

Switching Costs in the US Seed Industry: Technology Adoption and Welfare Impacts

Jinjing Luo, GianCarlo Moschini, Edward D. Perry

Working Paper 22-WP 639

December 2022

**Center for Agricultural and Rural Development
Iowa State University
Ames, Iowa 50011-1070
www.card.iastate.edu**

Jinjing Luo is Model Validation Senior Analyst, Citibank.

GianCarlo Moschini is Professor, Department of Economics, Iowa State University, Ames, Iowa 50011. E-mail: moschini@iastate.edu.

Edward D. Perry is Associate Professor, Department of Agricultural Economics, Kansas State University, Manhattan, Kansas 66506. E-mail: edperry@ksu.edu.

This publication is available online on the CARD website: www.card.iastate.edu. Permission is granted to reproduce this information with appropriate attribution to the author and the Center for Agricultural and Rural Development, Iowa State University, Ames, Iowa 50011-1070.

For questions or comments about the contents of this paper, please contact GianCarlo Moschini, moschini@iastate.edu.

Iowa State University does not discriminate on the basis of race, color, age, ethnicity, religion, national origin, pregnancy, sexual orientation, gender identity, genetic information, sex, marital status, disability, or status as a U.S. veteran. Inquiries regarding non-discrimination policies may be directed to Office of Equal Opportunity, 3410 Beardshear Hall, 515 Morrill Road, Ames, Iowa 50011, Tel. (515) 294-7612, Hotline: (515) 294-1222, email eooffice@iastate.edu.

Switching Costs in the US Seed Industry: Technology Adoption and Welfare Impacts

Jinjing Luo
Citibank

GianCarlo Moschini
Iowa State University

Edward D. Perry
Kansas State University

This version: 30 November 2022

ABSTRACT

We evaluate the role of brand and technology switching costs in the US soybean seed industry using a unique dataset of actual seed purchases by about 28,000 farmers from 1996 to 2016. Using a random coefficients logit model of demand, we estimate brand and technology switching costs, characterize the distributions of buyers' willingness to pay for seed brands and the glyphosate tolerance (GT) trait, and assess the implications of brand and technology switching costs for farmers' welfare, technology adoption, firm profits, and firm market shares. We find that farmers are willing to pay large premiums for brand labels, and even larger premiums for the GT trait, although there is considerable heterogeneity in these values. Switching costs play an important role in the soybean seed industry. Eliminating these costs would significantly increase buyers' welfare, reduce seed prices and firm profits, decrease adoption of the GT trait, and impact industry consolidation by expanding smaller firms' market shares.

Key Words. Brand inertia, Heterogeneity, Input demand, Mixed logit, State dependence, Switching Costs, Technology adoption.

JEL Codes. L13, L69, O13, Q12

Acknowledgements. This research was supported in part by the National Institute of Food and Agriculture, US Department of Agriculture, grant No. 2018-67023-27682.

1. Introduction

There is a long-standing recognition that buyers often prefer products they have experienced in the past, relative to products with similar or even identical characteristics. Beyond the role of idiosyncratic preferences, the resulting inertia in observed consumers' demand may reflect switching costs associated with moving away from past consumption choices, where switching costs can be broadly construed to include learning costs, contractual or transaction costs, and also psychological elements. A key challenge in this setting is to distinguish true *state dependence*—defined as the causal dependency of an individual's future choices on their current state (Heckman 1981a)—from the effects of unobserved preference heterogeneity (spurious state dependence). When present, true state dependence directly alters the returns to different choices and thus influences market equilibrium.

Understanding the nature and extent of switching costs has important implications for imperfectly competitive markets. By altering the responsiveness of consumer demand, switching costs directly affect equilibrium prices. Because these frictions make demand more inelastic, firms can exploit their loyal customer base by charging higher prices. Dynamically, however, switching costs also provide the opportunity for customer base expansion through aggressive pricing. How such contrasting harvesting and investment incentives actually play out remains an open question (Farrell and Klemperer 2007), and is ultimately an empirical matter that depends on the specific attributes of a given market.

In this paper, we study the nature and extent of switching costs in a novel context—soybean seed demand over a period characterized by the advent of genetically engineered (GE) varieties. This sweeping new technology led to early acquisitions and consolidation in the industry, as established brands had to adapt to the innovator's (Monsanto) entry and the tremendous success of GE soybeans with farmers. As such, it provides a useful context to empirically assess the extent to which inertia due to switching costs matters with respect to both brand and technology.

Our study adds to an extensive literature on demand inertia in both economics and marketing. Following Keane (1997), a large body of research, focused on the consumer-packaged goods (CPG) industry, has shown that the presence of state dependence can have important implications for the extent of market power and pricing behavior (Dubé, Hitsch, and Rossi 2010), market structure (Dubé, Hitsch, and Rossi 2009), the price effects of mergers (MacKay and Remer 2022),

the pro-competitive role of advertising (Shum 2004), and the persistence of brand shares (Bronnenberg, Dhar, and Dubé 2009; Bronnenberg, Dubé, and Gentzkow 2012; Grzybowski and Nicolle 2021). Among non-CPG applications, Sudhir and Yang (2014), and Train and Winston (2007), study the automobile industry; Handel (2013) and Yeo and Miller (2018) analyze the health insurance industry; Honka (2014) investigates automobile insurance; and, Barone, Felici, and Pagnini (2011) study local credit markets.

Unlike much of related existing work, which has privileged consumer demand, in this paper we study state dependence for competitive producers' (farmers') choices. Furthermore, the industry and period we study is characterized by a major innovation, which permits us to investigate both brand and technology inertia, as well as their interplay.

First introduced in 1996, GE varieties provided farmers with drastically new technological solutions for weed and pest management, which led to rapid widespread adoption (Moschini 2008). Among the major US crops, the soybean seed industry has perhaps undergone the largest transformation. The once-common farming practices of saving harvested soybeans for seed use, and/or purchasing publicly developed varieties, have been replaced by the almost complete reliance on new proprietary commercial soybean varieties that embed the GE trait for glyphosate tolerance (GT).¹ The GT trait was developed by Monsanto, who, as an outside proprietor, faced major challenges in converting their technology into a commercial product. Successful seed varieties required access to two complementary building blocks: the new GE traits, and elite germplasm in existing varieties (Graff, Rausser, and Small 2003). Monsanto made some early key seed brand acquisitions, providing them with direct access to farmers (Fernandez-Cornejo 2004). Furthermore, they also aggressively licensed the GT trait to other seed suppliers.

To characterize the pattern of farmers' seed choices and technology adoption, we estimate a random coefficients logit model of US soybean seed demand. The model characterizes the demand for seed brands and the GE trait, and allows for the presence of both state dependence and unobserved heterogeneity in farmers' preferences for brand labels and for the GT technology trait. For estimation, we draw on a unique dataset containing more than 200,000 seed purchase decisions by roughly 28,000 US soybean farmers from 1996 to 2016. These (unbalanced panel) data provide

¹ In 1970 about 70% of planted soybeans were public varieties (Fernandez-Cornejo 2004). Based on the data used in this paper, by 2016 this fraction was less than 1%.

information on seed purchase histories, seed characteristics, and prices. In developing and estimating the model, our main objectives are to: (a) identify the extent to which seed demand is affected by state dependence for brand labels and the GT trait; (b) characterize the importance of farmer heterogeneity for seed demand; (c) quantify farmers' willingness to pay (WTP) distributions for different brands, the GT trait, and switching costs (state dependence); and, (d) assess the consequences of switching costs, through empirical counterfactual scenarios, for farmers' welfare, GT trait adoption, and for the leading brands' market shares.

Overall, we find that farmers are willing to pay large premiums for brand labels and for the GE trait, and these WTP estimates can vary widely across farmers. For example, from 2010 to 2016, farmers' mean WTP for the GT trait was \$21.8/acre, with 10% of farmers valuing GT at \$40.2/acre or more, and another 10% of farmers valuing it at \$3.5/acre or less. We also find significant evidence of structural state dependence, even after controlling for persistent unobserved farm-level heterogeneity. The mean marginal switching cost associated with brands is estimated at \$13.1/acre; the corresponding estimates for trait switching costs are \$6.8/acre (conventional to GT products) and \$14.7/acre (GT to conventional products). We also uncover considerable heterogeneity across buyers for these WTP estimates, and for farmers' valuations of the various brands.

Using the model estimates and a differentiated product model of the supply side, we evaluate counterfactual scenarios that shed additional light on the implications of inertia for the industry. Removing all switching costs would decrease seed prices by 6%, on average, and increase buyers' welfare by \$7.7/acre/year, on average, with most of this value arising from the removal of brand switching costs. At an aggregate level, farmers' surplus would increase by \$463.5 million per year, whereas firm profits would decrease by \$225.9 million per year. Furthermore, without switching costs, the adoption of the GT trait would have been higher in the first few years, and lower thereafter, with the trend away from GT products more pronounced in recent years. The counterfactual simulations also provide evidence that inertia, in a context characterized by rapid innovation such as the US seed industry, can affect market shares in a substantial way.

The analysis and results of this paper contribute to the literature in several ways. First, we provide evidence on the importance of switching costs for competitive firms' choices over a key input. Most existing studies focus on the CPG industry and largely ignore the extent and implications of switching costs in a production context. Second, we contribute to the empirical literature on the

implications of state dependence, especially when combined with the introduction of a major technological innovation. Finally, we provide new evidence that the sizeable WTP estimates for GE seed trait innovation are characterized by considerable heterogeneity across farmers.

The rest of this paper is organized as follows. Section 2 provides background information on the US soybean seed market. Section 3 presents the data used in the econometric regression. Section 4 develops the demand model, discusses the identification strategy, and presents the estimation strategy. Section 5 presents the demand estimation results, and the computation of WTP distributions and demand elasticities. Section 6 uses counterfactual simulations to assess some empirical implications of state dependence and switching costs. Section 7 concludes.

2. Industry Background

The US seed industry has grown considerably over the last few decades.² Technological innovation, from expanded research and development (R&D) investments, has been a critical factor. A major development affecting seed markets was the introduction of GE traits in the mid-1990s. By using breakthrough recombinant DNA techniques of modern biology, it became possible to integrate certain foreign genes (from bacteria) into the germplasm of elite crop varieties. These genes confer traits to the resulting “transgenic” crops, such as herbicide tolerance and insect resistance, which are highly valued by growers (Barrows, Sexton, and Zilberman 2014).

The GE revolution in the seed industry also benefited from the general strengthening of intellectual property rights for biological innovations (Moschini 2010). This is particularly important for soybean seeds, the focus of this paper. Soybean varieties are self-pollinating, meaning that they reproduce true to type (unlike hybrid maize, for example). Thus, prior to the advent of GE varieties, farmers could rely on saved seeds (from the previous harvest) and could access varieties developed and released by public institutions (state universities), essentially at competitive prices (production costs). The introduction of patented GE traits, and the associated increased use of trade secrets and contracts, effectively permitted the industry to develop proprietary seed products (Clancy and Moschini 2017). This greatly increased the profitability of R&D in plant breeding, which led to new investments and a wave of industry consolidation (Fernandez-Cornejo 2004). By the year

² The size of the global commercial seed market in 2014 was estimated at about \$12 billion in the United States (ISAAA 2016), and around \$52 billion worldwide (Syngenta 2016).

2000, the two largest firms (Monsanto and DuPont) accounted for about 40% of the US soybean seed market, a combined share that has risen to about 60% in recent years.

Soybeans are the second-most planted crop in the United States. Unlike GE corn varieties, which can have several traits, GT is the only trait with major commercial relevance during our study period (1996–2016).³ Glyphosate, originally marketed under the trade name Roundup, is a powerful, broad-spectrum herbicide that can be used in combination with GT crops. It can kill approximately 99% of weeds without harming GT varieties (Wechsler, McFadden, and Smith 2018). By reducing the need for tillage, as well as multiple types of herbicides, GT varieties permit an extremely effective (and simplified) weed control strategy (Perry, Moschini, and Hennessy 2016). Because of this, GT soybeans were rapidly adopted, despite commanding a significant price premium (OECD 2018): first commercially introduced in 1996, GT varieties accounted for more than 50% of the market by 1999, and more than 90% by 2007. Previous research has found that US farmers' WTP for the GT trait far exceeds its cost, resulting in significant net economic gains (Ciliberto, Moschini, and Perry 2019).

The marketing of seed varieties relies heavily on long-standing and well-known brands, where multiple brands can be marketed by the same parent company. Over the period of analysis, DuPont has primarily sold varieties under the Pioneer brand, whereas Monsanto has marketed varieties under several brands such as Asgrow, DeKalb, and Channel. Each brand typically offers multiple distinct varieties that differ in characteristics such as herbicide tolerance, soybean cyst nematode resistance, and relative maturity. Most brands, over the period of study, offered both conventional and GT varieties.⁴

3. Data

The data used in this paper pertain to seed purchases by a large and representative sample of US soybean farmers. These data come from a proprietary dataset assembled by Kynetec USA, Inc., a market research company. The data span the 21-year period from 1996 to 2016. For each year, the

³ GE varieties tolerant to glufosinate did not achieve commercial relevance until very recently and, as in Ciliberto, Moschini, and Perry (2019), we do not distinguish between conventional and glufosinate tolerant varieties in our empirical analysis.

⁴ Two exceptions are Channel, a Monsanto brand, which only offers GT varieties, and public providers (e.g., state university programs) who offer only conventional varieties.

seed purchases of about 3,500 soybean farmers are recorded. The sample itself is constructed to be representative at the crop reporting district (CRD) level.⁵ Each soybean farmer in the sample is observed to make one or more seed purchases and the data contains detailed information on the nature of the purchase (e.g., variety name, brand, parent company, GE traits, price, quantity of seed, and acres planted). The raw data contain 213,062 distinct seed choices. After dropping choices pertaining to non-commercial transactions (e.g., saved soybeans), we are left with 204,607 seed purchases that permit a comprehensive characterization of the industry.⁶ The data at hand is an unbalanced panel data set, although a large portion of farmers are observed over multiple years (and individual farmers are often observed making multiple purchases in the same year).

3.1. Products

The finest possible product definition, in our setting, would be in terms of individual varieties. Analysis at the variety level, however, is not feasible. There are simply too many varieties—over the period 1996–2016, our dataset includes 18,420 distinct soybean varieties. Furthermore, individual varieties have limited geographic presence, as each is bred to be suited to specific agro-climatic conditions (e.g., latitude). In addition, new varieties are introduced every year, and the life cycle of any given variety is relatively short (four to five years, on average).

The value of modern soybean seed varieties primarily derives from two complementary sources: germplasm (i.e., the underlying genetics accumulated from past generations of selective breeding), and GE traits. Concerning germplasm, it is useful to think of varieties as forming “product lines” over time, as companies introduce improved new varieties that are built on the genetics of previous varieties. We presume that this continuity is captured by the “brand” (e.g., Asgrow). Varieties marketed by any one brand at different locations may differ, even considerably, but in any one local market one can expect varieties of the same brand to share common characteristics. Hence, we define a “product” as consisting of unique combination of the brand and trait. Specifically, we consider 13 distinct brands, listed in Table 1 (we treat public/university seeds as a single brand, named Public). These 13 brands account for about 70% of the US soybean seed market over the

⁵ CRDs are regions identified by National Agricultural Statistics Service of the US Department of Agriculture (USDA). Each US state contains multiple CRDs (for example, Iowa has 9), and each CRD includes multiple counties.

⁶ Table A1 in the Appendix provides more details on the data-cleaning steps that were performed.

period analyzed. For tractability reasons, all remaining varieties are aggregated into an “Others” group. Given the foregoing, each of these brands can be associated with two products, depending on whether or not the GT trait is embedded. Because the brand Channel only provides GT seeds, and Public varieties only consist of conventional seeds, there are 26 distinct products.

Table 1. Brand and GT Trait Market Shares

| Market share | 1996–98 | 1999–2002 | 2003–09 | 2010–16 |
|---------------------------------|---------|-----------|---------|---------|
| <i>Monsanto</i> | | | | |
| Asgrow | 0.120 | 0.171 | 0.154 | 0.212 |
| Channel | 0.000 | 0.000 | 0.000 | 0.033 |
| DeKalb | 0.094 | 0.056 | 0.048 | 0.004 |
| Kruger | 0.015 | 0.021 | 0.017 | 0.009 |
| <i>DuPont</i> | | | | |
| Pioneer | 0.191 | 0.195 | 0.242 | 0.273 |
| <i>Syngenta</i> | | | | |
| Golden | 0.037 | 0.061 | 0.045 | 0.002 |
| NK | 0.050 | 0.046 | 0.077 | 0.089 |
| <i>Dow</i> | | | | |
| Mycogen | 0.028 | 0.021 | 0.015 | 0.018 |
| <i>Smaller Companies</i> | | | | |
| Beck's | 0.008 | 0.012 | 0.017 | 0.031 |
| Croplan | 0.017 | 0.020 | 0.033 | 0.029 |
| Growmark | 0.016 | 0.017 | 0.013 | 0.010 |
| Stine | 0.047 | 0.040 | 0.027 | 0.023 |
| <i>Public</i> | 0.048 | 0.025 | 0.004 | 0.004 |
| <i>GT Trait</i> | 0.200 | 0.726 | 0.945 | 0.939 |

Note: Authors’ computations from Kynetec data. Calculation excludes records that do not involve a commercial transaction (e.g., saved seeds), specifically as per line 5 of Table A1 in the Appendix.

Table 1 reports average market shares for the 13 largest brands over the considered timespan. Brands are grouped by their parent companies.⁷ We separately categorize public seeds and the seeds sold by brands not owned by the “big four” parent companies (Monsanto, DuPont, Syngenta,

⁷ The company names in Table 1 reflect the industry configuration as of 2016, the last year of our data. Since then, major mergers and acquisitions have re-shaped the ownership structure of the industry—the acquisition of Syngenta by ChemChina in April 2017, the merger of Dow and DuPont in September 2017, and the acquisition of Monsanto by Bayer in June 2018. The agricultural concerns consolidated by the Dow-DuPont mergers were subsequently spun off as Corteva in 2019.

and Dow AgroSciences).⁸ Table 1 also illustrates some turnover in brands. Golden Harvest was phased out by Syngenta in 2012; and, in recent years, Monsanto has phased out DeKalb (this brand is now primarily used in the maize seed market).

3.2 Markets and choice sets

In any one choice situation, a farmer may not have access to all 26 alternatives (i.e., choice sets are market-specific). Following standard convention, a market is defined as a CRD-year combination. CRDs are spatial constructs that are reasonably homogeneous with respect to key growing conditions, such that residing farmers are expected to face the same set of choices. Differentiating markets by years is a natural extension, as commercialized varieties evolve over time, and a calendar year contains a natural planting window. In our dataset, we observe a total of 3,791 markets across 233 CRDs. The number of markets in selected years, and the average number of choice alternatives (i.e., products) available to farmers, are provided in Table 2. Note that, following the introduction of GT varieties, the number of products available to farmers initially increased as leading varieties were made available to farmers both with and without the GT trait, but eventually decreased as GT products crowded out conventional products.

Table 2. Markets and Products

| Year | Number of Markets | Average Number of Products | | |
|------|-------------------|----------------------------|------|--------------|
| | | Total | GT | Conventional |
| 1996 | 165 | 9.92 | 1.24 | 8.68 |
| 2000 | 182 | 16.18 | 8.05 | 8.13 |
| 2004 | 174 | 12.19 | 8.38 | 3.81 |
| 2008 | 178 | 9.59 | 7.94 | 1.65 |
| 2012 | 188 | 9.61 | 7.63 | 1.98 |
| 2016 | 189 | 9.83 | 7.19 | 2.64 |

Note: Authors' computations from Kynetec data. Calculation excludes records that do not involve a commercial transaction (e.g., saved seeds), specifically as per line 5 of Table A1 in the Appendix.

⁸ The ownership of each brand as reported in Table 1 also pertains to 2016, the last year of our data. Brands' affiliation with their parent company in some cases was the result of market consolidation that took place earlier in our sample. This is particularly true for Monsanto, which acquired Asgrow in 1997, DeKalb in 1998, and Kruger in 2006. Also, DuPont acquired Pioneer in 1999; Syngenta acquired NK in 2000 and Golden Harvest in 2004; and, Dow acquired Mycogen in 1998.

3.3 Inertia

To motivate our focus on switching costs in farmers' seed choices, Table 3 reports the purchase rate and repurchase rate (as percentages) for each brand and trait. The purchase rates are the unconditional probability of choosing each brand (market shares). Conditional on a purchase at time t , the repurchase rate is the frequency of the same brand/trait having been also purchased in the previous period. Table 3 shows that repurchase rates are considerably higher than the corresponding purchase rates. These data are certainly suggestive of persistence or inertia in brand (and trait) choices over time, but of course they offer no information on whether this inertia is due to state dependence or unobserved heterogeneity.

Table 3. Purchase and Repurchase Rates

| Brand/Trait | Purchase | Repurchase |
|----------------|----------|------------|
| Asgrow | 17.27 | 86.96 |
| Beck's | 2.03 | 87.23 |
| Channel | 1.17 | 76.29 |
| Croplan | 2.69 | 71.83 |
| DeKalb | 3.77 | 77.29 |
| Golden Harvest | 3.5 | 78.26 |
| Growmark | 1.31 | 85.78 |
| Kruger | 1.59 | 82.53 |
| Mycogen | 1.91 | 80.17 |
| NK | 6.96 | 79.77 |
| Pioneer | 23.36 | 92.06 |
| Public | 1.19 | 63.11 |
| Stine | 3.03 | 78.32 |
| Conventional | 17.15 | 92.12 |
| GT | 82.85 | 97.40 |

Note: Authors' computations from Kynetec data. Calculation based on the estimation sample, as per line 9 of Table A1 in the Appendix.

3.4 Prices

A common challenge in discrete choice models is that the researcher observes the price for the alternative chosen by the individual, but does not observe the prices of unchosen alternatives. A typical solution to this problem is to compute average transaction prices, in any given market, and

use these prices as the prices that individuals face for each alternative. This is the procedure that we follow here. A related but distinct issue concerns discounts, which are common in our data—about 63% of 204,697 total observed purchases have a discount. Similar to Goldberg (1995), we account for discounts by using the market average net price for each product. Finally, we note that the dataset spans 21 years, a long period during which prices changed considerably. Consistent with the homogeneity property of the per-acre profit function, described in what follows, we deflate prices by the USDA crop sector index of prices paid (index = 1 in 2011).

4. Model Specification

The seed demand model is derived under the assumption that farmers, on each of their plots, choose the seed alternative that maximizes expected profit. The structure of their payoff function depends on the production technology, the output price, and all input prices. Ciliberto, Moschini, and Perry (2019) show that, if the production function satisfies two reasonable properties—constant returns to scale in all inputs, and fixed proportions between land and seed—then per-acre expected profit is linear in the seed price. That is, the per-acre profit function dual to the production technology can be written as $\pi_{ij}^0(\mathbf{w}) - \alpha_i \tilde{p}_j$, where i indexes the plot to be planted (i.e., the “choice situation” in what follows); j indexes the seed product; \tilde{p}_j is the price of seed alternative j (per “bag” of seed); α_i denotes the amount of seed per acre (seed density) that is appropriate for plot i ; and, \mathbf{w} is the vector of all other prices—the soybean output price and the prices of all inputs used in production, other than seed (which appears separately as \tilde{p}_j) and land (implying that the per-acre profit in question can be interpreted as returns to land).

Whereas the optimal amount of seed per acre can be presumed constant for a given plot in a given market, it can vary spatially because of soil and climatic conditions and over time because newer varieties are supposed to be planted at slightly lower densities (Gaspar et al., 2020; Perry, Hennessy, and Moschini 2022). Given that, and the fact that the actual definition of a bag of seed has changed over our sample period,⁹ in what follows we define seed prices on a per-acre basis, as

⁹ Traditionally, a bag of soybean seeds amounted to 50 pounds. Starting with the 2010 growing season, the industry transitioned to defining a bag on a count basis—a bag of seeds was re-defined as containing 140,000 soybean kernels. In our data, farmers on average use 1.186 bags of soybean seeds per planted acre.

in Ciliberto, Moschini, and Perry (2019). Hence, per-acre profit is represented as

$$\pi_{ij} = \pi_{ij}^0(\mathbf{w}) - p_j \quad . \quad (1)$$

where p_j is the price of seed per unit of land (\$/acre). Given this objective function, the farmer's problem for plot i can be formulated in a standard discrete-choice framework, as that of choosing product j that yields the highest payoff π_{ij} .¹⁰

4.1 Switching costs

In our context, switching costs possibly exist whenever a farmer chooses a seed product with attributes (specifically, brand and technology) that involve a change relative to previous choices/experience. For each alternative j available in a market at time t , let B_t^j denote the brand and $T_t^j \in \{CV, GT\}$ denote the trait, where CV represents conventional (non- GT) products. The farmer's own past experience is encapsulated by their seed purchase history. As discussed earlier, we focus on choices made in the previous period. Hence, the relevant history for farmer f , denoted H_t^f , is the set of all brands and technology traits included in their immediately preceding period purchases, i.e., $H_t^f \equiv \{B_{t-1}^j, T_{t-1}^j \mid \forall j \text{ chosen by } f \text{ at } (t-1)\}$. Switching costs can now be captured by five indicator variables—three primitive indicators, and two interactions. Specifically:

$$I_{fjt}^B = 1 \text{ if } B_t^j \notin H_t^f \quad (2)$$

$$I_{fjt}^{GT} = 1 \text{ if } T_t^j = GT \notin H_t^f \quad (3)$$

$$I_{fjt}^{CV} = 1 \text{ if } T_t^j = CV \notin H_t^f \quad (4)$$

$$I_{fjt}^B \times I_{fjt}^{GT} \quad (5)$$

$$I_{fjt}^B \times I_{fjt}^{CV} \quad (6)$$

¹⁰ Similar to Train and Winston (2007), the model we develop is a conditional demand model—only soybean seed choices are considered, conditional on the farmer having chosen to purchase soybeans (i.e., there is no “outside option”).

The indicator I_{fjt}^B codes for a choice that, for decision maker f , entails a new brand (vis-à-vis their own history). The indicator I_{fjt}^{GT} codes for a choice that entails purchasing the GT trait when, in the previous period, only conventional soybeans were grown. Similarly, the indicator I_{fjt}^{CV} codes for choosing a conventional seed when, in the previous period, the farmer only used GT seeds.¹¹ The two interaction terms capture the simultaneous switching of *both* brand and technology. Note that we do not consider the interaction effect between traits because, of course, $I_{fjt}^{GT} \times I_{fjt}^{CV} = 0$.

4.2 The econometric model

The primitive data identify the chosen product for choice situation i in market m . The market has time dimension $t = t[m]$, and the choice is made by farmer $f = f[i]$ (who is typically observed in more than one choice situation). To make the framework outlined in the foregoing operational, we parameterize the profit function in equation (1). In addition to the seed price—which enters linearly, as discussed—we approximate the impact of other structural determinants of $\pi_{ij}^0(\cdot)$ by a set of seed-specific variables. We also explicitly include the switching cost indicator variables. Specifically, the per-acre profit for farmer f from product j , in choice situation i pertaining to market m , is:

$$\pi_{ijm} = \mathbf{x}_{jt}\boldsymbol{\beta}_f - \alpha_f p_{jm} + \mathbf{I}_{fjt}\boldsymbol{\gamma}_f + \xi_{jm} + h_{ff} + \varepsilon_{ijm}, \quad (7)$$

where \mathbf{x}_{jt} is a vector of seed characteristics that include a full set of GT trait and brand indicator variables. Product characteristics are varying over time because the GT and brand indicator variables are interacted with dummy variables identifying the sub-periods 1997–2002 and 2003–2009 (thereby allowing the mean value of the GT trait and brand preference to vary over time). Switching costs are captured by the vector \mathbf{I}_{fjt} , which consists of the five indicator variables defined in equations (2)–(6). The unobserved component of the profit equation consists of three terms: ξ_{jm} is a product- and market-specific unobserved demand shock, h_{ff} is unobservable

¹¹ Of course, most farmers can be presumed to already have experience with conventional soybean seeds (unless they began farming in the GE era). Postulating an adjustment cost when moving from *GT* to *CV* seeds, however, captures the fact that the new technology may involve specialized equipment (e.g., herbicide sprayers) the cost of which is sunk with the initial adoption of GT seeds.

heterogeneity at the farm-product level, and ε_{ijm} is an i.i.d. error term that is assumed to follow a Type 1 extreme value distribution. Note also that the coefficient of p_{jm} (the price of product j in market m), which equals 1 in the structure of equation (1), reflects the arbitrary scale parameter of the assumed extreme value distribution (Train 2009). Insofar as the variance of this unobserved component varies across farmers, this is captured by the farmer-specific coefficient α_f .

Some of the coefficients in (7) are specified as random across farmers, as denoted by the subscript f , while others are not. For the random components of the β and γ coefficients, we assume a normal distribution. To illustrate, following standard notation (e.g., Gandhi and Nevo 2021), if component r of vector β is random, then

$$\beta_f^{(r)} = \beta_0^{(r)} + \beta_v^{(r)} v_f^{(r)}$$

where $v_f^{(r)}$ is a standard normal random variable (such that $\beta_0^{(r)}$ represents the mean and $\beta_v^{(r)}$ represents the standard deviation of the random coefficient). For the price coefficient, on the other hand, we use the lognormal distribution:

$$\ln \alpha_f = \tilde{\alpha}_0 + \tilde{\alpha}_v v_f^{(0)}$$

where $v_f^{(0)}$ is a standard normal. The lognormal assumption is desirable for two reasons: first, it maintains the economic condition that profit is decreasing in the price of seed for all farmers; and second, as discussed later, it provides an elegant solution to the challenge of computing WTP estimates in random coefficient models (a task that involves the ratios of two random variables).

4.2 Identification

We face two identification challenges in the estimation of the seed demand model. The first issue is the well-known problem of price endogeneity in demand models of differentiated products (Berry 1994; Berry, Levinsohn, and Pakes 1995). Although our model is estimated using individual choices, which may alleviate some concerns about endogeneity (Goldberg 1995), there may still remain unobservable factors correlated with both price and demand. That is, the unobserved product-market term, ξ_{jm} , is potentially correlated with price. The inclusion of product fixed effects may alleviate this problem (Nevo 2000). Although we include brand and trait intercepts, and further interact these terms with time, any location-specific and/or market by

product-specific unobserved attributes, correlated with demand and observed by firms, may still lead to price endogeneity. The most common solution to this problem is to use two-stage least squares with instrumental variables (IVs). This approach, however, cannot be directly applied in non-linear individual-level discrete choice models (Train 2009). Therefore, we implement a control function approach (Petrin and Train 2010; Kim and Petrin 2019).

The basic ingredient of this procedure is to specify a reduced-form equation for the endogenous (price) variable. This is specified linearly, as follows:

$$p_{jm} = \mathbf{z}'_{jm}\boldsymbol{\omega} + v_{jm} \quad (8)$$

where \mathbf{z}_{jm} is a vector that includes market-level demand shifters plus a set of excluded IVs (in our case, cost shifters), and v_{jm} represents the unexplained portion of prices. Let \tilde{v}_{jm} denote the control residuals recovered from the least squares estimation of equation (8).¹² Then, following Kim and Petrin (2019), we replace the unobserved demand characteristic by a linear function of the control residuals: $\xi_{jm} = \lambda\tilde{v}_{jm}$.¹³

The selection of suitable IVs included in (8) is predicated on the structure of seed production—seed sold in year t is produced in year $t - 1$. Specifically, soybean seed companies typically contract out with specialized individual growers to produce their commercial seed supply, from the seed stock generated by their breeding program (Lamkey 2004). The terms of the contract are

¹² In addition to the instrumental variables discussed in the foregoing, the right-hand-side of the price equation includes the vector of characteristics of product j interacted with the dummies identifying the sub-periods of interest, as in equation (7), the previous year's market share for the brand associated with product j , the previous year's market share for the technology (conventional or GT) associated with product j , the size of market m , as well as year, CRD, and brand fixed effects. The inclusion of lagged share variables is meant to capture, in a reduced-form way, the dynamic aspects of firms' pricing decisions. The estimated coefficients for this equation are reported in Table A5 of the Appendix. Because the price equation conceivably represents equilibrium prices, the expected signs of the estimated coefficients are, in general, ambiguous and difficult to interpret (e.g., Pakes 2003).

¹³ As noted in Kim and Petrin (2019), the true control function may take a nonlinear form, and/or may depend on *all* residuals. We investigated the robustness of our approach by considering control functions where the impact of the own-product residual differed by time ($\lambda_t\tilde{v}_{jm}$) or by trait ($\lambda_j\tilde{v}_{jm}$). We also estimated models that included combinations of competing product residuals. In all cases, we found no significant difference in results across these alternative control functions. Our implementation of the control function approach is similar to the procedures used in Polyakova (2016) and Janssen (2022).

set such that the grower is paid at least what they could have obtained had they planted and sold their own soybeans as a commodity. The cost to the seed company of the “seed multiplication” stage thus depends on expected soybean output prices. A standard proxy for a commodity’s expected output price is the futures price corresponding to delivery in the month following harvest, as quoted at planting time (Kim and Moschini 2018). Because this applies to the year when the seed is actually grown under contract for the seed companies, we use the *previous year’s* soybean futures price as an instrument for the current year’s seed prices. Hence, this IV is highly correlated with seed companies’ costs, and also does not affect farmers’ relative demand for soybean seed products (i.e., it fulfills the exclusion restriction requirement).¹⁴ To allow for differential impacts across products, we interact futures prices with the brand and GT trait dummies, similar to the approach taken by Berto Villas-Boas (2007).

The second identification issue concerns the disentanglement of switching costs from unobservable heterogeneity. The basic issue is that inertia, or repeated purchases by a farmer, may stem from switching costs (structural state dependence) or unobservable heterogeneity (spurious state dependence). The inclusion of brand-specific and trait-specific random coefficients already controls for much of this heterogeneity. However, we must also address the so-called “initial condition” problem articulated by Heckman (1981b). The crux of the matter is that a farmer’s observed first choices in the sample are stochastically dependent on their individual heterogeneous preferences, which leads to a correlation between the switching cost variables and the unobserved heterogeneity term h_{fj} .

To deal with the initial condition problem we implement a simplified procedure along the lines of Wooldridge (2005), who models the distribution of unobserved effects by conditioning it directly on the initial observed values. That is, similar to Janssen (2022) and Rickert (2022), the alternative-specific individual payoffs in equation (7) are permitted to depend explicitly on farmers’ first observed choices (which are excluded from the estimation sample). For this purpose, we generate a set of indicator variables that, for each seed alternative j in choice situation i evaluated by farmer $f = f[i]$, code for how this alternative relates to farmer f ’s purchase(s) in their first observed time

¹⁴ If the previous year’s futures price is correlated with the current year’s futures price, it may correlate with a farmers’ decision of *which crop* to plant. Because the model we estimate is a *conditional* soybean demand model, however, such correlation is immaterial for the choice of which particular soybean seed to plant.

period.¹⁵ Specifically, we define three indicator variables to capture the effects associated with the initial choices of products, brands, and technologies: D_{ff}^P takes a value of 1 if farmer's f initial purchase set included the product of alternative j (value of zero otherwise); D_{ff}^B takes a value of 1 if farmer's f initial purchase set included the brand of alternative j (value of zero otherwise); and D_{ff}^T takes a value of 1 if farmer's f initial purchase set included the technology (GT or non-GT) of alternative j (value of zero otherwise). Conditioning on these three initial-condition variables, the unobserved heterogeneity term h_{ff} is represented as:

$$h_{ff} = \phi_1 D_{ff}^P + \phi_2 D_{ff}^B + \phi_3 D_{ff}^T$$

Thus, this parameterization of the initial condition maintains that when these variables take value 1, a farmer's expected profit for the corresponding product(s) is permanently shifted for all subsequent periods they are observed.¹⁶

Collecting the three initial condition indicator variables in the vector \mathbf{D}_{ff} , the per-acre profit of equation (7) that includes the initial period choices and the control function can be expressed as:

$$\pi_{ijm} = \mathbf{x}_{jt} \boldsymbol{\beta}_f - \alpha_f p_{jm} + \mathbf{I}_{ffjt} \boldsymbol{\gamma}_h + \lambda \tilde{v}_{jm} + \mathbf{D}_{ff} \boldsymbol{\phi} + \varepsilon_{ijm} \quad (9)$$

where the coefficients of the initial-condition variables and control function are estimated jointly with the structural parameters.

¹⁵ Recall that a farmer may make multiple purchases in a given year. As a result, the dummy variables pertaining to their first period choices can take value 1 for more than one product, brand, and/or technology.

¹⁶ Unobserved heterogeneity can also depend on the value of the exogenous variables and as well interactions of a farmer's first period choice with these values. In line with these possibilities, we explored several alternative specifications for h_{ff} . These included allowing the coefficient on the brand initial condition to differ for each brand, allowing for the first period choice to impact price sensitivity, allowing for the first period choice to impact the control function, and allowing the initial period price to impact returns. In all cases, we found minor differences in the estimation results and therefore elected to use the proposed simpler specification.

4.3 Estimation

The model is estimated using simulated maximum likelihood, as outlined in Hole (2007) and Train (2009). The profit function in (9) can be decomposed into the non-random and random components:

$$\pi_{ijm} = \underbrace{\mathbf{x}_{jt}\boldsymbol{\beta}_0 + \mathbf{I}_{fjt}\gamma_0 - \alpha_0 p_{jm} + \lambda \tilde{v}_{jm} + \mathbf{D}_{ff}\boldsymbol{\phi}}_{\delta_{fjm}} + \underbrace{(\mathbf{x}_{jt} \circ \mathbf{v})\boldsymbol{\beta}_v + (\mathbf{I}_{fjt} \circ \mathbf{v}^{(I)})\boldsymbol{\gamma}_v - \alpha_v p_{jm} v^{(0)}}_{\mu_{fjm}} + \varepsilon_{ijm} \quad (10)$$

where $(\mathbf{v}, \mathbf{v}^{(I)}, v^{(0)})$ is the vector of i.i.d. standard normal variates that characterize the random coefficients, and \circ denotes the Hadamard (element-by-element) product operator. Note that, similar to standard approaches, we have separated the mean components δ_{fjm} from the standard-deviation components μ_{fjm} .

For a *given* realization of μ_{fjm} , the probability that farmer f chooses alternative j in choice situation i is given by the familiar logit expression:

$$S_{ij}^f(\boldsymbol{\theta}_f) = \frac{\exp(\delta_{fjm} + \mu_{fjm})}{\sum_{n=1}^{J_m} \exp(\delta_{fjm} + \mu_{fjm})} \quad (11)$$

where $\boldsymbol{\theta}_f \equiv [\boldsymbol{\beta}_0, \alpha_0, \gamma_0, \lambda, \boldsymbol{\phi}, \boldsymbol{\beta}_v, \alpha_v, \boldsymbol{\gamma}_v]$ is the vector of coefficients to be estimated, and J_m denotes the number of alternatives in farmer f 's choice set (which is market-specific).¹⁷

For each farmer, we observe a sequence of choices, the probability of which is given by the product of the logits in equation (11). To obtain the unconditional probability of an individual's choice sequence, it is necessary to integrate out the random elements in $\boldsymbol{\theta}_f$. Because this integral has no closed-form solution, the simulated likelihood is constructed using Halton draws (we use 250 Halton draws). For estimation, we utilize the Stata "mixlogit" package written by Hole (2007).

¹⁷ Because of the assumed lognormal distribution of the price coefficient, discussed earlier, here $\alpha_0 \equiv E[\alpha_f] = \exp(\tilde{\alpha}_0 + \tilde{\alpha}_v^2/2)$ and $\alpha_v^2 \equiv Var(\alpha_f) = \exp(2\tilde{\alpha}_0 + \tilde{\alpha}_v^2)[\exp(\tilde{\alpha}_v^2) - 1]$. Note also that Stata output does not provide α_0 and α_v directly, it gives $\tilde{\alpha}_0$ and $\tilde{\alpha}_v$.

5. Results

The original data, after dropping observations pertaining to choices of non-newly purchased seed varieties, include 204,697 seed choices. Because the original data are from an unbalanced sample, it is important to ensure the estimation sample is suitable to identify switching costs. To this end, we require that observations included in the estimation sample in any year t pertain to farmers for whom observations are also available for year $t - 1$. This procedure permits the switching cost indicators to be consistently constructed from the preceding year's choice for all observations in the estimation sample. Following these data preparation steps (see Table A1 in the Appendix), the estimation sample includes 86,198 seed choices. The data is then expanded for each observed choice instance to include all possible market alternatives a farmer could have chosen, which results in the estimation dataset of 1,043,742 available choice alternatives. Table A2 in the Appendix shows that the average farmer chooses from about 8 seed alternatives, with a minimum of two and a maximum of 23 (recall that our product definition results in 26 possible alternatives). On average, a farmer selects about 1.65 products per year (and 2.4 products while present in the sample), and chooses slightly fewer brands than products. On average, a farmer contributes three years of data to the sample.

Table 4 summarizes the descriptive information for the primary variables used in estimation. The mean of the brand switching variable is 0.129—that is, about 13% of purchase instances consisted of a farmer choosing a brand they did not choose in the prior year. Switching was even less common among traits, with only 2.2% of cases involving a switch into the GT trait and 1.4% of cases involving a switch into a conventional variety. The low frequency of technology switches is the manifestation of the near complete adoption of the GT trait by the mid-2000s, with most switching occurring prior to the year 2000. For estimation, the trait and brand coefficients are allowed to differ over three sub-periods, specifically 1997–2002, 2003–2009, and 2010–2016. This permits the model to capture, in a parsimonious way, relative changes in the value of the various brands and the GT trait over time. The sub-period intervals themselves were chosen to roughly coincide with some of the major changes that took place in the industry. The first sub-period encompasses the earlier stages of rapid GT adoption, which averaged 61% in the period 1997–2002. The second period encompasses the mature stage of GT adoption, as well as the commodity price boom from 2006–08, and the final period covers the emergence of glyphosate weed resistance, which reportedly began to intensify after 2010.

Table 4. Descriptive Information on Main Model Variables

| Variable | Obs | Mean | Std. Dev. | Min | Max |
|--|--------|--------|-----------|-------|---------|
| Price (\$/acre) | 86,198 | 46.560 | 9.632 | 1.204 | 108.941 |
| Switching cost Indicators | | | | | |
| Brand | 86,198 | 0.129 | 0.335 | 0 | 1 |
| Conventional to GT trait | 86,198 | 0.022 | 0.145 | 0 | 1 |
| GT trait to conventional | 86,198 | 0.014 | 0.115 | 0 | 1 |
| Interaction brand × conventional to GT | 86,198 | 0.007 | 0.081 | 0 | 1 |
| Interaction brand × conventional to GT | 86,198 | 0.005 | 0.068 | 0 | 1 |
| GT technology adoption | | | | | |
| 1997-2002 | 29,445 | 0.612 | 0.487 | 0 | 1 |
| 2003-2009 | 26,020 | 0.945 | 0.228 | 0 | 1 |
| 2010-2016 | 30,733 | 0.938 | 0.242 | 0 | 1 |

Note: Authors' computations from Kynetec data. Calculation based on the estimation sample, as per line 9 of Table A1 in the Appendix.

5.1 Estimation results

Before discussing the main results for the mixed logit model, it is helpful to consider several alternative specifications for the simpler conditional logit model. The full mixed logit model requires significant estimation time, so the conditional logit model was used as a means for a preliminary assessment of the impact of different factors on model estimates. Table A3 in the Appendix presents estimation results for four specifications of the conditional logit model.

Looking across all specifications in Table A3, the estimated coefficients generally accord with expectations. The GT trait coefficients are positive and significant, and the switching cost coefficients are negative and significant, with the exception of the interaction coefficients, which are positive and statistically significant. The latter implies that the cost of switching both brand and trait is less than the sum of switching to each in isolation (in other words, the cost of switching both is sub-additive). The final column reports the most general model, which includes the switching cost variables, the control residual, and the initial conditions. The first column points to the consequences of excluding the switching cost variables. The log-likelihood is approximately 24% smaller compared to the full model. The importance of the control function approach is demonstrated by the results in the second column. The price coefficient is approximately zero and

statistically insignificant, suggesting that price endogeneity (biasing estimates towards zero when ignored) is indeed present. The positive and significant estimate for the control residual in the other columns also suggests the existence of unobservable factors positively correlated with demand, a finding similar to that in Petrin and Train (2010). The third column omits the initial conditions, which increases the brand switching cost coefficient and decreases the log-likelihood by approximately 1.3%. Overall, Table A3 provides support for the specification used in the mixed logit model.

Table 5 presents the mixed logit estimation results for the primary variables.¹⁸ The first two columns provide the estimated mean coefficients and their standard errors, and the final two columns provide the estimated standard deviation coefficients and their standard errors. These results demonstrate the importance of specifying random coefficients for switching costs, the GT trait indicator variable, and the brand variables. All standard deviation coefficients are statistically significant and indicate significant heterogeneity in trait value, price responsiveness, the magnitude of switching costs, and brand value. In addition, the log likelihood increases by approximately 6% compared to the full conditional logit specification.

Table 5 also includes the estimated coefficients for the three initial condition variables. The results indicate that estimated parameters are positive and statistically significant, implying that first period choices impact returns in every subsequent period. Not included in Table 5, for space reasons, are the estimated parameters for the brand-period interactions. These results, reported in Table A4 of the Appendix, indicate that, compared to the final 2010–2016 sub-period, all brands were valued less (relative to public varieties) in the first two sub-periods. This is consistent with the decline in demand for public varieties.

¹⁸ Recall that the price coefficient is modeled as log-normally distributed, whereas all the other random coefficients are assumed normally distributed.

Table 5. Mixed Logit Estimation Results

| | ---Mean--- | | ---Standard Deviation--- | |
|-------------------------------|-------------|---------|--------------------------|---------|
| | Coefficient | SE | Coefficient | SE |
| Price (\$/acre) | -1.594** | (0.091) | 0.087** | (0.018) |
| GT (1997–2002) | 3.760** | (0.270) | 1.070** | (0.071) |
| GT (2003–2009) | 5.019** | (0.313) | 1.866** | (0.164) |
| GT (2010–2016) | 4.415** | (0.313) | 2.863** | (0.235) |
| Switching Cost: Brand | -2.648** | (0.064) | 1.363** | (0.035) |
| Switching Cost: GT | -1.371** | (0.082) | 1.336** | (0.114) |
| Switching Cost: Non-GT | -2.974** | (0.154) | 1.904** | (0.158) |
| Switching Cost: Brand×GT | 0.408** | (0.086) | | |
| Switching Cost: Brand×Non-GT | 1.026** | (0.093) | | |
| Asgrow | 3.634** | (0.342) | 0.661** | (0.037) |
| Beck | 2.909** | (0.351) | 1.301** | (0.142) |
| Channel | 1.938** | (0.357) | 1.991** | (0.149) |
| Croplan | 3.078** | (0.358) | 0.733** | (0.068) |
| Dekalb | 0.928** | (0.272) | 0.327** | (0.097) |
| Golden | 1.355** | (0.375) | 0.931** | (0.114) |
| Growmark | 2.755** | (0.383) | 0.712** | (0.108) |
| Kruger | 2.123** | (0.417) | 1.054** | (0.120) |
| Mycogen | 1.577** | (0.309) | 1.234** | (0.102) |
| NK | 3.080** | (0.338) | 0.671** | (0.063) |
| Other | 2.953** | (0.297) | 0.818** | (0.036) |
| Pioneer | 3.203** | (0.302) | 0.629** | (0.048) |
| Stine | 1.872** | (0.282) | 0.681** | (0.094) |
| Control Residual | 0.205** | (0.019) | | |
| Initial Condition: Product | 0.531** | (0.036) | | |
| Initial Condition: Brand | 0.441** | (0.039) | | |
| Initial Condition: Technology | 0.349** | (0.057) | | |

Note: The omitted brand dummy is for “Public.” Standard errors (SEs), in parentheses, are clustered at the CRD level. $LL = -110,364$. $N = 1,043,742$. The estimated model is based on 86,198 seed choices. The model also includes brand indicators with the first two subperiod dummies (see Table A4 in the Appendix for these remaining estimated parameters). * $p < 0.05$, ** $p < 0.01$

5.2 Willingness-to-pay distributions

A useful characterization of the various economic factors that determine choice, given the estimated demand model, involves the concept of WTP. The marginal WTP for an attribute measures the maximum amount a buyer is willing to part with for that characteristic, and is obtained by dividing the coefficient of interest by the price coefficient (Train 2009). Here we focus on the WTP associated with the main variables of interest: the switching costs, and the GT trait.

Because we also allow farmers to have heterogeneous sensitivity to seed prices, the WTP for an attribute is given by the ratio of two random variables, which can be problematic. Our assumption that the price coefficient is log-normally distributed, however, turns out to be useful.¹⁹ The WTP to be estimated can be expressed as $E[c/\alpha]$, where c denotes the coefficient of the characteristic of interest, and α is the price coefficient. Our maintained assumptions are that $c \sim N(\mu_c, \sigma_c^2)$ and $\ln \alpha \sim N(\mu_{\ln \alpha}, \sigma_{\ln \alpha}^2)$. Thus, $E[c/\alpha] = E[c \cdot \exp(-\ln \alpha)]$. It follows that we can compute the expected value of the ratio of the random variables of interest in closed form. That is, the expected value of the WTP of interest is computed analytically as:

$$E[c/\alpha] = \mu_c \exp\left(-\mu_{\ln \alpha} + \frac{\sigma_{\ln \alpha}^2}{2}\right) \quad (12)$$

Table 6 contains the estimated WTP values based on the coefficient estimates from Table 5. The expected values are computed analytically according to equation (12). The 10th percentile, the median, and the 90th percentile of the estimated WTP distributions are computed numerically (based on 100,000 random draws). We note that the numerical percentiles indicate relatively symmetric distributions for the WTP estimates. This is because the estimated standard deviation of the (lognormally distributed) price coefficient is rather small (about 10% of the mean coefficient), such that the distribution of the WTP estimates primarily reflects the distribution of the parameters of the underlying characteristics, which are modeled as normal.

¹⁹ A normal distribution for the price coefficient, still common in the literature, has two main limitations. First, its unbounded support permits the decision maker to prefer higher prices to lower prices. Second, the implied WTP distribution is a Cauchy distribution, which has heavy tails and non-existent moments. Our choice of the log-normal distribution for the price coefficient has two desirable properties: the sign of the estimated price coefficient accords with economic theory (Hole and Kolstad 2012); and, the implied WTP distributions have well-defined moments (Hole and Kolstad 2012; Daly, Hess, and Train 2012).

Table 6. WTP Estimates

| | Expected value (analytical) | Distribution (numerical) | | |
|------------------------|--------------------------------|--------------------------|--------|-------|
| | | p10 | p50 | p90 |
| Switching costs | | | | |
| Brand | -13.09 | -21.85 | -13.00 | -4.42 |
| Trait: CV to GT | -6.78 | -15.29 | -6.73 | 1.73 |
| Trait: GT to CV | -14.70 | -26.92 | -14.57 | -2.59 |
| Brand & Trait CV to GT | 2.01 | 1.80 | 2.01 | 2.24 |
| Brand & Trait GT to CV | 5.07 | 4.52 | 5.05 | 5.65 |
| GT Trait | | | | |
| 1997–2002 | 18.59 | 11.63 | 18.42 | 25.78 |
| 2003–2009 | 24.81 | 12.83 | 24.64 | 36.93 |
| 2010–2016 | 21.82 | 3.49 | 21.60 | 40.21 |

Note: These estimates are based on the mixed logit results reported in Table 5. Entries in this table are interpreted as \$/acre (measured in 2011 dollars). The brand WTP values (relative to Public varieties) are for the period 2010–2016.

From Table 6, the average cost of switching to a new a brand is \$13.1/acre. Given that the mean price for a unit of soybeans is about \$47/acre, this cost is about 28% of the average price, a significant effect.²⁰ There is also considerable heterogeneity in brand switching costs. A substantial number of farmers (10%) face a cost of over \$21.8/acre if they switch to a new a brand, whereas some farmers (10%) incur a cost of \$4.4/acre or less. The trait switching costs are similarly large in magnitude and exhibit interesting asymmetries. On average, it costs about \$6.8/acre to switch to a GT variety, whereas it costs about \$14.7/acre to switch to a conventional variety. There are several reasons behind this asymmetry. First, it is relatively easy to switch to a GT system: farmers simply need to use glyphosate in combination with a GT variety. By contrast, a conventional soybean production system is more complex, as it requires multiple types of herbicides, applied at different times, as well as mechanical cultivation. For an individual who had not recently planted conventional soybeans, there may be substantial costs to doing so. Finally, the interaction terms indicate subadditivity in switching costs—the cost of switching both brand and

²⁰ As shown in Table 4, the average retail seed price is about \$46.56/acre. All prices are deflated by the crop sector index of prices paid (base year = 2011). To provide additional context for these estimates, in 2011 the soybean output price was \$12.5/bu. and the US average soybean yield was 42 bu./acre, implying an average total gross revenue of about \$525 per acre (based on USDA-NASS data).

technology is less than the sum of costs of switching each in isolation. Similar to brand switching costs, technology switching costs also exhibit significant heterogeneity.

The three WTP distributions for the GT trait illustrate interesting changes over time. In the first period (1997–2002), the average WTP is \$18.6/acre, which increases to \$24.8/acre in the second period (2003–2009), and then decreases to \$21.8/acre in the final period (2010–2016). As noted in Ciliberto, Moschini, and Perry (2019), the observed increase after the period 1997–2002 was likely the result of falling glyphosate prices (a complementary input to GT soybeans), rising output prices, and access to an increasing number of varieties with the GT trait. However, in the final period, the mean WTP for the GT trait decreases and its variance increases. This is consistent with the decline in the soybean output price after the 2007–2009 commodity price boom, as well as some loss of efficacy of GT soybeans. In recent years, glyphosate weed resistance has become increasingly problematic (Perry et al. 2016). Some farmers have responded by switching to non-GT varieties or by increasing glyphosate application rates (Perry, Hennessy, and Moschini 2019). Glyphosate weed resistance also varies considerably across regions, which may explain the increasing variance in farmers' WTP.

5.3 Price elasticity of demand

A major recognized advantage of the mixed logit framework, as contrasted with the basic logit model, is that it can capture richer substitution patterns (Berry, Levinsohn, and Pakes 1995). These patterns can be illustrated in terms of elasticities. In addition, we can also use our model to assess the impact of switching costs on elasticities by simulating price elasticities in a scenario without switching costs. Table 7 provides a summary of the (numerically computed) own- and cross-price elasticities in each scenario.

The base scenario provides price elasticities using the estimated coefficients from Table 5, whereas the scenario with no switching costs provides mean elasticities with the vector of all switching costs, denoted by \mathbf{I}_{jft} , set to zero (while still using the estimated coefficients from Table 5). We compute these elasticities separately by whether or not the seed product contained the GT trait, and, in the case of cross-price elasticities, whether brand and/or seed technology changes. In the baseline scenario, the mean own-price elasticity is -5.0 across all conventional products, and -6.0 for GT products (the higher elasticity likely reflecting the fact that the latter have a higher average

price). When removing switching costs, own-price elasticities increase in absolute value (to -6.3 for conventional products, and to -7.8 for GE products). As expected, removing switching costs results in a more elastic demand because it is now less costly to switch to a different variety.

Table 7. Estimated Elasticities

| | | Baseline | No Switching Costs |
|-----------------------|------------------------------------|----------|--------------------|
| Conventional Products | Own-price | -5.02 | -6.03 |
| | Cross-price, CV-CV (across brands) | 0.20 | 0.23 |
| | Cross-price, CV-GE (same brand) | 1.20 | 0.77 |
| | Cross-price, CV-GE (across brands) | 0.20 | 0.31 |
| GE Products | Own-price | -6.30 | -7.80 |
| | Cross-price, GE-GE (across brands) | 0.46 | 0.51 |
| | Cross-price, GE-CV (same brand) | 0.17 | 0.13 |
| | Cross-price, GE-CV (across brands) | 0.03 | 0.06 |

Note: These elasticities are computed, numerically, based on the estimated mixed logit model of Table 5.

The cross-price elasticities also exhibit systematic patterns. In the baseline, for conventional products the largest cross-price elasticity is for products within the same brand, even when switching technologies. On the other hand, for GE products, the largest substitutability is with other GE products. Given that own-price elasticities become more negative when switching costs are eliminated, cross-price elasticities must become more positive on average (because of the price homogeneity and adding-up effects). In fact, the cross-price elasticities across brands increase, whereas the cross-price elasticities across technologies and in the same brand decline somewhat.

6. Counterfactual Analysis

Having established that switching costs are significant in the soybean seed industry, in this section we assess the implications of switching costs for three important facets of the industry: buyers' welfare, technology adoption, and sellers' brand shares. To this end, we complement the estimated demand side with an explicit supply side model that postulates a standard static Nash-in-prices formulation (Nevo, 2001).

6.1. Bertrand-Nash equilibrium

The definition of products used in this paper is such that each brand is associated with two products, conventional and GT soybean seeds (except for “Public,” which is only sold as conventional, and “Channel,” a newer brand that only markets GT products). Furthermore, some firms own multiple brands. Let J_m^k denote the set of products that firm $k \in K$ sells in market m . The optimality conditions for profit maximization, which implicitly define the best response functions of the Bertrand-Nash game, are:

$$s_{jm}(\mathbf{p}_m) + \sum_{i \in J_m^k} (p_{im} - mc_{im}) \frac{\partial s_{im}(\mathbf{p}_m)}{\partial p_{jm}} = 0 \quad \forall j \in J_m^k, \forall k \in K \quad (13)$$

where $s_{jm}(\mathbf{p}_m)$ is the market share of product j in market m ; \mathbf{p}_m is the vector of all product prices in this market; and, mc_{jm} is the (constant) marginal cost of product j .

To express these conditions in matrix notation, let $\mathbf{S}(\mathbf{p}_m)$ be the $J_m \times J_m$ matrix of substitution terms, i.e., with elements $S_{jn} \equiv \partial s_{jm}(\mathbf{p}_m) / \partial p_{nm}$. Here, J_m is the total number of products in market m (if the full set of products is present, then $J_m = 26$). The information from the substitution matrix can then be combined with the $J_m \times J_m$ ownership matrix \mathbf{H}_m , which has elements $\mathbf{H}_{m,ij} = 1$ if products i and j are sold by the same firm, and $\mathbf{H}_{m,ij} = 0$ otherwise. This results in the $J_m \times J_m$ matrix $\mathbf{\Omega}(\mathbf{p}_m) \equiv -\mathbf{S}(\mathbf{p}_m) \circ \mathbf{H}_m$, where \circ denotes the Hadamard product of two matrices of the same dimension. The Bertrand-Nash equilibrium conditions in equation (13) can then be expressed as:

$$\mathbf{s}(\mathbf{p}_m) - \mathbf{\Omega}(\mathbf{p}_m) [\mathbf{p}_m - \mathbf{mc}_m] = 0 \quad (14)$$

where $\mathbf{s}(\mathbf{p}_m)$ and \mathbf{mc}_m are the vectors of market shares and marginal costs, respectively.

In the context of our paper, this standard setup is an approximation because of two sets of motives. First, it represents a static condition that does not capture dynamic pricing incentives that can arise in the presence of switching costs, leading to the (contrasting) “investment” and “harvesting” effects (e.g., Farrell and Klemperer 2007; Cabral 2016). The econometrics of the resulting stochastic dynamic models are data intensive and taxing from a computational point of view

(Akerberg et al. 2007), and outside the scope of this paper. We note, at any rate, that the reduced-form representation of equilibrium prices, used in the control-function approach to estimating the demand structure (Table A5 in the Appendix), indicates only a minor role for dynamic effects as captured by lagged shares. Furthermore, from equation (14) it is clear that equilibrium markups $\mathbf{p}_m - \mathbf{mc}_m = \mathbf{\Omega}(\mathbf{p}_m)^{-1} \mathbf{s}(\mathbf{p}_m)$ depend on price elasticities (as embedded in the substitution matrix used to define $\mathbf{\Omega}$). Insofar as switching costs affect elasticities, as documented in Table 7, the foregoing equilibrium relations still capture an important channel for inertia to affect equilibrium prices.

Second, the standard Nash equilibrium condition in equation (14) is only partially consistent with the licensing of the GT trait by Monsanto. For example, with fixed per-unit royalty rates, licensing would have two effects: it would raise the marginal production cost of Monsanto's competitors, and it may affect Monsanto's own pricing incentives. Whereas equation (14) implicitly accounts for the marginal cost effect of licensing, which is internalized by all firms (including the licensor), it does not account for the possible strategic pricing effects that could arise because Monsanto both licenses the GT trait and competes with the licensees (the precise nature of such effects, in any case, would depend on the undisclosed terms of licensing agreements).

Notwithstanding these limitations, the standard Bertrand-Nash equilibrium model can provide useful insights into the role of switching costs. Having estimated the demand model, we can exploit the equilibrium conditions to back out the vector of (otherwise unobserved) marginal costs, which are used in the counterfactual analyses.

6.2. Implementation of counterfactuals

The overarching goal of the counterfactual analysis concerns the effect of switching costs. Removing all switching costs is equivalent to setting $\gamma_0 = \gamma_v = 0$ in equation (11). Let $\tilde{s}_{jm}(\mathbf{p}_m)$ denote the counterfactual market shares one obtains under these parametric restrictions, at given market prices. The vector of counterfactual equilibrium prices $\tilde{\mathbf{p}}_m$ then satisfies:

$$\tilde{\mathbf{p}}_m = \mathbf{mc}_m + \tilde{\mathbf{\Omega}}(\tilde{\mathbf{p}}_m)^{-1} \tilde{\mathbf{s}}(\tilde{\mathbf{p}}_m) \quad (15)$$

To solve equation (15), we fix the vector of marginal costs, as estimated from the equilibrium

conditions in the baseline model in (14), and apply a fixed point algorithm.²¹

Because the foregoing representation of the Nash equilibrium applies at the market level, whereas our mixed logit estimates individual demands, we need to transition from individual choice probabilities, as represented in equation (11), to market shares. Specifically, the market-level vector $\tilde{s}_{jm}(\mathbf{p}_m)$ and matrices $\tilde{\Omega}(\tilde{\mathbf{p}}_m)$ are computed as averages, over all buyers in the market, from the corresponding individual level probabilities. Because the model entails random coefficients, individual probabilities are themselves averages over repeated random draws of the random coefficients.²² While implementing this procedure, we encountered some numerical problems with the standard fixed point iteration based directly on equation (15). Such issues were overcome by applying, instead, the fixed-point iteration developed by Morrow and Skerlos (2011). Their procedure entails the decomposition of the substitution matrix $\tilde{\mathbf{S}}(\mathbf{p}_m)$ into a diagonal matrix $\Lambda(\mathbf{p}_m)$ and a non-diagonal matrix $\Gamma(\mathbf{p}_m)$, such that $\tilde{\mathbf{S}}(\mathbf{p}_m) = \Lambda(\mathbf{p}_m) - \Gamma(\mathbf{p}_m)$, and the fixed point iteration only requires inversion of the diagonal matrix $\Lambda(\mathbf{p}_m)$ (rather than inversion of the dense matrix $\tilde{\Omega}(\tilde{\mathbf{p}}_m)$ in the standard procedure).²³

6.3. Welfare impact of switching costs

The estimated coefficients from the demand model indicate that switching costs are large in magnitude, but this does not directly reveal the actual welfare-reducing effect of switching costs. An individual's expected profit can only be directly impacted by switching costs insofar as their choice set contains products that were not purchased in the previous period. Consider, for example, a farmer that purchased all brands in the previous period. In this case, there would be no brand switching cost associated with any product. Thus, for this farmer, there would be no direct welfare effect from removing brand switching costs. The only potential effect would be due to the change

²¹ The notations $\tilde{\Omega}(\cdot)$ and $\tilde{\mathbf{s}}(\cdot)$ make it explicit that the market shares and matrix of substitution effects change not only because they are evaluated at a different price vector, but also because the counterfactual scenario is changing the structure of the choice probabilities.

²² For each simulation step we use 250 antithetic random draws from the standard normal distribution. All computations were coded in Matlab.

²³ Conlon and Gortmaker (2020) also report on the superiority of Morrow and Skerlos (2011) procedure.

in prices resulting from the removal of switching costs.

To estimate the actual welfare effects of switching costs we evaluate the change in the standard log-sum formula (Train, 2009), which expresses the expected profit per acre associated with a choice set. Recall from equation (10) that $(\delta_{fjm} + \mu_{fjm}) \equiv \psi_{fjm}$ denotes expected profit per acre.

For a given realization of the random coefficients, expected profit in the base scenario is given by:

$$\hat{\psi}_{fjm} = \mathbf{x}_{jt} \hat{\beta}_0 + \mathbf{I}_{fjt} \hat{\gamma}_0 - \hat{\alpha}_0 p_{jm} + \hat{\lambda} \tilde{v}_{jm} + \mathbf{D}_{ff} \phi + (\mathbf{x}_{jt} \circ \mathbf{v}) \hat{\beta}_v + (\mathbf{I}_{fjt} \circ \mathbf{v}^{(l)}) \hat{\gamma}_v - \hat{\alpha}_v p_{jm} v^{(0)} \quad (16)$$

The main counterfactual of interest removes all switching costs. This scenario is implemented by setting all coefficients of the indicator variables \mathbf{I}_{fjt} to zero. Expected profit in this counterfactual for any given draw of the random coefficient, denoted $\tilde{\psi}_{fjm}$, is computed from equation (16) under the constraint $\gamma_0 = \gamma_v = 0$ and by replacing observed prices with the counterfactual price \tilde{p}_m obtained by solving equation (15). The change in per-acre expected profit for choice situation i in year t is then given by:

$$\Delta CS_{it} = \frac{1}{\hat{\alpha}_f} \left[\ln \left(\sum_j \exp(\tilde{\psi}_{fjm}) \right) - \ln \left(\sum_j \exp(\hat{\psi}_{fjm}) \right) \right] \quad (17)$$

where, again, $f = f[i]$ denotes the farmer associated with choice situation i . Averaging over repeated draws from the vector of random coefficients provides an estimate of the change in surplus for each farmer from the elimination of all switching costs.

In addition to the main scenario discussed in the foregoing (no switching costs), we also evaluate the counterfactual situations where only brand switching costs are removed, or only technology switching costs are dropped. The welfare results are reported in Table 8. The first row provides the mean change in equilibrium prices. Overall, we find that the existence of switching costs has a significant impact on prices. Over all markets, the elimination of switching costs would result in a reduction of seed prices by 6.01%. Virtually all of this price effect is due to the removal of brand switching costs, as trait switching costs have a negligible impact on equilibrium prices.

The next four rows of Table 8 contain summary statistics for the change in farmer surplus. Both individual specific factors (e.g., initial conditions) and the random coefficients imply a different welfare effect for each farmer. Thus, Table 8 reports the mean of the distributions across all choice

situations during the sample period 1997–2016, along with the 5th percentile change, median change, and 95th percentile change. On average, eliminating all switching costs would result in an increase in farmers’ expected profit of \$7.69/acre, but there is significant heterogeneity in these impacts. The bottom 5% of farmers receive an increase of \$1.74/acre or less, whereas the upper 5% of farmers experience an increase in surplus of \$13.38/acre or more.

Table 8. Welfare Change from Removing Switching Costs, 1997–2016

| | Remove All Switching Costs | Remove Trait Switching Costs | Remove Brand Switching Costs |
|--|----------------------------------|------------------------------------|------------------------------------|
| Mean price change (percent) | -6.01% | 0.10% | -5.91% |
| Buyers’ unit surplus change (\$/acre) | | | |
| Mean | 7.69 | 0.41 | 7.15 |
| 5 th percentile | 1.74 | 0.00 | 1.60 |
| median | 7.57 | 0.24 | 6.99 |
| 95 th percentile | 13.38 | 1.22 | 12.29 |
| Buyers’ total surplus change (\$ million/year) | 463.5 | 26.1 | 428.8 |
| Industry total profit change (\$ million/year) | -225.9 | 3.5 | -226.8 |

Note: These estimates are calculated as per the counterfactual experiments described in the text. Entries in this table are average effects, over the period 1997–2016, expressed in \$/acre (measured in 2011 dollars).

The last two columns of Table 8 summarize the welfare impacts from removing each type of switching cost in isolation. Removal of brand switching costs has the largest impact on per-acre surplus, which increases on average by of \$7.15/acre, whereas removing trait switching costs only produces an average increase in surplus of \$0.41/acre. At first sight, this result may seem surprising given that trait switching costs are quite large in magnitude (Table 6), but it can be explained by GT adoption patterns and the rather large difference in value between the two traits. Specifically, removal of switching costs only affects an individual’s expected profit for products they did not purchase in the previous period. Because over 90% farmers purchased the GT trait for the majority of the sample, it was seldom the case they faced a switching cost associated with the purchase of a GT variety. Such costs were only present early on in the sample, and in fact most of the observed

increase in expected profit from removing trait switching costs occurred in the first period (1997–2002). It is, however, fairly common in the data for a farmer not to have purchased a CV product in the previous period, and thus removal of switching costs significantly increased the expected profit associated with CV products. But because the expected profit associated with GT varieties significantly exceeded the expected profit for CV products, the removal of such costs ultimately had little impact on overall expected welfare, at least for most farmers.

More generally, the results in Table 8 highlight an interesting feature of markets where switching costs are present. Switching costs only have significant welfare implications to the extent that: (a) a significant number of individuals do not purchase certain products regularly; and, (b) the difference in value between different products is not too large.

Table 8 also reports the total change in estimated surplus due to the elimination of switching costs, which is obtained by aggregating over all plots for each year of the 1997–2016 period. The effects are large, amounting to a total surplus gain by farmers of about \$464 million per year. The estimated market-level changes in prices and market shares from the counterfactuals also permits us to characterize the impact on seed sellers from eliminating switching costs. As noted, the counterfactual with no switching costs results in a price decline, on average, of 6.0%. Because we are using a conditional demand model, this is also the average percent decline in seed industry profits under the no-switching-cost scenario. When aggregated across the market, this amounts to a profit loss of about \$226 million per year.

6.4 Technology and brand shares

Beyond the welfare implications of switching costs for buyers, and for the equilibrium profits of firms, it is also of some interest to consider the consequences of switching costs for technology adoption and brand shares. This analysis is motivated by major changes that have taken place in the industry. In particular, adoption of the GT variety was relatively rapid, exceeding 90% of planted acres by 2007. However, in the final years of the sample, GT variety adoption declined somewhat, ostensibly due to weed resistance issues. A growing literature shows that glyphosate weed resistance is increasingly problematic because it significantly reduces the effectiveness of glyphosate as a weed control tool for GT crops. It is therefore instructive to consider whether farmers would have more readily switched back to conventional varieties in the absence of

switching costs. With respect to brand shares, the industry has become more concentrated over time. Yet, how switching costs impact brand shares and concentration has received scant attention. For each of the three scenarios considered—absence of all switching costs, absence of brand switching costs only, or absence of technology switching costs only—the structure of the estimated model permits calculation of counterfactual market shares. These shares differ from the baseline not only because of the removal of switching costs, but also because counterfactual prices differ from baseline prices. As was the case with the welfare calculation, simulation is required for the computation of these probabilities, which are then used to construct expected trait adoption under the baseline and counterfactual scenario.

Figure 1. Effects of Switching Costs on GT Technology Adoption




Figure 1 displays aggregate adoption rates for the GT trait under the baseline scenario and the counterfactual without switching costs. The primary finding that emerges from these predictions is that GT adoption, with exception of the first few years, would have been significantly lower in the absence of (all) switching costs. For example, in the final period we find that, without switching costs, conventional soybeans would have accounted for about 11% of acres, nearly double the rate under switching costs. These effects highlight the deep connection between switching costs and purchasing patterns. In this simple case of two technologies, the technology with the majority share benefits from switching costs, whereas the technology with the smaller share is penalized. The GT

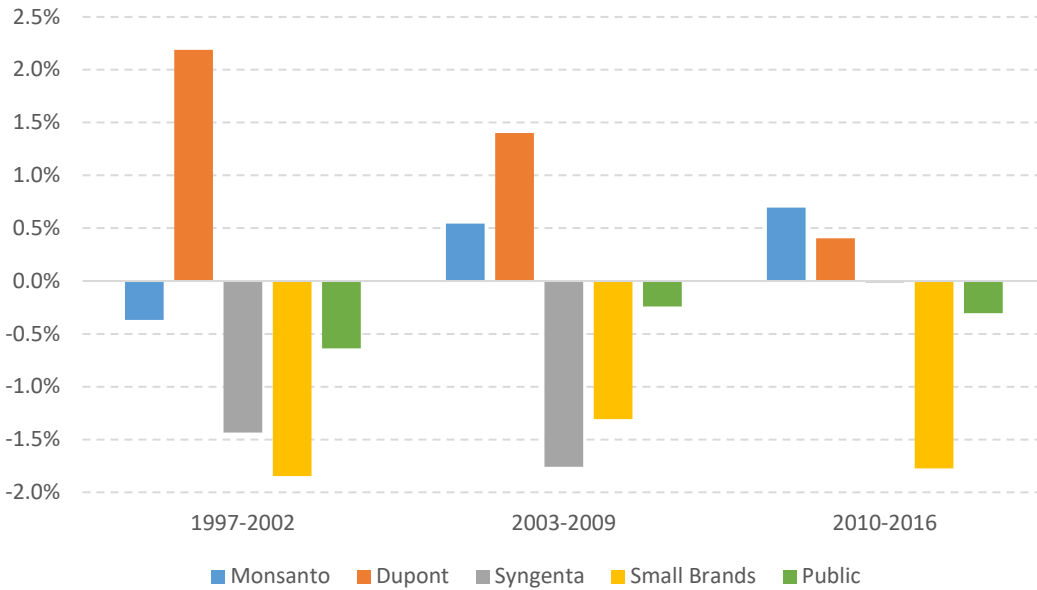
technology, having secured a larger customer base earlier in our study period, disproportionately benefited from the existence of switching costs. Analysis of the other two scenarios considered (absence of brand switching costs only, or absence of technology switching costs only) indicates that the pattern of Figure 1 is entirely attributable to technology switching costs: predicted GT adoption rates when only brand switching costs are dropped are virtually the same as the observed adoption rates of the baseline.

Next, we investigate how the existence of switching costs affect firms' market shares. From the estimated structural model, we compute product-level and brand-level predicted market shares for both baseline and counterfactual prices. These results are further aggregated to the firm (parent company) level, given the presumed structure of the profit maximization problem. The results are summarized in Figure 2. The individual parent companies that we explicitly report are Monsanto (Asgrow, Dekalb, and Channel brands), Syngenta (NK and Golden brands), and Dupont (Pioneer brand). The category "Small Brands" aggregates all small brands the model represents individually (Mycogen, Becks, Stine, Growmark, Croplan, and Kruger brands). We omit the "Other" group that appeared in the estimated model (an aggregate consisting of many minor brands that was defined mainly for the purpose of keeping estimation tractable). Because many of the estimated parameters were allowed to vary over time, the effects of switching costs on brand shares is reported for each of the three sub-periods considered in the model. Specifically, in Figure 2 the variable on the y-axis is the change in market share attributable to switching costs, as measured in actual percentage points.

The most salient feature that emerges from this figure is how the two largest companies—Monsanto and Dupont—have been differentially impacted by the existence of switching costs. Dupont, the incumbent firm with the largest market share at the dawn of the GE era, benefited considerably from switching costs in the earlier years. Monsanto, the upstart that developed the GE trait but needed to acquire a position in the seed market to market this revolutionary technology, was initially penalized by the existence of switching costs. Over time, however, Dupont's windfall benefits arising from switching costs dwindled, and Monsanto gradually benefited as the GT technology became dominant. In the last subperiod, Monsanto is the firm that benefited the most, in terms of market shares, from the existence of switching costs. In addition, because Monsanto also widely licensed the GT trait, the results presented in Figure 2 indicate that switching costs also indirectly benefited them through a larger share for the GT trait. Hence, the

results in Figures 1 and 2 together indicate that Monsanto is the company that benefited the most from switching costs.

Figure 2. Effects of Switching Costs on Market Shares



Consistent with expectations, Figure 2 also shows that the small brands individually considered in the model were negatively affected by the existence of switching costs, and the same applies to “public” varieties. Syngenta, the smallest of the large soybean seed companies, was negatively affected by switching costs on average, although it seems the impact was confined to the first two sub-periods considered.

7. Conclusion

In this paper, we add to the empirical study of inertia and switching costs in imperfectly competitive markets by contributing new evidence in a novel context using a unique dataset. Specifically, we develop and estimate a structural model of US soybean seed demand, by competitive farmers, over a period characterized by the introduction, and virtual complete adoption of a major new technology—genetically engineered GT varieties. A random coefficients logit model, estimated with an extensive and unique individual-level dataset, provides insights into the role of inertia and individual heterogeneity, and permits a quantitative assessment of the value

associated with the novel GT trait, as well as the extent of brand and technology switching costs. To our knowledge, this is one of the first studies to assess the welfare and share implications of switching costs in a production context.

The results show that structural state dependence is significant. The average WTP measure associated with switching costs is \$13.1/acre for brand, and \$6.8/acre or \$14.7/acre for technology switching costs (conventional to GT, and GT to conventional, respectively). We also find that state dependence is highly heterogeneous, with considerable variation across farmers. Farmers' WTP for the GT trait is sizeable, and varies over time, as well as over individuals. On average, the WTP for GT is \$18.6/acre during 1997–2002, goes up to \$24.8/acre during 2003–2009, and then declines to \$21.8/acre during 2010–2016. These results are similar in magnitude to the WTP estimates reported by Ciliberto, Moschini, and Perry (2019), despite clear differences in the modeling approach and period of analysis. Our finding of declining WTP for the GT trait in the last sub-period (2010–2016), which is not considered by the aforementioned study, is consistent with a declining soybean commodity price (a shifter of the estimated input demands), as well as the recent emergence of glyphosate-resistant weeds. We further show that farmers are heterogeneous in their valuation of the GT trait.

Estimated own- and cross-price demand elasticities show that seed input demand is highly elastic, and that farmers are more likely to substitute among products of the same brand and same trait. The existence of switching cost implies less elastic demand, as expected, and thus provides the opportunity for higher price premiums by oligopolistic firms.

The counterfactual analyses highlight several implications of switching costs. First, switching costs significantly reduce consumer welfare: on average, farmers' surplus increases by \$7.69/acre upon removing switching costs, amounting to \$464 million per year. Some of this surplus gain is due to the price effects uncovered in the counterfactual: on average, seed prices decrease by 6.0% from the elimination of switching costs. This price effect means that industry profits are positively impacted by switching costs (in aggregate, \$226 million per year). The majority of these welfare effects are due to the existence of brand switching costs, rather than technology switching costs. Indeed, most of the welfare increase from removing technology switching costs occurs in the first sub-period when GT trait adoption was still well below 90%.

Our model predicts that, absent switching costs, GT trait adoption would have been higher in the first few years during the initial diffusion phase, and lower thereafter. This is noteworthy because, as noted earlier, the emergence of glyphosate weed resistance in recent years has likely reduced the attractiveness of GT soybeans. Although the data indicate that some farmers have switched back to conventional soybeans, it seems that switching costs have dampened this adjustment process. From a policy standpoint, insofar as switching costs are a consequence of imperfect information, existing programs (e.g., state university extension services) should prioritize agricultural technologies with significant switching costs. Conventional soybeans are one such example, but other examples include cover cropping and reduced tillage. Ultimately, more information is important for reducing the extent of switching costs, even in the context of producer behavior, where economic theory typically assumes that profit maximization eventually eliminates persistent sources of inertia.

It is also of some interest to contrast these results with some of the marketing literature on switching costs. For example, Grzybowski and Nicolle (2021) find that the presence of switching costs in the smartphone industry benefitted the smaller IOS platform and hurt the more widely adopted Android platform. By contrast, our estimates indicate that it was in fact the more adopted technology (GT) that benefitted from switching costs. These differences can be ascribed to the differing asymmetries in switching costs. In our context, it is significantly more expensive to switch from GT to CV soybeans (by contrast, Grzybowski and Nicolle (2021) find that it was easier to switch to IOS than to Android).

Lastly, we find that switching costs tend to increase predicted shares for the largest brands and decrease predicted shares for smaller brands. Furthermore, we find that switching costs protected the market share of the large firm that did not innovate (Dupont), while it constituted an obstacle for the innovating firm in the early years (Monsanto). These effects wore out towards the end of the sample, as the innovative GE technology permitted Monsanto to overcome the initial disadvantage. The broader question of how switching costs relate to the exercise of market power remains an open question. The seed industry has been affected by considerable consolidation in the last few decades, and, as we have shown, seed demand exhibits significant switching costs. We find that the existence of switching costs, per se, has negatively affected the market share of small firms, thus possibly contributing to this consolidation.

References

- Akerberg, D., Benkard, C.L., Berry, S. and Pakes, A., 2007. Econometric tools for analyzing market outcomes. *Handbook of Econometrics*, Volume 6-A, 4171-4276. Elsevier.
- Barone, G., Felici, R. and Pagnini, M., 2011. Switching costs in local credit markets. *International Journal of Industrial Organization*, 29(6), 694-704.
- Barrows, G., Sexton, S., Zilberman, D. 2014. Agricultural biotechnology: The promise and prospects of genetically modified crops. *Journal of Economic Perspectives* 28; 99-119.
- Berry, S.T. 1994. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics* 242-262.
- Berry, S., Levinsohn, J., Pakes, A. 1995. Automobile prices in market equilibrium. *Econometrica* 841-890.
- Berto Villas-Boas, S. 2007. Vertical relationships between manufacturers and retailers: Inference with limited data. *The Review of Economic Studies* 74(2); 625-652.
- Bronnenberg, B.J., Dhar, S.K., Dubé, J.P.H. 2009. Brand history, geography, and the persistence of brand shares. *Journal of Political Economy* 117(1); 87-115.
- Bronnenberg, B. J., Dubé, J. P. H., Gentzkow, M. 2012. The evolution of brand preferences: Evidence from consumer migration. *American Economic Review* 102(6); 2472-2508.
- Bronnenberg, B. J., Dubé, J.-P., Moorthy, S. 2019. The Economics of Brands and Branding. In: Dubé, J. P. and Rossi, P.E., (Eds), *Handbook of the Economics of Marketing*. Elsevier B.V.: Amsterdam. ch. 6, p. 291-358.
- Cabral, L., 2016. Dynamic pricing in customer markets with switching costs. *Review of Economic Dynamics*, 20, 43-62.
- Ciliberto, F., Moschini, G., Perry, E. 2019. Valuing product innovation: Genetically engineered varieties in US corn and soybeans. *The RAND Journal of Economics* 50(3); 615-644.
- Clancy, M.S., Moschini, G. 2017. Intellectual property rights and the ascent of proprietary innovation in agriculture. *Annual Review of Resource Economics* 9; 53-74.
- Conlon, C. and Gortmaker, J., 2020. Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4); 1108-1161.
- Daly, A., Hess, S., Train, K. 2012. Assuring finite moments for willingness to pay in random coefficient models. *Transportation* 39(1); 19-31.
- Dubé, J.-P., Hitsch, G.J., Rossi, P.E. 2010. State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics* 41(3); 417-445.
- Dubé, J.-P., Hitsch, G.J., Rossi, P.E. 2009. Do switching costs make markets less competitive? *Journal of Marketing Research* 46(4); 435-445.
- Farrell, J., Klemperer, P. 2007. Coordination and lock-in: Competition with switching costs and network effects. In: Armstrong, M., and Porter, R., (Eds), *Handbook of industrial organization*, v3. Elsevier B.V.: Amsterdam. p. 1967-2072.

- Fernandez-Cornejo, J. 2004. The seed industry in US agriculture: An exploration of data and information on crop seed markets, regulation, industry structure, and research and development. United States Department of Agriculture, Economic Research Service.
- Gandhi, A. and Nevo, A., 2021. Empirical models of demand and supply in differentiated products industries. In *Handbook of Industrial Organization*, Vol. 4, No. 1, pp. 63-139). Elsevier.
- Gaspar, A.P., Mourtzinis, S., Kyle, D., Galdi, E., Lindsey, L.E., Hamman, W.P., Matcham, E.G. et al. 2020. Defining optimal soybean seeding rates and associated risk across North America. *Agronomy Journal* 112(3); 2103-2114.
- Goldberg, P. 1995. Product differentiation and oligopoly in international markets: The case of the US automobile industry. *Econometrica* 63(4); 891-951.
- Graff, G. D., Rausser, G. C., Small, A. A. 2003. Agricultural biotechnology's complementary intellectual assets. *Review of Economics and Statistics* 85(2); 349-363.
- Grzybowski, L., Nicolle, A. 2021. Estimating consumer inertia in repeated choices of smartphones. *The Journal of Industrial Economics* 69(1); 33-82.
- Handel, B.R. 2013. Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review* 103(7); 2643-82.
- Heckman, J.J. 1981a. Heterogeneity and State Dependence. In Rosen, S., (Ed), *Studies in labor markets*. Chicago: Chicago University Press. p. 91-140.
- Heckman, J.J. 1981b. The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process and some Monte Carlo evidence. In: Manski, C.F., and McFadden D.L. (Eds), *Structural analysis of discrete data with econometric applications*. The MIT Press: Cambridge; p. 179-195.
- Hole, A.R. 2007. Fitting mixed logit models by using maximum simulated likelihood. *The Stata Journal* 7(3); 388-401.
- Hole, A.R., Kolstad, J.R. 2012. Mixed logit estimation of willingness to pay distributions: A comparison of models in preference and WTP space using data from a health-related choice experiment. *Empirical Economics* 42(2); 445-469.
- Honka, E., 2014. Quantifying search and switching costs in the US auto insurance industry. *The RAND Journal of Economics*, 45(4), 847-884.
- International Service for the Acquisition of Agro-biotech Applications (ISAAA). 2016. *Global Status of Commercialized Biotech/GM Crops: 2016*. Brief 52.
- Janssen, A. 2022. Switching costs, brand premia and behavioral pricing in the pharmaceutical market. *Working Paper*. <https://aljoschajanssen.com/research/>
- Keane, M.P. 1997. Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics* 15(3); 310-327.
- Kim, H., Moschini, G. 2018. The dynamics of supply: US corn and soybeans in the biofuel era. *Land Economics* 94(4); 593-613.
- Kim, K., Petrin, A. 2019. Control function corrections for unobserved factors in differentiated product models. *Working paper*. <https://sites.google.com/site/kyookim/>.

- Lamkey, K. 2004. Seed production in corn and soybean. *Agronomy Reports*. <https://dr.lib.iastate.edu/handle/20.500.12876/5077>.
- MacKay, A., Remer, M. 2022. Consumer inertia and market power. *Working Paper*. <http://dx.doi.org/10.2139/ssrn.3380390>.
- Morrow, W.R. and Skerlos, S.J., 2011. Fixed-point approaches to computing Bertrand-Nash equilibrium prices under mixed-logit demand. *Operations Research*, 59(2), 328-345.
- Moschini, G. 2008. Biotechnology and the development of food markets: Retrospect and prospects. *European Review of Agricultural Economics* 35(3); 331-355.
- Moschini, G. 2010. Competition issues in the seed industry and the role of intellectual property. *Choices* 25(2); 1-12.
- Nevo, A. 2000. A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of Economics & Management Strategy* 9(4); 513-548.
- Nevo, A., 2001. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2), 307-342.
- Organization for Economic Co-operation and Development (OECD). 2018. *Concentration in seed markets: Potential effects and policy responses*. Paris: OECD.
- Pakes, A. 2003. A reconsideration of hedonic price indexes with an application to PC's. *American Economic Review* 93(5); 1578-1596.
- Perry, E.D., Ciliberto, F., Hennessy, D.A., Moschini, G. 2016. Genetically engineered crops and pesticide use in US maize and soybeans. *Science Advances* 2(8); e1600850.
- Perry, E.D., Moschini, G., Hennessy, D.A., 2016. Testing for complementarity: Glyphosate tolerant soybeans and conservation tillage. *American Journal of Agricultural Economics* 98(3); 765-784.
- Perry, E.D., Hennessy, D.A., Moschini, G. 2022. Uncertainty and learning in a technologically dynamic industry: Seed density in US maize. *American Journal of Agricultural Economics*. DOI: 10.1111/ajae.12276.
- Petrin, A., Train, K. 2010. A control function approach to endogeneity in consumer choice models. *Journal of Marketing Research* 47(1); 3-13.
- Polyakova, M. 2016. Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D. *American Economic Journal: Applied Economics* 8(3); 165-95.
- Rickert, D., 2022. When switching costs benefit consumers. Empirical Evidence from a Market with Lock-in Effects. Unpublished manuscript, March 9. <https://drive.google.com/file/d/1-N781e3iMCh5EOf47UGCafIwxTfxzH9P/view>
- Shum, M., 2004. Does advertising overcome brand loyalty? Evidence from the breakfast-cereals market. *Journal of Economics & Management Strategy*, 13(2), 241-272.
- Sudhir, K., Yang, N. 2014. Exploiting the choice-consumption mismatch: A new approach to disentangle state dependence and heterogeneity. Cowles Foundation *Discussion Paper* No. 1941.

- Syngenta. 2016. *Our industry 2016*. <https://www.scribd.com/document/317849502/Syngenta-Industry-Whats-Next> on Apr 3, 2020.
- Train, K.E. 2009. *Discrete choice methods with simulation*. Cambridge: Cambridge University Press.
- Train, K.E., Winston, C. 2007. Vehicle choice behavior and the declining market share of US automakers. *International Economic Review* 48(4); 1469-1496.
- Wechsler, S.J., McFadden, J.R., Smith, D.J. 2018. What do farmers' weed control decisions imply about glyphosate resistance? Evidence from surveys of US corn fields. *Pest Management Science* 74(5); 1143-1154.
- Wooldridge, J.M. 2005. Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics* 20(1); 39-54.
- Yeo, J. and Miller, D.P., 2018. Estimating switching costs with market share data: An application to medicare part D. *International Journal of Industrial Organization*, 61, 459-501.

Appendix

Table A1. Soybean Seed Purchase Data (1996–2016)

| Row number | Description | No. |
|------------|---|------------------|
| 0 | Farmers surveyed | 28,017 |
| 1 | Original soybean seed purchase choices | 213,062 |
| 2 | Drop purchase records if purchase source is not a standard commercial channel | 6,890 |
| 3 | Drop purchase records if the seed is not newly purchased | 1,367 |
| 4 | Drop purchase records if the product is “Public” with GT trait | 108 |
| 5 | Choices in analysis | 204,697 |
| 6 | Drop purchase records with zero net prices | 831 |
| 7 | Drop purchase records in markets with only one alternative | 673 |
| 8 | Drop purchase records of all farmers' one-year purchase when previous year's purchase record for the farmer is not observed | 116,995 |
| 9 | Choices in regression | 86,198 |
| 10 | Observations in regression | 1,043,742 |

Note: This table explicitly documents the steps taken in the data-cleaning process of the farm-level soybean seed purchase data (data source: Kynetec USA, Inc.). In particular:

- (a) Observations in line 1 include all distinct purchases in the original data. Note that a farmer may report more than one distinct choice (i.e., different varieties) in any given year.
- (b) Observations dropped in line 2 pertain to seed choices when the source is: “From my own farm,” “I’m a seed grower,” or “New seed that was left over from last year.”
- (c) Observations dropped in line 3 pertain to “saved seeds.”
- (d) Observations dropped in line 4 pertain to miscoded entries (no publicly released varieties containing the GT trait have been marketed).
- (e) Observations dropped in lines 6 and 7 do not affect an individual’s purchase history.
- (f) Observations dropped in line 8 do not affect the set of alternatives in any one market.
- (g) Line 8 ensures that all observations in year t remaining in the estimation sample pertain to farmers who also have documented purchases in year $t-1$.
- (h) Line 10 is fully equivalent to line 9, it simply expands each choice by available products in each local market.

Table A2. Descriptive Information on the Estimation Data

| Variable | Mean | Std. Dev. | Min | Max |
|--|------|-----------|-----|-----|
| No. of alternatives per market | 8.15 | 4.48 | 2 | 23 |
| No. of chosen brands per farmer | 2.10 | 1.70 | 1 | 10 |
| No. of chosen brands per farmer per year | 1.58 | 1.02 | 1 | 8 |
| No. of chosen products per farmer | 2.36 | 2.29 | 1 | 17 |
| No. of chosen products per farmer per year | 1.65 | 1.11 | 1 | 9 |
| No. of recorded years per farmer | 2.74 | 2.58 | 1 | 20 |

Note: Authors' computations from Kynetec data. Calculation based on the estimation sample, as per line 9 of Table A1 in the Appendix.

Table A3. Conditional Logit Estimation Results

| | (1) | (2) | (3) | (4) |
|------------------------------------|---------------------|---------------------|--------------------------|---------------------|
| Price (\$/acre) | -0.275** (0.017) | 0.001 (0.001) | -0.170** (0.016) | -0.153** (0.016) |
| Control Residual | 0.276** (0.017) | | 0.171** (0.016) | 0.155** (0.016) |
| GT (1997–2002) | 4.884** (0.244) | 0.426** (0.050) | 2.786** (0.236) | 2.658** (0.232) |
| GT (2003–2009) | 5.644** (0.242) | 1.268** (0.064) | 3.369** (0.228) | 3.227** (0.223) |
| GT (2010–2016) | 3.938** (0.181) | 0.756** (0.080) | 2.233** (0.167) | 2.111** (0.168) |
| Switching Cost Coefficients | | | | |
| Brand | | -2.667** (0.044) | -3.055** (0.047) | -2.665** (0.044) |
| GT | | -1.630** (0.065) | -1.707** (0.066) | -1.619** (0.066) |
| Non-GT | | -2.370** (0.063) | -2.529** (0.067) | -2.359** (0.063) |
| Brand×GT | | 0.566** (0.074) | 0.722** (0.074) | 0.565** (0.075) |
| Brand×Non-GT | | 1.178** (0.083) | 1.305** (0.085) | 1.187** (0.083) |
| Description | No Switch Costs | No IV | No Initial Conditions | Full |
| LL | -145,681 | -118,041 | -119,449 | -117,953 |

Note: This table reports the estimation results of four conditional logit model specifications, with differing levels of generality (as detailed in the “Description” row). In addition to the coefficients of the variables reported in this table, all models also include brand indicators, brand initial conditions, and interactions of the brand trait indicators with the first two sub-periods. Standard errors, in parentheses, are clustered at the CRD level. $N=1,043,742$. * $p < 0.05$, ** $p < 0.01$

Table A4. Additional Estimated Mixed Logit Parameters: Brand-Time Effects

| Variable | Coefficient | Standard error | | Coefficient | Standard error |
|---------------------------|-------------|----------------|---------------------------|-------------|----------------|
| <u>Brands (1997-2002)</u> | | | <u>Brands (2003-2009)</u> | | |
| Asgrow | -2.016** | (0.225) | Asgrow | -2.779** | (0.278) |
| Beck | -2.277** | (0.317) | Beck | -2.985** | (0.297) |
| Croplan | -1.985** | (0.234) | Croplan | -2.690** | (0.301) |
| Dekalb | 0.132 | (0.210) | Dekalb | -0.655** | (0.252) |
| Golden | -0.226 | (0.286) | Golden | -1.157** | (0.324) |
| Growmark | -1.432** | (0.244) | Growmark | -2.085** | (0.304) |
| Kruger | -2.316** | (0.303) | Kruger | -2.570** | (0.352) |
| Mycogen | -1.257** | (0.236) | Mycogen | -2.164** | (0.279) |
| NK | -2.105** | (0.224) | NK | -2.340** | (0.293) |
| Other | -1.709** | (0.228) | Other | -2.103** | (0.270) |
| Pioneer | -2.021** | (0.213) | Pioneer | -2.133** | (0.248) |
| Stine | -0.940** | (0.198) | Stine | -1.919** | (0.236) |
| LL | -110,364.4 | | | | |
| N | 1,043,742 | | | | |

Note: This table reports the estimated parameters of the mixed logit model not already included in Table 5 in the text. Estimates for Brands (1997-2002) and for Brands (2003-2009) pertain to shifters of the mean brand coefficients reported in Table 5. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$

Table A5. Price Equation for Control Function

| Variable | Coefficient | Standard error | Variable | Coefficient | Standard error |
|--------------------|-------------|----------------|---------------------|-------------|----------------|
| GT (1997–2002) | 11.520** | (1.122) | GT (1997–2002)×Cost | 0.286** | (0.109) |
| GT (2003–2009) | 10.733** | (1.378) | GT (2003–2009)×Cost | 0.205 | (0.145) |
| GT (2010–2016) | 4.007 | (2.450) | GT (2010–2016)×Cost | 0.433* | (0.221) |
| Asgrow | 13.043** | (2.999) | Asgrow×Cost | 0.488 | (0.257) |
| Beck | 6.502 | (3.997) | Beck×Cost | 1.059** | (0.351) |
| Channel | 8.615* | (4.364) | Channel×Cost | 0.817* | (0.379) |
| Croplan | 15.834** | (3.330) | Croplan×Cost | 0.317 | (0.286) |
| Dekalb | 8.007* | (3.333) | Dekalb×Cost | 0.449 | (0.275) |
| Golden | 9.513** | (3.603) | Golden×Cost | 0.703* | (0.285) |
| Growmark | 11.412** | (3.512) | Growmark×Cost | 0.672* | (0.301) |
| Kruger | 11.079** | (3.373) | Kruger×Cost | 0.618* | (0.287) |
| Mycogen | 6.321 | (3.331) | Mycogen×Cost | 0.821** | (0.287) |
| NK | 8.161** | (3.110) | NK×Cost | 0.879** | (0.267) |
| Other | 10.743** | (2.901) | Other×Cost | 0.469 | (0.249) |
| Pioneer | 10.378** | (2.909) | Pioneer×Cost | 0.527* | (0.250) |
| Stine | 6.785* | (3.131) | Stine×Cost | 0.706** | (0.271) |
| Market Size | 3.058** | (0.409) | | | |
| Lagged Brand Share | -0.945 | (0.592) | | | |
| Lagged Trait Share | -0.287 | (0.643) | | | |
| Constant | 22.753** | (1.732) | | | |

Note: This table reports the estimation results for the price equation (8) used for the control function. “Cost” is the IV cost shifter discussed in the text. In addition to the listed variables, the model also includes time and CRD fixed effects, as well as the brand indicators interacted with the first two sub-periods. Standard errors, in parentheses, are clustered at the CRD level.

$N=1,043,742$. * $p < 0.05$, ** $p < 0.01$