Technological Change and Productivity in the Rail Industry: A Bayesian Approach

Taylor K. McKenzie
University of Oregon
tmckenzi@uoregon.edu

November 7, 2016

Abstract

Productivity and its growth are central to the long-term growth, and long-term viability of firms and industries. Partial deregulation of railroads was led by concerns that existing regulation and changes to the industry led to stagnation in productivity. Policy changes made it easier for firms to increase productivity through broad organizational changes like mergers and abandoning unprofitable routes as well as specific technological innovation through the 1980s and early 1990s. However, as the industry has become increasingly consolidated and as more lines have been abandoned, firms may need to rely on technological change to increase productivity. I develop and estimate a model that separates changes in productivity due to innovation and those caused by non-innovative factors and use Bayesian estimation. This allows productivity and technology to evolve flexibly across firms and through time, allowing an examination of changes in railroad productivity and identification of its driving component. I find that every Class I railroad has experienced growth in productivity since 1999. Improvements in technology were the driving factor in the growth of BNSF, KCS, Soo Line, and UP, while CN, CSX, and NS saw significant growth due to broad organizational changes. Finally, I develop a metric that determines whether firms substitute inputs towards factors that innovation makes more productive. I estimate the probability that each firm takes that action to be around 50% with no discernible pattern over time, providing evidence that firms don't anticipate technological change or don't adjust input allocation to take advantage of innovations.

Keywords: Railroad, productivity, heterogeneous effects, time-varying effects, random coefficients, Bayesian estimation

1 Introduction

The productivity of firms is the amount of real output that can be produced with a marginal increase in real inputs. It has long been of interest to researchers, regulators, and industry analysts alike. Productivity allows firms to produce their products at lower cost and is also the source of long-term economic growth. Further, the level and growth of productivity informs regulators and is central in their regulatory mandate. Prior to its partial deregulation, the railroad industry was faced with many concerns of viability and low levels of productivity. While productivity growth was rapid through the mid-1900s, it had slowed dramatically by the mid-1970s due to the rise of competing modes of transportation and changes to the types of products being shipped. With the goal of reducing costs and increasing productivity, the industry was partially deregulated in 1980. The immediate effects of this policy have been studied extensively,[1] and it is clear that there has been rapid productivity growth since partial deregulation (Wilson, 1994).

The more recent effects of partial deregulation have not been examined. Immediately following partial deregulation, it was relatively easy for firms to merge (thereby taking advantage of economies of scale) and ceasing service on unprofitable routes (Bitzan and Wilson, 2007). The number of Class I railroads fell from 40 in 1980 to just 7 in 1999, and the total size of the network controlled by these carriers dropped from 164,822 miles in 1980 to 95,391 miles in 2013 (United States Surface Transportation Board, 2015). While these changes have dramatically reduced costs in the industry and improved its viability, there is relatively little room for to continue realizing productivity growth through these broad changes(Bitzan and

[1] For further reading, see Bitzan and Keeler (2007), Winston et al. (1990), and Barnekov and Kleit (1990).

Keeler, 2007). Instead, firms may need to improve their production technologies through innovation and substitute towards more productive inputs in order to realize continued growth, which is vital to the sustained viability of the industry.

In order to separately identify changes in productivity, I develop a model that incorporates inefficiency and also allows productivity and technology to vary across firms and evolve in a flexible way over time. Using a theoretical framework, I decompose changes in production into increased use of inputs, input substitution, increased productivity due to technological change, and increased productivity due to non-innovative factors. I then estimate my model using Bayesian methods. This allows me to identify and decompose productivity changes for each firm and each year. I am not aware of any published research that empirically separates growth in productivity due to innovation and that due to factors other than innovation. Productivity growth due to technological change becomes increasingly important as an industry matures and other methods of increasing productivity like merging with or acquiring other firms become less feasible. Further, firms can take potentially take greater advantage of innovations by substituting towards inputs that changing technology makes more productive. I identify a condition under which firms substitute towards more productive inputs and estimate the probability that firms take that action for each year. Finally, these models provide estimates of total factor productivity and its growth, which are key values for informing regulation and give insight into developments in the industry.

The models I estimate allow productivity growth to vary flexibly both across firms and over time by imparting structure on its dynamics; specifically, I allow productivity and technological parameters to follow random-walks with drift. Ignoring the effect of technological change, I find that the Canadian National (CN) railway showed the strongest productivity growth since 1999 at a rate of 3.551% per year. All other railroads exhibited modest productivity growth, between 0.235% and 2.474% per year. After accounting for technological innovation, I am able to identify how much of productivity growth is due to changing technology and how much is due to neutral shifts in the production technology. I find that CN, CSX, and Norfolk Southern (NS) railways experienced strong growth in productivity caused by factors other than technological growth; all other railroads showed decreasing productivity due to these non-innovative factors. Burlington Northern Santa Fe (BNSF), Kansas City Southern (KCS), Soo Line, and Union Pacific (UP) found significant increases in productivity due to technological change, with growth between 30% and 60% between 1999 and 2014. CN, CSX, and NS experienced smaller productivity gains due to changing technology. Overall, when considering total productivity due to all factors, CN and KCS have shown the strongest growth in productivity driven mostly by factors other than technological innovation, though both have seen decreases since 2011. BNSF showed modest total productivity growth due to technological change, and all other firms had constant total productivity. Finally, I find that firms have about a 50% chance of shifting resources towards inputs that innovation makes more productive. This provides evidence that firms don't anticipate technological changes, aren't able to substitute inputs fast enough to capitalize on innovations, or that input prices tend to offset changes in technology.

I estimate three different models with varying degrees of flexibility in the dynamics of productivity change and technological growth. The first model assumes the productivity of each firm follows a simple linear trend, the second allows productivity to follow a random-walk with drift while holding technology constant, and the third allows both productivity and technology to follow a random-walk with drift. Using Bayesian model selection, I find that the model allowing both productivity and technology to evolve flexibly over time has the greatest probability of being the true model, indicating the importance of controlling for technological change. Using estimated model probabilities, I conduct Bayesian model averaging of the results of each model and find that each firm likely experienced modest growth in productivity between 1999 and 2014, with median estimates ranging from 0.296% to 0.719% per annum. However, I estimate the probability that all firms experienced productivity growth is 46.783%, indicating that at least one firm likely saw a decrease in productivity over the sample period.

This paper begins with an overview of the railroad industry and its regulation. Following this, I provide a review of the relevant literature, covering both the methods used to measure productivity and how productivity has been studied in the railroad industry. I then develop my theoretical model and proceed to present the data used in this analysis. I cover each of the three empirical models presented in this paper, then show and explain my results. A conclusion of my findings follows.

2 Institutional Background

The railroad industry has been federally regulated since 1887. The Interstate Commerce Commission (ICC) was created in response to concerns of excessive rates, market power, and discriminatory pricing in the industry with the passage of the Interstate Commerce Act (ICA) of 1887. This policy gave the ICC control over collective rate making and oversight over mergers and provided a channel through which the reasonability of rates charged by railroads could easily been questioned by shippers. Through most of the 1900s, these regulations helped promote competition in the industry and kept shipping rates low while still allowing railroads to be profitable.

By the 1970s, the regulations that once promoted competition impeded firms in the industry. Not only had new competing modes of transportation such as air, barge, and trucking been developed and improved, but plastics, which are much less dense than goods shipped in the past, constituted a greater proportion of all goods shipped. Consequentially, railroad costs rose to the point that rate regulation prevented firms from being cost viable. In an effort to save the industry, railroads were partially deregulated with the passage of the 4R Act in 1976 and the Staggers Act in 1980. These policies allowed railroads to merge more easily to reduce costs through economies of scale, gave firms the ability to negotiate contracts and generally provided greater pricing flexibility, and allowed firms to more easily abandon lines on which operations were not profitable.[2]

Partial deregulation resulted in many drastic changes to the industry. The number of Class I railroads fell from 40 in 1980 to just 7 in 1999, mostly through acquisitions and mergers. The total size of the network controlled by Class I railroads fell from 164,822 miles in 1980 to 95,391 miles in 2013, largely through the abandonment and sale of unprofitable lines. The average length of haul increased from 615 miles in 1980 to 973 miles in 2013 (United States Surface Transportation Board, 2015). Overall, individual railroad networks were larger, the total size of the network grew smaller, and shipments were traveling longer distances. As a result of these changes, rail shipping rates fell dramatically, from $0.0646 per revenue-ton-mile in 1980 to $0.0329 in 2014 (United States Surface Transportation Board,

[2]There have been a considerable number of studies that describe these policies and their effects. See, for example, Bitzan and Wilson (2007), Transportation Research Board (2015), Wilson (1997), and Winston et al. (1990).

2015). The reduction in prices is largely reflective of a reduction in railroad costs and improvements in productivity (Bitzan and Keeler, 2007).

Following rapid changes that occurred in the railroad industry through the 1980s and early 1990s, the general structure of the industry has mostly remained constant since 1999. Only seven Class I railroads remained in 1999, and additional mergers have not occurred. Most of the lines on which operations were unprofitable have been abandoned or sold to short-line regional railways (Tretheway et al., 1997). As a result, there is little room for railroads to improve their productivity on those fronts. Thus, to remain viable, firms have turned towards other channels, such as technological progress, to realize productivity gains and further reduce costs. As an example, the elimination of cabooses, a remnant of the age of steam locomotives that required a crew to operate, resulted in a reduction in costs by between 5% and 8% between 1983 and 1997 (Bitzan and Keeler, 2003). Railroads have also invested $575 billion in infrastructure and equipment since 1980; recently, nearly 2700 new locomotives were purchased between 2008 and 2012, and many innovations have been made in safety, fault detection, and performing maintenance that preempts equipment failure (AAR, 2015). To my knowledge, there has been no published research that considers the effects of these recent innovations.

3 Literature Review

This research investigates productivity of the railroad industry using a stochastic frontier model. In this section, I provide a history and review of studies and methods used to estimate productivity in general. I then describe research that has investigated the productivity of railroads and the effects of the industry's partial deregulation. Finally, I describe stochastic frontier models and how they have been used to study productivity and separate it from inefficiency.

3.1 Total Factor Productivity

Productivity has rightfully long been a focal point in many branches of economics; various aspects of productivity can inform on the effectiveness with which inputs can be transformed into outputs as well as the overall efficiency of production. Total factor productivity has been studied extensively and provides a useful metric: The number units of real output a firm can produce with one unit of real inputs (Jorgenson and Griliches, 1967). The value of this measure can be easily seen; it can be used to evaluate economies of scale, trends provide information about growth rates, and heterogeneity across firms can shed light on factors that affect productivity and costs.

The notion of total factor productivity was created to explain economic growth. Growth can either be the result of increased use of inputs, usually called capital accumulation,[3] or growth in productivity. In light of limited resources, increases in productivity are the only way to sustainably promote economic growth.[4] Empirical findings have shown that productivity is the main cause of changes in economic growth; in his seminal paper, Solow (1957) found that between 1909 and 1949, approximately one-eighth of the variation in output was due to capital accumulation and seven-eighths was due to changes in productivity. Further,

[3] While non-capital inputs (e.g., labor) can also increase output, economists have historically not attributed long-term growth to those factors since the stock of those inputs tends to grow at a relatively slow rate.

[4] If resources are limited, capital cannot be endlessly accumulated, so growth must come from another source.

he estimated that annual productivity growth rates ranged from -7.6% to 7.2%, at an average of 1.5% per annum. Finally, Solow estimates several forms for the production function. Using a log-linear (i.e., Cobb-Douglas) form, he estimates that the level of productivity was approximately 0.482.

Productivity has been estimated using a variety of methods. Solow's seminal work on productivity suffered from a number of practical issues. Most notably, any deviations from the empirical model (i.e., residuals) were assumed to be the result of differences in productivity (Solow, 1957). Of course, there are many additional sources of error including differences in efficiency and measurement error. Further, Solow used a linear approximation in his analysis, which does not allow inputs to exhibit complementarity or substitutability and can result in large approximation errors. Several models have been developed and extended to address these issues and largely fit into two groups, either parametric or non-parametric.[5]

Parametric models assume a specific form for the production function and aim to decompose shifts in the production frontier into changes in productivity and efficiency and measurement error. There has been an abundance of research that estimate translog cost functions and infer changes in productivity from shifts in the cost function. Since translog cost functions are a second-order approximation of the true cost function, this method reduces the approximation error presents in Solow's work. Caves et al. (1981) used this framework to derive an expression for productivity growth that depends on the change in costs and change in output over time. The authors estimate these parameters in a cost function and use them to calculate productivity growth in the U.S. passenger and freight rail industry.

[5] Of course, semi-parametric models, which have some parametric and some non-parametric components, have also been used. For further reading, see Jondrow et al. (1982) and Park and Simar (1994).

Cost function frameworks similar to this have been used to study productivity in a variety of contexts.[6] While this framework is very flexible, it cannot separately identify productivity from inefficiency. Stochastic frontier (SF) models, which are further explained in Section 3.3, extend the standard translog estimation framework to include inefficiency; by noting that efficiency must lie between 0% and 100%, structure can be imparted on the model that allows productivity and efficiency to be separately identified (Aigner et al., 1977).

Non-parametric methods remain agnostic of the specific functional form of the production function and instead rely on non-parametric methods to infer its shape. Data envelopment analysis (DEA) is most commonly used to infer the production frontier. This method assumes that production plans lie on the frontier and uses linear programming techniques to trace out the exact location of the frontier (Friesner et al., 2006).[7] Lim and Lovell (2009) use DEA to investigate short-run profit changes in the rail industry. By using non-parametric methods to identify the production frontier, the authors decompose changes in profit into changes in price and productivity.

Both parametric methods like SF and non-parametric methods like DEA commonly appear in the literature. Eisenbeis et al. (1999) compares the two in the context of the banking industry and finds that while the level of estimated inefficiency is higher under DEA, the two measures are highly correlated, indicating they capture similar information. However, the authors also find that estimates from SF analysis more accurately capture efficiency in management and preferences for risk than do linear programming methods.

[6] For more examples of studies of the railroad industry that use translog cost functions, see Bitzan and Wilson (2007), and Bitzan and Keeler (2007).

[7] More specifically, the basic DEA model assumes *non-dominated* plans, for which no other plan produces more output with fewer inputs, lie on the frontier. Extensions of this model have been made to include inefficiency.

3.2 Stochastic Frontier Models

Stochastic frontier (SF) models extend the productivity estimation framework in two important ways. First, they assume that firms are not necessarily efficient; these analyses distinguish maximum possible output from actual output and term the deviation between the two inefficiency. The seminal work of Aigner et al. (1977) presents a commonly used formulation of the stochastic frontier framework:

$$q_i = f(x_i; \beta) + \varepsilon_i - \delta_i. \tag{1}$$

In equation (1), q_i represents log output, x_i are inputs, β are parameters describing production, ε_i is a productivity, and δ_i is inefficiency. Under usual error assumptions, it would be impossible to separately identify ε_i from δ_i; however, inefficiency, defined as deviation from maximum output, is inherently one-sided. Using this assumption and modeling δ_i as, for example, half-normal or log-normal allows productivity shocks and inefficiency to be separately identified.

Similar to other parametric methods of estimating production, the problem of estimating the form of the production function remains. Unlike non-parametric methods like DEA, parametric models require a the researcher to assume a specific form of the production function; in order to maintain minimal assumptions on the exact shape of the production function, researchers have historically utilized some form of functional approximation to address this problem. Early research into productivity, such as Solow (1957), used a first-order Taylor approximation on the log-production function, also known as the Cobb-Douglas form. While this form provides a good starting approximation, it does imply the assumption that production is additive in the inputs; that is, the Cobb-Douglas form assumes that the productivity of any given input depends only on the amount of that input being used and not on any other inputs. Christensen et al. (1973) test simple functional forms that assume additivity and constant returns to scale against a more flexible second-order Taylor approximation of the log-production function, also known as the translog form. The authors use data describing private production in the United States from 1929 to 1969 and find that the assumption of additivity is clearly not satisfied, leading to bias when a first-order approximation is used. As a result, it is safer to use a more flexible functional form such as translog or even higher-order approximations if the assumption of additivity is not clearly satisfied.

Stochastic frontier models have been extended in a number of ways and have been used to identify differences or changes in productivity and efficiency. Schmidt and Lovell (1979) separately estimate the production frontier and the cost function to separate technical inefficiency, which originates in the transformation of inputs into outputs, from allocative inefficiency, which occurs when inputs are not used in the optimal proportions. Applying this model to steam-electricity generation, there was evidence that both types of inefficiency were significant: Technical inefficiency raised costs by about 8.5% while allocative inefficiency raised costs by about 9.2%.

Kumbhakar (1988) adapts the technical/allocative inefficiency framework to panel data, assuming that productivity is constant across all firms and time but that technical inefficiency varies by firm. The author applies this model to Class I Railroads and estimates input demand to correct for possible endogeneity and separate technical and allocative inefficiency. As expected, the author finds sizable variation in inefficiency across firms.

Many researchers have found success in using Bayesian methods to estimate SF models. Generally, other estimation procedures such as maximum likelihood estimation can produce unstable estimates of SF model parameters (van Den Broeck et al., 1994). Further, the parameter uncertainty expressed by standard methods may not be accurate, especially for small sample sizes (Koop et al., 1995). Bayesian methods are able to properly express parameter uncertainty for large and small samples alike and tend to produce more stable estimates. Recent research by Yan et al. (2009) have introduced Bayesian estimation and extended the SF framework to analyze to analyze panel data and models productivity and inefficiency in a flexible manner. The authors assume productivity follows a deterministic trend shared by all firms and inefficiency is a random effect across firms with structural breaks across time. In using this model to analyze container ports, the author finds productivity increases by about 4.4% per year and that inefficiency showed heterogeneity both other firms and across time.

3.3 Productivity of Railroads

Productivity of railroads has long been a topic of interest to regulators and researchers. Worries about the efficiency and productivity of railroads was a major impetus for the partial deregulation of the railroad industry in the 1970s (Bitzan and Wilson, 2007). Proponents of deregulation argued that because of the development and improvement of other modes of shipping like planes, barges, and trucks and because of changes in the mix of products being shipped, existing regulation intended to promote competition for the majority of the 20th century were hindering efficiency and limiting the cost-viability of the industry (Winston et al., 1990). Following deregulation, firms were more easily able to take advantage of economies of scale by merging and could reduce costs by abandoning rail lines that were unprofitable (Bitzan and Keeler, 2007).

Of course, there has been much interest in how these changes have affected productivity. Additionally, in the light of the different characteristics and actions that firms took after deregulation, there is interest in how firms differentially progressed following deregulation and what factors led to those differences. Finally, the ultimate prospects for the industry remain unclear; recent declines in aggregate demand have further cut into firm profits and other modes of transportation continue to improve.

Many studies that examine railroad productivity have been conducted, and questions about many aspects of the industry have been addressed. Caves et al. (1981) began the investigation into changes in the industry and its effect on productivity and viability. The authors found that the industry was quickly becoming more productive prior to 1963; productivity growth was estimated to be 3.5% per year on average during the period from 1955 to 1963. However, in the following period from 1963 to 1974, productivity grew at a much slower rate, only 0.6% per year on average. The authors posit that in the early period, many firms began small and were able to exercise economies of scale as they grew through the late 1950s and early 1960s. The growth of the size of these firms slowed and most excess capacity was filled by the mid 1960s, leading to slower productivity growth through the mid 1970s. In accompaniment with changes in the industry, this slowing growth led many to worry about its ultimate survival and spurred its partial deregulation.

A crucial question for regulators is whether and by how much partial deregulation helped the industry. Tretheway et al. (1997) indirectly address this question by examining productivity and performance of Canadian railways, which underwent partial pricing deregulation in 1967, and compared with U.S. railroads, which were partially deregulated later in 1976 and 1980. While Canadian railroads had significantly higher productivity growth than U.S. railroads prior to the deregulation of the U.S. rail industry, U.S. railroads saw productivity growing between 1.3% and 1.5% per annum faster than Candian railways between 1981 and 1988. The authors conclude that this increased growth was due to reductions in the amount of inputs being used as well as higher traffic density in the U.S. While partial deregulation was not *necessarily* responsible for these changes, it did provide an environment where firms could more easily merge, thereby taking advantage of economies of scale, and had greater flexibility in abandoning routes, which could have lead railroads to find a more advantageous traffic density.

Further, while it was clear that some kind of intervention was needed to ensure the viability of the U.S. rail industry, there have been questions over exactly what type of intervention would be most beneficial to firms. Apart from deregulation, which aims to utilize free-market principles to improve the efficiency and viability of firms, the most commonly suggested intervention is public ownership of railways. Public ownership of the entire rail industry has not been investigated since such a program has not been enacted, but several studies have looked into public ownership of firms and appropriation of public funds towards private railways. Caves and Christensen (1980) compare the publicly-owned Canadian National Railway (CN) with the privately owned Canadian Pacific Railway (CP). Opponents of public ownership worry that firms won't face the proper incentives to minimize costs and improve efficiency. The authors found that competitive pressures both between CN and CP and from other modes of transportation were very strong; as a result, both railways experienced similar productivity growth. Due to an abundance of unprofitable lines, CN initially had lower productivity at the beginning of the sample in 1956; however, both CN and CP aggressively abandoned track through 1967 and saw their productivities converge and continue to grow at a similar rate through 1974.

Rather than owning railways outright, governments can appropriate funds towards supporting rail operations. Similar to completely publicly-owned firms, railroads that receive subsidies may not have the incentive to minimize costs and maximize efficiency absent sufficient competitive pressures. Oum and Yu (1994) consider railways in nineteen OECD countries[8] and investigate the effect of public funding and firms' autonomy from their public funders on efficiency. Since many of the firms did not see competition from other railroads or other forms of competition, efficiency tended to be higher for less publicly funded firms and for firms that had a greater degree of autonomy. In all, whether public funding or ownership is beneficial or detrimental is extremely dependent on whether firms will face competition; when they don't, railways will not have the incentive to improve and will tend not to do so as a result.

4 Conceptual Framework

Fundamentally, productivity research compares production plans and determines how much of the difference in output is due to increased use of inputs and how much is due to changes in productivity. An example of this decomposition is shown graphically in Figure 1. In this example, I consider an output that uses two inputs X_i and X_j. In practice, the researcher

[8] Contrary to many studies of the U.S. rail industry, the railways in this study transport mostly passengers rather than freight.

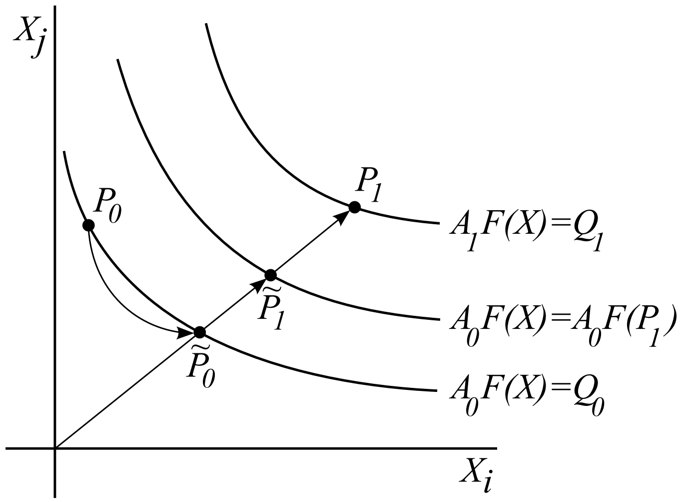

Figure 1: Input Substitution and Productivity

is faced with two production plans, given by P_0 and P_1 in the figure, and knows their associated levels of output Q_0 and Q_1. The two plans are on two different isoquant curves, given by $A_0F(X) = Q_0$ and $A_1F(X) = Q_1$. These differ only by a productivity factor, and the researcher's goal is to estimate the growth of productivity from A_0 to A_1. In changing production from P_0 to P_1, the firm can first change the composition of inputs it uses to most efficiently produce P_1; this is called input substitution and is shown in the graph by the shift from P_0 to $\widetilde{P}_0$. Since P_0 and $\widetilde{P}_0$ are on the same isoquant, they produce the same level of output, i.e., $A_0F(P_0) = A_0F(\widetilde{P}_0)$. The firm can also increase the amount of inputs it uses; this is shown in the graph by the movement from $\widetilde{P}_0$ to $\widetilde{P}_1$, where $A_0F(\widetilde{P}_1) = A_0F(P_1)$. Then, the proportional growth in output due to increased inputs is

$$\frac{A_0F(\widetilde{P}_1)}{A_0F(\widetilde{P}_0)} = \frac{F(P_1)}{F(P_0)}. \tag{2}$$

Finally, output can increase due to changes in productivity, shown in the graph by the shift from $\widetilde{P}_1$ to P_1. The proportional growth in output due to the increase in productivity is given by

$$\frac{A_1 F(P_1)}{A_0 F(\widetilde{P}_1)} = \frac{A_1 F(P_1)}{A_0 F(P_1)} = \frac{A_1}{A_0}. \tag{3}$$

Overall, the total proportional change in production is given by

$$\frac{Q_1}{Q_0} = \frac{A_1}{A_0}\frac{F(P_1)}{F(P_0)}. \tag{4}$$

To reiterate, the ratio A_1/A_0 represents the proportional increase in productivity while the quotient $F(P_1)/F(P_0)$ represents the proportional increase in output due to increased inputs.

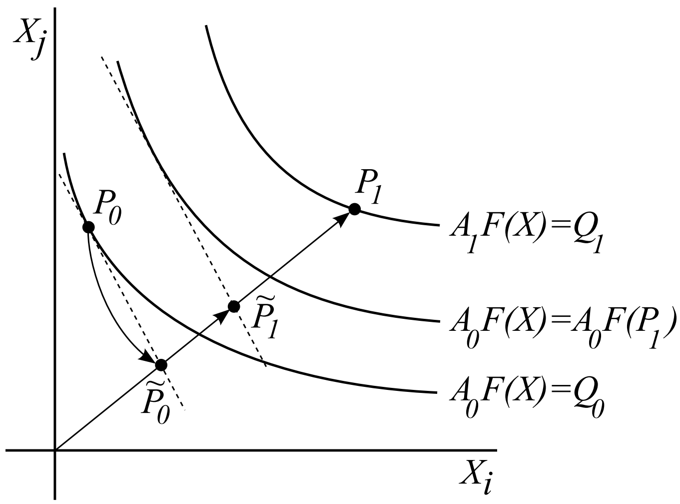

Figure 2: Estimating Productivity With Linearization

The above decomposition assumes the researcher knows the shape and position of the isoquants (and therefore also knows the production function). In practice, this is rarely true, and researchers have relied on a number of techniques to approximate or infer the shape of the production function. The use of Taylor approximations has been very prevalent in productivity estimation. These approximations have been popular largely because of their flexibility; apart from differentiability, Taylor approximations make no assumptions on the shape of the production function and as a result may be used in a variety of contexts. Further, by including more terms in the approximation, it can be made as accurate as desired, subject to data restrictions. To see how this affects estimation of productivity, consider Figure 2; the dotted line represents the linearization of the production function around the point P_0. In practice, the researcher does not observe the isoquant $A_0F(X) = Q_0$ but instead can only approximate its form. The researcher would then estimate input substitution as the shift from P_0 to $\widetilde{P}_0$. The proportional change in output due to increased use of inputs would be estimated as the shift from $\widetilde{P}_0$ to $\widetilde{P}_1$. Finally, the change in output due to increased productivity is estimated to be the shift from $\widetilde{P}_1$ to P_1. Comparing these results to Figure 1, we can see that the researcher would overestimate productivity in this example. Naturally, more complex Taylor approximations can be used, which would decrease the approximation bias in productivity estimates.

While having been used extensively in the literature, the simple framework presented above suffers from theoretical and practical issues. First, as illustrated above, productivity estimates could be biased due to errors in approximating the production function. Many have worked to reduce these errors by using higher order approximations, but unfortunately approximation error can never be eliminated because the approximations are never exact. As noted previously, some researchers have found success in using data envelopment analysis (DEA) to non-parametrically estimate the production frontier (and therefore production function).

Additionally, the above framework has encountered issues in estimation. Specifically, researchers have historically used a deterministic production function and have inferred that any deviations from that function (i.e., residuals) are due to differences in productivity. Of course, there are many channels through which errors can propagate. For example, in addition to approximation and measurement error, inefficiency, defined as deviation from maximum possible output, can affect the level of production independent of changes in productivity. As discussed in greater length in Section 3.3, stochastic frontier (SF) analysis works to separately identify these various sources of error by imparting structure on their form.

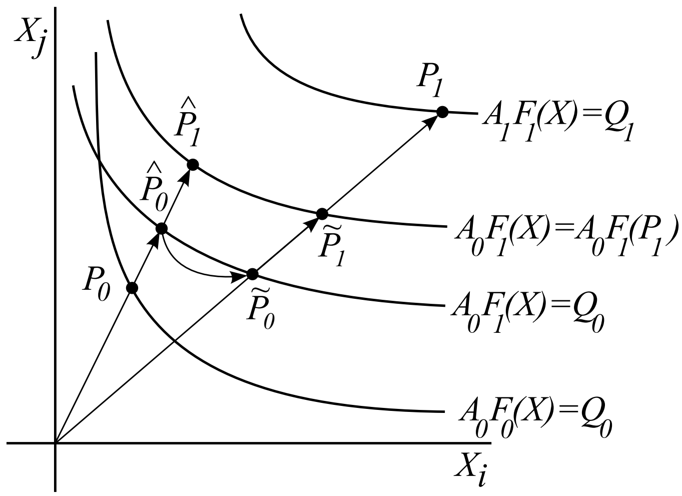

Figure 3: Productivity and Changing Technology

Finally, the method of identifying productivity described above ignores the possibility that the production technology could change. This will lead to a bias in the estimate of productivity. As an example, consider Figure 3, which shows two isoquants that describe different production technologies. The function F_1 represents the new production technology and F_0 represents the original. Since the isoquant is relatively less steep under F_1 than F_0, the input X_j is more productive under the new production technology. The firm increases production from Q_0 to Q_1 through a few different channels. First, the production technology changes, which changes the productivity of inputs, which in turn affects output. The proportional change in productivity due to innovation is shown in Figure 3 by the movement from P_0 to $\widehat{P}_0$. Notice that in this example, the production technology is becoming less efficient since it requires more of both inputs to produce the quantity Q_1, holding productivity constant. Then, the firm can substitute inputs to more efficiently produce Q_1, which is shown by the movement from $\widehat{P}_0$ to $\widetilde{P}_0$. Next, the firm can increase the amount of inputs it uses, shown in the shift from $\widetilde{P}_0$ to $\widetilde{P}_1$. In practice, the researcher cannot separately identify the change in technology from the increased use of inputs because those changes occur simultaneously. However, the researcher can observe the sum of these effects, shown in the graph by the movement from P_0 to $\widehat{P}_1$. The proportional change in output due to innovations and increases in inputs is quantified by

$$\frac{A_0 F_1(\widehat{P}_1)}{A_0 F_0(P_0)} = \frac{F_1(P_1)}{F_0(P_0)}. \tag{5}$$

Finally, the change in productivity due to factors other than technological change is shown by the movement from $\widetilde{P}_1$ to P_1, which can be written as

$$\frac{A_1 F_1(P_1)}{A_0 F_1(\widetilde{P}_1)} = \frac{A_1 F_1(P_1)}{A_0 F_1(P_1)} = \frac{A_1}{A_0}. \tag{6}$$

The total change in production can then be expressed as

$$\frac{Q_1}{Q_0} = \frac{A_1}{A_0}\frac{F_1(P_1)}{F_0(P_0)} \tag{7}$$

$$= \frac{A_1}{A_0}\frac{F_1(P_1)}{F_1(P_0)}\frac{F_1(P_0)}{F_0(P_0)}. \tag{8}$$

Here, A_1/A_0 represents the proportional change in productivity, $F_1(P_1)/F_1(P_0)$ represents the proportional change in output due to increasing inputs, and $F_1(P_0)/F_0(P_0)$ represents the proportional change in output due to changing technology using the inputs P_0. Thus, conditional on approximating the production function and how it evolves over time, the researcher can separately identify changes in output due to increased inputs, improvements in technology, and increases in productivity due to non-innovative factors.

In my analysis of changing productivity, there are a few key values of interest derived above. First, $F_1(P_0)/F_0(P_0)$ denotes the change in productivity due to innovation and A_1/A_0 represents the proportional change in productivity due to factors other than technological change. As a result, the product $A_1F_1(P_0)/A_0F_0(P_0)$ is the total change in productivity due to any factor, which I refer to as *technology-inclusive productivity growth*. I also focus on another value, $F_1(P_1)/F_0(P_1)$, which measures technology's contribution to productivity growth using the new inputs P_1. By comparing this value to $F_1(P_0)/F_0(P_0)$, inferences can be made about the benefits of input substitution:

- If $F_1(P_0)/F_0(P_0) > F_1(P_1)/F_0(P_1)$, the new technology increases output more for the original plan than for the new plan. Thus, the firm substituted towards inputs that innovation made less productive.

- If $F_1(P_0)/F_0(P_0) < F_1(P_1)/F_0(P_1)$, the new technology increases output more for the new plan than for the original plan. This indicates that the firm substituted towards factors that technology change made more productive.

Next, I turn to explaining how I approximate the shape of the production function. Consider the output of a firm i in year t, given by Q_{it}. Suppose that the firm's production technology is described by

$$Q_{it} = A_{it} F_t(X_{it}; \Phi_{it}) \Delta_{it}, \tag{9}$$

where α_{it} is a productivity factor, X_{it} is a vector of inputs, φ_{it} is a vector of network characteristics, and Δ_{it} is a constant between zero and one describing efficiency. While it is not possible to determine the exact shape of F_t, one can approximate it using the second order Taylor approximation of $\ln Q_{it}$ around zero:

$$\begin{aligned} q_{it} \approx \alpha_{it} - \delta_{it} + \sum_j \frac{\partial \ln F_t}{\partial \ln X^j} x_{it}^j + \sum_j \frac{\partial \ln F_t}{\partial \ln \Phi^j} \varphi_{it}^j \\ + \frac{1}{2!} \sum_j \sum_k \frac{\partial^2 \ln F_t}{\partial \ln X^j \partial \ln X^k} x_{it}^j x_{it}^k \\ + \frac{1}{2!} \sum_j \sum_k \frac{\partial^2 \ln F_t}{\partial \ln \Phi^j \partial \ln \Phi^k} \varphi_{it}^j \varphi_{it}^k \\ + \frac{1}{2!} \sum_j \sum_k \frac{\partial^2 \ln F_t}{\partial \ln X^j \partial \ln \Phi^k} x_{it}^j \varphi_{it}^k. \end{aligned} \tag{10}$$

Here, lower-case variables are log-transformed versions of upper case variables, and superscripts index vectors of variables. As an exception, $\Delta_{it} = \exp(-\delta_{it})$, and δ_{it} is restricted to be positive to ensure that the efficiency term Δ_{it} is between zero and one. I first assume that inputs and network characteristics are separable in production, so that $\frac{\partial^2 \ln F_t}{\partial \ln X^j \partial \ln \Phi^k} = 0$

for all j and k. Also, under a modest assumption on F_t,[9] the second derivatives of F_t will be symmetric.[10] Using these assumptions, equation (10) becomes

$$\begin{aligned} q_{it} \approx\ & \alpha_{it} - \delta_{it} + \sum_j \frac{\partial \ln F_t}{\partial \ln X^j} x_{it}^j + \sum_j \frac{\partial \ln F_t}{\partial \ln \Phi^j} \varphi_{it}^j \\ & + \frac{1}{2} \sum_j \frac{\partial^2 \ln F_t}{\partial (\ln X^j)^2} (x_{it}^j)^2 + \frac{1}{2} \sum_j \frac{\partial^2 \ln F_t}{\partial (\ln \Phi^j)^2} (\varphi_{it}^j)^2 \\ & + \sum_j \sum_{k>j} \frac{\partial^2 \ln F_t}{\partial \ln X^j \partial \ln X^k} x_{it}^j x_{it}^k \\ & + \sum_j \sum_{k>j} \frac{\partial^2 \ln F_t}{\partial \ln \Phi^j \partial \ln \Phi^k} \varphi_{it}^j \varphi_{it}^k \end{aligned} \tag{11}$$

Renaming partial derivatives with respect to inputs β^t and those with respect to network characteristics θ^t, equation (11) can be rewritten to arrive at the familiar translog form:

$$\begin{aligned} q_{it} \approx\ & \alpha_{it} - \delta_{it} + \sum_j \beta_j^t x_{it}^j + \sum_j \theta_j^t \varphi_{it}^j \\ & + \sum_j \beta_{jj}^t (x_{it}^j)^2 + \sum_j \theta_{jj}^t (\varphi_{it}^j)^2 \\ & + \sum_j \sum_{k>j} \beta_{jk}^t x_{it}^k x_{it}^j + \sum_j \sum_{k>j} \theta_{jk}^t \varphi_{it}^k \varphi_{it}^j \end{aligned} \tag{12}$$

Let x_{it} be the matrix of all log-inputs, all log-inputs squared, and all of the interactions between log-inputs (i.e., containing each x_{it}^j, $(x_{it}^j)^2$, and $x_{it}^j x_{it}^k$), and let φ_{it} be similarly defined. Then, equation (12) can be expressed in vector form as

$$q_{it} \approx \alpha_{it} - \delta_{it} + x_{it}\beta_t + \varphi_{it}\theta_t \tag{13}$$

[9] Specifically, I assume that the second derivatives of F_t are continuous in a neighborhood of zero.

[10] That is, $\frac{\partial^2 \ln F_t}{\partial \ln X^j \partial \ln X^k} = \frac{\partial^2 \ln F_t}{\partial \ln X^k \partial \ln X^j}$ for all j and k.

Naturally, there is some error incurred in the approximation and measurement of q_{it}. I label this approximation error ε_{it}, so that

$$q_{it} = \alpha_{it} - \delta_{it} + x_{it}\beta_t + \varphi_{it}\theta_t + \varepsilon_{it}. \tag{14}$$

5 Data

The data used in this analysis come from R1 forms, collected and presented by the United States Surface Transportation Board (STB). These forms are published annually and contain financial information and operating statistics for all Class I railroads, including aggregate output and input use and characteristics of each firm's network. The time span of the sample has been restricted to the period from 1999 to 2014; this analysis is interested in how productivity has evolved since Class I railroads fully merged in 1999. The Class I railroads in this sample are Burlington Northern Santa Fe (BNSF), the Canadian National Railway (CN), CSX Transportation (CSX), the Kansas City Southern Railway (KCS), the Norfolk Southern Railway (NS), the Soo Line Railroad (SOO),[11] and the Union Pacific Railroad (UP).

The dependent variable in this analysis is aggregate revenue-ton-miles, which are defined as one ton of product shipped one mile that generates revenue. Production of revenue-ton-miles is described by input use and network characteristics. I use amounts of locomotives and cars, quantity of fuel consumed, and total hours of labor worked, and investment per mile

[11] While Canadian Pacific Railway has owned the Soo Line Railroad since 1990, Soo changed in name to Canadian Pacific in the early 2000s; I will continue to refer to this railroad as SOO.

Table 1: Descriptive Statistics

	BNSF	CN	CSX	KCS	NS	SOO	UP	Total
Output								
Revenue ton-miles (millions)	589625	50471.232	232187.5	26437.171	190250	27312.274	526687.5	234710.097
	(77897.261)	(8422.489)	(14828.885)	(4909.945)	(11716.086)	(6972.383)	(29992.707)	(221554.046)
Inputs								
Locomotives	6180.875	541	3889.562	545.75	3787.688	421.438	7913.688	3325.714
	(895.528)	(107.54)	(248.982)	(53.432)	(266.951)	(90.79)	(646.264)	(2814.122)
Cars	83837.5	24652.75	94733.75	12846.688	96078.125	14690.812	90103.75	59563.339
	(7253.763)	(6450.657)	(20162.264)	(1585.578)	(11392.32)	(1985.792)	(15344.713)	(38578.147)
Investment per mile of road	1210.875	1425.349	1119.413	911.214	1137.301	612.706	1362.493	1111.336
	(127.323)	(289.596)	(221.539)	(192.393)	(233.919)	(68.127)	(139.185)	(319.765)
Gallons of fuel (millions)	1309.606	100.021	547.779	63.283	482.172	54.176	1235.647	541.812
	(113.96)	(24.114)	(53.895)	(6.05)	(33.314)	(12.274)	(127.394)	(504.872)
Labor hours (thousands)	81289.551	12899.74	57640.454	5726.081	57406.789	6274.362	96616.222	45407.6
	(5996.77)	(1162.888)	(5837.744)	(510.784)	(4212.059)	(1262.238)	(8104.72)	(35030.257)
Network characteristics								
Average length of haul (miles)	1046.849	288.825	533.266	363.193	464.402	440.834	921.168	579.791
	(67.523)	(22.837)	(40.975)	(43.623)	(21.087)	(38.505)	(24.854)	(271.546)
Average train speed (MPH)	18.925	20.5	18.411	17.504	19.066	19.612	22.355	19.482
	(1.74)	(1.085)	(0.706)	(2.571)	(0.7)	(1.904)	(1.566)	(2.135)
Miles of road	32482	5847.812	21783.938	3096.312	20889.375	4123.125	32447.312	17238.554
	(424.674)	(1353.077)	(1038.289)	(163.18)	(693.435)	(1325.3)	(575.021)	(12034.503)
Percent unit train	0.476	0.17	0.333	0.414	0.252	0.277	0.415	0.334
	(0.031)	(0.054)	(0.022)	(0.095)	(0.02)	(0.046)	(0.021)	(0.111)
Percent bulk shipments	0.196	0.295	0.194	0.161	0.156	0.294	0.183	0.211
	(0.009)	(0.1)	(0.011)	(0.012)	(0.007)	(0.034)	(0.014)	(0.068)
N	16	16	16	16	16	16	16	112

of road as inputs.[12] Following Spady and Friedlaender (1978), I opt to include investment in firms' networks as an input and include network size as a characteristic of output. The authors found that including network size as input results in negative output elasticities with respect to the network due to economies of density.

Characteristics of each firm's network are crucial in describing production, especially aggregate production, in the railroad industry. Tretheway et al. (1997) investigate the effect of aggregation on the estimates of productivity in the rail industry. The authors estimate productivity using both aggregate and disaggregate data and found significant differences. Using aggregate output assumes that the mix of products being shipped remains constant over time. If, for example, firms instead shift to shipping products that require fewer inputs, productivity would appear to increase even if productive capability remained constant. One would ideally use disaggregated data in their analyses; however, only aggregated data is publicly available, so it is important to control for other factors that could be correlated with aggregate input use but not necessarily with productive potential.

There are several network characteristics that are important to consider. First, as noted in Tretheway et al. (1997), traffic mix is a crucial feature of railroad networks. The mix of traffic is dependent both on the types of goods being shipped as well as types of shipments that traverse the network. In this analysis I include the percentage of shipments that carry bulk products[13] as a measure of product mix and the percentage of shipments that are unit train shipments as a measure of shipment mix. Bitzan and Keeler (2007) additionally find that

[12]Miles of road is defined as the total length of non-redundant track operated by a railroad. Investment was deflated using the GDP price deflator.

[13]I define bulk products as belonging to one of the following categories: Metallic ores, nonmetallic minerals (not fuels), waste/scrap metals, clay/concrete/glass/stone, farm products.

the shipment distance is important in describing costs.[14] The percentage of shipments that are unit train shipments partially captures aspects of shipment distances, and I additionally include the average length of haul into this analysis. Spady and Friedlaender (1978) find that network size is a crucial factor in transportation costs, so I also include miles of road for each firm. Finally, the quality of a railroad's track will determine how efficiently trains can traverse the network and can also influence maintenance costs. Following Wilson (1997), I use average locomotive speed as a measure of network quality.

Descriptive statistics for each of these variables are presented in Table 1. Means and standard deviations are given for each firm and as an average over all firms. The sample spans sixteen years, and the descriptive statistics are averaged over time.

6 Empirical Models

In the analyzing the productivity and efficiency of firms, we are most interested in estimating their respective parameters α_{it} and δ_{it} in equation (14). However, all of the parameters of equation (14) cannot be separately identified in a standard regression framework. There are a number of ways to manipulate this model so that productivity and efficiency can be identified, and in this section I describe the three models I use in this paper.

6.1 Deterministic Trend in Productivity, Constant Technology

This model first assumes that technology and the effect of network characteristics are constant across time, so that $\beta_t = \beta$ and $\theta_t = \theta$. I also assume that each firm has its own initial

[14]This comes at no surprise since, at the very least, short shipments require more fuel per revenue-ton-mile.

productivity which then follows a deterministic linear trend shared by all firms. Further, the model assumes that inefficiency is constant across time (but is allowed to vary by firm); as a result, $\delta_{it} = \delta_i$. Finally, recall that δ_i was restricted to be greater than zero; following the majority of the stochastic frontier literature, I assume δ_i has a half-normal distribution centered and truncated at zero. Inefficiency can be separately identified from productivity both because they have different dynamics across time[15] and because inefficiency is strictly greater than zero. The model can be expressed in the following relations.

$$
\begin{aligned}
q_{it} &= \alpha_{it} + x_{it}\beta + \varphi_{it}\theta - \delta_i + \varepsilon_{it} \\
\alpha_{it} &= \alpha_i + \tau t \\
\alpha_i &\sim N(\mu_\alpha, \sigma_\alpha) \\
\varepsilon_{it} &\sim N(0, \sigma_\varepsilon) \\
\delta_i &\sim N^+(0, \sigma_\delta)
\end{aligned}
$$

Conditional on having prior distributions over the parameters, this model can be estimated using Gibbs sampling. For a detailed description of the sampler, see Section 9.1.1 in the Appendix.

6.2 Random Walk in Productivity, Constant Technology

This model is similar to the previous model, but focuses on modeling productivity in a more flexible way than with a deterministic trend. Specifically, I assume that productivity follows a random walk with drift that is independent for each firm. This type of process has been used in a number of applications and can model many processes, especially those that exhibit

[15] That is, inefficiency is assumed to be static while productivity follows a linear trend.

persistence, flexibly and effectively.[16] Importantly, productivity likely exhibits persistence because firms don't tend to change their exact methods of production by a significant on an annual basis, and as a result, productivity in one year will be dependent on productivity in the previous year.

The previous model also assumed that productivity growth was constant across all firms; as a result, the estimates of productivity growth in that model are best viewed as the average productivity growth in the industry. It is more likely that each firm follows its own trend in productivity due to differences in how firms operate. This model relaxes the common trend assumption and allows each firm to have its own productivity trend.

I maintain all of the other assumptions of the model, which is written below.

$$
\begin{aligned}
q_{it} &= \alpha_{it} + x_{it}\beta + \varphi_{it}\theta - \delta_i + \varepsilon_{it} \\
\alpha_{it} &= \alpha_{it-1} + \tau_i + \eta_{it} \quad ; \quad t > 0 \\
\alpha_{i0} &\sim N(\mu_\alpha, \sigma_\alpha) \\
\varepsilon_{it} &\sim N(0, \sigma_\varepsilon) \\
\delta_i &\sim N^+(0, \sigma_\delta) \\
\eta_{it} &\sim N(0, \sigma_\eta)
\end{aligned}
$$

Given assumptions of the prior distributions of each parameter, which can be found in Section 9.1.2 of the Appendix, this model can be estimated using Gibbs sampling. Exact evaluation of the likelihood is complicated by the random-walk process in productivity, but is made possible via the Kalman filter. A review of the methodology for using the Kalman filter to estimate standard regression models with time-varying parameters in a Bayesian framework is given in Sarris (1973). While stochastic frontier models and time-varying

[16]For some examples of how time-varying parameters have been used in a variety of contexts, see Leybourne (1993), Mazzocchi (2003), and Del Negro and Otrok (2008).

parameter models have been estimated, I am not aware of any published research that combines the two to examine dynamic changes in productivity.

6.3 Random Walk in Productivity and Technology

This model presents an additional extension of the previous model. I maintain the assumption that productivity follows a random walk with drift but relax the assumption that technology remains constant over the time frame of the sample. There are a couple of perspectives that justify relaxing this assumption. First, firms are constantly striving to reduce costs and make innovations to their production technology to further that goal. As discussed in Section 2, firms have invested large amounts in improving improving their networks and pursuing innovation. Ignoring these innovations would lead to a bias in productivity, as discussed in Section 3.1.

A second line of reasoning refers back to the original definition of productivity: The marginal amount of output that can be produced using an additional unit of real resources. As the production technology changes, the combination of inputs that constitutes one unit of real resources will change; not only will the amount of output that can be produced with one unit of real inputs change, but firms will alter the composition of inputs they use as factor productivities change at differing rates. Assuming that technology remains constant over time, an increase in real expenditures will increase all inputs by a constant fixed amount, which will increase output by a fixed amount, after controlling for productivity.

Instead of assuming a technology that is constant across firms and time, I assume that all firms share the production technology (up to their multiplicative productivity) in a given year, but that technology is allowed to change over time. The primary estimating equation then becomes

$$q_{it} = \alpha_{it} + x_{it}\beta_t + \varphi_{it}\theta - \delta_i + \varepsilon_{it}. \tag{17}$$

The data prevent the separate identification of β_t for each year; instead, I propose that β_t follows a random walk with drift. Once again, I expect that β_t will exhibit persistence because new technology tends to adapt existing technology. The other assumptions of the model remain the same, which can be expressed in the following relations.

$$\begin{aligned}
q_{it} &= \alpha_{it} + x_{it}\beta_t + \varphi_{it}\theta - \delta_i + \varepsilon_{it} \\
\alpha_{it} &= \alpha_{it-1} + \tau_i + \eta_{it} \quad ; \quad t > 0 \\
\beta_t &= \beta_{t-1} + \rho + \psi_t \quad ; \quad t > 0 \\
\alpha_{i0} &\sim N(\mu_\alpha, \sigma_\alpha) \\
\beta_0 &\sim N(\mu_\beta, \Sigma_\beta) \\
\varepsilon_{it} &\sim N(0, \sigma_\varepsilon) \\
\delta_i &\sim N^+(0, \sigma_\delta) \\
\eta_{it} &\sim N(0, \sigma_\eta) \\
\psi_t &\sim N(0, \Sigma_\psi)
\end{aligned}$$

Once again, due to the random-walk process in productivity and technology parameters, exact evaluation of the likelihood function is difficult but is possible by using the Kalman filter; the general estimation procedure in a Bayesian context is described in Sarris (1973). Once prior distributions are assigned to each parameter, the model can be estimated using Gibbs sampling. A full description of the model, including prior assumptions, can be found in Section 9.1.3 in the Appendix.

7 Results

This section presents results for each of the three models detailed in Section 6. Each of these models was estimated using Gibbs sampling, a Bayesian estimation technique; unlike classical statistical methods which produce a point estimate for each parameter, Bayesian methods like Gibbs sampling produce a distribution for each parameter that is dependent on prior assumptions, the data, and the structure of the model. Consequentially, I present statistics describing the distribution of each parameter. The distributions of parameters, especially of productivity, have relatively high variance; since many of these parameters are then exponentiated to get their economically-intuitive value, their distributions show significant skew. As a result, I present only estimated medians for each parameter as these will give a better view of the central tendency of these distributions. Parameter estimates for each model estimated are presented in Table 2.

7.1 Deterministic Trend in Productivity, Constant Technology

Median estimates from the model with a deterministic trend in productivity and static technology are presented in the Baseline column of Table 2, and median estimates of annual productivity for each firm are plotted in Figure 4. While this is the most basic model presented in this paper, it provides some general insight into productivity and growth in the industry. First, I estimate that average productivity growth was modest over the period from 1999 to 2014, with median estimates of 1.2% growth per year on average across all firms. There is also heterogeneity in productivity across firms; CN, CSX, NS, and UP show the highest levels of productivity, between 1.65 and 1.708 in 2014, while KCS has the

Table 2: Posterior Medians and Median Absolute Deviations

	Baseline	Constant Technology	Changing Technology
Productivity in 2014 ($\exp(\alpha)$)			
BNSF	1.602 (2.345)	2.594 (0.622)	1.026 (1.513)
CN	1.708 (2.5)	4.962 (1.557)	1.004 (1.487)
CSX	1.682 (2.463)	4.077 (1.045)	0.979 (1.451)
KCS	1.52 (2.224)	2.95 (0.809)	0.873 (1.295)
NS	1.675 (2.453)	3.831 (1.058)	1.116 (1.655)
SOO	1.534 (2.245)	3.296 (1.023)	1.074 (1.592)
UP	1.65 (2.417)	3.82 (0.98)	0.891 (1.321)
Efficiency ($\exp(-\delta)$)			
BNSF	0.814 (0.187)	0.938 (0.056)	0.612 (0.357)
CN	0.864 (0.138)	0.934 (0.043)	0.611 (0.363)
CSX	0.851 (0.16)	0.949 (0.04)	0.611 (0.351)
KCS	0.814 (0.191)	0.925 (0.061)	0.63 (0.357)
NS	0.864 (0.147)	0.954 (0.043)	0.618 (0.356)
SOO	0.782 (0.213)	0.943 (0.054)	0.623 (0.35)
UP	0.85 (0.159)	0.919 (0.064)	0.62 (0.356)
Productivity Trend (τ)	0.012 (0.005)		
BNSF		0.002 (0.008)	-0.01 (0.582)
CN		0.037 (0.016)	0.008 (0.669)
CSX		0.028 (0.009)	0.001 (0.432)
KCS		0.007 (0.007)	-0.008 (0.646)
NS		0.021 (0.007)	-0.003 (0.455)
SOO		0.018 (0.008)	-0.01 (0.58)
UP		0.021 (0.008)	-0.007 (0.597)
Input Parameters			
Locomotives	-0.249 (1.8)	-2.522 (1.216)	0.161 (7.774)
Cars	-0.728 (1.894)	-1.076 (1.57)	0.245 (7.647)
Road	1.965 (1.82)	0.216 (1.62)	0.383 (7.359)
Fuel	1.354 (1.795)	2.928 (1.274)	0.474 (6.53)
Labor	-0.916 (2.088)	1.185 (2.208)	0.476 (6.963)
$(\text{Locomotives})^2$	-0.21 (0.144)	-0.46 (0.126)	0.049 (3.475)
$(\text{Cars})^2$	0.586 (0.145)	0.341 (0.303)	-0.173 (4.978)
$(\text{Road})^2$	-0.173 (0.114)	-0.206 (0.12)	0.22 (3.228)
$(\text{Fuel})^2$	0.165 (0.127)	0.034 (0.153)	0.072 (3.379)
$(\text{Labor})^2$	0.727 (0.401)	-0.287 (0.259)	0.102 (4.505)
(Locomotives):(Cars)	-0.009 (0.207)	-0.009 (0.206)	-0.008 (5.846)
(Locomotives):(Road)	-0.147 (0.183)	-0.486 (0.21)	0.123 (4.963)
(Locomotives):(Fuel)	0.199 (0.204)	0.792 (0.193)	0.059 (4.82)
(Locomotives):(Labor)	0.065 (0.332)	-0.022 (0.399)	-0.277 (5.417)
(Cars):(Road)	0.082 (0.2)	0.051 (0.179)	0.126 (5.262)
(Cars):(Fuel)	0.285 (0.285)	-0.811 (0.216)	-0.016 (5.669)
(Cars):(Labor)	-1.049 (0.534)	0.61 (0.386)	0.157 (6.147)
(Road):(Fuel)	0.076 (0.196)	0.137 (0.217)	-0.277 (4.45)
(Road):(Labor)	0.105 (0.367)	0.348 (0.332)	-0.219 (5.076)
(Fuel):(Labor)	-0.765 (0.264)	-0.121 (0.21)	0.068 (5.884)
Network Characteristics			
Avg. Length of Haul	1.592 (1.972)	-0.99 (2.241)	0.047 (2.939)
Avg. Speed	0.685 (1.982)	-1.438 (1.194)	0.035 (2.926)
Miles of Road	-0.788 (1.585)	-0.541 (1.344)	0.115 (2.88)
% Unit	-3.558 (1.307)	-2.04 (1.077)	-0.046 (2.891)
% Bulk	2.178 (1.387)	1.544 (1.049)	0.019 (2.983)
$(\text{Avg. Length of Haul})^2$	-0.287 (0.205)	0.027 (0.164)	0.015 (1.66)
$(\text{Avg. Speed})^2$	-0.447 (0.316)	-0.187 (0.21)	0.022 (2.497)
$(\text{Miles of Road})^2$	0.061 (0.083)	0.059 (0.066)	-0.014 (0.899)
$(\%\ \text{Unit})^2$	-0.069 (0.122)	0.004 (0.089)	-0.053 (2.268)
$(\%\ \text{Bulk})^2$	-0.061 (0.147)	0.014 (0.108)	-0.095 (2.618)
(Avg. Length of Haul):(Avg. Speed)	0.675 (0.447)	0.51 (0.277)	-0.002 (2.535)
(Avg. Length of Haul):(Miles of Road)	0.136 (0.167)	0.049 (0.119)	0.06 (2.103)
(Avg. Length of Haul):(% Unit)	0.776 (0.278)	0.508 (0.219)	0.021 (2.273)
(Avg. Length of Haul):(% Bulk)	-0.189 (0.215)	-0.267 (0.141)	-0.023 (2.351)
(Avg. Speed):(Miles of Road)	-0.338 (0.177)	-0.197 (0.09)	0.012 (1.908)
(Avg. Speed):(% Unit)	-0.113 (0.33)	-0.076 (0.148)	-0.048 (2.424)
(Avg. Speed):(% Bulk)	-0.462 (0.344)	-0.479 (0.152)	0.034 (2.672)
(Miles of Road):(% Unit)	-0.095 (0.1)	-0.047 (0.078)	0.035 (1.51)
(Miles of Road):(% Bulk)	0.026 (0.097)	0.178 (0.06)	-0.054 (1.654)
(% Unit):(% Bulk)	-0.014 (0.17)	0.12 (0.12)	0.012 (2.641)

lowest productivity at 1.52 in 2014. Mean productivity was estimated to be 1.661 in 2014. Estimates of efficiency range from 78.2% to 86.4%, with a mean of 81.4% across firms.

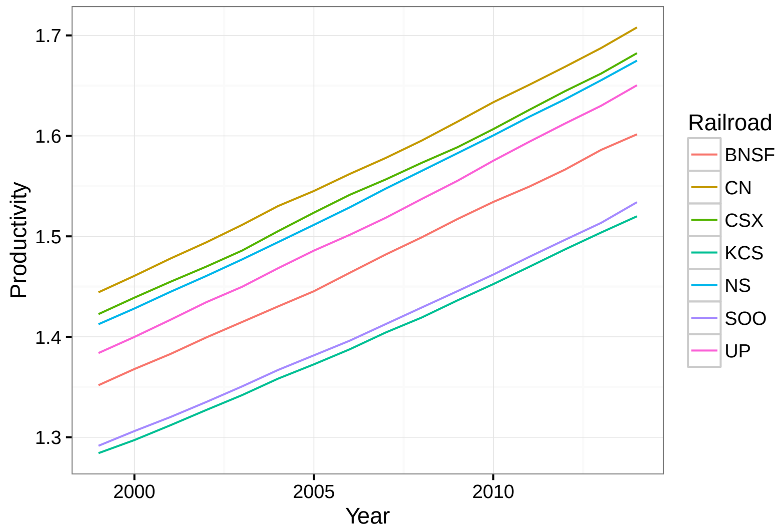

Figure 4: Railroad Productivity as a Deterministic Trend

As discussed in Section 4, the increase in output that results from a input substitution and increased input usage is given by $F(X_{it})/F(X_{i0})$, which is represented by $\exp((X_{it} - X_{i0})\beta)$ in the empirical model. A plot of the proportional increase in output due to changing input quantities is given in Figure 5. CN, KCS, and the Soo Line all saw increases in input use over the sample, CSX, NS, and UP all decreased input use, and BNSF saw little change in output due to changing inputs.

Using a Bayesian estimation framework permits direct evaluation of the probability that a parameter of interest, like productivity growth, lies within a given range. Specifically, the probability that a parameter ϖ lies within a set S is

$$\Pr(\varpi \in S) = \int I(\varpi \in S) p(\varpi | D) d\varpi, \tag{19}$$

where $I(\cdot)$ is an indicator function and $p(\varpi|D)$ is the posterior distribution of ϖ conditional on the data D. I estimate the probability that average productivity increased (i.e., $\Pr(\tau > 1)$) is 99.314%.

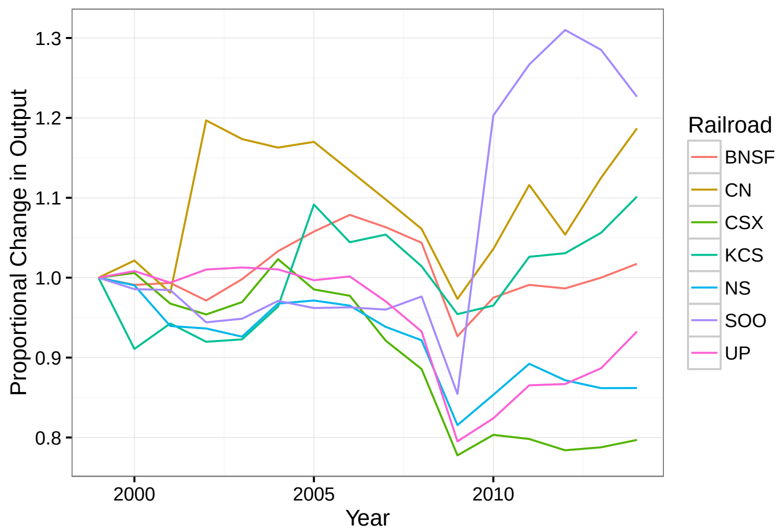

Figure 5: Increase in Output From Change in Inputs

In light of these results, one can safely conclude that there has been growth in average productivity since total consolidation in 1999. This is reassuring, the flexibility granted to firms by partial deregulation seems to have set the groundwork for continued long-term growth and sustained viability. However, while the industry appears to be growing on average, the performance of individual firms is not clear. To investigate how each railroad has progressed, I turn to my second model.

7.2 Random Walk in Productivity, Constant Technology

Median estimates from the random walk in productivity model are presented in the Constant Technology column in Table 2. There are several modest differences from the previous model's results. First, estimates of productivity and efficiency are lower for this model; mean productivity is estimated to be 3.684 in 2014 and mean efficiency was estimated to be 92.9%. I estimate that *effective productivity* (i.e., the product of productivity and efficiency) was 3.385 in 2014. Trends in productivity show marked heterogeneity across firms, and mean productivity growth is estimated to be 1.96% per year, similar to the 1.2% growth estimated by the previous model.

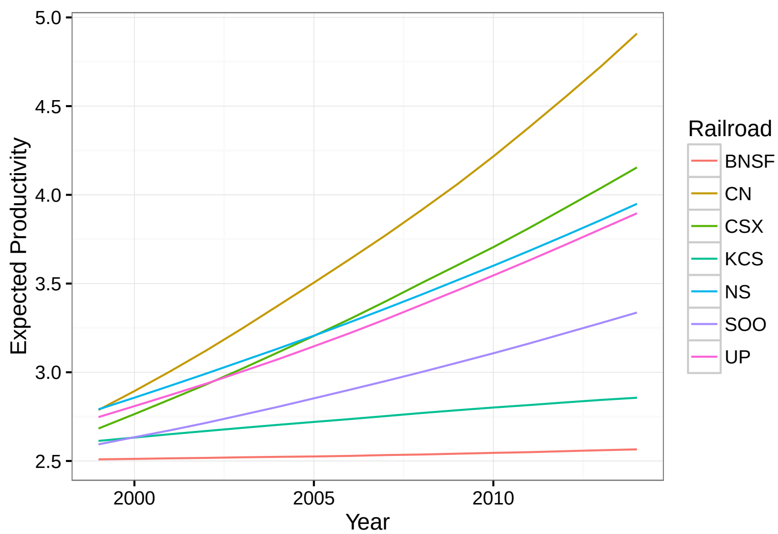

Figure 6: Expected Railroad Productivity as a Random Walk With Drift

Productivity for each firm over time is presented in Figures 6 and 7 in two ways. Figure 6 shows the expectation of firm productivity conditional on firm trends and information in the year 1999; this is identical to the deterministic part of productivity. Figure 7 shows estimated productivity, which includes both the trend as well as the random walk in productivity. The inclusion of the random walk is important because it reflects actual year-to-year variation in productivity that cannot be captured by a simple deterministic trend. There are many factors that could potentially affect productivity, and it would be neither technically feasible

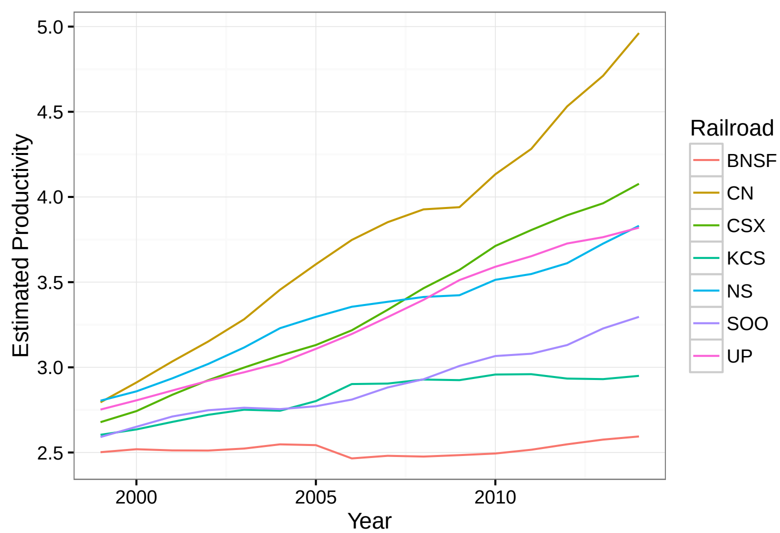

Figure 7: Estimated Railroad Productivity as a Random Walk With Drift

nor even possible given data restrictions to include them all into the model. By assuming productivity follows a random walk with trend, variation in productivity can be captured flexibly.[17]

One can quickly see that actual productivity differs from its expected value. As an example, NS was expected to have productivity growth of 2.1% per year over the course of the sample; in actuality, NS received two negative shocks to productivity in 2008 and 2009, which resulted in lower estimated productivity growth of 1.889% per year. This indicates that there are other factors influencing productivity that cannot be described by changes in input use and a simple trend. As explained previously, it can be difficult to attribute an exact cause to these fluctuations.

Estimates of the effect of changing input use on output are shown in Figure 8. Similar to

[17]Of course, this assumption comes at a cost. While variation in productivity can be identified, the exact source of this variation cannot. As a result, one can only use institutional knowledge to posit why firms see fluctuations in productivity; in order to empirically identify what factors drive changes in productivity, those factors must be explicitly included in the model.

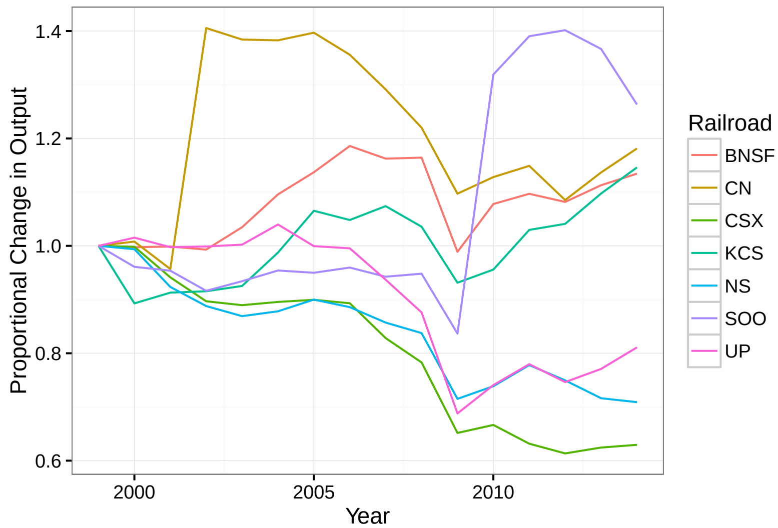

Figure 8: Increase in Output From Change in Inputs

the previous model, I find that BNSF, CN, KCS, and the Soo Line all increased the quantity of inputs used while CSX, NS, and UP saw decreases in input quantities. As noted earlier, CN and Soo Line also saw significant increases in productivity, indicating both matched increased demand with a combination of neutral factors and increased inputs.

I again estimate the probability that firms experience positive average growth in productivity as well as the median value of annualized average growth between 1999 and 2014, and the results are given in Table 3. Estimates suggest that each firm saw increases in productivity, with median estimates between 0.235% to 3.551% per annum, and probability of productivity growth between 65.7% and 100%. Further, I estimate the probability that *all* firms experience positive growth in productivity is 63.424%.

These results offer optimistic outcomes for some firms and a more modest outlook for others. Similar to Tretheway et al. (1997), which found that CN showed high productivity growth through 1991, I estimate that CN has the highest productivity in 2014 as well as

Table 3: Average Productivity Growth

Firm	Annual Productivity Growth	Probability of Increase
BNSF	0.235%	65.7%
CN	3.551%	100%
CSX	2.474%	99.908%
KCS	0.816%	94.102%
NS	1.889%	99.992%
SOO	1.657%	99.926%
UP	1.869%	99.996%

the highest growth rate over the sample. As noted by previous research, CN continued to operate on several less-profitable lines and had yet to fully take advantage of economies of density through the 1990s; as CN continued to improve on those frontiers, its productivity increased.

On the other hand, the remainder of the railroads exhibit modest growth in productivity. Some of these firms may no longer find it feasible to abandon lines and make improvements to economies of density and instead must turn towards innovating their production technology to realize higher growth. This model ignores the possibility of technological change; to investigate whether firms have been able to increase productivity through innovation, I turn to my third and final model.

7.3 Random Walk in Productivity and Technology

Median estimates from the changing technology model are presented in the Changing Technology column in Table 2, and median estimated productivity for each firm is plotted in Figure 9. As explained in Section 4, there are a few key values of interest in the analysis of productivity in the light of changing technology. First, the productivity growth estimated by this model does not include changes in productivity due to innovation but instead reflects growth from factors other than innovation. Both CN and CSX show growth in productivity due to non-innovative factors at 0.702% and 0.016% per year, respectively, indicating those firms may still be able to abandon lines and improve density to increase productivity. All other firms experience stagnating or even declining productivity due to non-innovative factors at rates between 0.283% and 0.901% per year, demonstrating that those other methods of increasing productivity are no longer feasible for these railroads.

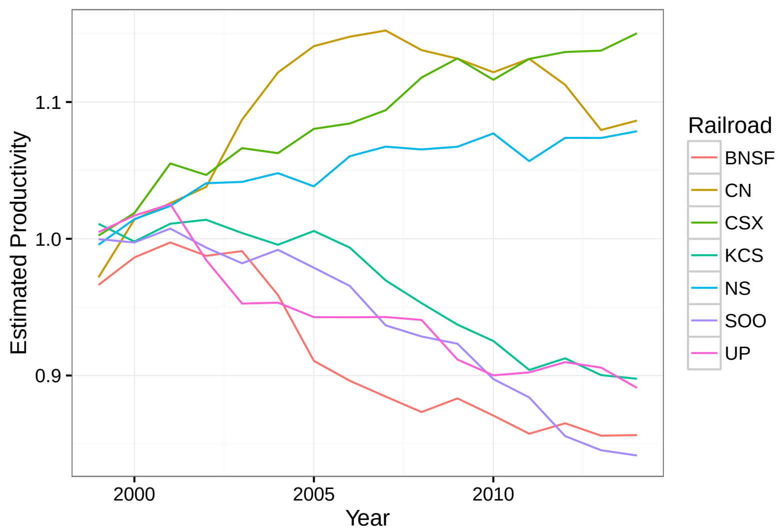

Figure 9: Estimated Railroad Productivity Accounting For Changing Technology

The growth in productivity due to technological change relative to the base year is given by $X_{i0}(\beta_t - \beta_0)$.[18] A plot of the estimated change in productivity due to innovations for each railroad is presented in Figure 10. Each railroad experienced positive average growth in productivity due to innovations over the course of the sample. BNSF, KCS, Soo Line, and UP saw the largest productivity gains due to chnages in technology. On the other hand, CSX, CN, and NS found more modest increases, indicating those firms have relied more heavily on

[18]Productivity growth due to innovation was given as $F_1(X_0)/F_0(X_0)$. Note that this is an analogous measure because $\ln(F_1(X_0)/F_0(X_0)) = f_1(X_0) - f_0(X_0) = X_0\beta_1 - X_0\beta_0 = X_0(\beta_1 - \beta_0)$.

other methods to increase their productivity. While I have a small sample of firms, I find that the large firms (i.e., BNSF and UP) are benefactors of changing technology; this correlates with the finding of Rose and Joskow (1990) where large firms are more likely to adopt innovations early due to risk preferences. A plot of technology-inclusive productivity, the sum of productivity due to innovation and that due to non-innovative factors, is presented in Figure 11. Overall, BNSF, CN, and Soo Line have shown the highest growth in total productivity. KCS and UP experienced modest gains in total productivity over the course of the sample, and all CSX and NS experienced relatively little productivity growth.

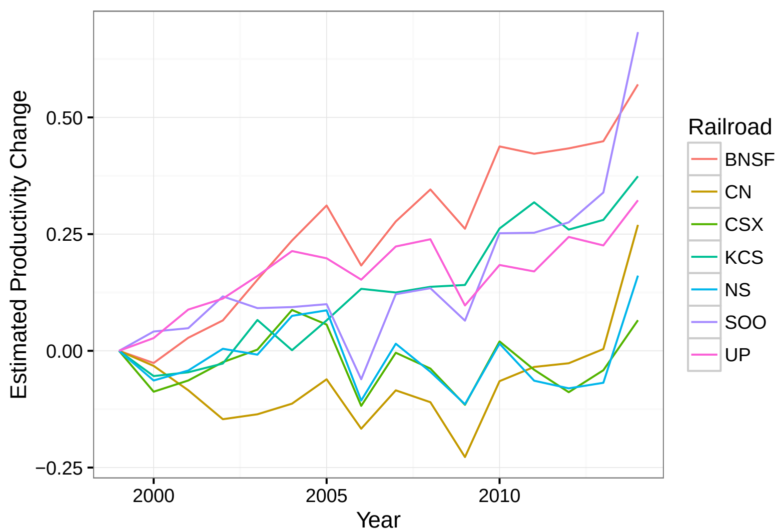

Figure 10: Estimated Change in Productivity Due to Innovations

Median estimates of average annual technology-inclusive productivity growth and the probability that each firm experienced positive growth in productivity are given in Table 4. Productivity growth estimates are much more modest than in previous models; CN saw the largest expected increase at 3.065% per annum, while NS experienced small decreases in productivity of 0.007% annually. I estimate the probability that *all* firms experienced an

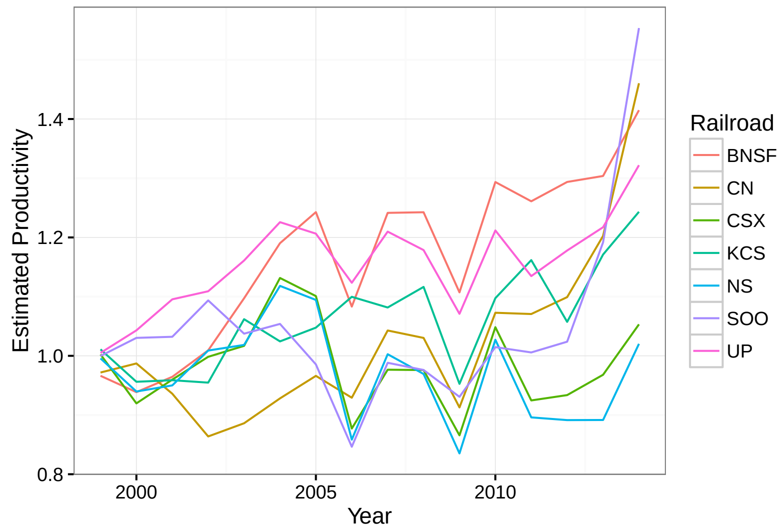

Figure 11: Estimated Productivity Including Effects of Innovation

increase in productivity between 1999 and 2014 is only 4.688%, indicating it is very likely that at least one firm saw a decrease in productivity over the sample period.

Table 4: Average Productivity Growth

Firm	Annual Productivity Growth	Probability of Increase
BNSF	2.272%	52.324%
CN	3.065%	51.792%
CSX	0.037%	50.036%
KCS	1.248%	51.156%
NS	-0.007%	49.992%
SOO	2.79%	53.064%
UP	1.623%	52.028%

The change in output due to input substitution and increased input usage is shown Figure 12. Most firms didn't increased or even decreased the amount of inputs they used over the sample, with the exception of CN. While CN saw decreases in productivity due to technological change, especially before 2010, it dramatically increased its input use during that time. This provides evidence that CN relied on increasing its input use rather than increasing productivity through innovation, especially before 2010.

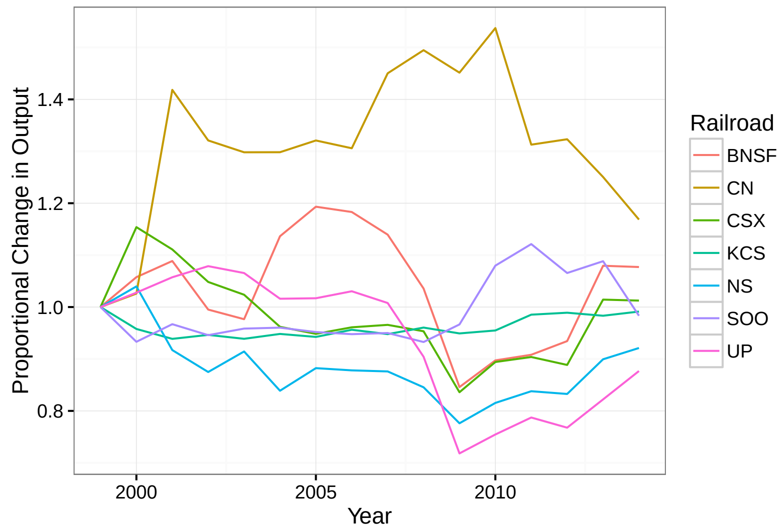

Figure 12: Increase in Output From Change in Inputs

The preceding analysis examines how innovations affect the productivity of railroads' original plan of production. To determine whether firms made changes that took advantage of changing technology, namely by allocating more resources towards more productive factors, one can examine how changing technology affects new production plans. In Section 4, I described the comparison of two measures: Technology's benefit to the original production plan, $F_t(X_{t-1})/F_{t-1}(X_{t-1})$, and the benefit of innovation to the new production plan, $F_t(X_t)/F_{t-1}(X_t)$. Since $F_t(X_s)$ is expressible as $\exp(X_s\beta_t)$ in this model, these two ratios can be calculated as $\exp(X_{t-1}\beta_t - X_{t-1}\beta_{t-1})$ and $\exp(X_t\beta_t - X_t\beta_{t-1})$, respectively. Recall that if the latter is greater than the former, then the firm allocated resources towards factors that innovation made more productive; otherwise, it made changes that clashed with technology change and must have found it cheaper to increase productivity through means other than input substitution. To exhibit these results, I calculate distributions of both values for each firm and year. Bayesian estimation methods make it possible to evaluate the probability that one of these values exceeds the other; I plot the probability that each firm substitutes towards more productive inputs[19] in each year in Figure 10. Every railroad has near 50% probability of substituting towards more productive inputs and no railroad appears to have consistent behavior. This shows that firms are largely not able to anticipate changes in the relative productivity of inputs and make allocative changes in response.

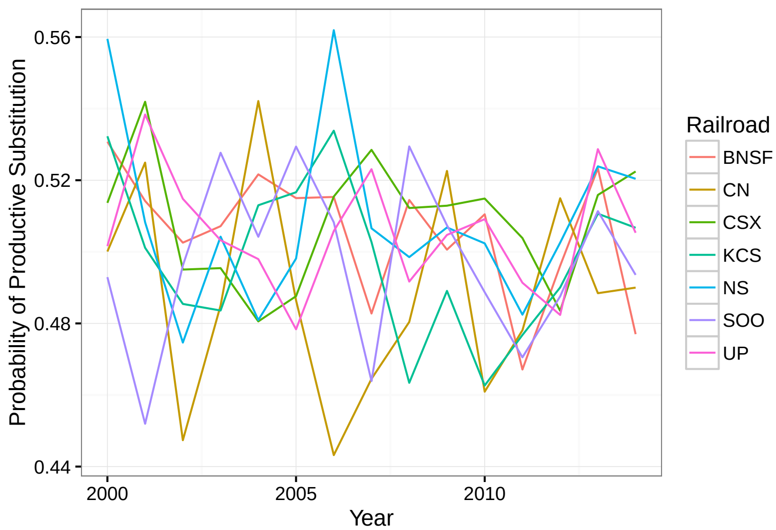

Figure 13: Probability of Substituting Towards More Productive Input

7.4 Bayesian Model Selection

An advantage of using a Bayesian estimation framework is that it allows for the direct computation of model probabilities, which can be used for model selection and averaging of results. Letting each of the previously discussed models be M_1, M_2, and M_3, respectively, the probability of model M_i being the correct model is

$$\Pr(M_k|D) = \frac{\Pr(D|M_k)\Pr(M_k)}{\Pr(D)} = \frac{\Pr(D|M_k)\Pr(M_k)}{\sum_j \Pr(D|M_j)\Pr(M_j)}, \tag{20}$$

[19] That is, I calculate the probability that $F_t(X_t)/F_{t-1}(X_t)$ is greater than $F_t(X_{t-1})/F_{t-1}(X_{t-1})$.

where $\Pr(D|M_k)$ is the marginal likelihood of the data D for model k and $\Pr(M_k)$ is the prior probability of model k, chosen by the researcher. Direct evaluation of the marginal likelihood is difficult in general, but can be computed using the methods described in Chib and Jeliazkov (2001).

I assume a uniform prior probability over the above three models, and posterior model probabilities are given in Table 3. The model that allows both productivity and technology

Table 5: Posterior Model Probabilities

Model	Prior Probability	Posterior Probability
Deterministic trend	1/3	0.24311
Random walk in productivity	1/3	0.08665
Random walk in productivity and technology	1/3	0.67024

to follow a random walk with drift has the highest probability of being the true model. The effects of changing technology have an important effect on productivity changes, as can be seen from the relatively low probability of the model that only allows a random walk in productivity and not in technological parameters.

Table 6: Average Productivity Growth

Firm	Annual Productivity Growth	Probability of Increase
BNSF	1.817%	64.907%
CN	2.636%	67.522%
CSX	0.513%	66.337%
KCS	1.181%	66.585%
NS	0.433%	66.315%
SOO	2.287%	68.368%
UP	1.524%	67.68%

Finally, I calculate median productivity for each firm over time, median productivity growth over the course of the sample, and the probability that firms experienced increases in productivity between 1999 and 2014 using Bayesian model averaging, which calculates parameters of interest for each model and weights them by their respective model probabilities. Specifically, for a statistic of interest ϖ, the average of ϖ over the models is

$$E[\varpi|D] = \sum_k \varpi_k \Pr(M_k|D),$$

where ϖ_k is the estimated value of ϖ in model k. Estimated productivity growth and the probability that each firm experienced growth in productivity over the course of the sample are given in Table 4, and figure 11 plots productivity averaged over the models for each firm over time.

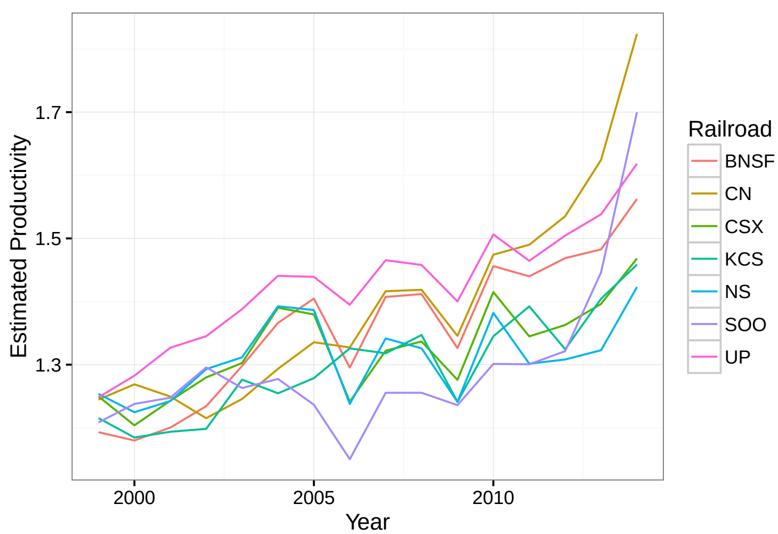

Figure 14: Model Average of Productivity Over Time

Model average results show that CN experienced an average annual productivity growth of 2.636% per year between 1999 and 2014, the greatest of all firms in the sample; CSX and NS showed the least growth, at 0.513% and 0.433%, respectively, while the remainder of the firms saw productivity growth between 1.181% and 2.287% annually. Each firm had a relatively high probability of experiencing positive growth, between 64.907% and 68.368%.

Finally, I calculate that the probability that *all* firms experienced positive growth 32.782%, meaning it is more likely than not that at least one firm saw a decrease in productivity between 1999 and 2014.

8 Conclusion

The level and growth of productivity offer important insight into the functioning of an industry, from its economies of scale to important factors of growth and even its long-term viability. The railroad industry has changed dramatically since partial deregulation 1980, which occurred largely because of worries about regulation and its effect on productivity growth. Many studies have examined the growth in productivity immediately following the industry's partial deregulation, but few have looked at how productivity has changed since the industry became more stable in 1999. Following the massive changes that occurred through the 1980s and early 1990s, it is unlikely that firms will be able to continue pursuing broad changes like line abandonment to increase productivity. Instead, they need to turn towards improving technology and substituting inputs towards more productive factors to increase their productivity.

Unfortunately, existing models of productivity fail to account for technological change. In this paper, I develop a model that flexibly accounts for changes in productivity and technology and use it to decompose changes in productivity into those caused by innovation and those caused by other factors. To my knowledge, no published research exists that separately identifies productivity growth due to technological change and that due to broad non-innovative changes. Further, this model allows productivity and technology to evolve flexibly over time and can produce estimates of the level of productivity and its growth, which can inform key values in regulation. Finally, I discover a metric that determines whether firms have allocated additional resources towards factors that innovation makes more productive in order to realize further productivity gains. I apply my model to the railroad industry to investigate the recent change in productivity and to determine whether it is being driven by technological change or factors other than innovation.

I find that each Class I railroad has likely seen productivity growth since 1999, but the driving forces behind this growth differ. BNSF, KCS, Soo Line, and UP have seen large increases in its productivity due to technological change, by as much as 60% between 1999 and 2014. On the other hand, CN, KCS, and NS saw much slower growth induced by changing technology. Instead, these railroads relied on other methods to increase productivity, such as continuing to abandon unprofitable lines. The probability that firms substituted inputs towards factors that technological change makes more productive is about 50%, with no discernible pattern across firms or over time, indicating that railroads are not able to anticipate innovations or simply aren't adjusting inputs to take advantage of technological change. Finally, I perform Bayesian model selection and find that the model that allows for flexibility in both productivity and technology has the highest probability of being the true model. Using Bayesian model averaging, I find that each firm experienced modest growth in productivity between 1999 and 2014, with median estimates ranging from 0.433% to 2.636% per year.

9 Appendix

9.1 Model and Sampling Specifications

9.1.1 Deterministic Trend in Productivity, Constant Technology

As stated in the Empirical Models section, this model assumes productivity follows a deterministic trend that is shared across all firms. The production technology is common across firms and is constant through time. The model can be expressed in the following relations.

$$\begin{aligned} q_{it} &= \alpha_{it} + x_{it}\beta + \varphi_{it}\theta - \delta_i + \varepsilon_{it} \\ \alpha_{it} &= \alpha_i + \tau t \\ \alpha_i &\sim N(\mu_\alpha, \sigma_\alpha) \\ \varepsilon_{it} &\sim N(0, \sigma_\varepsilon) \\ \delta_i &\sim N^+(0, \sigma_\delta) \end{aligned}$$

I use Gibbs sampling to draw inference on this model and estimate the posterior distribution of the parameters conditional on the data. This distribution is the likelihood of the data conditional on the parameters and the following prior assumptions over the model parameters:

$$\begin{aligned} \beta, \theta, \mu_\alpha &\sim N(0, 5) \\ \tau &\sim N(0, 1) \\ \sigma_\alpha, \sigma_\varepsilon, \sigma_\delta &\sim \Gamma(1.5, 1) \end{aligned}$$

These assumptions were chosen to be diffuse with respect to their real world values. For example, input elasticities are rarely estimated to be greater than five,[20] which is just one

[20] As an example, Solow (1957) estimated the elasticity of capital to be 0.353.

standard deviation of the prior distribution.

Conditional on a value for δ_i, this is a linear random-effects model, which can be estimated via Gibbs sampling. To draw values of δ_i conditional on other parameters, first notice that

$$\begin{aligned} & p(\{\delta_i\}|\{\alpha_i\}, \mu_\alpha, \sigma_\alpha, \tau, \beta, \theta, \sigma_\varepsilon, \sigma_\delta; q, x, \varphi) \\ & \propto p(q|\{\alpha_i\}, \tau, \beta, \theta, \sigma_\varepsilon, \{\delta_i\}; x, \varphi) \times p(\{\delta_i\}|\sigma_\delta) \\ & \propto \prod_i p(q_i|\alpha_i, \tau, \beta, \theta, \sigma_\varepsilon, \delta_i; x_i, \varphi_i) \times p(\delta_i|\sigma_\delta). \end{aligned} \tag{23}$$

Thus, each δ_i can be drawn in its own independent block. Then, the conditional distribution of δ_i is

$$\begin{aligned} & p(\delta_i|\alpha_i, \tau, \beta, \theta, \sigma_\varepsilon, \delta_i; q_i, x_i, \varphi_i) \\ & \propto p(q_i|\alpha_i, \tau, \beta, \theta, \sigma_\varepsilon, \delta_i; x_i, \varphi_i) \times p(\delta_i|\sigma_\delta). \end{aligned} \tag{24}$$

Next,

$$\begin{aligned} & p(q_i|\alpha_i, \tau, \beta, \theta, \sigma_\varepsilon, \delta_i; x_i, \varphi_i) \\ & \propto \exp\left(-\frac{1}{2\sigma_\varepsilon^2}\sum_t (q_{it} - (\alpha_i + \tau t + x_{it}\beta + \varphi_{it}\theta - \delta_i))^2\right) \\ & \propto \exp\left(-\frac{1}{\sigma_\varepsilon^2}\sum_t (q_{it} - (\alpha_i + \tau t + x_{it}\beta + \varphi_{it}\theta))\delta_i - \frac{1}{2\sigma_\varepsilon^2}\sum_t \delta_i^2\right) \\ & = \exp\left(\delta_i\left(-\frac{1}{\sigma_\varepsilon^2}\sum_t (q_{it} - (\alpha_i + \tau t + x_{it}\beta + \varphi_{it}\theta))\right) + \delta_i^2\left(-\frac{T}{2\sigma_\varepsilon^2}\right)\right). \end{aligned} \tag{25}$$

Further, since $\delta_i|\sigma_\delta^2 \sim N^+(0, \sigma_\delta)$, the normalizing constant of this half-normal distribution does not depend on δ_i. Thus,

$$p(\delta_i|\sigma_\delta) \propto \exp\left(-\frac{1}{2\sigma_\delta^2}\delta_i^2\right); \qquad \delta_i \geq 0. \tag{26}$$

So,

$$p(\delta_i|\alpha_i, \tau, \beta, \theta, \sigma_\varepsilon, \delta_i; q_i, x_i, \varphi_i)$$
$$\propto \exp\left(\delta_i \left(-\frac{1}{\sigma_\varepsilon^2}\sum_t (q_{it} - (\alpha_i + \tau t + x_{it}\beta + \varphi_{it}\theta))\right) + \delta_i^2\left(-\frac{T}{2\sigma_\varepsilon^2} - \frac{1}{2\sigma_\delta^2}\right)\right). \tag{27}$$

This expression can then be factored so that

$$p(\delta_i|\alpha_i, \tau, \beta, \theta, \sigma_\varepsilon, \delta_i; q_i, x_i, \varphi_i) \propto \exp(-\frac{1}{2s^2}(\delta_i - m)^2),$$

where

$$m = -\frac{\sigma_\delta^2}{T\sigma_\delta^2 + \sigma_\varepsilon^2}\sum_t (q_{it} - (\alpha_i + \tau t + x_{it}\beta + \varphi_{it}\theta))$$
$$s^2 = \frac{\sigma_\delta^2 \sigma_\varepsilon^2}{T\sigma_\delta^2 + \sigma_\varepsilon^2}. \tag{28}$$

This is the kernel of a normal distribution with mean m and standard deviation s; thus, $\delta_i|\alpha_i, \tau, \beta, \theta, \sigma_\varepsilon, \delta_i; q_i, x_i, \varphi_i \sim N^+(m, s)$, so this block can be sampled via rejection sampling or by directly sampling from a truncated normal distribution.

The posterior distribution of the parameters was estimated using 10,000 warmup iterations to achieve convergence of the Markov chain and 100,000 iterations to sample the posterior distribution. Convergence was checked by examining trace plots and autocorrelation factors. Prior distributions were also varied to ensure prior assumptions weren't driving results.

9.1.2 Random Walk in Productivity, Constant Technology

As discussed in the Empirical Models section, this model allows productivity to follow a more flexible process, a random walk with drift. Each firm is allowed to have its own trend in its productivity process. The production technology is still assumed to be constant across firms and time. The model can be expressed in the following relations:

$$
\begin{aligned}
q_{it} &= \alpha_{it} + x_{it}\beta + \varphi_{it}\theta - \delta_i + \varepsilon_{it} \\
\alpha_{it} &= \alpha_{it-1} + \tau_i + \eta_{it} \quad ; \quad t > 0 \\
\alpha_{i0} &\sim N(\mu_\alpha, \sigma_\alpha) \\
\varepsilon_{it} &\sim N(0, \sigma_\varepsilon) \\
\delta_i &\sim N^+(0, \sigma_\delta) \\
\eta_{it} &\sim N(0, \sigma_\eta)
\end{aligned}
$$

I once again use a Gibbs sampler to draw values from the posterior distribution of the parameters conditional on the data. This is complicated by the random walk process in productivity, but the procedure is outlined in Sarris (1973). Samples of inefficiency terms δ_i conditional on other parameters are taken from a half-normal distribution as described in Section 9.1.1. The posterior distribution is also dependent on prior assumptions, which are given below.

$$
\begin{aligned}
\beta, \theta, \mu_\alpha &\sim N(0, 5) \\
\tau_i &\sim N(0, 1) \\
\sigma_\alpha, \sigma_\varepsilon, \sigma_\delta, \sigma_\eta &\sim \Gamma(1.5, 1)
\end{aligned}
$$

The posterior distribution of the parameters was estimated using 10,000 warmup iterations and 100,000 sampling iterations. Convergence was checked using previously described methods, and various prior distributions were tested.

9.1.3 Random Walk in Productivity and Technology

This model allows each firm's productivity as well as the parameters describing the production technology to follow a random walk with drift. The production technology is assumed to be shared across firms, but is allowed to follow a flexible process over time. The model is expressed in the following relations:

$$
\begin{aligned}
q_{it} &= \alpha_{it} + x_{it}\beta_t + \varphi_{it}\theta - \delta_i + \varepsilon_{it} \\
\alpha_{it} &= \alpha_{it-1} + \tau_i + \eta_{it} \quad ; \quad t > 0 \\
\beta_t &= \beta_{t-1} + \rho + \psi_t \quad ; \quad t > 0 \\
\alpha_{i0} &\sim N(\mu_\alpha, \sigma_\alpha) \\
\beta_0 &\sim N(\mu_\beta, \Sigma_\beta) \\
\varepsilon_{it} &\sim N(0, \sigma_\varepsilon) \\
\delta_i &\sim N^+(0, \sigma_\delta) \\
\eta_{it} &\sim N(0, \sigma_\eta) \\
\psi_t &\sim N(0, \Sigma_\psi)
\end{aligned}
$$

I assume that Σ_ψ and Σ_β are diagonal and label the kth diagonal element of each ${\Sigma_\psi^{kk}}^2$ and ${\Sigma_\beta^{kk}}^2$. Once again, I draw inefficiency conditional on other parameters using the method described in Section 9.1.1. I use a Gibbs sampler with 5,000 warmup iterations and 50,000 sampling iterations to draw inference on the parameters. Prior assumptions over the parameters are given below.

$$
\begin{aligned}
\theta, \mu_\alpha, \mu_\beta &\sim N(0, 5) \\
\tau_i &\sim N(0, 1) \\
\rho &\sim N(0, I) \\
\sigma_\alpha, \Sigma_\beta^{kk}, \sigma_\varepsilon, \sigma_\delta, \sigma_\eta, \Sigma_\psi^{kk} &\sim \Gamma(1.5, 1)
\end{aligned}
$$

References

AAR (2015). 2015 outlook report. *American Association of Railroads*.

Aigner, D., C. K. Lovell, and P. Schmidt (1977). Formulation and estimation of stochastic frontier production function models. *Journal of Econometrics 6*, 21–37.

Barnekov, C. and A. Kleit (1990). The efficiency effects of railroad deregulation in the united states. *International Journal of Transport Economics / Rivista internazionale di economia dei trasporti 17*(1), 21–36.

Bitzan, J. D. and T. E. Keeler (2003). Productivity growth and some of its determinants in the deregulated us railroad industry. *Southern Economic Journal*, 232–253.

Bitzan, J. D. and T. E. Keeler (2007, February). Economies of density and regulatory change in the u.s. railroad freight industry. *The Journal of Law and Economics 50*(1), 157–180.

Bitzan, J. D. and W. W. Wilson (2007). Industry costs and consolidation: efficiency gains and mergers in the us railroad industry. *Review of Industrial Organization 30*(2), 81–105.

Caves, D. W. and L. R. Christensen (1980, October). The relative efficiency of public and private firms in a competitive environment: The case of canadian railroads. *Journal of Political Economy 88*(5).

Caves, D. W., L. R. Christensen, and J. A. Swanson (1981, December). Productivity growth, scale economies, and capacity utilization in u.s. railroads, 1955-74. *The American Economic Review 71*(5), 994–1002.

Chib, S. and I. Jeliazkov (2001). Marginal likelihood from the metropolis–hastings output. *Journal of the American Statistical Association 96*(453), 270–281.

Christensen, L. R., D. W. Jorgenson, and L. J. Lau (1973, February). Transcendental logarithmic production frontiers. *The Review of Economics and Statistics 55*(1), 28–45.

Del Negro, M. and C. Otrok (2008). Dynamic factor models with time-varying parameters: measuring changes in international business cycles. *FRB of New York Staff Report* (326).

Eisenbeis, R. A., G. D. Ferrier, and S. H. Kwan (1999, December). The informativeness of stochastic frontier and programming frontier efficiency scores: Cost efficiency and other measures of bank holding company performance. *Federal Reserve Bank of Atlanta Working Paper Series 99*(23).

Friesner, D., R. Mittelhammer, R. Rosenmane, et al. (2006). Inferring the latent incidence of inefficiency from dea estimates and bayesian priors. Technical report.

Jondrow, J., C. K. Lovell, I. S. Materov, and P. Schmidt (1982, August). On the estimation of technical inefficiency in the stochastic frontier production function model. *Journal of Econometrics 19*(2-3), 233–238.

Jorgenson, D. W. and Z. Griliches (1967, July). The explanation of productivity change. *The Review of Economic Studies 34*(3), 249–283.

Koop, G., M. J. Steel, and J. Osiewalski (1995). Posterior analysis of stochastic frontier models using gibbs sampling. *Computational Statistics 10*(4), 353–373.

Kumbhakar, S. C. (1988, November). On the estimation of technical and allocative inefficiency using stochastic frontier ffunction: The case of u.s. class 1 railroads. *International Economic Review 29*(4), 727–743.

Leybourne, S. J. (1993). Empirical performance of the aids model: Constant-coefficient versus time-varying-coefficient approaches. *Applied Economics 25*, 453–463.

Lim, S. H. and C. K. Lovell (2009, October). Profit and productivity of u.s. class i railroads. *Managerial and Decision Economics 30*(7), 423–442.

Mazzocchi, M. (2003). Time-varying coefficient in the almost ideal demand system: An empirical appraisal. *European Review of Agricultural Economics*.

Oum, T. H. and C. Yu (1994, May). Economic efficiency of railways and implications for public policy. *Journal of Transportation Economics*, 121–138.

Park, B. U. and L. Simar (1994). Efficient semiparametric estimation in a stochastic frontier model. *Journal of the American Statistical Association 89*(427), 929–936.

Rose, N. L. and P. L. Joskow (1990). The diffusion of new technologies: Evidence from the electric utility industry. *The RAND Journal of Economics 21*(3), 354–373.

Sarris, A. H. (1973). Kalman filter models: A bayesian approach to estimation of time-varying regression coefficients. *Annals of Economic and Social Measurement 2*(4), 501–523.

Schmidt, P. and C. K. Lovell (1979). Estimating technical and allocative inefficiency relative to stochastic production and cost frontiers. *Journal of Econometrics 9*, 343–366.

Solow, R. M. (1957, August). Technical change and aggregate production function. *The Review of Economics and Statistics 39*(3), 312–320.

Spady, R. H. and A. F. Friedlaender (1978). Hedonic cost functions for the regulated trucking industry. *The Bell Journal of Economics*, 159–179.

Transportation Research Board (2015). *Modernizing Freight Rail Regulation.* Transportation Research Board.

Tretheway, M. W., W. G. W. II, and A. K. Fok (1997, January). The total factor productivity of the canadian railways, 1956-91. *Journal of Transportation Economics and Policy 31*(1), 93–113.

United States Surface Transportation Board (2015). Complete r-1 railroad annual reports. Technical report.

van Den Broeck, J., G. Koop, J. Osiewalski, and M. J. Steel (1994). Stochastic frontier models: A bayesian perspective. *Journal of Econometrics 61*(2), 273–303.

Wilson, W. W. (1994, March). Market-specific effects of rail deregulation. *The Journal of Industrial Economics 42*(1), 1–22.

Wilson, W. W. (1997). Cost savings and productivity in the railroad industry. *Journal of Regulatory Economics 11*, 21–40.

Winston, C., T. M. Corsi, C. M. Grimm, and C. A. Evans (1990). *The Economic Effects of Surface Freight Deregulation.* The Brookings Institution.

Yan, J., X. Sun, and J. J. Liu (2009). Assessing container operator efficiency with heterogenous and time-varying production frontiers. *Transportation Research Part B 43*, 172–185.