find a buyer. This could explain why, for example, Gulfstream has a higher value of b_m than Cessna. b_m might also reflect differences in the generosity of buyback policies across manufacturers. For example, the prices paid to upgraders by manufacturers may differ from the "market price" observed in the data, leading to lower estimated values of b_m.

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.

I compute the baseline quantity demanded, $M_{kt}s_{kjt}$, of jet $j \in J_t$ from consumers who hold jet $k \in J_{t-1} \cup 0$ for each jet j at each year t at the estimated parameters and product unobservables, holding fixed prices and market sizes at the observed levels. These computed values are equal to the true quantities demanded up to numerical imprecision in the fixed point algorithm. I then compute quantity demanded with no buyback by setting $b_m = 0 \, \forall m$ and holding all other

Table 8: Demand Simulations

	Baseline Parameters			No Heterogeneity		
	No Buyback	Buyback	Δ	No Buyback	Buyback	Δ
(1) Upgrades to New	1413	1802	389	1462	1877	415
(2) Upgrades to Used	3754	3710	-44	3506	3489	-17
(3) Exits	14734	14710	-24	14618	14551	-67
Used Jet Supply to First Time Buyers = (1) + (3)	16147	16512	365	16080	16428	348
Manufacturer Revenue	6.593	7.066	0.473	6.475	7.066	0.591

Notes: Revenue is in billions of 2009 $ computed using "rental rate" prices. 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.5

parameters, prices, and market sizes fixed. Note that this is not an equilibrium counterfactual, as I do not allow prices to adjust, nor do I allow jet holdings to evolve endogenously over time. Rather, this exercise measures the shift in demand for new and used jets among replacement buyers at fixed prices.

The first three columns of table 8 illustrates how buyback shifts the demand for new and used jets from upgrading consumers at fixed prices. Buyback increases the demand for new units from upgrading consumers over the sample period by 389 units or about 28% 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 8 indicate that of the 389 additional new jets demanded, 11% comes from consumers substituting from used jets, 6% comes from substitution from market exits, and the remaining 83% 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,

$$SupplyUsed = UpgradesToNew + Exits. \tag{15}$$

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 15.

The implication of this decomposition is illustrated by the third line of Table 8. 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 94% 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 this, the second set of columns in Table 8 record analogous demand simulations using parameters estimated under the restriction of no-heterogeneity. That is, $\sigma_\nu = 0$ and $\delta = 0$.[10] Without preference heterogeneity, 16% of the increase in demand comes from substitution of market exits (compared to 6% at the estimated parameters), the increase in supply of used jets to first time buyers is only 84% of the increase in demand for new jets, and the change in revenue is 26.2% larger.

This comparison illustrates the importance of measuring preference heterogeneity in determining the effect of buyback on the first time buyers' market. At the estimated parameters, consumers that choose to exit the market are likely to have high values of α_i^p and β_i^0, and are therefore unlikely to be on the margin between exiting and upgrading to new that is affected by buyback. Buyback has a large effect of the supply of used units precisely because of the degree of preference heterogeneity among jet holders.

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

[10] Parameters estimated under this restriction are recorded in Appendix Table A.5.

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.

7.1 Computing Equilibrium Manufacturer Revenue

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. This requires a model of how new and used prices are set. I develop a two step procedure which sets equilibrium new jet prices for fixed used jet prices, and equilibrium used jet prices for fixed new jet prices. By iterating these two steps at a given parameter vector, I find an equilibrium in both new and used jet prices.

I assume new jet prices are set by profit maximizing firms. Let J_{ft} be the set of new jets offered by firm f in year t. Firm profit is given by

$$\Pi_{ft} = \sum_{j \in J_{ft}} D_{jt}(p_{jt} - c_{jt}). \tag{16}$$

Where total demand for jet j is $D_{jt} = \sum_{k \in J_{t-1} \cup 0} M_{kt} s_{kjt}(\hat{\theta}, \hat{\xi})$.[11] Rearranging the firm's first order condition yields the familiar markup expression

$$p_{jt} = c_{jt} - \frac{D_{jt} + \sum_{k \in J_{ft}, k \neq j} \frac{\partial D_{kt}}{\partial p_{jt}}}{\frac{\partial D_{jt}}{\partial p_{jt}}}. \tag{17}$$

I evaluate the markup term on the right hand side of equation 17 at the estimated parameters and unobservables obtained from the BLP fixed point algorithm, $(\hat{\theta}, \hat{\xi})$, for each jet-year pair, and use the observed prices to back out marginal costs, c_{jt}.[12] The distribution of implied markups, defined as $\frac{p_j - c_j}{cj}$, is illustrated in Figure 2. The mean markup is 27%.

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 new and used jet prices as follows. I

[11] I assume firms sell both year t and year $t-1$ models in year t, since much of the stock manufactured in year $t-1$ is likely to remain unsold until the subsequent year. I therefore adjust the demand faced by firms for year $t-1$ models subtracting the current stock of 1 year old jets held by consumers. Note also that firms are not forward looking - price is set optimally each year without considering how changes in the distribution of jet holdings will affect future demand.

[12] Throughout, I use the rental rates described in Section 4 in place of sticker prices. The marginal costs should therefore be interpreted as marginal costs in 'rental rate units'. If all jets depreciated at the same rate and the interest rate was constant, this would be equivalent to multiplying all prices by a constant.

Figure 2: Implied Markups

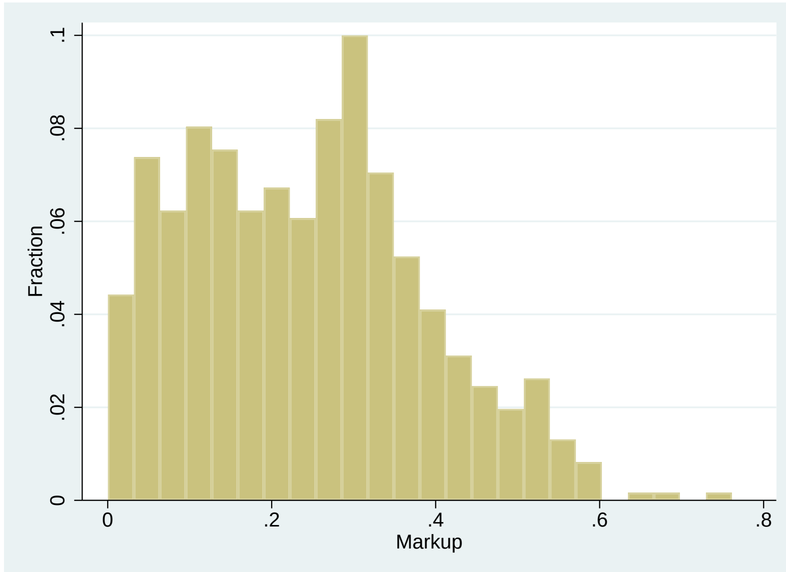

Notes: This figure is a histogram of the implied markups $\frac{p_j - c_j}{c_j}$, where p_j is the observed price and c_j is the implied marginal cost, among all new jets available in the estimation sample, where jets are defined by a manufacturer-year-segment. The figure drops the 0.8% of estimated markups that are negative and the 1.3% of estimated markups that are greater than 1. The marginal costs are backed out from the manufacturers' first order conditions as described in the text.

compute market shares according to equation 8. For new jets, I then compute the prices implied by equation 17. As usual (see Nevo, 2001), iteratively substituting price vectors into these equations yields equilibrium new jet prices for a fixed vector of used jet prices.

For used jet models (which are only available after the first year of the sample), I compute the excess demand as

$$ED_{jt}(p, \theta) = D_{jt}(p, \theta) - M_{jt}. \tag{18}$$

That is, the excess of the number of units demanded (where demand includes the decision of holders of jet j not to upgrade) over the existing stock of jet j held at the beginning of year t. 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.[13]

Starting with the first period in the sample, I alternate between solving for equilibrium new jet prices, fixing used jet prices, and equilibrium used jet prices, fixing new jet prices. When both new and used jet prices have converged, I record the resulting jet holdings and preference distributions at the end of period t and move to period $t + 1$.

[13]Since 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 18 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.

Table 9: Equilibrium Simulations

	Baseline Parameters			No Heterogeneity		
	No Buyback	Buyback	Δ	No Buyback	Buyback	Δ
Upgrades to New	1453	1802	349	1445	1877	432
Used Jet Supply to First Time Buyers	16133	16512	379	15980	16428	448
New Sales to First time Buyers	5737	5628	-109	5575	5553	-22
Average Used Jet Price	5.897	5.816	-0.081	5.848	5.816	-0.032
Manufacturer Revenue	6.738	7.066	0.268	6.520	7.066	0.546
Manufacturer Profit	1.330	1.393	0.063	0.829	0.877	0.048

Notes: Revenue is in billions of 2009 $ computed using "rental rate" prices. 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.5

7.2 The Effect of Buyback on Equilibrium Revenue

To illustrate the effect of buyback on new jet sales, manufacturer revenue, and profit 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 gain in revenue from additional upgrades to new jets 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 9 record total quantities sold, average prices, and total manufacturer revenue and profit in the buyback and no-buyback simulations at the estimated parameters for the entire sample period of 1961-2000.

Moving from the no-buyback equilibrium to the buyback equilibrium increases the number of new jets purchased as upgrades by 349, or 24%. 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 379 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 15, this implies that the number of market exits also increases in equilibrium when buyback is introduced. Note that the numbers here differ form those recorded in Table 8 because of equilibrium price effects. While exits fell in the fixed-price simulations recorded in Table 8, when prices are allowed to adjust, buyback causes a slight increase in market exits by pushing up the prices of new jets.

The increase in used jet supply lowers the average used jet "rental rate" price by 8% from $589.7

thousand in the no buyback equilibrium to \$581.6 thousand in the buyback equilibrium. This fall in price induces substitution towards used jets among first time buyers. Of the additional 379 used jets sold to first time buyers, 109 are from consumers who would have purchased new jets under the no-buyback equilibrium prices. Substitution away from new jets in the first time buyer market cannibalizes increased sales in the upgrade market, and reduces the change in new jet sales by about 31%. The change in total manufacturer revenue from introducing buyback is \$268 million. Comparing this figure to the revenue change of \$473 million recorded in Table 8, in which prices were held constant, indicates that this equilibrium cannibalization effect reduces the increase in manufacturer revenue from introducing buyback is reduced by 43%.

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 captured by the parameters σ^p and π. In the second set of three columns in Table 9, 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 448 units, but only reduces the number of new jets purchased by first time buyers by 22 units. Under no-heterogeneity, cannibalization only reduces the change in revenue by 7.6% relative to the fixed price simulation of table 8. 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.

7.3 Optimality of Buyback

Although these results show that buyback increases manufacturer profit, they does 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 level of costs associated with operating buyback.

To compute ranges of costs for which offering buyback is a dominant strategy for each firm, and

Table 10: Buyback Cost Ranges

Manufacturer	Dominant Strategy Costs		Prisoner's Dilemma Costs		$\frac{b_m}{E(\alpha_i^p \mid hold)}$
	Min	Max	Min	Max	
Bombardier	0	40.3	40.3	61.4	36.2
Cessna	0	21.8	21.8	37.4	32.4
Dassault	0	79.2	79.2	111.0	39.6
Gulfstream	0	139.7	139.7	170.3	57.0
Raytheon	0	94.5	94.5	111.6	74.7
Other	0	37.3	37.3	54.2	38.8

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. Mx prisoner's dilemma costs are computed by dividing the difference between buyback and unilateral deviation equilibrium profits by the number of buyback-eligible upgrades. Profits used in the calculations are recorded in Appendix Table A.6.

rarealitynges 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 equilbrium 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 upgraders from other brands with switch brand to m. Appendix Table A.6 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.

Table 10 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 column records the estimated value of $\frac{b_m}{E(\alpha_i^p \mid hold)}$, the per-unit reduction in transaction costs from buyback evaluated at the average price coefficient among jet holders. If buyback was a one to one transfer of transaction costs from consumers to manufacturers, then the cost of buyback would be equal to this reduction in transaction costs. The computed cost ranges are consistent with this hypothesis. For all but one of the manufacturers the reduction in transaction costs is within the range of dominant strategy costs. For Cessna, the reduction in transaction costs falls in the range of costs for which offering buyback is not a dominant strategy but is a best

Figure 3: Distribution of Jet Holders in 2000

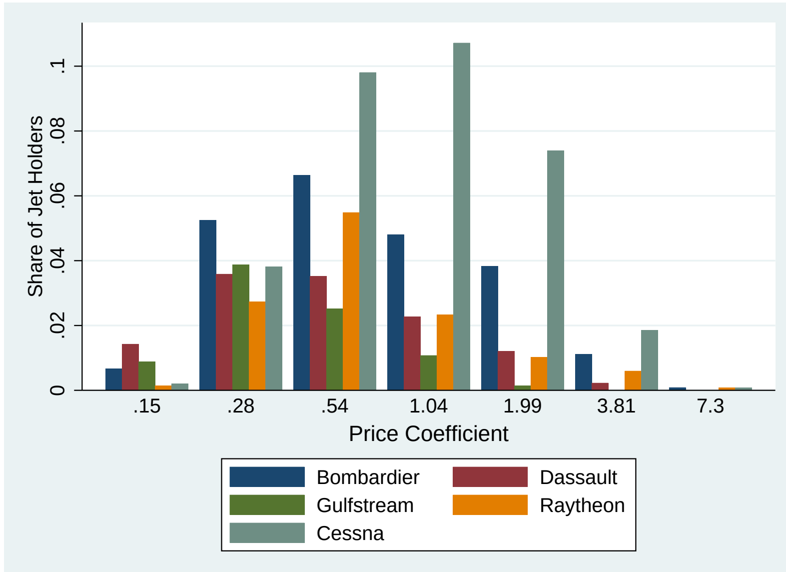

Notes: Bars record the distribution of α_i^p among holders of each jet brand in equilibrium simulations at the year 2000, using the estimated parameters. As explained in footnote 9, the distribution of α_i^p is discrete with 11 points of support. Here, I only display the smallest 7 points, as there is negligible weight on the upper tail among jet holders.

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.

Cessna does not benefit from buyback to the same extent as other firms because of the position of its products in the market. As recorded in Table 3, 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 Figure 3, which records the distribution of jet holders in 2000 across manufacturers and values of α_i^p. The modal value of α_i^p for holders of Cessna jets is 1.04, while the modal value for all other manufacturers is between 0.28 and 0.58. 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 58.3 fewer new jets to first time buyers in the no-buyback scenario compared to the buyback scenario. This is 56% of the increase in new Cessna jets sold to upgraders. The equivalent cannibalization rate for the entire market is 24%. Furthermore, the selection of buyers into Cessna ownership means that they are unlikely to switch to another,

Figure 4: Buyback Incentive and Transaction Cost Level

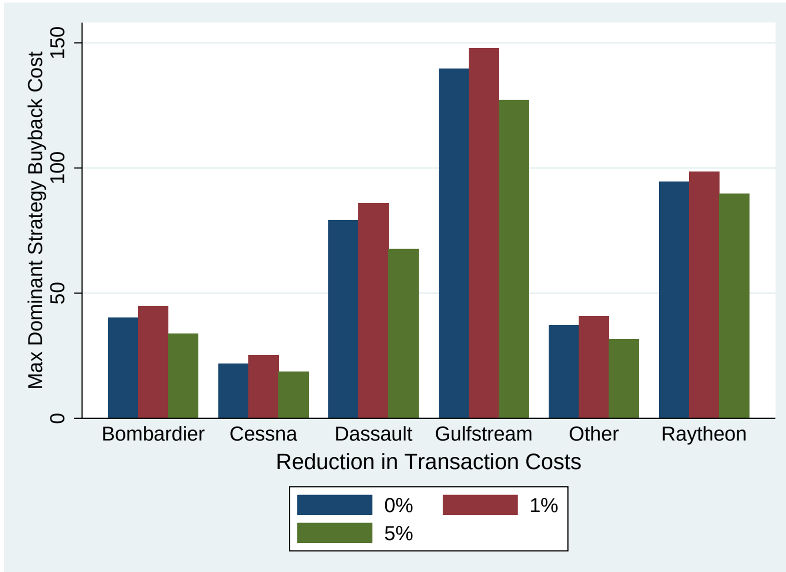

Notes: Bars record the computed maximum dominant strategy buyback costs for each firm, equivalent to those in the second column of Table 10, under simulations at the estimated parameters and with a 1% and 5% reduction in τ from the estimated value.

more expensive, brand. In the no-buyback scenario, Cessna receives 284 buyback-eligible upgrades, compared to between 51 and 204 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 operators.

How does the incentive for firms to accept trade-ins depend on the level of transaction costs faced by consumers? A reduction in transactions costs, τ, could be caused by the entry of more dealers willing to buy and resell jets, a reduction in information frictions from the advent of online platforms such as Controller.com and AvBuyer, or by the entry of leasing agents and fractional ownership companies such as Netjets amd Xojet who facilitate upgrades.[14] It is not obvious whether reduced transaction costs increase or decrease the incentive for manufacturers to engage with the secondary market. When the demand for upgrades is low, a reduction in τ will increase the elasticity of demand to changes in b_m, increasing the incentive to engage in buyback. On the other hand, lower values of τ increase the number of infra-marginal upgrades that happen under the no-buyback scenario, increasing the cost of implementing buyback. In the limit as $\tau \to 0$ consumers will upgrade every period, and elasticity of overall demand for upgrades to buyback will be 0.

Figure 4 illustrates how the maximum "dominant strategy" cost, as computed in the second column of Table 10, changes for different levels of τ. I simulate equilibrium profits under buyback and

[14] Gilligan (2005) documents the increasing share of leased aircraft since the mid-1990s. In the context of the model outlined in Section 4, in which the "rental rate" is the relevant price, one can think of a lease as an ownership arrangement that enables an owner to upgrade without finding a buyer for their used jet, thus reducing τ.

no-buyback scenarios for 1% and 10% reductions in τ from the estimated level. The bars record, for each firm, the maximum cost of buyback for which profit is higher in the buyback than in the no-buyback simulations. the blue bars record these values at the estimated level of τ, the red bars for a 1% reduction, and the green bars for a 5% reduction in τ. Notice that, for each firm, there is a non-monotonic effect of reducing τ on the maximum cost. A 1% reduction in transaction costs increases the maximum cost for which firms have an incentive to implement buyback, but a 5% decrease reduces the maximum cost from the baseline level. Buyback becomes more profitable for small reductions in τ because of the increased elasticity of demand to b_m, but for larger reductions in τ the increased cost of paying for a larger number of infra-marginal upgrades outweighs the benefit of increased demand.

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 estimate a 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 estimate transaction costs in the used jet market to be \$734 thousand, or 11% of the mean sale price. Manufacturer buyback reduce the transaction costs faced by owners by between 4.4% and 10.2%. At fixed prices, removing buyback from all manufacturers would cause the demand for new jets from upgrading consumers to fall by 28%.

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 revenue 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 manufacturer revenue by 43%.

The extent to which substitution away from new jets in the first time buyer market reduces the benefit of buyback to the manufacturer depends on the distribution of preferences among consumers. I show that when there is no heterogeneity in consumer preferences, the cannibalization effect only reduces the revenue gain from buyback by 7.6%.

I use these equilibrium simulations to 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 all firms except Cessna are in the dominant strategy region. The fact that Cessna does not benefit as much from buyback as other firms illustrates the importance of the substitutability of a firm's new units with used units among first time buyers to the effect of buyback on firm revenue. Finally, I study how the incentive to operate buyback changes with the baseline level of transactions costs in the industry and find a non-monotonic effect of reducing transaction costs on firms' incentive to buy back used units.

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. This type of policy is closely related to price discrimination between first time buyers and upgraders - it is similar to an effective price cut for upgrading consumers. The results suggest that the extent to which the firm is able to profitably engage in such price discrimination depends on the size of the cannibalization effect, in particular whether used units are close substitutes to the firm's new units. The framework developed in this paper could be used more broadly to analyze this type of price discrimination in durable goods industries.

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., 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, Tradeins, 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

Gowrisankaran, G., and M. Rysman. "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 Chocie Behavior." *Journal of Business Economics & Statistics* 15 (3): 310-327

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 Price Data

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. If one model's price is missing then the choice probabilities cannot be computed. 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.

A.2 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.3 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,

$$Q_N = S_N(p_N) = D_1^N(p_U, p_N, b) + D_0^N(p_U, p_N) \tag{19}$$

$$Q_U = S_U(p_N, p_U, b) = D_0^U(p_U, p_N). \tag{20}$$

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 19 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 20 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 $p'_U < p_U$. dollowing 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 20.

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

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.

Table A.1: Multiple Jet Ownership

	Share of Owner-Years		
	Dealers	Corportations	Other
Multiple Jets Held	40%	19%	26%
At Least One Jet Purchased	52%	25%	26%
Multiple Jets Purchased	20%	2%	4.6%

Notes: Sample includes all owner-years.

Table A.3: First Stage and 2SLS

	First Stage				OLS	2SLS
Dependent Variable:	Price				$log(s_{jt}) - log(s_{0t})$	
Price					-0.053***	-1.692***
					(0.003)	(0.190)
Substitute Jets (1)	-0.004***			0.002		
	(0.001)			(0.002)		
Substitute Jets (2)		-0.006***		-0.005***		
		(0.001)		(0.002)		
Models Available			-0.021***	-0.020***		
			(0.003)	(0.003)		
1st Stage F-Stat						25.747

Notes: Stars indicate statistical significance as follows: *** 1%, **5%, * 10%. 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 numebr 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.

A.4 Tables and Figures

Table A.2: Buyback Patterns for Different Sub-Periods

		Share of Buybacks		Share of Potential Buybacks
Manufacturer	Years	Upgrades to New	Own Brand	Sold to Manufacturer
All	1961-1980	78%	84%	20%
All	1981-1990	80%	95%	24%
All	1991-2000	88%	88%	50%

Notes: Statistics identical to those recorded in Table 4, recorded separately for transactions in three time periods. "All" manufacturers includes only the top five manufacturers as defined in the text.

Table A.4: Micro-Moments

Moment	Data	Model	Moment	Data	Model
1	0.048	0.049	12-14	8.818	8.870
				8.002	7.939
2	0.130	0.130		6.237	6.250
3	0.363	0.340	15-17	0.096	0.099
				0.126	0.128
4	0.270	0.272		0.163	0.152
5	0.526	0.531			
6-11	0.082	0.077			
	0.074	0.077			
	0.135	0.126			
	0.276	0.259			
	0.068	0.063			
	0.191	0.181			

Notes: Table records the values of the micro-moments used in estimation as defined in Table 6, both in the data, and as implied by the model at the estimated parameters.

Table A.5: No-Heterogeneity Parameter Estimates

Parameter	Estimate	SE	Parameter	Estimate	SE	Buyback Parameter b_m		
α^p	0.722	0.040	τ	8.769	0.039	Bombardier	0.358	0.159
σ^p	0	.	τ^{exit}	8.636	0.026	Cessna	0.394	0.360
α^{age}	-0.134	0.008	α^{sb}	1.471	0.059	Dassault	0.677	0.401
α^n	2.946	0.201	$\alpha^n_{upgrade}$	2.485	0.182	Gulfstream	1.163	0.883
α^{range}	0.004	0.002	α^0	-12.531	0.476	Raytheon	0.777	0.357
α^{power}	0.109	0.042	δ	0	.	Other	0.080	0.314
α^{weight}	-0.000	0.000	π	0	.			
$\alpha^{GDPgrowth}$	19.237	1.902						

Notes: Table reports estimated parameters and standard errors for the demand model under the restriction of no preference heterogeneity. Prices are in hundreds of thousands of 2009 dollars. Age is measured in years, range is measured in km, weight is measured in kg, and power is measured in kN. GDP growth is % change from the previous year.

Table A.6: Firm Profit

Manufacturer	BB units	No BB Profit	BB Profit	Unilateral Deviation Profit
Bombardier	281.666	0.32912	0.34047	0.32317
Cessna	407.263	0.28116	0.29005	0.27482
Dassault	124.150	0.22575	0.23558	0.2218
Gulfstream	105.550	0.22893	0.24368	0.22571
Raytheon	159.721	0.16364	0.17874	0.16091
Other	77.717	0.10171	0.10461	0.10040

Notes: Table records manufacturer profits in billions of 2009 dollars computed at "rental rate" prices 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.