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

Charles Hodgson[†]

July 8, 2016

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 also 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 removing buyback decreases demand for new jets among replacement buyers by 45% at fixed prices and raises used jet prices by 7% in equilibrium. Although the effect on used jet prices induces substitution towards new jets, manufacturer revenue is lower without buyback.

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 value for quality to sell depreciated goods to consumers with lower values for quality. 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 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 typically 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 consumers to substitute from buying used units to buying new units, conditional on upgrading. This increase in demand has an unambiguously positive effect on manufacturer revenue. 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 consumers to substitute *away* from new units, potentially lowering manufacturer revenue. The firm's decision to offer buyback depends on whether the benefits from directing trade towards their own new units outweighs the costs of increasing the volume of used goods traded in the secondary market[1].

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

---

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

manufacturers might offer buyback in equilibrium even though they would all have higher profits if they jointly agreed not to accept trade-ins.

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 1980 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 $1.929 million, or approximately 22% of the average jet price. I find that manufacturer buyback schemes eliminate between 7% and 18% 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 265, or 45%. I decompose this change in demand into several margins of substitution, and find that 86% can be attributed to substitution from upgrading to holding. That is, the decrease in the total number of upgrades, which corresponds to the change in the number of used jets supplied, is 86% of the decrease in new jet demand.

I simulate the equilibrium response of prices to these shifts in demand and supply and test whether manufacturer revenues are higher or lower in the no-buyback counterfactual equilibrium. I find that removing buyback raises used jet prices by 7%, inducing substitution from used to new jets among first time buyers. However, the direct effect of removing buyback on the demand for new jets more than offsets the effect on the supply of used jets in equilibrium. Total manufacturer revenues are 8% lower in the no-buyback equilibrium, with each of the major manufacturers individually worse off. This suggests that the equilibrium in which all manufacturers operate buyback is Pareto optimal, and that manufacturers would not benefit from jointly agreeing not to accept trade-ins.

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 re-selling 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. I differ from him in that I use transaction level data and do not use a dynamic demand model. 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

---

[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.

of demand for new and used jets. Section 5 describes the estimation and identification of the model, and results are described in Sections 6 and 7. Section 8 concludes.

2 Data and Setting

2.1 The Market for Business Jets

The market for used business jets is typical of a durable goods industry with active trade in used goods as well as prolonged holding. Between 1980 and 2000, the five leading business jet manufacturers - Cessna, Bombardier, Raytheon, Dassault, and Gulfstream - sold 7,182 new jets. Over the same period, there were 24,153 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. Manu-

facturers 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." The cost to the manufacturer of refurbishing and 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.

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.[3] 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 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 row of Table 2 records the mean and standard deviation of the prices of new jets of different sizes and production years. I express all prices in year 2000 dollars.

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.[4] 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 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.

---

[3]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.

[4]In particular, for the period 1995-2000, only 3.8% of purchases are of a jet with a missing price. In Appendix Table A.3 I repeat the main analysis using data from this period and obtain results very close to the main estimates reported in Section 5.

2.3   Estimation Sample

For the estimation of the model presented in Section 4, I restrict the sample to owners classified as corporations. This reduces the number of owners in the sample (excluding manufacturers and dealers) from 20,204 to 14,110. I define a time period in the data as one calendar year. I focus on the 20 years 1980-1999, as the market structure was relatively stable over this period. I drop owners who purchased their first jet before 1980, as early adopters of business jets may have systematically different preferences. This leaves me with 10,734 unique owners, or about 53% of the population of owners. 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*. The algorithm used to construct this panel is described in Appendix Section A.2. This procedure generates a panel of jet owners observed once a year, holding at most one jet each year. The mean owner is in the sample for 4.7 years and makes 0.214 upgrade purchases. The data used for estimation includes 50,858 owner-year observations.

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 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 how often these manufacturers and manufacturer affiliated dealers buy back used jets from corporations. 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 3 months. The first row indicates that in 40% of upgrades from used to new jets, the used jet is bought by a manufacturer. The corresponding figure is only 5% for used-used upgrades, as recorded in the second row. By way of comparison, the third row of the table records that 5.8% of *all* jets sold by corporations are bought by manufacturers, unconditional on whether the corporation upgraded. Columns two to six of Table 4 report these percentages conditional on the jet *sold* being of the given brand. The share of buybacks among used-new upgrades is between 24% and 49% for each of the major brands. These statistics confirm that manufacturers buy back used jets from owners who wish to upgrade to new units. Owners take advantage of buyback in a significant share of observed used-new upgrades. The 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, and that their involvement in the used jet market is largely through accepting trade-ins for used-new upgrades. Appendix Table A.1 repeats this exercise for different time periods, demonstrating that this pattern holds throughout the years 1980-1999.

The data also indicates that manufacturers buy back used jets of their own brand almost exclusively. Table 5 records own brand jets bought back as a share of all jets bought back for each of the five major manufacturers. The shares are over 75% for each manufacturer, and as high as 97% for Cessna. That is, 97% 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. Such a requirement could be rationalized as inducing customer loyalty. However, persistent heterogeneity in brand preference may also be responsible for this pattern. The fact that consumers who upgrade to new Cessna jets tend to trade in used Cessna jets may simply reflect the fact that some consumers have a persistent preference for Cessna aircraft. In the demand model estimated in Section 4 I allow for persistence in brand preference.

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 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 a 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 5 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 6 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 15 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 9.5 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 6 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 benefiting from buyback in either market. The difference in market shares between the two markets is 0.8 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 market shares is reduced from 9.5 to 4 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 first time buyers and upgrade buyers and between same brand and different brand upgrades reported in Table 6 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.[5]

During each year, t, the set of existing new and used jet models is J_t. A model $j \in J_t$ is described by its observable characteristics, x_{jt} and its price, p_{jt}. I assume that consumers hold at most one jet in any given period. Consumer i's flow utility from holding jet j in period t is given by:

$$u^i_{jt} \;\; = \;\; x_{jt}\alpha + \xi_{jt} + \epsilon_{ijt} \tag{1}$$

Consumer i's flow utility from buying jet j in period t, if they does not currently hold a jet, is given by:

$$\tilde{u}^i_{0jt} \;\; = \;\; x_{jt}\alpha - p_{jt}\alpha^p_i + \xi_{jt} + \epsilon_{ijt} \tag{2}$$

Where ξ_{jt} is unobserved product quality. ϵ_{ijt} is an i.i.d. type 1 extreme value shock to preferences.

---

[5]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".

If consumer i holds a jet of type k, then her flow utility from selling that jet and not replacing it is given by:

$$\tilde{u}^i_{k0t} \quad = \quad p_{kt}\alpha^p_i + \epsilon_{i0t} \tag{3}$$

Finally, if consumer i holds a jet of type k I allow her utility from selling that jet and buying a jet of type j to depend on a vector of observable characteristics of both jets, $\tilde{x}_{kjt}$. Her flow utility from upgrading from jet type k to jet type j is given by:

$$\tilde{u}^i_{kjt} \quad = \quad \tilde{x}_{kjt}\tilde{\alpha} - (p_{jt} - p_{kt})\alpha^p_i - \tau_{kj} + \xi_{jt} + \epsilon_{ijt} \tag{4}$$

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.

As I discuss futher 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. 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 "price" from here on.

The model yields the following probabilities for consumers who *hold a jet of type k*: the probability of buying a jet of type j,

$$P^i_{kjt} = \frac{exp(\tilde{x}_{kjt}\tilde{\alpha} - (p_{jt} - p_{kt})\alpha^p_i - \tau_{kj})}{exp(p_{kt}\alpha^p_i) + exp(x_{jt}\alpha + \xi_{jt}) + \sum_{l \in J_t} exp(\tilde{x}_{klt}\tilde{\alpha} - (p_{jt} - p_{lt})\alpha^p_i - \tau_{kl})} \tag{5}$$

the probability of continuing to hold a jet of type k,

$$P^i_{kkt} = \frac{exp(x_{jt}\alpha + \xi_{jt})}{exp(p_{kt}\alpha^p_i) + exp(x_{jt}\alpha + \xi_{jt}) + \sum_{l \in J_t} exp(\tilde{x}_{klt}\tilde{\alpha} - (p_{jt} - p_{lt})\alpha^p_i - \tau_{kl})} \tag{6}$$

and the probability of selling the currently held jet and not replacing it,

$$P^i_{k0t} = \frac{exp(p_{kt}\alpha^p_i)}{exp(p_{kt}\alpha^p_i) + exp(x_{jt}\alpha + \xi_{jt}) + \sum_{l \in J_t} exp(\tilde{x}_{klt}\tilde{\alpha} - (p_{jt} - p_{lt})\alpha^p_i - \tau_{kl})} \tag{7}$$

Conditional on buying some jet, the probability of buying a jet of type j for consumers who *do not hold* any jet is:

12

$$P^i_{0jt} = \frac{exp(x_{jt}\alpha - p_{jt}\alpha^p_i)}{\sum_{l \in J_t} exp(x_{lt}\alpha - p_{lt}\alpha^p_i)} \qquad (8)$$

Note that τ_{kj} does not enter into P^i_{0jt}. Buyers who do not have to sell an existing jet in order to upgrade are assumed not to pay any transaction costs and therefore not to benefit from any buyback schemes.

4.2 Econometric Specification

In the main specification, the vector x_{jt} contains an indicator for each of the jet size classifications (small, medium, large), an indicator for each manufacturer, the jet's maximum range (in km), total engine power (in kN), and maximum takeoff weight (in kg), the jet's age (in years), and an indicator for whether the jet is new, n_{jt}. Let α^n be the element of parameter vector α that multiplies n_{jt}, and α^{age} be the element that multiplies jet age.

Let $m(j)$ indicate the manufacturer of jet j. I assume that the term $\tilde{x}_{kjt}\tilde{\alpha}$, which enters the flow utility obtained by upgrading from a jet of type k to a jet of type j, can be written as:

$$\tilde{x}_{kjt}\tilde{\alpha} = x_{jt}\alpha + n_{jt}\beta^n + I(m(j) = m(k))\beta^{sb} \qquad (9)$$

The characteristics of jet j, x_{jt}, enter the flow utility of upgrading from jet k in the same way that they enter other utilities, and are multiplied by the same parameters, α. However, the utility obtained from upgrading to a new jet contains two additional terms. The inclusion of parameter β^n allows upgrade consumers to have a different preference for newness $(\alpha^n + \beta^n)$ to first time buyers (α^n). The parameter β^{sb} is the coefficient on an indicator for whether j and k are of the same brand. This term is included to separately identify persistent heterogeneous brand preferences or consumer inertia in brand choice.

For the main specification estimated below, I write τ_{kj} as:

$$\tau_{kj} = n_{jt}I(m(j) = m(k))b_{m(j)} - \tau \qquad (10)$$

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, n_{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.

Finally, I assume that the individual specific price coefficient, α_i^p, is drawn iid across owners according to the following lognormal specification:

$$\alpha_i^p = exp(\alpha^p + \sigma^p \nu_i) \tag{11}$$

$$\nu_i \sim N(0,1) \tag{12}$$

Note that the lognormal specification for the price coefficient restricts α_i^p to be greater than 0. The assumption that there is heterogeneity in the price coefficient α_i^p 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.

The parameters to be estimated are the mean utility parameters $\{\alpha, \beta^n, \beta^{sb}\}$, the distributional parameters of the individual specific price coefficient $\{\alpha^p, \sigma^p\}$, the transaction cost τ, and the buyback parameter b_m for each manufacturer m.

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 42% 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

14

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

The model describes jet owners' choices conditional on entering the market. In particular, P_{0jt}^i, defined above, is the probability of first time buyer i choosing jet j *conditional on choosing some jet* at date t. The estimation procedure below maximizes the likelihood of the observed decisions of each consumer, j, conditional on them buying some jet at the date they enter the market, t_i.

The model is estimated by simulated maximum likelihood (Train, 2003). Denote the jet owned by owner 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 $t_i \geq 1980$ as the date that owner i purchases their first jet and $T_i \leq 1999$ as the last observation of consumer i. Note that if an owner sells their jet and does not replace it, I treat them as having left the market forever. If $T_i < 1999$, then owner i sells their jet and does not replace it, exiting the market permanently at T_i. If $T_i = 1999$, then the owner may continue to hold their jet beyond the end of the sample period. Note that

15

since I observe data for the year 2000, I know whether owners sold, held, or upgraded their jets in 1999. For each owner, I observe a sequence of jets j_{it} with likelihood:

$$L_{i|\nu_i} = \prod_{t=t_i}^{T_i+1} \prod_{j \in J_t} \prod_{k \in J_{t+1}} P_{0kt}^{i1(j_{it}=0)1(j_{it+1}=k)} P_{jjt}^{i1(j_{it}=j)1(j_{it+1}=j)} P_{jkt}^{i1(j_{it}=j)1(j_{it+1}=k)} \qquad (13)$$

This is the likelihood of a sequence of decisions conditional on the owner choosing to buy a first jet. It is also conditional on the idiosyncratic part of the price coefficient, ν_i. Likelihood unconditional on ν_i is:

$$L_i = \int L_{i|\nu_i} df(\nu_i) \qquad (14)$$

Log Likelihood of the data is then:

$$LL = \ln\left(\prod_i L_i\right) \qquad (15)$$

Since the integral in the likelihood expression is difficult to compute, estimation is performed by simulated maximum likelihood. Replace L_i with:

$$\check{L}_i = \frac{1}{R} \sum_{r=1}^{R} L_{i|\nu^r} \qquad (16)$$

Where R is the number of simulations and v^r is a draw from the $N(0,1)$ distribution. For the estimates produced below, $R = 50$. Parameters are estimated by maximizing the simulated likelihood of the data.

An important issue in estimating these models is how to deal with the unobserved part of jet quality, ξ_{jt}. It is likely that p_{jt} is positively correlated with ξ_{jt} - that prices are higher when the unobserved demand shifter is higher. Ignoring the presence of these unobservables would likely bias down the estimates of the distribution of α_i^p. To deal wth this, I adopt the control function approach suggested by Petrin and Train (2009). In particular, I assume that ξ_{jt} is correlated only with p_{jt} and no other variable that enters consumer utility, and that p_{jt} can be expressed in 'reduced form' as:

$$p_{jt} = A(x_{jt}, z_{jt}) + e_{jt} \qquad (17)$$

Where x_{jt} includes jet all jet characteristics that enter the utility model (except for price) and z_{jt} is a vector of instruments, and e_{jt} is an unobservable which can be written as $e_{jt} = B(\xi_{jt})$. I assume that x_{jt} and z_{jt} are independent of ξ_{jt}. The function $A(\cdot,\cdot)$ can therefore be estimated consistently using least squares. Rewriting this equation to express

ξ_{jt} as a function of observables yields: $\xi_{jt} = B^{-1}(p_{jt} - A(x_{jt}, z_{jt}))$. This expression can then be substituted into the consumer's utility function in place of ξ_{jt}. In practice, I assume $A(\cdot, \cdot)$ is linear in its arguments and estimate it using OLS. I also assume $e_{jt} = \frac{1}{\lambda}\xi_{jt}$, and I substitute the proxy $\hat{\xi}_{jt} = \lambda\hat{e}_{jt}$ into the consumer's utility function, where $\hat{e}_{jt}$ are the residuals from the estimation of $A(\cdot, \cdot)$. λ is then a parameter in the utility function that is estimated using the simulated likelihood approach described above. This is the procedure suggested by Petrin and Train (2009) for a very similar model.

Ideally, the functional form $A(\cdot, \cdot)$ and the set of instruments should be motivated by economic theory. For example, firms' pricing equations for new units depend on observed and unobserved demand characteristics as well as supply shifters such as production costs. In this industry, the supply of both new and used jets in a given year depends on the number of units of each jet model currently held by owners. For used jets, 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 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 $t-1$. The demand for new jets of type j for an upgrading buyer depends on the jet model the buyer currently holds. Therefore the firms' pricing equation for new jets, discussed further in Section 7 below, depends on the share of owners who currently hold each type of jet. The optimal price set by a manufacturer will depend, for example, on the share of owners who hold that manufacturer's own models. These ownership shares are therefore supply shifters for new units which are also uncorrelated with ξ_{jt}. The vector of instruments z_{jt} therefore includes the number of owners of jet j at date t (which is 0 if j is a new model), and the number of owners of each manufacturer's jets.

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 6, 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 β^n and b_m, both of which shift the utility of new jets for replacement buyers. Similarly, the tendency of replacement buyers to buy jets of the same brand (manufacturer) as those they sell is captured by both β^{sb} and b_m. b_m is separately identified by the interaction of these two effects.

In particular, β^n is identified by the difference between the new jet share of first time purchases and the new jet share of replacement purchases of a different brand. That is, the probability that an owner of a jet of type k upgrades to a new jet of type j with $m(k) \neq m(j)$,

17

conditional on upgrading, is different that the probability that a first time buyer chooses jet j because of β^n - the additional preference for newness among replacement buyers. β^{sb} is identified by the difference between the share of replacement purchases of used jets of the same brand and the share of each brand in the first time market. For example, the extent to which the probability of a Cessna owner upgrading to another used Cessna, conditional on upgrading to some jet, exceeds the probability of a first time buyer purchasing a used Cessna is explained by β^{sb}. Since there is only one parameter for all brands, it is the average of this effect across brands that determines the magnitude of β^{sb}. Note that, as highlighted by Dubé et al. (2010), β^{sb} could measure either inertia in brand preference or unspecified persistent heterogeneity in brand preferences. The distinction is, however, not important for this paper.

The parameters b_m capture the preference for *new jets* of the *same brand* beyond the effects captured by β^n and β^{sb}. b_m is therefore identified for each brand m by the extent to which the share of new jets among all 'same brand' replacement purchases is greater than the share of new jets among all 'different brand' replacement purchases, as illustrated in panel C of Table 6. 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.

The transaction cost parameter, τ, is identified by the frequency with which owners in the sample upgrade. If $\tau = 0$, assuming prices do not change over time, 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 distribution of preferences over price, parameterized by α^p and σ^p, is identified by the extent to which consumers prefer lower priced jets on average (which pins down α^p), and the extent to which different consumers have persistent heterogeneity in preferences.

The parameters on age and newness, α^{age} and α^n, are identified by the extent to which consumers prefer newer jets, conditional on prices. Note that the inclusion of a parameter on newness and age allows for the depreciation in quality to be non-linear. The specification allows an initial drop (or rise) in quality when a jet transitions from being new to being used, after which depreciation is linear. Finally, the segment and brand dummies and the jet characteristic parameters in α are identified by the market shares of each model (defined by a brand, segment and year of manufacture), conditional on prices.

18

6 Results

6.1 Estimated Parameters

The main parameter estimates are presented in Table 7.[6] All parameters are statistically significant and of the expected sign. Columns 1 and 3 present simulated maximum likelihood estimates that ignore the presence of jet-year unobservable ξ_{jt}. Columns 2 and 4 present estimates of the model using the control function technique. The first stage regression of price on instruments is recorded in Appendix Table A.2. The coefficient on price which enters the utility function is $\alpha_i^p = exp(\alpha^p + \sigma^p \nu_i)$. When ignoring the demand unobservable, the expected marginal utility of one dollar is $E(\alpha_i^p) = 3.854 \times 10^{-6}$. Since the average price (rental rate) in the sample is \$707,000, it is more instructive to scale this parameter to the marginal utility of one million dollars, $1,000,000 \times E(\alpha_i^p) = 3.854$. The corresponding value estimated using the control function methodology is 4.329. As expected, including controls for ξ_{jt} increases the mean of the price coefficient distributions. This confirms the intuition that unobserved jet quality is positively correlated with price, and failing to control for ξ_{jt} will bias down the mean price coefficient. From here on, I therefore focus on the control function estimates.

The baseline model parameter estimates imply that the loss of utility from a jet aging one year, at the expected price parameter, $\frac{\alpha^{age}}{E(\alpha_i^p)}$, is equivalent to an increase in the price of \$17.0 thousand. The difference in utility between a new and used jet is equivalent to a change in price of \$346.7 thousand. Note that the drop in quality once a jet becomes used is over 20 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).

Converted into dollar units at the expected price parameter, $\frac{\tau}{E(\alpha_i^p)}$, the estimated transaction cost is \$1.929 million, approximately 2.7 times the average jet rental rate in the sample, and approximately 22% of the average jet 'sticker price' of \$6.4 million. This result is in line with Schiraldi's (2011) finding that transaction costs (defined in a similar manner) in the market for new and used cars are usually between 20% and 40% of the sale price. As discussed above, this parameter is identified by how frequently the owners in the sample upgrade. 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 also that the size of τ here should 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

---

[6]See Table A.3 in the Appendix for parameter estimates on $t \geq 1995$, for which there are fewer missing prices. The results are very similar to the parameter estimates on the whole sample

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. This would not be the case if 'sticker prices' were used instead of rental rates.

The estimate of β^{sb} is positive and both economically and statistically significant. The parameter estimate indicates the presence of either consumer inertia in brand choice or persistent heterogeneity in brand preferences. The parameter β^{n}, which identifies the extent to which replacement buyers have an additional preference for new units beyond the benefit of buyback, is negative and statistically significant. This suggests that replacement buyers are in fact less likely to buy a new jet than first time buyers, conditional on no buyback scheme being available. The magnitude of the estimated parameter indicates that a new jet is worth around $116.7 thousand less to a replacement buyer than a first time buyer, about 14% of the average price. This negative effect suggests that the fact that first time buyers are more likely to buy used jets than replacement buyers, as recorded in Table 6, is entirely due to the presence of buyback schemes.

Finally consider the manufacturer specific buyback parameters, b_m. As expected, they are positive and significant for each of the major manufacturers, reflecting the empirical fact that replacement buyers choose new jets more often than first time buyers. The estimates tell us that the impact of buyback schemes on demand is equivalent to a reduction of transaction costs of between $125,000, or 7% of transaction costs, for Dassault, and $348,000, or 18% of transaction costs for Cessna. Note that for jets with manufacturer classified as 'Other', the value of buyback is negative and not statistically significant. This is consistent with the composition of this category in the data, which comprises several small and unrelated manufacturers. It would be surprising if there were a strong tendency, for example, for owners of Saberliner jets to upgrade to new IAI jets - Sabreliner and IAI are both categorized as brand 'Other', but are unrelated companies.

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 impact of buyback schemes on demand and supply by comparing a simulation of the demand model, which assumes that all manufacturers operate buyback, at the estimated parameter values to a counterfactual simulation under which no manufacturers offer buyback.

I first simulate the model using the estimated parameters to evaluate model fit. I take as given the date of first purchase for each of the owners in the estimation sample, and then simulate first time purchases and hold or upgrade decisions for each owner for the years 1980 - 1999 using the estimated parameters. Table 8 reports summary statistics for the data in the first column, and for the simulated data at the estimated parameters in the second column. The statistics in the second column are the average of 100 simulations. The new jet market shares in the first time buyer and replacement markets and the share of new jets purchased among same brand and different brand replacements are the key moments for that identify the parameters β^n, β^{sb}, and b_m. These shares are reported in the first four rows of Table 8. The sixth row reports the mean number of replacements per owner, which identifies the transaction cost parameter, τ. For each of these statistics, the data matches the simulations closely.

To measure the impact of buyback on the number of new jets sold, I simulate demand in a market without buyback schemes. I set $b_m = 0 \forall m$ and all other parameters to their estimated values and simulate upgrade decisions. The third column of Table 8 reports the mean values of the statistics for 100 simulations under this counterfactual. The number of new jets purchased as replacements falls by about 45%, from 583 to 318, and the share of new jet purchases among replacement purchases falls from around 25% to 15%. The mean number of replacements per owner falls slightly from the baseline simulation, from 0.219 to 0.195. Notice that the share of new jets purchased as replacements falls below the share of new jets purchased by first time buyers. This reflects the fact that replacement buyers are estimated to have a lower taste for 'newness' than first time buyers. The higher share of new jets purchased by replacement buyers than by first time buyers observed in the data is attributed entirely to buyback schemes. There is some substitution from same brand upgrades to different brand upgrades - the share of new jets purchased among all different brand upgrades increases slightly when buyback is removed.

It is instructive to think of the effect of buyback on demand and supply working through two channels. First, buyback increases the frequency with which customers upgrade. More frequent upgrading increases the supply of used jets since units that are bought back are resold by manufacturers. Second, buyback increases the probability that, conditional on selling their used unit, a consumer will buy a new unit, and decreases the probability that a consumer will buy a used unit or exit the market. Thus, buyback induces substitution towards new units from used units and market exits, and demand for new units would increase even if the frequency of upgrades was fixed. The increase in the demand for new jets induced by buyback is therefore greater than or equal to the increase in the supply of used jets. To see this, note that the increase in supply can be decomposed mechanically as

21

follows:

$$\Delta Used\ Jets\ Supplied = \Delta New\ Jets\ Demanded + \Delta Used\ Jets\ Demanded + \Delta Market\ Exits$$
(18)

Where Δ indicates the buyback level less the no-buyback level, $\Delta Used\ Jets\ Demanded < 0$, $\Delta Market\ Exits < 0$, and $\Delta New\ Jets\ Demanded > 0$. The first coulmn of Table 9 reports this decomposition for the simulated demand. The change in each quantity is reported in levels and as a percentage of $\Delta New\ Jets\ Demanded$. The supply of used units increases by 229, which is 86.5% of the increase in the demand for new units. The number of used jets sold falls by 12.2 units when buyback is introduced, which is only 4.6% of the change in demand of new units. Similarly, the number of market exits falls by just 8.8% of the change in the number of new jets sold. This suggests that buyback policies in the business jet market increase demand mostly through increasing the frequency of replacement and not through substitution to new jets from owners who would otherwise upgrade to used jets or exit the market.

Note that although the number of used jets sold is higher with buyback, these results do not enable us to conclude that all manufacturers are better off operating buyback schemes than they would be in the no-buyback counterfactual. These counterfactual simulations hold all prices fixed at the levels observed in the data. They are best thought of as measuring the shift in demand and supply curves at a given price level, not as equilibrium simulations. In particular, the large 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 in equilibrium. In Section 7, I simulate the movement of prices in equilibrium in order to compare firm revenues with and without buyback.

7 Buyback Incentive in Equilibrium

7.1 Computing Equilibrium Manufacturer Revenue

In order to measure the effect of buyback on manufacturer revenue in equilibrium, I need to 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.

For new units, I specify a simple model of manufacturer optimization. In particular, I assume that each *new* jet type j is priced independently at each year, t. Profit for jet j at

time t is given by:

$$\Pi_{jt} = M_t^F P_{0jt}(p_j - c_j) + \sum_k M_t^k P_{kjt}(p_j - c_j) \qquad (19)$$

Where M_t^F is the market size, or number of potential first time buyers in year t, and M_t^k is the number of owners of jet type k in year t. Manufacturers face marginal costs, c_j and choose price p_j according to the first order condition:

$$0 = \left(M_t^F P_{0jt} + \sum_k M_t^k P_{kjt} \right) + (p_j - c_j) \left(M_t^F \frac{\partial P_{0jt}}{\partial p_j} + \sum_k M_t^k \frac{\partial P_{kjt}}{\partial p_j} \right) \qquad (20)$$

Notice that I assume firms price each jet independently, without considering the impact of their pricing decision on the profits from the other jets they manufacture. Rearranging the first order condition allows me to back out marginal costs using a variation of the usual markup formula:

$$p_j + \frac{M_t^F P_{0jt} + \sum_k M_t^k P_{kjt}}{M_t^F \frac{\partial P_{0jt}}{\partial p_j} + \sum_k M_t^k \frac{\partial P_{kjt}}{\partial p_j}} = c_j \qquad (21)$$

Note that p_j and M_t^k are observed in the data. P_{0jt} and P_{kjt} are market shares generated by the demand model. They are functions of the estimated parameters and the observed prices and characteristics of all jet models available at date t.[7] The number of potential first time buyers, M_t^F, is not directly observed - I only observe a owners that make purchases. The average number of first time buyers per year in the data is 536.7, the minimum is 277, and the maximum is 935. For these simulations, I set $M_t^N = 1000$ for all years.

Under these assumptions, I back out implied values of c_j that can be used to simulate equilibrium pricing decisions under counterfactual parameter values.[8] The distribution of implied markups, defined as $\frac{p_j - c_j}{c_j}$, at the estimated parameters is illustrated in Figure 1. The mean markup is 48% of price. Given the implied marginal costs, a parameter vector, and a vector of used prices, I solve for the equilibrium new jet prices by iterating the pricing equations until I find a fixed point in the vector of new prices at each period t.

To compute equilibrium prices for used units, I assume that all used units are supplied by owners who choose to upgrade or exit the market given the prevailing prices. I simulate the demand model at a given parameter vector, vector of new jet prices, and vector of initial used jet prices p, and calculate the total number of used jets purchased, $D(p)$, and the total

---

[7]Note that the calculation of these market shares requires a number of additional assumptions discussed in Section A.3 of the Appendix.

[8]Throughout, I use the rental rates described in Section 4 in place of sticker prices. The marginal costs should therefore be intererpreted 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.

number of used jets sold, $S(p)$ over the 20 year simulation. Let the excess demand ratio be $ED(p) = \frac{D(p)}{S(p)}$. If $ED(p) > 1$, I replace p with λp where $\lambda > 1$ is a scalar, and repeat the simulation. If $ED(p) < 1$, I choose $0 < \lambda < 1$ and repeat the simulation at λp. I repeat this prodcedure until I find λ^* such that $ED(\lambda^* p) \approx 1$.[9] Equilibrium used prices are then given by $p^* = \lambda^* p$. Note that this algorithm finds prices under which used markets are in equilibrium *on average*. I scale all used prices up or down by the same factor, and I do not require that the used jet market clears for each model, only that the used market clears on average across models and years. Since there are over 4,900 model-year markets in the data, finding an exact used market equilibrium is computationally demanding. The results below will therefore focus on the average effect of buyback on used prices.

Starting with the observed used and new prices, it is possible to iterate on the two procedures described above - the new price equilibrium and used price equilibrium - to find equilibrium prices for any alternative parameter vector.

7.2   The Effect of Buyback on Equilibrium Revenue

I simulate manufacturer revenue at the estimated parameters and under the no-buyback assumption to illustrate the effect of buyback on manufacturer revenue in equilibrium. Note that whether or not revenue is lower in the no-buyback counterfactual is an empirical question. The observation that all manufacturers operate buyback in reality 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. On the other hand it could be that the equilibrium in which all manufacturers operate buyback is Pareto optimal, and that at some or all manufacturers are worse off in the no-buyback counterfactual.

The first three columns of Table 10 record total revenue for each of the major manufacturers in three different equilibrium simulations. The first column takes the estimated parameters and prices as given, and simulates demand for the 20 year sample period. The second column simulates a no-buyback counterfactual which does not allow for used prices to adjust. I set $b_m = 0$ for all manufacturers m, simulate equilibrium new jet prices holding used jet prices fixed, and then simulate demand at these new jet prices. As illustrated in Section 6, removing buyback significantly lowers both the demand for new jets and the supply of used jets, but fixing used prices does not allow lower used jet supply to feed through into higher used

---

[9]In practice, I set $ED(p^*) \approx 1.155$ in equilibrium. This is the excess demand ratio when the model is simulated at the observed parameters and observed prices. Demand does not equal supply exactly because not all jet owners are included in the estimation sample, and consumers are only allowed to hold one jet at a time in the model.

unit prices. The difference in manufacturer revenues between the first and second column therefore captures the lower demand for new jets in the no-buyback counterfactual, but not the countervailing force of reduced used jet supply, which would induce substitution towards new jets if used prices were allowed to rise. As expected, revenue is lower for all manufacturers in column 2. Removing buyback without allowing used prices to adjust reduces the total revenue of the major manufacturers by 9.8%. For Cessna, the brand for which buyback is estimated to have the largest effect on transaction costs, revenue falls by 20%.

The third column of Table 10 simulates a no-buyback counterfactual equilibrium in which both new and used prices are allowed to adjust. In particular, I set $b_m = 0$ for all manufacturers m, and iterate the used price and new price equilibrium simulations until prices converge. In this counterfactual, the price of used units rises by 7.0% relative to the observed prices. Total manufacturer revenue is 2% higher in column 3 than in column 2. This is due to higher used jet prices inducing substitution towards new jets. To see how this works, consider a first time buyer that is on the margin between buying a new jet and a used jet. If buyback is removed, the first time buyer's demand is unaffected since, by assumption, they do not directly benefit from the buyback scheme. However, removing buyback decreases the supply of (and demand for) used jets, causing used jet prices to rise and encouraging the first time buyer to substitute towards the new jet. This mechanism increases revenue for each manufacturer. However, the effect of higher used prices on manufacturer revenue is not large enough to offset the fall in demand for used jets caused by removing buyback. Total manufacturer revenue in column 3 is still 8% lower than the revenue with buyback in column 1.

These findings suggest that each manufacturer is better off when all manufacturers operate buyback than when all manufacturers do not operate buyback. This implies that the observed equilibrium, in which all manufacturers accept trade-ins, is Pareto optimal. In particular, manufacturers are not in a prisoner's dilemma in which it is an equilibrium for all manufacturers to offer buyback, but manufacturers would all be better off if they jointly agreed not to accept trade-ins. Note that this conclusion relies on the assumption that it is costless for manufacturers to buy back and resell used jets. In reality, buyback reduces transaction costs for consumers by transferring the burden of these costs to the manufacturer. A high cost of buyback to the manufacturer could reduce profits in the buyback equilibrium sufficiently to make the no-buyback equilibrium more profitable for manufacturers. A back of the envelope calculation indicates that the per-jet cost of buyback would have to be at least $5.6 million[10] for total manufacturer profit less the cost of buyback to be higher in

---

[10]The difference in profits (revenue less marginal costs in rental rate units) between the baseline and no buyback equilibria is $100.1 million. The total number of jets bought back in the mean baseline demand simulations is 355.61. The difference in profit per jet bought back is $281.6 thousand. Dividing this by 0.05 provides an approximate conversion of the units from rental rates to sticker prices.

the no-buyback equilibrium. Since the transaction costs faced by the consumer are only $1.9 million on average, it seems unlikely that the manufacturer's cost of reselling used jets would be this high.

When might revenue be higher under the no-buyback equilibrium? As discussed above, the net effect of removing buyback on equilibrium revenue depends on the relative sizes of the direct effect on new jet demand, which lowers revenue, and the effect on used prices, which raises revenue by inducing substitution away from used jets. Removing buyback can only increase manufacturer revenue if the latter effect dominates. Intuitively, the effect on used prices should be larger when there is a larger effect of buyback on used jet supply[11].

To illustrate this, I find alternative parameters at which the effect of buyback on supply is larger than it is at the estimated parameters. In particular, I set $\tau = \hat{\tau} + 4$ and $\beta^n = \hat{\beta}^n + 5$, where $\hat{x}$ is the estimated value of x, and otherwise use the estimated parameters. At these parameters, consumers rarely upgrade to used jets since there is a strong preference for new jets among upgraders. Thus, the substitution between used and new jets induced by buyback is smaller relative to the substitution between holding and upgrading.

I first simulate equilibrium prices at these parameters using the iterative method described above. I then simulate demand at these parameters and prices. As with the simulations reported in Section 6, removing buyback (setting $b_m = 0$ for all m) lowers simulated demand for new jets from upgraders significantly. Column 2 of Table 9 reports the decomposition of the change in new jets demanded from removing buyback for these alternative parameters. The results indicate that the change in the supply of used jets at fixed prices is 90.3% of the change in demand for used jets, as opposed to 86.5% at the estimated parameters. Note also that, as expected, the change in used jets demanded is very small, at less than 0.1% of the change in new jets demanded. At these alternative parameters, almost all of the increase in demand for new jets comes from owners upgrading more frequently, rather than substituting from upgrading to used jets or exiting the market.

Column 4 of Table 10 reports manufacturer revenue at the alternative parameters and prices. Column 5 reports manufacturer revenue from a no-buyback simulation in which new jet prices are allowed to change but used jet prices are held fixed. Column 6 reports revenue in a no-buyback simulation in which both new and used jet prices are allowed to adjust. The difference between columns 5 and 6 is due to the effect of removing buyback on used jet prices. Used jet prices rise by 15.8% and total manufacturer revenue is 9.5% higher when used jet prices are allowed to adjust. The larger effect of removing buyback on used jet supply translates into a larger increase in used jet price, compared to that reported in column 3 for the estimated parameters. This large change in price leads to

---

[11]Indeed, in appendix section A.4 I show that, in a simplified model of demand, buyback will always increase revenue for a monopolist if there is no effect on used jet supply.

greater substitution towards used jets - recall that the revenues reported in column 3 are only 2% higher than those in column 2.

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 consumer substitution 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.

I present an identification strategy that relies on assumptions about the structure of these policies and allows me to estimate the effect of buyback on demand without observing exogenous variation in policies over time or across manufacturers. I estimate transaction costs in the used jet market to be $1.9 million on average, or 25% of the mean sale price. Manufacturer buyback reduce the transaction costs faced by owners by between 7% and 18%. Removing buyback from all manufacturers would cause the demand for new jets from upgrading consumers to fall by 45%. I also find that removing buyback has a significant effect on used jet supply - 86.5% of the decrease in demand at fixed prices is explained by a decrease in the frequency with which consumers upgrade and supply their used units to the market.

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 of removing buyback on new unit demand lowers revenue. This is counteracted by the equilibrium effect on used jet prices, which encourages substitution towards new jets. I quantify both effects and find that the effect on demand for new jets outweighs the equilibrium effect of the increase in used prices. Removing buyback therefore lowers equilibrium revenue for all manufacturers at the estimated parameters. I demonstrate how the magnitude of these effects change when demand parameters change. In particular I show that the effect of buyback on used unit prices is greater when buyback has a larger effect on used jet supply relative to its effect on new jet demand.

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. Future research could examine more systematically what features of an industry make buyback or other manufacturer interventions in secondary markets optimal, and why such policies are

common in some durable goods markets and not others. For example, a richer model could describe how heterogeneity in consumer preference for newness affects the firm's incentive to engage with the secondary markets.

Tables and Figures

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.

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)
Price ($ Millions)	Mean	8.987	17.896	5.168	10.379	27.053
	SD	(6.931)	(10.894)	(1.871)	(5.592)	(7.787)
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

Manufacturer	New Market Share 1980 - 2000			Used Market	
	Small	Medium	Large	Resale Ratio	Annual Used Sales
Cessna	59%	21%	0%	23.9%	233.14
Bombardier	22%	15%	33%	24.0%	133.32
Raytheon	14%	29%	0%	29.6%	131.00
Dassault	2%	18%	22%	22.4%	72.83
Gulfstream	0%	0%	42%	19.4%	118.00

Notes: Columns 1-3 record the market share of the top 5 manufacturer in new jet sales between 1980 and 2000 in each jet category. Column 4 records the average resale ratio between 1980 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 1980 and 2000.

Table 4: Used Jets Sold to Manufacturers

	All Jets	Cessna	Gulfstream	Bombardier	Dassault	Raytheon
Upgrades to New Jets	0.416	0.435	0.493	0.495	0.220	0.541
Upgrades to Used Jets	0.051	0.030	0.093	0.058	0.051	0.075
All Sales	0.060	0.054	0.115	0.069	0.044	0.076

Notes: This table reports the share of jets sold by corporations between 1980 and 2000 that are bought by manufacturers or dealers affiliated with manufacturers. Row 1 restricts to sales that take place at most 3 months before the corporation selling the jet buys a new jet. Row 2 restricts to sales that take place at most 3 months before the corporation selling the jet buys a used jet. Row 3 records these shares among all sales of jets by corporations. Columns 2 to 6 report these shares among all sales of each of the 5 major brands of jets.

Table 5: Own Brand Bought Back

Manufacturer	Own Brand Share of Buybacks
Raytheon	90%
Bombardier	78%
Cessna	97%
Dassault	87%
Gulfstream	85%

Notes: Table records the share of all jets purchased by a manufacturer that are of that manufacturer's own brand. Sample includes all sales to manufacturers between 1980 and 2000. Includes sales of jets to dealers that appear to be affiliated with manufacturers, as described in the text.

Table 6: Purchase Shares: New and Used

	(A) All Owners			(B) Excluding Buyback Users		
	Used	New	Diff	Used	New	Diff
First Jet	0.730	0.270	0.460	0.704	0.296	0.408
Upgrade	0.654	0.346	0.308	0.696	0.304	0.392
Diff	0.076	-0.076	0.152	0.008	-0.008	0.016
SE	(0.008)	(0.008)		(0.009)	(0.009)	

	(C) All Owners			(D) Excluding Buyback Users		
	Used	New	Diff	Used	New	Diff
Same Brand	0.631	0.369	0.262	0.692	0.308	0.384
Different Brand	0.726	0.274	0.452	0.732	0.268	0.464
Diff	-0.095	0.095	-0.190	-0.040	0.040	-0.080
SE	(0.039)	(0.039)		(0.045)	(0.045)	

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. Satndard errors reported are the stanrd error of the difference in means for each column.The sample used includes all purchases between 1980 and 2000.

Table 7: Parameter Estimates

Parameter	$\xi_{jt}=0$	$\xi_{jt}=\hat{\xi}_{jt}$	Buyback Parameter b_m	$\xi_{jt}=0$	$\xi_{jt}=\hat{\xi}_{jt}$
α^{age}	-0.041	-0.074	Cessna	1.348	1.506
	(0.002)	(0.004)		(0.117)	(.129)
α^n	1.156	1.501	Bombardier	0.683	0.713
	(0.031)	(0.052)		(0.152)	(0.154)
τ	8.365	8.316	Dassault	0.728	0.543
	(0.033)	(0.032)		(0.213)	(0.215)
α^p	-13.042	-12.642	Gulfstream	1.542	1.055
	(0.029)	(0.039)		(0.207)	(0.229)
σ^p	1.073	0.764	Raytheon	1.133	1.164
	(0.020)	(0.028)		(0.165)	(0.157)
β^n	-0.435	-0.505	Other	-0.440	-0.365
	(0.087)	(0.083)		(0.347)	(0.329)
β^{sb}	1.257	1.250			
	(0.048)	(0.043)			

Notes: Table reports estimated parameters and standard errors for the demand model. Parameters are reported for the simulated ML estimation that ignores unobserved jet characteristics in the columns labeled $\xi_{jt}=0$. Parameters are reported for the control function estimation described in the text in the columns labeled $\xi_{jt}=\hat{\xi}_{jt}$. Standard errors for these estimates are computed using 100 bootstrap replications of the two-step procedure. For each replication, a bootstrap sample over jet-years is used to estimate the first stage control function, and then a bootstrap sample over owners, i, is used to perform the second stage simulated ML estimation. The first stage estimates of the control function are reported in Appendix Table A.2. The first stage sample includes all jet-years for all jets available from 1980-1999. The second stage sample used includes all corporations who buy their first jet after 1979, as described in the text.

Table 8: Demand Simulations

|	Data	Baseline	No Buyback
New Share of First Time Purchases	0.176	0.169	0.169
New Share of Replacement Purchases	0.242	0.248	0.152
New Share Among Same Brand Replacements	0.337	0.345	0.151
New Share Among Different Brand Replacements	0.149	0.150	0.153
Number of New Jets Purchased as Replacements	556	583	318
Mean Number of Replacements per Owner	0.214	0.219	0.195

Notes: Column 1 of this table records the value of each statistic in the data used for estimation, which includes all corporations who buy their first jet after 1979 and forces each corporation to hold one jet at a time, as described in the text. Column 2 records the mean values of these statistics across 100 simulations of the demand model at the estimated parameters. The simulations take the date at which each manufacturer bought their first jet as given, and simulate demand for the years 1980-1999. Column 3 reports the mean values of these statistics across 100 simulations of the demand model with $b_m = 0$ for all m, and otherwise using the estimated parameters.

Table 9: Decomposition of Demand and Supply Changes

		Parameters	
		Estimated	Alternative
$\Delta Used\ Jets\ Supplied$	Level	229.0	521.1
	Share	86.5%	90.3%
$\Delta New\ Jets\ Demanded$	Level	264.6	577.1
	Share	100%	100%
$\Delta Used\ Jets\ Demanded$	Level	-12.2	-0.8
	Share	4.6%	0.0%
$\Delta Market\ Exits$	Level	-23.4	-55.2
	Share	8.8%	9.6%

Notes: Column 1 records the difference in the number of used jets supplied, new jets demanded, used jets demanded, and market exits between the baseline simulation and the no-buyback simulaion at the estimated parameters. The differences are reported as the baseline level minus the no-buyback level, and in absolute value as a percentage of the change in new jets demanded. New and used jets demanded are the number of new and used jets sold in the simulations. Used jets supplied is the number of sales of used jets by owners (when they upgrade or exit the market). Market exits is the number of times owners sell their ejt and do not upgrade. The second column records these differences at the alternative parameter simulations described in Section 7 of the text.

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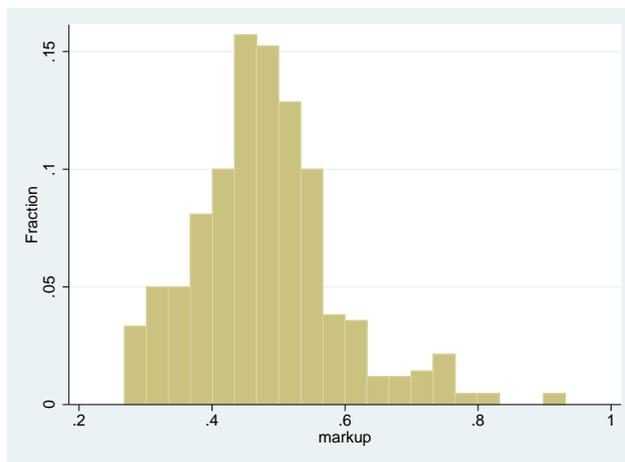
Notes: This figure is a histogram of the implied markups $\frac{p_j - c_j}{c_j}$, where p_j is the observed price and c_j is the implied marginal cost, among all new jets available in the estimation sample, where jets are defined by a manufacturer-year-segment. The marginal costs are backed out from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

Table 10: Equilibrium Revenue

Parameters	Estimated Parameters			Alternative Parameters		
Simualtion	Baseline	No Buyback		Baseline	No Buyback	
Used Prices		Fixed	Adjusted		Fixed	Adjusted
Total Revenue	2481.0	2238.1	2284.1	3071.5	2530.7	2770.9
Bombardier	586.8	543.9	549.2	691.3	610.0	660.8
Cessna	576.8	485.5	507.2	714.1	512.0	575.6
Dassault	463.3	439.4	446.8	589.1	527.7	577.2
Gulfstream	465.7	417.6	418.4	591.3	494.7	531.5
Raytheon	388.6	351.7	362.5	485.6	386.3	425.8
% Change in Used Prices			7.0%			15.8%

Notes: This table records manufacturer revenue for each of the major manufacturers under various simulations. The first three columns record revenue under simulations at the estimated parameters. The first column records revenue when demand is simulated at the estimated parameters and observed prices. The second column records revenue when $b_m = 0$ for all m, and new jet prices are in equilibrium as described in the text. The third column repeats the no-buyback simulation of column 2, but allows used jet prices to adjust as described in the text. The bottom row of the table records the percentage change in used prices between the baseline simulation and the simulation in column 3. Columns 4 - 6 report revenue for equivalent simulations at the alternative parameters (and alternative baseline prices) described in the text. Revenues in this table are in millions of year 2000 dollars, and are computed at rental rates, not sticker prices.

References

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

Fudenberg, D., and J. Tirole. 1998. "Upgrades, Tradeins, and Buybacks." *RAND Journal of Economics* 29 (2): 235-258

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

Hendel, I., and A. Lizzeri. 1999. "Interfering with Secondary Markets." *RAND Journal of Economics* 30 (1): 1-27

Petrin, A., and K. Train. 2009. "A Control Function Approach to Endogeneity in Consumer Choice Models." *Journal of Marketing Research* 47 (1): 3-13

Porter, R. H., and P. Sattler. 1999. "Patterns of Trade in the Market for Used Durables: Theory and Evidence." *NBER Working Paper*

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.

To test whether this procedure has a large effect on the estimated parameters, I run the estimation separately for $t \geq 1995$, for which only 10.0% of the prices of available jets are missing, and only 3.8% of purchases are of a jet with a missing price. 75% of the missing model-year price observations are for years before 1994. I estimate the demand model on this subset of the data uisng the control function method. The parameter estimates are presented in Table A.3. The estimated parameters are very close to those reported in columns 2 and 4 of Table 7. This suggests that much of the identifying variation is in the later part of the panel, when missing prices are not as much of a problem.

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*. 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. This procedure generates a panel of jet owners observed once a year, holding at most one jet each year.

A.3 Equilibrium Simulation Details

The equilibrium price simulations described in Section 7 use the following manufacturer first order condition:, which is described in the text

$$p_j + \frac{M_t^F P_{0jt} + \sum_k M_t^k P_{kjt}}{M_t^F \frac{\partial P_{0jt}}{\partial p_j} + \sum_k M_t^k \frac{\partial P_{kjt}}{\partial p_j}} = c_j \tag{22}$$

The markup term depends on the market share for jet type j among first time buyers, P_{0jt}, and the market share for jet type j among owners of jet type k, P_{kjt}. These market shares depaned on the prices and characteristics of all available jets, and are generated by the demand model. However, I must make two simplifying assumptions in order to generate these aggregate market shares in practice.

First, I let $P_{kjt} = \int P_{kjt}^i \phi(\nu_i) d(\nu_i)$, where $\phi(\cdot)$ is the density function of a standard normal distribution. Recall that the individual specific price coefficient, α_i^p, is given by $\alpha_i^p = exp(\alpha^p + \sigma^p \nu_i)$, and $\nu_i \sim N(0,1)$. Note however, that, the distribution of ν_i *conditional* on owning a jet of type k at date t need not be standard normal. For example, if jet k is a more expensive jet, then owners of that jet are likely to have low values of ν_i. The assumption used in the pricing equation that market shares are generated using the standard normal distribution is therefore equivalent to saying that the firm believes that ν_i is distributed standard normal among all groups of upgraders, regardless of what jet they own at date t, and does not correctly incorporate the heterogeneity in these distributions among different groups of potential upgraders.

Second, recall that the estimated model of demand is a model of consumer choice *conditional* on making a first time purchase. That is, the parameters describe the distribution of preferences among consumers who make at least one purchase in their lifetime. However, the firm's first order condition depends on P_{0jt}, which is the share of *potential* first time buyers who choose to buy a jet j. In order to estimate this share, and the marginal effect $\frac{\partial P_{0jt}}{\partial p_j}$, I write $P_{0jt} = P_t^{IN} \tilde{P}_{0jt}$. $\tilde{P}_{0jt}$ is the probability of choosing jet j conditional on buying *some* jet, implied by the model, $\tilde{P}_{0jt} = \int P_{0jt}^i \phi(\nu_i) d(\nu_i)$. P_t^{IN} is the probability of a potential first time buyer buying *some* jet at date t. I assume that each of the M_t^N potential buyers choose to buy some jet with probability $\frac{exp(IN_t)}{exp(\xi_t) + exp(IN_t)}$ where IN_t is the expected logit inclusive value of the jets available at date t and ξ_t is a date fixed effect. This model implies that potential consumers do not know their value of α_i^p or ϵ_{ijt} before

choosing to enter the market. I assume that $M_t^N = 1000$ for all years and calibrate ξ_t so that $1000 * \frac{exp(IN_t)}{exp(\xi_t)+exp(IN_t)}$ is equal to the observed number of first time buyers each period.

A.4 The Effect of Buyback on Equilibrium Revenue: Theory

To see why buyback can only increase revenue if there is an effect on used jet supply, consider the following a simplified model of equilibrium. A unit mass of consumers demands two goods: U and N. Let p_U be the price of good U, p_N be the price of good N, and b be the buyback parameter. Demand for good U is given by $P_U(p_U, p_N, b)$, demand for good N is given by $P_N(p_N, p_U, b)$, and demand for an outside good is $P_0(p_U, p_N, b)$. Suppose that $\frac{\partial P_i}{\partial p_i} < 0$, $\frac{\partial P_i}{\partial p_j} > 0$, and $\frac{\partial P_0}{\partial p_j} > 0$ for $i, j \in \{U, N\}$, $i \neq j$. That is, demand for each good is decreasing in own price and increasing the price of the other goods. The supply of good N is given by $S_N(p_N)$ with $\frac{\partial S_N}{\partial p_N} > 0$, and the supply of good U is given by $S_U(p_U, p_N, b)$ with $\frac{\partial S_U}{\partial p_U} > 0$ and $\frac{\partial S_U}{\partial p_N} < 0$. Prices are set by equating supply and demand of each good.

Finally, suppose that $\frac{\partial P_U}{\partial b} < 0$, $\frac{\partial P_0}{\partial b} < 0$, $\frac{\partial P_N}{\partial b} > 0$, and $\frac{\partial S_U}{\partial b} \geq 0$. That is, increasing the buyback parameter, b, reduces demand for used units and for the outside good, and increases demand for new units and the supply of used units. This simple model captures the important features of the econometric model presented in Section 4: used and new jets are substitutes, the supply of used jets depends on both the used and new jet prices, and buyback encourages substitution towards new jets and increases the supply of used jets.

Proposition 1. Let $R_N = S_N p_N$ be equilibrium manufacturer revenue. If $\frac{\partial S_U}{\partial b} = 0$, then $\frac{\partial R_N}{\partial b} \geq 0$.

Proof. Consider equilibrium under two values of the buyback parameter $b' > b$. Equilibrium prices are (p_N', p_U') and (p_N, p_U). Assume $R_N(b') < R_N(b)$. Since the supply of good N as a function of p_N, $S_N(\cdot)$ does not depend on b, it must be that $p_N' < p_N$, and therefore $P_N(p_N', p_U', b') < P_N(p_N, p_U, b)$, which implies $P_N(p_N', p_U', b') < P_N(p_N', p_U, b')$. This means that $p_U' < p_U$. For this to be the case it must be that $P_U(p_U', p_N', b') < P_U(p_U, p_N, b)$. Since $b' > b$, $p_N' < p_N$, and $p_U' < p_U$, it must be that $P_0(p_U', p_N', b') < P_0(p_U, p_N, b)$. This means that:

$$P_N(p_N', p_U', b') + P_U(p_U', p_N', b') + P_0(p_U', p_N', b') < P_N(p_N, p_U, b) + P_U(p_U, p_N, b) + P_0(p_U, p_N, b) \tag{23}$$

Which contradicts the assumption of a constant unit mass of consumers. It must therefore be that $R_N(b') \geq R_N(b)$. □

The intuition for the proof is that for revenue to fall, it must be that demand for 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, which requires demand for the used good to fall as well. But it can not be an equilibrium for demand for both the new and used good to be lower when buyback is increased, since lower prices and a higher buyback parameter also decrease the demand for the outside good.

A.5 Tables and Figures

Table A.1: Sales to Manufacturers as a Share of Corporate Sales: 5-Year Windows

	1980-1984	1985-1989	1990-1994	1995-1999
Sales Before New Purchase	0.159	0.256	0.565	0.426
Sales Before Used Purchase	0.082	0.032	0.077	0.039
All Sales	0.037	0.039	0.074	0.070

Notes: This table reports the share of jets sold by corporations that are bought by manufacturers or dealers affiliated with manufacturers for four periods between 1980 and 2000. Row 1 restricts to sales that take place at most 3 months before the corporation selling the jet buys a new jet. Row 2 restricts to sales that take place at most 3 months before the corporation selling the jet buys a used jet. Row 3 records these shares among all sales of jets by corporations.

Table A.2: First Stage Regression for Control Function Estimator

Dependent Variable: Price (in \$) of Jet Model j	
Number of Bombardier Models Available	156899.4
	(11568.2)
Number of Cessna Models Available	-77775.1
	(21948.6)
Number of Dassault Models Available	8414.6
	(6254.5)
Number of Gulfstream Models Available	-529646
	(79202.3)
Number of Raytheon Models Available	135412.9
	(14979.9)
Number of Other Models Available	-30035.0
	(15058.1)
Number of Model j Jets Owned	-2785.8
	(364.5)
R^2	0.825
N	4307

Notes: An observation is the price of a specific jet model (manufacturer-segment-year of manufacture) in a specific year. Control variables include manufacturer, jet size, and year of manufacture fixed effects, controls for jet age and whether the jet is new, as well as the jet characteristics included in the utility specification. The regressors defined by "models available" count all models on sale or owned in that year. I assume that consumers are not

Table A.3: Parameter Estimates on $t \geq 1995$

Parameter	Estimate	Parameter	Estimate	Parameter	Estimate	Parameter	Estimate
α^{age}	-0.097	σ^p	0.647	b Cessna	1.363	b Raytheon	1.291
	(0.007)		(0.037)		(0.165)		(0.217)
α^n	1.774	β^n	-0.512	b Bombardier	0.765	b Other	-0.760
	(0.070)		(0.127)		(0.205)		(0.722)
τ	8.538	β^{sb}	1.319	b Dassault	0.669		
	(0.049)		(0.071)		(0.296)		
α^p	-12.685			b Gulftstream	1.32		
	(0.055)				(0.287)		

Notes: Table reports estimated parameters and standard errors for the control function estimate of the demand model. Estimation sample restricted to choices in the years 1995-1999. Sample is otherwise identical to the main estimation sample. In particular, the sample used in this table includes the post-1995 decisions of corporations who bought their first jet before 1995. Standard errors are not adjusted for the two step procedure.