

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Economics of Research and Innovation in Agriculture

Volume Authors/Editors: Petra Moser, editor

Volume Publisher: NBER

Volume ISBNs: 978-0-226-77905-8 (cloth),
978-0-226-77919-5 (electronic)

Volume URL:
<https://www.nber.org/books-and-chapters/economics-research-and-innovation-agriculture>

Conference Date: May 17, 2019

Publication Date: September 2021

Chapter Title: Quantifying Heterogeneous Returns to Genetic Selection: Evidence from Wisconsin Dairies

Chapter Author(s): Jared Hutchins, Brent Hueth, Guilherme Rosa

Chapter URL:
<https://www.nber.org/books-and-chapters/economics-research-and-innovation-agriculture/quantifying-heterogeneous-returns-genetic-selection-evidence-wisconsin-dairies>

Chapter pages in book: p. 81 – 105

Quantifying Heterogeneous Returns to Genetic Selection

Evidence from Wisconsin Dairies

Jared Hutchins, Brent Hueth, and Guilherme Rosa

2.1 Introduction

Biological innovation is an important driver of productivity growth in the agricultural sector (Olmstead and Rhode 2008). This is especially so in the dairy sector, where milk yield has grown 3–4 percent per year during the past century; 50 percent of this growth is typically attributed to improvement in livestock genetics (Pryce and Veerkamp 2001). However, growth can be overattributed to genetic improvement when models ignore the fact that dairy farmers select genetics based on their farm’s returns to a given type of genetics. This is because in the dairy sector, the vast majority of “experimentation” undertaken to identify high-performing genetics takes place in *nonexperimental conditions*. Starting in 1908, the US Department of Agriculture (USDA) initiated a program, in partnership with land grant universities and local associations of dairy farmers, to measure and record animal-level performance (CDCB 2017). This partnership, which eventually came to be known as the Dairy Herd Improvement (DHI) program, continues to this day. Data from commercial dairy herds that participate in the DHI

Jared Hutchins is an assistant professor in the Department of Agricultural and Consumer Economics at the University of Illinois at Urbana-Champaign.

Brent Hueth is a branch chief of structure, technology and productivity at the US Department of Agriculture, Economic Research Service.

Guilherme Rosa is a professor in the Department of Animal and Dairy Sciences at the University of Wisconsin-Madison.

The findings and conclusions in this publication are those of the authors and should not be construed to represent any official USDA or US Government determination or policy. This research was supported in part by the US Department of Agriculture, Economic Research Service.

For acknowledgments, sources of research support, and disclosure of the authors’ material financial relationships, if any, please see <https://www.nber.org/books-and-chapters/economics-research-and-innovation-agriculture/quantifying-heterogeneous-returns-genetic-selection-evidence-wisconsin-dairies>.

program (roughly half of all US dairy herds) are used to estimate the “genetic merit” of individual male dairy cattle (sires). Specifically, USDA scientists use regression models to identify the impact of a sire on the performance of its offspring. Because of these *nonexperimental* conditions, these public estimates of sire productivity can contain both the quality of the genetics and the skill of the farmer in selecting them to the right environment. Without understanding the contribution of farmer management and selection to productivity, productivity growth in the dairy sector may be mistakenly attributed to the work of breeders instead of to both breeders and farmers.

Our work decomposes total productivity change on Wisconsin dairy farms due to genetics into separate effects for genetic improvement and endogenous selection. Using data from a large sample of Wisconsin dairy farms and national-level data on dairy sire rankings, we develop and estimate a model that accounts for selection behavior in the animal’s production function. We find that selection accounts for as much as 75 percent of the total productivity improvement in fat and protein yield in dairy cows. Our results provide evidence for positive assortative matching, whereby farmers who adopt above-average yield genetics also perform better than average for their chosen genetics. Further, we find that management behavior accounts for a significant portion of within-herd cow-level heterogeneity in genetic choice, suggesting that, contrary to previous evidence, farmers manage at the animal level and not solely at the herd level; this implies that previous regression models controlling only for herd-level variation do not adequately control for selection bias in the production function. Overall, our results indicate that a large portion of productivity growth in dairy farming can be explained by farmers’ ability to identify and select genetics that are well suited to their production environment—not solely the quality of the genetics they choose.

Using on-farm data to calculate the genetic merit of different sires, thus far the engine of genetic progress in the dairy industry, runs the risk of entangling the management savvy of dairy farmers with the quality of the genetics. Artificial insemination (AI) technologies, widely adopted beginning in the 1960s, expanded opportunities to track and identify the performance of sires. Modern AI technologies permit a single sire to produce hundreds of thousands of offspring, and each female offspring (birthed on the farm of a DHI program participant) contributes new sire evaluation data, improving estimates of genetic merit. Genetic merit in turn strongly influences the market price AI companies receive for their genetic material. However, the very thing that makes dairy unique also makes isolating the contribution of genetics tricky. Research by economists on technology adoption has shown that productivity gains can be overattributed to technology when there is “positive assortative matching,” meaning the very ones that benefit the most from technology are the ones who will adopt it. Hybrid seeds, for example, were adopted into their most productive environments in

the United States (Griliches 1957) as well as in Kenya (Suri 2011). Similarly, technologies such as fertilizer are often applied by farmers who have the highest returns from applying it (Foltz, Aldana, and Laris 2014).

Our work modifies the canonical modeling framework used by quantitative geneticists for determining genetic merit by accounting for the selection behavior of dairy farmers. Our research bridges two scientific domains, both indebted to the seminal work of Sewall Wright, that diverged early in the 20th century in their approaches to explaining the contribution of genetics in the farm production function. Using a control function approach, we find that the average returns to the adoption of high-yield genetics on dairy farms are as much as 75 percent lower after accounting for confounding factors. We find evidence of positive assortative matching at the cow level and the herd level, suggesting that dairy farmers manage their herds at the level of individual cows. Our work makes a novel contribution to the literature on technology adoption by investigating a level of detail in selection behavior that is previously unexplored in the agricultural sector.

We start with the idea that farmers select genetics based on their *ex ante* returns to the technology, which will cause them to “match” sire genetics to cows in a specific way. We assume further that some of the factors affecting match quality are observed by farmers but not by researchers. Using the framework of the correlated random coefficient (CRC) model, we explore the effect of this selection behavior on returns to production traits in dairy cattle and test whether correcting for selection behavior affects the estimation of the average effect. We use random variation in country-wide genetic evaluations as an instrument to identify the effect of choosing dairy sires with high predicted transmitting ability (PTA) indices for fat and protein yield in the cow’s production function, and we use the residual from the first stage to identify the heterogeneity in the effect. We find that the average return from increasing the index one unit, which is a one-unit increase in pounds of fat or protein production, is .6 and .4 for fat and protein without controlling for selection behavior. These estimates drop to .15 and .18 after controlling for selection behavior, which means that as much 75 percent of the return to high-yield genetics is explained by this matching behavior. Finally, we find that the heterogeneity in returns is over both farms and animals; this implies the high productivity gains are being driven by animal-level matching and not just farm-level matching. This changes the narrative of farm productivity in the dairy industry quite drastically: instead of the triumph of animal breeders and scientists alone, it is growth accomplished by a partnership between farmers and breeders.

2.2 Related Literature

Though seldom pointed out, the estimation of breeding values in quantitative genetics and production function estimation in economics share a

common history. The roots of both can be traced back to Sewall Wright, who pioneered work in population genetics, paving the way for the field of modern quantitative genetics (Gianola and Rosa 2015). Wright also conceptualized the notion of “path analysis,” which later evolved to be known as “instrumental variables,” and applied it to supply and demand systems as well as simultaneous equation models of commodity prices (Wright 1928; Wright 1925). From this common origin, economics and quantitative genetics diverged in point of focus with respect to estimating the following equation:

$$(1) \quad y = Z\mu + X\beta + \varepsilon,$$

where y is the output (e.g., butterfat production of a dairy cow or yield of a strain of maize), Z is an incidence matrix of genetic technologies or types, X is a matrix of “environment” covariates unrelated to genetics, and ε is the unexplained component of y .

In quantitative genetics, the parameters β are modeled as fixed, but the parameter μ is treated as the outcome of a genetic process and thus considered a random variable with a covariance matrix mapping all the relationships among genotypes. Nowhere has this genetic model developed quite the importance it has in animal breeding as a result of the work of C. R. Henderson (1953, 1973). Prior to Henderson’s work, there was no widely used method for attributing the performance of different livestock to its parents. The Henderson mixed model (HMM), still used in the US national DHI program, models breeding values as draws from the random variable μ (Henderson 1975).

The HMM has become integral to the dairy genetics industry because estimates of μ for each sire, $\hat{\mu}$, strongly influence market prices for dairy genetics. PTA, which is $\hat{\mu}/2$, is roughly interpreted as the value that a sire has for a particular trait y , which is predicted to be “transmitted” to the offspring (Van Vleck 1987). The national DHI program produces PTA values for a wide variety of traits including milk yield, fat yield, fertility, longevity, and “conformance” (elements of body structure such as udder size and height). Once published, these values influence adoption decisions, which then result in new data that feed back into the DHI program as raw data to create breeding values for new sires and to update estimates for breeding values of existing sires. Building from Sewall Wright, the HMM has become an important source of genetic progress for the dairy industry.

The field of economics developed in parallel to Henderson’s work but focused on a different set of estimation issues with respect to equation (1). In particular, the production function literature in economics has centered on the assumptions needed to identify estimate μ . If the adoption of certain genetics is associated with unobserved components of y , this means $\text{Cov}(Z, \varepsilon) \neq 0$, and standard regression approaches yield biased estimates of μ .

This bias was more generally referred to by Mundlak (1961) as “management bias,” defined as the presence of unobserved management decisions (or conditions of the decision environment) that influence input choice by farmers (genetic selection in this context). Griliches (1957) specifically suggested in the case of hybrid corn that genetic technology was historically adopted into the environment where it was the most profitable. Solutions to this problem have evolved from the simple fixed effects approach of Mundlak (1961) to invoking the wisdom of Sewall Wright and using exogenous variation to identify structural parameters of production functions (Griliches and Mairesse 1995).

More recently, labor economists have developed new frameworks for thinking about this identification issue. The Roy model (Roy 1951) posited that occupation decisions, much like technology adoption decisions, are not chosen randomly; instead, they are generated from behavior that takes into consideration *ex ante* idiosyncratic returns that are difficult to measure. This implies that measuring the returns to some decision on an outcome, such as the effect of the adoption of technology on firm output, is subject to a “selection bias” that must be dealt with in something like equation (1) (Heckman and Vytlačil 1998). A similar logic can be applied to the choice of genetics, since farmers likely observe or know something *ex ante*, unobserved to researchers, and affecting relative returns across relevant genetic profiles. Suri (2011) formalized the link between labor economics and production function estimation by using the Roy model to study selection bias in technology adoption. Her study found that farmers in Kenya adopted hybrid maize if their personal unobserved return was high, suggesting “positive assortative matching.” This has in turn helped spur a growing literature on quantifying the heterogeneous returns to agricultural technology adoption in other contexts (Foltz, Aldana, and Laris 2014; Michler et al. 2019; Zeitlin et al. 2010).

Our analysis circles back to an empirical question that has been studied for nearly 100 years: How do we evaluate the performance of animal genetics from observational data? We unite two divergent fields of study, economics and quantitative genetics, by bringing the insights and theory of economic analysis to the wealth of data on dairy animal performance and genetics and its associated modeling approaches. Returning to the basic structure of the HMM, we focus on estimating the effects of the genetic indices, PTAs, for production traits in dairy cattle and whether estimation suffers from selection bias. If PTAs are affected by selection behavior, this indicates that part of dairy farm productivity usually attributed to genetic progress should also be attributed to farmer skill at matching genetics to their environments.

In the next section, we provide a theoretical framework for thinking about heterogeneous returns to dairy genetics and how their effect on productivity can be investigated using Wooldridge’s (2015) CRC model.

2.3 Theory and Methodology

To begin, consider the case where choosing genetics is equivalent to choosing to increase or decrease a single trait by purchasing a sire with a particular PTA value. Every sire can be described as a vector of PTA index values for various traits, and in this case, we can think of genetic selection as the choice of a vector of index values. In reality, choosing a sire is a discrete decision, as the farmer faces some choice set of sires from various AI companies. We assume in what follows that the space of PTAs is “dense enough” that a farmer can choose any level of the trait they want from their choice set independently of other traits. We further assume this decision is only based on the trait itself and not on the sire’s identity or on the AI company that is offering it for sale.

Studying adoption via a continuous variable is preferable to the discrete approach in this case because it is not known which sires are in the farmer’s choice set. There are more than 10,000 unique sires in our data, and many more are actually available to farmers. Future analysis of the space of sires may be able to find a reduced dimensionality representation suitable for discrete choice analysis. As a first attempt, we study only the adoption of the traits via the PTA index to be able to apply a wider range of econometric tools.

2.3.1 Theoretical Framing

The following simple model demonstrates the role of farm- and animal-level heterogeneity in estimating the average returns to genetic investments via a continuous index. Unlike other input decisions, the decision to invest in genetics by choosing a certain sire happens three years before the animal starts producing. Assume that there is only one trait, z , which the producer has to choose three years before the cow begins production to maximize ex ante expected return:

$$\max_z \bar{\pi}(z, x, v) - wz,$$

where x and v are observable and unobservable management at the farm level, $\bar{\pi}$ is expected lifetime profit, and w is the price of purchasing one more unit of a trait.

In this model, the choice of z is only affected by *farm-level* heterogeneity, v and x . This is the level at which heterogeneity is usually analyzed based on the notion that management decisions operate at the level of an entire farm (Mundlak 1961; Suri 2011). What the above does not consider is that the characteristics of the mate—that is, the animal that is bred with the sire—should also affect returns to z .

Call these unobserved animal-level characteristics u . We can modify the above model only slightly to show why these characteristics are important.

Instead of z affecting $\bar{\pi}$ directly, it operates indirectly through a transmission function $f(z, u)$, which takes the traits of the father, or “sire” (z), and the traits of the mother, or “dam” (u), and maps to a new trait value, z' :

$$\begin{aligned} \max_z \bar{\pi}(z', x, v) - wz \\ \text{s.t. } z' = f(z, u). \end{aligned}$$

Now the optimal choice of z depends on the current period price that is to be paid versus the expected increase in profit weighted by how well the trait transmits. Adding this transmission function implies that unobserved heterogeneity affecting the adoption of z operates at the farm and animal levels. This is an important distinction and departure from the assumptions of both economic and animal science models of the returns to adoption. Economic models of the effect of technology adoption consider heterogeneity at the firm or farm level due to the assumption that confounding variation is from management behavior affecting all plots and animals. Animal science models refer to confounding variation at the animal level as “preferential treatment” and generally only control for farm-level effects because the literature does not find substantive evidence of animal-level decision-making that would bias evaluations (Graham, Smith, and Gibson 1991; Tierney and Schaeffer 1994).

Despite the lack of attention in the literature, animal-level heterogeneity can play its own part in biasing evaluations. If the manager observes components of u that are unobserved in data, then he or she may invest z with animals where the return is highest. Our next step is to investigate how the existence of unobserved u will affect our empirical evaluations of the returns to z in a production function.

2.3.2 Empirical Model

In data, we observe farms j and cows i during time period t . The PTA value of the sire chosen for an animal, z_{ij} , is time invariant. Using the above framework, we assume that there are animal-level (u_{ij}) and farm-level (v_j) “match quality” components that affect return to adoption of z_{ij} . Assume further that the total return is linearly separable, such that the total return is . Define this payoff as relative to some average expected return $\bar{\mu}$ so that u_{ij} and v_j are the farm’s known deviation from this average, $\tilde{\mu}_{ij} = u_{ij} + v_j$. The returns to z_{ij} for a given animal i and farm j are thus $\mu_{ij} = \bar{\mu} + \tilde{\mu}_{ij}$.

Heterogeneity in the production function manifests in the coefficients for z_{ij} :

$$y_{ijt} = \mu_{ij} z_{ij} + \beta X_{ijt} + \varepsilon_{ijt}.$$

Assuming a constant slope to identify $\bar{\mu}$ has the following effect on the equation:

$$\begin{aligned}
y_{ijt} &= (\bar{\mu} + \tilde{\mu}_{ij})z_{ij} + \beta X_{ijt} + \varepsilon_{ijt} \\
&= \bar{\mu}z_{ij} + \beta X_{ijt} + (\tilde{\mu}_{ij}z_{ij} + \varepsilon_{ijt}) \\
&= \bar{\mu}z_{ij} + \beta X_{ijt} + \xi_{ijt}.
\end{aligned}$$

Because $\tilde{\mu}_{ij}$ is unobserved match quality, then ordinary least squares will not identify an unbiased $\bar{\mu}$. It is biased because the variable return to a trait $\tilde{\mu}_{ij}z_{ij}$ is in the error term so that $E(z_{ij}\xi_{ijt}) \neq 0$. Instrumental variables will also not identify $\bar{\mu}$ because anything correlated with z must be correlated with ξ (Cornelissen et al. 2016).

Our identification strategy in this chapter uses instead the control function method and its specific approach to random coefficients, the CRC model. With this approach, we approximate input demand with a linear function of observed covariates plus an excluded variable and use the residual from the approximation to proxy for match quality in the production function. Wooldridge (2015) spells out two main conditions for the control function method to identify $\bar{\mu}$ and uncover heterogeneity in the effect of a trait. Defining η_{ij} as the residual term from a linear approximation of trait demand for z_{ij} , the two conditions for the CRC model are as follows:

$$\text{A1: } E(\varepsilon_{ijt} | \eta_{ij}) = \rho \eta_{ij}$$

$$\text{A2: } E(\tilde{\mu}_{ij} | \eta_{ij}) = \psi \eta_{ij}$$

These are both strict assumptions about how informative the residual η_{ij} is in capturing bias and match quality. A1 is a standard assumption for control function methods and says that selection bias takes a particular form: the conditional expectation of unobserved components of output is linear in η_{ij} . A2 says that the heterogeneous slope coefficient defined across cows must be proportional to the input demand residual η_{ij} . The unobserved components to technology adoption must include $\tilde{\mu}_{ij}$ if the manager considers their ex ante returns when choosing z_{ij} , but this assumption restricts their relationship to be proportional. Using A2, we can use η_{ij} interacted with z_{ij} to proxy for an estimate of $\tilde{\mu}_{ij}$. Either of these assumptions can be relaxed to be nonlinear, but explicit functional forms must be given so that we know how to include η in the production function. In our analysis, we maintain the linear forms.

We also need an exogenous shifter of z_{ij} that is uncorrelated with ε_{ijt} . Our instrument is the difference between the sire's PTA at the time it was chosen and its PTA value at its next evaluation four months after the adoption date (t'): $\Delta z_{ij} = z'_{ij} - z_{ij}$.¹ PTAs for every sire are updated by the Council on Dairy Cattle Breeding (CDCB) every four months using herd testing data from

1. Special thanks to our discussant, Paul Scott, for this suggestion.

around the country. The change in PTA from one evaluation to the next Δz_{ij} is linearly related to z_{ij} and so is a relevant predictor, but the size of the deviation has to do with the performance of the sire's daughters all across the country. This deviation is likely unrelated to unobserved production ε_{ijt} because it is based on the performance of other offspring of the sire before ε_{ijt} is ever realized. It is also unlikely that the updates happening right after the use of the genetics will somehow influence future management of that offspring; if this were the case, the PTA value of that sire at the time the offspring ij starts producing would be the more actionable information rather than the intermediate updates. For these reasons, we believe Δz_{ij} satisfies the exclusion restriction needed for an instrument. Our approach shares similarities to the control function approaches of Levinsohn and Petrin (2003) and Olley and Pakes (1996), which also use dynamic input lags as an exogenous source of variation to identify production function parameters.

Using A1 and A2 and the instrument Δz_{ij} , we can adjust the production function for the bias resulting from heterogeneous returns. Defining $\hat{\eta}_{ijt}$ as the estimated residual from the first-stage input demand, we now write our empirical model as two equations:

$$z_{ij} = \alpha_0 + \gamma \Delta z_{ij} + \beta_0 X_{ijt} + \eta_{ijt}$$

$$y_{ijt} = \alpha_1 + \bar{\mu} z_{ij} + \rho \hat{\eta}_{ijt} + \psi \hat{\eta}_{ijt} \times z_{ij} + \beta_1 X_{ijt} + \varepsilon_{ijt}$$

- y_{ijt} : dairy cow performance for butterfat/protein in a given lactation
- z_{ij} : value of PTA butterfat/protein of the sire chosen at the time of adoption
- Δz_{ij} : deviation in trait value in the next updated evaluation
- X_{ijt} : time-varying management decisions affecting y (a full list can be found in appendix A)
- η_{ijt} : input demand residual

Our main research question has to do with the parameters $\bar{\mu}$, ρ , and ψ . The hypothesis of “perfect transmission” of a trait is that $\bar{\mu} = 1$, so a one-unit increase in PTA causes a one-unit increase in the offspring's performance (Kearney et al. 2004). How much this parameter differs from one before and after our bias correction indicates whether unobserved management decisions affect the average return to a trait in our sample. If ρ is statistically different than zero, this rejects the null hypothesis that the models with and without the correction are equivalent (Wooldridge 2015).

Finally, ψ indicates the relationship between match quality and returns to z . If $\psi > 0$, then cows matched with a sire that has higher-than-expected PTA will also have a higher marginal return to PTA in their production function. This is consistent with the “positive assortative matching” story, which is that farmers adopt traits that work particularly well on their farms.

2.3.3 Heterogeneity Distribution

An output of the above model is an estimate of μ_{ij} , $\hat{\mu}_{ij} = \bar{\mu} + \psi\hat{\eta}_{ij}$.² Assuming our theoretical framework from before, this estimate contains both μ_{ij} and v_j . The farm-specific component v_j has been the focus of most studies in economics and is controlled for in animal science using fixed effects (termed “contemporary groups” in the animal science literature). However, we may also be interested in how much of the distribution in returns is driven by the animal-specific component u_{ij} . If there are heterogeneous returns at the animal level, then this means that the returns to the adoption of genetics are diverse even within a given farm environment. It also implies that sire evaluation models using farm fixed effects do not completely control for confounding factors and that there is evidence of managers matching specific genetics to specific animals, which could bias estimates of the return to genetics.

After estimating the parameters $\bar{\mu}$, ρ , and ψ , we estimate the distribution $\hat{\mu}_{ij}$ using the CRC model with three different specifications: no fixed effects, herd fixed effects, and herd-by-time fixed effects. The first specification estimates a distribution that contains both u_{ij} and v_j , and the second nets out v_j . The third specification mimics the fixed effects strategy of many genetic evaluation models, which use a herd fixed effect interacted with the time of the observation to soak up dynamic management decisions affecting the returns to genetics.

2.4 Data

As described above, the market for dairy sires makes heavy use of CDCB evaluations, which are calculated from DHI data. In addition to milk yield and somatic cell count, the DHI program tracks the number of times per day each cow is milked (usually two, sometimes three), their calving and birth dates, and their “lactation number” (the number of lactation cycles a cow has been through at a given point in time). Unfortunately, no other management decisions are observed. Our current data set covers DHI herds served by one dairy records processing center in the state of Wisconsin from June 2011 to January 2015, which is representative of about 40 percent of Wisconsin dairy herds. At the lactation level, there about 1 million lactation records for approximately 277,000 dairy cows on 1,500 dairy farms.

Because of the lack of management decisions observed in DHI data, the HMM includes in X a number of fixed effects to attempt to control for the confounding impact of management on genetics. In this model, we control

2. Note that this is in contrast to HMM, which would assume a normal distribution for such an effect and center it at zero. We gain flexibility with the distribution of the coefficient only because we specify exactly what determines the distribution, which is the unobserved variation in input demand.

Table 2.1 Records description

Herds	1,459
Sires	7,628
Dairy cows	277,695
Number of lactations	424,910
Lactation records	1,065,308

for lactation length, lactation number, and proportion of lactation milked three times a day in our specification. This is also a “cohort” effect, which is an interaction between herd and test month, which is a herd-specific time trend. There are also biological factors such as birth year, calving month, and breed that are included as fixed effects. In addition to these controls, we also include prices such as the milk price, the ration cost, and the price of replacement heifers. In our main specification, we use only herd fixed effects but also estimate the model with herd-by-time effects as a robustness check when looking at relevant distributions.

Every cow that shows up in DHI data has an ID that connects back to a sire and associated evaluation available from the CDCB. The CDCB updates evaluations three times a year, and these PTA values are the ones that will appear to the farmer when choosing genetics. Sire evaluations are publicly available on the CDCB’s website and are reported by AI companies when selling sires. These evaluations are updated four times a year. Using the sire IDs in our data set, we recovered all available records of these sires throughout time and matched them to cow records. Thus for each cow in our data, we know the PTA value of its parent sire at the time the choice of sire was made. The “time they were chosen” is calculated as 10 months prior to the cow’s birth date to account for the gestation period of a dairy cow. Our data set contains more than 7,000 unique dairy sires matched to our 1 million lactation records.

We use the following price covariates in our model. For output and input price, we use “income over feed cost,” a relationship between milk price and ration cost determined by the 2006 farm bill. We also include the price of 16 percent dairy ration as an additional proxy for feed cost. Finally, we include the cost of replacement, which we calculate as the beef price per pound times 1,400, the typical weight of a dairy cow, minus the cost of a replacement heifer.

In addition to issues discussed thus far, the analysis of dairy cow lactation records is complicated by survival bias for cows on their second lactation onward. Managers may remove cows from their herd if they do not meet some threshold of production during the first lactation. This selection issue is discussed in detail in Henderson (1975), and often lactation records of cows past lactation one are not used in sire evaluations for this reason. Keeping this in mind, we implement practices commonly followed in dairy

Table 2.2 Covariate description

	Mean	SD
Continuous variables		
PTA fat	28.79	27.12
PTA protein	21.46	20.48
Proportion milked 3×	0.58	0.49
Lactation length	310.44	23.48
Herd size	157.35	232.99
Binary variables (%)		
Lactation number		
1st	45.73	
2nd	28.71	
3rd	15.35	
4th	7.34	
5th	2.87	

science literature when analyzing lactation data. We do not consider cows that are lactation six or higher (about .1 percent of the data), and we analyze “primiparous” (first lactation) cows separately from “multiparous” cows. Primiparous cows should not be subject to survival bias, while multiparous cows are a subset of the first group that was not culled. It should be the case that multiparous cows are more subject to the management bias we discuss, and we analyze this group separately to see how our bias correction works differently in this subset. If bias is severe in multiparous cows, this suggests an interaction between the behavior affecting genetic selection and the behavior affecting culling decisions.

Using our matched data, we graph the kernel densities of PTA values for butterfat and protein chosen in this sample with dotted lines indicating their average. Recall that the HMM used to produce PTA measures fixes the distribution to be normal with a mean of zero for the relevant population. The densities are not centered at zero and are not symmetric; both densities have very long left tails, which shift averages to the left.

This does not give any indication of at what level the selection is occurring, however. For example, do farmers simply choose the same PTA value for all their animals in a month? At what level is there variation in the chosen traits? A quick calculation of within-group sum of squares can shed light on how variable each selected trait is within a given herd versus between herds. For example, if a farmer simply chooses the same trait value for all his or her cows, then the sum of squares within herds should be zero. The farmer may choose the same trait in all time periods or, within a certain month, choose the same value for all cows. The difference between herd-month and herd essentially approximates the importance of time-variant factors, such as prices and other economic factors. We calculate the ratio between the within




Fig. 2.1 Distributions of PTAs in the data

Table 2.3 Proportion of SST explained

	PTA fat	PTA protein
Herd	0.954920	0.955163
Herd by time	0.811062	0.797823

sum of squares for these two groups (herd-adoption month and herd) and the total sum of squares.³

The proportion of the sum of squares total explained by within-herd variation in choices of both traits is quite large: about 80 percent for herd-adoption month and around 95 percent for herd. This is evidence that the largest amount of variation in trait choices is within a herd and not between herds, or a large variation at the cow level. This is not consistent with a model where heterogeneity in selection behavior is driven at the herd level. Given this indicative evidence, we proceed to our empirical model to explore the impact of this heterogeneity in selection on the average return to high-yield genetics.

2.5 Results

We study the traits protein and butterfat, which are the components of milk that are most important to profitability for dairy farmers in Wisconsin. The PTA index is in units of pounds of fat and protein and represents the

3. Calculated as $\sum_{g=1}^G \sum_{i=1}^I (y_{ig} - \bar{y}_g)^2 / \sum_{g=1}^G \sum_{i=1}^I (y_{ig} - \bar{y})^2$, where G is either herd groups or herd adoption month group.

expected increase in yield of a dairy cow using the given sire relative to a base sire (whose PTA is zero). Specifically, they are predictions of a statistical model, the HMM, which are interpreted as the increase in fat or protein for the specific sire that is chosen. Since the outcome y_{ijt} is measured in the same units as PTA, if $\bar{\mu} = 1$, then increases in sire ability correspond one-to-one with increases in the offspring's ability.

For each trait, we estimate several different specifications to examine how the coefficient on z changes with different corrections. We estimate ordinary least squares (OLS), OLS with herd fixed effects, two-stage least squares, two-stage least squares with herd fixed effects, and correlated random coefficients (i.e., including $\hat{\eta}_{ijt}$ and $\hat{\eta}_{ijt} \times z_{ij}$ as regressors in a fixed effects regression). If heterogeneity is only at the herd level, then according to Wooldridge (2005), herd fixed effects alone should identify the average treatment effect. Including the control function terms ($\hat{\eta}_{ijt}$ and $\hat{\eta}_{ijt} \times z_{ij}$) in the fixed effects model identifies the heterogeneity within herds specifically, so the difference between these specifications provides evidence regarding the importance of cow heterogeneity in determining the average effect.

Finally, noting that the marginal benefit of a trait is given by

$$\hat{\mu}_{ijt} = \bar{\mu} + \hat{\psi} \hat{\eta}_{ijt},$$

we can graph the resulting distribution to examine the variability of returns across the entire sample. In each specification, we analyze three samples based on the lactation year: all lactations, first lactations, and later lactations. The first-lactation cows are studied separately because they are not subject to survival bias, as later-lactation cows possibly are. Estimates of the first stage of the model from which the input demand residual is calculated are presented in appendix B. Standard errors are calculated clustered at the herd level and cluster bootstrapped for the CRC model.

2.5.1 Fat

For both OLS and fixed effects, the average return to increasing the butterfat of a sire is positive and different than 0. It is around 0.6, meaning a one-unit increase in a pound of PTA causes a 0.6-pound increase in offspring. The correction, however, attenuates the effect toward 0 by a large amount. When using instrumental variables and a constant coefficient on z_{ij} , the coefficient is near 0, implying that the correction takes away most of the productivity gain that would otherwise be (mis)attributed to the choice of PTA. The average effect identified in the CRC model is higher, about 0.14, with a positive and significant ρ (meaning there was significant selection bias in the OLS specification). The CRC specification also tells us that cows with a higher than predicted amount of the trait have a higher marginal return to the trait—that is, $\psi > 0$. At all levels, we reject the hypothesis that $\rho = \psi = 0$, which indicates that the model with the correction is statistically different than the model without it. This suggests that the instrument was necessary

Table 2.4 All lactations

	OLS (1)	FE (2)	IV (3)	IV + FE (4)	CRC + FE (5)
PTA fat	0.604*** (0.0321)	0.544*** (0.0134)	0.0325** (0.00723)	0.0355*** (0.00434)	0.149*** (0.0083)
$\hat{\eta}$					0.563*** (0.0236)
$\hat{\eta} \times$ PTA fat					0.0066***
N	1,065,308	1,065,308	1,065,308	1,065,308	1,065,308
Adj. R^2	0.351	0.562	0.345	0.557	0.564

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.5 Across lactation cows

	All (1)	First lactation (2)	Later lactation (3)
PTA fat	0.14933*** (0.00827)	0.10593*** (0.00829)	0.19799*** (0.01303)
$\hat{\eta}$	0.56271*** (0.02364)	0.53914*** (0.02447)	0.58808*** (0.03091)
$\hat{\eta} \times$ PTA fat	0.00662*** (0.00037)	0.00602*** (0.00042)	0.00724*** (0.00044)
N	1,065,308	511,446	553,859
Adj. R^2	0.564	0.537	0.514

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

to correct the estimates for endogeneity. These results imply that the average effect of investing in high-fat genetics is reduced by 75 percent after taking into account confounding factors. More than half of the return to PTA fat is explained by unobserved confounding variables.

Table 2.5 shows a similar pattern to the whole sample and further indicates large differences in first- and later-lactation cows for average returns. There is strong evidence of selection behavior that affects the returns to genetics, as the average return for first-lactation cows is half that of later-lactation cows. If this difference is generated by culling, it indicates that farmers cull cows in their first lactation that have low marginal return to the high-yield genetics.

2.5.2 Protein

Similar to fat, the average returns to protein are much lower when accounting for confounding factors. Using simple OLS, the return to protein is 0.427 and indistinguishable from 0 when using instrumental variables. Using the CRC model, the effect is different than 0 but is less than half of the OLS

Table 2.6 All lactations

	OLS (1)	FE (2)	IV (3)	IV + FE (4)	CRC + FE (5)
PTA protein	0.427*** (0.0343)	0.358*** (0.0109)	0.00883 (0.00603)	0.0165*** (0.00381)	0.21724 *** (0.011)
$\hat{\eta}_{ijt}$					0.25252*** (0.0102)
$\hat{\eta}_{ijt} \times$ PTA protein					0.00983***
N	1,065,308	1,065,308	1,065,308	1,065,308	1,065,308
Adj. R^2	0.451	0.669	0.448	0.667	0.671

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.7 Across lactation cows

	All (1)	First lactation (2)	Later lactation (3)
PTA protein	0.21724*** (0.011)	0.19005*** (0.01104)	0.23834*** (0.0134)
$\hat{\eta}$	0.25252*** (0.02248)	0.2472*** (0.02271)	0.26652*** (0.02875)
$\hat{\eta} \times$ PTA protein	0.00983*** (0.00041)	0.00997*** (0.00042)	0.00972*** (0.0005)
N	1,065,308	511,446	553,859
Adj. R^2	0.671	0.632	0.622

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

coefficient when controlling for confounding factors: the estimate changes from 0.43 to 0.18. The direction of ψ suggests positive assortative matching for the adoption of high-protein genetics, just as for high-fat genetics. When looking at different lactations, there is less evidence of culling based on returns to protein. There are slightly lower returns for first-lactation cows than later-lactation cows, but it is a much smaller difference compared to the differences for fat.

2.5.3 Distributions

Here we estimate the resulting distributions from the CRC specification, $\hat{\mu}_{ijt} = \bar{\mu} + \hat{\psi}\hat{\eta}_{ijt}$. Figure 2.2 shows the distribution of returns across all lactations for fat and protein. Figure 2.3 shows the distributions of both traits across different lactations and across different levels of fixed effects. In addition to using herd fixed effects, we also include a combination of three fixed effects for herd, test month, and calving month. The product of these three indicator variables is typically referred to as a “contemporary group effect”

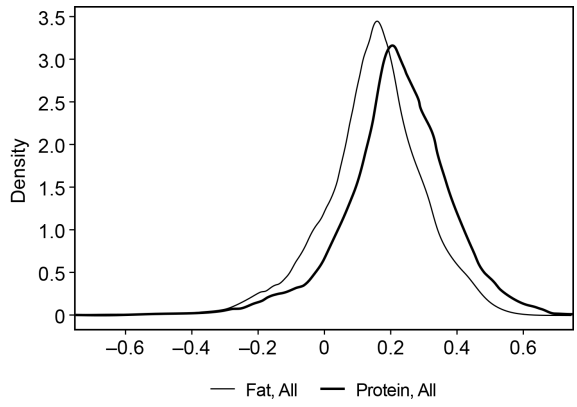



Fig. 2.2 Distributions of returns

A. PTA Fat



B. PTA Protein




Fig. 2.3 Fat and protein distributions

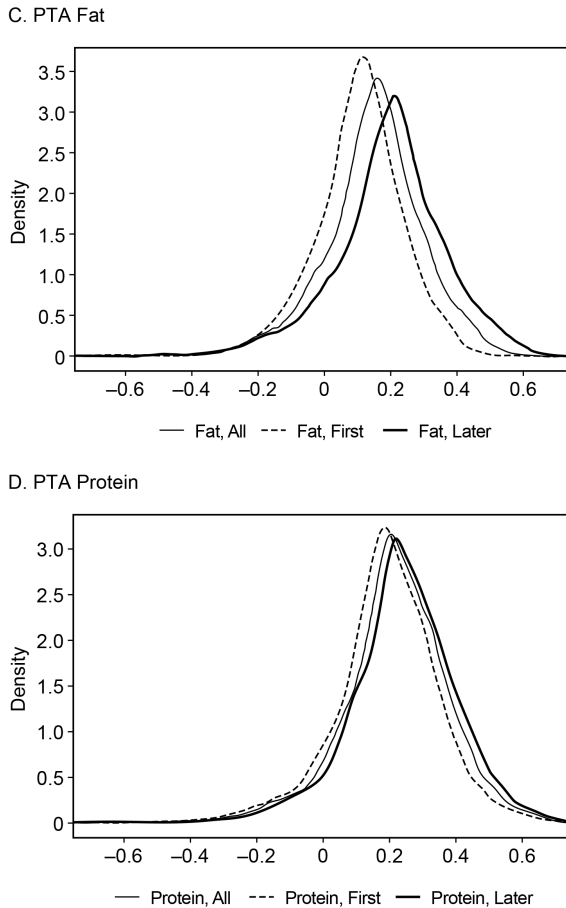


Fig. 2.3 (cont.)

in a genetic evaluation model. Table 2.8 shows the means and standard deviations of all of the distributions.

While different in their average effect, both traits have about the same standard deviation. For fat, a part of the distribution actually has a negative coefficient for adoption. For that farm or animal, the returns from adopting the technology may in fact be negative because of a combination of management environment and unobserved animal-level factors. In this case, however, it is hard to justify why increasing a trait in one parent would actually decrease that same trait in the offspring. The reason for this may be that high-fat or high-protein genetics are correlated with another trait that may negatively affect milk production in certain environments. We made

Table 2.8 Distributions of marginal returns

	Fat		Protein	
	Mean	Std. dev	Mean	Std. dev
All lactations	0.150	0.151	0.218	0.164
First lactations	0.106	0.133	0.191	0.161
Later lactations	0.199	0.168	0.239	0.166
No fixed effects	0.170	0.185	0.249	0.197
Herd fixed effects	0.150	0.151	0.218	0.164
Herd-by-time fixed effects	0.132	0.146	0.187	0.158

the assumption that traits could be chosen independent of one another, but in reality, traits have genetic correlations. For example, high milk yield and health are negatively correlated, so increasing production traits could negatively impact health, causing a decrease in phenotypic yield. A more advanced model of adoption would need to find a way to model their adoption together and explicitly include these correlations as part of the choice problem.

Figure 2.3 shows the differences in returns across lactations and across different fixed effects specifications. The distributions across lactations support the results in tables 2.5 and 2.7. The goal of using different fixed effects specifications was to look at the effects of parsing out herd-level unobserved factors v_j versus animal-level unobserved factors u_{ij} from the distribution $\hat{\mu}_{ij}$. Without any fixed effects, the effect sizes are 0.17 and 0.25 for fat and protein, respectively. After netting out time-invariant herd effects, the effect sizes drop to 0.15 and 0.218. Finally, when netting out *all* herd-level variation using herd-by-time effects, the effect sizes drop to 0.132 and 0.187. The resulting distributions are entirely generated by variation in input demand at the animal level, meaning dairy farm managers observe animal-level returns that they use to choose genetics. These results imply that, on average, these returns make up 77 percent and 75 percent of the variable returns to investment in high-yield genetics. A sizable portion of the selection process, therefore, appears to happen at the level of animals rather than at the level of farms.

2.6 Discussion and Conclusion

We examine the effect of economic selection behavior on the returns to the adoption of genetic technology for dairy farms in Wisconsin. Previous literature has attributed a large amount of productivity growth on dairy farms to improvements in genetics without considering the possibility that traits are selected into environments because of unobserved (to the researcher)

farm- or animal-specific returns to a given trait. Using the theoretical framework of the Roy model, we argue that farm- and animal-level heterogeneity may bias estimates in the returns to genetic traits such as butterfat and protein through the selection process. We use herd testing association data on dairy cows and the evaluations of their sires at the time they were chosen to estimate a CRC model. Our approach permits us to examine the effects of correcting for bias when estimating the impact of genetic improvement on productivity, the effect of selection behavior as a source of heterogeneity in returns to traits, and the relative importance of cow-level heterogeneity versus farm-level heterogeneity.

We find that correcting for selection bias lowers the estimated contribution of genetic improvement to productivity differences across cows by an average of 50 percent. We estimate average returns to adopting genetics with 1 more pound of fat or protein to be about 0.6 pounds and 0.4 pounds before the correction and 0.15 pounds and 0.2 pounds after correction for selection bias. Our model also indicates positive assortative matching, meaning farms with the highest return to adopting a given set of traits are the ones that adopt, giving credence to the upward bias in the coefficients. We also find that first-lactation cows have the lowest average return to high-yield genetics, indicating that farmers tend to cull cows with lower-than-expected ex post marginal return to the traits. Analyzing the distributions of returns, we find that up to 75 percent of the heterogeneity in returns generated from the input demand residual were at the animal level instead of the farm level. This implies that the factors confounding the returns to genetics are also at the animal level; cows with different trait investments are managed differently in ways that are not controlled for using farm-level fixed effects.

This study has several limitations that should be addressed in future work. First, we model trait adoption as though each trait could be chosen independent of other traits. This may not be a reasonable assumption given the extent to which traits are correlated with one another. Accounting for this possibility would require using a system of equations with cross-equation restrictions limiting trait-selection possibilities to those that are feasible given relevant empirical context. Such a model might be able to explain why there are negative returns to high-yield genetics for some animals and farms. Second, we treat trait investment as continuous even though farms choose sires discretely. This assumes that the trait values are dense enough to treat the variable as continuous, whereas the adoption decision is discrete over a choice set of individual bulls. One way to model this as a discrete problem while also taking into account correlations between traits would be to use a lower dimensionality representation of sires determined from the data. Unsupervised machine learning methods such as K-means clustering could be used to characterize an implied grouping of sires that have certain traits in common. The problem of choosing traits would then become one of

choosing a basket of traits represented by a certain grouping of bulls. The matching decision itself also needs further elaboration. In future work, we hope to turn our attention to data on breeding decisions where we have additional data that may permit a more detailed investigation of the selection decision. We also aim to develop a more sophisticated modeling approach that will take into account that farmers select a portfolio of traits rather than choosing one at a time.

Despite these shortcomings, our results point to new possibilities for studying technology adoption and suggest the need for a reinterpretation of and further research on the expansive literature that examines the contribution of genetic progress to productivity growth in the dairy sector. For the economics field tackling technology adoption, animal-level heterogeneity is important and should not be overlooked. Appendix B contains the first-stage results of the model, which show that cows with higher trait investment are milked more frequently and survive to more lactations. If such behavior happens at the animal level, it is important to take this into account when thinking of sources of farm productivity that farmers may act on in the context of the Roy model. Previous studies of farm productivity usually identify “unobserved” returns at the farm level, and for this reason, many papers studying dairy farms or animal operations sum production to the herd level. This assumption also suggests a reevaluation of extension programming developed to advise farmers about herd-level management. Our work shows that a large amount of the heterogeneity in returns is driven by animal-level variation, meaning there are ample opportunities to increase productivity by emphasizing management on this level. Agricultural data are becoming more granular, and there is no doubt there will be increasing opportunities for economics research to take selection of genetics by farmers into account. We consider only animal agriculture, where every animal must be bred, but the approach we develop here may also be used at some scale in crop agriculture.

Overall, we find that selection behavior biases estimates of the effect that genetic improvement alone has on productivity growth. An important component of productivity change depends on farmers choosing genetics that work particularly well in conditions that are idiosyncratic to their individual farming operations. This changes the narrative regarding the source of farm productivity in the dairy industry from one where science alone is the source of gains from new technology to one where growth is the result of complementary inputs provided by farmers and scientists. Indeed, the success of the dairy industry thus far depends on collaboration among farmers and scientists via institutions often taken for granted, such as the DHI program, land grant universities, and a variety of industry collaborators (represented collectively by the CDCB). The interplay among these organizations and the remarkable record of success (as measured by productivity growth)

they have achieved make the dairy industry a unique model of research and innovation in agriculture that merits further analysis and critique in the economics field.

Appendix A

Regression Controls

To select controls for the animal equation, we draw on the animal science literature to inform controls we include in the model.

The vector X_{ijt} contains the following variables:

- Economic Controls
 - cost of 16 percent dairy ration
 - income over feed cost
 - replacement cost (beef price \$/lb \times 1,400 – cost of replacement heifer)
 - time trend
- Biological Controls
 - calving month (indicator)
 - test month (indicator)
 - birth year (indicator)
 - lactation number (indicator)
 - Holstein (indicator)
- Management Controls
 - proportion of lactation milked three times a day
 - herd size (deviations from average)
 - lactation length (days in milk of record)

Appendix B

First-Stage Estimates

The first-stage equation for our model uses past variation in a sire's evaluation, which occurs at the national level as a source of exogenous variation:

$$z_{ij} = \alpha_{0j} + \gamma\Delta z_{ij} + \beta_0 X_{ijt} + \eta_{ijt}.$$

While this prediction is time invariant (the selection occurs only once), the residual η_{ijt} will still be time variant because of the term X_{ijt} . Due to the presence of X_{ijt} , the first stage essentially treats the same cow at different points in time as entirely separate cows who happen to have the same values of z_{ij} . This means that when we examine the PTA investment for one cow at two

Table 2.9 First stage regression

	PTA fat		PTA protein	
	OLS	FE	OLS	FE
Δz_{ij}	0.499*** (0.000648)	0.500*** (0.000613)	0.495*** (0.000550)	0.496*** (0.000514)
Lactation no. = 2	1.589*** (0.265)	0.368* (0.207)	1.793*** (0.193)	0.954*** (0.157)
Lactation no. = 3	2.220*** (0.519)	-0.195 (0.409)	2.789*** (0.376)	1.132*** (0.310)
Lactation no. = 4	3.164*** (0.766)	-0.393 (0.594)	3.675*** (0.569)	1.254*** (0.461)
Lactation no. = 5	3.839*** (0.999)	-0.787 (0.786)	4.718*** (0.746)	1.610*** (0.605)
Proportion milked 3×	1.528*** (0.481)	-0.310 (0.687)	1.210*** (0.351)	-0.0692 (0.382)
Herd size	0.0000154 (0.000258)	-0.000717 (0.000572)	-0.0000608 (0.000198)	-0.000657 (0.000519)
Lactation length (days)	0.00640*** (0.00108)	0.00561*** (0.000913)	0.00559*** (0.000781)	0.00522*** (0.000678)
Holstein	1.475 (1.475)	3.191*** (0.561)	3.260*** (0.791)	2.707*** (0.656)
Observations	1,641,022	1,641,022	1,641,022	1,641,022
Adjusted R^2	0.249	0.303	0.281	0.333

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

different lactations, it essentially treats these as two adoption decisions; by deciding to let the animal keep producing, the manager implicitly adopts the genetics again. This is not necessarily problematic, but we must estimate the model on these surviving animals separately to understand how the culling decision interacts with the adoption of genetics.

One implication of this approach is that the first stage will help us understand the trait investments for animals that survive. Table 2.9 shows the results of the first stage and the coefficients on animal-level variables. Both OLS and fixed effects are shown to get a sense of what level of variation is important. For example, both lactation length and lactation number are significant in predicting z , which implies that cows that have a larger trait investment are milked longer and are more likely to not be culled in their first year. Milking the cow three times per day is significant in the OLS specification but not in the fixed effects, implying that farms that choose higher investments in production traits also milk more intensively at the herd level. Holstein cows are also most likely to have the highest investment in production traits, which is to be expected given their comparative advantage in high-volume production.

Differences across production traits are mostly seen in the culling deci-

sion. Without herd fixed effects, cows that are kept past the first lactation have higher trait investment for both fat and protein. Once herd fixed effects are used, fewer differences are seen across lactations considering only animal-level variation. For fat, only second-lactation cows have marginally more fat investment than first-lactation cows. For protein, all later-lactation cows have higher investments in protein (on the order of one pound more). One thing that can be learned from these results is that adoption and other management decisions are inextricably linked. Specifically, cows that have a high PTA investment are more likely to be kept, milked longer, and milked more intensively.

References

- Cornelissen, T., C. Dustmann, A. Raute, and U. Schonberg. 2016. "From LATE to MTE: Alternative Methods for the Evaluation of Policy Interventions." *Labour Economics* 41:47–60.
- Council on Dairy Cattle Breeding (CDCB). 2017. *History of USDA Dairy Evaluations*. Accessed September 24, 2019. https://queries.usdccb.com/aip/history/hist_eval.htm.
- Foltz, J., U. Aldana, and P. Laris. 2014. "The Sahel's Silent Maize Revolution: Analyzing Maize Productivity in Mali at the Farm Level." In *African Successes*, vol. 4, *Sustainable Growth*, edited by Sebastian Edwards, Simon Johnson, and David N. Weil, 111–36. Chicago: University of Chicago Press.
- Gianola, D., and G. J. M. Rosa. 2015. "One Hundred Years of Statistical Developments in Animal Breeding." *Annual Review of Animal Biosciences* 3 (1): 19–56.
- Graham, N. J., C. Smith, and J. P. Gibson. 1991. "Investigation of Preferential Treatment for Milk Yield in Canadian Holsteins." *Canadian Journal of Animal Science* 71 (1): 21–27.
- Griliches, Z. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica, Journal of the Econometric Society* 25 (4): 501–22.
- Griliches, Z., and J. Mairesse. 1995. "Production Functions: The Search for Identification." NBER Working Paper No. 5067. Cambridge, MA: National Bureau of Economic Research.
- Heckman, J., and E. Vytlacil. 1998. "Instrumental Variables Methods for the Correlated Random Coefficient Model: Estimating the Average Rate of Return to Schooling When the Return Is Correlated with Schooling." *Journal of Human Resources* 33 (4): 974–87.
- Henderson, C. R. 1953. "Estimation of Variance and Covariance Components." *Biometrics* 9 (2): 226–52.
- . 1973. "Sire Evaluation and Genetic Trends." *Journal of Animal Science* 1973 (issue symposium): 10–41.
- . 1975. "Best Linear Unbiased Estimation and Prediction under a Selection Model." *Biometrics* 31 (2): 423–47.
- Kearney, J. F., M. M. Schutz, P. J. Boettcher, and K. A. Weigel. 2004. "Genotype \times Environment Interaction for Grazing Versus Confinement. I. Production Traits." *Journal of Dairy Science* 87 (2): 501–9.

- Levinsohn, J., and A. Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies* 70 (2): 317–41.
- Michler, J. D., E. Tjernström, S. Verkaart, and K. Mausch. 2019. "Money Matters: The Role of Yields and Profits in Agricultural Technology Adoption." *American Journal of Agricultural Economics* 101 (3): 710–31.
- Mundlak, Y. 1961. "Empirical Production Function Free of Management Bias." *Journal of Farm Economics* 43 (1): 44–56.
- Olley, G. S., and A. Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263.
- Olmstead, A. L., and P. W. Rhode. 2008. *Creating Abundance: Biological Innovation and American Agricultural Development*. Cambridge, MA: Cambridge University Press.
- Pryce, J., and R. Veerkamp. 2001. "The Incorporation of Fertility Indices in Genetic Improvement Programmes." *BSAP Occasional Publication* 26 (1): 237–49.
- Roy, A. D. 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers* 3 (2): 135–46.
- Suri, T. 2011. "Selection and Comparative Advantage in Technology Adoption." *Econometrica* 79 (1): 159–209.
- Tierney, J., and L. Schaeffer. 1994. "Inclusion of Semen Price of the Sire in an Animal Model to Account for Preferential Treatment." *Journal of Dairy Science* 77 (2): 576–82.
- Van Vleck, L. 1987. "Contemporary Groups for Genetic Evaluations." *Journal of Dairy Science* 70 (11): 2456–64.
- Wooldridge, J. M. 2005. "Fixed-Effects and Related Estimators for Correlated Random Coefficient and Treatment-Effect Panel Data Models." *Review of Economics and Statistics* 87 (2): 385–90.
- . 2015. "Control Function Methods in Applied Econometrics." *Journal of Human Resources* 50 (2): 420–45.
- Wright, P. G. 1928. *Tariff on Animal and Vegetable Oils*. New York: Macmillan.
- Wright, S. 1925. *Corn and Hog Correlations*. Department Bulletin No. 1300, US Department of Agriculture. <https://www.biodiversitylibrary.org/bibliography/108042#/summary>.
- Zeitlin, A., S. Caria, R. Dzene, P. Jansky, E. Opoku, and F. Teal. 2010. "Heterogeneous Returns and the Persistence of Agricultural Technology Adoption." CSAE Working Paper Series 2010–37, University of Oxford Centre for the Study of African Economies.