

**Genetically engineered varieties in the U.S. maize and soybean seed markets: Production
impacts, environmental implications, and welfare effects**

by

Edward D. Perry

A dissertation submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Major: Economics

Program of Study Committee:
GianCarlo Moschini, Co-Major Professor
David Hennessy, Co-Major Professor
Otavio Bartalotti
John Beghin
John Schroeter

Iowa State University

Ames, Iowa

2016

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	iii
CHAPTER 1 INTRODUCTION	1
CHAPTER 2 TESTING FOR COMPLEMENTARITY: GLYPHOSATE TOLERANT SOYBEANS AND CONSERVATION TILLAGE.....	5
Introduction	6
Modeling Complementarity	10
Data	22
Empirical Results	27
Conclusion	34
References	36
Appendix	41
CHAPTER 3 GENETICALLY ENGINEERED CROPS AND PESTICIDE USE IN U.S. MAIZE AND SOYBEANS.....	47
Introduction.....	48
Results	49
Discussion and Conclusion	54
Materials and Methods	55
References	68
Appendix	71
CHAPTER 4 A DISCRETE CHOICE MODEL OF SEED DEMAND: GENETICALLY ENGINEERED VARIETIES IN U.S. CORN AND SOYBEANS	99
Introduction.....	100
Data.....	105
Demand Model	112
Results	117
Conclusion.....	123
References.....	125

ACKNOWLEDGEMENTS

First and foremost, I would like thank my advisors GianCarlo Moschini and David Hennessy. Many hours have been spent in GianCarlo's office discussing the ideas and content of this dissertation. His guidance and advice have been invaluable, and I cannot thank him enough for all of the time and support he has provided during my time at Iowa State University. I am also grateful to David for his encouragement and advice in writing this dissertation. I would also like to thank Federico Ciliberto. I owe much to him for his feedback and advice on the third and fourth chapters of this dissertation.

I am also indebted to the rest of my committee – Otavio Bartalotti, John Beghin, Brent Kreider, and John Schroeter - for their time and helpful comments.

Finally, I am grateful to my family, in particular my parents, for their unwavering support, guidance, and encouragement.

CHAPTER 1

INTRODUCTION

Since their commercial introduction in 1996, genetically engineered (GE) corn and soybean varieties have become the primary type of seed planted by U.S. farmers. In every year since 2007, the share of land planted to GE corn and soybean varieties has exceeded 75% and 90%, respectively. Despite such popularity among farmers, GE crops have been surrounded by controversy. Concerns raised have ranged from the safety of GE crops for consumption to the market power possessed by the large firms that develop and market GE crops. Amidst these concerns a large body of research has arisen that attempts to ascertain the many consequences of GE crops. Unsurprisingly, this research has shown that there are both benefits and costs associated with their adoption.

This dissertation contributes to the literature in the form of three essays that explore how the widespread adoption of GE varieties has impacted three different facets of the agricultural landscape. The first two essays investigate the impact of GE variety adoption on two related farming practices: conservation tillage adoption and pesticide use. These two practices have not only served an instrumental role in farming, but they also have significant implications for the environment. Conservation tillage has long been advocated because it reduces soil erosion and chemical runoff, and pesticide use has implications for human health and biodiversity. The third essay focuses on the monetary benefits for the farmers that adopt GE varieties. Although GE varieties have been adopted quickly and loyally, consolidation and sharply rising prices in the corn and soybean seed markets have prompted many to question whether farmers have continued to benefit. We investigate these concerns by estimating farmers' willingness to pay for GE varieties in the U.S. corn and soybean seed markets. In what follows, I discuss each of these three essays in detail.

The adoption of conservation tillage in soybean farming has been historically limited by the inability of conventional herbicide programs to fully substitute for tillage as a means of weed control. Glyphosate tolerant soybeans, introduced in 1996, potentially solved this problem by allowing for more effective post-emergent weed control. In the first essay, "Testing for

Complementarity: Glyphosate Tolerant Soybeans and Conservation Tillage”, we explore this possibility by developing a discrete choice model of joint technology adoption in which farmers choose the most profitable option from a set of four mutually exclusive soybean-tillage systems: (i) conventional soybeans and intensive tillage, (ii) glyphosate tolerant soybeans and intensive tillage, (iii) conventional soybeans and conservation tillage, (iv) glyphosate tolerant soybeans and conservation tillage. Within this framework we develop a test, based on the theory of supermodularity, for whether glyphosate tolerant soybeans and conservation tillage are complementary practices.

This essay contributes to the literature in two significant ways. First, the model is estimated with a large, unbalanced panel of 82,056 soybean and tillage choices over the 1998-2011 period. This is in contrast to previous analyses which have had to rely on either state-level data, or individual data that spanned a single year. Second, the developed model possesses two properties essential for the identification of complementarity between glyphosate tolerant soybeans and conservation tillage. The first property is that the decision problem for a given farmer is framed as consisting of a single choice between all four possible soybean-tillage systems. Previous work has modeled farmers as making two distinct, albeit possibly correlated, choices. The limitation of this approach is that complementarity is either ruled out from the beginning (as is the case in the bivariate probit model, e.g.), or poorly characterized. The second property that the developed model possesses is that it distinguishes between the correlation induced by unobserved and correlated tastes versus the correlation that results from genuine structural complementarity. Existing work ignores this distinction, and thus potentially incorrectly concludes in favor of (or against) complementarity.

We find that glyphosate tolerant soybeans and conservation tillage are complementary practices, and this finding is robust to a variety of alternative specifications. In the baseline specification, the marginal benefit of the complementarity effect is about \$1.69 per acre. In addition, the results indicate conservation tillage is more likely to be adopted on large farms, highly erodible land, in drought-like conditions, and when fuel prices are high. The estimates are also used to simulate a counterfactual in which glyphosate tolerant soybeans were not available as a choice. The simulation indicates that the availability of glyphosate tolerant

soybeans led to an increase in the adoption of conservation tillage and no-tillage by about 10% and 20%, respectively.

In the second essay, titled “Genetically Engineered Crops and Pesticide Use in U.S. Maize and Soybeans”, we examine the relationship between the adoption of GE corn and soybeans and pesticide use. Since their introduction, there has been little doubt that GE varieties have impacted which types and how much pesticide farmers use, but the extent of those impacts as well as their environmental implications remain unclear. In the case of insecticides, the existing literature has attributed significant reductions to the adoption of insect resistant varieties. For herbicides, however, there is less agreement. This is in part owed to the more complex relationship between herbicide resistant crops and herbicide use. Whereas with insect resistant varieties there is a clear substitution effect, with herbicide tolerant crops farmers predictably increase the herbicide to which the crop is tolerant, but decrease the use of other herbicides. Perhaps more importantly, there has simply been a lack of good data with which to analyze these questions. Most studies have used aggregate data to compare average pesticide usage decisions between GE and non-GE adopters, and in many cases these studies have resorted to imputing non-GE adopter herbicide use by using recommended conventional herbicide programs that would achieve a level of weed control on par with GE variety herbicide programs. The problem with this approach is that the usage rates associated with these recommended programs significantly exceed observed usage rates prior to the GE era.

We estimate the impact of GE varieties on pesticide use with data on the seed and pesticide use decisions of 86,736 plots in soybeans and 134,264 plots in maize over the 1998-2011 period. We control for potentially important omitted variables by including grower-specific fixed effects, time fixed effects, and CRD-specific time trends. We find that glyphosate tolerant soybean and corn adopters used about 28% (0.30 kg/ha) more and 1.2% (0.03 kg/ha) less than non-adopters, respectively. Insect resistant corn adopters used about 11.2% (0.013) less insecticide than non-adopters. When weighted by an index called the environmental impact quotient these results are modified to 0%, -9.8% and -10.4%, respectively. Perhaps most interestingly, we find that the relationship between GE and non-GE adopters changed significantly over time. In particular, glyphosate tolerant corn and soybean adopters gradually

used more herbicide relative to non-adopters over time. We show that this trend is in part attributable to the emergence of glyphosate weed resistance.

The final essay, “A Discrete Choice Model of Seed Demand: Genetically Engineered Varieties in U.S. Corn and Soybeans”, develops a discrete choice model of seed demand and applies it to a dataset of corn and soybean seed purchases from 1996-2011. We then use the model to retrieve farmers’ willingness-to-pay (WTP) for GE varieties. Importantly, we permit the WTP estimates to structurally vary over the three sub-periods – 1996-2000, 2001-2006, and 2007-2011. These sub-periods correspond to the expiration of Monsanto’s glyphosate patent in 2000 and the sharp increase in corn and soybean output prices that began in 2007. There is strong reason to expect that both events altered the relative profitability of GE varieties. In particular, the expiration of Monsanto’s glyphosate patent was followed by a decline in the price glyphosate, one of the main inputs associate glyphosate tolerant corn and soybean varieties; and the increase in corn and soybean in output prices in 2007 raised the relative profitability of higher yielding, GE varieties. It is also worth noting that by allowing the estimates to vary over time, other dynamic factors that potentially affected farmers’ WTP are captured. Two examples include learning effects and the gradual release of more and better yielding GE varieties.

We find that corn and soybean growers were almost always willing to pay a premium for GE varieties, with the extent of that willingness increasing significantly over time. The results also show that the benefits associated with varieties that contain multiple GE traits are sub-additive (i.e., the premium a grower is willing to pay for two GE traits is less than the sum of the premiums he is willing to pay for each of the GE traits in isolation), a finding consistent with previous work in this area. Finally, a comparison of the WTP estimates to the actual premiums charged by seed firms suggests that farmers gained the most from GE varieties in the final sub-period, also the sub-period during which seed prices rose the most.

CHAPTER 2

TESTING FOR COMPLEMENTARITY: GLYPHOSATE TOLERANT SOYBEANS AND CONSERVATION TILLAGE

Edward D. Perry, GianCarlo Moschini, and David. A Hennessy

Abstract

Many decisions in agriculture are made over combinations of inputs and/or practices that may form a technology system linked through complementarity. The presence of complementarity among producer decisions can have far-reaching implications for market outcomes and for the effectiveness of policies intended to influence them. Identifying complementarity relations, however, is made difficult by the presence of unobserved heterogeneity. Drawing on recent methodological advances, in this paper we develop a test for complementarity between glyphosate tolerant soybeans and conservation tillage that overcomes certain limitations of previous studies. Specifically, we develop a structural discrete choice framework of joint soybean-tillage adoption that explicitly models both complementarity and the correlation induced by unobserved heterogeneity. The model is estimated with a large unbalanced panel of farm-level choices spanning the 1998–2011 period. We find that glyphosate tolerant soybeans and conservation tillage are complementary practices. In addition, our estimation shows that farm operation scale promotes the adoption of both conservation tillage and glyphosate tolerant seed, and that all of higher fuel prices, more droughty conditions and soil erodibility increase use of conservation tillage. We apply our results to simulate annual adoption rates for both conservation tillage and no-tillage in a scenario without glyphosate tolerant soybeans available as a choice. We find that the adoption of conservation tillage and no-tillage have been about 10% and 20% higher, respectively, due to the advent of glyphosate tolerant soybeans.

Introduction

Decision variables in many real-world problems are often best viewed as combined in clusters, e.g., bundles of goods or sets of practices. This clustering naturally arises when the payoff associated with the level of one variable is increasing in the level of another variable; that is, when they are complements. The underlying supermodular structure of the decision makers' objective function constitutes the essence of such situations (Milgrom and Roberts 1990). Complementary choices are ubiquitous and appear in consumption problems, production contexts, dynamic choices, and organizational design (Berry et al. 2014). They are relevant in an agricultural setting as well, where farmers' decisions increasingly pertain to choices of "systems" composed of alternative combinations of inputs or practices. For example, the choices of which crop to produce, what rotation to use and type of tillage to employ are often intertwined with mechanical equipment investments and the choices of an array of chemical inputs and genetics. An accurate characterization of such choices—that is, determining whether they form a technology system linked through complementarity—is crucial for both policy analysis and the evaluation of alternative hypotheses. Indeed, many policy interventions entail spillover effects and unintended consequences, which are often the result of unaccounted-for complementarities between targeted and other variables.

An open question in agriculture that can benefit from a focus on complementarity relates to the environmental implications of genetically engineered (GE) crop varieties. Since their introduction in 1996, GE crops have been both commercially successful and controversial (Moschini 2008). Environmental concerns have ranged from the possibility that adoption of GE crops facilitates monoculture to the detriment of desirable rotations, to the incentive that herbicide tolerant crops provide for the increased use of certain herbicides, and the risk of resistance build-up among the weeds and insects targeted by GE traits. Potential environmental benefits have been posited as well, however, such as a reduction in the use of certain insecticides and a reduction in agriculture's footprint (Barrows, Sexton and Zilberman, 2014). An additional important hypothesized impact, which at present remains unresolved, is that the

adoption of glyphosate tolerant (GT) crops induces the adoption of environmentally beneficial tillage methods.

Tillage is an important part of farming. It aids in seedbed preparation and has historically provided a critical means for weed control both before and after the crop has emerged (Givens et al. 2009). It has nonetheless been associated with several negative effects, including increased soil erosion (Blevins and Frye 2003), chemical runoff (Fawcett, Christensen and Tierney 1994), and the carbon footprint of agriculture (Kern and Johnson 1993; West and Marland 2002). Conservation tillage (CT), defined as a tillage system that leaves at least 30% of crop residues on the soil surface, has long been advocated as a way to reduce these detrimental effects (Holland 2004). Even before the introduction of GT crops, the use of CT had increased significantly in the second half of the twentieth century, largely due to the adoption of chemical herbicides that allowed growers to reduce their reliance on tillage for weed control. Despite this, the chemical-induced diffusion of CT was limited by several factors. First, to be effective some herbicides need to be applied at levels that can injure the crop; for high-residual chemicals, those injuries can potentially extend to future crops. In addition, the range of weeds that a typical chemical can treat is narrow, the post-emergence application window for many chemicals is highly sensitive to the environment, and there is often antagonism between grass and broad-leaf herbicides. In this setting, the advent of GT soybeans, introduced in the United States in 1996, was a game changer. Glyphosate is an effective broad-spectrum, low-residual herbicide, and GT crops can be treated with glyphosate with little to no injury (Carpenter and Gianessi 1999).

Because the combination of glyphosate and GT crops provides such an effective and convenient post-emergent weed control strategy, it can change farmers' propensity to adopt CT. Indeed, previous evidence indicates a positive correlation between GT crops and CT: cropped acreage under 'no-tillage' systems has increased considerably in the United States, Argentina, and Canada since the introduction and widespread adoption of GE varieties (Barrows, Sexton and Zilberman 2014, Fernandez-Cornejo et al. 2014). To investigate whether these correlations indicate a complementary relationship, previous research has employed econometric models

that estimate whether the adoption of GT varieties induces the adoption of CT, and also whether the adoption of CT induces the adoption of GT varieties. For cotton, Roberts et al. (2006), Frisvold, Boor, and Reeves (2009), and Kalaitzandonakes and Suntornpithug (2003) conclude in favor of complementarity, whereas Banerjee et al. (2009) fail to reject the null hypothesis that CT and GT cotton are independent. For soybeans, Fernandez-Cornejo, Klotz-Ingram and Jans (2002) and Fernandez-Cornejo et al. (2013) present evidence support a causal relationship between CT and GT soybeans, whereas the results in Fernandez-Cornejo et al. (2003) partially reject the presence of complementarities.

Overall, the evidence in favor of complementarity between GT crops and CT outweighs that against it, but data limitations and certain methodological assumptions restrict the generality of the existing findings. With respect to data, we note that, because of its nature, complementarity is best studied at the level of individual choices. Yet, three of the papers cited above (Roberts et al. 2006, Frisvold, Boor, and Reeves 2009, and Fernandez-Cornejo et al. 2013) rely on state-level data rather than farm-level choices. The three studies that do rely on farm-level data (Fernandez-Cornejo et al. 2003, Kalaitzandonakes and Suntornpithug 2003, and Banerjee et al. 2009) have access to just a single cross-section. Regarding methodology, two important features for the identification of complementarity have been neglected by previous studies. First, an appropriate test for complementarity requires a choice-set defined over all possible combinations of the available practices (Gentzkow 2007). For example, a grower facing the choice between two binary technologies would be modelled as choosing between four technology systems. When this is not true—as is the case for the bivariate probit or logit models often used in existing studies in this area—complementarity is either ruled out or inadequately characterized (Miravate and Pernías 2010, Gentzkow 2007). The second important modelling feature is allowance for the possibility that the unobserved returns are correlated across practices. This is because the clustering (or lack thereof) of the observed practices may be the result of correlated unobserved tastes, rather than complementarity. Restricting the unobserved returns across practices to be uncorrelated—as done by nearly all existing studies dealing with the complementarity between GT crops and CT—can lead to accepting complementarity when

it is absent, or rejecting complementarity when it is present (Athey and Stern 1998; Cassiman and Veuglers 2006).

In this paper we reconsider the problem of testing for complementarity between GT crops and conservation tillage. The novelty of our contribution relates to both the data used, which are considerably more extensive than in previous applications, and the econometric methodology that we apply, which draws on recent econometric advances. Concerning data, we use a representative farm-level dataset that spans the period 1998–2011 and contains the seed and tillage choices of 29,518 soybean growers. Because GT soybeans were commercially introduced in 1996, our data covers much of the period during which growers transitioned from conventional (CV) soybeans to GT soybeans. Moreover, while our dataset is not a balanced panel it does contain repeated observations over time for a subset of the individuals (on average, 43% of farmers sampled in any given year are re-sampled the next year). As a result, for many farmers we observe whether or not their tillage choice changed upon switching to GT soybeans, thus aiding in distinguishing complementarity from the correlation among unobserved returns. Regarding methodology, our empirical framework is based on a structural model with a single choice set for farmers that includes all four possible combinations of adoption decisions over GT soybeans and CT. This contrasts with previous farm-level tillage studies, where a grower is modelled as making two simultaneous, albeit distinct, adoption decisions. In such models, complementarity is not directly estimated and consequently the results can be difficult to interpret.¹ Our model also controls for the correlation induced by unobserved heterogeneity by estimating the full covariance matrix associated with the individual random effects.²

¹ For example, Fernandez-Cornejo et al. (2003) found that the adoption of GT soybeans did not induce the adoption of CT, but that the adoption of CT did induce the adoption of GT soybeans. It seems difficult to provide a structural interpretation to such an asymmetric adoption interaction, and it is unclear what one ought to conclude about whether CT and GT soybeans are complementary.

² An example of the type of unobserved factors we have in mind is a farmer's education. If producers with more education are both more likely to use CT and adopt GT soybeans, then the unconditional

Our results indicate that GT soybeans and CT are indeed complementary, a conclusion supported by several robustness checks. We also use our results to investigate the counterfactual scenario in which soybean growers did not have the option of choosing GT soybeans. We find that the adoption rates for both CT and no-tillage have increased by about 10% and 20%, respectively, due to the availability of GT soybeans. One of the implications of this result is that soil erosion was potentially lowered by 27 million tons per year during the 1998-2011 period. An approximate dollar value for this reduction is \$189 million per year.

The rest of this paper proceeds as follows. We first develop the model to be estimated, beginning with an exposition on the challenges associated with the econometric analysis of complementarity. We then specify the model and present the econometric procedure, with an explicit discussion of the identification conditions. This is followed by a description of the data, and a presentation and discussion of the empirical results. The paper concludes with a brief investigation of some counterfactual scenarios and a discussion of possible policy implications.

Modeling Complementarity

The definition of complementarity between two activities is that the marginal return to each activity is increasing in the level of the other activity. The relevant return to focus on depends on the objective function of the decision maker. For technology adoption, it is natural to focus on the profit associated with the various potential choices. The characterization of the notion of complementarity is best expressed by the property of supermodularity of the objective function (Brynjolfsson and Milgrom 2013). Consider two technologies or practices that a producer can choose to adopt separately, together, or not at all. Let $d_j = 1$ and $d_j = 0$ denote, respectively, the adoption and non-adoption of practice j , for $j \in \{1, 2\}$. The profit from using

correlation between CT and GT soybeans would be greater than the correlation that conditions on education.

any one of the four possible combinations of practices can therefore be expressed as $\tilde{\pi}(d_1, d_2)$.

Practices d_1 and d_2 are said to be complementary if profits are supermodular, i.e., if

$$(1) \quad \gamma \equiv [\tilde{\pi}(1,1) - \tilde{\pi}(1,0)] - [\tilde{\pi}(0,1) - \tilde{\pi}(0,0)] \geq 0.$$

When two practices are complementary, therefore, adoption of one while using the other has a larger effect on profits than adopting the practice in isolation. The structural representation in (1) provides the basis for testing hypotheses about complementarity.

Depending on the type of data at hand, there are two ways to proceed. First, given access to firm-level profit data, γ can be directly estimated (see Cassiman and Veugelers 2006). Often, however, data on profits (or other suitable performance measures) are not available—this is the case in our study. Alternatively, the hypothesis of equation (1) can be tested using adoption data only. The presumption is that a producer chooses the combination of practices that maximizes returns, thereby revealing information about the interaction between those practices.

Two significant challenges arise, however, in testing for complementarity with adoption data. First, the empirical framework needs to explicitly distinguish between complementarity and the correlation induced by unobserved heterogeneity. A common reduced-form approach taken by past studies, for example, has been to test for complementarity by estimating the correlation between two activities after controlling for firm characteristics (Arora and Gambardella 1990; Arora 1996; Cassiman and Veugelers 2006). The main limitation of this approach is that one can rarely control for all relevant characteristics; thus, finding a conditionally positive correlation will, at best, indicate that complementarity might be present. Alternatively, Athey and Stern (1998) outline a structural framework in which γ can be directly estimated (while still controlling for unobserved heterogeneity). Several papers have since used such a framework to test for complementarity in different environments. For example, Miravete and Pernías (2006) use a version of the multinomial probit model to test for complementarity among production and innovation strategies, and Gentzkow (2007) uses a mixed logit model to test for complementarity between print and online newspapers. Although these two papers pursue different modelling frameworks, an essential element of both is that the choice-set

includes all possible combinations of available practices.³ This permits the sign of γ to be directly estimated. Furthermore, both papers control for the potential correlation among the unobserved returns. Miravete and Pernías (2006) estimate the covariance between the unobserved returns to each practice. Similarly, Gentzkow (2007) allows the normally distributed error terms in his mixed logit framework to be correlated. This is in contrast to multinomial logit models, where the errors are assumed to be independently and identically distributed (IID) across alternatives.⁴

The second significant challenge to testing for complementarity with adoption data is sufficient identifying variation. The basic problem is that the observed clustering of two practices could result either from unobserved heterogeneity or true complementarity (as defined in (1)). For example, observing that two practices are almost always adopted together could be entirely due to individuals simply having a high preference for both practices, rather than the presence of an interaction effect. Additional information or identifying restrictions are thus required to distinguish between these two alternative explanations. In estimating our model we draw on three sources for identification. One source is exclusion restrictions, i.e., the inclusion of variables that affect the returns to some practices but not others (Gentzkow 2007).

³ In general, if there are n available practices then the choice-set would consist of 2^n alternatives. As Berry et al. (2014) note, the fact that the choice set grows exponentially can be a serious limitation to the types of problems that can be studied using this approach.

⁴ Two related studies in the agricultural literature deserve mention. Wu and Babcock (1998) use a multinomial logit model to explore the environmental implications of three farming practices. The choice-set they specify consists of all eight possible combinations of the three practices. However, because of computational considerations, they do not allow for correlation among the unobserved returns. Moreover, the objective of their study was not to test for complementarity (e.g., they do not try to estimate γ). Dorfman (1996) uses a multinomial probit model to study two technology adoption decisions by U.S. apple growers. He also specifies the choice-set over all combinations of decisions, and his model allows for the unobserved returns to be correlated (by estimating the covariance matrix). However, he does not attempt to identify structural complementarity.

The intuitive basis of exclusion restrictions is that changing a variable that only directly affects one practice will have no impact on the adoption of another practice unless they are inter-related.⁵ A second source of identification is panel data (Gentzkow 2007). Repeated observations for an individual indicate whether (s)he, upon changing one practice, is more (less) likely to choose another practice, thereby indicating that the practices are complements (substitutes). A third source of identification, the intuition behind which is similar to the idea of exclusion restrictions, is exogenous variation in choice-sets (Nevo 2000, p. 529). If some growers lack access to a certain practice, e.g., GT seed, then observing that they are less (more) likely to adopt another practice, e.g., CT, would indicate the presence of complementarity (substitutability). We return to how identification conditions apply specifically in our setting after we have provided details on the model.

The Model

We implement a variant of the mixed logit model similar to Gentzkow's (2007) framework. Let soybean growers be indexed by $i \in \{1, \dots, N\}$, a year by $t \in \{1, \dots, T\}$, and a field by $f \in \{1, \dots, F_{it}\}$. The formal unit of analysis is a farm-field-year combination. On each field in a given year, a soybean grower makes a discrete choice for two practices: the type of seed to plant, denoted by d_s ; and the type of tillage to employ, denoted by d_τ . For seed, a grower may choose conventional seed ($d_s = CV$) or glyphosate tolerant seed ($d_s = GT$); for tillage, he may choose intensive tillage ($d_\tau = IT$) or conservation tillage ($d_\tau = CT$). With two practices, there are four mutually exclusive systems (d_s, d_τ) :

$$(2) \quad \Omega_0 \equiv \{(CV, IT), (GT, IT), (CV, CT), (GT, CT)\}.$$

Denote the choice set for each grower in each year by Ω_{it} . For the most part, $\Omega_{it} = \Omega_0$. That is, we assume that nearly all growers in all years can choose among all four systems. However, a

⁵ Keane (1992) demonstrated via simulation that the covariance matrix of a multinomial probit model is not well-identified without exclusion restrictions.

handful of crop reporting districts (CRDs) early on in the sample have no observed GT soybean purchases.⁶ For these districts-years the presumed choice-set is: $\Omega_{it} = \{(CV, IT), (CV, CT)\}$.

Rather than directly specifying the normalized returns for each pair of choices, as done in Gentzkow (2007), in our setting it is instructive to start with the (unobserved) per-acre profit associated with system (d_s, d_τ) , denoted by $\tilde{\pi}_{itf}(d_s, d_\tau)$. For each of his/her field in each time period, grower i chooses system (d_s, d_τ) such that $\tilde{\pi}_{itf}(d_s, d_\tau) > \tilde{\pi}_{itf}(d'_s, d'_\tau)$, for all $(d'_s, d'_\tau) \in \Omega_{it}$ where $(d'_s, d'_\tau) \neq (d_s, d_\tau)$. For each system, the per-acre returns are specified to depend on a number of observable and unobservable variables, as follows:

$$(3) \quad \begin{aligned} \tilde{\pi}_{itf}(CV, IT) = & \tilde{\beta}_0^{CV, IT} + \beta_1 p_{CV, t} + \beta_2 r_{CV, t} + (\tilde{\beta}_3^{CV} + \tilde{\beta}_3^{IT}) Size_{it} + \tilde{\beta}_4^{IT} Fuel_t + \tilde{\beta}_5^{IT} Futures_t \\ & + \tilde{\beta}_6^{IT} EI_i + \tilde{\beta}_7^{IT} Palmer_{it} + (\tilde{\beta}_8^{CV} + \tilde{\beta}_8^{IT}) Trend_t + \tilde{v}_i^{CV} + \tilde{v}_i^{IT} + \tilde{\varepsilon}_{itf}^{CV, IT}, \end{aligned}$$

$$(4) \quad \begin{aligned} \tilde{\pi}_{itf}(GT, IT) = & \tilde{\beta}_0^{GT, IT} + \beta_1 p_{GT, t} + \beta_2 r_{GT, t} + (\tilde{\beta}_3^{GT} + \tilde{\beta}_3^{IT}) Size_{it} + \tilde{\beta}_4^{IT} Fuel_t + \tilde{\beta}_5^{IT} Futures_t \\ & + \tilde{\beta}_6^{IT} EI_i + \tilde{\beta}_7^{IT} Palmer_{it} + (\tilde{\beta}_8^{GT} + \tilde{\beta}_8^{IT}) Trend_t + \tilde{v}_i^{GT} + \tilde{v}_i^{IT} + \tilde{\varepsilon}_{itf}^{GT, IT}, \end{aligned}$$

$$(5) \quad \begin{aligned} \tilde{\pi}_{itf}(CV, CT) = & \tilde{\beta}_0^{CV, CT} + \beta_1 p_{CV, t} + \beta_2 r_{CV, t} + (\tilde{\beta}_3^{CV} + \tilde{\beta}_3^{CT}) Size_{it} + \tilde{\beta}_4^{CT} Fuel_t + \tilde{\beta}_5^{CT} Futures_t \\ & + \tilde{\beta}_6^{CT} EI_i + \tilde{\beta}_7^{CT} Palmer_{it} + (\tilde{\beta}_8^{CV} + \tilde{\beta}_8^{CT}) Trend_t + \tilde{v}_i^{CV} + \tilde{v}_i^{CT} + \tilde{\varepsilon}_{itf}^{CV, CT}, \end{aligned}$$

$$(6) \quad \begin{aligned} \tilde{\pi}_{itf}(GT, CT) = & \tilde{\beta}_0^{GT, CT} + \beta_1 p_{GT, t} + \beta_2 r_{GT, t} + (\tilde{\beta}_3^{GT} + \tilde{\beta}_3^{CT}) Size_{it} + \tilde{\beta}_4^{CT} Fuel_t + \tilde{\beta}_5^{CT} Futures_t \\ & + \tilde{\beta}_6^{CT} EI_i + \tilde{\beta}_7^{CT} Palmer_{it} + (\tilde{\beta}_8^{GT} + \tilde{\beta}_8^{CT}) Trend_t + \tilde{v}_i^{GT} + \tilde{v}_i^{CT} + \tilde{\varepsilon}_{itf}^{GT, CT}. \end{aligned}$$

In these equations, $p_{CV, t}$ and $p_{GT, t}$ represent the year t seed prices for CV and GT soybeans, respectively. Similarly, $r_{CV, t}$, and $r_{GT, t}$ denote the prices of herbicides used on these two types of varieties. $Size_{it}$ is a dummy variable indicating whether the grower grew more than 500 acres in soybeans. $Fuel_t$ is a price index for diesel fuel, $Futures_t$ is the average soybean futures price in January for the next November contract, EI_i is an index that measures soil

⁶ CRDs are regions—each representing a collection of counties—used by the USDA for statistical reporting of certain data. It is also the finest level at which our seed and tillage data are representative.

erodibility, $Palmer_{it}$ is a drought severity index and $Trend_t$ is a time trend.⁷ The ν_i terms are time-invariant, practice-specific normally distributed unobservables. They represent individual characteristics we do not observe, such as education, which may affect the returns to the different practices. As we discuss further below, we allow for the ν_i to be correlated across systems. The terms $\tilde{\varepsilon}_{itf}^{d_s, d_\tau}$ are system-specific IID type I extreme value errors.⁸ Their inclusion captures the fact that growers with the same characteristics and the same environment may still choose a different system.

The remaining symbols in equations (3)-(6) are parameters to be estimated. The intercepts $\tilde{\beta}_0^{d_s, d_\tau}$ are alternative-specific constants that capture the mean unobserved returns to each system. The superscripts of the other parameters indicate whether, and how, the associated variables are presumed to have a practice-specific effect. For example, we assume that El_i , which is invariant across systems, will differ in its impact on profits depending on the type of tillage used. If this were not the case, i.e., if the effect of El_i was the same across systems then it would have no effect on the grower's choices (the term would drop out upon differencing the equations). This highlights the additional fact that not all of the parameters in equations (3)-(6) are identified. Only parameters that contribute to differences in per acre returns are estimable (Train 2009).

To clarify which parameters are identified, as well as how the model nests a test for complementarity, we normalize returns relative to a base system, which is taken to be the

⁷ Further details on and summary statistics for each of these variables are provided in the Data section below.

⁸ Per standard practice, the variance of the extreme value distribution is normalized to $\pi^2/6$. Thus, the model coefficients are identified relative to the unobserved scale parameter (see, e.g., Kurkalova, Kling, and Zhao 2006).

(CV, IT) system. Defining $\pi_{itf}(d_s, d_\tau) \equiv \tilde{\pi}_{itf}(d_s, d_\tau) - \tilde{\pi}_{itf}(CV, IT)$, normalized returns can then be written as follows:

$$(7) \quad \pi_{itf}(CV, IT) = 0$$

$$(8) \quad \pi_{itf}(GT, IT) = \beta_0^{GT} + \beta_1(p_{GT,t} - p_{CV,t}) + \beta_2(r_{GT,t} - r_{CV,t}) + \beta_3^{GT} Size_{it} + \beta_8^{GT} Trend + v_i^{GT} + \varepsilon_{itf}^{GT}$$

$$(9) \quad \begin{aligned} \pi_{itf}(CV, CT) = & \beta_0^{CT} + \beta_3^{CT} Size_{it} + \beta_4^{CT} Fuel_t + \beta_5^{CT} Futures_t + \beta_6^{CT} EI_i + \beta_7^{CT} Palmer_{it} \\ & + \beta_8^{CT} Trend_t + v_i^{CT} + \varepsilon_{itf}^{CT} \end{aligned}$$

$$(10) \quad \pi_{itf}(GT, CT) = \pi_{itf}(GT, IT) + \pi_{itf}(CV, CT) + \gamma + \varepsilon_{itf}^\gamma,$$

where, for each system, the parameters' superscript now denotes the practice that is different relative to the base system (CV, IT) (e.g., $\beta_0^{GT} \equiv \tilde{\beta}_0^{GT, IT} - \tilde{\beta}_0^{CV, IT}$). Furthermore:

$$(11) \quad \gamma \equiv \left(\tilde{\beta}_0^{GT, CT} - \tilde{\beta}_0^{GT, IT} \right) - \left(\tilde{\beta}_0^{CV, CT} - \tilde{\beta}_0^{CV, IT} \right)$$

$$(12) \quad \varepsilon_{itf}^\gamma \equiv \left(\tilde{\varepsilon}_{itf}^{GT, CT} - \tilde{\varepsilon}_{itf}^{GT, IT} \right) - \left(\tilde{\varepsilon}_{itf}^{CV, CT} - \tilde{\varepsilon}_{itf}^{CV, IT} \right).$$

Hence, the sum $\gamma + \varepsilon_{itf}^\gamma$ captures whether GT soybeans and CT are complementary. To see this, note that, in terms of the un-normalized returns, we have:

$$(13) \quad \gamma + \varepsilon_{itf}^\gamma = \left(\tilde{\pi}_{itf}(GT, CT) - \tilde{\pi}_{itf}(GT, IT) \right) - \left(\tilde{\pi}_{itf}(CV, CT) - \tilde{\pi}_{itf}(CV, IT) \right).$$

Equation (13) revisits the relation in equation (1), which determines whether the two choices of interest are complementary. However, this relation is now adjusted for the presence of unobserved heterogeneity; complementarity can vary over the population through ε_{itf}^γ .

Because $E[\varepsilon_{itf}^\gamma] = 0$, it follows that γ is best interpreted as a measure of *mean* complementarity in the population. If our estimate for γ is statistically significantly greater (less) than zero, then GT soybeans and CT are, on average, complements (substitutes). Note also that, in this framework, γ does not vary on the basis of the observable characteristics. This is a consequence of our assumption that the observable variables have practice-specific effects rather than

system-specific effects. This assumption is primarily rooted in our goal of obtaining a straightforward test for complementarity, as encapsulated by γ . In this regard we follow Miravete and Pernías (2006), Gentzkow (2007), and Kretschmer, Miravete, and Pernías (2012), who also specify the observable variables as having practice-specific effects rather than system-specific effects.⁹

To control for the correlation induced by unobserved heterogeneity, we allow for ν_i^{GT} and ν_i^{CT} to be correlated. Specifically, we assume that $(\nu_i^{GT}, \nu_i^{CT}) \sim N(0, \Sigma)$, where

$$(14) \quad \Sigma = \begin{pmatrix} \sigma_{GT}^2 & \sigma_{GT,CT} \\ \sigma_{GT,CT} & \sigma_{CT}^2 \end{pmatrix}.$$

By estimating $\sigma_{GT,CT}$, we control for unobserved factors that contribute simultaneously to the returns of $\pi_{itf}^{GT,IT}$ and $\pi_{itf}^{CV,CT}$. For example, if ν_i^{CT} is large (small) whenever ν_i^{GT} is large then these two terms will be positively (negatively) correlated. Without controlling for such correlation, estimates of γ would be biased upward (downward). Some of the specific kinds of unobserved variables that we have in mind include the grower's education, attitude towards new technologies, and degree of risk aversion. For example, better educated individuals may face lower adoption costs and so may be more likely to use both GT soybeans and CT. Similarly, individuals that are generally more open to new technologies (so-called early adopters) may be more likely to use both GT soybeans and CT. If a person is very risk averse, on the other hand, the opposite may hold true: GT soybeans may be viewed as less risky than CV soybeans whereas CT may be viewed as more risky than IT, leading to a negative correlation between the unobserved returns.

Because we have differenced out the returns to the (CV, IT) system, the model as written in equations (7)-(10) makes explicit which parameters are identified. The parameters on variables that enter all of the equations are identified relative to the (CV, IT) system. For

⁹ See Athey and Stern (1998) for a more detailed discussion of these issues.

example, the sign of the estimate for β_1^{GT} will indicate whether a large farm is more likely to adopt GT soybeans *relative to* CV soybeans. The parameters on the alternative-specific variables, such as prices, indicate how changes in the differences of those variables affect returns. For example, β_1 is the effect of a change in the price of GT seed relative to the price of CV seed.

Identification

The model as presented is formally identified: there are more moments than parameters. However, as noted previously, the precise identification of the parameters, in particular the complementarity and covariance parameters, requires additional sources of variation and information that go beyond the basic formal requirements. The issue is that the patterns of adoption generated by a model with positively correlated unobserved returns ($\sigma_{GT,CT}$ is high) and practices that are substitutes ($\gamma < 0$) can be very similar to the adoption patterns generated by a model with negatively correlated unobserved returns ($\sigma_{GT,CT}$ is low) and practices that are complements ($\gamma > 0$). Thus, to distinguish between correlated tastes and complementarity requires some form of variation in the data that would occur because of only one of these effects, while holding the other constant.¹⁰

As noted earlier, one source of identification is exclusion restrictions. To illustrate the role of these identifying restrictions in the context of the model just presented, suppose that the price of GT soybean seeds relative to that of CV soybeans directly affects the seed choice but not the tillage choice (i.e., the relative seed price is an excluded variable). Further, suppose that there is a shock to this relative price, for example it decreases. Then some producers will switch from CV soybeans to GT soybeans. If GT soybeans and CT are independent, then there should be no change in the adoption of CT since the seed price does not directly affect it. If they complement, however, then we would also observe an increase in the use of CT. Some of the producers that previously chose CV soybeans with IT would switch to using GT soybeans with

¹⁰ For a more comprehensive discussion of these issues, see Gentzkow (2007).

CT. Intuitively, the switch to GT soybeans (based on the price change) would shift up the return to CT, thus also leading to its adoption.

The variables that fulfil the exclusion restrictions in our model are those that affect the seed choice – i.e., variables in equation (8) – but not the tillage choice (equation (9)), and vice versa. The specific variables that we assume directly affect the seed choice but not the tillage choice include the difference in seed prices ($p_{GT,t} - p_{CV,t}$) and the difference in herbicide prices ($r_{GT,t} - r_{CV,t}$) (i.e., these variables enter the second equation but not the third). Differences in relative seed prices should have no effect on the relative return to the different tillage operations. With regard to herbicide prices, previous studies by Bull et al. (1992), Fawcett et al. (1994), and Fuglie (1999) do not find a significant difference in pesticide use between CT and IT systems; thus we assume it does not directly affect the tillage choice.¹¹

The variables assumed to directly affect the tillage choice but not the seed choice include $Fuel_t$, $Futures_t$, EL_i , and $Palmer_{it}$. The variable $Fuel_t$ is included to capture the argument that CT generally requires less fuel (Triplett and Dick 2008). For a given tillage method, however, there will be little difference in fuel usage for different seed types. Similarly, the EL_i only enters the tillage equation because the degree of erodibility will not have a differential effect on the seed choice (holding the tillage-type constant). The same argument applies for $Palmer_{it}$, which is included because CT leaves more ground cover in place and may be chosen to conserve moisture in dry years. Finally, $Futures_t$ is included to capture changes in relative returns due to yield differences between the tillage options. Previous research has generally indicated that there is no significant yield difference between GT and CV soybeans (Qaim 2009). Rather, the primary reason farmers prefer GT soybeans is because they provide easier weed control and a reduction in management time (Qaim 2009).

¹¹ As part of robustness checks reported later, we do allow for herbicide prices to differ in their impact by the type of tillage employed. We find that it does not affect our complementarity result.

The contribution of panel data to identification occurs through the estimation of the distribution of the time-invariant random effects. Intuitively, if the adoption of GT soybeans and CT are correlated because of a high covariance parameter $\sigma_{GT,CT}$ (rather than complementarity), then the adoption of GT and CT for a given individual would be uncorrelated over time. Individuals may have a high propensity to use both, but conditional on changing one practice, they would be no more (or less) likely to use another. On the other hand, if we observe across time periods that whenever a given individual adopts GT, he or she is more likely to adopt CT, then this would imply the presence of synergies. Regarding choice-set variation, as noted previously, early on in our sample we do not observe any purchases of GT varieties in certain crop reporting districts (CRDs). We interpret this to mean that they were not available as an option, and thus we exclude them from the choice-sets of individuals within that region.¹²

Estimation

The model is estimated by the method of simulated maximum likelihood (SML) (Train 2009). To simplify the notation, let j denote system (d_s, d_τ) , that is, $j \in \Omega_{it}$. Furthermore, rewrite equations (7)-(10) as:

$$(15) \quad \pi_{itf}^j = x_{itf}^j \beta^j + v_i^j + \varepsilon_{itf}^j,$$

where x_{itf}^j is the vector of explanatory variables pertaining to system j , and β^j is the associated parameter vector (note that $\pi_{itf}^{CV,IT} = 0$, as above). Let θ denote the vector of all parameters to be estimated (this includes the vector of all β parameters, which implicitly also define the complementarity parameter γ , as well as the parameters of the covariance matrix Σ

¹² Because only a small number of CRDs do not have observed GT seed purchases (early on in our sample), this type of identification plays a small role in our analysis.

). Then, for a *given* realization v_i^j , the probability of choosing system j is provided by the standard logit expression:

$$(16) \quad L_{itf}^j(v; \theta) = \frac{e^{x_{itf}^j \beta^j + v_i^j}}{\sum_{k \in \Omega_{it}} e^{x_{itf}^k \beta^k + v_i^k}}.$$

Let $j_{itf} \in \Omega_{it}$ denote the actual system choice of grower i for field f in year t , and define $\zeta_i \equiv \{j_{itf}\}$ as the set of all actual choices in the sample for grower i . Given v_i^j , the probability of ζ_i is given by the product of the corresponding logits:

$$(17) \quad L_{\zeta_i}(v; \theta) = \prod_{j \in \zeta_i} L_{itf}^j(v; \theta).$$

The unconditional probability is given by the integral over all v that generate ζ_i :

$$(18) \quad P_{\zeta_i} = \int L_{\zeta_i}(v; \theta) f(v) dv.$$

Since P_{ζ_i} is an integral it can be estimated via simulation. For each individual, multiple draws of the v_{ij} are taken, L_{ζ_i} is computed, and then averaged. Specifically, let R denote the number of draws of v_i^j for each individual. Then P_{ζ_i} is approximately given by:

$$(19) \quad P_{\zeta_i} \approx \frac{1}{R} \sum_{r=1}^R L_{\zeta_i}(v_r; \theta).$$

The SML estimator is therefore given by:

$$(20) \quad \hat{\theta} = \arg \max_{\theta} \sum_i \left[\ln \frac{1}{R} \sum_{r=1}^R L_{\zeta_i}(v_r; \theta) \right].$$

The statistical package that we use is the Stata user-written mixlogit package developed by Hole (2007) (for further details see also Cameron and Trivedi, 2010, p. 523). In simulating the likelihood function we use 250 Halton draws, which is well above the minimum

recommendation of 100 (Hensher, Rose, and Greene 2005, p. 616).¹³ It is also important to re-emphasize that the estimated parameters, $\hat{\theta}$, are identified relative to the unobserved variation of the IID extreme value unobservables, which are implicitly normalized prior to estimation (see earlier footnote 8). Thus, for example, instead of estimating the complementarity parameter γ , the model actually estimates γ/ϕ (where ϕ is the unobserved scale parameter for the extreme value type I distribution). For simplicity, and slightly abusing notation, we continue to use the same parameter symbols (e.g., γ) for the remainder of the paper.

Data

The model is estimated with farm-level seed and tillage data from the survey company GfK.¹⁴ These data, which are designed to be representative at the CRD level, span 1998–2011 and include about 4,982 farmers per year (each farmer can have multiple fields). As noted above, about 43% of growers sampled in any given year are also sampled the next year. In total, our sample contains 82,056 farm-field-year observations across 235 CRDs in 31 states (with the largest soybean states being the most heavily represented). Among the variables previously defined, those that come from the GfK data include tillage and seed choices (i.e., the endogenous variables), seed and herbicide prices, and the variable for farm size. With respect to the tillage choice, in our data each plot is identified as using one of three following alternatives: “Intensive Tillage,” “Conservation Tillage,” or “No-Till.” For our baseline specification, we treat intensive tillage as a distinct category, and combine the plots identified as “Conservation Tillage” and “No-Till” into the model’s CT category. However, we also consider an alternative

¹³ Train (2000) demonstrated that the SML estimates for a mixed logit model using 100 Halton draws outperform the SML estimates using 1,000 random draws. The practical benefit of this is that estimation time is decreased by a factor of ten while simultaneously increasing accuracy. For a further discussion of Halton sequences see Train (2009).

¹⁴ Specifically, we use data from GfK’s AgroTrak® and Soybean TraitTrak™. See the company’s website (<http://www.gfk.com/us>) for a brief description of these proprietary data products.

aggregation procedure where the model's CT category is associated only with the plots classified as "No-Till" (NT) in the dataset, and combine the remaining two classifications into the model's IT category. Where applicable, we make explicit which definition is being used.

The shares for each seed-tillage system are provided in table 1, where the distribution of system choices over time are disaggregated into three sub-periods. From 1998–2001, CV soybeans still accounted for about 40% of the observations, but from 2002–2006 they only made up about 13%, and for the final sub-period just over 5%. Overall, systems with GT soybeans accounted for about 80% of all observations, whereas systems with CT accounted for about 62% of all observations.

Table 1 also shows that about 67% of acres planted to GT soybeans use CT whereas about 50% of acres planted to CV soybeans use CT. This is generally consistent with previous work based on different data sources (e.g., Fernandez-Cornejo et al. 2014). The correlation coefficient between GT soybeans and CT is 0.125 and is significant at a 1% level. Changes over time also show a positive correlation. Figure 1 contains U.S. annual adoption rates for GT soybeans, CT, and NT from 1998–2011. GT soybean adoption increased from just under 40% of acres in 1998 to about 95% of acres in 2011. Over the same period, CT increased from just under 60% of acres in 1998 to nearly 70% of acres in 2011. NT increased even more, from 32% in 1998 to 45% in 2011 (and a peak of 53% in 2008).

With regard to the remaining variables, the *EI* data were obtained from the National Resources Inventory (a survey conducted by the National Resources Conservation Service), soybean futures were obtained at www.quandl.com, diesel fuel prices were obtained from Quick Stats at the USDA-NASS website, and the Palmer Z-Index was obtained from www.ncdc.noaa.gov. Below we provide additional details, as well as a discussion of their expected effects, for each of the regressors. Table 2 provides a summary of their distributions.

Farm Size is a dummy variable that indicates whether a grower planted more than 500 acres in soybeans. The arbitrary cut-off of 500 acres was a natural choice given the available data, in which each farm is classified into one of five categories: (i) <100 soybean acres, (ii) 100–

249 acres, (iii) 250–499 acres, (iv) 500–999 acres, and (v) 1,000 or more acres.¹⁵ We include *Farm Size* for both the seed and tillage choices to control for scale effects. Past studies have noted that the use of CT, in particular no-tillage, can require large fixed costs in the form of better adapted machinery (Knowler and Bradshaw 2007). Given this, we expect that larger farms will be more likely to adopt CT. With regard to the seed choice, we have no strong prior expectations. Fernandez-Corenjo et al. (2002) find that larger farms are more likely to adopt GT soybeans, whereas Fernandez-Cornejo et al. (2003) find no size effect. The latter argue that since the adoption of GT soybeans does not require significant fixed costs, there should not be significant differences in adoption between large and small farms.

Futures is the Chicago Mercantile Exchange mean soybean futures price in the month of January for same year November contract. It is included as a proxy for the output price that is expected by producers. We use January because that is a common time at which practice decisions are made, and we use November because it is the closest month after harvest. We include it as an explanatory variable for the tillage choice because there might be yield differences between IT and CT. Previous studies, however, are inconclusive on the effect of output prices on CT (Knowler and Bradshaw 2007).

Fuel Price is an annual index for diesel fuel prices (as noted above, it is obtained from USDA-NASS). We use the mean index from September to May as this is the period during which most tillage decisions are made. The index is included to control for potential differences in fuel usage between CT and IT operations. From 1998–2011, real fuel prices rose significantly and thus could explain some of the variation in tillage trends. Since CT tends to use less fuel, our expectation is that higher prices will increase the likelihood of using CT.

EI is a county-specific, time-invariant index of soil erodibility due to water events. It measures a soil's potential to erode. A higher index indicates that greater investment is required to maintain the sustainability of the soil under intensive cultivation. The National Resources

¹⁵ For robustness checks we also estimate the model with the *Size* variable cutoffs set at 250 and 1,000 acres. Overall, the results remained unchanged.

Inventory considers scores of 8 or above to indicate highly erodible land. The *EI* is included for a couple of reasons. First, the 1985 Farm Bill requires a producer that grows crops on highly erodible land to meet certain minimum conservation requirements (Stubbs 2012) in order to be eligible for some government payments. An acceptable way to comply is to use CT. Second, a grower may be more likely to use CT on highly erodible land in order to preserve the soil's productivity into the future (Soule, Tegene, and Wiebe 2000). Given these two rationales, as well as previous findings, we expect that the *EI* will have a positive sign; i.e., a grower will be more likely to use CT on more erodible land.

Palmer's Z is the mean Palmer's Z-Index for the month of September in the prior year (this variable is at the climate-division level; see Xu et al., 2013, for more details). This index indicates how dry a locality is relative to normal conditions. Negative values indicate drier conditions, whereas positive values indicate wetter conditions. We include Palmer's Z-Index because the presence of drought may increase the likelihood of adopting CT. For instance, Ding, Schoengold, and Tadesse (2009) find that drought is associated with a greater likelihood of using no-till and other CT practices.

The *Seed Price* term ($p_{GT,t} - p_{CV,t}$) is the difference between mean annual U.S. GT and CV soybean seed prices (\$/50 lb bag). In our data we observe the transaction prices for each individual, but we do not observe the price for the type of seed they did not buy (e.g., if a grower purchased CV seeds, we do not know the price they would have paid for GT seeds). Thus, as a proxy for that price, we average over all individuals within a given year. We aggregate to the national level because, beyond 2003, there are very few observations for CV seed purchases. As a result, averaging at a finer level would introduce considerable sampling variation. Figure 2 presents GT and CV seed prices from 1998–2011.

Prior to 2009 there was comparatively little movement in both relative prices and overall prices. The increase in soybean output prices in 2008 led to a significant rise in seed prices in 2009. In terms of expectations, the higher the price of GT seed relative to CV seed, the smaller the return for GT seeds. Thus, a negative sign is expected. It is worth noting, however, that previous studies have found a positive sign for seed price (see, e.g., Fernandez Cornejo et al.

2002). This is likely because of the rapid diffusion of GT soybeans that coincided with a slight increase in relative prices; we control for this process with a time-trend.

The *Herbicide Price* term ($r_{GT,t} - r_{CV,t}$) is the difference between the annual U.S. price indices for glyphosate and for a group of seven post-emergence conventional herbicides. Our assumption is that the glyphosate price is the main herbicide price a grower looks at when considering the adoption of GT soybeans. For CV soybeans, the matter is less straightforward. As noted earlier, many of the herbicides used on CV soybeans are only effective against specific weed species. In addition, only some of these herbicides can be applied post-emergence. We chose to use only the prices from post-emergence herbicides because they are what primarily differentiate CV soybeans from GT soybeans.¹⁶ In terms of calculation, glyphosate prices are annual volume-weighted averages in dollars per pound. The price for CV soybeans is a Laspeyres Index: each year, the index is a weighted average of the ratio of current prices to base prices. For the base, we use the mean prices and shares of the seven herbicides for the entire 1998–2011 period, and the resulting index is re-scaled to equal 1 for the year 1998. Figure 3 presents these indices for the 1998–2011 period. For comparison, both the glyphosate and CV herbicide prices are normalized to equal 1 in 1998. The price of glyphosate has fallen considerably and almost uniformly since 1998. This is primarily due to the expiration of Monsanto’s patent in 2000. The exception to the trend decline occurred during 2008–2009, when prices rose significantly. During 2008–2009 commodity prices, and in turn land used for cropping, were very high. This, combined with a growing demand for GT corn, led to shortages in glyphosate and an associated price increase.

The time *Trend* variable is included to capture the impact of other factors that contributed to the diffusion of GT soybeans and CT. This was particularly important for GT soybeans, and adoption rose from 38% to 86% over the period 1998–2003. This adoption pattern

¹⁶ The seven herbicides we include are Raptor®, Flexstar® 1.88L, Fusion®, FirstRate®, Select® 2 EC, Cobra®, and Pursuit® 2 EC. We selected these herbicides because they were the most frequently used post-emergence herbicides applied on CV soybeans.

was driven by a variety of factors that are not captured by our model. We expect that the adoption of GT soybeans will be positively associated with this trend variable. For CT, we have no strong prior expectations.

Empirical Results

Table 3 contains our baseline specification. Overall, the results are consistent with expectations. The alternative-specific constant for GT seed varieties is positive and significant. Conversely, the constant for CT is negative and significant. This is unsurprising given that a large number of farms continued to adopt IT despite the presence of synergies between GT soybeans and CT (as indicated by the result for γ). Higher prices for GT seed (relative to CV seeds) and glyphosate (relative to substitute herbicides) are associated with a lower likelihood of using GT soybeans. Larger farms are more likely to use both GT soybeans and CT. Also, the relative size of the parameter for CT is significantly larger, suggesting that farm size plays a larger role for the tillage decision. The linear time trend is significant and positive for both GT soybeans and CT, though significantly larger for GT soybeans, as would be expected. Among the variables exclusive to the tillage decisions, there are some interesting results. Higher soybean futures prices are associated with a lower likelihood of using CT, though the effect is only significant at 5%. This suggests that there may be a small perceived yield-loss associated with the use of CT. For some soils the formal agronomy literature provides evidence to support this perception (Triplett and Dick 2008). Higher fuel prices, on the other hand, significantly increase the likelihood of using CT. We also find that more drought-like conditions, as captured by Palmer's Z-Index, increase the likelihood of using CT, corroborating the finding by Ding, Schoengold and Tadesse (2009). Finally, a higher *EI* is also found to be associated with a significantly higher probability of CT use.

For the unobservables, we find significant evidence of unobserved variation in preferences for both GT soybeans and CT. The unobserved variance for CT is particularly large, suggesting that a variety of omitted individual characteristics are important for determining the

best tillage practice. This seems intuitive given the relatively large adoption rates for both CT and IT throughout the sample period. Unobserved variation in tastes is also important for the seed choice, though relatively less so. This is probably a reflection of the fact that later on GT soybeans are adopted by nearly everyone, and thus a relatively smaller variance can rationalize the small share of farms that still use CV soybeans. The covariance across the errors is also significant. The implied correlation is about 0.105. Thus, farmers who have a strong preference for GT soybeans (i.e., a large v_i^{GT}) are more likely to have a strong preference for CT (i.e., a large v_i^{CT}) and vice versa. Finally, the estimate for complementarity, γ , is highly significant and positive, indicating that GT soybeans and CT are indeed complementary practices.

What is the economic significance, to the farmer, of the estimated complementarity effect? One measure is provided by a grower's willingness to pay (WTP) for it. In a standard discrete choice random utility model, the WTP for an attribute is given by the ratio of that attribute's coefficient to the absolute value of the coefficient for the price variable (note that the ratio of the two estimated coefficients will be independent of the unidentified scaling parameter). In our model, the objective function is profit per acre. As a result, the estimated coefficient for the seed price represents the number of soybean bags planted per acre (relative to the unidentified scaling parameter). Dividing an attribute's coefficient by the absolute value of the coefficient for seed price thus gives the WTP per bag of soybeans for that attribute. For γ , this implies a WTP of about \$1.41 per bag of soybean seeds. Given that a typical density for soybeans is 1.2 bags/acre, the WTP of a typical farmer for the synergies provided by complementarity between GT seeds and CT is \$1.69/acre.

Because the coefficients are identified relative to the scale parameter ϕ , only their sign is directly interpretable. To get a better idea of the importance of each of the variables we simulate the change in the adoption of GT soybeans and CT in response to a change in the value of each of the exogenous variables. This exercise also serves to highlight the role of complementarity for the impacts of each of the independent variables. Table 4 contains the Average Marginal Effects (AMEs) for GT and CT adoption with respect to each of the regressors. The AME of a variable is

the average change in the probability of adopting a practice, e.g., GT soybeans or CT, in response to a change in that variable. With the exception of *Size*, we compute elasticities. As an example, to calculate the effect of a change in the *EI* on GT soybean adoption, we first simulate and compute for each individual:

$$(21) \quad \psi_{itf}^{GT,EI} \equiv \frac{\Delta \Pr(GT)}{\Delta EI} \frac{EI}{\Pr(GT)},$$

where $\psi_{itf}^{GT,EI}$ denotes the elasticity of the probability of GT soybean adoption with respect to the *EI*. The result reported in table 4 is the average of these elasticities over all individuals, time periods, and fields. The superscripts “I” and “D” indicate whether the impact of the variable on the practice is indirect or direct, respectively.

Overall, the results indicate that the seed price plays the largest role among the variables. For example, a 1% increase in the price of GT soybeans relative to that of CV soybeans results in a slightly-more-than 1% direct decrease in the probability of adopting GT soybeans. Through the complementarity effect, it also indirectly decreases the probability of adopting CT by 0.08%. The impacts of the other continuous variables can be interpreted in a similar manner. Because the variable *Size* is binary, an elasticity cannot be computed; instead, we compute the percent difference in the probability of adopting a practice between growers with more than 500 soybean acres and growers with less than 500 soybean acres. Note also that the impacts for *Size* are made up of both direct and indirect effects. The simulation indicates that a farm with 500 or more soybean acres is 6.9% more likely to adopt CT and 2.1% more likely to adopt GT soybeans.

Complementarity Under Alternative Specifications

Certain variations on our specification, such as allowing herbicide prices to directly impact the relative profitability of CT, are also plausible, which may be important for the complementarity finding. In addition to testing for robustness across these alternatives, this section serves to highlight the role of certain assumptions, such as admitting non-zero correlation between the unobserved returns, for the estimate of γ . Table 5 contains estimates of γ for several different specifications. Allowing for the *Herbicide Price* variable to directly impact

the tillage choice reduces the coefficient somewhat, but does not alter the finding of complementarity.

The next specification demonstrates the effect of not allowing unobserved tastes to be correlated (i.e., $\sigma_{GT,CT} = 0$). In this case the estimate for γ increases as it captures some of the effect that is actually the result of correlated tastes. We also estimate the model when ignoring the fact that some individuals have repeated observations (i.e., we assume that the ν terms are IID across fields and time for the same individual). This substantially increases both the estimate and the standard error for γ , which suggests that when using the mixed logit model, it is important to utilize the panel aspect of the data. The “Basic Logit” specification not only ignores the panel aspect of the data but also does not allow for unobserved heterogeneity (i.e., the ν terms are set to 0). In this case, the estimate for γ is actually closer to the original model than the estimate that ignored the panel aspect of the data.

We also estimated the model with data from the Central Corn Belt (CCB) only (the states we include are IA, IL, and IN). These three states account for nearly 35% of U.S. soybean land alone. Our result for γ in this case is less than before. However, since γ is estimated on a different sample, it is not directly comparable to the estimate obtained from our baseline specification. Because the parameters are identified relative to the scale parameter, a different value could indicate that complementarity between GT soybeans and CT is less in this region, but it could alternatively indicate that the IID portion of unobserved variation is larger in the CCB (relative to the rest of the country).

The final specification changes the way the tillage choice is structured. Instead of specifying the tillage choice for the farmer as being between CT and IT, we instead specify it as being between no-tillage (NT) and tillage (i.e., some positive level of tillage). We expect the complementarities between NT and GT soybean to be even stronger than between CT and GT. Intuitively, the improved efficiency and convenience of weed control offered by GT varieties will be especially beneficial when making the leap to a NT system. This is weakly confirmed by the correlation coefficient between GT soybeans and NT, which is slightly larger at 0.139

(compared to 0.125). The estimate for γ presented in table 4 indicates that NT and GT soybeans are complementary, and the magnitude of γ is indeed larger than it was for the CT specification. As was noted for the case of the CCB specification, the estimates for complementarity are not directly comparable. Nonetheless, the fact that the estimates of the parameters for the GT variables – the constant, the seed price, and the herbicide price – remain essentially unchanged relative to the base specification, suggests that the larger estimate for $\hat{\gamma}$ is in fact the result of stronger complementarity, rather than smaller variation in the IID portion of unobserved tastes.

Conservation tillage without GT varieties

A natural question that arises from our model is what CT adoption rates would have been if GT soybeans were never introduced into the market. To answer this question, we calculate the following: (i) the annual predicted CT adoption rates using the estimates from table 3 (i.e., the predicted rates based on having GT soybeans as part of the choice-set); and (ii) the annual predicted CT adoption rates after removing GT soybeans from the choice-set for all individuals (also using the parameter estimates from table 3). To arrive at the first set of adoption rates, we first compute for each farm-field-year combination the vector of predicted probabilities of choosing systems with CT (which requires simulation). Specifically,

$$(22) \quad \hat{L}_{itf}^j(\hat{\theta}) = \frac{1}{R} \sum_{r=1}^R \frac{e^{x_{itf}^j \hat{\beta}^j + v_{i,r}^j}}{\sum_{k \in \Omega_{it}} e^{x_{itf}^k \hat{\beta}^k + v_{i,r}^k}} \quad , \quad j \in \{(CV, CT), (GT, CT)\} .$$

The predicted probability for choosing CT is then given by: $\hat{L}_{itf}^{CT} = \hat{L}_{itf}^{CV,CT} + \hat{L}_{itf}^{GT,CT}$. To move from this expression to annual adoption rates we use a variable in our dataset that consists of the number of acres that each farm-field-year represents in the population for that year. Denote this quantity by A_{itf} . The predicted share of CT acres in year t is then given by

$$(23) \quad \hat{S}_t^{CT} = \frac{\sum_{i=1}^{I_t} \sum_{f=1}^{F_{it}} A_{ift} \hat{L}_{ift}^{CT}}{\sum_{i=1}^{I_t} \sum_{f=1}^{F_{it}} A_{ift}}.$$

To compute the predicted annual shares for CT when GT soybeans are not available, we follow essentially the same steps, except that the predicted probability of using CT now just consists of a singleton, denoted by $\tilde{L}_{ift}^{CV,CT}$ (i.e., the only choice being made concerns which the tillage practice to use). We calculate this probability according to

$$(24) \quad \tilde{L}_{ift}^{CV,CT} = \frac{1}{R} \sum_{r=1}^R \frac{e^{x_{ift}^{CV,CT} \hat{\beta}^{CV,CT} + v_{i,r}^{CV,CT}}}{1 + e^{x_{ift}^{CV,CT} \hat{\beta}^{CV,CT} + v_{i,r}^{CV,CT}}}.$$

Note that, as compared with (22), the denominator inside of the summation in (24) does not include the terms for GT choices. The predicted adoption rates for CT when GT soybeans are not available can then be computed as in (23), but with $\tilde{L}_{ift}^{CV,CT}$ replacing \hat{L}_{ift}^{CT} . Table 6 contains these predicted adoption rates for each year of the 1998–2011 period. In 1998 the adoption rate for CT is 3.3 percentage points less in a world without GT soybeans as an option. This difference increases steadily up until 2003, at which point it begins to level off and approach 6 percentage points (or about 10% of the no-GT soybean scenario). This is a reflection of the diffusion of GT soybeans, which also began to level off in 2003. Note also that the predicted rate for CT increases considerably over the period, by about 10 percentage points, even when GT soybeans are not available. The implication of our model is that such an increase would have been driven mainly by steadily rising fuel prices, an overall increase in farm size, and other unknown factors captured by the trend variable. The simulation is also performed for NT. In this case the gains from complementarity are even greater. In 1998, the difference is about 4 percentage points

more when GT soybeans are available, and by 2011 the difference is 9 percentage points or a bit over 20% relative to the scenario without GT soybeans.¹⁷

An Application to Soil Erosion

Conservation tillage or no tillage are not necessarily desirable, *per se*. Rather, interest in these practices is motivated by the fact that they affect a variety of environmentally-relevant outcomes. Exploring all such implications is beyond the scope of this paper. As suggested by a reviewer, however, it may be desirable to provide an illustration of one such impact. To do so, we compute the implied impact of GT soybeans, through their impact on NT adoption, on soil erosion. We base our computation on Montgomery (2007, p. 13270), which compiles and presents results from 1,673 measurements of erosion rates under different settings.¹⁸ The median erosion rate from these measurements under conventional agriculture is about 1.5 mm/year, which is roughly 20 times the median erosion rate of 0.08 mm/year for conservation agriculture. The difference of ≈ 1.4 mm/year is equivalent to ≈ 6.8 tons/acre per year (assuming a soil bulk density of 1,200 kg/m³). Using the percent differentials for CT from table 6, and total annual U.S. acres planted to soybeans (source: Quick Stats at the USDA-NASS website), this implies a mean reduction in soil loss of 27 million tons per year. For context, estimated total soil erosion for U.S. cropland in 2007—assuming a mean erosion rate of 0.95 mm/year (Montgomery

¹⁷ Whereas in the text we have presented a constructive procedure to compute predicted adoption rates if GT soybeans were not available, we note that one could obtain the same results by considering the counterfactual in which CT and GT soybeans are independent. That is, the simulated adoption rates in table 6 are identical to those one would obtain by putting $\gamma = 0$ while maintaining the full choice set. The intuition for the equivalence is that when the seed and tillage practices are entirely independent, then each is chosen separately without regard to the other.

¹⁸ We alternatively considered computing implied soil loss using the Universal Soil Loss Equation (USLE), a widely used model for this purpose. However, use of this model requires detailed information (e.g., slope length and slope steepness) that are not available to us. Moreover, there are acknowledged problems with estimating soil loss based on the USLE (Trimble and Crosson 2000; Montgomery 2007).

2007, p. 13271) and given a total U.S. cropland of 408 million acres (USDA-ERS)—can be estimated at about was 1.9 billion tons. To assess the monetary value of these savings in soil erosion we use the USDA/NRCS estimated benefits of \$4.93 per ton in water quality improvements and \$1.93 per ton in saved fertilizer (USDA/NRCS 2009). Thus, the value of the benefits associated with the implied soil savings is \$189 million per year.

Conclusion

Complementarity is arguably a common feature among many of the inputs and practices chosen by agricultural producers. A possible instance of complementary in agriculture that has attracted considerable interest concerns the interaction between herbicide tolerant crops and conservation tillage practices. In this paper we have developed a new discrete choice model of joint practice adoption in which soybean producers choose among four tillage-soybean systems, and use it to investigate the existence and significance of complementarity between GT soybeans and CT practices. Our model explicitly incorporates both unobserved heterogeneity and complementarity, thus allowing for a direct test of whether GT soybeans and CT are complements. Using a large unbalanced panel dataset on individual farmers' choices spanning the period 1998–2011, we find that GT soybeans and CT are indeed complementary practices. This finding is robust to multiple specifications. Moreover, by ignoring unobserved heterogeneity, the degree of complementarity is overestimated. We further find that GT soybeans and no-till are likely stronger complements than GT soybeans and CT. In addition to the complementarity findings, our results indicate that highly erodible land, drought-like conditions, and higher fuel prices increase the likelihood of choosing CT. We also simulate annual adoption rates for CT and NT in a world without GT soybeans. The simulations indicate that CT adoption and NT adoption have been about 10% larger (or 6 percentage points) and 20% larger (9 percentage points), respectively, than what they would have been as a result of the availability of GT soybeans (holding total acreage fixed).

Whereas the framework of analysis that we propose and illustrate in this paper has broader methodological applicability to many issues in the economics of agricultural production, some policy implications follow immediately from our finding that GT soybeans and CT are complements. When complementarities are present, policy shocks that directly affect one activity will also indirectly affect complementary activities and will do so in the same direction. In recent years, for example, glyphosate weed resistance has become increasingly problematic in certain parts of the world (Powles 2008). As a result, there has been an initiative to slow that resistance in order to preserve the viability of glyphosate. Because GT soybeans and CT complement one another, such efforts also indirectly preserve the use of CT systems. A similar type of reasoning can be applied to the recent de-regulation of other herbicide tolerant crops (e.g., Dicamba resistant crops). To the extent that these crops also promote the use of CT, then their overall benefits are potentially under-estimated.

Concerning future research, an important question that remains unanswered relates to the effect of herbicide tolerant crops on herbicide use. Our framework could potentially be extended to look at this question by also incorporating the choice of how much herbicide to use. More generally, our framework could be used to consider relationships between a multitude of other agricultural choices, such as crop-rotation, farm size, row-spacing, and the type of machinery to purchase. For example, economies of scope at the farm level, rooted in the possible submodularity of a farm's cost structure (and so supermodularity of profits), represent an important possible application of our framework of analysis. Given the concerns associated with crop specialization and monoculture practices, especially vis-à-vis sustainability considerations, a deeper understanding of the complementarity relations that promote or hinder such trends would be valuable.

References

- Arora, A. 1996. Testing for Complementarities in Reduced-Form Regressions: A note. *Economics Letters*, 50(1), 51-55.
- Arora, A., and A. Gambardella. 1990. Complementarity and External Linkages: the Strategies of the Large Firms in Biotechnology. *The Journal of Industrial Economics*, 361-379.
- Athey, S., and S. Stern. 1998. An Empirical Framework for Testing Theories About Complementarity in Organizational Design (No. w6600). National Bureau of Economic Research.
- Banerjee, S., S. W. Martin, R. K. Roberts, J. A. Larson, R. J. Hogan Jr, J. L. Johnson, and J. M. Reeves. 2009. Adoption of Conservation-Tillage Practices and Herbicide-Resistant Seed in Cotton Production. *AgBioForum*, 12(3&4), 258-268.
- Barrows, G., S. Sexton, and D. Zilberman. 2014. Agricultural biotechnology: The Promise and Prospects of Genetically Modified Crops. *The Journal of Economic Perspectives*, 28(1), 99-119.
- Berry, S., A. Khwaja, V. Kumar, A. Musalem, K. Wilbur, G. Allenby, B. Anand, P. Chintagunta, M. Hanemann, P. Jeziorski and A. Mele. 2014. Structural Models of Complementary Choices. *Marketing Letters*, 25: 245-256.
- Blevins, R. L., and W. W. Frye. 1993. Conservation Tillage: an Ecological Approach to Soil Management. *Advances in Agronomy*, 51, 33-78.
- Brynjolfsson, E., and P. Milgrom. 2013. Complementarity in organizations. In R. Gibbons and J. Roberts, eds., *The Handbook of Organizational Economics*, Chapter 1, pp. 11-55. Princeton, NJ: Princeton University Press, 2013.
- Bull, L., H. Delve, C. Sandretto, and B. Lindamood. 1993. Analysis of Pesticide Use by Tillage System in 1990, 1991, and 1992 Corn and Soybeans. *Agricultural Resource Situation and Outlook Report 32*. Economic Research Service, U.S. Department of Agriculture, Washington DC, October 1993.
- Cameron, A. C., and P. K. Trivedi. 2010. *Microeconometrics Using Stata, Revised Edition*. Stata Press.

- Carpenter, J., and L. Gianessi. 1999. Herbicide Tolerant Soybeans: Why Growers are Adopting Roundup Ready Varieties. *AgBioForum*, 2(2), 65-72.
- Cassiman, B., and R. Veugelers. 2006. In Search of Complementarity in Innovation Strategy: Internal R&D and External Knowledge Acquisition. *Management Science*, 52(1), 68-82.
- Ding, Y., K. Schoengold, & T. Tadesse. 2009. The Impact of Weather Extremes on Agricultural Production Methods: Does Drought Increase Adoption of Conservation Tillage Practices? *Journal of Agricultural and Resource Economics*, 395-411.
- Dorfman, J. H. 1996. Modeling Multiple Adoption Decisions in a Joint Framework. *American Journal of Agricultural Economics*, 78(3), 547-557.
- Fawcett, R. S., B. R. Christensen, and D. P. Tierney. 1994. The Impact of Conservation Tillage on Pesticide Runoff into Surface Water: a Review and Analysis. *Journal of Soil and Water Conservation*, 49(2), 126-135.
- Fernandez-Cornejo, J. C. Hallahan, R. Nehring, and S. Wechsler. 2013. Conservation Tillage, Herbicide Use, and Genetically Engineered Crops in the United States: the Case of Soybeans. *AgBioForum*, 15(3): 231-241.
- Fernandez-Cornejo, J., C. Klotz-Ingram., and S. Jans. 2002. Farm-level Effects of Adopting Herbicide-Tolerant Soybeans in the USA. *Journal of Agricultural and Applied Economics*, 34(1), 149-164.
- Fernandez-Cornejo, J., C. Klotz-Ingram, R. Heimlich, M. Soule, W. McBride, and S. Jans. 2003. Economic and Environmental Impacts of Herbicide Tolerant and Insect Resistant Crops in the United States. In *The Economic and Environmental Impacts of Agbiotech* (pp. 63-88). Springer US.
- Fernandez-Cornejo, J., Wechsler, S.J., Livingston, M., L. Mitchell. 2014. *Genetically Engineered Crops in the United States*. Economic Research Report No. (ERR-162) 60 pp., February 2014. Washington, DC: USDA ERS.
- Frisvold, G., A. Boor, and J. Reeves. 2009. Simultaneous Diffusion of Herbicide Resistant Cotton and Conservation Tillage. *AgBioForum*, 12(3&4), 249-257.
- Fuglie, K.O. 1999. Conservation Tillage and Pesticide Use in the Cornbelt. *Journal of Agricultural and Applied Economics*, 31(1), 133-147.

- Gentzkow, M. 2007. Valuing New Goods in a Model with Complementarity: Online Newspapers. *American Economic Review*, 97(3), 713-744.
- Givens, W. A., D. R. Shaw, G. R. Kruger, W. G. Johnson, S. C. Weller, B. G. Young, R. G. Wilson, M.D.K. Owen, and D. Jordan. 2009. Survey of Tillage Trends Following the Adoption of Glyphosate-Resistant Crops. *Weed Technology*, 23(1), 150-155.
- Hensher, D. A., J. M. Rose, and W. H. Greene. 2005. *Applied Choice Analysis: a Primer*. Cambridge University Press.
- Hole, A. R. (2007). Estimating Mixed Logit Models Using Maximum Simulated Likelihood. *Stata Journal*, 7(3), 388-401.
- Holland, J. M. 2004. The Environmental Consequences of Adopting Conservation Tillage in Europe: Reviewing the Evidence. *Agriculture, Ecosystems & Environment*, 103(1), 1-25.
- Kalaitzandonakes, N., and P. Suntornpithug. 2003. Adoption of Cotton Biotechnology in the United States: Implications for Impact Assessment. In N. Kalaitzandonakes (Ed.), *The Economic and Environmental Impacts of Agbiotech: A Global Perspective* (pp. 103-118). New York: Kluwer Academic.
- Keane, M. P. 1992. A Note on Identification in the Multinomial Probit Model. *Journal of Business & Economic Statistics*, 10(2), 193-200.
- Kern, J. S. and M. G. Johnson. 1993. Conservation Tillage Impacts on National Soil and Atmospheric Carbon Levels. *Soil Science Society of America Journal*, 57(1), 200-210.
- Knowler, D., and B. Bradshaw. 2007. Farmers' adoption of conservation agriculture: A review and synthesis of recent research. *Food Policy*, 32(1), 25-48.
- Kretschmer, T., E. J. Miravete, and J.C. Pernías. 2012. Competitive Pressure and the Adoption of Complementary Innovations. *American Economic Review*, 102(4), 1540-70.
- Kurkalova, L., C. Kling, and J. Zhao. 2006. "Green Subsidies in Agriculture: Estimating the Adoption Costs of Conservation Tillage From Observed Behavior." *Canadian Journal of Agricultural Economics* 54: 247-267.
- Milgrom, P., and J. Roberts 1990. The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *The American Economic Review*, 80(3), 511-528.

- Miravete, E. J., and J. C. Pernías. 2006. Innovation Complementarity and Scale of Production. *The Journal of Industrial Economics*, 54(1), 1-29.
- Miravete, E. J., and J. C. Pernías. 2010. Testing for Complementarity When Strategies are Dichotomous. *Economics Letters*, 106(1), 28-31.
- Montgomery, D. R. 2007. Soil Erosion and Agricultural Sustainability. *Proceedings of the National Academy of Sciences*, 104(33), 13268-13272.
- Moschini, G. 2008. Biotechnology and the Development of Food Markets: Retrospect and Prospects. *European Review of Agricultural Economics*, 35(3), 331-355.
- 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.
- Powles, S. B. 2008. Evolved Glyphosate-Resistant Weeds Around the World: Lessons to be Learnt. *Pest Management Science*, 64(4), 360-365.
- Qaim, M. 2009. The Economics of Genetically Modified Crops. *Annual Review of Resource Economics*, 1(1), 665-694.
- Roberts, R. K., B. C. English, Q. Gao and J. A. Larson. 2006. Simultaneous Adoption of Herbicide-Resistant and Conservation Tillage Cotton Technologies. *Journal of Agricultural and Applied Economics*, 38, 629-643.
- Soule, M. J., A. Tegene, and K. D. Wiebe. 2000. Land Tenure and the Adoption of Conservation Practices. *American Journal of Agricultural Economics*, 82(4), 993-1005.
- Stubbs, M. 2012. Conservation Compliance and U.S. Farm Policy. R4259, *Congressional Research Service*.
- Train, K. 2000. Halton Sequences for Mixed Logit. Working Paper No. E00-278, Department of Economics, University of California, Berkeley.
- Train, K. E. 2009. *Discrete Choice Methods with Simulation*. Cambridge University Press.
- Trimble, S. W., and P. Crosson. 2000. US Soil Erosion Rates--Myth and Reality. *Science*, 289(5477), 248.
- Triplett, G. B., and W. A. Dick. 2008. No-Tillage Crop Production: A Revolution in Agriculture! *Agronomy Journal*, 100(Supplement 3), S-153.

- USDA. 2009. Interim Final Benefit-Cost Analysis for the Environmental Quality Incentives Program (EQIP). Report by the National Resource Conservation Service, United States Department of Agriculture, January 9.
- West, T. O., and G. Marland. 2002. A Synthesis of Carbon Sequestration, Carbon Emissions, and Net Carbon Flux in Agriculture: Comparing Tillage Practices in the United States. *Agriculture, Ecosystems & Environment*, 91(1), 217-232.
- Wu, J., and B. A. Babcock. 1998. The Choice of Tillage, Rotation, and Soil Testing Practices: Economic and Environmental Implications. *American Journal of Agricultural Economics*, 80(3), 494-511.
- Xu, Z., D.A. Hennessy, K. Sardana and G. Moschini, "The Realized Yield Effect of Genetically Engineered Crops: U.S. Maize and Soybean," *Crop Science*, 53(2013):735-745.

Appendix




Figure 1. Conservation Tillage and GT Adoption Rates for U.S. Soybeans (percent of acres)




Figure 2. U.S. Soybean Seed Prices, 1998-2011 (\$/50lb)

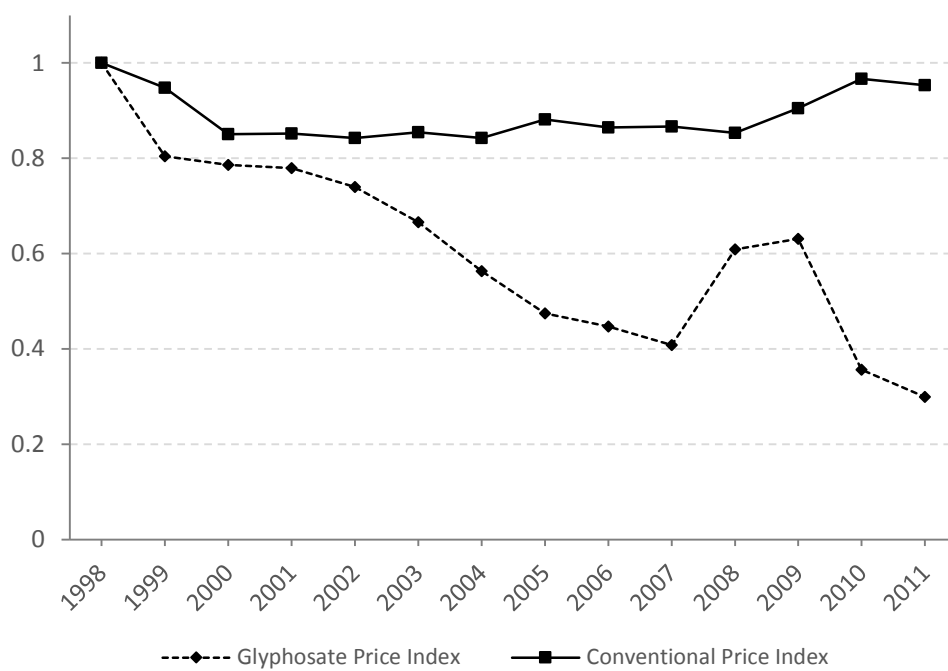


Figure 3. U.S. Soybean Herbicide Prices, 1998-2011

Table 1. Distribution of Tillage and Seed Systems (% of observations)

System	1998-2001	2002-2006	2007-2011	1998-2011
(CV,IT)	20.73	6.34	2.26	10.18
(GT,IT)	21.53	30.41	29.38	27
(CV,CT)	20.3	6.63	3.01	10.35
(GT,CT)	37.44	56.61	65.34	52.47
Observations	28,701	29,240	24,115	82,056

Table 2. Regressor Summary Statistics

Variable	Mean	S.D.	Min	0.25	Median	0.75	Max
<i>Size (>500 acres)</i>	0.33	0.47	0	0	0	1	1
<i>Futures (\$/bu)</i>	7.3	2.78	4.48	5.2	6.37	9.6	13.13
<i>Fuel Price Index</i>	49.96	24.33	19.6	29	43.8	65.4	91.2
<i>Erodibility Index</i>	8.36	9.49	0.29	2.67	5.2	11.32	192.07
<i>Palmer's Z-Index</i>	0.29	2.47	-4.93	-1.46	-0.11	1.48	11.84
<i>Seed Price (\$/50lb bag)</i>	8.98	1.92	6.34	7.46	8.67	9.84	12.41
<i>Herbicide Price Index</i>	-0.28	0.2	-0.65	-0.42	-0.26	-0.1	0

Table 3. Simulated Maximum Likelihood Results

Parameter (variable)	Coefficient	Standard Error
GT Adoption		
β_0^{GT} (constant)	1.5973***	(0.2060)
β_1 (Seed Price)	-0.3262***	(0.0249)
β_2 (Herbicide Price)	-0.9837***	(0.1565)
β_3^{GT} (Size)	0.1192***	(0.0418)
β_8^{GT} (Trend)	0.4420***	(0.0098)
CT Adoption		
β_0^{CT} (Constant)	-0.5710***	(0.1317)
β_3^{CT} (Size)	0.2850***	(0.0556)
β_4^{CT} (Fuel)	0.0069***	(0.0021)
β_5^{CT} (Futures)	-0.0255**	(0.0124)
β_6^{CT} (Erodibility)	0.0786***	(0.0117)
β_7^{CT} (Palmer)	-0.0237**	(0.0097)
β_8^{CT} (Trend)	0.0436***	(0.0093)
Other parameters		
γ	0.4609***	(0.0405)
σ_{GT}^2	2.2200***	(0.1097)
σ_{CT}^2	3.9186***	(0.2225)
$\sigma_{GT,CT}$	0.3094***	(0.0846)

Notes:

Number of observations = 82,056. Standard errors are clustered at the CRD level. Except for the covariance parameters, the coefficients are identified relative to ϕ , the scale parameter for ε_{itf}^j .

The covariance parameters are identified relative to ϕ^2 .

***Significant at the 1% level. **Significant at the 5% level.

Table 4. Average Elasticities

	GT(%)	CT(%)
<i>Seed Price</i>	-1.14 ^D	-0.08 ^I
<i>Herbicide Price</i>	-0.11 ^D	-0.01 ^I
<i>Soy Futures</i>	-0.002 ^I	-0.04 ^D
<i>Fuel Price</i>	0.002 ^I	0.07 ^D
<i>Palmer Z</i>	-0.0004 ^I	-0.01 ^D
<i>Erodibility Index</i>	0.01 ^I	0.10 ^D

Notes:

The reported effects are elasticities, i.e., the % change in the probability of adopting GT (CT) given a 1% change in the respective variable. See text for additional discussion.

^D = Direct Effect ; ^I = Indirect Effect.

Table 5. Alternative Estimates for Complementarity

Alternative Specifications	γ Coefficient	Standard Error
Include <i>Herbicide Price</i> in CT Variables	0.4143***	(0.0395)
No Correlation: $\sigma_{GT,CT} = 0$	0.5849***	(0.0322)
Ignore Panel Aspect of Data	1.3610**	(0.6699)
Basic Logit	0.5473***	(0.0333)
Restrict Sample to Central Corn Belt Only ^a	0.3039***	(0.0519)
No-Till or Till for Tillage Choice ^b	0.6514***	(0.0414)

Notes:

^a Includes Iowa, Illinois, and Indiana, for which there are 26,304 observations in all.

^b This variation specifies the tillage choice as being between no-till or a positive amount of tillage (rather than between conservation tillage and intensive tillage).

***Significant at the 1% level. **Significant at the 5% level.

Table 6 – Tillage Predicted Adoption Rates (percent of acres)

	Conservation Till Predicted					
	Rates			No-Till Predicted Rates		
	With GT	Without GT	Difference	With GT	Without GT	Difference
1998	53.9	50.6	3.3	31.0	27.0	4.0
1999	55.9	52.1	3.8	32.8	28.2	4.6
2000	57.6	53.3	4.2	35.8	30.5	5.4
2001	59.6	54.6	5.0	38.4	32.0	6.4
2002	59.9	54.7	5.3	38.6	31.8	6.8
2003	62.1	56.3	5.9	40.6	33.0	7.6
2004	62.0	56.0	6.0	40.4	32.6	7.7
2005	64.8	59.1	5.8	45.2	37.3	7.9
2006	66.6	60.7	5.9	48.0	39.8	8.2
2007	66.7	60.6	6.1	47.6	39.1	8.5
2008	68.7	62.7	6.0	48.8	40.2	8.6
2009	66.7	60.5	6.1	45.6	37.2	8.4
2010	68.7	62.6	6.1	48.1	39.4	8.7
2011	69.4	63.3	6.2	49.4	40.5	8.9

CHAPTER 3

GENETICALLY ENGINEERED CROPS AND PESTICIDE USE IN U.S. MAIZE AND SOYBEANS

Edward D. Perry, Federico Ciliberto, David A. Hennessy, and GianCarlo Moschini

Abstract

The widespread adoption of genetically engineered (GE) crops has clearly led to changes in pesticide use, but the nature and extent of these impacts remain open questions. We study this issue with a unique, large and representative sample of plot-level choices made by U.S. maize and soybean farmers from 1998 to 2011. On average, adopters of GE glyphosate tolerant (GT) soybeans used 28% (0.30 kg/ha) more herbicide than non-adopters, adopters of GT maize used 1.2% (0.03 kg/ha) less herbicide than non-adopters, and adopters of GE insect resistant (IR) maize used 11.2% (0.013 kg/ha) less insecticide than non-adopters. When pesticides are weighted by the environmental impact quotient (EIQ), however, we find that (relative to non-adopters) GE adopters used about the same amount of soybean herbicides, 9.8% less of maize herbicides, and 10.4% less of maize insecticides. In addition, the results indicate that the difference in pesticide use between GE and non-GE adopters has changed significantly over time. For both soybean and maize, GT adopters used increasingly more herbicides relative to non-adopters whereas adopters of IR maize used increasingly less insecticides. The estimated pattern of change in herbicide use over time is consistent with the emergence of glyphosate weed resistance.

Introduction

One of the most salient developments in global agriculture in the past twenty years has been the introduction of genetically engineered (GE) crop varieties (1–5). In the United States, in 2015, GE varieties accounted for 94% of planted soybean and 93% of planted maize (6). Adoption of this new technology was rapid: first introduced in 1996, GE soybean varieties embedding the glyphosate-tolerant (GT) trait have exceeded 80% of planted hectares since 2003. The share of planted maize using GE varieties—embedding GT and/or insect-resistant (IR) traits—has exceeded 80% since 2008. GT varieties are complementary inputs with glyphosate and their adoption has inevitably led to substitution away from other herbicides (7). Conversely, IR varieties can substitute for the use of insecticides, conceivably leading to lower pesticide use. Because pesticides have implications for human health and ecological diversity, factors that impact their use are of considerable policy interest (8–10). The nature and extent of the impact of GE variety adoption on pesticides use, however, remain open questions.

The prevailing consensus is that IR crops have significantly reduced insecticide use, but for herbicides the literature is divided (11, 12). Because most studies have lacked extensive survey data (11), a key issue has been how to impute counterfactual herbicide use for GE adopters. Some have used rates based on recommended conventional herbicide programs (13–15). Such recommended rates, however, are much larger than average observed herbicide usage rates prior to the advent of GE crops (9,10), so that, unsurprisingly, this method suggests large reductions in herbicide use due to GE adoption. Studies that instead rely on observed herbicide usage rates have hitherto been limited to one or two years of data, and in the earlier stages of GE crop adoption (16–18). As such, the generality of their results is limited, and they cannot shed light on whether the impact of GE variety adoption on pesticide use has changed over time. In particular, there has been little data to assess whether the recent development of glyphosate resistant weeds has eroded whatever herbicide use benefits there may have been from GT crops (11).

Our analysis relies on a unique, large farm-level dataset that spans the period 1998-2011. The data have been assembled annually by GfK Kynetec, a unit of a major market research organization that specializes in the collection of agriculture-related survey data. For each year, the samples are designed to be representative at the crop reporting district (CRD) level and include an annual average of 5,424 farmers for maize and 5,029 farmers for soybeans (Table S1). Based on these data, for each farmer we match the amount of pesticide used with the size of the corresponding plot, and the attributes of seed planted on that plot (including the type of GE traits embedded). Some farmers make more than one chemical/seed choice in any one year (i.e., they have more than one plot), and some (but not all) are observed for more than one year (Fig. S1). We are thus able to estimate the impact of GE crops on pesticide use by means of a fixed-effects regression analysis with observations on a large number of individual plot-level choices.

Results

Data on pesticide use and GE crop adoption in U.S. soybeans and maize are illustrated in Fig. 1. For maize, the share of varieties containing the GT trait (whether alone or stacked with IR traits) is reported separately from the share of varieties embedding one or more IR traits (henceforth *Bt* maize) (Fig. 1A). The rate of use of insecticides applied to maize fell from 0.2 kg/ha in 1998 to about 0.05 kg/ha in 2011, a 75% decrease (Fig. 1B). Since 1998, the most striking trend has been an increase in the use of glyphosate (Figs. 1C and 1D). By 2011, glyphosate dominated the soybean herbicide market with just over 80% of total herbicide applied, and in maize it accounted for nearly 40% of applied herbicide (a near twenty-fold increase from 1998). Increased glyphosate use came at the expense of other herbicides, though for soybeans there was also an increase in total herbicide use that began in 2007 and steadily rose through 2011.

The average rates in Fig. 1 are constructed by adding the amount of active ingredient (a.i.) of a large number of different chemicals. A concern with this (common) procedure is that the total weight associated with a bundle of heterogeneous chemicals is a poor measure of

environmental impact (19, 20). There is no agreed-upon superior procedure to aggregate heterogeneous pesticides. Following other studies (13, 14, 21), we use the Environmental Impact Quotient (EIQ) (22) as an alternative benchmark. Specifically, each a.i. is weighted by its EIQ value (23), and the resulting weighted sum is normalized so as to have the same overall mean as the unweighted total. Despite certain shortcomings (24), the EIQ's appeal in our context is that it converts an array of attributes specific to each pesticide into a single value meant to summarize the toxicity of the chemical. In general, re-weighting chemicals by their EIQ score does not significantly affect overall trends in pesticide use, except for soybeans where, from 1998 to 2005, the herbicide rate slightly increased but declined in the EIQ-weighted amount (Fig. 1).

To further investigate the impact of GE variety adoption on pesticide use, we use our plot-level data to estimate the fixed-effects regression model outlined in the Materials and Methods section. We consider two different measures of the amount of pesticides per unit of land applied by growers: unweighted sum of all a.i. used (kg/ha), and EIQ-weighted sum. The model is estimated separately for soybean herbicides, maize herbicides, and maize insecticides. For soybeans, we have a total of 86,736 plot-level observations, whereas for maize we have a total of 134,264 observations.

To assess the average impact over the entire 1998-2011 period, the fixed-effects model is first estimated under the restriction that the impact of GE varieties is constant over time, i.e., $\beta_t = \beta, \forall t$ (Table 1; full results in Table S2). Overall, GT soybeans increased the quantity of herbicides used by 0.30 kg/ha (a 28% increase relative to the average use by non-GT growers over the entire period). When herbicides are weighted by their EIQ score, however, the coefficient of the adoption variable is not significantly different from zero, reflecting the relatively lower EIQ values for glyphosate. For maize, GT adopters used about 0.03 kg/ha less herbicide (a 1.2% decline relative to the average overall use by non-GT growers). In EIQ terms, the savings were larger at 9.8%, again reflecting the relatively low EIQ values for glyphosate. With respect to insecticides, GE adopters of IR varieties used about 0.013 kg/ha less insecticide than non-adopters (a 11.2% decline relative to the average overall use by non-Bt adopters), a difference that is essentially unaffected by EIQ weighting.

The EIQ index is composed of three subcomponents: farmworker EIQ, which accounts for farmer exposure to dermal and chronic toxicity; consumer EIQ, which captures exposure to chronic toxicity and potential groundwater effects; and ecology EIQ, which captures the impacts of chemicals on fish, birds, bees, and beneficial arthropods (22). To gain further insight into the EIQ result in Table 1, we decompose the G_i coefficient into these three subcomponents. For all of soybean herbicides, corn herbicides, and maize insecticides, the farmworker and consumer components were lower on account of GE variety adoption. For the ecology component, maize herbicides and insecticides were improved by GE adoption, but for soybean herbicides, GE adoption had a detrimental effect (Table 2). Given that leaching potential and dermal toxicity are specific to the farmworker and consumer components, these results are broadly consistent with previous work that finds that herbicide usage patterns associated with GE varieties are beneficial (16, 18).

Next, we estimate the model where the β_t parameters are allowed to vary over time. The full results are reported in Table S3; here we graph the estimated β_t coefficients, along with their 95% confidence interval (Fig. 2). The impact of GT variety adoption on herbicide use has changed dramatically over time. In all periods, GT soybean adopters used more herbicide than non-adopters, and this difference increased considerably over time. By 2011, the amount applied by GT adopters was 0.66 kg/ha greater than non-adopters, an increase of 0.49 kg/ha from 1998. Moreover, although the total amount applied by a GT user was initially less harmful (as measured by the EIQ), from 2003 onward the reverse applied. The estimated trend for the impact of GT adoption for maize herbicides shows a similar pattern: over time GT adopters gradually employed more herbicide relative to conventional users, and by 2008 this difference was positive and statistically significantly greater than zero. Even when weighted by the EIQ impact, by 2011 GT adopters used more herbicide per hectare than non-adopters.

As for the impact of GE maize varieties embedding Bt traits, GE adopters used less insecticides than conventional growers for all years since 2000 (Fig. 2). The reduction in insecticide use attributable to the adoption of GE varieties increases (in absolute value) and becomes more significant (statistically) over time, possibly because of the diffusion of GE maize

varieties with multiple Bt traits (e.g., conveying resistance to corn rootworm, in addition to the European corn borer). In interpreting these results, however, one should bear in mind the possibility that Bt adoption might reduce the need for insecticide use by non-adopters as well, via an area-wide suppression effect, a conjecture supported by some evidence (25, 26).

Whereas Fig. 2 illustrates the estimated differential pesticides use by GE adopters relative to non-adopters, it is also of interest to investigate the underlying time trend of pesticide use by non-adopters. This information is conveyed by the year-specific intercepts of the estimate model. Fig. 3 graphs the estimated α_t coefficients (full results are in Table S3). For maize herbicides, there was a steady downward trend in herbicide use per hectare. Much of this downward trend can be explained by the decline of certain high-rate herbicides. For example, the active ingredient Metolachlor was supplanted by the lower-rate S-Metolachlor, and Cyanazine was phased out by the FDA (in cooperation with Dupont) by 2002 (Fig. S2). Other low-rate herbicides like Mesotrione also gained market penetration over the study period. For soybean herbicides, a downward trend also occurred early on, but the trend inverted in 2006. For maize insecticides, the use by non-adopters declined steadily up to 2007, stabilizing thereafter. This is broadly consistent with stylized facts concerning insecticide use in U.S. agriculture (8-10). More specifically, even before the introduction of Bt crops, there was a trend towards products with lower application rates. One class of low-rate insecticides that has been widely adopted recently are neonicotinoids, which are applied in the form of seed treatments. By 2011, our data indicate that nearly 50% of applied weight in insecticides took the form of seed treatments (Fig. 1B).

The robustness of the results obtained from the baseline model was investigated by considering several variations, including: the alternative where farmers' heterogeneity is instead represented by a random-effect model (Table S4); explicit accounting for the expansion of no-tillage practices (Table S5); explicit representation of plot-specific weed pressure (Table S6); accounting for selection bias due to the possible role of un-observed plot-level heterogeneity (Tables S7 and S8); and omission of choices associated with zero pesticide use (Tables S9 and S10). Details for each of these variations, and additional discussion, are provided

in the Supplementary Materials. Overall, the results of interest are essentially unchanged under these alternative specifications.

A clear result that emerges from our analysis is the change in differential herbicide use by GT adopters relative to non-GT adopters over time. What are the sources of such significant and persistent upward trends? Explanations such as the expansion of no-tillage practices, or unobserved plot-level heterogeneity, can be ruled out on the basis of the alternative specifications noted above. Part of the trend can be explained simply by the fact that non-adopters, particularly in maize, transitioned to lower-rate herbicides. But this cannot explain the sharp increase in later periods (specifically, 2007-2011). One explanation not ruled out by our investigations concerns the possible role of weed resistance. This is of particular interest, as glyphosate weed resistance has recently emerged as a significant concern (27–29). With GT crops, growers can apply glyphosate multiple times in a relatively short time span. Furthermore, the simultaneous availability of GT soybeans and GT maize has led to maize-soybean rotations that use glyphosate exclusively, thus significantly reducing the degree of chemical heterogeneity faced by weed populations (an important factor for preventing the emergence of herbicide tolerance) (29).

Making a direct link between our results and weed resistance is difficult because the data do not contain a plot-level variable that correlates with glyphosate weed resistance. To pursue an indirect inference route, however, we decompose the results in Fig. 2 by estimating the fixed-effects regressions separately for glyphosate and non-glyphosate herbicides. The underlying rationale for this procedure is that one of the early indicators of resistance would be a relative increase in the use of non-glyphosate herbicides by GE adopters. We find that for both soybeans and maize there has been a significant increase in non-glyphosate herbicides applied by GT adopters (relative to non-GT users) (Table S11). In soybeans, a GT adopter in 1998 used about 0.71 kg/ha less in non-glyphosate herbicides relative to a conventional user; by 2011, the difference was just 0.48 kg/ha (Fig. 4A). In maize, GT adopters went from using 1.31 kg/ha less in non-glyphosate herbicides in 1998 to only 0.32 kg/ha less in 2011 (Fig. 4B). The role of glyphosate weed resistance is also supported by data on the fraction of GT plots that relied

exclusively on glyphosate for weed control. Up to 2006, more than 70% of land planted with GT soybeans, and more than 40% of land planted with GT maize, were treated exclusively with glyphosate. But since then, these rates have dropped significantly, reaching lows of 41% (soybeans) and 19% (maize) (Fig. 4C).

Discussion and Conclusion

The role of GE crops in shaping the patterns of pesticide use remains a controversial topic. Over the period 1998-2011, our results show that GE variety adoption reduced both herbicide and insecticide use in maize, while increasing herbicide use in soybeans. Weighting pesticides by the EIQ, however, lowers the difference in herbicide use by GT soybean adopters (such that the estimated average impact over the study period is statistically indistinguishable from zero). Adoption of Bt maize, on the other hand, is associated with a clearer decline in insecticide use. This is broadly consistent with previous work (11-13, 17), although we find a smaller reduction. For herbicides, our results confirm the critical role of increased glyphosate use, but again we find less optimistic conclusions than other studies (13-15). These differences not only reflect the data that we use, but also the methodology of our study: unlike much of the existing work, our analysis relies on directly observed herbicide use for plots using GE and non-GE varieties, rather than arbitrarily constructed counterfactual use rates.

The richness of the data that we use, together with the methodology that we propose—with year-specific GE adoption effects, while controlling for the possible confounding effects of omitted variables via farmer fixed effects, year fixed effects, and regional trends—also permit us to characterize the time path of the GE variety adoption effects. Interestingly, we find clear evidence of increasing herbicides use by GT variety adopters over time for both soybeans and maize, a finding that we attribute in part to the emergence of glyphosate weed resistance. No such pattern appears for maize insecticide use over time, consistent with the evidence that non-Bt maize refugia have been broadly effective as a means to prevent the onset of pest resistance (30).

Materials and Methods

Data

The data used in this study come from AgroTrak®, a large, farm-level commercial dataset assembled by GfK Kynetec. Iowa State University acquired limited access to these proprietary data via a marketing research agreement with GfK Kynetec. Each year GfK Kynetec conducts surveys throughout the United States of randomly sampled farmers about decisions pertaining to seed and pesticide choices. The samples constructed for AgroTrak® are representative at the crop reporting district (CRD) level. Each CRD is a multi-county area identified by the National Agricultural Statistics Service of the USDA (Fig. S3). Agrotrak® is widely considered to be the most comprehensive source for these data and has been used in several other studies (31-33).

The subset of Agrotrak® utilized in this analysis pertains to pesticide use by U.S. soybean and maize farmers during the 14-year period 1998-2011. Over this period, on average the surveys included 5,029 farmers per year for soybeans and 5,424 farmers per year for maize. For each crop, respondents indicate how much land is planted, with what seed trait, and the type of tillage used. A grower's land planted with the same seed trait (e.g., GT soybeans) and with the same tillage method (conventional, conservation or no till) defines a "plot" for the purpose of our analysis. Over the 14-year period, we identify a total of 86,736 plots for soybeans, and a total of 134,264 plots for maize. For each of these plots, Agrotrak® provides sufficient information to reconstruct the amount of all commercial pesticides products applied by the farmer. By utilizing the table providing each product's a.i., also in the dataset, we calculate the total amount of pesticides used on each plot.

We use two measures of pesticide use for each plot. The first measure is the total amount of all active ingredients used on the plot. Specifically, if Q_i^k denotes the quantity of commercial products k applied on plot i , with a per-unit content a_j^k of a.i. j , and L_i denotes the land size of plot i , our first plot-level measure of pesticide use (kg/hectare) is defined as

$$y_i^A = \frac{1}{L_i} \sum_k \sum_j Q_i^k a_j^k \quad (25)$$

The second measure of total pesticide use per plot is meant to address the composition heterogeneity of commercial pesticides by weighing active ingredients by their EIQ values. The latter are obtained from the list in (23), updated in 2012. Specifically, if E_j is the EIQ value associated with a.i. j , the EIQ measure of plot-level pesticide use is defined as

$$y_i^E = \frac{\kappa}{L_i} \sum_k \sum_j Q_i^k E_j a_j^k \quad (26)$$

where κ is a normalizing constant chosen such that y_i^E and y_i^A have the same overall mean (this facilitates comparison of regression coefficients obtained from these two alternative measures of pesticide use).

Table S1 and Fig. S1 contain some summary statistics of the structure of the AgroTrak® data used in this study. An important feature of the GfK dataset is that it contains repeated observations across time for a subset of the growers. Of the 38,693 farmers in the sample, over 50% were sampled two or more years, and more than 30% were sampled for at least three years. This is a key element that permits us to estimate a model that controls for the possible impact of unobserved farmer-level heterogeneity.

Model

The main results of the analysis are based on the following fixed-effects regression model, which is estimated separately for herbicides and insecticides and for each of the two crops of interest (maize and soybeans):

$$y_i = \alpha_{t[i]} + \beta_{t[i]} G_i + \gamma_{r[i]} T_{t[i]} + \phi_{f[i]} + e_i, \quad i = 1, 2, \dots, N \quad (27)$$

where i indexes the plot, N is the total number of observations (thus, $N = 86,736$ for soybeans and $N = 134,264$ for maize), $t[i]$ identifies the year in which data for plot i are observed, $r[i]$ denotes the region (i.e., the CRD) of the plot and $f[i]$ indicates the farmer to whom the plot belongs (the notation follows (34)). As noted, we consider two different measures for the dependent variable, and thus either $y_i = y_i^A$ or $y_i = y_i^E$. The main independent variable of

interest, G_i , is a binary variable that equals one if plot i was planted with a GE variety, and zero otherwise. For soybean and maize herbicides, $G_i = 1$ if the variety embeds a GT trait, and for maize insecticide $G_i = 1$ if the variety contained one or more IR traits (i.e., Bt maize). The year-specific β_t parameters, our main focus, capture the impact of adopting GE crops on pesticide use. This impact is estimated relative to the underlying benchmark of pesticide use on non-GE plots captured by the time fixed effects α_t . The remaining terms are grower-specific effects, denoted by ϕ_f , and CRD-specific time trends, denoted by $\gamma_r T_t$ (here T_t is a linear time trend, suitably demeaned so that the estimated α_t can be interpreted as the average use of pesticide on non-GE plots).

The identifying assumption for estimation is that—conditional on the fixed effects and regional trends— G_i is exogenous with respect to e_i , that is $E[e_i | G_i] = 0$. To illustrate the statistical basis of this assumption, consider the simplest possible specification of time-varying GE-adoption effects:

$$y_i = \delta + \beta_{t[i]} G_i + u_i, \quad i = 1, \dots, N,$$

where i indexes the plot and N is the total number of plots in the sample (the notation $t[i]$ signifies that observation of plot i pertains to year t), and u_i is the error term. Unbiased estimation of the β_t parameters of interest would require the mean independence condition $E(u_i | G_i) = 0$. This is likely to be violated because the term u_i may reflect unobserved variables that are correlated with G_i . To deal with this problem, our strategy is to decompose the error term as follows:

$$u_i \equiv \phi_{f[i]} + (\alpha_{t[i]} - \delta) + \gamma_{r[i]} T_{t[i]} + e_i,$$

where $\phi_{f[i]}$ is a time-invariant farmer fixed effect, $\alpha_{t[i]}$ is a time-specific fixed effect, $\gamma_{r[i]} T_{t[i]}$ is a CRD-specific time trend, and e_i is a residual error term. The term $\phi_{f[i]}$ captures time-invariant unobserved farm-level variables (such as, for example, a grower's education, location, and attitude towards the environment). The terms $\alpha_{t[i]}$ capture the impact of variables that may

affect pesticide use but that can be safely presumed constant in a given year (such as herbicide prices and crop output prices). The $\gamma_{r[i]}T_{t[i]}$ capture the effect of possible location-specific variables that change over time. By using the panel structure of our data to explicitly estimate these terms, we then only require the more plausible assumption $E(e_i | G_i) = 0$. Because the e_i terms now reflect variables that are plot-specific and have no trend at the regional level over time, this assumption seems reasonable. We note that this approach is common in comparable existing analyses (e.g., Kathage and Qaim, *PNAS* 2012). Moreover, we report a number of sensitivity analyses in which we investigate how the $\beta_{t[i]}$ estimates change when additional plot-specific variables are added to the model (such as the type of tillage used or the type of weed targeted). Overall, the estimates are largely unaffected by these changes (see tables S5-S9).

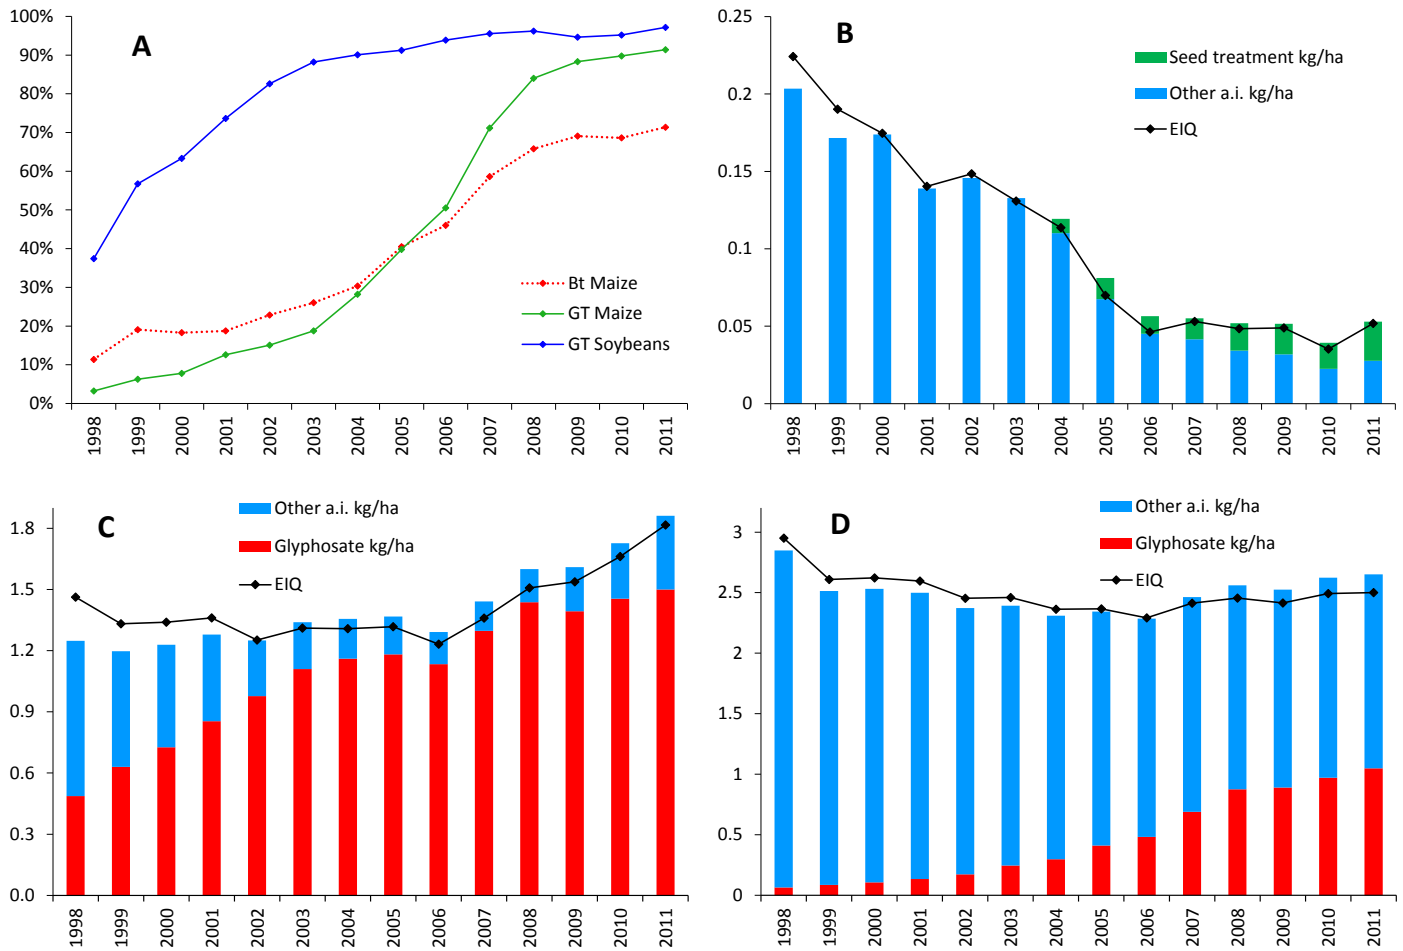


Fig. 1. GE variety adoption and pesticide use, maize and soybeans in the United States, 1998-2011. (A) Adoption rates of glyphosate tolerant (GT) soybeans, GT maize and Bt maize (embedding one or more genes from *Bacillus Thuringiensis*). (B) Insecticide use in maize, kg/ha and Environmental Impact Quotient (EIQ) weights. (C) Herbicide use in soybeans, kg/ha and EIQ weights. (D) Herbicide use in maize, kg/ha and EIQ weights. Adoption rates and a.i. use (kg/ha) are reported in tables S12 and S13









Fig. 2. Estimated β_t parameters from the fixed-effects model. (A) Year-specific impacts of glyphosate tolerant (GT) soybeans on herbicide use, kg/ha and Environmental Impact Quotient (EIQ) weights. (B) Year-specific impacts of GT maize on herbicide use, kg/ha and EIQ weights. (C) Year-specific impacts of Bt maize on insecticide use, kg/ha and EIQ weights. For all panels, vertical bars denote 95% confidence intervals.




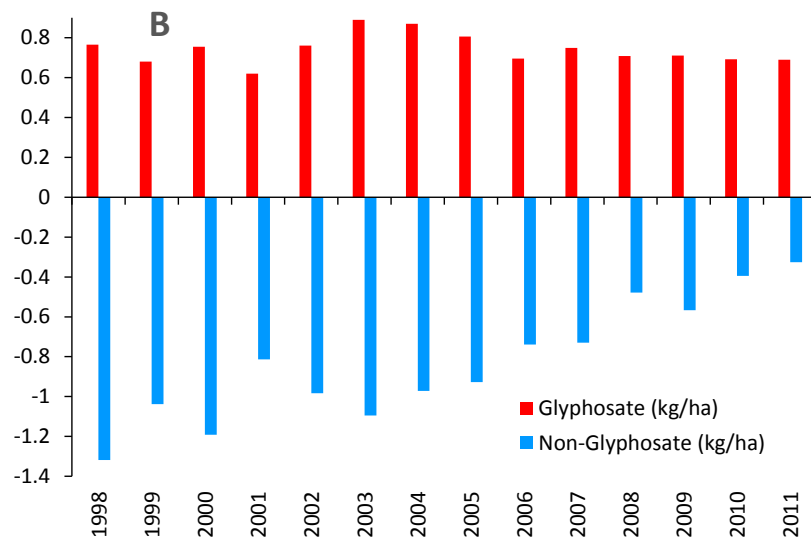



Fig. 3. Estimated α_t parameters from the fixed-effects model. (A) Year-specific herbicide use by non-GT soybean adopters, kg/ha and Environmental Impact Quotient (EIQ) weights. (B) Year-specific herbicide use by non-GT maize adopters, kg/ha and EIQ weights. (C) Year-specific insecticide use by non-Bt maize adopters, kg/ha and EIQ weights. For all panels, vertical bars denote 95% confidence intervals.






Fig. 4. Decomposition of year-specific impacts of GE variety adoption. (A) Differences in herbicide use between glyphosate tolerant (GT) soybean adopters and non-adopters, kg/ha, glyphosate (blue bars) and all other herbicides (red bars). (B) Differences in herbicide use between GT maize adopters and non-adopters, kg/ha, glyphosate (blue bars) and all other herbicides (red bars). (C) Fraction of hectares planted to GT varieties that use exclusively glyphosate.

Table 1. Estimated impact of GE varieties on pesticide use, average impact over 1998-2011

(assumes $\beta_t = \beta, \forall t$). N = number of observations. Standard errors (in parentheses) are clustered at the farmer level. The model includes time fixed effects, CRD-specific time trends and individual (farmer) fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	Soybean Herbicides		Maize Herbicides		Maize Insecticides	
	a.i.	EIQ	a.i.	EIQ	a.i.	EIQ
	kg/ha	kg/ha	kg/ha	kg/ha	kg/ha	kg/ha
G_i	0.3021*** (0.0097)	0.0045 (0.0122)	-0.0329* (0.0150)	-0.2590*** (0.0156)	-0.0129*** (0.0014)	-0.0122*** (0.0014)
N	86,736	86,736	134,264	134,264	134,264	134,264
R^2	0.067	0.028	0.022	0.027	0.039	0.051

Table 2. Estimated impact of GE varieties on the farmer, consumer and ecology components

of EIQ-weighted pesticide use, average impact over 1998-2011 (assumes $\beta_t = \beta, \forall t$). N = number of observations. Standard errors (in parentheses) are clustered at the farmer level. The model includes time fixed effects, CRD-specific time trends and individual (farmer) fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	Soy Herbicide EIQ			Maize Herbicide EIQ			Maize Insecticide EIQ		
	Farmer	Consumer	Ecology	Farmer	Consumer	Ecology	Farmer	Consumer	Ecology
G_i	-0.0081*** (0.0021)	-0.0281*** (0.0013)	0.0407*** (0.0091)	-0.0301*** (0.0024)	-0.0534*** (0.0017)	-0.1755*** (0.0116)	-0.0019*** (0.0003)	-0.0003*** (0.0001)	-0.0100*** (0.0011)
N	86,736	86,736	86,736	134,264	134,264	134,264	134,264	134,264	134,264
R^2	0.034	0.051	0.027	0.029	0.048	0.025	0.041	0.027	0.053

References

1. M. Qaim, D. Zilberman, Yield effects of genetically modified crops in developing countries. *Science*. **299**, 900-902 (2003).
2. G. Moschini, Biotechnology and the development of food markets: retrospect and prospects. *Eur. Rev. Agric. Econ.* **35**, 331-355 (2008).
3. J. Kathage, M. Qaim, Economic impacts and impact dynamics of Bt (*Bacillus thuringiensis*) cotton in India. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 11652-11656 (2012).
4. Z. Xu, D.A. Hennessy, K. Sardana, G. Moschini, The realized yield effect of genetically engineered crops: US maize and soybean. *Crop Sci.* **53**, 735-745 (2013).
5. G. Barrows, S. Sexton, D. Zilberman, Agricultural biotechnology: The promise and prospects of genetically modified crops. *J. Econ. Perspect.* **28**, 99-119 (2014).
6. USDA-ERS, Genetically engineered varieties of corn, upland cotton, and soybeans, by State and for the United States, 2000-15, www.ers.usda.gov/data-products/adoption-of-genetically-engineered-crops-in-the-us.aspx.
7. J. Fernandez-Cornejo, S. Wechsler, M. Livingston, L. Mitchell, "Genetically engineered crops in the United States" (USDA-ERS Econ. Res. Rep. 162, 2014).
8. National Research Council, *Impact of Genetically Engineered Crops on Farm Sustainability in the United States*, Washington, DC: National Academies Press, 2010.
9. C. D. Osteen, J. Fernandez-Cornejo, Economic and policy issues of US agricultural pesticide use trends. *Pest Manag. Sci.* **69**, 1001-1025 (2013).
10. J. Fernandez-Cornejo, R. Nehring, C. Osteen, S. Wechsler, A. Martin, A. Vialou, "Pesticide use in US agriculture: 21 selected crops, 1960-2008" (USDA-ERS, Econ. Info. Bull. 124, 2014; www.ers.usda.gov/publications/eib-economic-information-bulletin/eib124.aspx).
11. J.E Carpenter, Peer-reviewed surveys indicate positive impact of commercialized GM crops. *Nature Biotechnol.* **28**, 319-321 (2010).
12. W. Klümper, M. Qaim, A meta-analysis of the impacts of genetically modified crops. *PLoS One*, 9(11), pp. 111629 (2014).

13. G. Brookes, P. Barfoot, Environmental impacts of genetically modified (GM) crop use 1996–2013: Impacts on pesticide use and carbon emissions. *GM Crops Food*. **6**, 103-133 (2015).
14. G. A. Kleter, R. Bhula, K. Bodnaruk, E. Carazo, A. S. Felsot, C. A. Harris, A. Katayama, A. H. Kuiper, D. K. Racke, B. Rubin, Y. Shevah, G. R. Stephenson, K. Tanaka, J. Unsworth, R. D. Wauchope, S. S. Wong, Altered pesticide use on transgenic crops and the associated general impact from an environmental perspective. *Pest Manag. Sci.* **63**, 1107-1115 (2007).
15. S. R. Johnson, S. Strom, K. Grillo, “Quantification of the impacts on US agriculture of biotechnology-derived crops planted in 2006” (National Center for Food and Agricultural Policy, 2007).
16. M. Qaim, G. Traxler, Roundup Ready soybeans in Argentina: farm level and aggregate welfare effects. *Agr. Econ.* **32**, 73-86 (2005).
17. M.G. Cattaneo, C. Yafuso, C. Schmidt, C.Y. Huang, M. Rahman, C. Olson, C. Ellers-Kirk, B.J. Orr, S.E. Marsh, L. Antilla, P. Dutilleul, Farm-scale evaluation of the impacts of transgenic cotton on biodiversity, pesticide use, and yield. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 7571-7576 (2006).
18. A. Wossink, Z.S. Denaux, Environmental and cost efficiency of pesticide use in transgenic and conventional cotton production. *Agr. Syst.* **90**, 312-328 (2006).
19. J. Carpenter, GM crops and patterns of pesticide use. *Science*, **292**, 637-638 (2001).
20. S.O. Duke, S.B. Powles, Glyphosate: a once-in-a-century herbicide. *Pest Manag. Sci.* **64**, 319-325 (2008).
21. T. Brimner, G. Gallivan, G. Stephenson, Influence of herbicide-resistant canola on the environmental impact of weed management. *Pest Manag. Sci.* **61**, 47-52 (2005).
22. J. Kovach, C. Petzoldt, J. Degni, J. Tette, “A method to measure the environmental impact of pesticides” (NY and Life Sci. Bull. 139, 1992).
23. B. Eshenaur, J. Grand, J. Kovach, C. Petzoldt, J. Degni, J. Tette, Environmental Impact Quotient: “A Method to Measure the Environmental Impact of Pesticides.” (New York

Integrated Pest Management Program, Cornell Cooperative Extension, Cornell University, 1992-2015, <http://www.nysipm.cornell.edu/publications/EIQ/>).

24. A. R. Kniss, C. W. Coburn, Quantitative evaluation of the environmental impact quotient (EIQ) for comparing herbicides. *PLoS ONE*. **10**, e0131200. (2015).
25. W.D. Hutchison, E.C. Burkness, P.D. Mitchell, R.D. Moon, T.W. Leslie, S.J. Fleischer, M. Abrahamson, K.L. Hamilton, K.L. Steffey, M.E. Gray, R.L. Hellmich, Areawide suppression of European corn borer with Bt maize reaps savings to non-Bt maize growers. *Science*. **330**, 222-225 (2010).
26. Y. Lu, K. Wu, Y. Jiang, Y. Guo, N. Desneux, Widespread adoption of Bt cotton and insecticide decrease promotes biocontrol services. *Nature*. **487**, 362-365 (2012).
27. S.B. Powles, Evolved glyphosate-resistant weeds around the world: lessons to be learnt. *Pest Manag. Sci.* **64**, 360-365 (2008).
28. N. Gilbert, A hard look at GM crops. *Nature*. **497**, 24-26 (2013).
29. M. Livingston, J. Fernandez-Cornejo, J. Unger, C. Osteen, D. Schimmelpfennig, T. Park, D. Lambert, "The Economics of Glyphosate Resistance Management in Corn and Soybean Production" (USDA-ERS, Econ. Res. Rep. 184, 2015; <http://www.ers.usda.gov/publications/err-economic-research-report/err184>).
30. B.E. Tabashnik, Delaying insect resistance to transgenic crops. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 19029-19030 (2008).
31. S. Gangwal, D.M. Reif, S. Mosher, P.P. Egeghy, J.F. Wambaugh, R.S. Judson, E.A.C. Hubal, Incorporating exposure information into the toxicological prioritization index decision support framework. *Sci. Total Environ.* **435**, 316-325 (2012).
32. G.P. Thelin, W.W. Stone, "Estimation of annual agricultural pesticide use for counties of the conterminous United States, 1992-2009" (US Department of the Interior, US Geological Survey, Sci. Inv. Rep. 2013-5009, 2013).
33. P.D. Mitchell, Market-level assessment of the economic benefits of atrazine in the United States. *Pest Manag. Sci.* **70**, 1684-1696 (2014).
34. A. Gelman, J. Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge, UK: Cambridge University Press, 2007.

35. M. Qaim, The economics of genetically modified crops. *Annu. Rev. Resour. Econ.* **1**, 665-694 (2009).
36. M. Qaim, A. De Janvry, Bt cotton and pesticide use in Argentina: Economic and environmental effects. *Environ. Dev. Econ.* **10**, 179-200 (2005).
37. S. Bonny, Genetically modified glyphosate-tolerant soybean in the USA: adoption factors, impacts and prospects. A review. *Agron. Sustain. Dev.*, **28**, 21-32 (2008).
38. J. Fernandez-Cornejo, C. Klotz-Ingram., S. Jans, Farm-level effects of adopting herbicide-tolerant soybeans in the USA. *J. Agr. Appl. Econ.*, **34**, 149-164 (2002).
39. C. M. Benbrook, Impacts of genetically engineered crops on pesticide use in the US—the first sixteen years. *Environ. Sci. Eur.* **24**, (2012).
40. G. Brookes, J. Carpenter, A. McHughen, A review and assessment of 'Impact of genetically engineered crops on pesticide use in the US—the first sixteen years: Benbrook C (2012)'. (2012); available at: <http://www.ask-force.org/web/Pesticides/Brookes-Carpenter-Hughes-Rebuttal-Benbrook-2012.pdf> (accessed 3-5-2016)
41. O. Sydorovych, M.C. Marra, A genetically engineered crop's impact on pesticide use: A revealed-preference index approach. *J. Agr. Resour. Econ.* **32**, 476-491 (2007).
42. G.C. Nelson, D.S. Bullock, Simulating a relative environmental effect of glyphosate-resistant soybeans. *Ecol. Econ.*, **45**, 189-202 (2003).
43. J. Fernandez-Cornejo, S. Jans, Quality-adjusted price and quantity indices for pesticides. *Am. J. Agr. Econ.* **77**, 645-659 (1995).
44. J. Fernandez-Cornejo, C. Hallahan, R. Nehring, S. Wechsler, A. Grube, Conservation tillage, herbicide use, and genetically engineered crops in the United States: the case of soybeans. *AgBioForum* **15**, 231-241 (2012).
45. E.D. Perry, G. Moschini, D.A. Hennessy, Testing for complementarity: glyphosate tolerant soybeans and conservation tillage. *Am. J. Agr. Econ.* **98**, 765-784 (2016).

Appendix

Content of this Appendix:

Supplementary Text

Fig. S1. Number of Years Sampled for Growers in AgroTrak® dataset

Fig. S2. Maize Herbicide Use by non-GT Adopters (selected herbicides, kg/ha)

Fig. S3. Crop Reporting Districts (CRD)

Fig. S4. Trends in Glyphosate and Expected Crop Output Prices, 1998-2011

Table S1. Summary Statistics for AgroTrak® Dataset

Table S2. Full Results Corresponding to Table 1 in the Main Text

Table S3. Full Results Corresponding to Fig. 2 and Fig. 3 in the Main Text

Table S4. Random Effects replace Farmers Fixed Effects

Table S5. Model Estimates with the No-Till Binary Variable Included

Table S6. Targeted weeds and impact of GE variety adoption on herbicide use (kg/ha of a.i.)

Table S7. Model excludes growers that plant both GE and non-GE varieties within a given year

Table S8. Model excludes growers that plant both GE and non-GE varieties within a given year

Table S9. Model excludes farmers that never used pesticides (on any of their plots)

Table S10. Model excludes farmers that never used pesticides (on any of their plots)

Table S11. Full Set of Results Corresponding to Fig. 4 in the Main Text

Table S12. U.S. GE Adoption Rates (% of planted hectares), 1998-2011

Table S13. U.S. Pesticide Rates (kg/ha), 1998-2011

Table S14. Correlation Between State-Level GE Adoption Rates from USDA and GfK Data

Table S15. Summary Statistics by Adoption Choice

Supplementary text

The AgroTrak® data

The data used in this study comes from AgroTrak®, a large, farm-level commercial dataset assembled by GfK Kynetec. Iowa State University acquired limited access to these proprietary data via a marketing research agreement with GfK Kynetec. Each year GfK Kynetec conducts surveys throughout the United States of randomly sampled farmers about decisions pertaining to seed and pesticide choices. The samples constructed for AgroTrak® are representative at the crop reporting district (CRD) level. Each CRD is a multi-county area identified by the National Agricultural Statistics Service of the USDA (Fig. S2). Table S1 and Fig. S1 contains some summary statistics of the structure of the AgroTrak® data used in this study. An important feature of the GfK dataset is that it contains repeated observations across time for a subset of the growers. Of the 38,693 farmers in the sample, over 50% were sampled 2 or more years, and more than 30% were sampled for at least 3 years.

Agrotrak® is widely considered the most comprehensive source for pesticide use data and has been used in several other studies, including Gangwal et al. (31), Thelin and Stone (32), and Mitchell (33). Concerning farmers' use of GE varieties, also documented in Agrotrak®, we note that estimates of GE crop variety adoption have been independently reported by the USDA (based on National Agricultural Statistics Service surveys) since 2000 (6, 7). This provides the opportunity for an additional external validation of some of the proprietary data used in this study. To do so, we compared state-level GE crop adoption rates reported by the USDA to state-level GE crop adoption rates computed from AgroTrak®. In the manuscript we use adoption rates for varieties that contain the GE trait(s) of interest (e.g., varieties that contain the GT trait). Some of those varieties may incidentally contain other GE traits as well, e.g., Bt traits. The USDA does not report adoption rates for maize varieties that have GE herbicide tolerance, whether alone or stacked with another GE trait (e.g., Bt). Rather, they report the adoption rate for maize varieties with GE herbicide-tolerance *only*. They also report the adoption rates for all GE varieties. As a result, we compute what we believe to be the comparable adoption rates from

the GfK data. Table S14 reports the correlation between these two types of adoption rates at the state level for US maize and soybeans. Overall, they are highly correlated.

Literature review: GE variety adoption and pesticide use

Given the breadth of the literature on GE crops and pesticide use, this section only focuses on those studies most relevant to our analysis. Somewhat more comprehensive literature reviews can be found in Carpenter (11), Klümper and Qaim (12), and Qaim (35). We consider the prior literature from three different perspectives: (i) their findings, (ii) the data used, and (iii) the methods they employ. The literature relevant to our analysis for maize insecticides is discussed first.

Because Bt crops do not relate to any one particular insecticide, conclusions about their environmental impact are fairly straightforward: if they reduce insecticide use the environment is the better for it (and vice versa). Overall, most studies have found that Bt crop adopters use less insecticide than non-adopters (11-13, 17, 35, 36). Drawing on large number of studies Klumper and Qaim (12) find that these savings are on average 37%. A less studied issue has been the potential benefits reaped by non-Bt growers from Bt adopters. Results in Hutchison et al. (25) reveal that non-Bt growers benefited from Bt adopters through the associated suppression of the European Corn Borer population. Whether this has led to a reduction in insecticide use, however, has not been studied. A basic statistical trend in favor of this effect is that non-Bt maize adopters significantly reduced insecticide use as the adoption of Bt maize rose (7).

The complementary relation between GT crop varieties and glyphosate use implies a more complex characterization with respect to environmental impact. At the initial stages of commercialization of GT crops the basic question could be reduced to whether the increase in glyphosate use exceeded the decrease in a number of more narrow-spectrum herbicides, and whether that net change was better or worse for the environment. Some early studies found that adopters of GT soybeans and/or GT maize used less herbicide than non-GT adopters (see Table 4 in Fernandez-Cornejo et al. (7)). In more recent years, however, that trend seems to have reversed with GT growers typically using more herbicide in terms of weight (37). By most

environmental measures, however, that greater amount of herbicide – in particular glyphosate – was an improvement over the lesser amount used by non-adopters (16, 18, 38). Whether this has remained true more recently (beyond 2006) has been less studied. Moreover, the question has been complicated by the emergence of glyphosate weed resistance, which has brought back the use of some previously abandoned, narrow-spectrum herbicides (37). More recently, Benbrook (39) finds that GT soybeans are sprayed with significantly more herbicide than non-GT soybeans; however, certain limitations of these findings have been noted by Brookes, Carpenter, and McHughen (40). In brief, Benbrook (39) relies on USDA data that does not disaggregate pesticide use by GE trait, and thus his findings critically depend on somewhat arbitrary assumptions about how that use is broken down.

The most widely cited source on this issue has been a series of studies conducted by Brookes and Barfoot, the most recent of which is Brookes and Barfoot (13). These studies are of particular interest to our analysis because they use some of the same data that we employ. In general, they report significant reductions in herbicide use from GE crops, even during some of the more recent periods in which glyphosate weed resistance has reportedly intensified. We note two important limitations of their analysis. First, they do not control for unobserved heterogeneity across farmers: in general, they compare unconditional annual average herbicide usage rates between GE and non-GE adopters. Second, the procedure they use to compute those average rates in part relies on strong assumptions about counterfactual pesticide use. For years during which non-GE adopters comprised less than 50% of the population, rather than use observed herbicide usage rates by non-GE adopters for the counterfactual they use the rates implied by various recommended conventional herbicide programs that would achieve a level of weed control similar to that in GT crops (this method is also used, e.g., in Kleter et al. (14) and Johnson, Strom, and Grillo (15)). As we note in the paper, the usage rates implied by these programs significantly exceed average herbicide rates observed prior to the GE era (based on USDA data). The discrepancy between the recommended rates and the historically observed rates is likely due to the fact that the profit maximizing amount of herbicides for a non-GT user is less than the amount that would achieve the same level of weed control for a GT crop. In

general, the appropriate counterfactual should be one in which the adoption of GT crops does not exert any indirect or direct influence on the choice (otherwise it would be part of the effect).

With regard to data, most survey-based studies use samples that are restricted to one or two years prior to 2006 (16–18, 37, 39, 41). Exceptions are Benbrook (39) and Brookes and Barfoot (13), but both of these studies conduct analyses that are not at the farm level. As a result, there has not yet been a farm-level, survey-based study that extends from the beginning of the GE crop era into the early stages of glyphosate weed resistance (Kathage and Qaim (3) conduct a multi-year farm-level analysis of Bt crops in India, but pesticide use is not one of the variables they consider).

The EIQ is one among several methods to aggregate and/or measure the environmental impact of pesticides. Various studies have employed alternative procedures or measures (16, 18, 41–44). Two of these aggregation procedures, however, do not explicitly capture external environmental impacts (41, 43), and were thus not considered in this study. Among the remaining studies, Nelson and Bullock (42) use the LD50 dose for rates, Wossink and Denaux (18) use leaching potential, and Qaim and Traxler (16) break herbicides down by toxicity class. Each of these measures is to some extent captured in the various components that make-up the EIQ (e.g., leaching potential and dermal toxicity are in consumer and farmworker components), and depending on the analysis, one may be more desirable to use than another. Below we show and discuss how some of these finer measures are impacted by GE crops.

Details for the results reported in the main text

Table S2 contain the full set of estimates for Table 1, Table S3 contains the full set of estimates for Fig. 2 and Fig. 3, and Table S11 contains the full set of estimates for Fig. 4A and Fig. 4B.

Supplementary Results

Individual (farmer) random effects

Table S4 contains results for the same specification as in Table S2 (which provides details for the results of Fig. 2 in the paper), but with the farmer-specific fixed effects replaced by random effects. With individual fixed effects, growers who are sampled only once and with only one

plot (this accounts for 13.5% of the observations for soybeans and 6.9% of the observations for maize) do not contribute to estimating the β_t coefficients. With the random effects model, all observations contribute to estimating the β_t coefficients. The limitation of random effects is that if they are correlated with the observables then the estimated coefficients are not consistent. Comparison of Table S4 with Table S2 indicates that the estimated β_t coefficients are hardly affected by the choice of how one models individual heterogeneity. However, the Hausman test, which compares the difference in coefficient estimates for all variables (also reported in Table 4), rejects the random effects model in favor of the fixed effects model.

Herbicide prices

The results provided in Table S11 indicate that both GT and non-GT adopters increased their use of glyphosate over time. These trends can in part be explained by changes in herbicide and crop output prices over time (Fig. S4). In 2000, Monsanto's patent on glyphosate expired and as a result glyphosate prices fell relative to non-glyphosate prices from 2001 onward. In addition, the commodity boom that began in the mid-2000s led to rising maize and soybean prices, which in turn encourage the use of yield-enhancing inputs like glyphosate.

No tillage

An additional important variable that could potentially confound our estimate for the impact of GM crops on pesticide use, is the adoption of no tillage (NT). Previous work has shown that NT and GE crops are complementary practices (45). NT may also use more herbicide relative to a conventional tillage operation. Thus, the greater use of herbicides observed for GE adopters may in part be attributable to the fact that they are more likely to adopt no tillage (from 1998-2009, no-till adoption increased from about 32% to 53% of land). Table S5 reports the results for soybean and maize herbicides when a binary variable for no tillage is included. We find that, although no tillage significantly increased herbicide use – by about 0.16 kg/ha in both maize and soybeans – it does not significantly alter the coefficients for the GE trait binary variable G_i .

Weed pressure and plot heterogeneity

A possible alternative explanation for the estimated pattern reported in Fig. 2 of the main paper is that the quantity of herbicides applied on a given plot is affected by weed pressure, and the latter may be related to the farmer's decision to adopt a GE variety. Insofar as weed pressure on plots belonging to the same farmer is highly correlated (e.g., it is a time-invariant attribute of the given farmer's location), the inclusion of a farmer-specific fixed effect in the estimating model provides a measure of control. However, insofar as there is additional unobserved plot-specific heterogeneity, the estimated coefficients on the GE variable G_i may reflect the impact of an implicit plot selection process.

To explore this possibility, we first investigate the sensitivity of our results to the inclusion of a set of control variables that capture the types weeds targeted by growers on each plot. Table S15 reports descriptive statistics for some major targeted weeds, separately for soybeans and maize, and for GT adopters and non-GT adopters. For the most part, there are no major differences in the frequency of weeds targeted between GT-adopters and non-GT adopters. To systematically explore the effects of weed pressure in the fixed effects regression model we add a set of indicator variables, where each variable takes the value one if the corresponding weed is targeted on that plot (and value zero otherwise). The results of this extended model are reported in Table S6. It turns out that the farmer's reporting of targeting each one of these major weeds does increase the amount of herbicides applied to that plot, for both crops and for all weeds. The estimated β_i that capture the differential impact of GE variety adoption, however, are robust to the inclusion of these weed pressure control variables.

Another way to investigate the impact of plot-specific heterogeneity is to estimate the fixed effects model on the subset of growers that plant either exclusively GT or exclusively non-GT varieties (i.e., exclude all growers that plant both GT and non-GT varieties within a given year). This procedure effectively eliminates potentially confounding plot specific factors. The results of this estimation are presented in tables S7 and S8. Qualitatively, the results are largely unchanged, but there is a small change in magnitude to the estimates. For herbicides, the GT

coefficient(s) are slightly smaller in both cases (they becomes less positive for soybeans and more negative for maize).

We can also test for the presence of an implicit plot selection process by using a simple model to generate predictions about the dynamics of herbicide use and compare those predictions to what we observe in the data. The following analysis illustrates.

Suppose that there is a continuum of plots, each of which is indexed by the degree of weediness $w \in [\underline{w}, \bar{w}]$, where w is distributed according to a continuous distribution function $F(w)$. Higher values of w represent higher weed pressure. Suppose that this factor was the only element in determining the sequence of GT variety adoption, and let $z \in [0, 1]$ denote the GT adoption rate. Then, for a given adoption rate $\hat{z} \in [0, 1]$ there is a weediness threshold $\hat{w} \in [\underline{w}, \bar{w}]$ such that all plots with $w \geq \hat{w}$ adopt GT varieties, and plots with $w < \hat{w}$ adopt conventional varieties. The threshold \hat{w} is determined by $F(\hat{w}) = 1 - \hat{z}$. Next, suppose that a plot's herbicide application rate depends on both the type of crop (GT or conventional) and the degree of weediness, and represent these amounts by $a_G(w)$ and $a_T(w)$ for GT and conventional (traditional) varieties, respectively.

In this setting we are interested in computing the expected (average) herbicide rate for conventional and GT varieties for any given adoption rate \hat{z} . Let $y_G(\hat{z})$ and $y_T(\hat{z})$ denote these average application rates. Then:

$$y_G(\hat{z}) = \int_{\hat{w}}^{\bar{w}} a_G(w) dF(w)$$

$$y_T(\hat{z}) = \int_{\underline{w}}^{\hat{w}} a_T(w) dF(w)$$

The coefficients α_t from the fixed effect regression model in the main paper essentially estimate the difference $\Delta(\hat{z}) \equiv y_G(\hat{z}) - y_T(\hat{z})$ as the term $\alpha_t G_{ift}$. How the foregoing conjectured adoption driver impacts these estimates cannot be established without further assumptions on the shape of the functions $a_G(w)$ and $a_T(w)$, and of the distribution function $F(w)$. To illustrate, suppose that $F(w)$ is a uniform distribution, such that $\hat{w} = \underline{w} + 1 - \hat{z}$, and that

$$a_G(w) = \alpha_G + \beta_G w$$

$$a_T(w) = \alpha_T + \beta_T w$$

Then:

$$y_G(\hat{z}) = \int_{\hat{w}}^{\bar{w}} (\alpha_G + \beta_G w) \frac{1}{(\bar{w} - \hat{w})} dw$$

$$y_T(\hat{z}) = \int_{\underline{w}}^{\hat{w}} (\alpha_T + \beta_T w) \frac{1}{(\hat{w} - \underline{w})} dw$$

Performing the integration:

$$y_G(\hat{z}) = \alpha_G + \frac{1}{2} \beta_G (\hat{w} + \bar{w})$$

$$y_T(\hat{z}) = \alpha_T + \frac{1}{2} \beta_T (\underline{w} + \hat{w})$$

and

$$\Delta(\hat{z}) = (\alpha_G - \alpha_T) + \frac{1}{2} (\beta_G (\hat{w} + \bar{w}) - \beta_T (\underline{w} + \hat{w})) = (\alpha_G - \alpha_T) + \frac{1}{2} (\beta_G \bar{w} - \beta_T \underline{w}) + \frac{1}{2} (\beta_G - \beta_T) \hat{w}$$

Several cases are possible, depending on the relative magnitudes of the intercepts α_i and the slopes β_i , $i=G, T$.

Testable Implication: Regardless of the relative magnitudes of parameters α_i and β_i , $i \in \{G, T\}$,

it is clear that, once the adoption rate \hat{z} stops increasing, such that \hat{w} is constant, then the difference in herbicide quantity used on GT and conventional plots, $\Delta(\hat{z})$, should converge to a constant. This suggests a testable implication for the estimated fixed effects regression model. For soybeans, in particular, the adoption rate \hat{z} has stabilized in the last part of the sample (for the last six years, 2006-2011, this rate has hovered between 94% and 97%). Hence, if the process being investigated was the primary explanation for the estimated pattern reported in Fig. 2, we should expect the estimated parameters α_t to be constant over these years. This null hypothesis $H_0 : \alpha_t = \alpha$ for all $t \in [2006, 2011]$, however, is rejected by the appropriate F statistics (F-statistic of 4.90, p -value = 0.0002).

Additional implication – Special case 1: $\alpha_G = \alpha_T$ and $\beta_G = \beta_T \equiv \beta$. In this case, the average herbicide applications on GT is greater than that on conventional variety plots, i.e.,

$$\Delta(\hat{z}) = \frac{1}{2} \beta (\bar{w} - \underline{w}) > 0, \text{ which would explain the paper's finding for soybeans reported in Table}$$

1. In this case, however, the difference does not change as the adoption rate changes, which is contrary to the paper's finding that $\Delta(\hat{z})$ increases with time (which is strongly positively correlated with the adoption rate) for both soybeans and maize, as reported in Fig. 2.

Additional Implication – Special case 2: $\alpha_G = \alpha_T = 0$ and $\beta_G \geq \beta_T$. In this case, again, the average herbicide applications on GT is greater than that on conventional variety plots, which would explain the paper's finding for soybean reported in Table 1. In this case, however, the difference $\frac{1}{2}(\beta_G \bar{w} - \beta_T \underline{w}) + \frac{1}{2}(\beta_G - \beta_T) \hat{w}$ is increasing in \hat{w} , and therefore decreasing in the adoption rate \hat{z} . Hence, over time, as adoption \hat{z} increases we should expect that the difference in average herbicide rates decreases, which again is contrary to the pattern uncovered for both soybean and maize, as reported in Fig. 2.




Fig. S1 – Number of Years Sampled for Growers in AgroTrak® dataset




Fig. S2. Maize Herbicide Use by non-GT Adopters (selected herbicides, kg/ha)

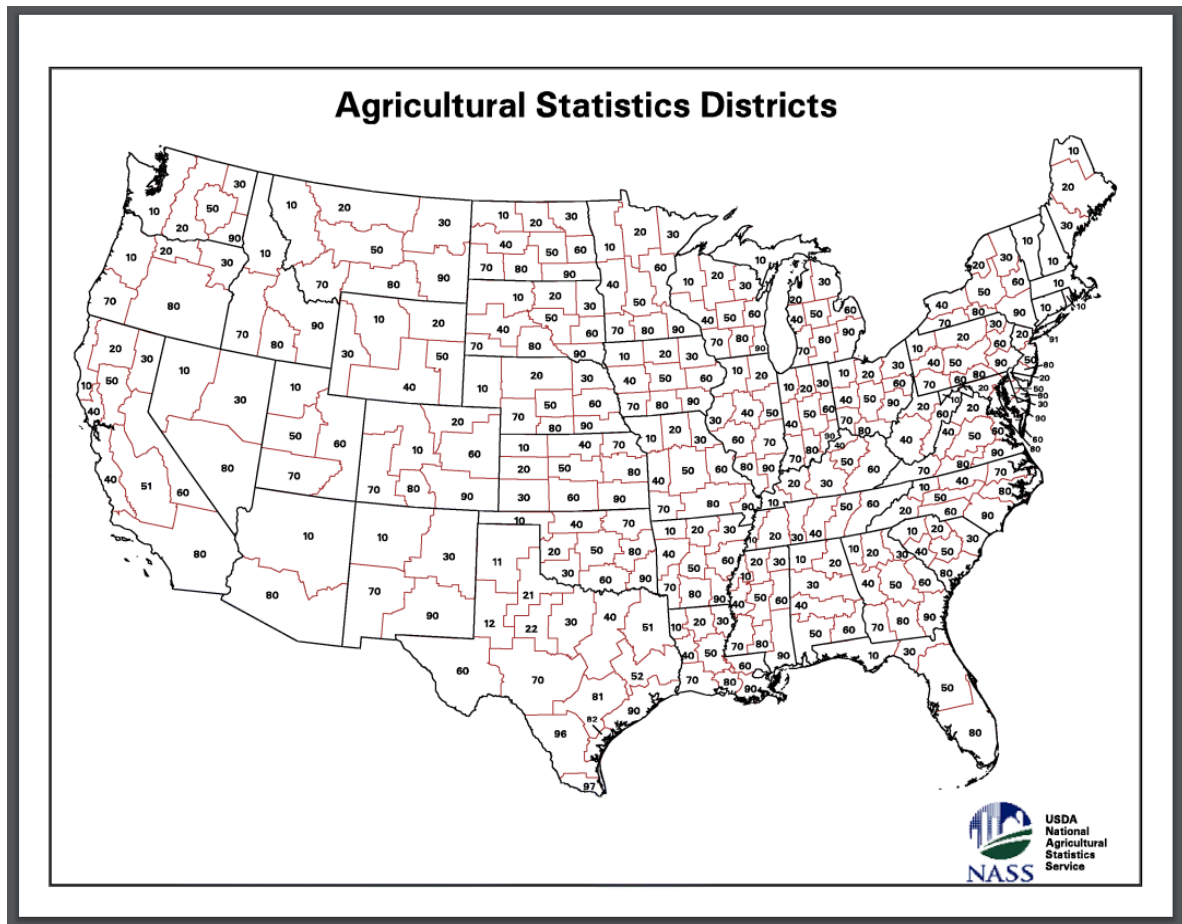


Fig. S3. Crop Reporting Districts (CRD), National Agricultural Statistics Service, U.S. Department of Agriculture.




Fig. S4. Trends in Glyphosate and Expected Crop Output Prices, 1998-2011

Table S1. Summary Statistics for AgroTrak® Dataset

	Annual Averages	
	Maize	Soybeans
Number of growers	5,424	5,029
Number of plots per farmer	1.77	1.23
Number of CRDs	248	197
Number of states	41	29
Number of herbicide a.i.	48	42
Number of insecticide a.i.	27	-

Table S2. Full Results Corresponding to Table 1 in the Main Text

	Soybean Herbicides		Maize Herbicides		Maize Insecticides	
	a.i. kg/ha	EQ kg/ha	a.i. kg/ha	EQ kg/ha	a.i. kg/ha	EQ kg/ha
G_t	0.3021*** (0.0097)	0.0045 (0.0122)	-0.0329* (0.0150)	-0.2590*** (0.0156)	-0.0129*** (0.0014)	-0.0122*** (0.0014)
1999	-0.1314*** (0.0161)	-0.1569*** (0.0197)	-0.2692*** (0.0302)	-0.2731*** (0.0314)	-0.0265*** (0.0064)	-0.0368*** (0.0075)
2000	-0.0946*** (0.0171)	-0.1156*** (0.0212)	-0.3318*** (0.0327)	-0.3458*** (0.0335)	-0.0267** (0.0082)	-0.0417*** (0.0090)
2001	-0.1029*** (0.0175)	-0.1423*** (0.0213)	-0.3925*** (0.0329)	-0.3924*** (0.0339)	-0.0485*** (0.0074)	-0.0652*** (0.0086)
2002	-0.1419*** (0.0178)	-0.2218*** (0.0212)	-0.4725*** (0.0343)	-0.4817*** (0.0352)	-0.0486*** (0.0092)	-0.0674*** (0.0098)
2003	-0.0819*** (0.0188)	-0.1791*** (0.0218)	-0.4910*** (0.0355)	-0.5078*** (0.0366)	-0.0575*** (0.0084)	-0.0731*** (0.0097)
2004	-0.0780*** (0.0185)	-0.1868*** (0.0218)	-0.5340*** (0.0362)	-0.5382*** (0.0372)	-0.0796*** (0.0083)	-0.1024*** (0.0096)
2005	-0.0661*** (0.0191)	-0.1760*** (0.0222)	-0.5603*** (0.0386)	-0.5570*** (0.0394)	-0.0928*** (0.0083)	-0.1158*** (0.0094)
2006	-0.1246*** (0.0193)	-0.2382*** (0.0223)	-0.5583*** (0.0398)	-0.5691*** (0.0410)	-0.1265*** (0.0083)	-0.1478*** (0.0093)
2007	-0.0005 (0.0208)	-0.1224*** (0.0237)	-0.4703*** (0.0416)	-0.4855*** (0.0428)	-0.1441*** (0.0100)	-0.1643*** (0.0105)
2008	0.1127*** (0.0210)	-0.0180 (0.0237)	-0.3675*** (0.0409)	-0.3857*** (0.0418)	-0.1323*** (0.0085)	-0.1505*** (0.0096)
2009	0.1520*** (0.0203)	0.0265 (0.0235)	-0.3346*** (0.0405)	-0.3500*** (0.0415)	-0.1304*** (0.0082)	-0.1513*** (0.0094)
2010	0.2421*** (0.0212)	0.1142*** (0.0239)	-0.2922*** (0.0400)	-0.3282*** (0.0409)	-0.1436*** (0.0088)	-0.1647*** (0.0098)
2011	0.2994*** (0.0220)	0.1698*** (0.0249)	-0.2641*** (0.0433)	-0.3137*** (0.0442)	-0.1400*** (0.0095)	-0.1580*** (0.0103)
Constant	1.1297*** (0.0142)	1.4439*** (0.0173)	2.8727*** (0.0272)	2.9758*** (0.0278)	0.1890*** (0.0063)	0.2060*** (0.0074)
N	86,736	86,736	134,264	134,264	134,264	134,264
R^2	0.067	0.028	0.022	0.027	0.039	0.051

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S3. Full Results Corresponding to Fig. 2 and Fig. 3 in the Main Text

	Soybean Herbicides		Maize Herbicides		Maize Insecticides	
	a.i. kg/ha	EQ kg/ha	a.i. kg/ha	EQ kg/ha	a.i. kg/ha	EQ kg/ha
$G_i \times 1998$	0.1714*** (0.0201)	-0.1487*** (0.0247)	-0.5527*** (0.1121)	-0.7993*** (0.1143)	0.0011 (0.0079)	0.0059 (0.0099)
$G_i \times 1999$	0.2089*** (0.0196)	-0.1067*** (0.0245)	-0.3574*** (0.0636)	-0.5778*** (0.0670)	0.0130* (0.0063)	0.0138* (0.0067)
$G_i \times 2000$	0.2422*** (0.0196)	-0.0472 (0.0253)	-0.4363*** (0.0639)	-0.6754*** (0.0664)	-0.0022 (0.0073)	-0.0033 (0.0074)
$G_i \times 2001$	0.3017*** (0.0215)	0.0112 (0.0271)	-0.1933*** (0.0507)	-0.3796*** (0.0523)	-0.0011 (0.0072)	-0.0030 (0.0070)
$G_i \times 2002$	0.3060*** (0.0236)	0.0094 (0.0293)	-0.2228*** (0.0468)	-0.4464*** (0.0479)	-0.0169* (0.0079)	-0.0195** (0.0072)
$G_i \times 2003$	0.4444*** (0.0293)	0.1512*** (0.0356)	-0.2060*** (0.0441)	-0.4706*** (0.0453)	-0.0067 (0.0065)	-0.0072 (0.0068)
$G_i \times 2004$	0.4280*** (0.0303)	0.1139** (0.0388)	-0.1027* (0.0418)	-0.3618*** (0.0426)	-0.0114 (0.0060)	-0.0085 (0.0062)
$G_i \times 2005$	0.3977*** (0.0317)	0.1014* (0.0395)	-0.1210** (0.0376)	-0.3701*** (0.0387)	-0.0082* (0.0039)	-0.0101** (0.0039)
$G_i \times 2006$	0.3936*** (0.0381)	0.1403** (0.0474)	-0.0442 (0.0347)	-0.2569*** (0.0359)	-0.0081* (0.0039)	-0.0062 (0.0036)
$G_i \times 2007$	0.4525*** (0.0512)	0.1768** (0.0642)	0.0190 (0.0385)	-0.2131*** (0.0404)	-0.0168** (0.0063)	-0.0149** (0.0053)
$G_i \times 2008$	0.5399*** (0.0521)	0.2915*** (0.0652)	0.2301*** (0.0395)	0.0247 (0.0410)	-0.0337*** (0.0039)	-0.0304*** (0.0037)
$G_i \times 2009$	0.4917*** (0.0442)	0.2304*** (0.0545)	0.1438*** (0.0406)	-0.0639 (0.0423)	-0.0277*** (0.0034)	-0.0248*** (0.0033)
$G_i \times 2010$	0.6428*** (0.0492)	0.3760*** (0.0587)	0.2973*** (0.0436)	0.0894 (0.0458)	-0.0207*** (0.0032)	-0.0218*** (0.0033)
$G_i \times 2011$	0.6604*** (0.0552)	0.4262*** (0.0638)	0.3639*** (0.0567)	0.1535** (0.0590)	-0.0191*** (0.0038)	-0.0182*** (0.0040)
1999	-0.1306*** (0.0217)	-0.1548*** (0.0278)	-0.2637*** (0.0303)	-0.2684*** (0.0315)	-0.0301*** (0.0063)	-0.0398*** (0.0074)

Table S3. continued

2000	-0.1090*** (0.0232)	-0.1439*** (0.0300)	-0.3171*** (0.0330)	-0.3307*** (0.0338)	-0.0268*** (0.0078)	-0.0408*** (0.0087)
2001	-0.1534*** (0.0255)	-0.2067*** (0.0324)	-0.3869*** (0.0330)	-0.3914*** (0.0341)	-0.0488*** (0.0072)	-0.0642*** (0.0084)
2002	-0.1961*** (0.0279)	-0.2856*** (0.0353)	-0.4606*** (0.0348)	-0.4707*** (0.0358)	-0.0451*** (0.0102)	-0.0625*** (0.0103)
2003	-0.2537*** (0.0327)	-0.3635*** (0.0402)	-0.4758*** (0.0364)	-0.4870*** (0.0376)	-0.0568*** (0.0088)	-0.0714*** (0.0098)
2004	-0.2391*** (0.0336)	-0.3423*** (0.0433)	-0.5312*** (0.0374)	-0.5289*** (0.0387)	-0.0775*** (0.0082)	-0.1004*** (0.0093)
2005	-0.2023*** (0.0349)	-0.3221*** (0.0439)	-0.5452*** (0.0407)	-0.5354*** (0.0417)	-0.0920*** (0.0084)	-0.1133*** (0.0092)
2006	-0.2600*** (0.0409)	-0.4231*** (0.0513)	-0.5640*** (0.0420)	-0.5803*** (0.0435)	-0.1263*** (0.0086)	-0.1475*** (0.0094)
2007	-0.1931*** (0.0533)	-0.3444*** (0.0674)	-0.5116*** (0.0473)	-0.5236*** (0.0493)	-0.1392*** (0.0117)	-0.1595*** (0.0115)
2008	-0.1648** (0.0538)	-0.3513*** (0.0679)	-0.5801*** (0.0490)	-0.6146*** (0.0508)	-0.1168*** (0.0087)	-0.1359*** (0.0096)
2009	-0.0771 (0.0472)	-0.2454*** (0.0588)	-0.4894*** (0.0495)	-0.5200*** (0.0513)	-0.1187*** (0.0083)	-0.1404*** (0.0094)
2010	-0.1305* (0.0515)	-0.2965*** (0.0622)	-0.5914*** (0.0531)	-0.6436*** (0.0555)	-0.1362*** (0.0089)	-0.1555*** (0.0099)
2011	-0.0902 (0.0569)	-0.2886*** (0.0667)	-0.6317*** (0.0638)	-0.6957*** (0.0661)	-0.1335*** (0.0095)	-0.1510*** (0.0103)
Constant	1.1812*** (0.0164)	1.5043*** (0.0205)	2.8851*** (0.0273)	2.9887*** (0.0279)	0.1865*** (0.0063)	0.2028*** (0.0072)
<i>N</i>	86,736	86,736	134,264	134,264	134,264	134,264
<i>R</i> ²	0.071	0.032	0.026	0.031	0.040	0.051

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S4. Random Effects replace Farmers Fixed Effects (compare to table S3)

	Soybean Herbicides		Maize Herbicides		Maize Insecticides	
	a.i. kg/ha	EIQ kg/ha	a.i. kg/ha	EIQ kg/ha	a.i. kg/ha	EIQ kg/ha
$G_i \times 1998$	0.1605*** (0.0174)	-0.1726*** (0.0212)	-0.6704*** (0.0988)	-0.9345*** (0.1002)	-0.0050 (0.0078)	-0.0028 (0.0099)
$G_i \times 1999$	0.1987*** (0.0174)	-0.1288*** (0.0218)	-0.4138*** (0.0599)	-0.6418*** (0.0634)	0.0120 (0.0063)	0.0124 (0.0068)
$G_i \times 2000$	0.2313*** (0.0176)	-0.0804*** (0.0228)	-0.5255*** (0.0571)	-0.7775*** (0.0594)	-0.0042 (0.0073)	-0.0055 (0.0074)
$G_i \times 2001$	0.2918*** (0.0195)	-0.0152 (0.0248)	-0.2572*** (0.0472)	-0.4587*** (0.0489)	-0.0066 (0.0072)	-0.0090 (0.0070)
$G_i \times 2002$	0.3100*** (0.0217)	-0.0016 (0.0271)	-0.3153*** (0.0428)	-0.5534*** (0.0440)	-0.0220** (0.0085)	-0.0228** (0.0077)
$G_i \times 2003$	0.4439*** (0.0267)	0.1348*** (0.0334)	-0.2640*** (0.0409)	-0.5426*** (0.0421)	-0.0100 (0.0059)	-0.0114 (0.0065)
$G_i \times 2004$	0.4122*** (0.0276)	0.0868* (0.0357)	-0.1892*** (0.0383)	-0.4618*** (0.0393)	-0.0114 (0.0059)	-0.0087 (0.0061)
$G_i \times 2005$	0.3837*** (0.0295)	0.0659 (0.0377)	-0.1734*** (0.0346)	-0.4345*** (0.0356)	-0.0086* (0.0040)	-0.0107** (0.0039)
$G_i \times 2006$	0.4106*** (0.0338)	0.1454*** (0.0431)	-0.0765* (0.0324)	-0.3003*** (0.0337)	-0.0078* (0.0034)	-0.0057 (0.0032)
$G_i \times 2007$	0.5277*** (0.0460)	0.2459*** (0.0579)	0.0294 (0.0360)	-0.2124*** (0.0379)	-0.0166*** (0.0047)	-0.0143** (0.0045)
$G_i \times 2008$	0.6226*** (0.0486)	0.3620*** (0.0613)	0.2556*** (0.0372)	0.0437 (0.0387)	-0.0299*** (0.0037)	-0.0277*** (0.0035)
$G_i \times 2009$	0.5368*** (0.0420)	0.2529*** (0.0526)	0.1929*** (0.0388)	-0.0210 (0.0405)	-0.0250*** (0.0032)	-0.0229*** (0.0032)
$G_i \times 2010$	0.6821*** (0.0458)	0.3950*** (0.0557)	0.3742*** (0.0415)	0.1615*** (0.0436)	-0.0194*** (0.0030)	-0.0211*** (0.0033)
$G_i \times 2011$	0.6875*** (0.0498)	0.4305*** (0.0581)	0.4794*** (0.0528)	0.2663*** (0.0550)	-0.0175*** (0.0035)	-0.0169*** (0.0038)
N	86736	86736	134264	134264	134264	134264
R^2	0.065	0.025	0.012	0.016	0.019	0.024
Hausman test	607.4***	641.2***	2,615***	3,004***	3,164***	4,335***

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects and CRD-specific time trends. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S5. Model Estimates with the No-Till Binary Variable Included

	Soybean Herbicides		Maize Herbicides	
	a.i. kg/ha	EQ kg/ha	a.i. kg/ha	EQ kg/ha
No Till	0.1600*** (0.0081)	0.1408*** (0.0090)	0.1531*** (0.0162)	0.1240*** (0.0168)
$G_i \times 1998$	0.1609*** (0.0201)	-0.1580*** (0.0247)	-0.5588*** (0.1117)	-0.8043*** (0.1140)
$G_i \times 1999$	0.1986*** (0.0195)	-0.1158*** (0.0245)	-0.3606*** (0.0635)	-0.5804*** (0.0670)
$G_i \times 2000$	0.2334*** (0.0196)	-0.0549* (0.0254)	-0.4394*** (0.0639)	-0.6779*** (0.0664)
$G_i \times 2001$	0.2906*** (0.0215)	0.0014 (0.0271)	-0.1989*** (0.0506)	-0.3841*** (0.0523)
$G_i \times 2002$	0.2956*** (0.0236)	0.0003 (0.0293)	-0.2238*** (0.0468)	-0.4472*** (0.0479)
$G_i \times 2003$	0.4346*** (0.0292)	0.1426*** (0.0356)	-0.2095*** (0.0440)	-0.4734*** (0.0452)
$G_i \times 2004$	0.4197*** (0.0302)	0.1066** (0.0387)	-0.1102** (0.0417)	-0.3678*** (0.0425)
$G_i \times 2005$	0.3862*** (0.0317)	0.0913* (0.0395)	-0.1244*** (0.0376)	-0.3728*** (0.0387)
$G_i \times 2006$	0.3803*** (0.0382)	0.1286** (0.0475)	-0.0458 (0.0347)	-0.2583*** (0.0359)
$G_i \times 2007$	0.4488*** (0.0511)	0.1736** (0.0641)	0.0172 (0.0385)	-0.2146*** (0.0405)
$G_i \times 2008$	0.5298*** (0.0522)	0.2826*** (0.0653)	0.2272*** (0.0395)	0.0223 (0.0410)
$G_i \times 2009$	0.4801*** (0.0443)	0.2201*** (0.0547)	0.1404*** (0.0406)	-0.0666 (0.0423)
$G_i \times 2010$	0.6341*** (0.0491)	0.3683*** (0.0588)	0.2922*** (0.0435)	0.0853 (0.0458)
$G_i \times 2011$	0.6420*** (0.0548)	0.4100*** (0.0634)	0.3578*** (0.0568)	0.1486* (0.0591)
N	86,736	86,736	134,264	134,264
R^2	0.079	0.037	0.028	0.032

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S6. Impact of GE variety adoption on herbicide use (kg/ha of a.i.)
Results including indicator variables for major targeted weeds

	Soybean Herbicides		Maize Herbicides	
	a.i. kg/ha	a.i. kg/ha	a.i. kg/ha	a.i. kg/ha
G_i	0.3176*** (0.0096)		-0.0245 (0.0147)	
$G_i \times 1998$		0.1976*** (0.0199)		-0.4943*** (0.1097)
$G_i \times 1999$		0.2275*** (0.0194)		-0.2978*** (0.0618)
$G_i \times 2000$		0.2571*** (0.0195)		-0.3906*** (0.0624)
$G_i \times 2001$		0.3192*** (0.0213)		-0.1674*** (0.0497)
$G_i \times 2002$		0.3146*** (0.0233)		-0.1845*** (0.0461)
$G_i \times 2003$		0.4562*** (0.0289)		-0.1727*** (0.0430)
$G_i \times 2004$		0.4366*** (0.0299)		-0.0790 (0.0414)
$G_i \times 2005$		0.4097*** (0.0313)		-0.1075** (0.0370)
$G_i \times 2006$		0.3999*** (0.0379)		-0.0455 (0.0342)
$G_i \times 2007$		0.4483*** (0.0491)		0.0160 (0.0378)
$G_i \times 2008$		0.5417*** (0.0508)		0.2177*** (0.0394)
$G_i \times 2009$		0.5019*** (0.0430)		0.1317** (0.0405)
$G_i \times 2010$		0.6548*** (0.0475)		0.2733*** (0.0426)
$G_i \times 2011$		0.6530*** (0.0544)		0.3260*** (0.0564)

Table S6. continued

Cocklebur	0.0672*** (0.0086)	0.0654*** (0.0086)	0.1168*** (0.0188)	0.1149*** (0.0188)
Foxtail	0.1005*** (0.0079)	0.0990*** (0.0079)	0.2935*** (0.0157)	0.2888*** (0.0157)
Lambsquarters	0.0617*** (0.0091)	0.0599*** (0.0091)	0.1070*** (0.0169)	0.1030*** (0.0169)
Pigweed	0.0728*** (0.0090)	0.0718*** (0.0090)	0.1876*** (0.0163)	0.1848*** (0.0163)
Ragweed	0.0517*** (0.0094)	0.0514*** (0.0093)	0.1334*** (0.0183)	0.1302*** (0.0183)
Velvetleaf	0.0306*** (0.0089)	0.0290** (0.0088)	0.1573*** (0.0168)	0.1549*** (0.0168)
Waterhemp	0.0831*** (0.0107)	0.0839*** (0.0107)	0.0767*** (0.0233)	0.0752** (0.0233)
Morningglory	0.0797*** (0.0144)	0.0775*** (0.0144)	0.2045*** (0.0313)	0.2028*** (0.0313)
Johnson grass	0.0857*** (0.0152)	0.0871*** (0.0152)	0.0641* (0.0294)	0.0687* (0.0294)
<i>N</i>	86,736	86,736	134,264	134,264
<i>R</i> ²	0.080	0.083	0.043	0.046

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S7. Model excludes growers that plant both GE and non-GE varieties within a given year (compare with Table 1 in the main text)

	Soybean Herbicides		Corn Herbicides		Corn Insecticides	
	a.i. kg/ha	EIQ kg/ha	a.i. kg/ha	EIQ kg/ha	a.i. kg/ha	EIQ kg/ha
G_i	0.2559*** (0.0189)	-0.0515* (0.0226)	-0.0956** (0.0295)	-0.3480*** (0.0308)	-0.0106 (0.0070)	-0.0101 (0.0069)
<i>N</i>	71,239	71,239	108,327	108,327	62,265	62,265
<i>R</i> ²	0.060	0.037	0.027	0.032	0.030	0.042

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S8. Model excludes growers that plant both GE and non-GE varieties within a given year (compare with Table S3)

	Soybean Herbicides		Corn Herbicides		Corn Insecticides	
	a.i. kg/ha	EQ kg/ha	a.i. kg/ha	EQ kg/ha	a.i. kg/ha	EQ kg/ha
$G_i \times 1998$	0.0602 (0.0364)	-0.2774*** (0.0422)	-0.9996*** (0.2887)	-1.2646*** (0.2933)	0.0011 (0.0351)	0.0036 (0.0468)
$G_i \times 1999$	0.1358*** (0.0340)	-0.1985*** (0.0412)	-0.8637*** (0.1609)	-1.1141*** (0.1738)	-0.0180 (0.0292)	-0.0322 (0.0327)
$G_i \times 2000$	0.1690*** (0.0319)	-0.1425*** (0.0389)	-0.5778*** (0.1182)	-0.8486*** (0.1211)	0.0099 (0.0231)	0.0272 (0.0271)
$G_i \times 2001$	0.2276*** (0.0365)	-0.0860 (0.0455)	-0.4668*** (0.0921)	-0.6952*** (0.0948)	0.0704 (0.0679)	0.0540 (0.0576)
$G_i \times 2002$	0.2509*** (0.0379)	-0.0591 (0.0475)	-0.3760*** (0.0774)	-0.6292*** (0.0784)	0.0010 (0.0245)	-0.0110 (0.0218)
$G_i \times 2003$	0.4036*** (0.0459)	0.1162* (0.0558)	-0.2470*** (0.0661)	-0.5122*** (0.0670)	-0.0129 (0.0225)	-0.0054 (0.0229)
$G_i \times 2004$	0.3770*** (0.0452)	0.0740 (0.0580)	-0.1321* (0.0605)	-0.3954*** (0.0619)	0.0157 (0.0382)	0.0217 (0.0407)
$G_i \times 2005$	0.4047*** (0.0456)	0.1211* (0.0554)	-0.1272* (0.0596)	-0.3910*** (0.0608)	-0.0052 (0.0155)	-0.0099 (0.0178)
$G_i \times 2006$	0.3448*** (0.0578)	0.0967 (0.0715)	-0.1408* (0.0567)	-0.3775*** (0.0586)	-0.0082 (0.0129)	-0.0060 (0.0120)
$G_i \times 2007$	0.4456*** (0.0666)	0.1756* (0.0809)	-0.1416* (0.0644)	-0.4053*** (0.0679)	-0.0198 (0.0226)	-0.0215 (0.0190)
$G_i \times 2008$	0.5352*** (0.0656)	0.2680** (0.0827)	0.1822** (0.0645)	-0.0579 (0.0669)	-0.0336*** (0.0101)	-0.0244* (0.0095)
$G_i \times 2009$	0.5135*** (0.0638)	0.2473** (0.0789)	0.0409 (0.0671)	-0.1995** (0.0702)	-0.0255* (0.0113)	-0.0211 (0.0119)
$G_i \times 2010$	0.6811*** (0.0662)	0.3947*** (0.0799)	0.2575*** (0.0721)	0.0086 (0.0748)	-0.0201 (0.0117)	-0.0251* (0.0126)
$G_i \times 2011$	0.6311*** (0.0762)	0.3868*** (0.0840)	0.2363** (0.0894)	-0.0250 (0.0930)	-0.0196 (0.0124)	-0.0199 (0.0122)
N	71,239	71,239	108,327	108,327	62,265	62,265
R^2	0.064	0.042	0.030	0.034	0.030	0.043

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S9. Model excludes farmers that never used pesticides (on any of their plots)
(compare with Table 1 in the main text)

	Soybean Herbicides		Corn Herbicides		Corn Insecticides	
	a.i. kg/ha	EIQ kg/ha	a.i. kg/ha	EIQ kg/ha	a.i. kg/ha	EIQ kg/ha
G_i	0.3023*** (0.0097)	0.0045 (0.0122)	-0.0331* (0.0150)	-0.2594*** (0.0157)	-0.0140*** (0.0015)	-0.0131*** (0.0016)
N	85,932	85,932	132,824	132,824	106,256	106,256
R^2	0.067	0.028	0.022	0.028	0.042	0.055

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S10. Model excludes farmers that never used pesticides (on any of their plots)
(compare with Table S3)

	Soybean Herbicides		Corn Herbicides		Corn Insecticides	
	a.i. kg/ha	EIQ kg/ha	a.i. kg/ha	EIQ kg/ha	a.i. kg/ha	EIQ kg/ha
$G_i \times 1998$	0.1709*** (0.0202)	-0.1499*** (0.0247)	-0.5537*** (0.1132)	-0.8029*** (0.1154)	-0.0037 (0.0116)	0.0015 (0.0145)
$G_i \times 1999$	0.2093*** (0.0197)	-0.1071*** (0.0247)	-0.3598*** (0.0640)	-0.5817*** (0.0674)	0.0179* (0.0089)	0.0191* (0.0096)
$G_i \times 2000$	0.2422*** (0.0197)	-0.0477 (0.0255)	-0.4366*** (0.0640)	-0.6761*** (0.0664)	-0.0049 (0.0103)	-0.0064 (0.0105)
$G_i \times 2001$	0.3021*** (0.0217)	0.0109 (0.0274)	-0.1931*** (0.0507)	-0.3796*** (0.0524)	-0.0008 (0.0098)	-0.0032 (0.0095)
$G_i \times 2002$	0.3071*** (0.0239)	0.0101 (0.0297)	-0.2228*** (0.0469)	-0.4468*** (0.0480)	-0.0224* (0.0110)	-0.0254* (0.0100)
$G_i \times 2003$	0.4472*** (0.0298)	0.1539*** (0.0363)	-0.2061*** (0.0441)	-0.4709*** (0.0453)	-0.0078 (0.0085)	-0.0084 (0.0088)
$G_i \times 2004$	0.4301*** (0.0309)	0.1159** (0.0396)	-0.1027* (0.0419)	-0.3623*** (0.0427)	-0.0126 (0.0071)	-0.0090 (0.0073)
$G_i \times 2005$	0.3982*** (0.0323)	0.1021* (0.0402)	-0.1210** (0.0377)	-0.3703*** (0.0387)	-0.0079 (0.0043)	-0.0097* (0.0043)
$G_i \times 2006$	0.3981*** (0.0392)	0.1458** (0.0488)	-0.0440 (0.0348)	-0.2571*** (0.0360)	-0.0067 (0.0042)	-0.0044 (0.0039)
$G_i \times 2007$	0.4569*** (0.0541)	0.1827** (0.0678)	0.0189 (0.0387)	-0.2137*** (0.0406)	-0.0145* (0.0067)	-0.0124* (0.0056)
$G_i \times 2008$	0.5433*** (0.0557)	0.2981*** (0.0697)	0.2311*** (0.0400)	0.0258 (0.0415)	-0.0337*** (0.0041)	-0.0303*** (0.0038)
$G_i \times 2009$	0.4907*** (0.0465)	0.2305*** (0.0573)	0.1439*** (0.0412)	-0.0638 (0.0430)	-0.0273*** (0.0035)	-0.0242*** (0.0034)
$G_i \times 2010$	0.6425*** (0.0529)	0.3776*** (0.0632)	0.3024*** (0.0447)	0.0952* (0.0470)	-0.0195*** (0.0033)	-0.0206*** (0.0034)
$G_i \times 2011$	0.6602*** (0.0579)	0.4295*** (0.0669)	0.3697*** (0.0588)	0.1599** (0.0611)	-0.0185*** (0.0039)	-0.0176*** (0.0041)
N	85,932	85,932	132,824	132,824	106,256	106,256
R^2	0.072	0.032	0.026	0.031	0.043	0.055

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S11. Full Set of Results Corresponding to Fig. 4 in the Main Text

	Soybeans		Maize	
	Glyphosate (kg/ha)	Non-Glyphosate (kg/ha)	Glyphosate (kg/ha)	Non-Glyphosate (kg/ha)
$G_i \times 1998$	0.8772*** (0.0117)	-0.7057*** (0.0176)	0.7651*** (0.0263)	-1.3178*** (0.1114)
$G_i \times 1999$	0.8686*** (0.0111)	-0.6596*** (0.0175)	0.6808*** (0.0240)	-1.0382*** (0.0659)
$G_i \times 2000$	0.8481*** (0.0115)	-0.6059*** (0.0183)	0.7547*** (0.0261)	-1.1910*** (0.0644)
$G_i \times 2001$	0.8908*** (0.0131)	-0.5891*** (0.0192)	0.6197*** (0.0206)	-0.8130*** (0.0512)
$G_i \times 2002$	0.8831*** (0.0151)	-0.5771*** (0.0200)	0.7607*** (0.0195)	-0.9835*** (0.0461)
$G_i \times 2003$	0.9882*** (0.0185)	-0.5438*** (0.0241)	0.8892*** (0.0179)	-1.0952*** (0.0427)
$G_i \times 2004$	1.0075*** (0.0188)	-0.5795*** (0.0251)	0.8695*** (0.0172)	-0.9722*** (0.0395)
$G_i \times 2005$	0.9548*** (0.0188)	-0.5572*** (0.0266)	0.8060*** (0.0153)	-0.9270*** (0.0361)
$G_i \times 2006$	0.8344*** (0.0223)	-0.4408*** (0.0333)	0.6953*** (0.0142)	-0.7394*** (0.0333)
$G_i \times 2007$	0.8794*** (0.0351)	-0.4268*** (0.0408)	0.7490*** (0.0134)	-0.7300*** (0.0375)
$G_i \times 2008$	0.9715*** (0.0338)	-0.4316*** (0.0438)	0.7085*** (0.0148)	-0.4784*** (0.0385)
$G_i \times 2009$	0.9671*** (0.0283)	-0.4754*** (0.0370)	0.7104*** (0.0147)	-0.5666*** (0.0389)
$G_i \times 2010$	1.1295*** (0.0321)	-0.4868*** (0.0380)	0.6920*** (0.0173)	-0.3947*** (0.0431)
$G_i \times 2011$	1.1278*** (0.0370)	-0.4674*** (0.0434)	0.6898*** (0.0222)	-0.3260*** (0.0547)

Table S11. continued

1999	-0.0229* (0.0091)	-0.1077*** (0.0200)	0.0008 (0.0054)	-0.2645*** (0.0300)
2000	0.0249* (0.0104)	-0.1339*** (0.0217)	0.0175** (0.0067)	-0.3347*** (0.0326)
2001	0.0484*** (0.0121)	-0.2018*** (0.0234)	0.0267*** (0.0067)	-0.4136*** (0.0326)
2002	0.0876*** (0.0148)	-0.2837*** (0.0249)	0.0381*** (0.0071)	-0.4988*** (0.0344)
2003	0.0787*** (0.0179)	-0.3324*** (0.0282)	0.0534*** (0.0076)	-0.5292*** (0.0359)
2004	0.0781*** (0.0182)	-0.3173*** (0.0292)	0.0523*** (0.0079)	-0.5835*** (0.0369)
2005	0.1427*** (0.0185)	-0.3450*** (0.0306)	0.0743*** (0.0098)	-0.6194*** (0.0392)
2006	0.2103*** (0.0224)	-0.4702*** (0.0367)	0.1221*** (0.0100)	-0.6861*** (0.0413)
2007	0.2855*** (0.0349)	-0.4785*** (0.0438)	0.1576*** (0.0109)	-0.6692*** (0.0465)
2008	0.2996*** (0.0334)	-0.4644*** (0.0465)	0.2279*** (0.0135)	-0.8080*** (0.0485)
2009	0.3118*** (0.0285)	-0.3889*** (0.0405)	0.2174*** (0.0137)	-0.7068*** (0.0481)
2010	0.2371*** (0.0323)	-0.3676*** (0.0408)	0.2891*** (0.0166)	-0.8805*** (0.0526)
2011	0.2493*** (0.0370)	-0.3395*** (0.0457)	0.3312*** (0.0210)	-0.9628*** (0.0620)
Constant	0.1476*** (0.0079)	1.0336*** (0.0150)	0.0625*** (0.0061)	2.8226*** (0.0270)
<i>N</i>	86,736	86,736	134,264	134,264
<i>R</i> ²	0.357	0.238	0.440	0.108

Notes: Standard errors (in parentheses) are clustered at the farmer level. Model includes time fixed effects, CRD-specific time trends and individual fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table S12. US GE Adoption Rates (% of planted hectares), 1998-2011

Year	Bt Maize	HT	
		Maize	HT Soy
1998	11%	3%	37%
1999	19%	6%	57%
2000	18%	8%	63%
2001	19%	13%	74%
2002	23%	15%	83%
2003	26%	19%	88%
2004	30%	28%	90%
2005	40%	40%	91%
2006	46%	51%	94%
2007	59%	71%	96%
2008	66%	84%	96%
2009	69%	88%	95%
2010	69%	90%	95%
2011	71%	91%	97%

Table S13. US Pesticide Rates (kg/ha), 1998-2011

Year	Maize	Maize	Soybean
	Insecticides	Herbicides	Herbicides
1998	0.203	2.848	1.248
1999	0.172	2.514	1.197
2000	0.174	2.532	1.228
2001	0.139	2.498	1.279
2002	0.146	2.372	1.250
2003	0.133	2.392	1.339
2004	0.119	2.310	1.356
2005	0.081	2.344	1.367
2006	0.056	2.285	1.292
2007	0.055	2.464	1.441
2008	0.052	2.560	1.600
2009	0.051	2.525	1.609
2010	0.039	2.624	1.726
2011	0.053	2.652	1.862

Notes: these rates are calculated by adding up the total amount of active ingredients used in each year and dividing by the total number of hectares planted to the corresponding crop.

Table S14. Correlation Between State-Level GE Adoption Rates from USDA and GfK Data

	Maize	Soybeans
All GE Varieties	0.987	0.901
GT Only Varieties	0.956	0.901

Table S15. Summary Statistics by Adoption Choice

Variable	Soybeans		Maize	
	non-GT	GT	non-GT	GT
Weed Targeted				
Cocklebur	37.7%	28.2%	30.2%	24.7%
Foxtail	57.0%	48.3%	60.5%	57.3%
Lamb's Quarters	26.3%	23.1%	29.1%	32.0%
Pigweed	27.8%	23.0%	27.3%	26.6%
Ragweed	28.3%	24.2%	25.6%	24.1%
Velvet Leaf	26.0%	25.8%	34.0%	34.6%
Water Hemp	16.9%	19.6%	17.0%	21.4%
Morning Glory	14.6%	10.7%	7.5%	6.5%
Johnson Grass	11.2%	11.0%	7.5%	7.8%
Hectares Grown	198.5	204.6	175.9	217.9

CHAPTER 4**A DISCRETE CHOICE MODEL OF SEED DEMAND: GENETICALLY ENGINEERED
VARIETIES IN U.S. CORN AND SOYBEANS**

Federico Ciliberto, GianCarlo Moschini, and Edward D. Perry*

Abstract

In the two decades since they were commercially released in the United States, genetically engineered (GE) corn and soybean varieties have become a staple of U.S. farming, exceeding 90% of planted land since 2011. However, consolidation and rising seed prices in the agrochemical and seed industries, particularly since 2007, have increasingly generated concern about the benefits for the farmers that now almost universally rely on GE varieties. In this paper, we develop a discrete choice model of corn and soybean seed varieties and use it to estimate U.S. farmers' willingness-to-pay for GE traits over the period 1996-2011. Importantly, we allow the returns to GE traits to structurally vary over the three sub-periods that correspond to the expiration of Monsanto's glyphosate patent in 2000 and the sharp increase in commodity prices that began in 2007. We find that the marginal benefits to GE varieties were small early on and large in the final sub-period. A comparison of these benefits to the average premiums charged by seed firms suggests that the gains from GE varieties to farmers were in fact the greatest during the sub-period in which seed prices rose the most.

*Senior authorship not assigned.

Introduction

The systematic application of modern biotechnology techniques to crop improvement has been one of the most salient developments affecting global agriculture in the last 20 years. The most impactful creation to emerge from these efforts are genetically engineered (GE) crop varieties, first introduced commercially in 1996. Despite being limited to four main crops (maize, soybean, cotton, and canola) and extensively grown in only a handful of countries (United States, Brazil, Argentina, India, and Canada accounted for more than 90% of GE crop planting in 2015), as of 2015, GE varieties were grown on more than 444 million acres worldwide (James, 2015). The United States has been at the forefront of these developments. U.S. seed companies have been leading innovators, and U.S. farmers have been ready and loyal adopters. In 2015, GE varieties were planted on more than 175 million acres of U.S. farmland (James, 2015), nearly 95% of which was maize and soybeans, and for all years since 2007, GE adoption rates for corn and soybeans have exceeded 75% and 90%, respectively (Fernandez-Cornejo et al., 2014).

The development and diffusion of GE varieties have given rise to a number of economic issues that have been addressed by a considerable body of research (Shoemaker et al., 2001; Moschini, 2008; Barrows, Sexton, and Zilberman 2014). Despite the productivity enhancing potential of GE crops, the use of GE technology in agriculture has been marred by controversy (Carter, Moschini, and Sheldon, 2011). Concerns raised include the fear that GE products are harmful to human health and/or the environment, ethical objections related to human intervention in the DNA of living plants and animals, and mistrust of the ownership interests of multinational corporations that commercialize GE products.

Many of these concerns have been allayed. In fact, the environmental impacts of GE varieties appear to be generally positive. The National Research Council study (NRC, 2010) concluded that GE varieties result in reduced pesticide use and the use of lower toxicity products. These conclusions were largely confirmed in Bennett et al. (2013), although it is important to recognize that the effects of GE variety adoption on pesticide use have changed

over time (Perry, Ciliberto, Hennessy, and Moschini 2016). Other research has demonstrated a range of positive spill-over effects, including the area-wide suppression of crop-destroying pests by insect resistant corn (Hutchison et al. 2010) and the increased adoption of no-tillage systems by herbicide resistant soybean adopters (Perry, Moschini, and Hennessy 2016). Positive human health impacts have also been documented: for example, insect resistant varieties improve farm workers' health by reducing exposure to insecticides (Huang et al., 2005).

Despite much work, there is less conclusive evidence on the monetary benefits associated with the main economic parties involved in GE crops (firms, farmers, and consumers). This is particularly important given that many view GE varieties negatively because of the belief that they mostly benefit the large firms that have developed and marketed them, but not the end users (farmers and, indirectly, consumers). The rapid adoption of GE crops is *prima facie* evidence of producer profitability. However, in recent years the culmination of high output prices and increasing market concentration have occurred alongside sharply increasing seed prices, prompting increased discussion about the role of market power in the seed industry (Moschini, 2010; Moss, 2010). The extent to which farmers have continued to benefit from GE varieties thus depends on how their marginal benefits have evolved relative to increasing costs. The answer is not clear. Farmers also benefit from rising output prices and the expiration of Monsanto's patent on the herbicide glyphosate in 2000 has led to significantly lower herbicide costs for glyphosate tolerant seed adopters.

Ultimately, an economic assessment of the net economic value of such efficiency-enhancing innovations requires both an explicit structural model and detailed data. In this paper we adapt the theory of discrete choice in a differentiated product setting (Anderson, De Palma and Thisse 1992) to farmers' choice of seed varieties. The model is estimated by using a large farm-level dataset of seed choices by U.S. corn and soybean farmers over the 1996-2011 period. The econometric estimates that we obtain are used to estimate the dynamic marginal benefit of GE variety adoption accruing to farmers. Among other things, these estimates can be compared with the average premiums charged by seed firms, which provides an initial

assessment of how the *ex post* returns of GE innovation are allocated between seed companies and farmers.

In developing and estimating our model we contribute to the extant literature in several ways. First, we draw on large and detailed dataset of plot-level choices of corn and soybean growers over the 1996-2011 period. In the past, the availability of detailed data on GE seed varieties has been a serious impediment to the kinds of analyses that could be conducted on these issues. Earlier studies that considered the welfare effects of GE varieties avoided this limitation by developing computable partial equilibrium frameworks. They found sizeable increases in welfare due to GE crop adoption, with the benefits shared by farmers, consumers and the sellers of improved crop varieties (Falck-Zepeda, Traxler, and Nelson, 2000; Moschini, Lapan, and Sobolevsky, 2000), a result requiring some qualifications when allowance is made for the induced product differentiation and costly segregation and identity preservation activities (Bullock and Desquilbet, 2002; Sobolevsky, Moschini, and Lapan, 2005). But, by design, these studies could not speak to individual level estimates of the benefits associated with GE crops, nor on what the more recent developments in the seed markets imply for growers.

The most recent direct econometric evidence has been provided by a series of papers that use a subset of the proprietary data used in this paper. The basic framework of these studies is presented in Shi, Chavas, and Stiegert (2010), with extensions in Shi, Stiegert, and Chavas (2011) and Shi et al. (2012). In general, they estimate hedonic regressions with seed prices as the dependent variable and the various GE traits (and their combinations) as the independent variables (they also include various other controls). For the period 2000-2007, they report positive premiums for most traits in both corn and soybeans, and they also find strong evidence of sub-additive pricing (i.e., the premium charged for a typical seed with two GE traits is less than the sum of the premiums charged for the respective single GE trait seeds).

Our work differs from these studies in that our econometric framework incorporates many of the recent advances made in the empirical IO literature over the past two decades

(Berry 1994; Berry et al. 1995; Nevo 2001). Our demand model is built from individual discrete choices on the basis of profit maximization, and we appropriately account for the presence of imperfect competition by including instrumental variables for seed prices. There is thus no ambiguity about the meaning of our estimates. This is in contrast to Shi et al (2010) and Shi et al. (2012) where it is to some degree unclear as to whether their estimates are capturing marginal benefits or marginal costs.

In addition, we specify and estimate a unified framework of corn and soybean seed demand; that is, within our framework farmers are modeled as choosing between all corn and soybean varieties at the same time. This is important for at least two reasons. First, estimating the demand for each crop in isolation, as done by all previous studies, implicitly assumes that there is no significant short-run substitution between crops. Recent work by Hendricks, Smith, and Sumner (2014) clearly demonstrates that there can be significant short-run substitution between corn and soybeans. This was particularly evident in 2007, when a sharp increase in the corn futures price led to significant substitution of corn for soybeans. The practical problem associated with modeling demand for each crop in isolation is that the estimated elasticities will be biased towards zero. This assumption also precludes the analysis of the setting where seed firms jointly choose corn and soybean prices. Indeed, large firms such as Monsanto and Dupont may enjoy an additional markup by virtue of the fact that they have strong positions in both the corn and soybean seed markets.

Lastly, we permit our estimates for the marginal benefits associated with GE traits to differ over three important sub-periods: (i) 1996-2000, (ii) 2001-2006, and (iii) 2007-2011. By doing so, we allow the estimates to reflect certain important institutional characteristics, as well as two major events. As with any new innovation, there is a learning period, and in the case of corn and soybean seed varieties, the incorporation of new GE traits does not occur instantaneously. Over time, a wider array and better yielding set of varieties receive the technology. The presence of these factors are in part what produces the classic adoption patterns discussed by Griliches in his paper on corn hybrids (Griliches 1957). What these factors also imply is that the benefits associated with GE varieties change over time.

As noted, there were also two important events during the period we study. The first was the expiration of Monsanto's glyphosate patent in 2000. This subsequently led to a fall in glyphosate prices of more than 70% over the next ten years. Because glyphosate is so closely tied to GE glyphosate tolerant soybeans, their marginal benefits will have certainly reflected this. The second major event was the sharp increase in corn and soybean output prices that began in 2007. Because an increase in output prices raises the relative value of higher yielding inputs, GE varieties, whose yield has been shown to be superior, particularly for insect resistant varieties (Xu et al. 2013), will become more valuable relative to non-GE varieties. This in turn will be reflected in their estimated marginal benefits.

Our results generally confirm these facts. We find that farmers' willingness to pay for GE traits tended to be small early on and high in the final sub-period. In general, there were significant differences across time, with all values increasing by at least 100%. For example, for corn varieties resistant to the European Corn Borer, we estimate that farmers were willing to pay just \$2.15 per acre in the first sub-period and \$14.22 per acre in the final sub-period. In addition, we also find that the benefits for multiple GE traits (in a single seed) were larger compared to single-trait seeds but sub-additive, a finding consistent with the work of Shi et al. (2012). Perhaps most interestingly, a comparison of the WTP estimates to the actual average premiums charged by seed firms shows that the difference was largest in the final sub-period. For example, soybean farmers were willing to pay \$19.34 for glyphosate tolerance in final sub-period, whereas seed firms charged an average of just \$6.98 for it, a difference of \$11.36. By contrast, the difference was just \$.59 in the first sub-period. We interpret this as suggesting that, counterintuitively, the gains for GE varieties were in fact greatest during the period in which seed prices rose the most. The caveat to this interpretation is that it hinges on the assumption that non-GE prices would not have been significantly different in the counterfactual without GE varieties. This provides an avenue for future work in which the prices of non-GE varieties are simulated.

The rest of this paper proceeds as follows. We first motivate the problem at hand with a discussion of the data, as well some important details and trends for the corn and soybean

industries. This is followed by the setup of the demand model and the econometric specification. We then present the results, followed by concluding remarks.

Data

The main sources of data are the soybean and corn TraitTrak® datasets, two large, farm-level commercial datasets assembled by GfK Kynetec. GfK constructs the TraitTrak® data from annual surveys of randomly sampled farmers in the United States. The samples are developed to be representative at the crop reporting district (CRD) level, a multi-county area identified by the National Agricultural Statistics Service of the USDA (Fig. S1). Because this is the finest level at which the data are representative, as well as the level at which agro-climatic conditions roughly vary, we define a market as a CRD-year combination.¹⁹ This definition determines the choice-set for farmers, as well the level at which prices and market shares are computed.

In the survey, farmers are questioned about the types, amounts, and cost of the seed they purchase. In particular, for each seed variety we observe whether it contains one or multiple GE traits (e.g., glyphosate tolerance (GT)), the parent company (e.g., Monsanto), and the brand (e.g., Asgrow). We define a variety (or product) as a crop-brand-GE trait combination (e.g., Asgrow-GT is a variety in our analysis).²⁰ Combined with our market definition, this results in a total of 38,009 variety-market observations.

The primary product level variables to be calculated are market shares and prices. Market shares are computed as the ratio of total acres grown (for a given variety) to the sum of all soybean and corn acres grown within the market. Following Goeree (2008) and Eizenberg

¹⁹ Corn and soybean varieties are bred to possess characteristics that match a particular agro-climatic region. For example, the more northern U.S. regions have shorter growing seasons, so seeds bred for those regions have shorter maturity time frames. This, combined with the fact that the data is representative at the CRD-level, made it an obvious first choice for our definition of a market.

²⁰ We actually observe the varietal number of each seed, which is at an even finer level than the brand level. Defining a variety at this level, however, would result in more than 10,000 varieties.

(2014), we construct prices for each variety in each market by dividing total sales by total acres planted.

We also have data on the Chicago Mercantile Exchange soybean and corn futures prices, both obtained from www.quandl.com, as well data on the consumer price index (CPI), which was obtained from Quickstats on the USDA-NASS website. The soybean prices are the average futures prices in the month of January for delivery in November, and the corn prices are the average prices in January for delivery in September.

The Corn and Soybean Seed Industries

GE varieties commercialized so far include two families of traits: herbicide tolerance (HT) and insect resistance (IR). These traits are attractive to farmers because they address two of the most important sources of yield losses: weed pressure and insect damage. GE varieties were first introduced in the United States by Monsanto as Roundup Ready® soybeans in 1996. The Roundup Ready® trait confers tolerance to the effective and broad-spectrum herbicide glyphosate, permitting over the top applications after the soybean plants have emerged (henceforth, Roundup Ready soybeans will be referred to as glyphosate tolerant (GT) soybeans). In contrast to maize, where multiple GE traits have been commercialized and widely adopted, the GT trait has essentially been the only widely adopted transgenic trait in soybeans. The adoption of GT soybeans occurred rapidly, surpassing even the rate at which corn hybrids were adopted (Griliches 1957). By 2003, over 90% of land was planted to GT soybeans, and from then on adoption never fell below 90% (Figure. 1). The basis for such rapid and widespread adoption is primarily due to improved weed control and a reduction in management time (Qaim 2009; Barrows, Sexton and Zilberman 2014).

For maize, both GT and numerous IR traits have been widely adopted, either alone or stacked. IR traits use genes from *Bacillus thuringiensis* (Bt), a soil-dwelling bacterium known for its biological pesticide properties. Earlier Bt maize varieties were resistant only to the European corn borer (CB), first introduced in 1996, but Bt varieties resistant to rootworm (RW) were introduced subsequently in 2003 (GT corn was released in 1998). The attractiveness of Bt

varieties is that they increase expected yields and reduce yield volatility—in particular, yield losses are greatly reduced when pest pressure is high (Fernandez-Cornejo et al., 2014). The right panel of Figure 1 reports adoption rates for the various types of GE corn varieties. Note that “CB” and “GT”, as well as all of the other GE trait combinations in Figure 1, represent seeds that contain that exact configuration; i.e., with the exception of the green line, the adoption rates sum to one.

As noted, in corn more than one trait can be inserted into a seed: these seeds are referred to as stacked seeds. Stacked seeds have been very prominent in the corn seed industry, especially in recent years. Though corn with just CB trait was initially the most widely adopted, by 2011, 70% of corn varieties had multiple GE traits. By that time, the only widely adopted variety with a single GE trait was GT corn at 18%, and GE corn varieties on the whole accounted for more than 90% of purchased seed. In our econometric analysis, we consider a set of seven mutually exclusive crop-trait configurations. For soybeans, there are two possibilities: GT and non-GT varieties. For corn, there are five: CB, GT, GT-CB, GT-CB-RW, and non-GE varieties. For use later on, denote an element of this set by κ .²¹ Note that in corn we also observe RW, CB-RW, and GT-RW, so in principle we could have three more crop-trait combinations. However, these varieties were very rarely adopted in our sample (around 1% of observations each), so we drop them from the analysis.

²¹ This is important for our analysis later on because in our econometric framework we use dummy variables for GE traits that are not mutually exclusive.




Figure 1. Soybean (left) and Corn (right) Adoption Rates (% of planted acres)




Figure 2. Soybean (left) and Corn (right) Seed Prices (nominal \$/acre)

We can get an idea of the degree to which farmers value GE traits by looking at their prices. Indeed, the fast adoption of GE varieties occurred even despite significant premiums charged relative to non-GE varieties (Figure 2). In both corn and soybeans, the premium charged for a single trait, whether it be the GT trait or the CB trait, averaged a bit over \$9 (Table

1). Having both the GT and CB traits in corn cost an average of \$14.36 and to get all three traits, \$28.33. Perhaps the starkest trend in prices was the sharp increase that began 2007. Over the period 2007-2011, seed prices increased by more than 50% in soybeans and 70% in corn.

Table 1. Average GE Variety Premiums (difference relative to non-GE varieties, nominal \$/acre)

	Soy GT	Corn GT	Corn CB	Corn GT-CB	Corn GT-CB-RW
Premium	9.55	9.81	9.35	14.36	28.33

These trends cannot be fully appreciated without looking at three other major related factors: glyphosate prices, crop output prices, and firm market shares. Each of these variables may not only in part explaining seed prices, but they also clearly indicate the need for time-varying marginal benefits for GE traits, as well as a statistical framework that accounts for the presence of imperfect competition.

Figure 3 documents the trends in glyphosate and crop output prices from 1998-2011. The price of glyphosate is important because variations in its magnitude lead to variations in the benefits associated with GT varieties: when glyphosate prices are lower, for a given seed price, the benefits associated with GT varieties rise. In 2000, Monsanto's patent ended, leading to a steady decline in the price per pound (with the exception of 2007-2008). From 1998-2011, prices fell by roughly 70%. This motivates our choice, discussed above, to permit the marginal benefits associated with GT varieties to be different post-2000.

The other major important related event was the sharp increase in crop output prices that began in 2007. In just two years – 2007 and 2008 – corn and soybean prices doubled. The effect of such large shocks to output prices is to raise farmers demand for the required inputs, and so it is unsurprising that seed prices went up so much. What is interesting is that in the glyphosate market, which can be largely described as competitive by 2007, the shock to prices was not persistent: eventually supply caught up and glyphosate prices came back down. By contrast, seed prices did not come back down, raising the question of how large a role imperfect competition plays in the industry.

Since the introduction of GE crops there have been major changes to organization of the seed industry. From 1996-2011, the four-firm concentration ratio increased from 35.7% to 71.7% in soybeans and 61.6% to 79.9% in corn. To understand these changes some background on each industry is necessary.




Figure 3. Trends in nominal Glyphosate and Expected Crop Output Prices, 1998-2011

Before the introduction of GT soybeans, the soybean seed market could be characterized as essentially competitive. Commercial soybeans have the property that they reproduce “true to type”; i.e., if a farmer plants the yield from a given soybean seed it will grow to be an exact replica of its parent (note that is not the case in maize). As a result, the ability to save seed – using your own yield for next season’s seed – or purchase seed locally from last year’s yield (so-called “bin run” seed), created a competitive fringe that significantly limited the ability of seed producers to charge a price in excess of marginal cost. Whatever difference there was between the cost of seed for planting and the cost of soybeans on the output market, most of it could be imputed to cleaning, conditioning, and certification costs (Fernandez-Cornejo et al. 2004).

Table 2. Two Year Average Market Shares for the U.S. Soybean Industry, 1996-2011

Year	1996- 1997	1998- 1999	2000- 2001	2002- 2003	2004- 2005	2006- 2007	2008- 2009	2010- 2011
Monsanto ¹	1.8	19.5	21.2	21.2	20.4	25.5	28.9	27.1
Dupont/Pioneer ²	15.3	14.7	17.8	20.7	23.3	25.3	26.9	31.5
Syngenta ³			2.1	4.7	9.3	11.5	10.6	10.3
Dow Agro		1.1	2.0	1.8	1.6	1.5	1.6	2.2
AgReliant ⁴			0.7	1.5	1.6	2.0	1.7	1.9
Beck's	0.4	0.7	1.0	1.2	1.2	1.4	1.5	2.7
Croplan	1.0	1.1	1.7	3.0	2.9	3.6	3.1	2.6
Stine	3.1	3.3	3.4	2.7	2.3	2.2	2.5	2.0
Pub./Saved ⁵	27.8	19.6	15.5	7.6	3.5	1.8	1.8	1.6
CR4	35.7	42.4	46.0	50.5	55.8	65.9	69.5	71.7
HHI	0.043	0.067	0.083	0.096	0.109	0.146	0.170	0.187

¹Monsanto acquired DeKalb in 1997 and Asgrow in 1998.

²Pioneer was fully purchased by Dupont in 1999.

³Novartis and Astra-Zeneca combined to form Syngenta in 2000.

⁴AgReliant was formed in 2000 by a joint venture between Limagrain and KWS.

⁵Pub./Saved seed includes all acres that were planted with saved or public seed.

Data Source: GfK Kynetec.

Table 3. Two Year Average Market Shares for the U.S. Corn Industry, 1995-2011

Company	1995- 1997	1998- 1999	2000- 2001	2002- 2003	2004- 2005	2006- 2007	2008- 2009	2010- 2011
Monsanto ¹		8.4	10.7	11.8	15.6	26.9	34.5	33.8
Dupont/Pioneer ²	38.6	36.2	36.7	34.7	32.4	30.1	30.1	32.4
Syngenta ³			3.2	6.2	9.6	11.0	7.7	7.3
Dow Agro ⁴		1.8	5.0	5.1	3.8	3.3	3.6	4.4
AgReliant ⁵			1.6	3.5	3.8	5.7	5.8	6.4
Beck's	0.3	0.8	1.0	1.0	1.1	1.4	1.4	1.5
Croplan	0.6	0.6	1.2	2.3	1.7	2.4	2.2	1.8
CR4	61.6	58.7	57.2	57.9	61.3	73.6	78.1	79.9
HHI	0.175	0.156	0.158	0.149	0.144	0.181	0.221	0.232

¹Monsanto acquired DeKalb in 1997 and Asgrow in 1998.

²Pioneer was fully purchased by Dupont in 1999.

³Novartis and Astra-Zeneca combined to form Syngenta in 2000.

⁵AgReliant was formed in 2000 by a joint venture between Limagrain and KWS.

Data Source: GfK Kynetec.

With GT soybeans, however, growers sign contracts preventing them from saving seed. Moreover, because the GT trait is covered by an enforceable utility patent, growers cannot purchase seed from their local bin off-contract and then indefinitely replant part of their yield. This is arguably the primary driving force behind the doubling of concentration levels in the soybean seed industry.

By contrast, commercial seed corn varieties are hybrids, which means that if a farmer plants his own yield in the following season, the results will be significantly less robust and uniform. As a result, there is a natural barrier to competition in corn seed, which has historically permitted a more concentrated industry. This is what accounts for the much larger CR4 in corn compared to soybeans in 1996. Nonetheless, the widespread adoption of GE corn varieties has led to even further concentration.

For our purposes, there are two main takeaways from the market share data. The first is that in both corn and soybean seed markets there is imperfect competition. The implication for our demand model is that, because prices are set by firms, they are likely endogenous. Our solution to this is to use instrumental variables (discussed below). The second is that the industry was considerably more concentrated later on, which begs the question of whether firms captures a larger share of the rents. In particular, we are interested in how the GE trait premiums charged by seed firms over time are related to the respective marginal benefits for farmers.

Demand Model

The demand for corn and soybean varieties is modeled using a random-coefficient-logit specification (Nevo 2001). Following Train and Winston (2007), we do not include a general outside good in the model. As they note, there are certain limitations associated with including a general outside good. One problem is that the way in which “utility” (per-acre profit, in our case) is typically specified for the outside good is not structural: because the general outside

good encompasses many possible options, it is hard or not meaningful to specify it as depending on various attributes such as price. As a result, its inclusion may produce bias in the coefficient estimates. An additional problem with having a general outside good is that it precludes the inclusion of demographic explanatory variables for which the population distribution is unknown. For example, in our case we observe a variable for farm-size in our sample of corn and soybean seed buyers, but we do not observe the distribution of farm-size for farmers that chose not to plant corn or soybeans. Without additional outside information about this distribution, any estimation that solely relied on the farm-size distribution of the farmers in our sample would likely produce biased coefficients.

The primary limitation of not including an outside good is that if growers substitute readily to something else besides corn and soybeans, then our price coefficients will be biased towards zero. Since our analysis is skewed towards corn and soybean farmers in the central corn belt, we believe there is little to be lost by assuming that, once a farmer has decided to plant corn or soybeans, they will always choose one or the other. One way to assess the importance of this assumption is to ask how extreme conditions would need to be for farmers to choose something else besides corn or soybeans. In our view, there would need to fairly dramatic changes for farmers to re-consider an option different from corn or soybeans. Overall, we believe the downsides associated with including a general outside option outweigh the benefits.

Concerning the model, we denote a market (CRD-year combination) by m . In each market there are $j = 0, 1, \dots, J_m$ available varieties (crop-brand-trait combination) and $i = 1, \dots, I_m$ plots, where a plot is synonymous with a choice situation (note that growers may have multiple plots). For each plot, a farmer chooses the seed variety that maximizes the following per-acre return objective function:²²

²² Our framing of the problem as one in which a grower makes a discrete choice on each plot is in many respects similar to the framework developed in Caswell and Zilberman (1985).

$$(1) \quad \pi_{ijm} = \underbrace{\beta' z_{jm} - p_{jm} + \xi_{jm}}_{\delta_{jm}} + \underbrace{\phi' x_{ijm} + \sigma' w_{ijm}}_{\psi_{ijm}} + \mu \varepsilon_{ijm}$$

where z_{jm} is a $K \times 1$ vector of seed characteristics, p_{jm} is the seed cost per acre of variety j in market m , and ξ_{jm} is a product-market specific demand shock that is unobserved by the econometrician. The latter represents factors such as how well a particular market is matched to a given seed variety. For the purposes of exposition, the mean return associated with variety j in market m is denoted by δ_{jm} . The remaining variables capture the portion of returns that vary by plot. The vector x_{ijm} contains farmer characteristics, e.g., farm-size, interacted with seed characteristics. By including these variables, we aim to capture how farmer characteristics impact the mean valuation of each of the seed attributes (in particular, the GE traits). Unobserved heterogeneity is captured by w_{ijm} , which contains the interaction of seed characteristics and mean-zero normally distributed error terms. The σ vector of parameters thus captures whether there is unobserved heterogeneity in preferences for the different seed attributes. In particular, they represent the variance of the normally distributed errors. The term ε_{ij} is an IID Type-I Extreme Value residual; it captures unobserved plot specific factors. The parameter μ characterizes the variance of ε_{ij} and is typically referred to as the scale parameter (Train 2009). This parameter plays an important role in our analysis. As we show below, the estimated parameter on the price variable is the inverse of the scale parameter. Anderson, de Palma, and Thisse (1992) interpret μ as a measure of taste heterogeneity and Caswell and Zilberman (1985) call its inverse the “water cost-saving” coefficient in their analysis of irrigation technologies. We use this parameter to obtain WTP estimates for the various GE traits.

As with all discrete choice models, only differences in returns matter, so the return to some variety needs to be normalized (since we are not including a general outside good). We index this variety by $j = 0$ and call it the reference good. We use Dupont-Pioneer non-GE corn varieties as the reference good since they are observed in the largest number of markets. Operationally, this means subtracting the observed values of the attributes for the reference good from the values of each of the inside goods. For example, in market m the price of variety

j is $p_{jm} = \tilde{p}_{jm} - \tilde{p}_{0m}$, where \tilde{p}_{jm} and \tilde{p}_{0m} are the observed average seed costs of variety j and the reference good, respectively.

Let v_i denoted the vector of zero-mean normally distributed error terms for a particular plot. Conditional on a realization of v_i , the probability that variety j is chosen is given by the logit expression:

$$(2) \quad L_{ijm}(v) = \frac{e^{(\delta_{jt} + \phi'x_{ijm} + \sigma'w_{ijm})/\mu}}{1 + \sum_k e^{(\delta_{kt} + \phi'x_{ikm} + \sigma'w_{ikm})/\mu}}$$

where the inverse of μ now scales the per acre returns in the exponential terms. The unconditional probability is given by the integral over all possible realizations of v_i :

$$(3) \quad P_{ijt} = \int L_{ijt}(v_i) f(v_i) dv_i.$$

Because this expression has no closed-form, the standard practice is to approximate it using simulation. For each individual, we compute:

$$(4) \quad \tilde{P}_{ijt} = \frac{1}{R} \sum_{r=1}^R L_{ijt}(v_r)$$

where R is the total number of draws. In this application we elect to estimate the simpler logit model.²³ This is equivalent to assuming that $\psi_{ijt} = 0$. This collapses the predicted probabilities into a simple expression, which can be algebraically manipulated into the basic linear logit equation derived in Berry (1994):

$$(5) \quad \ln(s_{jm}) - \ln(s_{0m}) = \tilde{\beta}'z_{jm} - \alpha p_{jm} + \tilde{\xi}_{jm}$$

²³ Future work will include an extension of the econometric framework to accommodate both observed and unobserved heterogeneity.

where $\alpha = 1/\mu$, $\tilde{\beta} = \beta/\mu$, $\tilde{\xi}_{jm} = \xi_{jm}/\mu$; and s_{jm} and s_{0m} are the shares of variety j and the reference good, respectively, in market m . As noted, α is the parameter for calculating farmers' WTP: dividing $\tilde{\beta}$ by α gives the un-scaled coefficient vector β .

Variables

The main variables included in the vector z_{jm} include brand fixed effects and sub-period-specific dummy variables for each of the possible GE trait configurations. Brand-fixed effects are included to control for unobserved quality differences in brands. For example, Pioneer varieties are typically perceived as being of higher quality than other brands; to the extent that quality is correlated with price, their exclusion would result in bias. For the GE traits we include a total of five different types of dummies, denoted by d_τ , where $\tau \in \Omega$:

$$(6) \quad \Omega = \{\text{soy-GT, corn-GT, corn-CB, corn-GT-CB, corn-GT-CB-RW}\}$$

For ease of interpretation, these dummy variables merely indicate the *presence* of a certain trait combination and thus are not mutually exclusive. For example, the dummy variables for corn varieties with the double stack of GT and CB would take on the values of:

$$d_{\text{corn-GT}} = 1, d_{\text{corn-CB}} = 1, d_{\text{corn-GT-CB}} = 1$$

with all other dummies equal to zero. Each possible combination of dummies corresponds to an element κ , noted earlier. The primary reason we specify it this way is that the coefficient on $d_{\text{corn-GT-CB}}$ now conveys whether returns are sub-additive in the GT and CB traits. In keeping with the previous discussion, we estimate these dummies separately for each of the three sub-periods: (i) 1996-2000, (ii) 2001-2006, and (iii) 2007-2011.

In certain specifications, we also include a soybean dummy variable, annual soybean and corn futures prices, crop-year specific dummy variables and crop-year specific CRD effects. The soybean dummy variable controls for time-invariant unobserved differences between soybean and corn seed. The futures prices and crop-year specific dummy variables are included to control for crop-specific differences in the returns over time. Crop-year specific CRD effects

control for unobserved differences in how well particular regions are matched to each of the two crops.

Instruments

Because seed firms observe ξ_{jm} , prices are likely correlated with ξ_{jm} (Trajtenberg 1989). For example, if ξ_{jm} captures the degree to which a market matches variety j , then seed firms may charge higher prices for varieties that are well-matched to a particular region.²⁴ As a result, the OLS estimator for α will be biased towards zero. The solution to this problem is to use instrumental variables for the price variable. Following Berry et al. (1995) we use the sums of the characteristics in competing varieties as our instruments. Since the GE traits are the main characteristics that vary over seed varieties, this amounts to counting up the unique number of GE varieties. Specifically, we calculate four sets of sums. They are the total number of competing varieties with a particular trait configuration by: (i) market, (ii) company, (iii) crop, and (iv) company and crop. With seven possible crop-trait configurations, there are twenty-eight total instrumental variables.

Results

Basic summary statistics for the main included variables are provided in table 4. During estimation, all prices are deflated by the CPI. Average seed costs over the entire sample were \$36.47 per acre, and average soybean and corn futures prices were \$7.25 and \$3.20 per bushel, respectively. Soybean varieties accounted for 43% of the sample, with GT varieties accounting for 72% of soybean varieties. In corn, 49% of varieties were conventional, 36% contained the CB trait, 37% had the GT trait, 22% had the GT-CB combination, and 13% had the GT-CB-RW combination.

²⁴ For example, seeds with longer maturity time-frames will be less demanded in more northern regions, and so may also command a lower price (relative to the price commanded in more southern regions). Evidence in Stiegert et al. (2011) supports this presence of spatial price variation.

Table 4. Descriptive statistics for selected variables in the estimated model

Variable	Mean	Std. Dev.	Min	Max
Seed Cost (\$/acre)	36.47	18.11	1.11	130.27
Soybean Futures (\$/bu)	7.25	2.52	4.48	13.13
Corn Futures (\$/bu)	3.20	1.04	2.31	5.68
Dummy Variables				
<u>Soy</u>	0.43	0.49	0.00	1.00
GT	0.71	0.46	0.00	1.00
<u>Corn</u>	0.57	0.49	0.00	1.00
non-GE	0.49	0.50	0.00	1.00
CB	0.36	0.48	0.00	1.00
GT	0.37	0.48	0.00	1.00
GT-CB	0.22	0.41	0.00	1.00
GT-CB-RW	0.13	0.34	0.00	1.00

We present coefficient estimates for four different specifications. All specifications included brand fixed effects; column (2) adds soybean and corn futures prices, column (3) replaces futures prices with crop-specific time effects, and column (4) adds crop-specific CRD effects. The results generally accord with expectations. The coefficient on seed price is negative and significant in all four cases with average own price elasticities ranging from -1.95 (column 3) to -3.34 (column 2). The coefficient estimates for the futures prices are both positive, indicating that profit per acre for both corn and soybean varieties increases when corn and soybean futures prices rise.

The most interesting results are for the GE coefficient estimates. For the basic GE effects – soy-GT, corn-GT, corn-CB – the estimated coefficients are positive and significant in most sub-periods. The exception is the negative coefficient in the first sub-period for corn-GT varieties, which suggests poor availability of good hybrids and/or lack of awareness of the GT trait. This contrasts with GT soybeans, for which there was a positive coefficient in all periods. This is consistent with the lack of alternative herbicide options in soybeans (relative to corn). In most cases, the GT-CB coefficients are negative and significant. This means that the return to varieties with both the GT and CB traits is less than the sum of their return; i.e., profits are sub-additive in the GT-CB stack. This is consistent with the finding of sub-additive prices in Shi et al. (2012).

The corn GT-CB-RW is positive and significant in the final sub-period, indicating that growers value the RW trait. Note that because we do not estimate dummies for the RW trait on its own, we cannot say what its value is in isolation.

In all specifications there is a clear upward trend in growers' willingness to pay for GE traits. As discussed previously, this is potentially indicative of learning effects, better GE varieties, falling glyphosate prices, and rising output prices. The increase in the coefficients is particularly large for corn GE traits in the final sub-period, which is consistent with a perceived yield effect: the sharp increase in corn output prices raised the differential monetary gain to GE varieties. By contrast, the increase in the soy GT coefficient is largest from the first sub-period to the second sub-period and more muted from the second sub-period to the final. Overall, this is unsurprising given the observed adoption patterns. The relative profitability of the various GE traits is pinned down by the shares of the varieties that contain them (controlling for prices, brand effects, etc.).

Willingness-to-pay for GE traits

While informative, the coefficient estimates do not convey the monetary magnitude of the various GE traits. We calculate willingness-to-pay (WTP) estimates for the various GE trait configurations using the coefficient estimates from column 4 in table 5. For a given trait configuration, κ , the WTP is given by the ratio of the sum of the estimated GE coefficients to the price coefficient:

$$(7) \quad WTP_{\kappa} = \frac{1}{\alpha} \sum_{\tau \in \Omega} \tilde{\beta}_{\tau} d_{\tau}$$

where $\tilde{\beta}_{\tau}$ is the coefficient on the crop-trait combination τ . For the calculations, we use column 4 because of its superior fit and the fact that, through its additional controls, it is the most robust to omitted variables. Nonetheless, the WTP estimates we would obtain from the other specifications are quite similar.

WTP estimates for each of the possible GE-trait combinations are provided in table 6. Most of the estimates are positive and tightly estimated. All of the estimates are generally reasonable and in line with what might be expected given knowledge of seed prices and the observed adoption patterns by farmers. Marginal benefits range from as high as \$41.88 in the 2007-2011 sub-period for GT-CB-RW seed corn to as low as -\$3.50 for GT-only corn varieties in the first sub-period. Notably, the WTP for GT soybean varieties is significant and positive in all sub-periods, reaching a high of \$19.34 in the 2007-2011 period. As was the case for coefficient estimates, there is a clear and significant upward time-path.

To provide additional context for the estimates, in particular for their changes over time, table 6 also contains statistics for the average premium – the difference between the average price of non-GE varieties and the average price for the respective GE varieties – charged by seed firms, as well as the difference between those premiums and farmers' WTP. What's particularly interesting is that the average premiums charged by seed firms do not actually change as much as one might expect. In fact, the average premium charged for GT varieties in the 2007-2011 period was less than in both previous sub-periods (recall that all prices are deflated by the CPI; in nominal terms, the premiums are all higher in later periods). Under certain assumptions, this implies that growers increasingly benefited from GE varieties over time.²⁵ In the 2007-2011 period, the share of differential marginal benefits captured by seed firms ranged from 36% for GT soybeans to 63% for CB corn. In the other cases, the ratio was between 50% and 60%. We interpret this as indicating that, although seed firms enjoyed significant market power, farmers gained significantly from GE varieties, and, perhaps counterintuitively, the majority of those gains occurred during the sub-period in which seed costs rose the most.

²⁵ This statement is premised on the assumption that conventional seed prices would have been the same in the counterfactual where GE varieties did not exist. This is an admittedly strong assumption, particularly given the presence of imperfect competition. To test the accuracy of this conclusion, future work will need to combine the demand side with the supply side in order to simulate what non-GE prices would have been in the absence of GE varieties.

Table 5. Demand model 2SLS Regression Results (dependent variable is $\ln(s_{jm}) - \ln(s_{0m})$)

Variable	(1)	(2)	(3)	(4)
Price (\$/acre)	-0.0939*** (0.0063)	-0.0949*** (0.0064)	-0.0555*** (0.0069)	-0.0671*** (0.0089)
Soy Dummy	-0.9025*** (0.0467)	-0.7320*** (0.0929)		
Soy Futures (\$/bu)		0.2011*** (0.0222)		
Corn Futures (\$/bu)		0.5356*** (0.0750)		
Soy GT Effects				
1996-2000	0.9275*** (0.0629)	0.9512*** (0.0624)	0.5522*** (0.0665)	0.6207*** (0.0787)
2001-2006	1.3404*** (0.0502)	1.3543*** (0.0478)	1.1136*** (0.0631)	1.1525*** (0.0746)
2007-2011	2.2634*** (0.0353)	2.3724*** (0.0538)	1.2444*** (0.0868)	1.2970*** (0.0965)
Corn GT Effects				
1996-2000	-0.1006 (0.0976)	-0.0949 (0.0978)	-0.3040*** (0.0914)	-0.2348* (0.0945)
2001-2006	0.5031*** (0.0509)	0.5089*** (0.0512)	0.2739*** (0.0517)	0.3551*** (0.0612)
2007-2011	1.8428*** (0.0608)	1.8511*** (0.0613)	1.5415*** (0.0631)	1.6385*** (0.0771)
Corn CB Effects				
1996-2000	0.3327*** (0.0629)	0.3409*** (0.0632)	0.0523 (0.0637)	0.1439 (0.0749)
2001-2006	0.5901*** (0.0508)	0.5962*** (0.0512)	0.3398*** (0.0526)	0.4221*** (0.0635)
2007-2011	1.0864*** (0.0583)	1.0909*** (0.0586)	0.8876*** (0.0568)	0.9534*** (0.0630)
Corn GT-CB Effects				
1996-2000	0.1750 (0.3940)	0.1678 (0.3946)	0.3127 (0.3600)	0.2709 (0.3532)
2001-2006	-0.2962*** (0.0640)	-0.2993*** (0.0642)	-0.1494* (0.0602)	-0.2019** (0.0630)
2007-2011	-0.8863*** (0.0700)	-0.8882*** (0.0701)	-0.7992*** (0.0645)	-0.8321*** (0.0645)
Corn GT-CB-RW Effects				
2001-2006	0.5151*** (0.1245)	0.5233*** (0.1249)	0.2412* (0.1169)	0.3169** (0.1221)
2007-2011	1.2689*** (0.0692)	1.2780*** (0.0696)	0.9609*** (0.0701)	1.0489*** (0.0830)
Fixed Effects	Brand	Brand	Brand, Year-Crop	Brand, Year-Crop, CRD-Crop
<i>N</i>	38,009	38,009	38,009	38,009
<i>R</i> ²	0.156	0.154	0.297	0.325

Notes: Standard errors are in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6. Willingness to Pay (WTP) for GE Seed Varieties (\$/acre)

	WTP Estimate	Average Premium	Difference	Ratio
Soy GT				
1996-2000	9.25*** (0.53)	8.66	0.59	94%
2001-2006	17.18*** (1.42)	7.76	9.43	45%
2007-2011	19.34*** (1.87)	6.98	12.35	36%
Corn GT Only				
1996-2000	-3.50* (1.69)	5.68	-9.18	-162%
2001-2006	5.29*** (0.46)	4.99	0.30	94%
2007-2011	24.43*** (2.23)	11.86	12.57	51%
Corn CB Only				
1996-2000	2.15* (0.88)	8.38	-6.23	391%
2001-2006	6.29*** (0.39)	7.13	-0.84	113%
2007-2011	14.22*** (1.36)	5.21	9.01	63%
Corn GT-CB				
1996-2000	2.68 (5.14)	9.65	-6.97	360%
2001-2006	8.58*** (0.58)	8.94	-0.37	104%
2007-2011	26.24*** (2.11)	12.80	13.43	51%
Corn GT-CB-RW				
2001-2006	13.30*** (1.66)	17.30	-4.00	130%
2007-2011	41.88*** (1.96)	25.04	16.83	60%

Notes: WTP estimates are computed using the estimated coefficients from column 4 of Table 5.

Each estimate is the ratio of the respective GE total effect to the estimated Price coefficient. All prices are deflated by the CPI. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

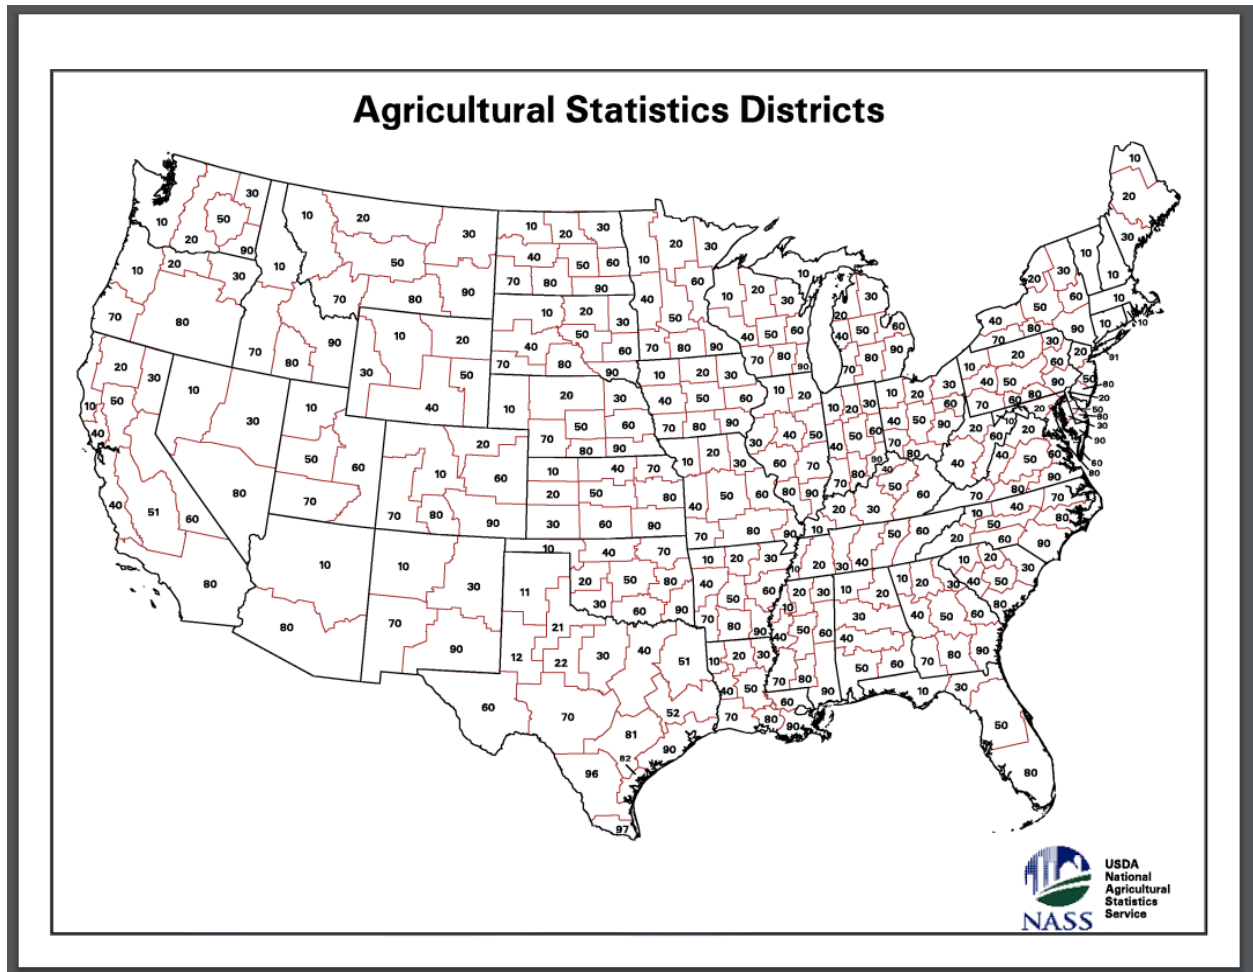
Conclusion

In this paper we develop and estimate a structural discrete choice model of corn and soybean seed demand for the period 1996-2011. The model is specified and estimated to take into account imperfect competition in the seed industries. We examine farmers' WTP across two important events: the expiration of Monsanto's glyphosate patent and the sharp increase in commodity prices that began in 2007. We find that farmers were almost always willing to pay a significant premium for GE varieties and the extent of that willingness increased significantly over time. We attribute this to learning effects, the gradual release of a wider array and better yielding GE varieties, decreasing glyphosate prices, and increase in output prices. The increase in corn and soybean outprices seems to have had a particularly large effect on the marginal benefits of corn GE varieties, suggesting a strong perceived yield premium. In addition, we find that the WTP estimates for stacked GE varieties are sub-additive, a finding consistent with Shi et al. (2012).

We also compare farmers' WTP for GE varieties to the average premiums charged by seed firms. Interestingly, we find that the WTP estimates for GE varieties exceeded average premiums by the most in the final sub-period, the period during which seed prices increased sharply. Why might this be? One possibility is that there's a lag or institutional stickiness to seed prices, so that when corn and soybean futures prices suddenly go up, and farmers are willing to pay substantially more for the better varieties, seed firms don't immediately fully adjust.

It is also important to note the limitations of our findings. One limitation is that we do not estimate firms' marginal costs, so we cannot ascertain the exact distribution of *ex post* GE returns between seed firms and farmers. Nor do we simulate the counterfactual in which GE varieties are not available, which would permit us to ascertain the extent to which non-GE prices have been impacted by GE varieties (through the induced re-structuring of the seed industries). Both of these extensions constitute important avenues for future research.

Fig. S1. Crop Reporting Districts (CRD), National Agricultural Statistics Service,
U.S. Department of Agriculture.



References

- Anderson, S. P., De Palma, A., & Thisse, J. F. (1992). *Discrete choice theory of product differentiation*. MIT press.
- Barrows, G., Sexton, S., & Zilberman, D. (2014). Agricultural Biotechnology: The Promise and Prospects of Genetically Modified Crops. *The Journal of Economic Perspectives*, 28(1), 99-119.
- Bennett, A. B., Chi-Ham, C., Barrows, G., Sexton, S., and Zilberman, D. (2013). "Agricultural Biotechnology: Economics, Environment, Ethics, and the Future." *Annual Review of Environment and Resources* 38: 249–279.
- 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: Journal of the Econometric Society*, 841-890.
- Berry, S., Levinsohn, J., & Pakes, A. (2004). Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market. *Journal of Political Economy*
- Bullock, D., and Desquilbet, M. (2002). "The Economics of Non-GMO Segregation and Identity Preservation." *Food Policy* 27: 81–99.
- Carter, C., Moschini, G. and Sheldon, I., eds. (2011). *Genetically Modified Food and Global Welfare*. Bingley, U.K.: Emerald Group Publishing Limited.
- Caswell, M., & Zilberman, D. (1985). The choices of irrigation technologies in California. *American Journal of Agricultural Economics*, 67(2), 224-234.
- Eizenberg, A. (2014). Upstream innovation and product variety in the us home pc market. *The Review of Economic Studies*, 81(3), 1003-1045

- Falck-Zepeda, J. B., Traxler, G. and Nelson, R. G. (2000). "Surplus Distribution from the Introduction of a Biotechnology Innovation." *American Journal of Agricultural Economics* 82: 360–69.
- Fernandez-Cornejo, J. (2004). *The Seed Industry in U.S. Agriculture: An Exploration of Data and Information on Crop Seed Markets, Regulation, Industry Structure, and Research and Development*. Washington DC: USDA, AIB No. 786
- Fernandez-Cornejo, J., Wechsler, S. J., Livingston, M., and Mitchell, L. (2014). *Genetically Engineered Crops in the United States*. Economic Research Report Number 162, U.S. Department of Agriculture, February 2014.
- Goeree, M. S. (2008). Limited information and advertising in the US personal computer industry. *Econometrica*, 76(5), 1017-1074.
- Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 501-522.
- Huang, J., Hu, R., Rozelle, S. and Pray, C. (2005). "Insect-Resistant GM Rice in Farmers' Fields: Assessing Productivity and Health Effects in China." *Science* 308: 688–690.
- James, C. 2015. 20th Anniversary (1996 to 2015) of the Global Commercialization of Biotech Crops and Biotech Crop Highlights in 2015. ISAAA Brief No. 51. ISAAA: Ithaca, NY.
- Moschini, G., (2008). "Biotechnology and the Development of Food Markets: Retrospect and Prospects." *European Review of Agricultural Economics* 35(2008): 331–355.
- Moschini, G. (2010). Competition issues in the seed industry and the role of intellectual property. *Choices*, 25(2), 1-12.
- Moschini, G., Lapan, H. and Sobolevsky, A. (2000). "Roundup Ready Soybeans and Welfare Effects in the Soybean Complex." *Agribusiness* 16: 33–55.

- Moss, D. L. (2010). Transgenic seed platforms: competition between a rock and a hard place? Addendum. *American Antitrust Institute (AAI)* [http://www. antitrustinstitute. org/Archives/seed. ashx](http://www.antitrustinstitute.org/Archives/seed.ashx).
- National Research Council (NRC). 2010. *The Impact of Genetically Engineered Crops on Farm Sustainability in the United States*. Washington, D.C.: National Academies Press.
- Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2), 307-342.
- 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., Ciliberto, F., Hennessy, D. A., & Moschini, G. (2016). Genetically Engineered Crops and Pesticide Use in U.S. Maize and Soybeans. Unpublished manuscript.
- Petrin, A. K. (2002). Quantifying the Benefits of New Products: The Case of the Minivan. *Journal of Political Economy*, 110(4), 705-729.
- Qaim, M. (2009). The economics of genetically modified crops. *The Annual Review of Resource Economics*, 1.
- Qaim, M., & Zilberman, D. (2003). Yield effects of genetically modified crops in developing countries. *Science*, 299(5608), 900-902.
- Shi, G., Chavas, J.P., Stiegert, K., (2010). An analysis of the pricing of traits in the U.S. corn seed market. *American Journal of Agricultural Economics*, 92(5): 1324--1338.
- Shi, G., Stiegert, K., Chavas, J.P., (2011). An analysis of bundle pricing in horizontal and vertical markets: The case of the U.S. cottonseed market. *Agricultural Economics*. 42, supplement: 77-88.
- Shi, G., Chavas, J.P., Stiegert, K., (2012). An analysis of bundle pricing: the case of biotech seeds. *Agricultural Economics*. 43, supplement: 125-139.

- Shoemaker, R., Harwood, J. L., Day-Rubenstein, K. A., Dunahay, T., Heisey, P. W., Hoffman, L. A., Klotz-Ingram, C., Lin, W. W., Mitchell, L., McBride, W. D., and Fernandez-Cornejo, J. (2001). *Economics Issues in Agricultural Biotechnology*. Agricultural Information Bulletin No. 762, Economic Research Service, U.S.D.A., February.
- Stiegert, K., Shi, G., Shavas, J.P., (2011). Spatial Pricing of Genetically Modified Hybrid Corn Seeds. In: Carter, Moschini, and Sheldon (eds.), *Genetically Modified Food and Global Welfare* (Frontiers of Economics and globalization, Volume 10), Emerald Group Publishing Limited.
- Train, K. (2009). *Discrete choice methods with simulation*. Cambridge university press.
- Train, K. E., & Winston, C. (2007). Vehicle Choice Behavior and the Declining Share of US Automakers. *International Economic Review*, 48(4), 1469-1496.
- Trajtenberg, M. (1989). The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners. *Journal of Political Economy*, vol. 97, no. 2.
- Xu, Z., D.A. Hennessy, K. Sardana and G. Moschini, "The Realized Yield Effect of Genetically Engineered Crops: U.S. Maize and Soybean," *Crop Science*, 53(2013):735-745.