

Investment and misallocation in infrastructure networks: The case of U.S. natural gas pipelines

Philip Solimine¹ and Paul Schrimpf¹

¹*The University of British Columbia*

October 17, 2024

[Click here for the most recent version]

Abstract

This paper investigates regulatory distortion in the incentives for firms to invest in expanding transmission capacity in the United States natural gas pipeline network. Transmission rates for interstate gas pipelines are tightly controlled by federal regulators, who set them so that firms earn a fixed rate of return on capital under projected demand and cost conditions. This rate of return regulation removes the incentive of pipeline operators to exercise local market power by withholding capacity; however, it also distorts firms' incentives to invest in expanding the network. Misallocated capital in a transmission network can lead to inefficiencies and congestion. To ensure that new capital will be desirable, the regulator subjects new capital investment to a stringent approval process. Leveraging a detailed dataset of pipeline regulatory filings, we estimate a dynamic model of the firms' investment incentives using a debiased nonparametric machine learning approach. We then construct and estimate a structural measure of the marginal social value of capital that is based on a dynamic model of optimal network investment by a social planner. This measure ties the social value of pipeline capacity to differences in prices across state borders that exceed the marginal cost of transmission, indicating excess demand. We find that in most areas, the incentives of firms to invest in the pipeline network under fixed rates of return exceed the social value of capital. This suggests that some costly approval process surrounding pipeline investment is indeed necessary to realign firms' incentives to expand the network. While overall the implied investment costs have been close to optimal, at a disaggregated level we find evidence of some systematic deviations from the optimal policy both spatially and intertemporally. We suggest that a welfare-improving reallocation of regulatory costs would be one that streamlines investment approval in the northeast but increases regulatory stringency in the southeast and parts of the west.

1 Introduction

This paper examines how investment approval, in tandem with price regulation, shapes firms' incentives to invest in expanding infrastructure for resource transmission between regional energy markets. In the United States, natural gas pipelines provide an ideal setting for this investigation; as a classic example of a natural monopoly, interstate gas transmission is tightly regulated by the Federal Energy Regulatory Commission (FERC). These nearly three million miles of pipeline infrastructure transport trillions of cubic feet of gas every year¹, and the resource accounts for more than 40% of electricity generation and

¹<https://www.eia.gov/energyexplained/natural-gas/natural-gas-pipelines.php>, Archived version; October 07, 2024

over a third of total energy consumption in the U.S.^{2 3} FERC regulations require that pipeline owners provide open access to unbundled transportation services at fixed, regulated rates. Natural gas marketers contract with pipeline owners at these regulated rates for rights to pipeline capacity, and then purchase gas from producers, trade amongst themselves, and sell it to consumers, all at unregulated prices. This system of separating transmission operators from unregulated marketers allows market based prices to be used to efficiently allocate gas throughout the country, and reduces the incentive for pipeline owners to exercise monopoly power by withholding idle capacity.

Although this system of regulated rates for transmission and market-based prices for allocation leads to a more efficient utilization of the existing transmission network, flat transmission rates across the country may not properly encourage firms to invest in areas where capacity is most needed. Because regulated transmission rates are insulated from market prices, firms face distorted incentives to invest in expanding pipeline capacity. FERC sets transmission rates using a fixed rate-of-return regulation; rates are set to cover pipeline's costs plus a reasonable rate of return on their capital. In addition to de-incentivizing firms from investing in regions with the highest demand, this rate of return regulation may lead pipeline firms to overinvest in capital more generally. This is because increasing capital directly increases rates the pipeline can charge—leading to an excess return from capital. This incentive problem is called the Averch and Johnson (1962) effect. In part to mitigate this possibility, pipelines must apply to FERC for approval for any investment project. This paper quantifies the investment incentives of pipelines created by this system of regulation. Using detailed pipeline accounting information, we estimate pipelines' expected marginal profit from capital. We then use pipelines' investment Euler equations to recover the marginal cost of investment that rationalizes pipelines' observed investment decisions. We interpret the estimated marginal cost of investment (beyond the accounting cost of investment) as the shadow cost of meeting regulatory constraints. We find that the expected marginal profit of capital is about 10% per year. This relatively high rate of return is comparable with previous estimates (e.g. Von Hirschhausen, 2008) and consistent with FERC guidelines for setting transmission rates. More notably, we find relatively little geographic or temporal variation in the return to capital. Actual investment does show considerable geographic variation. Our estimates imply that this variation in investment is largely driven by variation in the pipeline approval process.

To analyze the efficiency of the observed investment allocation, we develop a model of a social planner choosing investment. We use the model to compare a social planner's investment incentives with the estimated investment incentives of pipelines. We extend the theoretical framework of Cremer, Gasmi, and Laffont (2003), to show that a social planner sets the marginal cost of capacity equal to the difference in market prices across nodes in the pipeline network when the capacity constraint binds, and derive a structural measure of the marginal social value of capital. This measure is derived the planner's Euler equation, which implies that the marginal cost of investment should equal the marginal change in capacity times the shadow price of relaxing capacity constraints. Recall that while pipelines are regulated as a natural monopoly, there are many gas marketers that buy and sell gas at various locations throughout the pipeline network at unregulated prices. Assuming that the gas trading market is competitive, the social planner's shadow price of capacity equals the difference in observed market prices. Thus, we can use market prices and data on the relationship between capital and capacity to estimate the components of the planner's Euler equation. We use these estimates to check whether the observed investment satisfies the Euler equation, a necessary condition for optimality.

Comparing our estimates of the social value of investment and pipelines' investment incentives, we draw two main conclusions. First, the allowed rate of return to capital is high compared to the marginal

²<https://www.eia.gov/tools/faqs/faq.php?id=427&t=3>, Archived version; October 06, 2024

³<https://www.eia.gov/energyexplained/natural-gas/use-of-natural-gas.php>, Archived version; October 07, 2024

social value of investment. This suggests that a stringent approval process is important for avoiding over-investment. Second, there is clear and persistent geographic variation in the marginal social value of investment, but our estimates of pipelines marginal regulatory investment cost do not have as clear of a geographic pattern. FERC may be able to improve the allocation of investment by streamlining investment applications in the northeast, and being more stringent toward applications in the south. This is evidenced by the lack of investment in the northeast relative to the outsized marginal product of capital earned by firms and the large periodic price divergences that characterize gas prices in the region. On the other hand, states in the southeast see the opposite—firms invest beyond the socially optimal level. Underregulation of investment is made evident by tightly integrated prices compared to the relatively large amount of investment in the region, suggesting that the social value of investment in this region is reasonably low. This misallocation of capital across the network has become particularly pronounced as the shifting landscape of gas production in the United States has moved away from the Gulf of Mexico and toward shale drilling in the northeast.

The remainder of this paper is organized as follows: the rest of Section 1 contains a more in-depth examination of the relevant literature, and an overview of the structure of the natural gas market. This section will also provide some background on the evidence around rate-of-return regulation and its alternatives, how congestion is measured, and the details of the regulatory approval procedure. Section 2, describes the data we use in this study, and provides some stylized facts about how the market and pipeline network have changed over time. In Section 3, we lay out a theoretical framework to model investment by the pipeline, and extend the model of Cremer, Gasmi, and Laffont (2003) to characterize optimal investment and the social value of capital. Section 4 lays out our identification arguments, along with our nonparametric, debiased machine learning estimation approach. Section 5 presents the results of these estimations, and Section 6 concludes.

1.1 Related literature

Our work contributes to the understanding of natural gas infrastructure, resource transmission networks, and more broadly, the economics of transportation networks. Methodologically, our paper is related to the literature on estimating Euler equations, semiparametric estimation, and debiased machine learning for structural models.

Economics of natural gas infrastructure Natural gas is a crucial industry in North America, and there is a large body of work investigating market structure and regulation policy in this setting, beginning with Wellisz (1963). Our study is primarily concerned with the regulatory regime that came into place following FERC’s 1996 restructuring, which functionally separated gas transmission from marketing and implemented fixed rates of return for transmission based on pipelines’ cost of service. Oliver and Mason (2018) gives an overview of the history of regulations in these markets, which we will summarize in section 1.2. Oliver (2019) took a theoretical approach to examine firm investment behavior under fixed rate of return regulation. In simulations, they found that relaxing fixed rates to a price cap would marginally improve outcomes. Adamson (2018) compared the structure of regulations across the electric and natural gas industries in the US, emphasizing the relative strength of pipeline network growth under the current regime when compared to the electric grid.

Hausman and Muehlenbachs (2019) studied price regulation as well, and found some evidence of the Averch-Johnson effect on the intensive margin, as firms trade off repair costs for capital investment. While FERC maintains a functional no-conduit policy that separates employees by function into marketing and transmission, forbidding communication of private information between the two groups, Marks et al. (2017) found some evidence that vertically integrated gas companies can withhold capacity, by

scheduling more than they will need and releasing it too late to be repurchased. While this behavior is not illegal, it underscores the importance of understanding market structure, regulation policy, and firm incentives in this setting. In a relevant study, Brito and Rosellón (2011) characterized the optimal regulation policy for transmission pricing, emphasizing how the simple economics of pipeline network expansion lead to a relatively simple characterization of the optimal policy for price regulation. However, they emphasize that the inelastic demand for natural gas, which plays a central role in this characterization, can also lead to large price spikes following relatively small shock to demand.

In another group of studies, works collected in MacAvoy (2007) conducted a number of empirical studies of how the natural gas market had fared since its partial deregulation in 1996. As part of this line of work, Marmer, Shapiro, and MacAvoy (2007) characterized how persistent price differences between connected locations can be used to identify areas of congestion in the pipeline network. In particular, they found that the northeast and midwest regions, along with California, were not well integrated. Similarly, later studies by Oliver, Mason, and Finnoff (2014) and Avalos, Fitzgerald, and Rucker (2016) found evidence of bottlenecks and congestion both in a specific pipeline route through the Rocky Mountain region as well as into California and Florida. We will show that while there has been substantial investment in the pipeline network, this new investment has failed to integrate the northeast into the rest of the network, suggesting that new capital investment has not been placed in regions where it is most needed.

Recent work has also emphasized the importance of understanding regulation policy to the efficiency of energy transitions. Scott (2023) highlighted how expansions to pipelines encourage generators to switch from natural gas to coal, while Gowrisankaran, Langer, and Reguant (2024) highlights how fixed rate regulation may not properly incentivize pipeline companies to develop new pipeline capital in the areas where it would be useful to this energy transition. Prior to these, some other papers also investigated the pipeline network, with an eye on efficiency and design. These include Cremer and Laffont (2002); Cremer, Gasmí, and Laffont (2003), whose theoretical approach characterized how the social value of pipeline capacity is tied to price differences across regions. Many studies have also examined the design of the network and regulation policy particularly in the European context, where investment in international gas transmission is more carefully planned. These include Ejarque (2011); Cobanlı and Hubert (2014), and Neumann, Rosellón, and Weigt (2015). Dott (1999) adopted more of an engineering design approach, characterizing engineering details such as the optimal placement and density of central facilities in the design of the network.

Economics of transportation networks Our work contributes to the literature on transportation networks, which dates back to Meyer (1959) and Vickrey (1969). Recent work in this area includes Fajgelbaum and Schaal (2020), who developed a spatial equilibrium model of optimal transportation networks. Their model proceeds in a way that is similar to ours—in particular, prices arise as Lagrange multipliers in a static social planner problem, which is nested in a model of optimal network growth. Our model proceeds in a similar manner but with a number of important differences; primarily that we allow the rate regulation of interstate pipelines to distort firm incentives and decouple profits from the structure of the network, potentially driving a wedge between profits and social value that could lead to excess returns and misallocation of capital investment.

In the space of recent empirical work, a large number of recent papers have examined optimal development and congestion in road networks. Allen and Arkolakis (2022) developed a gravity-type model of transportation with congestion in order to discuss counterfactual network design policies for the U.S. highway system and Seattle road network. Jaworski and Kitchens (2019) focused on the interaction between regulation and incentives in the highway network and its effects on integrating regional markets

highlighting how regionally targeted policies focused on network improvement can reduce inequality and improve welfare through connectivity. Jaworski, Kitchens, and Nigai (2023) extended a similar model to the U.S. highway system and studied its impacts on globalization and welfare gains from trade. However, the setting of road networks differs from coupled energy markets in a number of ways, in part because the road network setting requires careful attention to equilibrium between multiproduct firms. Furthermore, since the highway network was developed by planners, there is no general friction between social value and the incentives of the network designers. Since transportation through trade networks can be expensive, these papers also spend significant effort on understanding the structure and role of transportation costs. On the contrary, marginal costs of transmission in energy markets are generally very low with capacity and congestion taking more of a central role in the design problem. In energy markets, recent work by Preonas (2024) has examined market structure in transportation. They focus on coal, which is transported primarily by railroad, and highlight how market power in transportation can create inefficiencies; in this case, by limiting the impact of policies to reduce carbon emissions.

Recent studies have also extended similar models to study the impact of connectivity on economic outcomes, along with network design and optimality, in a broad class of settings with varying cost structures and regulatory environments. For example, in the contexts of internet and communication infrastructure (Caoui and Steck, 2023; Goldstein, 2024), rail networks (Keeler, 1974; Degiovanni and Yang, 2023; Chen, 2024), ports and cargo shipping (Brancaccio, Kalouptsidi, and Papageorgiou, 2020; Ducruet et al., 2024), air transportation (Yuan and Barwick, 2024), and beyond. A particularly relevant example is Brancaccio, Kalouptsidi, and Papageorgiou (2024), who showed the value of strategically targeting infrastructure investment in the setting of ports and oceanic transportation of commodities. Redding and Turner (2015) provide comprehensive reviews of some literature in these areas. In particular, the setting of energy transmission and particularly natural gas pipelines differs from these in a number of ways, first through the single-product nature of the transmission network, along with the specific forms of rate regulation used in these industries. Our primary contribution to this literature is characterizing the wedge between social value and firm profits created by cost-of-service regulation, and of how regulators can use a stringent approval process to realign these incentives and produce socially optimal network growth.

Estimation from Euler equations Estimation of Euler equations has a long history in economics, dating back at least to Hansen and Singleton (1982). Early estimates of Euler equations used parametric specifications of preferences. Recent contributions have focused on nonparametric identification and inference. Escanciano et al. (2021) show nonparametric identification of marginal utilities in the consumption-based asset pricing Euler equation. Our paper considers investment Euler equations and identification of marginal cost functions follows from similar arguments. Escanciano et al. (2021) proposes a kernel density based estimator and derive the convergence rate of their estimator and show \sqrt{n} asymptotic normality of functionals of the estimator. Cui, Hong, and Li (2021) analyze a spline estimator for consumption Euler equations and derive its convergence rate and pointwise asymptotic distribution. Our estimator uses a combination of a reproducing kernel Hilbert space embedding to estimate a conditional expectation operator and neural networks to estimate marginal cost functions. We focus on inference on low dimensional functionals of the estimands. We use the automatic debiased machine learning approach of Chernozhukov et al. (2021) for estimation and inference on our functionals of interest. Like the estimation method of Escanciano et al. (2021), automatic debiased machine learning ensures that our estimates are \sqrt{n} asymptotically normal. However, automatic debiased machine learning considerably weakens to assumptions placed on the nonparametric conditional expectation and marginal cost estimators. Moreover, debiasing of our functionals of interest, ensures that the parameters of interest are estimated with efficient influence, so our estimates attain the semiparametric efficiency bound.

Semiparametric estimation for structural models All of our estimation results are completely non-parametric. We place no parametric restrictions on the cost or profit functions of firms, nor on the distribution of any unobservables. This is contrast to most empirical structural papers, which typically impose parametric restrictions. To accomplish this nonparametric estimation, we make use of recent developments in debiased machine learning by Chernozhukov et al. (2021) (which builds on Chernozhukov, Newey, and Singh (2022), Chernozhukov et al. (2022) and earlier work by the same authors) to estimate functionals of firm and regulator incentives. The estimation approach consists of three steps. First, various conditional expectation functions and conditional expectation operators are flexibly estimated using neural networks and reproducing kernel Hilbert space embeddings. Next, Riesz representators for the derivative of our functionals of interest are flexibly estimated. Third, the Riesz representor and original functional of interest are combined to give an orthogonal moment, which is used for inference. Among our functionals of interest, are both linear (such as the average expected marginal product of capital) and nonlinear (such as the average marginal cost) functionals. This paper is among the first to apply debiased machine learning to a structural model. It is also among the first to apply *automatic* debiased machine learning to nonlinear functionals.

1.2 Background on natural gas pipelines

Natural gas transmission is widely considered to be a straightforward case of natural monopoly. Ever since its implementation, the cost-of-service method that is used to regulate pricing has been the subject of economic scrutiny.

Oliver and Mason (2018) provides a detailed account of the rich history of the US federal regulation in the natural gas pipeline industry. They explain that the current system of transmission price regulation is the result of substantial restructuring of the natural gas market. This restructuring began in 1985 with FERC Order No. 436. The order was the result of years-long efforts to restructure the natural gas market, which was preceded by an era of ‘wellhead’ price controls that were managed by the Federal Power Commission (FPC). Wellhead pricing was initially spatially disaggregated, but became increasingly coarse over time due to the practical costs of its implementation. This attempt to integrate prices efficiently was widely considered an abject failure, and the FPC was dissolved in 1978. The purpose of Order 436 was to encourage the decoupling of the market for transmission from that for the sale of the final product; moving away from the “merchant carrier” model and into the current era of common carriers and independently operating natural gas marketers.

Following the success of Order 436, FERC’s next move in shaping the contemporary system was to implement FERC Order No. 636. Order 636 reinforced the restructuring initiated by Order 436, by disallowing any bundling of transmission services with the sale of the final product. Its other innovation was introducing cost-of-service rate setting through the Straight Fixed-Variable (SFV) method. The SFV method requires a strict delineation of fixed and variable costs of pipeline operation to determine individual rates of return for capacity reservation and transmission. This new method of price control is credited with driving a stark convergence in nationwide natural gas prices which followed immediately from its implementation. Together, these orders form the contemporary regulatory framework for the pricing of natural gas transmission through the pipeline network, the backbone of what has been considered to be a largely successful restructuring of the natural gas market.

A key feature of natural gas pipelines is the separation between pipeline owners and gas marketers. As required by FERC, pipeline owners sell their services at regulated rates and must treat all customers the same. Pipelines sell two main forms of transmission service. Firm transportation service contracts sell the guaranteed right to transmit a certain volume of gas per day. These contracts are usually long term, lasting well over five years. These contracts typically specify a large fixed price per dekatherm of

gas reserved, and a small, sometimes zero, additional charge per dekatherm of gas actually transmitted.

Most pipelines sell firm transportation service contracts for all of their available capacity. However, holders of firm transportation service often do not use all of their reserved capacity. In this case, the pipeline can sell the unused as well as any unreserved capacity as interruptible transportation service. Interruptible transportation service is sold in spot markets through pipeline companies' websites.

FERC uses cost of service regulation to set the maximum price of pipeline services. FERC designs the rates to cover pipelines' operating costs and allow a reasonable rate of return on investment. The large reservation price and small additional transmission price of firm transportation service reflects the fact that the majority of pipeline costs are fixed costs. According to its guidelines, FERC typically sets the price of interruptible service to the sum of the reservation and transmission prices of firm service.

Pipelines are allowed to negotiate rates lower than the maximum set by FERC. However, this rarely occurs. In addition to regulating prices, FERC regulates pipeline construction, decommission, and sales. Under section 7(c) of the Natural Gas Act, companies must apply for FERC approval before building, decommission, buying, or selling any interstate pipelines and related infrastructure. The entire process of planning and building a pipeline takes one to three years. A company interested in building a pipeline begins by holding a (generally nonbinding) "open season exercise" to solicit buyers of firm transportation service. The company will try to obtain binding agreements to purchase long-term (5-10) firm transportation service contracts if the proposed pipeline is built. If there is sufficient interest, the company then files an application of public convenience and necessity with FERC. FERC reviews the application, holds public hearings, does an environmental assessment, and eventually rejects or accepts, often with some conditions, the application. The actual construction of a new pipeline takes about six months. FERC accepts nearly all section 7(c) applications. A complete diagram of the U.S. natural gas pipeline network is shown in Figure 1.

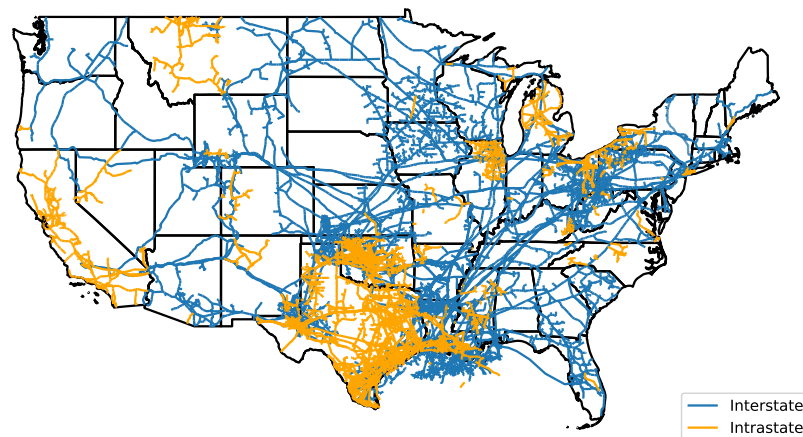


Figure 1: Map of the U.S. natural gas pipeline network

Pipeline services are primarily purchased by natural gas marketers. These marketers purchase gas from producers, transmit it through pipelines, and sell it to consumers. Consumers of natural gas require differing levels of reliability of service and feature different variability of volume. For example, local distribution companies, who provide gas to homes for heating and cooking, require highly reliable delivery of a varying volume of gas. The variation of volume is partly predictable—for example due to seasonality—but also unpredictable—for example due to the weather. As climate change continues to produce unpredictable and extreme weather patterns, this unpredictable component will only continue to increase. Industrial users of natural gas may demand a more constant volume and might be capable

of tolerating some delivery interruptions. Marketers combine the various services of multiple pipelines to meet the varying needs of end-users.

During off-peak times, gas can be injected into underground storage facilities. Pipeline companies often own storage facilities as well. Storage facilities help to meet peak demand. Gas is extracted when demand is high. Liquefied natural gas facilities can also be used to meet peak demand, but they are primarily built for export.

1.3 Rate regulation and misallocation

Rates for gas transmission are set according to a cost-of-service policy that allows investors a fixed rate of return on their capital. Capital depreciates according to a straight-line method. This means that in the absence of new capital investment, pipelines fully depreciate to zero capital within a finite, predetermined time horizon (usually around 25 years). Rates are set after observing depreciation and additional investment, meaning that firms can increase their rates through any investment that increases capital.

Oliver (2019) theoretically examined how the existing regulatory pricing rules (specifically the SFV method) may distort firms' incentives to invest in the development of the transmission network. Using their theoretical model, they then conducted computational experiments to test for differences in optimal network investment under varying regulatory regimes. They evaluate a policy counterfactual that tweaks transmission pricing rules to allow for moderately more flexibility in pricing; introducing a single price cap on capacity reservation and transmission in contrast the current SFV standard of fixing each price separately. In their model, this partial relaxation of the price control could create welfare gains and maintain the ability of the regulation to control market power. However, such policies are sensitive to implementation and have proven difficult to successfully implement in practice Cox and Isaac (1987); Isaac (1991). The prevailing consensus is that the current cost-of-service system produces reasonably efficient outcomes and there is no immediate need for change overall. We argue that part of the reason the existing cost-of-service setup has been so successful can be attributed to the costly approval process creating barriers for new investment.

On the other hand, there is increasing empirical evidence that some aspects of the existing regulatory regime could be improved. For example, Marks et al. (2017) find that pipeline companies have an incentive to downschedule capacity—to oversubscribe and subsequently release the excess capacity too late for it to sell on the secondary market. While pipeline firms have denied that they engage in this behavior,⁴ the authors estimate that the practice had cost New England ratepayers dramatically; to the tune of \$3.6 billion over a three year period. They suggest that reducing regulation to allow additional expansions to the network may not be as important as tweaking aspects of the current regulatory design. In agreement with their findings, we find evidence that the process of regulatory approval was close to optimal—with the possible exception of the northeast region.

Once commonly cited conceptual issue with cost-of-service regulation is that it can create perverse incentives for firms to overinvest in capital expenditure. This type of distortion is called the Averch-Johnson effect Averch and Johnson (1962). However, theoretical and empirical studies searching for this effect have come up short (Dechert, 1984; Joskow, 2005). Our results suggest that such effects may exist, but are masked or mostly eliminated by the additional regulatory cost of added capital.

Other previous work has also found mixed results regarding the effects of cost-of-service regulation on network infrastructure investment in the natural gas pipeline. For example, by comparing weighted-average costs of capital to benchmark levels, Von Hirschhausen (2008) concludes that the evidence does not support the hypothesis that fixed rates of return reduced overall investment in the pipeline. On the

⁴<https://www.utilitydive.com/news/eversource-avangrid-artificially-constrained-gas-pipeline-capacity-for-year/507018/>

contrary, they suggest that there is some evidence pointing toward these firms overinvesting in pipelines, in response to the additional marginal value of capital provided by cost-of-service.

In another particularly relevant study, the idea that overcapitalization may be appearing in a way that is partially masked from aggregated patterns but becomes more clear on the right margin. In particular, this is supported by Hausman and Muehlenbachs (2019), who found empirical evidence that regulated rates lead pipeline firms to overcapitalize on the intensive margin, by trading off repair costs for replacements that increase capital stock. Critically, they find that the trade-off away from repair and toward capital has led to a massive excess of damaging methane leaks. They emphasize that the problem is so severe that repair costs were substantially less than the commodity value of leaked gas during their period of observation. However, they also suggest that changes to safety regulations, which took place in 2010, were at least partially effective at mitigating the issue. Similarly, Gowrisankaran, Langer, and Reguant (2024) found that the similar regulatory environment of the electric grid has encouraged overinvestment in coal infrastructure, limiting the progress of utilities in transitioning away from coal and toward natural gas-based generation. We suggest another avenue where such an effect may appear, which is through a spatial misallocation of capital expenditure by overspending on capital investment in areas where it is not as necessary, while neglecting areas with a higher social value. Are results show how investment approval regulation is crucial to realign incentives to invest in regions with persistent price gaps. We recommend adjustments to the spatial targeting of investment regulation stringency that could help to realign firm's incentives at this level.

1.4 Congestion in the pipeline network

Beyond a simple understanding of how the price regulation will increase or decrease investment, we are interested in how firms may strategically target their investment to maximize their profit in the face of regulation, and how this targeting impacts the efficiency of the network structure. In order to understand this, it is important to understand how congestion is measured in the pipeline network. The study of congestion in networked transmission dates back to Vickrey (1969), who studied bottlenecks in both transportation and public utilities. They define bottlenecks as areas of the network where capacity constraints on links are substantially lower than up- and down-stream demand, and recognized that “the information provided by a system of congestion control through pricing has an essential role to play in the evaluation of investments designed to afford relief from congestion.” A framework for measuring congestion in the pipeline network was developed by Marmer, Shapiro, and MacAvoy (2007). This method uses deviations in regional spot market prices to identify bottlenecks in the transmission network. Their results suggest that the network structure contributes to a segmented market, characterized by connected prices within region, but a marked failure of the network to cointegrate these prices into a single common market at the national scale. We will provide a microfoundation for this method by explicitly tying price divergences to the social value of capital, and suggest policy interventions that could better align firm incentives with the goal of integrating nationwide prices.

Following FERC's major restructuring of the natural gas market in 1996, a number of studies emerged addressing the properties and practical feasibility of transmission pricing based on congestion, which aim to price-in marginal externalities of network usage. Nasser (1999) treats this framework theoretically, and places its efficiency on a natural scale between perfect competition and monopoly, with monopoly operators able to raise profits by restricting network supply. Rosenberg (2000) explains that by the start of the 21st century, regional operators in New York and California had already begun implementing congestion-based pricing. For firms who relied on predictable prices, the option of a long-term contract for firm's transmission rights would provide a hedge against the fluctuations induced by congestion pricing. However, long-term transmission contracts were not without controversy; FERC was skeptical of

these arrangements from the beginning (Rosenberg, 2000), as it was feared that they may introduce a mechanism for firms to artificially restrict supply (Bushnell, 1999).

MacAvoy (2007) further develops the idea that price differentials between start nodes and end nodes can be used to identify features of the market structure. A notable takeaway from their book is the role of nodal storage capacity in reducing the strain of congestion on spot-price arbitrage. Oliver, Mason, and Finnoff (2014) studies the interaction between storage and basis price differentials in detail, and explains how these long-term contracts for capacity have lead to a thriving, largely unregulated secondary market for unused capacity. They show how this friction, along with the limits on transmission price that manifest as bottlenecks and congestion in supply, leads to an inefficient wedge between hub prices. They go on to show that this wedge is driven by congestion increasing transmission prices in the secondary market. Our results provide a demonstration that firms do not necessarily invest in ways that reduce congestion, and in fact that their incentives do not generally align with this objective.

The theoretical frameworks used in many of these previous papers comes from Cremer and Laffont (2002) and Cremer, Gasmi, and Laffont (2003). Cremer and Laffont (2002) develops a theory to understand how transmission rights lead to a natural monopoly setting, in which firms can exploit market power by artificially restricting supply to raise prices. As a solution, they study the effect of a regulation that adds excess capacity to the pipeline, and find that this would be theoretically effective in mitigating this market power. Cremer, Gasmi, and Laffont (2003) further develops this framework by characterizing optimal consumer, producer, and transmission prices in a competitive network. This results in a natural optimality condition, in which the price differential between nodes is equal to the marginal cost of transmission; a result we will use to test whether the structure and capacity of the network are optimal. As part of this paper, we will extend some of these theoretical results to account for a fully general pipeline network with arbitrary structure.

1.5 Pipeline development procedure

Pipeline operators looking to invest in developing the infrastructure or increasing capacity must follow a time-consuming and potentially costly procedure to obtain regulatory approval. Figure 2, reproduced from the website for the US Energy Information Administration (EIA)⁵ explains the procedure of obtaining permits and approval that precedes breaking ground on pipeline development or expansion projects. The procedure begins with a 1-2 month “open season exercise”, where the pipeline developer allows potentially interested firms to sign up to reserve a portion of the new capacity. These agreements are nonbinding, but signal to the developer whether there is enough demand for this new capacity to warrant their investment.

If the public announcement and open season signal enough interest, then the developer will continue, otherwise, they will either place the project on hold indefinitely until demand shifts in their favor, or cancel the project altogether. If the project proceeds, the next step is to develop plans for the pipeline design, or an expansion plan for the existing pipeline. The most common types of expansion are those which add or upgrade the compressors along the pipeline, which allows the existing pipes to carry a higher capacity. Another option is looping, which amounts to adding another pipeline in parallel to an existing one, allowing the developer to take advantage of the other infrastructure in place along the pipeline. Development projects may also involve replacing and upgrading older parts, adding a “lateral” pipeline extending from an existing one, reversing the direction of flow in an existing pipeline, or converting a pipeline to natural gas from another petroleum product such as oil.

⁵https://www.eia.gov/naturalgas/archive/analysis_publications/ngpipeline/develop.html Archived version; June 01, 2023

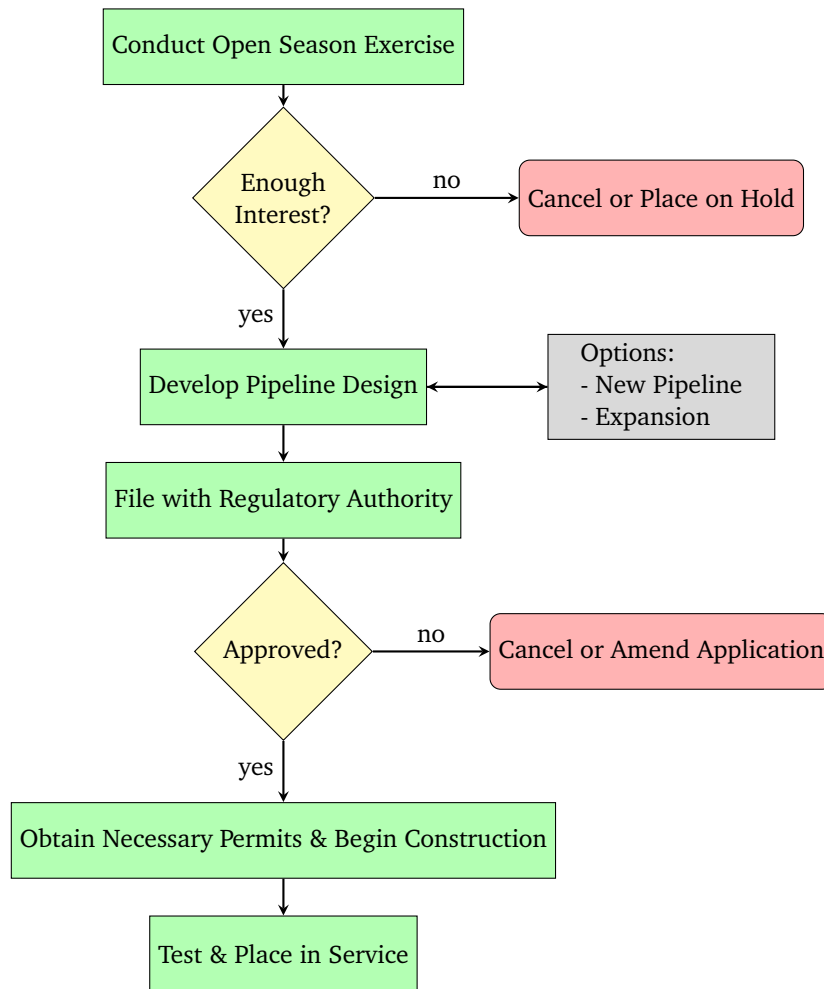


Figure 2: Development and expansion process for pipeline projects, reproduced from the EIA website

During the early stages of design, if the development project involves an interstate pipeline (that is, if the project falls under the jurisdiction of FERC,) then the developer has the option to request a pre-filing review from FERC. This pre-filing review allows FERC the opportunity to review the potential environmental impacts of the project, streamlining the eventual process of regulatory approval.

Once the design plans, and potential pre-filing reviews, are complete, the developer submits their proposal to FERC (or the relevant state regulatory body) for a full review. FERC’s review process takes between 5-18 months, (with an average of 15,) after which FERC will either reject the project, or conditionally accept it. If the project is rejected, the developer has the option to amend the project and resubmit at a later date. Otherwise, they may choose to cancel the project outright.

After obtaining regulatory approval, developers then enter the permitting and construction phase, which generally takes between 6 and 18 months, depending on the scale of the project. Following a short testing period, the pipeline then becomes active.

2 Data and stylized facts on network development

We utilize data on natural gas pipelines from a number of sources. Our primary source of data on pipeline activity come from FERC Forms 2 and 2A. Form 2 records financial and operational information of major interstate pipelines. A less detailed Form 2A records financial and operational information of

non-major interstate pipelines. A major pipeline is defined as having combined gas transported or stored that exceeds 50 million dekatherms.

Form 2/2A is reported annually and data is available from 1991 to the present. Prior to 1996, only data for major pipeline companies is available. In 2021, FERC changed its data collection and retention format from Visual FoxPro to XBRL. Data in the new XBRL format is not available as a bulk download and is more difficult to parse. As a result, we focus on 1996-2021. In 2016 there were 128 active pipeline companies. Not all companies are active in all years. Forms 2 and 2A include detailed information about each company's revenue, expenses, capital, and transmission volume.

FERC Form 2 primarily records data at the company level. Many companies operate over a large geographic area; to determine the areas in which these companies operate, we incorporate data from the Energy Information Agency (EIA) Form 176. Form 176 is designed to collect data on gas supplies, disposition, and certain revenues by state, which covers the data from 1997 to 2015. We merge data from EIA Form 176 and FERC Form 2 by matching company names. Discrepancies in spelling, typos, and changes in companies names required that this matching be checked manually. After doing so, we were able to match company names that collectively own 97% of pipeline mileage.

Additionally we use data from the EIA about natural gas reserves and prices. They provide monthly average prices to both commercial and residential consumers from 1989 to 2023 at the state level. EIA also provides monthly data on statewide extraction and production of natural gas, commercial and residential consumption, statewide storage capacity, and net deposits and withdrawals from storage. The reserve data dating back from 1979 until 2013 included reserves which are wet after lease separation and dry natural gas. The price data also includes wellhead price and citygate price. Wellhead price is calculated by "dividing the total reported value at the wellhead by the total quantity produced", which also includes all the cost before shipping. Citygate price is referred to the price at a point or a measure station where a distributing gas utility receives gas from a natural gas pipeline company or transmission system.⁶

2.1 Development and structural change in the pipeline network

The left panel of Figure 4 shows a scatterplot of pipeline development projects, comparing their costs with their added capacity and length on a log-log scale. It is clear that the cost of a network upgrade is highly correlated with its ability to carry more gas over longer distances. Furthermore the relationship appears roughly linear on this scale. In the right panel, a stacked bar chart shows how these investments were allocated to different types of projects over time. This shows that expansions of existing pipelines are reasonably consistent over time, while the development of new pipelines has been highly concentrated in a few years, from 2008-2011, and 2017-2020. With the fracking boom in the late 2010's, we also see a significant number of projects that reverse the flow of existing pipelines (especially those that flowed from the Gulf to the Northeast region), reflecting a nationwide shift in where gas was being extracted.

As the extraction of natural gas from the Gulf of Mexico began to decline with the boom in shale gas production in the early 2010s, production shifted away from the Gulf Coast and southeastern United States to the Marcellus and Utica shale plays in the northeast. These changes in production across regions can also be seen in the left panel of Figure 5, which shows monthly production levels in top producing states, as well as the Gulf of Mexico. Notably, the downward trend in production in the Gulf of Mexico is more than offset by the rapid increase in production in the Marcellus and Utica shale plays in Pennsylvania and Ohio. The rightmost panel of Figure 5 also displays patterns of consumption over time across the top consuming states. From this, we can see that consumption loosely follows patterns of extraction in

⁶http://www.eia.gov/dnav/ng/TblDefs/ng_pri_sum_tbldef2.asp, Archived version, April 03, 2023

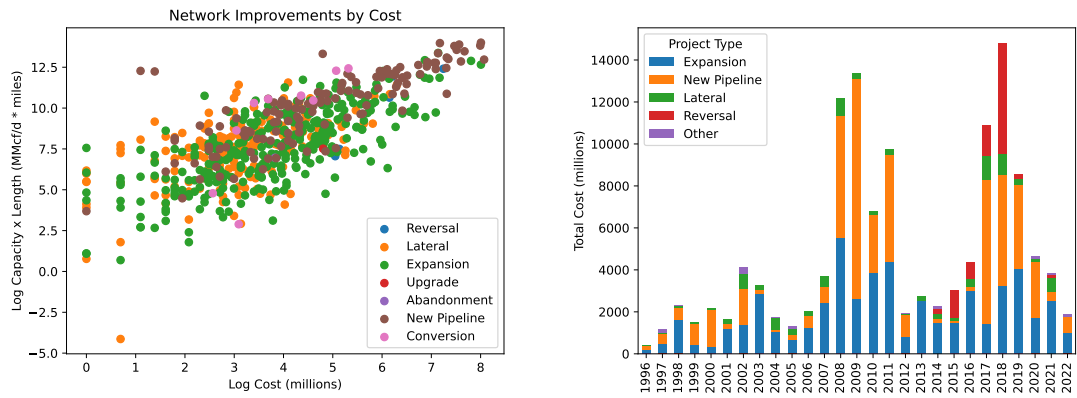


Figure 3: Scatterplot of pipeline development projects, comparing their costs with their added capacity and length on a log-log scale

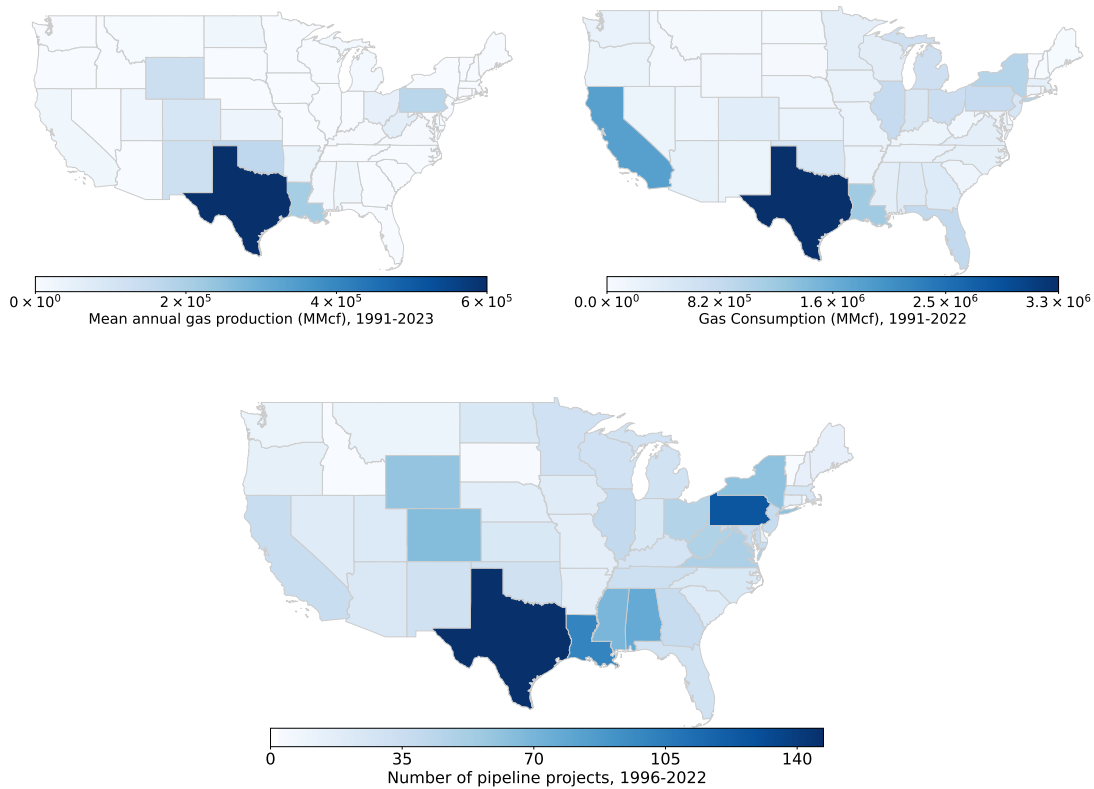


Figure 4: Maps showing the spatial layout of the U.S. gas market

states that produce natural gas, and has remained relatively stable in states that do not.

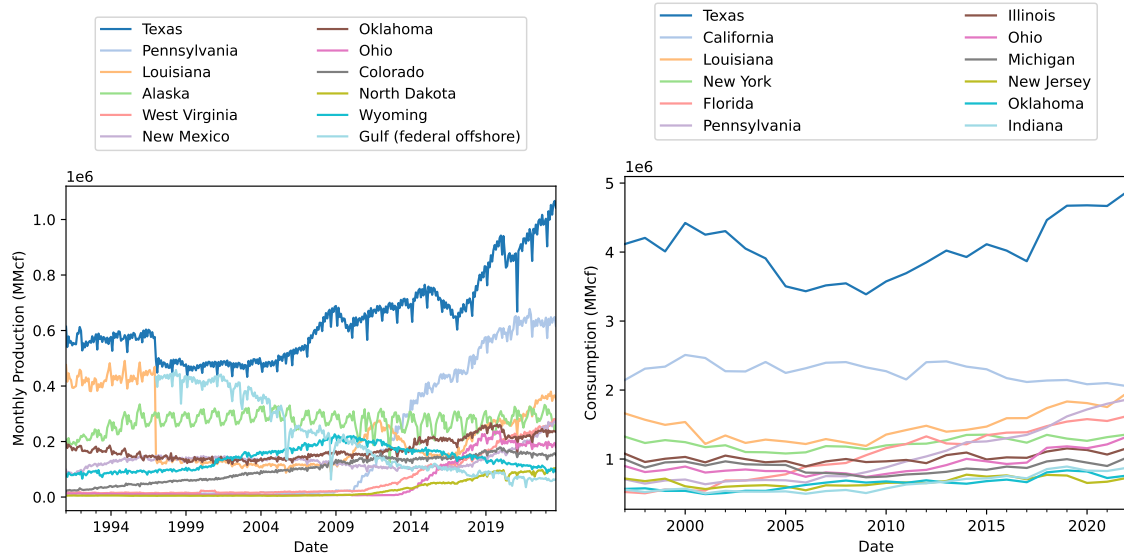


Figure 5: Production and consumption trends in top states

In order to analyze how pipeline projects have changed connectivity, we first investigate how development of interstate gas transmission has affected the connectivity and integration of the network over time. We first examine the spectrum of the network’s normalized Laplacian matrix in 2002 and 2022. The spectrum of a graph’s Laplacian matrix contains information about cut sizes; relatively how difficult it would be to cut the graph into a certain number of components. The difficulty of cutting or segmenting the network into a given number of components is measured as the amount of capacity that would need to be severed to completely disconnect a region from the rest of the network. Due to the duality between the min-cut and max-flow problems, these cuts also give us information about regions in which prices should synchronize reasonably quickly on the intensive margin but where may take longer to integrate across regional boundaries.

Figure 6 suggests a particular cluster structure in network developments. The first large gap in the spectrum of the Laplacian matrix is a measure of the difficulty of cutting the graph into that many components. For that reason, the first large gap between the eigenvalues represents a natural segmentation of the network. A generic clustering algorithm such as K-Means can then be used to identify these clusters by using the eigenvectors of the Laplacian matrix, which is the basis for spectral clustering.

Due to this interpretation of the gaps between eigenvalues of the Laplacian matrix, we can see from Figure 6 that the pipeline network became slightly more difficult to segment into 4 or more components. On the other hand, the difference between the first 3 eigenvalues is reasonably small, indicating that the network is still relatively easy to segment into 3 components. This suggests that developments to the network have been focused within the groups identified by these clusters, while it generally has not served to integrate the network much across the boundaries of these clusters.

Figure 8 shows schematic diagrams of interstate gas transmission capacity in 2002 and 2022 with state centroids colored by their region in this natural cluster structure. In this network, each node represents a state, and states are linked by directed edges that represent the capacity to transport gas from one state to another. While some of the clusters change from 2002 to 2022, the differences between the two are minor, and mostly focused in the boundary between the west and the midwest regions. The northeastern cluster, however, is relatively stable with the exception of Rhode Island, which became more integrated with the southern and midwestern network. It is clear, however, that developments to the network have

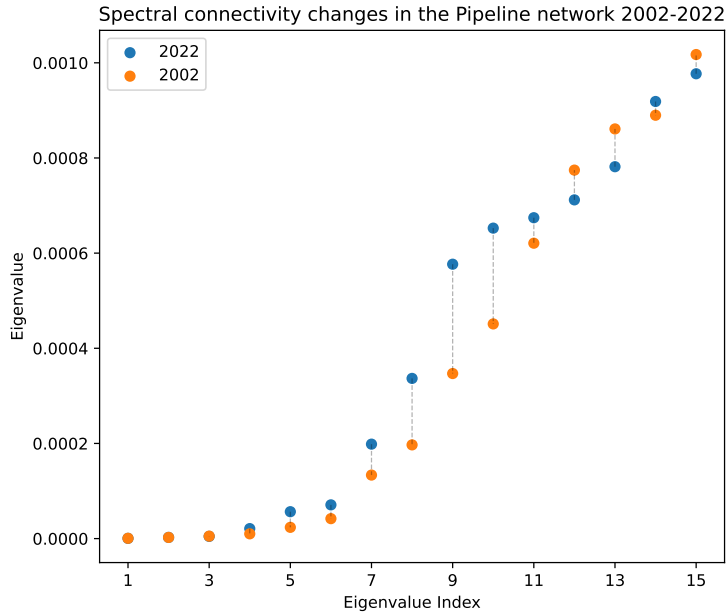


Figure 6: Spectrum of the graph Laplacian in 2002 and 2022

not been focused on expanding the bottleneck to states in northern New England. This indicates that network upgrades have not been targeted at improving connectivity between these regions, although within-region connectivity has improved substantially.

Next, we examine the change in network concentration over time as new pipeline projects are added. For this, we turn to a measure of network centralization which has been called the Von-Neumann entropy (Simmons, Coon, and Datta, 2018). The Von-Neumann entropy is a natural measure of network concentration, and is defined as the entropy of the eigenvalues of the network’s normalized Laplacian matrix. This measure is maximized when the network is completely connected and balanced, and is smaller when the network can be easily separated into any number of components.

Figure 7 shows the change in network centralization due to each pipeline project. The x-axis shows the year in which the project was completed, and the y-axis shows the negative change in the Von-Neumann entropy. The scale is made negative to highlight that large values increase the centralization of the transmission network so that it is more concentrated in a certain area. For each point, we computed the Von-Neumann entropy of the network in the year before the project was completed, and then added the new capacity from each project to the network. We then recomputed the Von-Neumann entropy of the network with the additional capacity, and took their difference.

Figure 7 further emphasizes the finding that pipeline projects do not appear particularly targeted at increasing connectivity between disjoint regions. While some projects have had this effect, (and the net effect of all the changes is decreased centralization,) many projects have been added in areas that have already been well connected. Two particular projects stand out as having particularly increased the centralization of the network. These are the Cheniere Southern Trail Pipeline in 2008 (later renamed the Creole Trail Pipeline) and the Midcontinent Express Pipeline in 2009. These projects are both focused in the Gulf Coast region, which was already well-connected by that time. On the other hand, some projects have had the effect of reducing centralization, with the most notable being the Alliance Project in 2000, which connects Canada to the Midwestern United States, and the Ruby Pipeline Project, which expanded capacity in the American northwest.

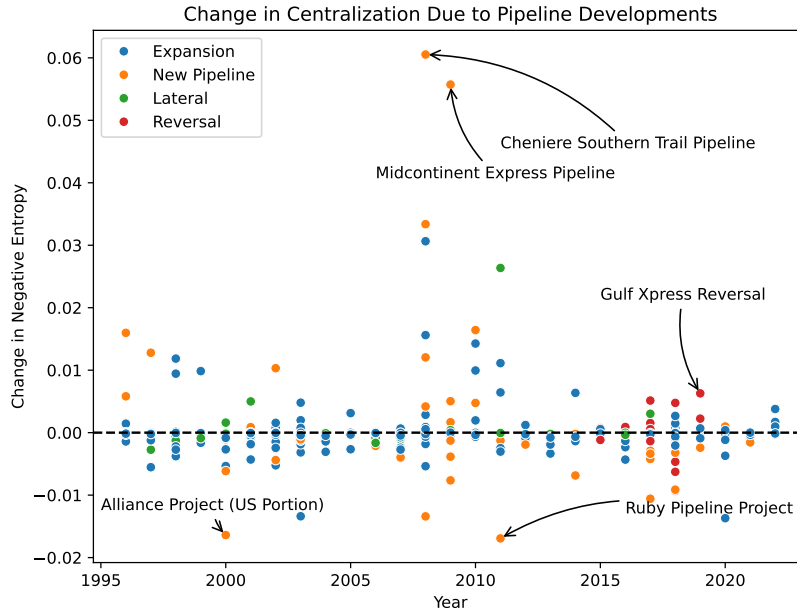


Figure 7: Changes in network centralization due to pipeline projects

The shift in production areas is reflected in pipeline development, as a large number of Reversal projects emerged to accommodate the shift in production region⁷. These include, notably, the reversal of the Gulf Xpress pipeline in 2019 which had previously carried gas from the Gulf Coast to the northeast.

3 Models of investment and optimal capacity

This section describes a model of optimal usage of and investment in the pipeline network. Each period, a social planner optimally transports gas through the pipeline network. The model implies that shadow price of capacity along each segment is equal to the difference in prices between the ends of the segment while the segment is fully utilized. We then characterize optimal investment as satisfying an Euler equation that relates the marginal cost of capacity to price differentials. A key empirical implication is that data on price differences and estimates of the marginal cost of capacity can be used to identify pipeline segments where investment is most and least valuable.

To understand how regulation shapes pipeline companies' investment decisions we also develop a model of firms' investment decisions subject to regulatory constraints. We derive an Euler equation relating firm's marginal profits from investment to the marginal total cost of investment, including the shadow cost of meeting a regulatory constraint. We estimate this Euler equation to evaluate the relative importance financial incentives and regulatory compliance in shaping observed investment patterns.

3.1 Price gaps and consumer surplus

We present a simple theoretical model of optimal pipeline capacity by extending the framework of Cremer, Gasmı, and Laffont (2003) to a fully general transportation network. We will then nest it in a multiple time-scale model—in which markets set prices monthly, and investments are planned annually.

⁷<https://www.naturalgasintel.com/northeast-inspiring-pipe-flow-reversals-capacity-additions/>
Archived version; September 01, 2022

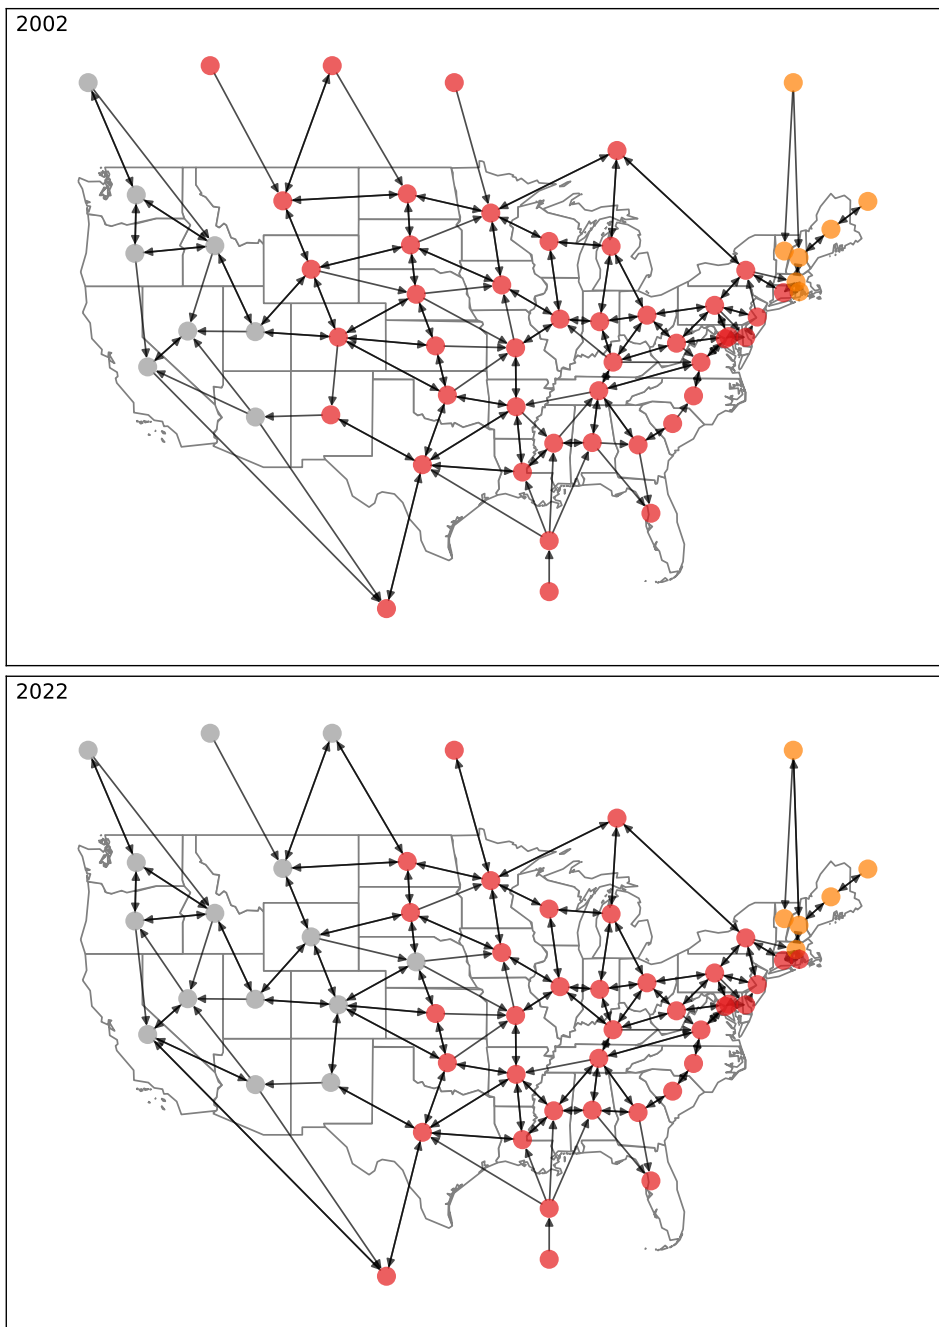


Figure 8: Schematic diagram of interstate gas transmission capacity in 2002 and 2022 with state centroids colored by spectral clustering.

The primitives of the problem include a set of regional markets \mathcal{A} and a matrix Φ of flows, with ϕ_{ij} representing the flow of gas from location i to location j . The matrix C is comprised of c_{ij} , which are the marginal costs of transmission for gas from i to j . Vectors \mathbf{q} and \mathbf{d} represent supply and demand, respectively. Let $u_i(d_i)$ represent the aggregate demand at location i and define $\mathbf{u}(\mathbf{d})$ as a vector with $[\mathbf{u}(\mathbf{d})]_i = u_i(d_i)$. Perfectly competitive extraction firms face total costs $c_i^e(q_i)$ with $\mathbf{c}(\mathbf{q})$ defined similarly. We assume for all locations i that both $u_i(d_i)$ and $c_i^e(q_i)$ are continuously differentiable functions. Finally, the matrix of capacity constraints is K where $\kappa_{ij} \geq 0$ is the maximum volume of gas that can be transferred from location i to location j .

First we characterize a condensed economic dispatch problem for a social planner whose goal is to maximize consumer surplus. The social planner's static problem is to solve

$$\text{maximize}_{\mathbf{q}, \mathbf{d}, \Phi} \sum_{i=1}^n (u_i(d_i) - c_i^e(q_i)) - \sum_{i=1}^n \sum_{j=1}^n c_{ij} \phi_{ij} \quad (1)$$

$$\text{subject to } q_i, d_i \geq 0, \quad \forall i \in \mathcal{A} \quad (2)$$

$$0 \leq \phi_{ij} \leq \kappa_{ij}, \quad \forall i, j \in \mathcal{A} \quad (3)$$

$$q_i + \sum_{\ell=1}^n \phi_{i\ell} = \sum_{\ell=1}^n \phi_{\ell i} + d_i, \quad \forall i \in \mathcal{A}. \quad (4)$$

The Lagrangian for this optimization problem can be written as

$$\begin{aligned} \mathcal{L} = & \sum_{i=1}^n (u_i(d_i) - c_i^e(q_i)) - \sum_{i=1}^n \sum_{j=1}^n c_{ij} \phi_{ij} - \sum_{i=1}^n \sum_{j=1}^n \lambda_{ij} (\phi_{ij} - \kappa_{ij}) \\ & + \sum_{i=1}^n \sum_{j=1}^n \gamma_{ij} \phi_{ij} + \sum_{i=1}^n \mu_i q_i + \sum_{i=1}^n \eta_i d_i \\ & + \sum_{i=1}^n p_i \left(q_i - d_i + \sum_{\ell=1}^n \phi_{i\ell} - \sum_{\ell=1}^n \phi_{\ell i} \right), \end{aligned} \quad (5)$$

in which \mathbf{p}_i are the Lagrange multipliers on the constraint that describes conservation of flow, λ_{ij} is the shadow price of the capacity constraint on each leg of the pipeline, η_i and μ_i are multipliers on the non-negativity constraints. For simplicity, we can rewrite this Lagrangian in vectorized form as

$$\begin{aligned} \mathcal{L} = & \mathbf{1}^\top \mathbf{u}(\mathbf{d}) - \mathbf{1}^\top \mathbf{c}(\mathbf{q}) - \text{tr}(C^\top \Phi) - \text{tr}(\Lambda^\top (\Phi - K)) \\ & + \text{tr}(\Gamma^\top \Phi) + \boldsymbol{\mu}^\top \mathbf{q} + \boldsymbol{\eta}^\top \mathbf{d} + \mathbf{p}^\top (\mathbf{q} - \mathbf{d} + (\Phi - \Phi^\top) \mathbf{1}). \end{aligned} \quad (6)$$

This Lagrangian is straightforward to differentiate; the relevant gradients are given by the following expressions

$$\nabla_{\mathbf{d}} \mathcal{L} = \nabla_{\mathbf{d}} \mathbf{u}(\mathbf{d}) - \mathbf{p} + \boldsymbol{\eta} \quad (7)$$

$$\nabla_{\mathbf{q}} \mathcal{L} = -\nabla_{\mathbf{q}} \mathbf{c}(\mathbf{q}) + \mathbf{p} + \boldsymbol{\mu} \quad (8)$$

$$\nabla_{\Phi} \mathcal{L} = -C - \Lambda + \Gamma + \mathbf{p} \mathbf{1}^\top - \mathbf{1} \mathbf{p}^\top \quad (9)$$

$$\nabla_{\Lambda} \mathcal{L} = -(\Phi - K) \quad (10)$$

$$\nabla_C \mathcal{L} = \Phi \quad (11)$$

$$\nabla_{\mathbf{p}} \mathcal{L} = \mathbf{q} - \mathbf{d} + (\Phi - \Phi^\top) \mathbf{1}. \quad (12)$$

The first order conditions $\nabla_{\mathbf{d}} \mathcal{L} \equiv \mathbf{0}$ and $\nabla_{\mathbf{q}} \mathcal{L} \equiv \mathbf{0}$ determine that the conservation of flow through

the network implies that $\mathbf{p} = \nabla_{\mathbf{d}}\mathbf{u}(\mathbf{d}) = \nabla_{\mathbf{q}}\mathbf{c}(\mathbf{q})$, or that the marginal utility is equal to the marginal extraction cost, in any region with both positive consumption and positive production. The first order condition for Φ asserts that the shadow prices, Λ , of capacity constraints K are determined by the equation $\Lambda - \Gamma = \mathbf{p}\mathbf{1}^\top - \mathbf{1}\mathbf{p}^\top - C$. Using that $[\mathbf{p}\mathbf{1}^\top]_{ij} = p_i$ and $[\mathbf{1}\mathbf{p}^\top]_{ij} = p_j$, the shadow price of the individual constraint on capacity from location i to location j is pinned down by $\lambda_{ij} - \gamma_{ij} = p_i - p_j - c_{ij}$.

Dual feasibility requires that Λ , Γ , \mathbf{p} , $\boldsymbol{\mu}$, and $\boldsymbol{\eta}$ are elementwise nonnegative. Finally, the complementary slackness conditions for the capacity constraints (10) hold that $\text{tr}(\Lambda^\top(\Phi - K)) = \text{tr}(\Gamma^\top\Phi) = 0$. Together, these conditions imply that $\lambda_{ij} = \max\{p_i - p_j - c_{ij}, 0\}$ and $\gamma_{ij} = \max\{p_j - p_i + c_{ji}, 0\}$.

Denote by $v(K)$ the value function for the static social planner problem (1) given the supply and demand functions \mathbf{u} and \mathbf{c} . By the envelope theorem, we have that $\frac{\partial v}{\partial \kappa_{ij}} = \lambda_{ij} = \max\{p_i - p_j - c_{ij}, 0\}$. In other words, the marginal social value of expanding a capacity constraint is simply determined by the magnitude of a price difference that exceeds the marginal cost of transmission.

The quantities and capacities are chosen at different time scales, with quantities chosen monthly while investments in new capacity are chosen annually to maximize annual expected surplus. The planner's value function over investment is given by

$$V^s(\mathbf{k}_t, s_t) = \max_{\mathbf{i}_t} \mathbb{E} \left[\sum_{m=1}^{12} v_{mt}(K(\mathbf{k}_t)) \mid s_t \right] - c(\mathbf{i}_t, \mathbf{k}_t) + \beta \mathbb{E} [V^s(K(\mathbf{k}_t + \mathbf{i}_t), s_{t+1}) \mid s_t, \mathbf{k}_t], \quad (13)$$

where $K(\mathbf{k}) : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$ maps capital at the firm level to a matrix of interstate capacity constraints. The planner's Euler equations, necessary conditions for the social optimality of investment, are given by

$$\nabla_{\mathbf{i}_t} c(\mathbf{i}_t, \mathbf{k}_t) = \beta \mathbb{E} \left[\nabla_{\mathbf{k}} \mathbb{E} \left[\sum_{m=1}^{12} v_{m,t+1}(K(\mathbf{k}_{t+1})) \mid s_{t+1} \right] + \nabla_{\mathbf{i}_t} c(\mathbf{i}_{t+1}, \mathbf{k}_{t+1}) - \nabla_{\mathbf{k}} c(\mathbf{i}_{t+1}, \mathbf{k}_{t+1}) \mid s_t, \mathbf{k}_{t+1} \right]. \quad (14)$$

Which implies that at the firm level, socially optimal investment would satisfy

$$\begin{aligned} \frac{\partial c}{\partial \mathbf{i}}(i_t, k_t, s_t) - \beta \mathbb{E} \left[\frac{\partial c}{\partial \mathbf{i}}(i_{t+1}, k_{t+1}, s_{t+1}) - \frac{\partial c}{\partial \mathbf{k}}(i_{t+1}, k_{t+1}, s_{t+1}) \mid s_t, k_{t+1} \right] = \\ \beta \sum_{m=1}^{12} \mathbb{E} \left[\frac{\partial v_{m,t+1}}{\partial \mathbf{k}}(K(k_{t+1})) \mid s_t, k_{t+1} \right]. \end{aligned}$$

Finally, substituting the envelope theorem identity from above, we have that $\frac{\partial v_{mt}}{\partial \mathbf{k}} = \sum_{j=1}^n \sum_{\ell=1}^n \frac{\partial \kappa_{j\ell}}{\partial \mathbf{k}} \max\{p_{jmt} - p_{\ell mt} - c_{j\ell}, 0\}$. This term takes the right hand side of the expression above,

$$\begin{aligned} \frac{\partial c}{\partial \mathbf{i}}(i_t, k_t, s_t) - \beta \mathbb{E} \left[\frac{\partial c}{\partial \mathbf{i}}(i_{t+1}, k_{t+1}, s_{t+1}) - \frac{\partial c}{\partial \mathbf{k}}(i_{t+1}, k_{t+1}, s_{t+1}) \mid s_t, k_{t+1} \right] = \\ \beta \sum_{m=1}^{12} \sum_{j=1}^n \sum_{\ell=1}^n \mathbb{E} \left[\frac{\partial \kappa_{j\ell}}{\partial \mathbf{k}} \max\{p_{jmt+1} - p_{\ell mt+1} - c_{j\ell}, 0\} \mid s_t, k_{t+1} \right], \quad (15) \end{aligned}$$

and therefore we see that in a socially optimal setting, firms would align the marginal cost of investment to the expected difference in price across a link in the network. However, a key consideration is that pipeline owners' incentives do not align perfectly with social value, and therefore the regulator can introduce additional regulatory costs as a control to steer pipeline investment toward the social optimum.

3.2 Investment model

Next we model the firm's choice of how much to invest in gas transportation infrastructure in the presence of regulation. We treat investment as a single continuous choice; in year t , a pipeline with capital of k_t

earns profits from transporting gas of $\pi(k_t, s_t)$, where s_t is a vector of state variables outside of pipeline p 's control that affect profits.⁸ This will include the capital of other pipelines, the locations of natural gas wells, and the locations where gas is demanded. The pipeline chooses how many dollars to invest in capital that will become operable next period. Investing i_t costs $c(i_t, k_t, s_t)$. This cost function may simply be equal to the dollars of investment, i_t , but it may also include adjustment or other costs. A regulator must approve the pipeline's investment amount. The regulator will approve investment if $R(i_t, k_t, s_t) \leq 0$. The pipeline has annual discount factor β and has information set spanned by s_t when choosing investment. The Bellman equation for the pipeline's value of capital k_t in state s_t is

$$V(k_t, s_t) = \max_{i_t} \pi(k_t, s_t) - c(i_t, k_t, s_t) + \beta \mathbb{E}[V(k_t + i_t, s_{t+1}) | s_t, k_t + i_t] \\ \text{s.t. } R(i_t, k_t, s_t) \leq 0.$$

A necessary condition for the profit maximizing investment is given by an Euler equation.

$$\frac{\partial c}{\partial i}(i_t, k_t, s_t) + \lambda_t \frac{\partial R}{\partial i}(i_t, k_t, s_t) = \tag{16}$$

$$\beta \mathbb{E} \left[\begin{array}{c} \frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) + \frac{\partial c}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) + \lambda_{t+1} \frac{\partial R}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) \\ - \frac{\partial c}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) - \lambda_{t+1} \frac{\partial R}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) \end{array} \middle| s_t, k_{t+1} \right]. \tag{17}$$

3.3 Optimal regulation

We can now characterize the optimal regulation policy by combining the regulator's and firms' Euler equations (15) and (16). First note that (16) can be rearranged to

$$\frac{\partial c}{\partial i}(i_t, k_t, s_t) - \beta \mathbb{E} \left[\frac{\partial c}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) - \frac{\partial c}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) \middle| s_t, k_{t+1} \right] = \\ \beta \mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) + \lambda_{t+1} \frac{\partial R}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) - \lambda_{t+1} \frac{\partial R}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) \middle| s_t, k_{t+1} \right] - \lambda_t \frac{\partial R}{\partial i}(i_t, k_t, s_t). \tag{18}$$

Noting that the left hand side of (18) is the same as that of (15), it follows that the optimal regulation $R^*(i, k, s)$ should satisfy

$$\beta \mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) + \lambda_{t+1} \frac{\partial R^*}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) - \lambda_{t+1} \frac{\partial R^*}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) \middle| s_t, k_{t+1} \right] - \lambda_t \frac{\partial R^*}{\partial i}(i_t, k_t, s_t) = \\ \beta \sum_{m=1}^{12} \sum_{j=1}^n \sum_{\ell=1}^n \mathbb{E} \left[\frac{\partial \kappa_{j\ell}}{\partial k} \max\{p_{jmt+1} - p_{\ell mt+1} - c_{j\ell}, 0\} \middle| s_t, k_{t+1} \right],$$

or, by rearranging,

$$\beta \mathbb{E} \left[\lambda_{t+1} \frac{\partial R^*}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) - \lambda_{t+1} \frac{\partial R^*}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) \middle| s_t, k_{t+1} \right] - \lambda_t \frac{\partial R^*}{\partial i}(i_t, k_t, s_t) = \\ \beta \mathbb{E} \left[\left(\sum_{r=1}^{12} \sum_{j=1}^n \sum_{\ell=1}^n \frac{\partial \kappa_{j\ell}}{\partial k} \max\{p_j - p_\ell - c_{j\ell}, 0\} \right) - \frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) \middle| s_t, k_{t+1} \right]. \tag{19}$$

Because the purpose of the regulation is to align pipeline incentives with social value, optimal regulatory costs simply move in tandem with the difference between marginal social value and marginal product of capital.

⁸These state variables will include the capital and investment of other pipelines. The pipelines are choosing investment simultaneously in a dynamic game. We assume that the data comes from a single equilibrium of this game, but otherwise we do not explicitly use the game or strategic interactions in our estimation. Therefore, we will abstract from the details of the equilibrium.

If pipelines were price takers and faced price $p_j - p_\ell$ for transporting gas from j to ℓ , then pipeline's marginal product of capital would equal the marginal social value of capital,

$$\frac{\partial \pi}{\partial k} = \frac{\partial \kappa_{j\ell}}{\partial k} \max\{p_j - p_\ell - c_{j\ell}, 0\},$$

and there would be no need for regulation. However, this is unlikely to be the case. Pipelines have local market power and are unlikely to behave as though $\frac{\partial}{\partial k}(p_j - p_\ell) = 0$. Also, to discourage withholding capacity, FERC regulates the price received by pipelines and sets it not equal to $p_j - p_\ell$. Let $r_{j\ell}(k)$ be the regulated price for transporting gas from j to ℓ . It is explicitly a function of capital due to the rate of return regulation imposed by FERC. To more clearly see the interaction of regulation and investment incentives, suppose annual profits from operating a pipeline from j to ℓ are given by

$$\pi(k, s) = \mathbb{E} \left[\sum_{m=1}^{12} (r_{j\ell}(k) - c_{j\ell}) \left(\kappa_{j\ell} \mathbf{1}\{p_{j,m} - p_{\ell,m} > r_{j\ell}(k)\} + \mathbf{1}\{p_{j,m} - p_{\ell,m} = r_{j\ell}(k)\} Q_m(r_{j\ell}(k)) \right) \right]$$

In months where the market price gap exceeds the shipping cost, the pipeline will sell its full capacity. In other months, the pipeline will sell transportation until $p_{j,m} - p_{\ell,m} = r_{j\ell}$, and transport quantity $Q_m(r)$, which is a decreasing function of r . The marginal product of capital is then:

$$\frac{\partial \pi}{\partial k} = \mathbb{E} \left[\sum_{m=1}^{12} (r_{j\ell}(k) - c_{j\ell}) \frac{\partial \kappa_{j\ell}}{\partial k} \mathbf{1}\{p_{j,m} - p_{\ell,m} > r_{j\ell}(k)\} + \frac{\partial r_{j\ell}}{\partial k} \kappa_{j\ell} \mathbf{1}\{p_{j,m} - p_{\ell,m} > r_{j\ell}(k)\} + \frac{\partial r_{j\ell}}{\partial k} \mathbf{1}\{p_{j,m} - p_{\ell,m} = r_{j\ell}(k)\} Q_m(r_{j\ell}(k)) + (r_{j\ell}(k) - c_{j\ell}) \mathbf{1}\{p_{j,m} - p_{\ell,m} = r_{j\ell}(k)\} \frac{\partial Q_m}{\partial r} \frac{\partial r_{j\ell}}{\partial k} \right]$$

Two special cases help illustrate the role of regulation. First, if a pipeline is never at capacity, then only the second line in the prior expectation is non-zero. Such a pipeline may want to increase its capital to raise its regulated price above marginal cost. Second, if a pipeline is always fully utilized, then only the first line in the expectation is non-zero. Importantly, the marginal product of capital does not depend on the magnitude of $p_{j,m} - p_{\ell,m}$ (as long as it exceeds $r_{j\ell}$). Thus, pipeline companies with multiple segments at capacity do not receive a clear financial signal about where additional capacity is most valuable to society. The regulator could still achieve an efficient allocation of investment by ensuring that $\frac{\partial \pi}{\partial k}$ is high enough that companies want to invest and only approving projects with high social value.

4 Identification and estimation

Our goal is to assess to what extent FERC's rate of return price regulation and pipeline approval process leads to an efficient level of investment. We do this by comparing the marginal social value⁹ of investment reflected by gas market prices and the relationship between pipeline capacity and investment (i.e. the right hand side of (15)) with the marginal value of investment to pipeline owners. To estimate the social value of capital along segments of the pipeline network, we use data on city-gate gas prices in each state and the observed relationship between pipeline capacity and cost. To estimate the marginal value of investment for pipeline owners, we use pipeline accounting data to estimate the marginal profits of investment. We later use pipelines' observed investment choices to estimate the shadow cost of FERC approval from investment Euler equations.

⁹Note that our measure of welfare does not account for emissions or environmental externalities except to the extent that they are priced into local markets. Such externalities are likely to further increase the need for a costly approval process.

4.1 Comparing firms' financial incentives and the social value of investment

To assess where and when the pipeline approval process should be more or less stringent, we compare firms financial incentive to invest with the marginal social of investment. That is, we will estimate

$$\mathbb{E} \left[\sum_{m=1}^{12} \sum_{j,\ell} \frac{\partial \kappa_{j\ell}}{\partial k} \max\{p_{jmt+1} - p_{\ell mt+1} - c_{j\ell,t+1}, 0\} - \frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) | s_t, k_{t+1} \right] \equiv \Delta(s_{t-1}, k_t) \quad (20)$$

When $\Delta(s_{t-1}, k_t) > 0$, firms have less incentive to invest than a social planner, and there is no need for a costly approval process. Conversely, when $\Delta(s_{t-1}, k_t) < 0$, firms' financial incentive to invest exceeds the social value of investment, so a costly approval process may be useful.

As discussed in section 2, we observe most pipeline information at the company-year level. Let π_{it} be the profits from transportation of company i in year t , and $\kappa_{ij\ell t}$ be the capacity of company i from state j to ℓ in year t . We assume that

$$\pi_{it} = \pi(k_{it}, s_{it}) + \epsilon_{it}.$$

Given the rich set of financial and network information that we observe, we assume that $\mathbb{E}[\epsilon_{it} | k_{it}, s_{it-1}] = 0$.

Similarly, we let

$$\kappa_{ij\ell t} = \kappa_{j\ell}(k_{it}) + u_t$$

with $\mathbb{E}[u_t | k_{it}, j, \ell] = 0$. We are also assuming that the relationship between pipeline capacity and cost depends on location $j\ell$, but not any other state variables.

To facilitate interpretation and inference instead of the entire $\Delta(s_{t-1}, k_t)$ function, we focus on recovering the projection of $\Delta(s_{t-1}, k_t)$ onto location and year. That is, we report

$$\mathbb{E}[\Delta(s_{t-1}, k_t) | t, \text{location}].$$

4.2 Identifying regulatory costs

We identify firms' marginal cost of investment, including regulatory costs, from the Euler equation for investment.

Let $\tilde{c}(i, k, s) = c(i, k) + \lambda R(i, k, s)$ ¹⁰. We will call $\frac{\partial \tilde{c}}{\partial i}$ the effective marginal cost of investment. With this notation, we can rewrite the firm's Euler equation as

$$\frac{\partial \tilde{c}}{\partial i}(i_t, k_t, s_t) - \beta \mathbb{E} \left[\frac{\partial \tilde{c}}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) - \frac{\partial \tilde{c}}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) | s_t, k_{t+1} \right] = \beta \mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) | s_t, k_{t+1} \right]. \quad (21)$$

Given the detailed accounting information that we have, we assume that we can consistently estimate pipelines' profit function, and thus know the right hand side of (21). We then want to recover the effective investment cost function, \tilde{c} , from the Euler equation. There are two difficulties. One is that (21) is a partial differential equation. Second is that (21) lacks any shock unobserved by the econometrician. We discuss each of these in turn.

The Euler equation is a partial differential equation, so without some additional restrictions, \tilde{c} is not identified. To see this, suppose $\tilde{c}(i, k, s) = c_0 + c_i i + c_k k$ is linear. Clearly, c_0 cannot be identified from (16), since it only involves derivatives of \tilde{c} . Less trivially, (16) also cannot separately recover c_i from c_k .

¹⁰ λ is also a function of k, s .

Some algebra will show that it implies

$$\beta c_k + (1 - \beta)c_i = \beta \mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) | s_t \right],$$

from which c_k and c_i cannot be separately identified.

Fortunately, a natural additional restriction on \tilde{c} is that $\tilde{c}(0, k, s) = 0$ for all k and s . This will be the case if $R(0, k, s) < 0$ for all k, s and $c(0, k) = 0$ for all k . In other words, the regulator should always allow pipeline owners to invest zero and this should come with zero cost. Equation (21) is a first order linear partial differential equation. With the boundary condition, $\tilde{c}(0, k, s) = 0$, equation (21) has a unique solution.

To make (21) a viable statistical model, we also introduce an unobserved investment cost shock. Specifically, we assume that

$$\tilde{c}(i, k, s, \eta) = (1 + \eta)i + c_r(i, k, s).$$

Recall that investment is measured in dollars, so c_r can be interpreted as the cost of regulation along with any unreported adjustment or other costs. η is a marginal cost shock, and could include both unobserved cost variation and optimization errors by the pipeline. With this assumption (21) can be written as

$$1 + \eta_t + \frac{\partial c_r}{\partial i}(i_t, k_t, s_t) - \beta \mathbb{E} \left[1 + \eta_{t+1} + \frac{\partial c_r}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) - \frac{\partial c_r}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) | s_t, k_{t+1} \right] = \beta \mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) | s_t, k_{t+1} \right].$$

We will assume that $E[\eta_{t+1} | k_{t+1}, s_t] = 0$ and that η_t is independent of k_t and s_t . Furthermore, we assume that investment is monotonically decreasing in η_t . This implies that

$$\eta_t = G_\eta^{-1} (F_{i|k,s}(i_t | k_t, s_t))$$

where G_η is the CDF of η , and $F_{i|k,s}$ is the conditional CDF of investment. Substituting this expression into the Euler equation gives

$$G_\eta^{-1} (F_{i|k,s}(i_t | k_t, s_t)) + \frac{\partial c_r}{\partial i}(i_t, k_t, s_t) - \beta \mathbb{E} \left[\frac{\partial c_r}{\partial i}(i_{t+1}, k_{t+1}, s_{t+1}) - \frac{\partial c_r}{\partial k}(i_{t+1}, k_{t+1}, s_{t+1}) | s_t, k_{t+1} \right] = \beta \mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) | s_t, k_{t+1} \right] - 1 + \beta$$

The right side of this equation is either estimable or will be assumed to be known. $F_{i|k,s}$ is observable, but the distribution of η and cost function are unknown. A further restriction is needed to generate separate variation in these two unknown functions.

We assume that some components of the state variables are excluded from c_r , but do help predict profits. These variables must shift $\mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_t) | s_t, k_{t+1} = k_t + i_t \right]$, so that they are relevant for i_t , but do not affect the regulators decision to approve pipelines. For exclusions we use variables related to the operating cost of pipelines, such as wages. These are clearly relevant for profits. Pipeline approval is largely based on demonstrating sufficient need for the pipeline. The regulator does not directly consider the pipeline owner's operating costs when approving investment projects.

Let $s_t = (\tilde{s}_t, z_t)$ where c_r might depend on \tilde{s}_t , but not z_t . The Euler equation is then

$$\begin{aligned} G_\eta^{-1}(F_{i|k,s}(i_t|k_t, \tilde{s}_t, z_t)) + \frac{\partial c_r}{\partial i}(i_t, k_t, \tilde{s}_t) - \beta \mathbb{E} \left[\frac{\partial c_r}{\partial i}(i_{t+1}, k_{t+1}, \tilde{s}_{t+1}) - \frac{\partial c_r}{\partial k}(i_{t+1}, k_{t+1}, \tilde{s}_{t+1}) \middle| \tilde{s}_t, z_t, k_{t+1} \right] = \\ = \beta \mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) \middle| \tilde{s}_t, z_t, k_{t+1} \right] - 1 + \beta \end{aligned}$$

This equation and the restrictions discussed are sufficient to identify G_η and c_r . However, to simplify estimation and inference we place an additional restriction on c_r . We assume that $\frac{\partial c_r}{\partial i}$ and $\frac{\partial c_r}{\partial k}$ do not depend on i or k . The boundary condition then implies that $\frac{\partial c_r}{\partial k} = 0$, and the Euler equation simplifies to

$$\begin{aligned} \eta_t + \frac{\partial c_r}{\partial i}(\tilde{s}_t) - \beta \mathbb{E} \left[\frac{\partial c_r}{\partial i}(\tilde{s}_{t+1}) \middle| \tilde{s}_t, z_t, k_{t+1} \right] = \\ = \beta \mathbb{E} \left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) \middle| \tilde{s}_t, z_t, k_{t+1} \right] - 1 + \beta, \end{aligned} \quad (22)$$

and $\frac{\partial c_r}{\partial i}(\tilde{s}_t)$ can be estimated from the conditional moment restriction that $\mathbb{E}[\eta_t | k_t, \tilde{s}_t, z_t] = 0$. Estimation details are described in section 4.4.

To help understand this estimation strategy, consider what happens if \tilde{s}_t is nothing, so $\frac{\partial c_r}{\partial i}$ is constant. In that case a consistent estimate is simply a rescaled and shifted average derivative,

$$\widehat{\frac{\partial c_r}{\partial i}} = \frac{\beta}{1 - \beta} \frac{\partial}{\partial k} \overline{\mathbb{E}[\pi_{t+1} | s_t, k_{t+1}]} - 1.$$

From this, we see that the level of marginal investment costs cannot be separately identified from the discount factor. Moreover, the level of \hat{c}_i will be quite sensitive to the choice of discount factor, as $\frac{\beta}{1-\beta}$ varies rapidly with β near 1.

We will pay special attention to the discount factor that would be implied by a regulator who was optimizing on average. Doing so will allow us to focus on changes that the regulator could make on the intensive margin, rather than in levels. Note that if the regulator were approximately optimizing, their Euler equation (15) should hold. The partial derivatives on the left hand side of (15) do not include these additional regulatory costs. If concrete investment costs are linear, then the planner's Euler equation reduces to

$$1 + \eta_t - \beta = \beta \sum_{m=1}^{12} \sum_{j=1}^n \sum_{\ell=1}^n \mathbb{E} \left[\frac{\partial \kappa_{j\ell}}{\partial k} \max\{p_{jmt+1} - p_{\ell mt+1} - c_{j\ell}, 0\} \middle| s_t, k_{t+1} \right].$$

Then following the same argument as above, and using the moment condition that $\mathbb{E}[\eta_t | s_t, k_{t+1}]$ the discount rate should satisfy

$$\beta \approx \left(\sum_{m=1}^{12} \sum_{j=1}^n \sum_{\ell=1}^n \mathbb{E} \left[\frac{\partial \kappa_{j\ell}}{\partial k} \max\{p_{jmt+1} - p_{\ell mt+1} - c_{j\ell}, 0\} \middle| s_t, k_{t+1} \right] + 1 \right)^{-1}. \quad (23)$$

Recalling (from Figure 10) that marginal social value of capital closely follows the lagged risk-free rate of return presents some motivation for this assumption. For this reason, while we will present results for a range of discount factors, we will focus most attention on $\beta_s = 1.0361^{-1} \approx 0.965$.

4.3 Details of DoubleML estimation

We recover the profit function using a deep neural network, and then form a double robust estimate of $\mathbb{E}[\Delta(s_{t-1}, k_t) | t, \text{location}]$ using the automatic double machine learning approach of Chernozhukov,

Newey, and Singh (2022) and specifically a two-stage version of the deep RieszNet architecture described by Chernozhukov et al. (2022). In choosing deep neural networks to represent the profit function, we are inspired by recent work highlighting how inductive bias in these models aligns well with solutions to dynamic problems in economics (Kahou et al., 2024). The deep learning architecture we used to learn and debias the profit function is shown in Figure 9. In order to recover estimates of Δ , we first estimate each sub-expectation separately.

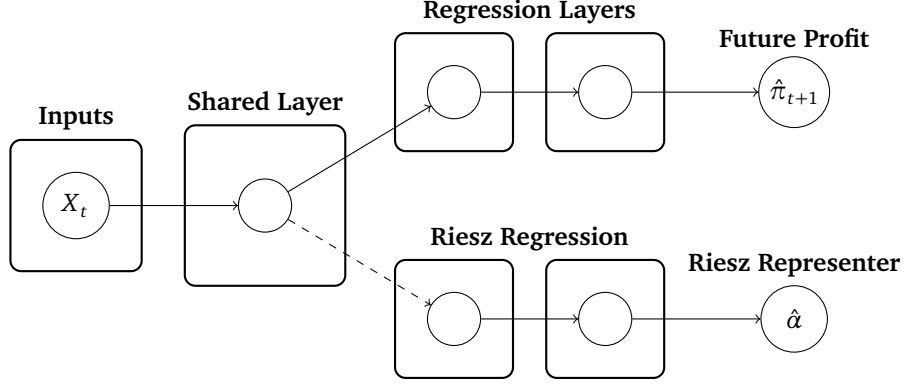


Figure 9: Graphical description of the Auto-DML architecture used to recover and debias the profit function.

We train the neural network using the Adam optimizer, in two separate stages; for the first 1000 epochs, we train the network to predict future profits using a standard mean-squared-error loss function and estimate $\hat{\pi} = \arg \min_{\pi} E[(\pi_{t+1} - \pi_0)^2 | k_{t+1}, s_t]$. After the network is trained, we then use an inner hidden layer of the network as inputs to a separate neural network, which learns the Riesz representation of the moment function over another 1000 epochs. For the profit function, we are interested in the moment function for average derivative estimation, $m(W, g) = \frac{\partial}{\partial k} \pi_{t+1}$. The Riesz network trains to solve $\hat{\alpha} = \arg \min_{\alpha} E[\alpha(X)^2 - 2m(W, \alpha)]$ plus some elastic net (ℓ_1 and ℓ_2 norm) regularization. We then recover “debiased” estimates of $E[\frac{\partial}{\partial k} \pi_{t+1} | k_{t+1}, s_t]$ using the debiased moment formula $E[\frac{\partial}{\partial k} \pi_{t+1} | k_{t+1}, s_t] = E[\frac{\partial}{\partial k} \hat{\pi}_{t+1} + \hat{\alpha}(k_{t+1}, s_t)(\pi_{t+1} - \hat{\pi}_{t+1}) | k_{t+1}, s_t]$. We use cross-fitting with five folds when forming $E[\frac{\partial}{\partial k} \pi_{t+1} | k_{t+1}, s_t]$. Averages and averages conditional¹¹ on k_{t+1}, s_t being in sets of positive measure of $E[\frac{\partial}{\partial k} \pi_{t+1} | k_{t+1}, s_t]$ are consistent and asymptotically normal with the same asymptotic variance as if π and α were known.

The social value of investment is estimated using a similar process but the moment of interest is a bilinear function of two conditional expectations. We follow Chernozhukov et al. (2021) to allow for this nonlinearity. Our parameters of interest are averages and conditional averages of the marginal social value of investment,

$$\theta_0 = E \left[\frac{\partial \kappa_{j\ell}}{\partial k} E[\max\{p_{j_{t+1}} - p_{\ell_{t+1}} - c_{j\ell}, 0\} | s_t, k_{t+1}] \mathbb{1}_A(s_t) / P(A) \right].$$

We assume that observed pipeline capacity is given by

$$\kappa_{j\ell t} = \kappa(k_{j\ell t}, x_{j,\ell}) + u_{j\ell t}$$

with $E[u_{j\ell t} | k_{j\ell t}, x_{j,\ell}] = 0$, and where $x_{j,\ell}$ are location characteristics. To reduce notation, we will assume

¹¹The behavior of the estimates for conditional averages follows from the fact that as long as X does not include k , $m(W, g) \mathbb{1}\{X \in A\} / P(A) = m(W, g \mathbb{1}\{X \in A\}) / P(A)$, so the Riesz representer for $E[m(W, g) | X \in A]$ is simply $\mathbb{1}\{X \in A\} / P(A) \alpha_0(X)$.

that x is a subvector of s . Mimicking the notation of Chernozhukov et al. (2021), let

$$\gamma_{\kappa,0}(k, x) = E[\kappa_{j\ell t} | k, x],$$

$$\gamma_{p,0}(s, k) = E[\max\{p_{j_{t+1}} - p_{\ell_{t+1}} - c_{j\ell}, 0\} | s_t, k_{t+1}],$$

and

$$m(s, k, \gamma_\kappa, \gamma_p) = \left(\frac{\partial}{\partial k} \gamma_\kappa(k, x) \right) \gamma_p(k, s).$$

Then $\theta_0 = E[m(s, k, \gamma_{\kappa,0}, \gamma_{p,0})]$. Moreover, we can form an orthogonal moment condition as

$$\theta = E\left[m(s, k, \gamma_\kappa, \gamma_p) + \alpha_{\kappa,0}(k, x)(\kappa - \gamma_\kappa(k, x)) + \alpha_{p,0}(k, s)(\max\{p_j - p_\ell - c_{j\ell}, 0\} - \gamma_p(k, s)) \right]$$

where

$$\alpha_{\kappa,0} = \arg \min_{\alpha_\kappa} E[-2m(s, k, \alpha_\kappa, \gamma_{p,0}) + \alpha_\kappa(k, x)^2]$$

and

$$\alpha_{p,0} = \arg \min_{\alpha_p} E[-2m(s, k, \gamma_{\kappa,0}, \alpha_p) + \alpha_p(k, s)^2]$$

are Riesz representors for the derivatives of m with respect to each γ . The particular form for the objectives yielding $\alpha_{p,0}$ and $\alpha_{\kappa,0}$ is a consequence of m being bilinear. For estimation, we first fit neural networks to recover $\hat{\gamma}_\kappa$ and $\hat{\gamma}_p$. We then estimate $\hat{\alpha}_\kappa$ and $\hat{\alpha}_p$ by solving the empirical analog of the above minimization problems.

4.4 Details of regulatory cost estimation

Estimation of the marginal regulatory cost is based on the simplified Euler equation (22). This equation involves three unknown objects to be estimated: the expected marginal product of capital, the conditional expectation operator, and the marginal cost function. The expected marginal product of capital, $E\left[\frac{\partial \pi}{\partial k}(k_{t+1}, s_{t+1}) | \tilde{s}_t, z_t, k_{t+1}\right]$, is estimated using the neural network approach described in section 4.3. Denote the resulting estimate by $E\left[\widehat{\frac{\partial \pi}{\partial k}}(k_{t+1}, s_{t+1}) | s_t, k_{t+1}\right]$.

We use a reproducing kernel Hilbert space (RKHS) embedding approach to estimate the conditional expectation operator, $E[\cdot | s_t, k_{t+1}]$. Song et al. (2009) uses a similar approach to model dynamic systems, and Grünewälder et al. (2012b) uses an RKHS embedding to model transition dynamics in a Markov decision problem. Here, we sketch out the theory of RKHS embeddings. More detailed and rigorous descriptions can be found in Grünewälder et al. (2012a) and Park and Muandet (2020). We suppose that $\frac{\partial c_t}{\partial i} \in \mathcal{X}$, a reproducing kernel Hilbert space with known kernel k and inner product $\langle \cdot, \cdot \rangle$. The elements of \mathcal{X} are functions from S to \mathbb{R} . The kernel is a function from $S \times S$ to \mathbb{R} , and $\langle f, k(s, \cdot) \rangle = f(s)$. For each $x \in S$, $E[\cdot | s = x]$ is a linear function from \mathcal{X} to \mathbb{R} . Thus, for each x , there is a Riesz representer, $\mu(x, \cdot)$ such that $E[f(s') | s = x] = \langle f, \mu(x, \cdot) \rangle$. We estimate this μ , and hence $E[\cdot | s]$.

To motivate the estimator, note that the conditional expectation minimizes $\min_g E[(f(s') - g(s))^2]$, and

$$\begin{aligned} E\left[(f(s') - \langle f, \mu(s, \cdot) \rangle)^2\right] &= E\left[\langle f, k(s', \cdot) - \mu(s, \cdot) \rangle^2\right] \\ &\leq \|f\|^2 E[\|k(s', \cdot) - \mu(s, \cdot)\|^2]. \end{aligned}$$

Therefore, we estimate μ by minimizing

$$\min_{\mu} \frac{1}{N(T-1)} \sum_{i=1}^N \sum_{t=1}^{T-1} \|k(s_{it+1}, \cdot) - \mu(s_{it}, \cdot)\|^2 + \lambda \|\mu^2\|$$

The minimizer is

$$\hat{\mu}(s, s') = k(s, \mathbf{s}_t)(K + \lambda I)^{-1} k(\mathbf{s}_{t+1}, s')$$

where K is an $N(T-1) \times N(T-1)$ matrix with entries $k(s_{it}, s_{jr})$, $k(s, \mathbf{s}_{t+1})$ is a $1 \times N(T-1)$ vector with elements $k(s, s_{it+1})$, and $k(\mathbf{s}_t, s')$ is a $N(T-1) \times 1$ vector with elements $k(s_{it}, s')$. With this $\hat{\mu}$, the estimate of the conditional expectation is then

$$\begin{aligned} \widehat{E[f(s')|s]} &= \langle f, \hat{\mu}(s, \cdot) \rangle \\ &= k(s, \mathbf{s}_t)(K + \lambda I)^{-1} f(\mathbf{s}_{t+1}). \end{aligned}$$

In our estimates, we standardize each component of s to have zero mean and unit variance and then use a Gaussian kernel, $k(s, s') = e^{-\|s-s'\|^2}$, and set $\lambda = 1$.

Finally, we represent $\frac{\partial c_r}{\partial i}$ by a neural network and minimize squared violations of the empirical Euler equation,

$$\min_{\frac{\partial c_r}{\partial i}} \frac{1}{N(T-1)} \sum_{i,t}^{N,T-1} \left(\frac{\partial c_r}{\partial i}(\tilde{s}_{it}) - \beta k(s_{it}, \mathbf{s}_t)(K + \lambda I)^{-1} \frac{\partial c}{\partial i}(\tilde{\mathbf{s}}_{t+1}) - E \left[\frac{\partial \pi}{\partial k}(\widehat{k_{t+1}, s_{t+1}} | s_t, k_{t+1}) \right] \right)^2.$$

We use a network with 2 layers and 200 hidden units for the marginal cost function. We train using the Adam optimizer for 2000 epochs. As in section 4.3, we use cross-fitting with five folds for $\widehat{\frac{\partial c_r}{\partial i}}$.¹²

5 Results

In this section we present the results of our estimation. We first present the results for Δ_t , the wedge between profits and social value created by fixed cost-of-service rate regulation. We find that some investment regulation is necessary to align incentives in nearly all states across the time period of observation. Next, we identify the magnitude of the actual investment costs and compare them to the optimal benchmark regulatory strategy.

5.1 Wedge between investment incentives and social value

FERC does not collect capacity data. Such data are provided by the EIA, but firm names are recorded differently. To link the firms across both datasets, we use the LinkTransformer package (Arora and Dell, 2023), with OpenAI's text embedding model (text-embedding-3-large) as the backend. After linking the two datasets, we recover projections of our estimated $E[\Delta(s_{t-1}, k_t | t, location)]$.

We estimate this object according to the procedure described in Section 4.3. Figure 10 shows the results, averaged by state, over time, along with 95% confidence intervals and relevant benchmark values. The top left frame shows how the expected average price gap changes over time, compared to the risk-free rate. The top right frame shows estimates of the marginal profit from capital, compared to previous

¹²The current results do not use debiased moment conditions for the functionals of $\frac{\partial c_r}{\partial i}$ that we report below. We plan to incorporate debiasing for $\frac{\partial c_r}{\partial i}$ in a future revision. See appendix A for details.

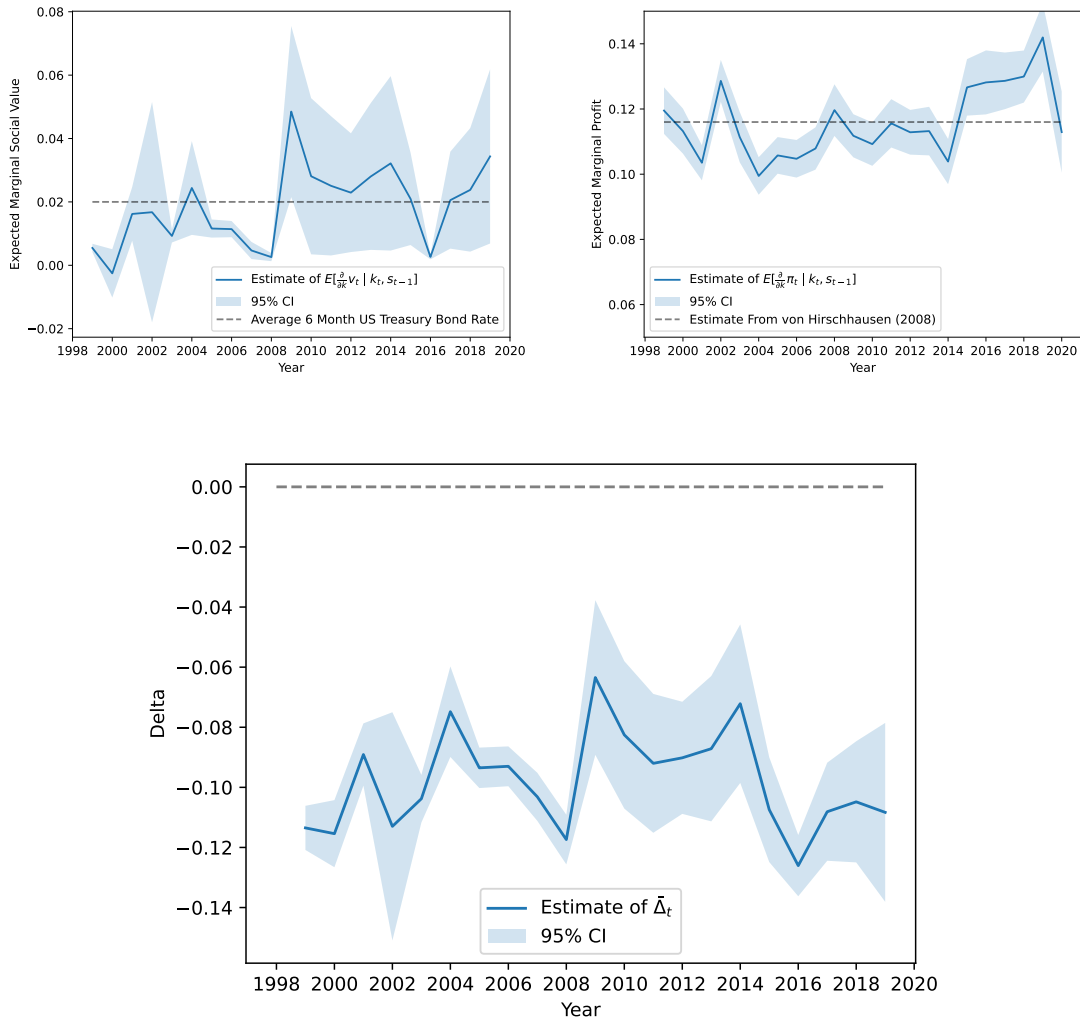


Figure 10: Marginal profits and social value of capital.

estimates from Von Hirschhausen (2008) By differencing these estimates, the bottom panel of 10 shows the same projections of Δ_t over time. To understand how this varies by location, Figure 11 shows averages of $\hat{\Delta}_t$ over the entire time period, grouped by state. As we can see from the figures, estimates are consistently negative in almost all states and the overall time average is always less than zero, to a high degree of statistical significance, across the entire time period. This suggests that a costly approval process for new investments is necessary at the federal level to align profits with social value.

Finally, in Figure 12, we disaggregate these estimates even further, showing how they vary by state and time. As we can see from the figure, the average value of Δ_t is less than or not significantly different from zero in all states and time periods after 2010, further highlighting the importance of investment barriers as well as the substantial spatiotemporal variation in value added from regulatory barriers to investment which are the primary focus of our analysis. In this figure, the left panel shows the 10 states where the incentive gap from $\hat{\Delta}_t$ is highest on average over the time period of observation. As we can see, these results suggest that the need for investment barriers is smallest in the northeast and mountain west; notably in Vermont, New Hampshire, Maine, and Massachusetts. Prior to 2010, we can see that the social value of investment was substantially larger than the marginal profit from capital in Nevada, indicating a bottleneck in the network on which development may have been hampered by the federal

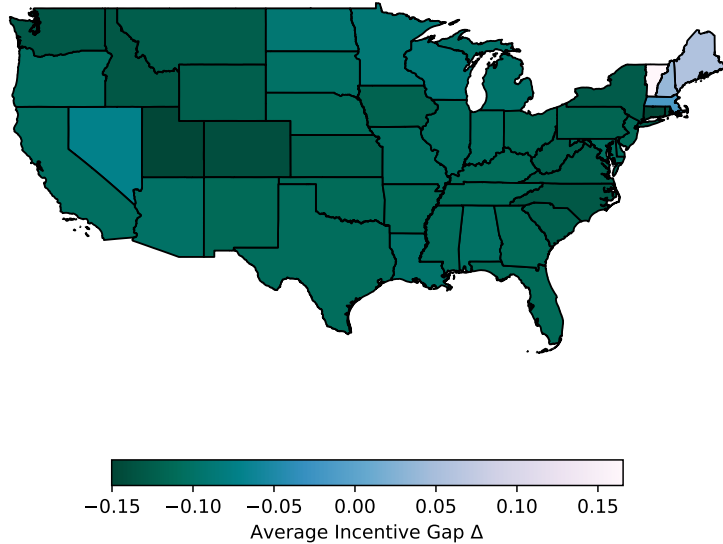


Figure 11: State average incentive gaps

approval process. After 2010, however, Vermont is the only state where this incentive gap is periodically greater than zero (and generally not by very much). This suggests that a stringent approval process is capable of improving efficiency of the network in virtually all of the states, especially between 2010 and 2019.

The results in Figure 12 also agree with our findings from the exploratory analysis in Section 2.1, and particularly the results of spectral clustering. As we can see in Figure 8, the largest dimension on which investment failed to integrate the network separated the map into the West, New England, and a large cluster around the midwest/southeast. Developments over the time period of interest appeared to integrate the mountain west with the midwest, however, New England was not substantially integrated outside of Rhode Island. Indeed, we can see that investments in Rhode Island and Connecticut were substantial, but had low social value—this additional capacity helps to carry gas into the region, but the additional capacity still hits a bottleneck into Massachusetts and fails to reach the northern states, limiting the social value of Rhode Island’s additional capacity.

In the southeast and much of the mountain west, however, investment appears to exceed what would be expected if firm incentives were aligned with social value. This suggests that a costly approval process is most important to limit over-investment in these areas.

5.2 Estimating regulatory costs

We begin by examining the average marginal cost of investment and how it depends on the discount factor. Table 1 shows the estimated marginal “regulatory” cost of investment for a range of discount factors. The estimated marginal product of capital of just over 10% is comparable to what has been found elsewhere. For example, using a very different methodology, von Hirschhausen (2008) found that regulated rates of return average 11.6% for pipeline projects between 1996 and 2003.

The estimation methodology assumes that the marginal return to capital must, on average, equal the marginal cost. Hence, the average “regulatory” marginal cost will be either positive or negative, depending on the discount factor.

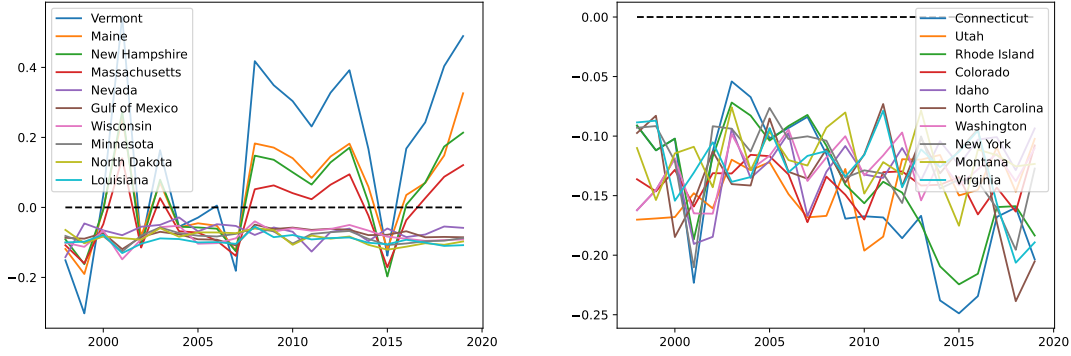


Figure 12: Estimated incentive gap by state and time, in the 10 states where it is largest and smallest.

Table 1: Constant Marginal Cost Estimates

β (fixed)	0.900	0.920	0.940	0.960	0.980	0.990
\hat{c}_i	-0.050	0.213	0.653	1.532	4.170	9.446
	(0.045)	(0.058)	(0.078)	(0.120)	(0.245)	(0.495)

Estimates of average “regulatory” marginal cost of investment. Standard errors clustered on pipeline in parentheses.

We can gain more insight into pipeline’s incentives by allowing investment costs to vary across pipelines. We estimate $\frac{\partial \hat{c}_i}{\partial i}$ as described in section 4.4. Table 2 reports the percentiles of investment costs and marginal returns to capital. From the first row, we see that returns are fairly stable. This is not surprising given that FERC’s regulation of transmission prices is targeted to give a stable return to capital. In the second row, we see that the model implies more variation in marginal costs in investment. Figure 13 shows the distribution of marginal returns to capital and marginal investment costs. We see that overall, the regulatory investment cost is applied evenly across different project types. In the northeast, however, investment costs are higher for projects that are slightly less profitable, creating a downward trend. Since all new capital, once operationalized, should generate fixed rates of return, this suggests that the missing cluster of “low-profit” capital investment reflects projects that take a few years to develop and operationalize. There appears to be a lack of such projects in the northeast, which we can attribute to the higher regulatory costs on such investment. For more profitable projects, however, this investment cost appears smaller on average than in other regions. This suggests that faster projects, which are easier to operationalize, are less subject to these high regulatory costs. Overall, this indicates

Table 2: Variation in Marginal Costs

	Percentile					
	5	10	25	50	75	95
$\frac{\partial}{\partial k} E[\pi_{t+1} s_t, k_{t+1}]$	0.03	0.04	0.06	0.10	0.14	0.21
$\frac{\partial \hat{c}_i}{\partial i}$	0.10	0.11	0.13	0.17	0.23	0.37
Correlation	0.02					

Percentiles of estimated expected marginal product of capital and marginal “regulatory” cost of investment across firms with $\beta = 0.98$.

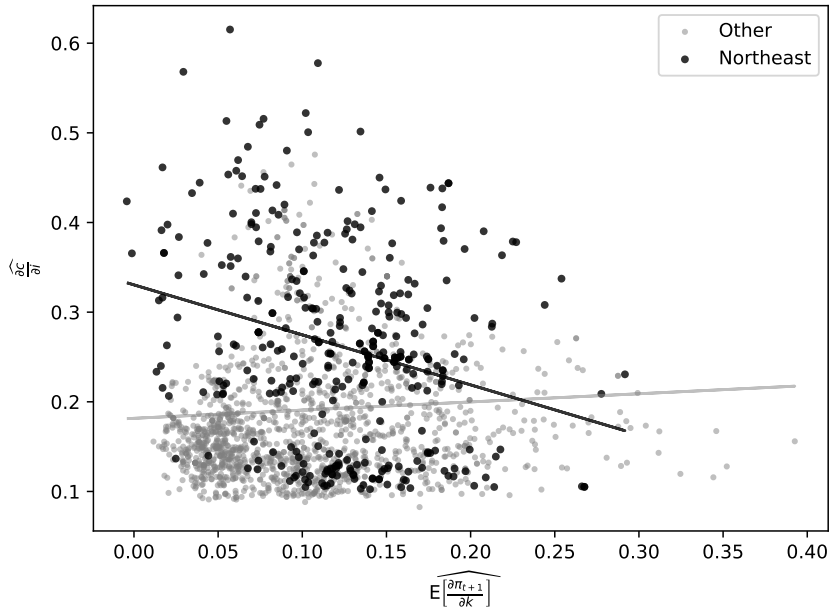


Figure 13: Distribution of marginal costs and marginal returns.

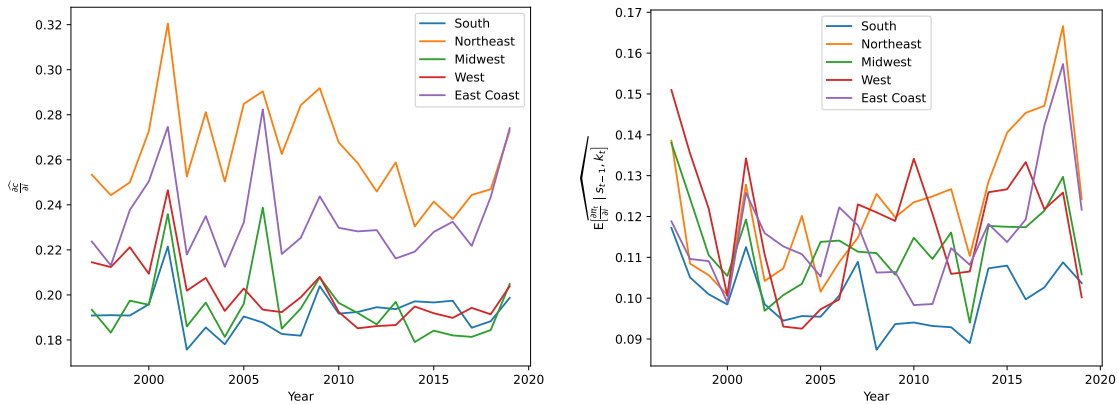


Figure 14: Regional investment incentives over time.

a pattern whereby the higher regulatory costs have created entry barriers for new pipeline infrastructure in the area and gives incumbent firms the opportunity to make small, profitable upgrades at low cost.

Figure 14 shows regulatory costs and marginal profits as they vary over time, and grouped by region. As we can see, implied investment costs are substantially higher in the northeast than in other regions. While profits are slightly higher, there is not a noticeable difference in the average costs of any region over time. In southern states, both marginal profits and marginal investment costs are lower on average.

6 Discussion

In this paper, we estimated the shadow cost of investment approval regulation in the U.S. interstate natural gas pipeline network. By comparing investment incentives to social value, we showed that regulatory approval is necessary for the growth of the network to properly align with efficient outcomes. We found that the existing system of regulatory approval, implemented by FERC, has been approximately optimal

over the past two decades. However, while it is approximately right in levels, we find evidence of some spatiotemporal misallocation of capital investment. This suggests that a more detailed spatial targeting of regulatory stringency could improve consumer welfare. Such a reallocation would reduce regulatory constraints in the northeast, but would also tighten them in the coastal south and mountain west.

References

- Adamson, Seabron. 2018. "Comparing Interstate Regulation and Investment in US Natural Gas and Electric Transmission." *Economics of Energy & Environmental Policy* 7 (1):7–24.
- Allen, Treb and Costas Arkolakis. 2022. "The Welfare Effects of Transportation Infrastructure Improvements." *The Review of Economic Studies* 89 (6):2911–2957. URL <https://doi.org/10.1093/restud/rdac001>.
- Arora, Abhishek and Melissa Dell. 2023. "LinkTransformer: A Unified Package for Record Linkage with Transformer Language Models." *arXiv preprint arXiv:2309.00789* .
- Avalos, Roger, Timothy Fitzgerald, and Randal R Rucker. 2016. "Measuring the effects of natural gas pipeline constraints on regional pricing and market integration." *Energy Economics* 60:217–231.
- Averch, Harvey and Leland L Johnson. 1962. "Behavior of the firm under regulatory constraint." *The American Economic Review* 52 (5):1052–1069.
- Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. "Geography, transportation, and endogenous trade costs." *Econometrica* 88 (2):657–691.
- . 2024. "Investment in Infrastructure and Trade: The Case of Ports." Tech. rep., National Bureau of Economic Research.
- Brito, Dagobert L and Juan Rosellón. 2011. "Lumpy investment in regulated natural gas pipelines: an application of the theory of the second best." *Networks and Spatial Economics* 11:533–553.
- Bushnell, James. 1999. "Transmission rights and market power." *The Electricity Journal* 12 (8):77–85.
- Caoui, El Hadi and Andrew Steck. 2023. "Network Diversity, Market Entry, and the Global Internet Backbone." *SSRN eLibrary* .
- Chen, Yanyou. 2024. "Network Structure and the Efficiency Gains from Mergers: Evidence from US Freight Railroads." Available at SSRN 4742990 .
- Chernozhukov, Victor, Whitney Newey, Victor M Quintas-Martinez, and Vasilis Syrgkanis. 2022. "Riesznet and forestriesz: Automatic debiased machine learning with neural nets and random forests." In *International Conference on Machine Learning*. PMLR, 3901–3914.
- Chernozhukov, Victor, Whitney K Newey, Victor Quintas-Martinez, and Vasilis Syrgkanis. 2021. "Automatic debiased machine learning via neural nets for generalized linear regression." *arXiv preprint arXiv:2104.14737* URL <https://arxiv.org/pdf/2104.14737.pdf>. Revised 2024.
- Chernozhukov, Victor, Whitney K Newey, and Rahul Singh. 2022. "Automatic debiased machine learning of causal and structural effects." *Econometrica* 90 (3):967–1027.
- Cobanli, Onur and Franz Hubert. 2014. "Pipeline Power: A Case Study of Strategic Network Investments." Tech. rep., Kiel und Hamburg: ZBW-Deutsche Zentralbibliothek für

- Cox, James C and R Mark Isaac. 1987. “Mechanisms for incentive regulation: theory and experiment.” *The RAND Journal of Economics* :348–359.
- Cremer, Helmuth, Farid Gasmi, and Jean-Jacques Laffont. 2003. “Access to Pipelines in Competitive Gas Markets.” *Journal of Regulatory Economics* 24 (1):5–33. URL <http://dx.doi.org/10.1023/A:1023943613605>.
- Cremer, Helmuth and Jean-Jacques Laffont. 2002. “Competition in gas markets.” *European Economic Review* 46 (4-5):928–935. URL <http://www.sciencedirect.com/science/article/B6V64-44W42T9-3/2/1a1d33358e00c05f2810096d6933ae1b>.
- Cui, Liyuan, Yongmiao Hong, and Yingxing Li. 2021. “Solving Euler equations via two-stage non-parametric penalized splines.” *Journal of Econometrics* 222 (2):1024–1056. URL <https://www.sciencedirect.com/science/article/pii/S0304407620302323>.
- Dechert, W Davis. 1984. “Has the Averch-Johnson effect been theoretically justified?” *Journal of Economic Dynamics and Control* 8 (1):1–17.
- Degiovanni, Pedro and Ron Yang. 2023. “Economies of Scale and Scope in Railroading.” Tech. rep., Working Paper.
- Dott, Dawn Robin. 1999. *Optimal network design for natural gas pipelines*. MQ-31387 UMI. University of Calgary.
- Ducruet, César, Réka Juhász, Dávid Krisztián Nagy, and Claudia Steinwender. 2024. “All aboard: The effects of port development.” *Journal of International Economics* 151:103963.
- Ejarque, Joao Miguel. 2011. “Evaluating the economic cost of natural gas strategic storage restrictions.” *Energy Economics* 33 (1):44–55.
- Escanciano, Juan Carlos, Stefan Hoderlein, Arthur Lewbel, Oliver Linton, and Sorawoot Srisuma. 2021. “Nonparametric Euler Equation Identification and Estimation.” *Econometric Theory* 37 (5):851–891.
- Fajgelbaum, Pablo D. and Edouard Schaal. 2020. “Optimal Transport Networks in Spatial Equilibrium.” *Econometrica* 88 (4):1411–1452. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15213>.
- Goldstein, Ezra G. 2024. “Communication costs in science: evidence from the National Science Foundation Network.” *Industrial and Corporate Change* 33 (4):785–807.
- Gowrisankaran, Gautam, Ashley Langer, and Mar Reguant. 2024. “Energy transitions in regulated markets.” Tech. rep., National Bureau of Economic Research.
- Grünewälder, Steffen, Guy Lever, Luca Baldassarre, Sam Patterson, Arthur Gretton, and Massimiliano Pontil. 2012a. “Conditional mean embeddings as regressors.” In *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML’12*. Madison, WI, USA: Omnipress, 1803–1810. URL <https://icml.cc/2012/papers/898.pdf>.
- Grünewälder, Steffen, Guy Lever, Luca Baldassarre, Massimiliano Pontil, and Arthur Gretton. 2012b. “Modelling transition dynamics in MDPs with RKHS embeddings.” In *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML’12*. Madison, WI, USA: Omnipress, 1603–1610.

- Hansen, Lars Peter and Kenneth J. Singleton. 1982. "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models." *Econometrica* 50 (5):1269–1286. URL <http://www.jstor.org/stable/1911873>.
- Hausman, Catherine and Lucija Muehlenbachs. 2019. "Price regulation and environmental externalities: Evidence from methane leaks." *Journal of the Association of Environmental and Resource Economists* 6 (1):73–109.
- Isaac, R Mark. 1991. "Price cap regulation: a case study of some pitfalls of implementation." *Journal of Regulatory Economics* 3 (2):193–210.
- Jaworski, Taylor, Carl Kitchens, and Sergey Nigai. 2023. "Highways and globalization." *International Economic Review* 64 (4):1615–1648.
- Jaworski, Taylor and Carl T Kitchens. 2019. "National policy for regional development: Historical evidence from Appalachian highways." *Review of Economics and Statistics* 101 (5):777–790.
- Joskow, Paul L. 2005. "Regulation and deregulation after 25 years: Lessons learned for research in industrial organization." *Review of Industrial Organization* 26:169–193.
- Kahou, Mahdi Ebrahimi, James Yu, Jesse Perla, and Geoff Pleiss. 2024. "How Inductive Bias in Machine Learning Aligns with Optimality in Economic Dynamics." *arXiv preprint arXiv:2406.01898* .
- Keeler, Theodore E. 1974. "Railroad costs, returns to scale, and excess capacity." *The Review of Economics and Statistics* :201–208.
- MacAvoy, Paul W. 2007. *Natural Gas Networks Performance after Partial Deregulation: Five Quantitative Studies*, vol. 5. World Scientific.
- Marks, Levi, Charles F Mason, Kristina Mohlin, and Matthew Zaragoza-Watkins. 2017. "Vertical market power in interconnected natural gas and electricity markets." Tech. rep., Center for Economic Studies & Ifo Institute.
- Marmer, Vadim, Dmitry Shapiro, and Paul MacAvoy. 2007. "Bottlenecks in regional markets for natural gas transmission services." *Energy Economics* 29 (1):37–45.
- Meyer, John Robert. 1959. *The economics of competition in the transportation industries*. 107. Harvard University Press.
- Mollenhauer, Mattes and Péter Koltai. 2020. "Nonparametric approximation of conditional expectation operators." *ArXiv abs/2012.12917*. URL <https://api.semanticscholar.org/CorpusID:229371102>.
- Nasser, Thomas-Olivier. 1999. *Congestion pricing and network expansion*. The World Bank.
- Neumann, Anne, Juan Rosellón, and Hannes Weigt. 2015. "Removing cross-border capacity bottlenecks in the European natural gas market—a proposed merchant-regulatory mechanism." *Networks and Spatial Economics* 15:149–181.
- Oliver, Matthew E. 2019. "Pricing flexibility under rate-of-return regulation: Effects on network infrastructure investment." *Economic Modelling* 78:150–161.
- Oliver, Matthew E and Charles F Mason. 2018. "Natural gas pipeline regulation in the United States: Past, present, and future." *Foundations and Trends® in Microeconomics* 11 (4):227–288.

- Oliver, Matthew E, Charles F Mason, and David Finnoff. 2014. “Pipeline congestion and basis differentials.” *Journal of Regulatory Economics* 46 (3):261–291.
- Park, Junhyung and Krikamol Muandet. 2020. “A Measure-Theoretic Approach to Kernel Conditional Mean Embeddings.” In *Advances in Neural Information Processing Systems*, vol. 33, edited by H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin. Curran Associates, Inc., 21247–21259. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/f340f1b1f65b6df5b5e3f94d95b11daf-Paper.pdf.
- Preonas, Louis. 2024. “Market power in coal shipping and implications for US climate policy.” *Review of Economic Studies* 91 (4):2508–2537.
- Redding, Stephen J and Matthew A Turner. 2015. “Transportation costs and the spatial organization of economic activity.” *Handbook of regional and urban economics* 5:1339–1398.
- Rosenberg, Alan E. 2000. “Congestion pricing or monopoly pricing?” *The Electricity Journal* 13 (3):33–41.
- Scott, Jonathan B. 2023. “Positive spillovers from infrastructure investment: How pipeline expansions encourage fuel switching.” *Review of Economics and Statistics* :1–17.
- Simmons, David E, Justin P Coon, and Animesh Datta. 2018. “The von Neumann Theil index: characterizing graph centralization using the von Neumann index.” *Journal of Complex Networks* 6 (6):859–876.
- Song, Le, Jonathan Huang, Alex Smola, and Kenji Fukumizu. 2009. “Hilbert space embeddings of conditional distributions with applications to dynamical systems.” In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML ’09*. New York, NY, USA: Association for Computing Machinery, 961–968. URL <https://doi.org/10.1145/1553374.1553497>.
- Vickrey, William S. 1969. “Congestion theory and transport investment.” *The American Economic Review* :251–260.
- Von Hirschhausen, Christian. 2008. “Infrastructure, regulation, investment and security of supply: A case study of the restructured US natural gas market.” *Utilities Policy* 16 (1):1–10.
- von Hirschhausen, Christian. 2008. “Infrastructure, regulation, investment and security of supply: A case study of the restructured US natural gas market.” *Utilities Policy* 16 (1):1 – 10. URL <http://www.sciencedirect.com/science/article/pii/S0957178707000598>.
- Wellisz, Stanislaw H. 1963. “Regulation of natural gas pipeline companies: An economic analysis.” *Journal of Political Economy* 71 (1):30–43.
- Yuan, Zhe and Panle Jia Barwick. 2024. “Network Competition in the Airline Industry: An Empirical Framework.” Tech. rep., National Bureau of Economic Research.

A Double ML for Cost Estimates

As with the marginal return to capital, our main results about costs are based on averages and conditional averages of the marginal cost. However, unlike section 4.3, The average of marginal cost is a nonlinear function of conditional expectations. We must account for this nonlinearity while debiasing. We follow the approach of Chernozhukov et al. (2021) with some technical modifications to allow for our parameter

of interest depending on a conditional expectation operator in addition to conditional expectations. Let $\gamma_0(s) = E[\pi(k', s')|s]$ and $\mu_0(f)(s) = E[f(s')|s]$. The approach of Chernozhukov et al. (2021) relies on γ_0 and μ_0 being elements of Hilbert spaces, so that the Riesz representation theorem can be used. We assume $\gamma_0 \in \mathcal{L}^2(P_S)$. A natural space for μ_0 is $BL(\mathcal{L}^2(P_S), \mathcal{L}^2(P_S))$, but this is not a Hilbert space. As in the estimation of μ , reproducing kernel Hilbert spaces provide a convenient solution. We assume that $\frac{\partial c_r}{\partial i} \in \mathcal{K}$, an RKHS with kernel k . Mollenhauer and Koltai (2020) show that the completion of

$$\text{span}\{k(s, \cdot)h(\cdot) : s \in S, h \in \mathcal{K}\}$$

with innerproduct

$$\langle k(s, \cdot)h, k(s', \cdot)g \rangle_{\mathcal{G}} = k(s, s') \langle h, g \rangle_{\mathcal{K}}$$

is isometric to the space of Hilbert-Schmidt operators from $\mathcal{K} \rightarrow \mathcal{K}$. We will denote this space as \mathcal{G} and assume that $\mu_0 \in \mathcal{G}$.

For parameters of interest, we will focus on weighted averages of the marginal cost function. Using the notation of section 4.3, we take

$$\theta = E[\underbrace{w(s)(I - \beta\mu)^{-1}T(\gamma)(s)}_{\equiv m(s, \gamma, \mu)}]$$

where w is a known weighting function and T is known linear transformation (e.g. the derivative of γ with respect to k).

The directional derivative of m with respect to μ at γ_0 and μ_0 ,

$$\frac{d}{dt} E[w(s)m(s, \gamma_0, \mu_0 + t\Delta)]|_{t=1} = E[-w(s)(I - \beta\mu)^{-1}\beta\Delta(I - \beta\mu_0)^{-1}T(\gamma_0)(s)] = E[D_\mu(s, \Delta, \gamma_0, \mu_0)]$$

is a continuous linear functional of $\Delta \in \mathcal{G}$. Therefore, there exists a unique $\alpha_{0, \mu} \in \mathcal{G}$ such that

$$E[w(s)(I - \beta\mu)^{-1}\Delta(I - \beta\mu_0)^{-1}T\gamma_0(s)] = \langle \alpha_{0, \mu}, \Delta \rangle_{\mathcal{G}}$$

Also, m is linear in γ , so there exists $\alpha_{0, \gamma}$ such that

$$E[m(s, \gamma, \mu_0)] = E[\alpha_{0, \gamma}(s)\gamma(s)].$$

An orthogonal moment condition for estimating θ is then

$$\theta_0 = E[m(s, \gamma, \mu) + \alpha_{0, \gamma}(s)(\pi_{t+1} - \gamma(s))] + \langle \alpha_{0, \mu}, \mu_0 - \mu \rangle_{\mathcal{G}}$$

As in Chernozhukov et al. (2021) and section 4.3,

$$\alpha_{0, \gamma} = \arg \min_{\alpha \in \mathcal{L}^2(P_S)} E[\alpha(s)^2 - 2m(s, \alpha, \mu_0)],$$

and an empirical analog can be used for estimation. Similarly for $\alpha_{0, \mu}$,

$$\alpha_{0, \mu} = \arg \min_{\alpha \in \mathcal{G}} \|\alpha\|_{\mathcal{G}}^2 - 2E[D_\mu(s, \alpha, \gamma_0, \mu_0)].$$

A regularized empirical analog is

$$\hat{\alpha}_\mu = \arg \min_{\alpha \in \mathcal{G}} (1 + \lambda_\alpha) \|\alpha\|_{\mathcal{G}}^2 - 2 \frac{1}{NT} \sum_{i,t} D_\mu(s_{it}, \alpha, \hat{\gamma}, \hat{\mu}). \quad (24)$$

As typical in regularized empirical loss minimization on RKHS, note that $G_n = \text{span}\{\sum_{i,t} h_{it}(\cdot)k(s_{it}, \cdot) : h_{it} \in \mathcal{X}\}$ is a linear subspace of \mathcal{G} . For any $\alpha \in \mathcal{G}$, there is a unique $\alpha_n \in G_n$ and $\alpha_n^\perp \in G_n^\perp$ such that $\alpha = \alpha_n + \alpha_n^\perp$, and

$$\|\alpha\|_{\mathcal{G}} = \|\alpha_n\|_{\mathcal{G}} + \|\alpha_n^\perp\|_{\mathcal{G}}.$$

Moreover, $D_\mu(s_{it}, \alpha, \hat{\gamma}, \hat{\mu})$ only depends α_n , not α_n^\perp . It follows that the minimizer in (24) must have $\hat{\alpha}_{\mu,n}^\perp = 0$. We can find $\hat{\alpha}$ by solving

$$\min_{\{h_{it+1}\} \in \mathcal{X}^{NT}} (1 + \lambda_\alpha) \sum_{i,t=1}^{N,T} \sum_{j,r=1}^{N,T} k(s_{it}, s_{jr}) \langle h_{it}, h_{jr} \rangle_{\mathcal{X}} + \frac{1}{NT} \sum_{i,t=1}^{N,T} D_\mu(s_{it}, \sum_{i,t} h_{it}(\cdot)k(s_{it}, \cdot), \hat{\gamma}, \hat{\mu}).$$

By repeating a similar argument, we can further conclude that the minimizing h_{it} must be of the form $h_{it} = \sum_{j,r=1}^{N,T} a_{jr}^{it} k(s_{jr+1}, \cdot)$ where $a_{jr}^{it} \in \mathbb{R}$. Thus, $\hat{\alpha}$ can be written as $k(\cdot, \mathbf{s}_t) A k(\mathbf{s}_{t+1}, \cdot)$, much like $\hat{\mu}$ from section 4.4. We estimate $\hat{\alpha}$ by minimizing over the a_{jr}^{it} .