# Who Bears the Risk of Extreme Weather? Capacity Allocation in U.S. Natural Gas Pipelines *

Yanyou Chen
University of Toronto

Emmanuel Murray-Leclair
University of Melbourne

Adam Wyonzek
The Brattle Group

August 2026

**Abstract**

Who bears the risk of extreme weather when transportation capacity must be allocated before the weather is known? We study the U.S. natural gas pipeline network, where cold waves increase heating demand from both gas utilities and gas-fired electricity generators, who must secure pipeline capacity and organize point-to-point transport logistics. We develop a model of ex-ante capacity allocation and ex-post trade, accounting for the need to secure pipeline capacity in advance and for peak-day capacity requirements imposed on gas utilities. We show that the market implements the efficient allocation if and only if the ranking of users' marginal valuations does not change with weather. This condition fails when gas-fired generators are dispatched only during extreme weather periods. Peak-day regulation, intended to protect residential gas customers, induces utilities to keep too much capacity during cold waves, shifting weather risk to generators and households relying on electric heating. Using the universe of U.S. interstate pipeline capacity release contracts, we find that the price of short-term capacity releases more than doubles during ten major cold waves from 2014–2023, while quantity released doesn't change. Prices, rather than quantities, adjust, leaving users without long-term contracts facing a thin and expensive market when their demand for pipeline capacity rises.

**Keywords:** Ex-Ante Contracting; Regulatory Distortions; Natural Gas Pipelines; Extreme Weather
**JEL Codes:** L95, L51, Q41, Q48

*We are grateful to Victor Aguirregabiria, Heski Bar-Isaac, Yenan Wang, and seminar participants at the University of Toronto, Western University, the 2024 Canadian Economic Association Conference (CEA), and IIOC 2025 for helpful discussions and comments. We also thank Anna Li, William Qi, Eden Zheng, Shuxin Tan, and Vincent Zhang for excellent research assistance. Chen gratefully acknowledges financial support from the Social Sciences and Humanities Research Council of Canada through the University of Toronto SSHRC Institutional Grant. All errors are our own.

# 1 Introduction

A regulation designed to protect households from extreme weather can shift the risk onto other households. In U.S. natural gas markets, state regulators require gas utilities to secure enough pipeline capacity to serve residential customers on peak-demand days. These obligations give gas-heated households priority access to scarce transportation, but the same pipelines also supply gas-fired power plants that generate electricity for households, including those that heat with electricity. When a cold wave raises both demands, capacity reserved by gas utilities may be unavailable to generators just when the latter need it most. The resulting distributional tension is stark: a rule intended to ensure reliable gas heating can increase the exposure of electrically heated households to winter fuel shortages. This paper asks how capacity contracts and regulation determine who bears the risk of extreme weather when transportation must be allocated before the weather is known.

The problem arises because pipeline rights are difficult to reallocate after demand changes. Natural gas pipelines require large sunk investments and are financed primarily through long-term contracts, often lasting a decade or more, that confer priority in pipeline scheduling. A secondary market allows holders to release capacity temporarily to other shippers, but transportation is decentralized across pipeline segments and systems, and interruptible service is rarely available when demand is high. In our capacity release data, transportation rights are typically secured about a week in advance, while the severity of a cold wave remains uncertain. Cold waves are among the largest shocks to gas demand, are difficult to forecast (Safavi, Wu, Lo Prete and Ünel, 2025), and are becoming more frequent (Wang et al., 2021; Bathiany et al., 2018). They raise demand directly among the roughly half of U.S. households that heat with gas and indirectly through gas-fired electricity generation. The shock can therefore change which shipper values gas most after contractual priority has already been assigned.

We focus on two frictions that turn this timing problem into a distributional distortion. The first is a *deliverability constraint*: because shippers must arrange transportation in advance, pipeline capacity cannot readily move to its highest-valued use after the weather is realized. Unlike electricity markets, where a central operator coordinates transportation, the gas market relies on decentralized scheduling and prearranged priority rights.[1] The second is a *capacity obligation*: local distribution companies (LDCs), which sell gas to households at regulated prices that adjust slowly, are required to plan for peak-day demand and can be held responsible for failing to serve it. Without such an obligation, an LDC facing a sharp

[1]Priority contracting is not unique to gas pipelines. Hoy et al. (2016) study priority and spot contracts for cloud computing capacity; priority arrangements also exist in telecommunications, air freight, and bulk water. Appendix A and the appendices that follow describe the institutions of the gas market in detail.

rise in spot gas prices could have an incentive to curtail customers. With it, the LDC has an incentive to reserve more transportation capacity and to release less of that capacity to others. Since pipeline contracts determine scheduling priority, a rule that protects residential gas customers can crowd out gas-fired generators, whose gas needs rise sharply during extreme cold. Generators have reported fuel shortages during major cold waves, and grid operators have identified reliance on non-firm gas transportation as a winter reliability risk (Safavi, Wu, Lo Prete and Ünel, 2025; Freeman, Apt and Moura, 2020; Borenstein, Busse and Kellogg, 2012).

We develop a model of ex-ante capacity allocation and ex-post gas trade that formalizes these two frictions, and we test its predictions with data covering the universe of U.S. interstate pipeline contracts and capacity releases. The model features a single congested pipeline, two types of shippers—a power plant and an LDC—competitive gas supply, and an aggregate weather shock. Capacity is traded before the shock is realized, in a market that resembles the capacity release market. Gas is traded afterward, at a spot price that clears the market given each shipper's capacity holdings. The first-best benchmark equalizes marginal valuations of gas across shippers in every weather state. The model delivers two results.

First, the ex-ante capacity market implements the first best in every weather state if and only if the ranking of users' marginal valuations does not change with the weather. We call this the state-invariant marginal-utility ranking condition. When it holds, the shipper who should receive priority is the same in every state, so a state-independent capacity contract can implement the efficient allocation. When it fails, no allocation chosen before the weather is known can be efficient in every state, even in a perfectly competitive capacity market. This failure is especially relevant during cold waves: gas-fired peaking plants may value gas little in mild weather but highly in extreme weather, whereas residential heating demand rises more smoothly with the cold. Thus, the deliverability constraint alone can prevent efficient ex-post reallocation, without market power or other contractual imperfections.

Second, the capacity obligation distorts the allocation even when the ranking condition holds. We represent the obligation as a penalty proportional to the consumer surplus lost by an LDC's customers when they are curtailed, with a parameter governing the stringency of regulation. The penalty drives a wedge between the LDC's private valuation of capacity and the true marginal value of gas to its customers, a wedge that increases with weather severity. As a result, the LDC over-contracts capacity and the power plant holds less than in the first best. The incidence varies with the weather: residential gas customers are over-served in mild and moderate states, while the power plant is under-served in extreme states. When the ranking condition also fails, the distortions compound. Regulation therefore does

more than insure gas customers against peak demand: by changing contractual priority and discouraging capacity release, it shifts part of the cost of that insurance onto generators and the households that depend on electricity for heat.

Our empirical analysis examines major cold waves over 2014–2023 rather than a single event. We identify cold waves using daily temperatures in the 3,233 counties of the contiguous United States, standardized relative to each county's climatology, and determine when each cold wave reaches each of the 129 economic areas that define our markets. The procedure, calibrated on four documented cold waves, identifies 19 cold waves, ten of which are major, covering at least 40 percent of the country at their peak. We then follow capacity release transactions by the week in which they are signed and the market they serve. In a balanced panel of market-weeks, we compare a market during its own cold weeks with that market before the shock and with markets that are not cold in the same week.

The evidence shows little reallocation of transportation rights once a major cold wave arrives. One-month releases, which account for about 80 percent of released capacity and are typically signed at month-end, show no response in either quantity or transaction probability. In the market for releases lasting one to seven days, adjustment occurs in prices rather than quantities: the probability of a short-term release falls by 3.2 percentage points, about 40 percent of its baseline, while the tariff on releases that do occur rises by 0.84 log points, more than doubling. Released quantities do not rise for any contract type, and there is no anticipatory response in the weeks before the cold wave. Capacity holders retain their priority rights as their own demand increases, leaving non-priority users facing a thinner and more expensive market at the moment of greatest need. We do not find comparable responses to minor cold waves.

Our paper is most closely related to the economics literature on the allocation of pipeline capacity and access to network infrastructure (Avalos, Fitzgerald and Rucker, 2016; Oliver, Mason and Finnoff, 2014; Brown and Yücel, 2008; Marmer, Shapiro and MacAvoy, 2007; Cremer, Gasmi and Laffont, 2003). This work has emphasized physical bottlenecks and the rules governing pipeline access. We show that even conditional on a given network, the timing of capacity allocation and asymmetric regulatory obligations can prevent scarce transportation rights from reaching the users who value them most during a demand shock. Our contribution is to characterize when an ex-ante capacity market can implement the efficient allocation and to document how its secondary market responds when demand spikes.

We also contribute to the literature on reliability requirements in energy markets and the regulation of utilities (Borenstein, Bushnell and Mansur, 2023; Bushnell, Flagg and Mansur, 2017; Joskow and Tirole, 2007). Peak-day obligations are intended to protect households that buy gas from regulated utilities. Our analysis identifies an indirect incidence of these

requirements outside the regulated customer base: by increasing utilities' incentives to acquire and retain priority capacity, they can reduce fuel availability for electricity generators and thereby shift weather risk toward households relying on electric heating. Finally, our results speak to evidence on gas-fired generators' exposure to winter fuel shortages (Freeman et al., 2020; Borenstein et al., 2012) and to the gas–electricity market coordination problems documented in the policy report by Safavi et al. (2025). We provide a contractual mechanism linking ex-ante priority, regulatory obligations, and the limited adjustment of the secondary market during extreme weather. More broadly, the paper informs the design of markets in which long-term rights finance infrastructure while short-run shocks call for reallocating its use.

The rest of the paper is organized as follows. Section 2 describes how pipeline capacity is allocated and introduces the two frictions. Section 3 describes the data. Section 4 documents the frictions: the prevalence of cold waves, the stickiness of residential prices behind the capacity obligation, and the lead times behind the deliverability constraint. Section 5 develops the model and Section 6 its empirical counterpart. Section 7 presents the evidence from the cold waves of 2014–2023, and Section 8 concludes.

# 2 Industry Background and Key Frictions

This section describes how transportation capacity on U.S. interstate natural gas pipelines is allocated and introduces the two frictions at the center of the paper. Both frictions exist in an environment with priority contracting, in which shippers use capacity contracts to have priority when scheduling gas on a pipeline. The first friction, the *deliverability constraint*, underlies a fundamental need for planning in this market. Shippers must organize the entire transportation logistics to move gas from an origin to a destination, which may involve passing through dozens of pipeline segments and even through multiple pipelines. Beyond logistical challenges, capacity contracts confer scheduling priority, so last-minute interruptible service is seldom available during high-demand periods. As a result, shippers secure capacity before they know how much gas they will need, typically at least one week in advance. The second friction, the *capacity obligation*, is imposed by regulated gas utilities. State regulators require local distribution companies (LDCs) to plan for peak-day demand, and the only way to meet that obligation is to hold contracted capacity. We describe the relevant institutions here. Section 4 then documents both frictions in our data.[2]

[2]In the appendix we provide a more detailed and comprehensive institutional background: the market structure and regulation of local distribution companies (Appendix A), the rate regulation of interstate pipelines (Appendix B), how the capacity release market works in practice (Appendix C), and our interviews with practitioners (Appendix D).

## 2.1 How Pipeline Capacity Is Allocated

Natural gas transmission pipelines provide the bulk of natural gas transportation in North America, moving gas hundreds or thousands of kilometers from production basins to three main demand markets: electricity generation, local distribution companies serving residential and commercial heating, and industrial consumers. We focus on interstate transmission pipelines regulated by the Federal Energy Regulatory Commission (FERC). During our sample period, from 2014 to 2023, the set of operating pipelines was stable, with 69 interstate pipelines in 2014 and 70 in 2023, and only three entries and two exits, all involving small regional systems. Figure 1 maps the pipeline network. Different colors represent different interstate pipeline networks. Given the large number of companies shown, we omit the legend identifying individual company names.[3]

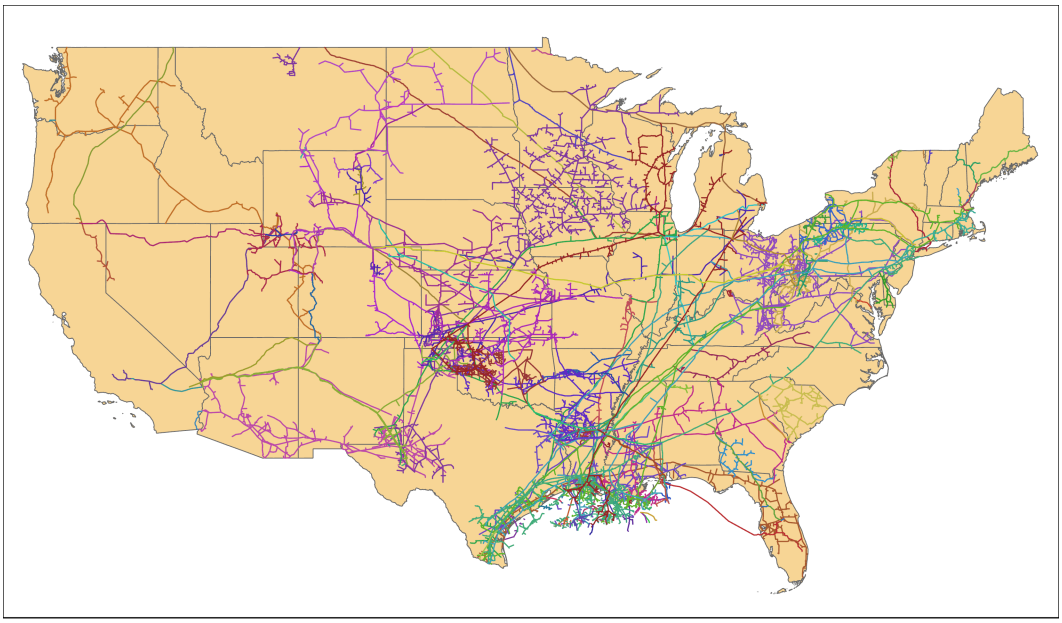


**Figure 1:** Natural Gas Pipelines in the U.S.

Source: The U.S. Energy Information Administration

**Long-term contracts in the primary market.** Transmission pipelines are extremely costly to build, averaging $5 million per kilometer, and take ten to twenty years to break even. To share the risk of these sunk investments, pipelines sign long-term contracts, typically of ten years or more, with shippers prior to construction. These contracts entitle the shipper to a reserved quantity of capacity, usually between specific receipt and delivery points, in

[3]The geographic information for the pipeline system is obtained from the U.S. Energy Information Administration. We infer entry and exit from the quarterly index of customers that FERC-regulated pipelines are required to publish. Section 3 describes the data.

exchange for a reservation charge that is due whether or not the capacity is used.[4] Pipeline rates are regulated by FERC under cost-of-service principles: rates are set so that the pipeline recovers prudently incurred costs and earns a reasonable return on its rate base. Appendix B provides further details.

**Decentralized logistics and priority scheduling.** Unlike electricity markets, in which a centralized grid operator organizes transportation by consolidating supply and demand bids, shippers on natural gas pipelines are themselves responsible for organizing transportation logistics, including the exact route that gas takes from origin to destination. Two features explain this market structure. First, supply and demand need not be continuously equal on a pipeline: a pipeline can be in net inflow or net outflow for several hours within a well-defined pressure window, whereas an electricity grid requires a system operator to balance supply and demand continuously. Second, natural gas marketers have built a business out of arranging transportation for industrial consumers, power plants and other end users, and have an interest in maintaining this structure.

In this decentralized system, long-term contracts function as priority contracts. Each gas day, shippers nominate the quantities they want scheduled in a day-ahead cycle and in several intraday cycles. Shippers holding capacity from long-term contracts, or from the secondary market described below, have priority in every cycle. Shippers without contracted capacity can only use interruptible service, which is scheduled on whatever capacity remains after firm nominations and is curtailed first when the pipeline is constrained. Interruptible service is inconsistent in normal times and is typically unavailable during periods of high demand (Safavi et al., 2025).[5]

**The secondary market for capacity release.** While long-term contracts reduce investment risk, they make the allocation of capacity rigid relative to short-run fluctuations in demand. The secondary market, known as capacity release, emerged to address this rigidity. In this market, holders of long-term contracts can sell all or part of their contractual rights for durations ranging from a few days to the remaining life of the contract. Capacity obtained through a release carries exactly the same scheduling priority as capacity from a long-term contract. Figure 2 illustrates the relationship between the two markets: primary-contract holders who sell capacity are called releasers, and their counterparties are called replacers. Appendix C describes the mechanics of a release, including the posting and bidding rules

[4] A firm transportation contract gives a shipper the right to have the pipeline receive a specified quantity of gas at one or more receipt points and deliver it at one or more delivery points over a specified period. Capacity is either specific as to both location and quantity, or general as to location and specific as to quantity.

[5] The industry, regulatory and government documents behind the institutional facts in this section are cited in Appendix A.

mandated by FERC.

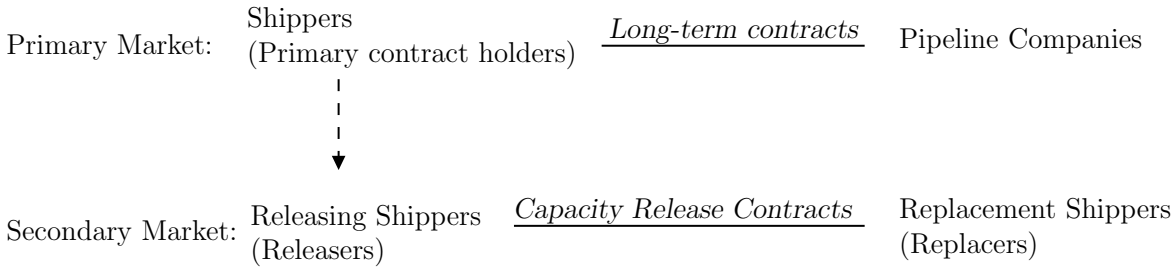


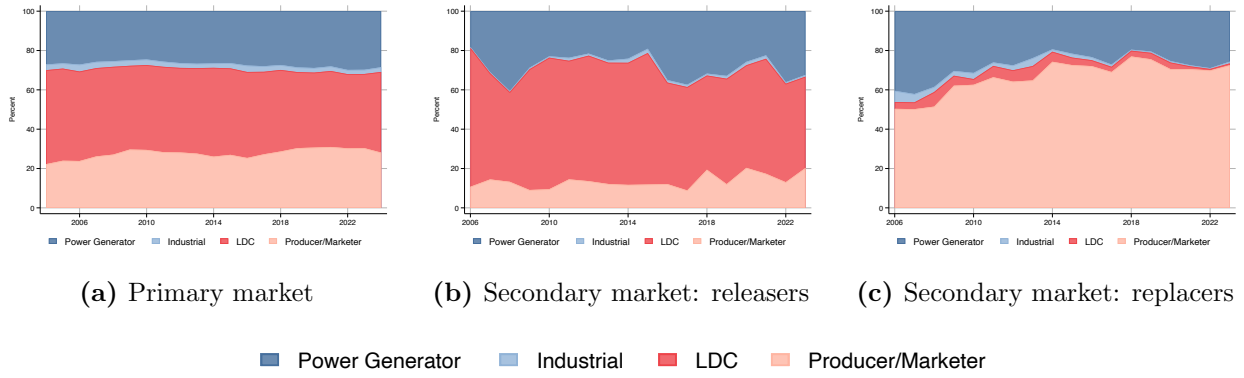

**Figure 2:** Relationships of the primary and the secondary market

**Who holds capacity and who trades it.** Figure 3 shows the composition of contracted capacity by shipper type in the primary market and on each side of the secondary market. Gas utilities hold the majority of primary-market capacity and are the dominant releasers. Marketers are the dominant replacers. Many natural gas end users, in particular electric generators, industrial facilities and gas retailers, hold little or no capacity of their own, primarily because the credit requirements for firm service are steep and because take-or-pay reservation charges appear as debt-like obligations on their balance sheets.[6] These users rely on marketers to arrange both gas supply and transportation, and marketers in turn rely on the secondary market for capacity. Because gas utilities do not need their full capacity for most of the year, they release it to marketers, who ultimately transport gas for power generators and other end users. This division of roles is stable over time and is central to the paper. On the sell side of the secondary market are regulated gas utilities that face the capacity obligation described below. On the buy side are marketers serving gas-fired generators that, during a cold wave, provide heating to the same households through electricity: electricity supplies roughly 43% of residential heating services in the U.S., compared with 47% for natural gas.

## 2.2 Friction 1: The Deliverability Constraint

During cold waves, scheduled gas flows increasingly rely on long-term or capacity-release firm contracts. Interruptible services, which would normally accommodate last-minute demand, are among the first to be curtailed as demand from firm contract holders rises. A shipper that wants to schedule gas delivered during a cold wave must therefore hold contracted

[6] A typical credit requirement for firm service is a security guarantee equal to three months of reservation charges at the maximum tariff rate for the full contracted volume, provided as a cash deposit, a standing letter of credit, a security interest in collateral, or a guarantee from a creditworthy parent. Because firm contracts are take-or-pay, financial institutions treat them as debt obligations, which affects the credit metrics of shippers that hold large amounts of capacity.




**(a)** Primary market **(b)** Secondary market: releasers **(c)** Secondary market: replacers

**Figure 3:** Composition of Contracted Capacity by Shipper Type

*Notes:* Each panel shows the share of contracted capacity (measured at delivery points) held by each shipper type over time. Panel (a) uses long-term contracts from the index of customers. Panels (b) and (c) use capacity release contracts, split by releasing and replacement shipper. Shippers that we cannot classify are excluded. Section 3 describes the classification.

capacity through firm contracts. In this section, we show that shippers must sign these firm contracts before the cold wave arrives. We call this need for advance planning the deliverability constraint.

Two institutional features underlie a need for planning. First, capacity release transactions are multidimensional. Pipeline capacity is defined over narrow segments between receipt and delivery points, and a release must specify the segments, the quantity, the duration, the rate, and options such as recall rights. Gas that crosses several pipeline systems requires many such contracts. Buyers and sellers must agree on all of these dimensions, and the market for any particular combination can be thin. Second, capacity releases involve administrative and coordination frictions. Arranging the appropriate path, quantity, duration, and contractual terms also requires coordination. Section 4 shows that the median lead time between the signing of a release and its effective date is about one week, and that shippers adjust their positions in the week before a cold wave reaches their market. In the model of Section 5, the deliverability constraint requires the market to allocate capacity ex ante.

These deliverability constraints have two related consequences. Capacity is allocated one to two weeks before demand is realized, at a time when the severity of a cold wave is uncertain. Major recent cold waves, including Winter Storm Uri in February 2021 and Winter Storm Elliott in December 2022, were under-forecast by between 9 and 13 percent (Safavi et al., 2025). This uncertainty could cause a mismatch between who needs gas the most and who has pipeline capacity. Moreover, the shippers most exposed to the constraint are those that hold the least capacity. Gas-fired generators are frequently dispatched only when electricity demand peaks, and many of them cannot recover the fixed cost of firm

transportation through their electricity market bids, so they rely on interruptible or short-term arrangements. Grid operators have repeatedly identified this reliance as a source of winter reliability risk: the New York Independent System Operator reports declining winter capacity margins once generation without firm fuel arrangements is excluded, and the four largest regional transmission organizations note that gas supply becomes illiquid for generators precisely during cold periods and holiday weekends, when large buyers have already secured their needs in advance.

## 2.3 Friction 2: Capacity Obligations on Gas Utilities

The second friction concerns the largest holders of pipeline capacity, the gas utilities/local distribution companies (LDCs). Three facts about LDCs matter for our analysis. Appendix A documents each of them State by State. First, LDCs are natural monopolies. In every state we study, a handful of investor-owned utilities (IOUs) serve non-overlapping territories and account for the large majority of residential and commercial volumes. Second, LDCs are rate-regulated by State Public Utility Commissions. Retail rates must be "just and reasonable," are set through cost-of-service proceedings, and can only be changed after notice and hearing. Natural gas costs can be passed through to customers, but with a lag and subject to prudence review, so residential prices move slowly relative to wholesale prices (Davis and Muehlegger, 2010). Third, because retail prices are sticky while wholesale prices can spike during cold waves, residential consumers do not see the wholesale price signal and their demand may not match aggregate market conditions. An unregulated LDC would have an incentive to curtail its customers to avoid losses. Regulators respond with capacity obligations to ensure residential heating demand is met.

**Design-day planning and curtailment priority.** The capacity obligation takes two forms. The first is a planning requirement. Regulators require each LDC to demonstrate, typically through an annual filing, that it holds enough firm supply and transportation capacity to serve its firm customers on a design day, a hypothetical day of extreme cold defined by a return period such as one in thirty or one in ninety years. Between 2014 and 2018, for example, the New York Public Service Commission reviewed each utility's supply and capacity plan through annual winter supply reviews, and in 2020 it opened a formal gas planning proceeding requiring long-term plans on a three-year cycle. New Jersey utilities set their own design-day standards until 2019, when conflicting capacity assessments led the Board of Public Utilities to commission an independent study and, in 2022, to adopt a standardized emergency playbook. Illinois requires utilities to report their forecasted design-

day demand, available peak-day supply, and reserve margin in every rate case.[7]

The second form is a curtailment priority order. Utility tariffs and State rules specify the order in which customers are curtailed when supply is short. Residential and other "human needs" customers are placed at the top in every jurisdiction we have examined (Safavi et al., 2025). The priority order matters for our analysis because it defines whom the obligation protects, namely residential and other human-needs customers: an LDC cannot meet its obligation by curtailing these customers and selling gas at the spot price. Gas-fired generators that take service inside an LDC's city gate are also subject to this order, and an LDC may curtail deliveries to such a generator, even one holding firm service, in order to serve residential demand.[8]

The main takeaway is that, while well-intentioned, these curtailment priorities may, in fact, protect some "human needs" customers at the expense of other "human needs" customers, who get heating services from electricity generated with natural gas. At the extreme, curtailment to gas generators may even prevent residential gas heaters from turning on if natural gas shortages in the electricity market are causing blackouts.

**Enforcement.** Formal penalties for failing to meet peak demand are rare. What disciplines the utility is the combination of prudence review, cost-recovery risk, and political exposure that follows a curtailment or a purchase of gas at extreme spot prices. As Borenstein et al. (2012) document, the severity of that review depends on whether the utility can show it made a good-faith effort to secure supply in advance, and a utility that has released capacity before a cold wave cannot make that argument. The aftermath of Winter Storm Uri, documented in Appendix A.2, illustrates the stakes: utilities in Texas, Oklahoma, and Kansas could recover their extraordinary procurement costs only through special legislation enacted after the storm. In the model of Section 5, we summarize these enforcement channels with a single stringency parameter that scales the loss the LDC internalizes when its customers are curtailed.

The stringency of the obligation varies across states and over time. Before 2021, New York, New Jersey and Illinois operated formal design-day reviews, whereas Texas, Oklahoma and Kansas imposed no capacity mandate on gas utilities and, in the case of Texas, maintained an explicit policy of market-based supply determination. After Uri, all of these states tightened their regimes, though in different ways. Appendix A classifies the regimes and their changes, and we exploit this variation in the empirical analysis.

[7] Appendix A.3 documents these rules state by state and over time, with the orders, statutes and studies behind them.

[8] These generators do not contract for interstate pipeline capacity and do not appear in our data, so we abstract from them. Our analysis concerns generators, and the marketers serving them, that contract directly for pipeline capacity and compete with LDCs in the capacity market.

## 2.4 From Capacity Obligations to Regulation-Induced Priority

Combining the two frictions yields the mechanism that the rest of the paper studies. Because of the deliverability constraint, the only way for an LDC to satisfy its capacity obligation is to hold contracted capacity when a cold wave arrives. Capacity obligations raise the value the LDC attaches to capacity, and it does so most in the states of the world in which capacity is scarce. We expect two behaviors to follow. The LDC over-contracts: it holds more capacity than it would absent the obligation, as a precaution to avoid curtailing customers and paying a regulatory penalty. In the secondary market, the LDC under-releases: it releases less capacity to marketers than it would absent the obligation as a cold wave approaches, and it retains recall rights on the capacity it does release.[9]

The shippers that are crowded out are those with the least capacity, namely marketers and the gas-fired generators they serve. Industry participants describe similar outcomes from recent cold waves. Following Uri, generators told the Railroad Commission of Texas that curtailment rules had prioritized LDCs and left insufficient gas for power plants. After Winter Storm Elliott, electric cooperatives testified to FERC that LDCs prioritized retail customers over generators, and that generators face a choice between paying for gas they may not use and paying penalties for failing to generate. Grid operators in New England attribute winter fuel scarcity in part to competing demand from LDCs, and the joint report of ISO New England, MISO, PJM and SPP calls on states to reconsider the definition of "human needs" customers on the grounds that households served by electric heat depend on gas-fired generation for the same health and safety reasons.[10]

We emphasize that this mechanism does not require any strategic withholding by utilities. A 2017 study argued that New England utilities had withheld pipeline capacity to raise prices (Marks et al., 2017); a subsequent FERC staff inquiry found no evidence of such behavior. Our argument is different. LDCs respond to a regulatory obligation in the way the obligation intends, and the allocation that results shifts the risk of extreme weather onto other pipeline users. Whether this shift is costly depends on how the marginal value of gas compares across users in extreme states, which is the question the model in Section 5 addresses. Before turning to the model, Section 4 documents both frictions, and the resulting over-contracting and under-release, in our data.

[9] Recall provisions are common in capacity release contracts, and the probability that a releaser exercises them is highest during scarcity periods such as winter storms. Pipelines do not report whether a release was recalled (Safavi et al., 2025). Appendix C discusses recall and reput options.

[10] These comments, testimonies and reports are cited in Appendix A.4.

# 3 Data

We use five datasets. First, the index of customers lists every contract on a pipeline system: the shipper that holds it, its start and end dates, the type and path of the service, the amount of capacity the shipper may use, and the contract rate. Pipelines regulated by the Federal Energy Regulatory Commission (FERC) must publish a quarterly index of customers, from which we obtain these data. Second, the capacity release data from each pipeline record every secondary-market transaction, including the transaction date, the type and duration of the contract, the rate, and options such as recall and reput rights.[11] Third, we observe daily spot prices of natural gas at all U.S. hubs, compiled by Capital IQ Pro. Fourth, we obtain from Capital IQ Pro the full history of daily throughput between each origin and destination point on each interstate pipeline, which gives the daily utilization of physical capacity within each system.[12] These four datasets cover 2014 to 2023. Last, we use geographic data from the U.S. Energy Information Administration (EIA) on the interstate pipeline network and the location of major gateways and hubs, and daily average temperature by U.S. county from the National Oceanic and Atmospheric Administration (NOAA).

For the index of customers, capacity release and throughput data, we identify the precise origin and destination points of each contract and aggregate them to U.S. counties, to economic areas, and to the nearest natural gas hub. This spatial mapping lets us merge each origin and destination point with the spot price at its nearest hub and with county-level daily temperature, from which we construct the extreme weather events used in the analysis (Appendix F).

We classify each shipper as a gas utility (local distribution company), a producer or marketer, a power generator, or an industrial user. Figure 3 uses this classification; shippers that we cannot classify are excluded there. In the secondary market, we also classify each contract by the history of its two parties: contracts between shippers that have met fewer than five times in the past are new relationships, and contracts between shippers that have met five or more times are long-term relationships.

Table 1 reports summary statistics for both markets. In the primary market, the average long-term contract lasts about eleven years, the median ten, and a typical pipeline has about a hundred shippers. Secondary-market contracts are short: the median contract lasts one month, some last half a year or a year, and few last longer. The secondary market

[11] With the advent of the capacity release market, FERC required pipelines to post the deals that their service requesters sought to transact (FERC Order No. 636, et al.).

[12] Although these data are public, most pipelines periodically delete past transactions, such as contracts signed in the capacity release market. We purchase the full history of capacity releases, throughput and hub spot prices from Capital IQ Pro.

is nonetheless economically important. Capacity obtained through a release carries the same scheduling priority as a long-term contract, and across pipeline-quarters the quantity released averages 32 percent of the quantity held under long-term contracts, with a median of 7 percent. Participation is broad: 46 percent of the 3,719 shippers that hold long-term contracts release capacity at least once.

**Table 1:** Summary Statistics: Primary and Secondary Markets

*Panel A: Primary market (index of customers)*

| | Shippers per pipeline | Contract duration (months) | Shippers per O–D hub pair |
|---|---|---|---|
| Mean | 97 | 140 | 30 |
| 5th percentile | 2 | 10 | 1 |
| 25th percentile | 12 | 41 | 3 |
| Median | 64 | 120 | 10 |
| 75th percentile | 110 | 220 | 32 |
| Observations | 72 pipelines<br>3,739 shippers | 382,983 contracts | 38 hubs<br>547 O–D pairs |

*Panel B: Secondary market (capacity releases)*

| | Releasing shippers per pipeline | Replacement shippers per pipeline | Contract duration (days) |
|---|---|---|---|
| Mean | 87 | 113 | 125 |
| 5th percentile | 6 | 8 | 2 |
| 25th percentile | 29 | 35 | 29 |
| Median | 61 | 56 | 30 |
| 75th percentile | 129 | 163 | 30 |
| Observations | 35 pipelines<br>1,710 releasing shippers | 35 pipelines<br>1,297 replacement shippers | 491,797 contracts |

*Notes:* All data start in 2014. Each contract specifies several origin and destination points, so the primary-market data contain 1.9 million contract-point observations and the secondary-market data 1.7 million. Origin-destination (O–D) hub pairs assign each origin and destination point to its nearest natural gas hub. Seventy-four percent of primary-market and 67 percent of secondary-market contracts are spatially identified; the remainder are pooled contracts without specific locations or contracts that we cannot locate. The statistics include all contracts, whether spatially identified or not.

# 4 Growing Concerns and Motivating Facts

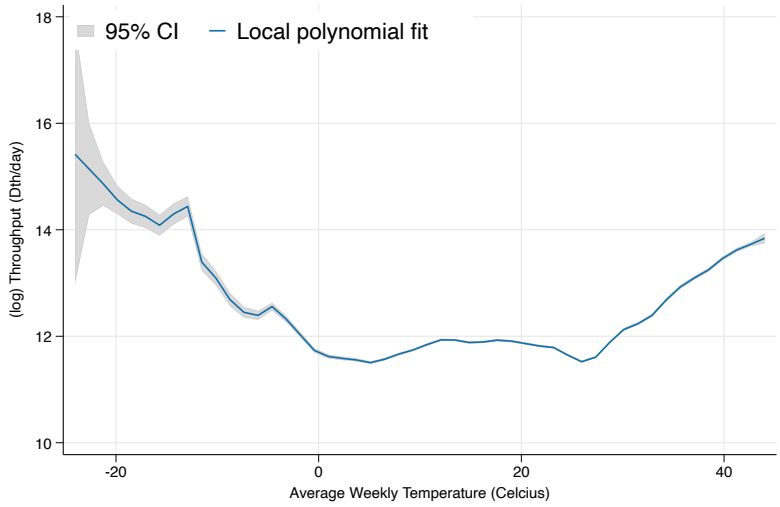
In this section, we first document why cold waves are a growing concern for this industry. We then use our data to document the two frictions introduced in Section 2. Whereas Section 2 described the institutions that generate the deliverability constraint and the capacity

obligation, here we show what they look like in the data: how far in advance capacity is contracted, how far the spot price departs from the residential price during cold waves, and how the positions of gas utilities and gas-fired generators move as a cold wave approaches.

## 4.1 Prevalence of Extreme Cold Waves

Extreme cold waves have become increasingly common in the U.S. due to climate change (Wang, Biasutti, Byrne, Castro, Chang, Cook, Fu, MGrimm, Jaha, Hendon et al., 2021; Bathiany et al., 2018). Cold waves put multiple pressures on natural gas demand. First, almost 50% of residential heating comes from natural gas. Almost all the remainder comes from electricity, and 45% of electricity in the U.S. is generated with natural gas. Less than 10% of residential heating does not depend directly or indirectly on natural gas.

**Figure 4:** Relationship Between Gas Throughout and Temperature




Additionally, extreme cold waves are volatile, with uncertain realizations in magnitude across space and over time. Many major cold waves in recent history, including the February 2021 Winter Storm Uri that caused 200 deaths and the 2022 Winter Storm Elliott, were associated with significant under-forecast error, between 9.2% and 12.5% (Safavi et al., 2025). When temperatures are already low, even marginal changes in temperature can cause large shifts in natural gas demand because heating demand is most sensitive at those temperature. See Figure 4, where we show the typical U-shaped relationship between temperature and gas throughput.

## 4.2 The Capacity Obligation in the Data: Sticky Residential Prices

Section 2.3 described the capacity obligation and its rationale: residential consumers buy gas from rate-regulated LDCs at prices that adjust slowly, so an unregulated LDC would have an incentive to curtail them when the spot price rises. Figure 5 shows how large the gap can be. It plots the spot price at the hub serving Massachusetts against the residential price paid by consumers in the state. Outside of cold waves the two series track each other. During cold waves the spot price rises far above the residential price, and the residential price responds only with a lag, consistent with the pass-through mechanics documented in Appendix A.2.

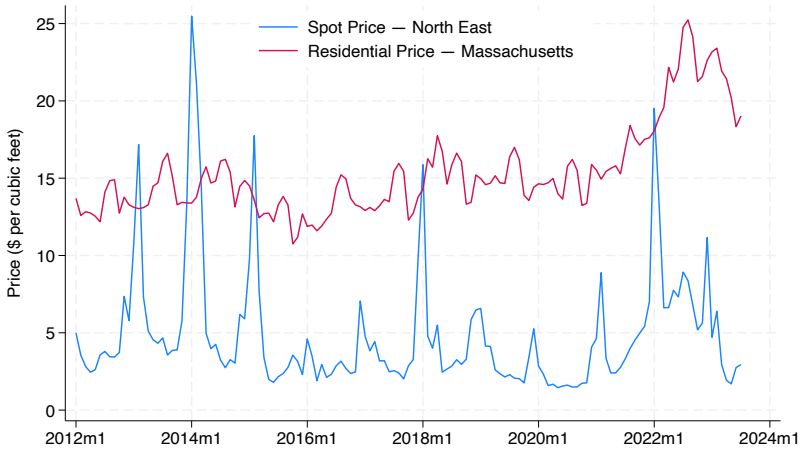



**Figure 5:** Gas Prices (Spot vs. Residential) — Massachusetts

In such states of the world, an LDC that sold gas at the residential price would lose money on every unit delivered, and would recover the loss, if at all, only after a prudence review. This is the situation in which the capacity obligation binds: the LDC must hold enough contracted capacity to serve its firm customers on a design day, and it is held responsible if it does not. As documented in Section 2, states have tightened these requirements since 2019, and after Winter Storm Uri several states that previously imposed no capacity mandate introduced new obligations on gas utilities.

## 4.3 The Deliverability Constraint in the Data: Lead Times

While the capacity obligation may protect some residential consumers, it may induce gas utilities to over-contract pipeline capacity as a precaution and to under-release it in the secondary market, giving them de facto priority and effectively transferring risk to other market participants, particularly other residential consumers who get heating services from

electricity. Whether it does so depends on the deliverability constraint: the market for pipeline capacity is not liquid ex post. On the day that the cold wave hits, the pipeline allocation is essentially locked. Interruptible services are seldom available in periods of high demand (Safavi et al., 2025), and almost all pipeline scheduling comes from long-term and secondary market (capacity release) contracts. We now document how far in advance those contracts are signed.

In Table 2, we show that the lead time between the date a contract is signed and the date it comes into effect averages 7 days. These capacity release contracts are multidimensional. Pipeline capacity is auctioned off in narrow segments, and contracts often feature dozens of those segments, necessary for gas to reach a destination of interest. Additionally, buyers and sellers must agree on quantity, length of time capacity will be released, and specific delivery and receipt locations. The market for what a given participant wants may be very thin, and we observe significant planning as a result.

**Table 2:** Distribution of Capacity Release Delay

| | $10^{th}$ | $25^{th}$ | Median | $75^{th}$ | $90^{th}$ |
|---|---|---|---|---|---|
| Delay (Days) | 2 | 4 | 7 | 9 | 11 |

In this context, where pipeline capacity allocations are made one week ahead, two potential issues can arise in anticipation of a cold wave. First, regulated gas utilities, which own primary market contracts, may be excessively risk-averse and unwilling to release capacity to other users, raising the price for pipeline capacity and preventing smaller gas power plants from securing enough capacity for themselves. During the 2021 Winter Storm Uri, 20% of all unplanned power outages were attributed to gas-fired power generators not getting sufficient gas supply (Safavi et al., 2025). Thus, peak day demand regulation may effectively displace the health risk that cold waves pose to other households rather than reduce it.

Second, and more generally, capacity release contracts are state-independent contracts with linear pricing. To the extent that demand profiles may vary across realizations of a cold wave, capacity release contracts may not allocate capacity to shippers who need it the most. For example, consider a market with many gas peaking power plants. These power generators may have a large need for pipeline capacity relative to other shippers if electricity demand from the cold wave is higher than expected, but a low need for pipeline capacity relative to other shippers if electricity demand turns out lower than expected. By virtue of being state-independent, capacity release contracts cannot simultaneously allocate capacity efficiently under both scenarios. In the next section, we formalize those ideas in a theoretical model to study the implications of these frictions on allocative efficiency in the market for

pipeline capacity and the market for natural gas. We then test the predictions of our model by leveraging the staggered rollout of major recent cold waves and study alternative market designs.

# 5 A Model of Capacity Allocation

In this section, we develop a parsimonious model of pipeline capacity allocation. We ask whether a capacity-release market that allocates capacity before a weather shock is realized can implement the socially optimal allocation of gas once the weather shock is realized. We first lay out the primitives: market structure, participants, and preferences. We then define the first-best benchmark (Section 5.1) and the market's two frictions (Section 5.2), before turning to the decentralized market allocation with and without those frictions (Section 5.3 and 5.4).

**Participants and Preferences.** Allocations are defined relative to an aggregate demand (weather) shock $A \in \Omega_A$ with density $f_A$. We assume a single-line, single-origin, single-destination network with total physical capacity $K$. We assume a one-to-one correspondence between gas consumption ($q$) and pipeline capacity ($\kappa$). Both are needed to ship and consume gas. We define three types of market participants:

1. **Power plants (type-1 agents).** They represent the unregulated shippers: a composite of gas-fired power plants, marketers, industrial users, and retailers. They respond to spot market prices.

2. **Gas Utilities / LDCs (type-2 agents)**. They are regulated natural monopolies that serve end-use residential and commercial customers. They offer fixed-price contracts and face capacity obligation in the form of peak-day demand regulation.

3. **Gas Suppliers.** They form a competitive fringe that can produce gas quantity $Q$ at cost $C(Q)$ with strictly increasing marginal cost $c(Q) \equiv C'(Q) > 0$, $C'' > 0$.

We index the two representative types by $i \in \mathcal{I} = \{1, 2\}$ (power plant $i = 1$, LDC $i = 2$). However, the analysis extends to $N$ heterogeneous shippers. Agent $i$ has period utility $u_i(q, A)$ that is continuous, strictly increasing, strictly concave, and invertible in gas consumption $q$, and increasing in the shock. We write

$$P_i(q, A) \equiv \frac{\partial u_i(q, A)}{\partial q} > 0$$

for agent $i$'s marginal utility (equivalently, its inverse demand for gas). Strict concavity makes $P_i$ strictly decreasing in $q$, and the shock raises the willingness to pay: $P_i$ is strictly increasing in $A$. For LDCs, preferences should be interpreted as the preferences of their end-use customers. This distinction becomes important when introducing market frictions, as LDCs and retail customers may not share the same incentives, which we discuss in detail in Section 5.4. For the purpose of establishing a first-best allocation, LDCs and their end-use customers can be used interchangeably, and there is no difference between type-1 and type-2 agents.

Throughout the exposition, we use the constant relative risk aversion (CRRA) family of utility functions as a leading parametric example,

$$u_i(q, A) \;=\; A^{\theta_i}\,\frac{q^{\,1-\gamma_i}}{1-\gamma_i}, \qquad \Longrightarrow \qquad P_i(q, A) \;=\; A^{\theta_i}\,q^{-\gamma_i}, \tag{1}$$

where $\theta_i$ governs the sensitivity of marginal utility to the weather shock and $\gamma_i$ the curvature (risk aversion).

## 5.1 The First-Best Benchmark

We begin by characterizing the planner's problem, which serves as the first-best benchmark. After $A$ is realized, a social planner allocates deliverable gas to maximize total surplus net of production costs, subject only to the physical pipeline capacity constraint $K$, ensuring that gas consumption never exceeds it. Within this constraint, the planner can freely allocate pipeline capacity $\kappa_i(A)$ and gas quantity $q_i(A)$ across market participants for different weather shocks $A$. There is no wedge between capacity and gas quantity, and $\kappa_i^*(A) = q_i^*(A) \quad \forall i, A$.

$$W^{\mathrm{FB}}(A) \;\equiv\; \max_{\{q_i \geq 0\}} \sum_i u_i(q_i, A) \;-\; C\Big(\sum_i q_i\Big) \qquad \text{s.t.} \qquad \sum_i q_i \leq K. \tag{2}$$

Let $\lambda(A) \geq 0$ be the multiplier on the capacity constraint. The first-order conditions characterize the first-best allocation $q^*(A) = \{q_i^*(A)\}_i$:

$$P_i\big(q_i^*(A), A\big) \;=\; \underbrace{c\big(Q^*(A)\big) + \lambda(A)}_{\equiv\; p^*(A)} \quad \forall i, \qquad Q^*(A) = \sum_i q_i^*(A), \qquad \lambda(A)\,\big[K - Q^*(A)\big] = 0. \tag{3}$$

The shadow price in equation (3) splits into two economically distinct pieces, $p^*(A) = c(Q^*(A)) + \lambda(A)$. The first, $c(Q^*(A))$, is the *resource cost*: the marginal cost of physically

producing and delivering the last unit of gas, which rises with $c' > 0$. The second, $\lambda(A)$, is the *scarcity rent of the pipeline*: the multiplier on $\sum_i q_i \leq K$, i.e. the social value of one more unit of transmission capacity in state $A$.

In low-demand states, the pipe is slack ($Q^* < K$), capacity has no scarcity value, and the efficient price is pure resource cost, $p^* = c(Q^*)$. In high-demand states, demand at resource-cost price exceeds what the pipe can carry, so capacity binds, and quantity is pinned at $Q^* = K$. Figure 6 draws the two regimes.

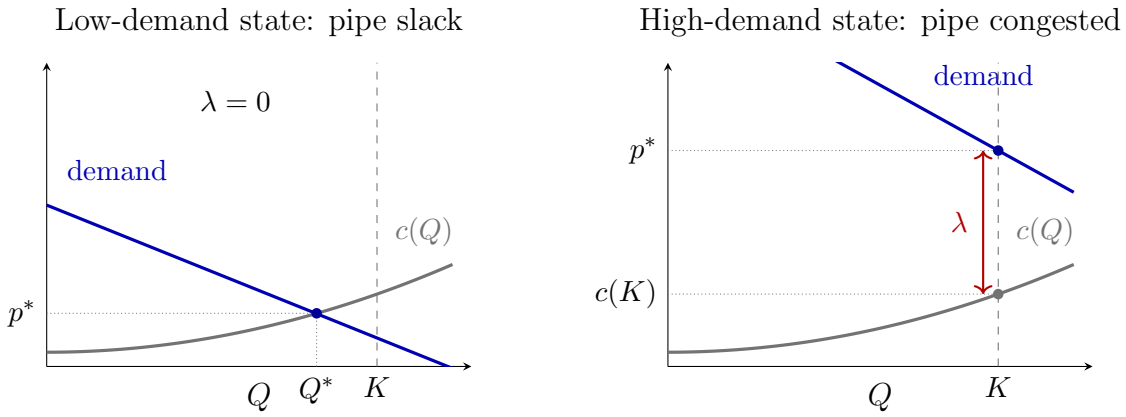



**Figure 6:** The first best in state $A$. *Left:* in a low-demand state, demand meets marginal cost before the pipe fills, so $Q^* < K$, the capacity constraint is slack ($\lambda = 0$), and the efficient price is pure resource cost, $p^* = c(Q^*)$. *Right:* in a high-demand state, desired demand would cross $c(Q)$ beyond the pipe, so quantity is pinned at $K$ and the efficient price rises above resource cost by the scarcity rent, $p^* = c(K) + \lambda(A)$. The decomposition $p^*(A) = c(Q^*(A)) + \lambda(A)$ is the vertical gap on the right.

The first best (i) *equalizes marginal utilities across shippers* at the shadow price $p^*$, and (ii) sets the *aggregate* quantity such that marginal utilities equal the marginal cost of production $p^* = c(Q^*)$ when the pipe is slack and equal the marginal cost of production at capacity plus the scarcity rent $p^* = c(K) + \lambda$ when it binds. In the absence of the frictions discussed in the next Section, it is well known that this allocation can be decentralized with standard scarcity pricing (Cremer et al., 2003).

## 5.2 Decentralized Allocation and Key Frictions

We now define the market structure under which these participants operate and discuss potential market frictions that may prevent the market from reaching the first-best outcome. For any decentralized allocation, we define the welfare loss in state $A$ as

$$\mathcal{L}(A) \;\equiv\; W^{\mathrm{FB}}(A) - W^{\mathrm{dec}}(A) \;\geq\; 0,$$

where $W^{\text{dec}}$ evaluates the decentralized allocation at weather shock $A$, and we report expected welfare loss $\int_{\Omega_A} \mathcal{L}(A)\, f_A(A)\, dA$.

The market we consider contains two key frictions: deliverability constraints and capacity obligations. We proceed in two steps. Section 5.3 studies the distortion generated by the deliverability constraint alone and discusses conditions to restore the first-based allocation. We treat this as an underlying market characteristic that cannot be fully eliminated. Section 5.4 then asks how the market allocation gets further distorted once we introduce capacity obligations imposed on LDCs.

**Friction 1: Deliverability Constraint.** The first fundamental friction underlies the need for planning, which we refer to as the deliverability constraint. For reasons discussed at length previously, pipeline capacity is typically allocated via contracts in two markets: a primary capacity market and a secondary capacity market. Shippers can sign long-term contracts with pipelines to obtain pipeline capacity. Then, typically one week in advance, shippers can release and obtain capacity in a secondary market. Pipeline capacity allocation in both primary and secondary markets provide priority scheduling in the day-ahead and intraday markets. The deliverability constraint comes this priority scheduling. Shippers must procure capacity in advance if they want to secure gas in the future, without knowing how much they will need.[13]

In our model, firm $i$ enters with pipeline capacity from its primary-market (long-term) contracts and can adjust this holding by trading in the capacity-release market. We summarize the outcome of both in a single object: the *net forward position* $\kappa_i$, the total capacity firm $i$ carries into the delivery period. We model the capacity-release market as an *ex-ante* market in which firms adjust their net forward positions before the weather shock $A$ is realized. Once $A$ is realized, pipeline scheduling and gas consumption decisions take place in an *ex-post* market:

1. ***Ex-ante* (capacity-release) market.** Each participant $i$ trade pipeline transmission capacity and chooses a forward capacity position $\kappa_i \geq 0$ at the forward price $\tau_p$, before $A$ is realized.

2. ***Ex-post* (spot gas) market.** The shock $A \in \Omega_A$ is realized, and the desired consumption is revealed. The gas market clears at price $p_g(A)$, but each agent can consume at most its forward-contracted capacity, $q_i(A) \leq \kappa_i$.

We assume a one-to-one correspondence between gas and pipeline capacity. Both are needed to ship and consume gas. All shippers are price takers in both markets. We introduce

[13] In practice, pipelines can use interruptible services to allocate unused pipeline capacity to shippers without contracts. However, in periods of high demand, interruptible services are seldom available.

the deliverability constraint by imposing that shippers can consume only up to their forward position in the ex-post market and cannot trade with each other ex-post. In any state $A$ shipper $i$ may consume at most $\kappa_i$, so

$$q_i(A) = \min\{\kappa_i,\ D_i(p_g(A), A)\}, \qquad D_i(p, A) \equiv \underbrace{P_i^{-1}(p, A)}_{\text{Unconstrained demand}}.$$

Total capacity is fixed at $\sum_i \kappa_i = K$.

**The ex-post gas market.** We model the ex-post spot market as a competitive market for delivered gas served by the price-taking supply fringe. Given some ex-ante capacity allocation $\{\kappa_i\}$ and a weather realization $A$, the spot market clears when aggregate demand $D(p, A)$ equals supply $S(p)$. Aggregate supply is just the inverse of the marginal cost of production. Aggregate demand is the sum of individual truncated demands:

$$D(p, A) \ \equiv \ \sum_i \min\{\kappa_i,\ D_i(p, A)\}.$$

The competitive ex-post price $p_g(A)$ is the unique solution to

$$c^{-1}\big(p_g(A)\big) = D\big(p_g(A), A\big). \tag{4}$$

The solution $p_g(A)$ exists and is unique.[14]

**The ex-ante market for capacity.** Taking $\tau_p$ and the ex-post rule $p_g(\cdot)$ as given, shipper $i$ chooses its forward position to solve

$$\begin{aligned} U_i = \max_{\kappa_i \geq 0} \Big\{ \ \tau_p \left(K_i - \kappa_i\right) + \int_{\Omega_A} \big[\, u_i\big(q_i(A), A\big) - p_g(A)\, q_i(A) \,\big]\, f_A(A)\, \mathrm{d}A \Big\}, \\ \text{where} \quad q_i(A) = \min\big\{\, \kappa_i,\ D_i\big(p_g(A), A\big) \,\big\}. \end{aligned} \tag{5}$$

Differentiating with respect to $\kappa_i$ yields the first-order condition

$$\tau_p = \int_{\Omega_A} \underbrace{\mathbf{1}\{D_i(p_g(A), A) > \kappa_i\}}_{\text{Binding capacity constraint}} \times \big[\ \underbrace{P_i(\kappa_i, A) - p_g(A)}_{\text{Value of additional capacity}}\ \big] f_A(A)\, dA; \tag{FOC-1}$$

This first-order condition is intuitive. It says that the price of capacity equals its scarce value, integrated over all states of the world in which capacity is scarce. A marginal increase

[14] See Lemma 1 in Appendix E.1 for the proof.

in capacity only has value when it is binding, which is why the indicator only selects states of the world in which the shipper is capacity-constrained.

## 5.3 When Can the Decentralized Market Restore the First Best?

We now ask how far the ex-ante capacity-release market alone can go toward decentralizing the first best of Section 5.1.

Once weather shock $A$ is realized, the ex-post market falls into one of three configurations. First, if no firm's capacity constraint binds, firms consume their unconstrained demands and the spot price equals marginal cost. This allocation is first best because marginal utilities are equalized across firms. Second, if some firms are capped while others retain unused capacity, the allocation is never first best: capped firms have higher marginal utility than unconstrained firms, but non-transferable capacity prevents a surplus-improving reallocation. Third, if all firms consume at capacity and aggregate consumption equals $K$, the first best may or may not be restored. Efficiency then depends on whether the ex-ante capacity allocation $\kappa_i$ divides the pipe as the planner would.

We now show that the ex-ante capacity market implements the first-best allocation in fully congested states if and only if first-best consumption shares across shippers are state independent. In other words, shipper $i$ must value gas by the same factor relative to shipper $j$, regardless of $A$. If so, a decentralized market with deliverability constraints can achieve the first best.

Below we define a class of utility functions with this property:

**Definition 1** (SIMR)**.** The family $\{P_i(\cdot,\cdot)\}_{i\in\mathcal{I}}$ satisfies the *state-invariant marginal-utility ranking* (SIMR) property if, for any given pair of agents $(i,j)$ and capacities $(\kappa_i,\kappa_j)\in\mathbb{R}^2_+$,

$$\big[P_i(\kappa_i,A)-P_j(\kappa_j,A)\big]\times\big[P_i(\kappa_i,A')-P_j(\kappa_j,A')\big]\;\geq\;0,\quad\forall A,A'\in\Omega_A.$$

Intuitively, SIMR is a no-crossing condition. The weather shock changes how much gas is worth, but never *who* values it more. Fix any pair of holdings $(\kappa_i,\kappa_j)$, definition 1 requires that the two marginal-utility differences never have opposite signs: if agent $i$'s marginal unit is worth more than agent $j$'s in some state, it is worth weakly more in every state.

A simple sufficient condition for SIMR is that marginal utility factorizes into a state component common to all agents and an agent-specific component,

$$P_i(\kappa,A)=\chi(A)\,\psi_i(\kappa),\qquad\chi(A)>0.$$

**Proposition 1** (First-best implementation)**.** *Suppose that the capacity constraint binds with positive probability. Then, under competitive ex-post pricing, the decentralized allocation coincides with the first best in* every *state $A$ if and only if SIMR holds.*

*Proof.* See Appendix E.1. □

To illustrate the intuition behind Proposition 1, we use an example with two shippers with CRRA preferences and contrast what happens when SIMR holds and when it is violated.

The spot price moves with the shock: while the pipe is slack, $p_g(A)$ equals the marginal cost of the cleared quantity and rises with demand; once every cap binds, the cleared quantity is fixed at $K$ and the price settles at $c(K)$.

**SIMR holds (common $\theta$).** In the CRRA specification in (1), SIMR holds when $\theta_1 = \theta_2 = \theta$:

$$P_i(q, A) \;=\; A^{\theta}\, q^{-\gamma_i}, \tag{6}$$

the marginal utility curves share the shape $A^{\theta}$ and differ only by the agent-specific scalar $q^{-\gamma_i}$. As discussed before, the key configuration in which restoring the first best is in question is full congestion, where all firms consume at capacity and aggregate consumption equals $K$; we therefore plot each shipper's marginal utility at capacity, $P_i(\kappa_i, A) = A^{\theta}\kappa_i^{-\gamma_i}$, against the state.

In equilibrium, the ex-ante market clears at:

$$\tau_p = \underbrace{\int \left[A^{\theta}\kappa_1^{*-\gamma_1} - p_g(A)\right]^{+} dF(A)}_{\text{power plant's excess value}} = \underbrace{\int \left[A^{\theta}\kappa_2^{*-\gamma_2} - p_g(A)\right]^{+} dF(A)}_{\text{LDC's excess value}}. \tag{7}$$

Therefore, in equilibrium, equation (7) indicates that $\kappa_1^{*-\gamma_1} = \kappa_2^{*-\gamma_2} = \Gamma$: each expected excess value is strictly increasing in the shipper's own scalar, so the two integrals can be equal only if the scalars coincide. The marginal utility curves at capacity then collapse onto a single curve, $A^{\theta}\Gamma$, across shippers. Capacity is thus allocated so that, for any realization of $A$, either all shippers are constrained or none of them is. Partial congestion never occurs in this case. Panel (a) of Figure 7 shows an off-equilibrium allocation in which the two curves differ in level; Panel (b) shows the equilibrium outcome.

**SIMR fails (heterogeneous $\theta$).** We focus on the case $\theta_1 > \theta_2$: the power plant's marginal utility is more sensitive to the weather shock than residential demand. This ordering reflects how the two types use gas. Gas-fired generators sit at the top of the electricity supply curve and are dispatched only when electricity demand is high, so their gas consumption is

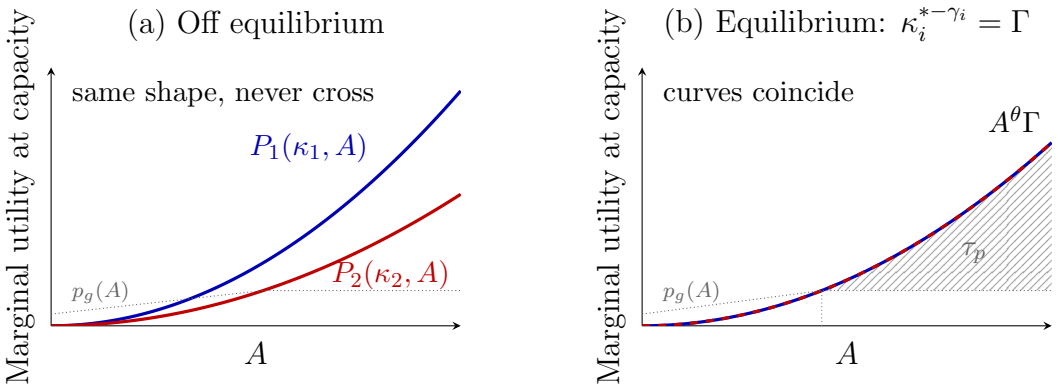



**Figure 7:** SIMR (common exponent $\theta_1 = \theta_2 = \theta$), with increasing marginal cost: the spot price $p_g(A)$ (dotted) rises while the pipe is slack and settles at $c(K)$ once every cap binds. The two curves are the same convex shape $A^\theta$, so the ex-ante market can rescale capacities until they coincide on $A^\theta\Gamma$, and the decentralized allocation is first-best. The capacity price $\tau_p$ is each shipper's expected excess value over the spot price (shaded area), a pure scarcity rent.

close to zero in low- and moderate-demand states and spikes in extreme ones. Residential heating demand, by contrast, rises smoothly with the severity of the weather.[15] When $\theta_1 > \theta_2$ the curves have different shapes and cross once, at $A^*$: the LDC values capacity more in low-demand states, the power plant in high-demand ones (Figure 8). No rescaling can remove a crossing between curves of different shape, so they never coincide.

The ex-ante FOC equalizes the two expected excesses:

$$\tau_p = \int \big[P_2(\kappa_2, A) - p_g(A)\big]^+ dF(A) = \int \big[P_1(\kappa_1, A) - p_g(A)\big]^+ dF(A). \tag{8}$$

Subtracting the part common to both integrals (the grey shaded area in Figure 8), equilibrium equivalently equalizes the two shaded difference-regions: the LDC's advantage in moderate states (the red shaded area) against the power plant's advantage in high-demand ones (the blue shaded area).

The figure now displays all three ex-post configurations of Section 5.3 as $A$ rises: below $A_2^c$ neither cap binds and both shippers consume at $P_i(q_i, A) = p_g(A)$, which rises with demand (interior, first-best); between $A_2^c$ and $A_1^c$ only the LDC is capped while the power plant holds slack capacity that cannot be transferred, and the price keeps rising as the power plant's demand grows (partial congestion); beyond $A_1^c$ both are capped, the pipe is full, and the

[15] Gas-fired peaking units are dispatched last in the merit order and run only when electricity demand peaks, while residential and commercial heating loads track heating degree days. During the January 2024 cold snap, U.S. natural gas consumption set an all-time record as gas-fired generation surged on top of heating demand; during Winter Storm Uri in February 2021, gas-fired generators unable to procure fuel accounted for the majority of unplanned generation outages. The sources are cited in Appendix A.5 and, for the Uri outages, Appendix A.3.

price settles at $c(K)$ (full congestion). Equalizing *expected* excess values still leaves realized marginal utilities apart in almost every congested state, and a single fixed vector $(\kappa_1^*, \kappa_2^*)$ cannot track a first-best split that rotates with $A$.

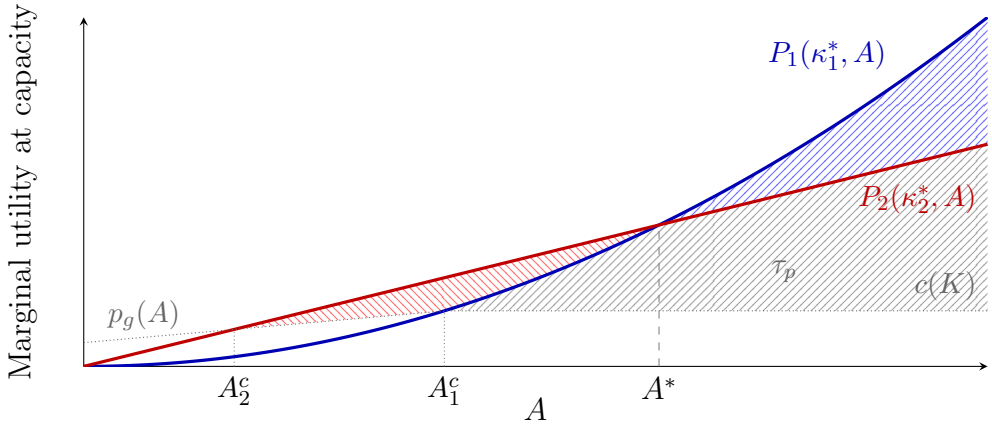



**Figure 8:** SIMR fails ($\theta_1 > \theta_2$), with increasing marginal cost: the spot price $p_g(A)$ (dotted) rises while the pipe is slack and settles at $c(K)$ once every cap binds. Each shipper's ex-ante value of capacity is the expected area between its own curve and the spot-price path; equilibrium (8) equalizes these, i.e., the two shaded difference-regions have equal expected area (red: the LDC's advantage; blue: the power plant's). The grey area is the part of the excess value common to both shippers, so the capacity price satisfies $\tau_p = \text{grey} + \text{red} = \text{grey} + \text{blue}$.

Regarding the loss relative to the first best, the loss is zero in the interior region and at the crossing state $A^*$, where the fixed split coincides with the planner's, and positive in every other congested state. In each such state, the welfare gain from reallocating the marginal unit of gas toward the under-served shipper is the vertical gap between its marginal utility and the other shipper's opportunity cost of gas. In expectation, the loss has two parts:

$$\mathbb{E}\big[\mathcal{L}(A)\big] = \underbrace{\int_{A_2^c}^{A_1^c} \big[P_2(\kappa_2^*, A) - p_g(A)\big]\, dF(A) + \int_{A_1^c}^{A^*} \big[P_2(\kappa_2^*, A) - P_1(\kappa_1^*, A)\big]\, dF(A)}_{\text{LDC under-served: the red shaded region}} + \underbrace{\int_{A^*}^{\infty} \big[P_1(\kappa_1^*, A) - P_2(\kappa_2^*, A)\big]\, dF(A)}_{\text{power plant under-served: the blue shaded region}}$$

The first part accrues in the states of the red shaded region in Figure 8: between $A_2^c$ and $A_1^c$ the LDC's marginal utility exceeds the spot price at which the power plant consumes, and between $A_1^c$ and $A^*$ it exceeds the power plant's marginal utility at capacity, so the planner would reallocate gas toward the LDC. The second part accrues in the states of the blue shaded region, where the ranking reverses and the planner would reallocate gas toward the power plant.

## 5.4 Capacity Obligation

We now turn to the second friction introduced in Section 2: the capacity obligation imposed on LDCs.

Many final consumers do not purchase gas directly from the spot market. Instead, local distribution companies (LDCs) purchase gas in wholesale markets and resell it to households and other final consumers. As discussed previously, these consumers often face fixed or slow-moving prices that do not vary with the spot market. Davis and Muehlegger (2010) reject marginal-cost pricing for residential gas in all fifty U.S. states, with residential markups averaging over 40%. Residential rates may track the average cost of gas but are less volatile than spot market prices. As a result, residential demand does not fully reflect the economic signal from the wholesale market, which is a heavily debated topic in energy markets. In this section, we let the LDC serve its consumers at a fixed residential price $p_c$ while procuring gas at the spot price $p_g(A)$.

In addition, distribution networks lend themselves to natural monopolies, so LDCs are regulated by State Commissions. In the absence of regulation, LDCs have an incentive to curtail their customers whenever the spot market price rises above the residential price $p_g(A) > p_c$, which may occur in high-demand periods when gas is most needed. Because the consequences of even a brief interruption of residential heating service are severe, regulators impose capacity obligations, mandating that shippers have enough capacity to meet peak-demand. The details of those regulations vary across the U.S. and were discussed at length in previous sections.

We model capacity obligations as a penalty $\mu\,\ell$ that the LDC must pay if it doesn't meet residential demand at price $p_c$. This becomes a model of curtailment in which the LDC optimally decides how much to curtail its customers during high-demand periods, similar to Joskow and Tirole (2007). The penalty has two terms. First, a loss term $\ell$ that is proportional to the loss of consumer surplus that consumers incur due to curtailment. If consumers demand $D_2(p_c, A)$ and the LDC only provides $q_2$, the loss of consumer surplus is:

$$\ell(q_2; p_c, A) \;\equiv\; \int_{q_2}^{D_2(p_c,A)} \big[P_2(x, A) - p_c\big]\, dx. \tag{9}$$

The idea of (9) underlies our expectation that peak-day demand regulation provides consumers the grounds to sue the LDC if their demand is not met. The amount consumers ask in compensation should be roughly proportional to the damage they incurred. Second, the penalty is indexed by a parameter $\mu \geq 0$, which varies the stringency of the regulation. Finally, under rate regulation, LDCs in most states are mandated to charge "just and rea-

sonable" rates for their services, which keeps the regulated residential price low relative to wholesale prices in the high-demand periods on which our empirical analysis focuses. We therefore assume that $p_c$ never exceeds $p_g(A)$ on the domain of $A$ we study.[16]

We now consider the optimal curtailment choice given weather realization $A$ and the capacity constraint $\kappa_2$. The LDC's problem is:

$$\max_{q_2(A)\leq\kappa_2} \Big\{(p_c - p_g(A))q_2(A) - \mu\ell(q_2(A); p_c, A)\Big\}. \tag{10}$$

Because the spot price is above the residential rate, the LDC wants to curtail: it delivers only $D_2(\tilde{p}_g(A,\mu), A) < D_2(p_c, A)$, residential demand evaluated at an "effective price" $\tilde{p}_g(A,\mu)$ derived below. Depending on whether capacity is binding at that level, the LDC might curtail even more. In either case, $q_2^*(A) < D_2(p_c, A)$ and $\ell > 0$.

By the Leibniz rule, the LDC wants to equalize the marginal utility of its consumers to the "effective price". When capacity doesn't bind:

$$P_2(q_2(A), A) = \underbrace{\frac{1}{\mu}p_g(A) + \Big(1 - \frac{1}{\mu}\Big)p_c}_{\text{"Effective" price } \tilde{p}_g(A,\mu)}.$$

Whenever capacity binds, $P_2(\kappa_2, A) > \tilde{p}_g(A,\mu)$. That is to say, the LDC would deliver more gas if it had more capacity, and the gap $P_2(\kappa_2, A) - \tilde{p}_g(A,\mu)$ is its marginal value of capacity at $A$. This effective price varies with the stringency of regulation in an intuitive way:

1. $\lim_{\mu\to 0} \tilde{p}_g(A,\mu) = \infty$. As the LDC becomes unregulated, it has no incentive to procure gas because it sells it at a loss, so it chooses full curtailment.

2. $\mu = 1 \to \tilde{p}_g(A,1) = p_g(A)$. In this case, the LDC curtails gas supply so that consumers consume what they would want if they faced the spot market price. Joskow and Tirole (2007) shows that this scenario can be achieved in equilibrium by unregulated but perfectly competitive LDCs. It is as close as the regulator can get to the first-best.

3. $\lim_{\mu\to\infty} \tilde{p}_g(A,\mu) = p_c$. As regulations become increasingly stringent, the LDC brings consumers closer and closer to what they want at the fixed price $p_c$.

Taken together, we can use this ex-post behavior to determine the optimal pipeline

[16] For completeness, Appendix E.2 characterizes the LDC's optimal curtailment choice over the complete domain of $A$, covering both the "out of the money" regime ($p_g(A) \geq p_c$) studied here and the "in the money" regime ($p_c > p_g(A)$), in which the LDC earns a positive margin and curtails only when its capacity binds.

capacity in the ex-ante market.

$$U_2 = \max_{\kappa_2 \geq 0} \Big\{ \tau_p(K_2 - \kappa_2) + \int_{\Omega_A} [(p_c - p_g(A))q_2(A) - \mu \ell(q_2(A); p_c, A)] f_A(A) dA \Big\},$$
$$q_2(A) = \min\big\{\kappa_2,\ D_2\big(\tilde{p}_g(A,\mu), A\big)\big\}$$

Capacity is valuable to the LDC at the margin only when it binds, which leads to the following FOC:

$$\tau_p = \int_{\Omega_A} \mathbf{1}\left\{D_2(\tilde{p}_g(A,\mu), A) > \kappa_2\right\} \left[P_2(\kappa_2, A) - \tilde{p}_g(A,\mu)\right] f_A(A) dA$$

To better understand the tradeoff generated by peak-day demand regulation, we explain its benefit and cost in more detail. First, a more stringent obligation delivers more gas to residential customers. The higher the value of $\mu$, the higher is expected residential consumer surplus: a higher $\mu$ both lowers the effective price toward $p_c$, raising delivery for a given capacity, and raises the LDC's equilibrium capacity.

This is the benefit side of the regulation. Intuitively, raising $\mu$ shifts gas toward residential customers in the high-demand states where the spot price would otherwise curtail them; but, as the rest of this section shows, it does so by inducing the LDC to over-contract capacity, distorting the cross-sectional allocation away from the first best. The parameter $\mu$ therefore trades residential reliability against allocative efficiency, and which effect dominates depends on the social value of residential reliability.

**When SIMR is satified.** Recall from Section 5.3 that the market can implement the first best only when the ranking of marginal valuations across firms is the same in every state (SIMR satisfied). Regulation breaks this property even when true preferences satisfy it. To see why, compare the two firms' marginal values of capacity in a state where their constraints bind. For the power plant, a marginal unit of capacity is worth its marginal utility net of the gas price, $P_1(\kappa_1, A) - p_g(A)$. For the LDC, a marginal unit earns the retail margin $p_c - p_g(A)$ and avoids the penalty $\mu\big[P_2(\kappa_2, A) - p_c\big]$, so it is worth

$$v_2 \equiv \mu P_2(\kappa_2, A) - (\mu - 1)\, p_c - p_g(A) \ = \ \underbrace{P_2(\kappa_2, A) - p_g(A)}_{\text{unregulated valuation}} + \underbrace{(\mu - 1)\big[P_2(\kappa_2, A) - p_c\big]}_{\text{regulatory wedge}}. \quad (11)$$

The first term is the valuation of an unregulated firm with the same preferences; the second is the wedge created by the obligation, which amplifies every deviation of consumers' marginal utility from the retail price by the factor $\mu - 1$.

**Proposition 2** (Precautionary over-contracting by LDCs)**.** *Suppose SIMR holds and the ex-*

*post market clears competitively. Then equilibrium LDC capacity $\kappa_2(\mu)$ is strictly increasing in $\mu$ and the power plant's $\kappa_1(\mu)$ is strictly decreasing; hence $\kappa_2(\mu) > \kappa_2^*$ for all $\mu > 1$.*

*Proof.* See Appendix E.1. □

To illustrate how LDC regulation affects the equilibrium, we return to the CRRA example of Figure 7: common exponent $\theta_1 = \theta_2 = \theta$ and increasing marginal cost. Absent regulation ($\mu = 1$), the ex-ante market sets $\kappa_1^{-\gamma_1} = \kappa_2^{-\gamma_2} = \Gamma$, both marginal-utility curves collapse onto $A^\theta\Gamma$, and the allocation is first-best. Now let firm 2 be an LDC with retail price $p_c > 0$ and obligation stringency $\mu > 1$. The obligation makes the LDC's valuation of gas exceeds its consumers' true marginal utility by the regulatory wedge. Equivalently, the obligation lowers the price the LDC effectively pays for gas from $p_g(A)$ to $\tilde{p}_g(A, \mu) < p_g(A)$.

The wedge is positive in every binding state and grows with the demand state, so regulation makes the LDC behave as if it were more weather-sensitive than its customers actually are: Panel (a) of Figure 9 shows the regulated valuation $v_2$ lying above the consumers' true marginal-utility curve, by more the higher the state.

In equilibrium, the two ex-ante first-order conditions hold at the common capacity price:

$$\tau_p = \underbrace{\int \left[P_1(\kappa_1, A) - p_g(A)\right]^+ dF(A)}_{\text{power plant's excess value}} = \underbrace{\int \left[P_2(\kappa_2, A) - \tilde{p}_g(A, \mu)\right]^+ dF(A)}_{\text{LDC's excess value}}. \tag{12}$$

At equal capacities ($\kappa_1^{-\gamma_1} = \kappa_2^{-\gamma_2} = \Gamma$), the LDC's excess value strictly exceeds the power plant's because it faces the lower effective price in every binding state. Restoring equality requires $\kappa_2^{-\gamma_2} < \Gamma < \kappa_1^{-\gamma_1}$: the LDC expands its position and the power plant is crowded out, with $\kappa_2 > \kappa_2^*$ and $\kappa_1 < \kappa_1^*$. This is precautionary over-contracting, as in Proposition 2: the obligation gives the LDC a larger ex-ante position and, because contracts confer scheduling priority, regulation-induced priority in the physical allocation of gas. Panel (b) of Figure 9 shows the equilibrium: relative to the unregulated equilibrium $A^\theta\Gamma$, the LDC's marginal-utility curve rotates down and the shipper's rotates up.

The resulting misallocation is one-directional. Because the true marginal-utility curves are proportional and never cross (SIMR is satisfied), the marginal unit of gas is worth more to the power plant than to the LDC's consumers in every congested state. The distortion is entirely manufactured by the obligation: at $\mu = 1$ both curves collapse back onto $A^\theta\Gamma$ and the loss vanishes.

Regarding the loss relative to the first best, unlike in the unregulated case, it arises even in states in which neither cap binds. Whenever the LDC's cap is slack, its consumers purchase gas at the effective price $\tilde{p}_g(A, \mu) < p_g(A)$, causing residential demand to be over-served

(a) Regulation amplifies the LDC's valuation (b) Equilibrium: the LDC over-contracts

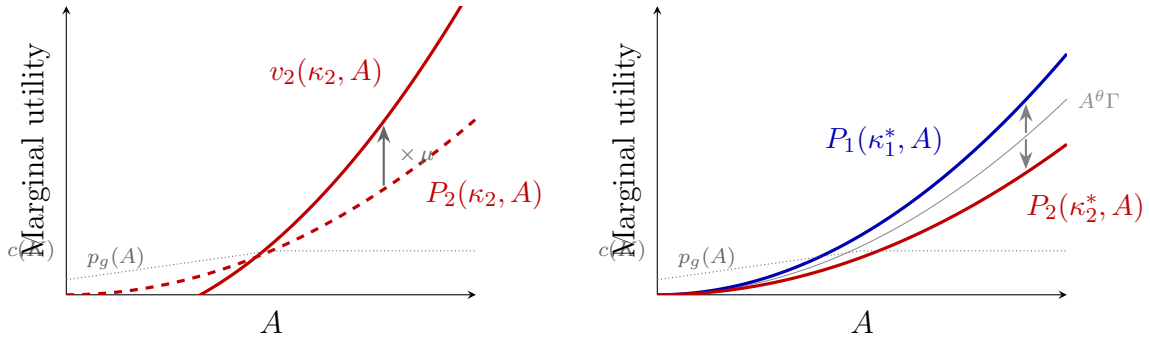



**Figure 9:** LDC regulation when SIMR holds, with increasing marginal cost: the spot price $p_g(A)$ (dotted) rises while the pipe is slack and settles at $c(K)$ once every cap binds. *Panel (a):* the obligation amplifies the LDC's valuation of gas, equation (11). *Panel (b):* ex-ante equilibrium allocation. Relative to the unregulated equilibrium $A^\theta\Gamma$ (grey), the LDC over-contracts ($\kappa_2(\mu) > \kappa_2^*$) and the power plant is crowded out ($\kappa_1(\mu) < \kappa_1^*$).

relative to the first best.

Once the power plant's cap binds (at $A_1^c$), the power plant is in addition under-served: the marginal unit of gas is worth more to the power plant than to the LDC's consumers. In expectation, the loss has two parts:

$$\mathbb{E}\big[\mathcal{L}(A)\big] \;=\; \underbrace{\int_0^{A_1^c} \big[p_g(A) - \tilde{p}_g(A,\mu)\big]\, dF(A)}_{\text{residential over-served: the red shaded region}} + \underbrace{\int_{A_1^c}^{A_2^c} \big[P_1(\kappa_1^*, A) - \tilde{p}_g(A,\mu)\big]\, dF(A) \;+\; \int_{A_2^c}^{\infty} \big[P_1(\kappa_1^*, A) - P_2(\kappa_2^*, A)\big]\, dF(A)}_{\text{power plant under-served: the blue shaded region}}$$

Figure 10 shows the two parts, together with why the equilibrium takes this form. The grey line is the LDC's regulated valuation of gas, $v_2 = \mu P_2(\kappa_2^*, A) - (\mu - 1)p_c$: the obligation scales up the consumers' true curve $P_2(\kappa_2^*, A)$ so that $v_2$ overtakes even the power plant's curve in the highest-demand states, which is why the LDC outbids the power plant for capacity and over-contracts. The power plant's cap binds beyond $A_1^c$, where its curve crosses the spot-price path; the LDC's cap binds beyond $A_2^c$, where $v_2$ crosses the spot-price path.

The red shaded region, between the spot-price path and the effective-price path, is the over-service part of the loss: in these states the planner would curtail residential consumption back to the spot price. The blue shaded region, between the power plant's curve and the LDC's marginal utility is the crowding-out part: the planner would reallocate gas from the LDC's consumers to the power plant. Both parts vanish at $\mu = 1$, when the effective price

coincides with the spot price and the equilibrium positions collapse back to the unregulated ones.

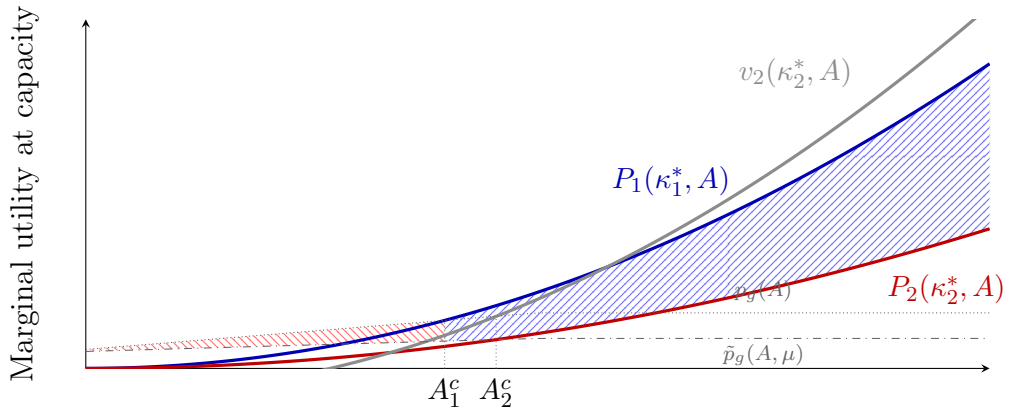



**Figure 10:** The loss under LDC regulation when SIMR holds, drawn at the regulated equilibrium positions $(\kappa_1^*, \kappa_2^*)$. The grey line is the LDC's regulated valuation $v_2 = \mu P_2(\kappa_2^*, A) - (\mu - 1)p_c$, which overtakes the power plant's curve in the highest-demand states: this is why the ex-ante market assigns the LDC its large position. The LDC's consumers purchase at the effective price $\tilde{p}_g(A, \mu)$ (dash-dotted), below the spot price $p_g(A)$ (dotted); the power plant's cap binds beyond $A_1^c$, and the LDC's beyond $A_2^c$, where $v_2$ crosses the spot-price path. The red shaded region, between the two price paths, is the marginal loss from over-serving residential consumers while the LDC's cap is slack; the blue shaded region, between the power plant's curve and the LDC's marginal utility, is the marginal loss from crowding out the power plant. Both regions vanish at $\mu = 1$.

**When SIMR is not satisfied.** Last we turn to the case in which SIMR fails. By Proposition 1, the deliverability constraint alone then already prevents the market from implementing the first best: the planner's preferred split of the pipe varies with the state, while the market's ex-ante positions are fixed across states.

The capacity obligation adds a second, independent distortion on top. As in the previous subsection, the regulatory wedge raises the LDC's valuation of gas above its consumers' true marginal utility in every binding state. The ex-ante positions must therefore absorb two distortions at once: a state-dependent first-best split that no fixed vector can track, and an inflated LDC demand for capacity that tilts the fixed vector itself toward the LDC.

Figure 11 shows what the shift does in the CRRA example. At the equilibrium positions $(\kappa_1^*, \kappa_2^*)$, the *true* curves cross once at $A_1$. This crossing is the analogue of $A^*$ in Figure 8, but the two differ: the ex-ante market readjusts positions in equilibrium under regulation, so the crossing of the true curves shifts as well. Regulation raises the LDC's valuation of gas above its consumers' true marginal utility, and the wedge is positive in every binding state. At the original equilibrium crossing $A_1$, the LDC's valuation therefore still exceeds the power plant's, so the market-relevant crossing moves right to $A^* > A_1$.

On the shaded band $A \in (A_1, A^*)$, the first best favors the power plant while regulation favors the LDC: the power plant's true marginal utility already exceeds that of the LDC's consumers, yet the LDC's regulated valuation still exceeds the power plant's, so the ex-ante market keeps the marginal unit of capacity with the LDC.

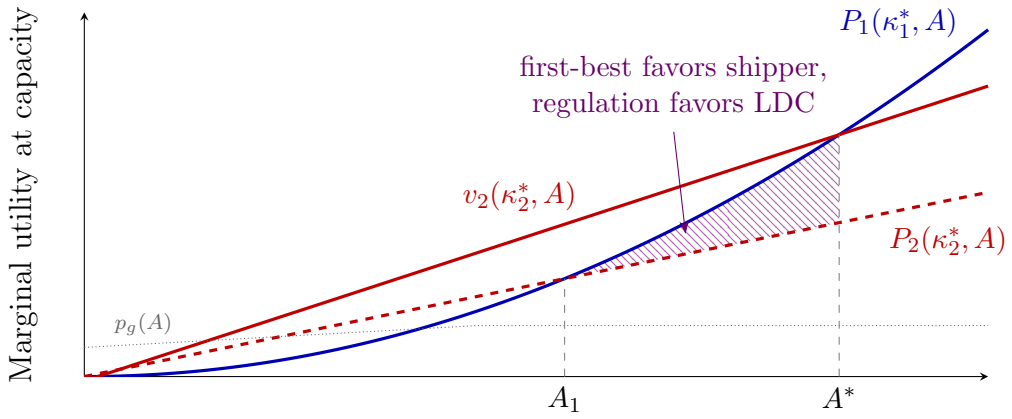



**Figure 11:** LDC regulation when SIMR fails. The power plant's curve (blue) crosses the LDC's *true* curve (dashed red) at $A_1$. The regulatory wedge lifts the LDC's valuation to the solid red line $v_2$, and moves the market-relevant crossing to $A^* > A_1$. On the shaded band $(A_1, A^*)$ the first best favors the power plant while regulation favors the LDC.

Regarding the loss relative to the first best, both channels now operate, and the misallocation is again two-sided. In mild states the LDC's consumers are over-served: they purchase at the effective price $\tilde{p}_g(A, \mu)$, and even after the LDC's cap binds at $A_2^c$ their marginal utility remains below the spot price until the state $A_0$, where $P_2(\kappa_2^*, A_0) = p_g(A_0)$. Between $A_0$ and the true crossing $A_1$ the ranking flips: the LDC's consumers' marginal utility exceeds the spot price at which the power plant consumes—and, once the power plant's cap binds at $A_1^c$, the power plant's marginal utility at capacity—so residential demand is now *under*-served. Beyond $A_1$ the power plant is the under-served party, and the gap grows with the state. In expectation, the loss has three parts:

$$\mathbb{E}\big[\mathcal{L}(A)\big] = \underbrace{\int_0^{A_2^c} \big[p_g(A) - \tilde{p}_g(A,\mu)\big]\, dF(A) + \int_{A_2^c}^{A_0} \big[p_g(A) - P_2(\kappa_2^*, A)\big]\, dF(A)}_{\text{residential over-served: the red shaded region}}$$
$$+ \underbrace{\int_{A_0}^{A_1^c} \big[P_2(\kappa_2^*, A) - p_g(A)\big]\, dF(A) + \int_{A_1^c}^{A_1} \big[P_2(\kappa_2^*, A) - P_1(\kappa_1^*, A)\big]\, dF(A)}_{\text{residential under-served: the violet shaded region}}$$
$$+ \underbrace{\int_{A_1}^{\infty} \big[P_1(\kappa_1^*, A) - P_2(\kappa_2^*, A)\big]\, dF(A)}_{\text{power plant under-served: the blue shaded region}}$$

Figure 12 presents the three parts in separate panels, each over the relevant range of weather states. Panel (a) covers the mild states, in which residential consumers are over-served. The cause is regulation itself: the obligation lowers the price the LDC effectively pays from $p_g(A)$ to $\tilde{p}_g(A, \mu)$, so its consumers purchase beyond the first-best level even though capacity is not scarce. Panel (b) covers the moderate states, in which residential demand is under-served. The cause is the SIMR violation, by exactly the logic of Figure 8: a fixed position cannot deliver as much gas to the LDC's consumers as the planner would in the states where their valuation is relatively high. Regulation, however, makes this region *smaller*: over-contracting gives the LDC a larger position and shifts the true crossing $A_1$ to the left. Panel (c) covers the extreme states, in which the power plant is under-served. Here regulation makes the loss worse: the power plant is crowded out by the LDC's over-contracting, so the gap between its marginal utility and the LDC's widens. Most importantly, because the power plant's valuation is the more weather-sensitive one, this gap grows with the state, so the power plant is starved the most exactly when it needs gas the most. Regulation therefore not only adds a loss of its own; it relocates the misallocation loss into the extreme states: gas going to the wrong firm, now occurs over a larger set of high-demand states and a smaller set of moderate states, moving the burden into precisely the states in which gas is most valuable.

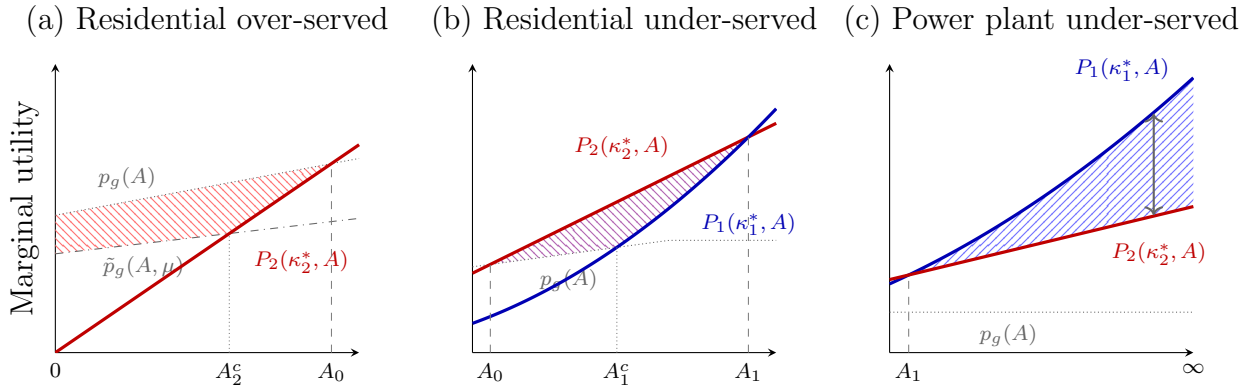



**Figure 12:** The loss when SIMR fails and the LDC is regulated, drawn at the regulated equilibrium positions $(\kappa_1^*, \kappa_2^*)$ and split by the range of the weather shock; note the different axis scales across panels. *Panel (a), mild states:* the LDC's consumers purchase at the effective price $\tilde{p}_g(A, \mu)$ (dash-dotted), below the spot price, and are over-served (red region); the LDC's cap binds at $A_2^c$, and the loss vanishes at $A_0$. *Panel (b), moderate states:* residential demand is under-served (violet region), as in Figure 8; over-contracting shrinks this region by shifting the true crossing $A_1$ leftward. *Panel (c), extreme states:* the power plant is under-served (blue region), and the gap between the two curves grows with the state: the power plant is starved the most in the most extreme states.

**Summary:** Here we summarize the main theoretical implications of our model. First, the deliverability constraint requires the market to assign pipeline capacity before the weather shock is realized. Without regulation, and when SIMR holds, the ex-ante market implements

the first best: capacity is scarce in high-demand states, but curtailment is simultaneous and efficient, and no participant bears a misallocation loss in any state. When SIMR fails, the loss is confined to congested states, but its incidence varies with the severity of the weather shock: in moderate states the less weather-sensitive firm is under-served, while in extreme states the power plant is starved of gas precisely when its valuation peaks.

Regulation changes both the size and the incidence of the loss. The capacity obligation lowers the price the LDC effectively pays for gas, which over-serves residential consumers in every state where the LDC's capacity is slack and induces the LDC to over-contract capacity ex ante. When SIMR holds, the resulting misallocation is one-directional: in every binding state, the marginal unit of gas is worth more to the power plant than to the LDC's consumers. When SIMR fails, over-contracting additionally shifts the crossing of the true marginal-utility curves to the left, shrinking the range of states in which residential consumers are under-served and expanding the range in which the power plant is under-served.

The incidence pattern across weather states is therefore systematic. In low-demand states, either no one bears a loss (without regulation) or residential consumers are over-served at the expense of productive efficiency (with regulation). In extreme states, whenever any friction binds, the under-served party is the unregulated power plant.

The model shows that the risk of extreme weather ultimately falls on unprotected pipeline users, and that reliability regulation further shifts this risk toward them: by insulating residential consumers, it grants LDCs regulation-induced priority in the pipeline and crowds out the unregulated users.

Table 3 summarizes, for each case, whether the first best is implemented, where the loss arises, and who is mis-served in each range of weather states; each case is evaluated against its own first best. The incidence entries reflect our maintained assumption that gas-fired power plants are more weather-sensitive than residential heating demand: generators are dispatched only when electricity demand peaks, so their gas consumption is concentrated in extreme states, whereas heating demand rises smoothly with the severity of the weather. The predictions otherwise require only that $p_g(A)$ be non-decreasing in $A$ and that capacity binds in extreme states. We examine these assumptions and the theoretical predictions in the empirical analysis in the next section.

# 6 Empirical Model and Demand Estimation

We now expand the model to match our data. For each shipper $i$ of type $k \in \{1, 2\}$ located in market (BEA of destination) $j$, we define the following period utility function of consuming $q$ units of gas:

**Table 3:** The Loss and Its Incidence across Cases (each case relative to its own first best)

| | $\mu = 1$ **(no LDC reg.)** | | $\mu > 1$ **(LDC reg.)** | |
|---|---|---|---|---|
| | SIMR holds | SIMR fails | SIMR holds | SIMR fails |
| First best implemented? | Yes | No | No | No |
| Loss in uncongested states? | No | No | Yes | Yes |
| LDC capacity (vs. $\mu = 1$) | — | — | Over-contracts | Over-contracts |
| **Who is mis-served, by realization of $A$:** | | | | |
| Low-demand states | None | None | Residential (over) | Residential (over) |
| Moderate states | None | Residential (under) | Power plant (under) | Residential (under) |
| Extreme states | None | Power plant (under) | Power plant (under) | Power plant (under) |

*Notes:* Each case is evaluated against the first best under its own preferences. "Over" and "under" denote over- and under-service relative to the first-best allocation.

$$u_i(q, A_{jt}) = \Gamma_i A_{jt}^{\theta_{kj}} \frac{q^{1-\gamma_{jk}}}{1-\gamma_{jk}}, \tag{13}$$

where $A_{jt}$ is an aggregate demand shock associated with weather, which affects region $j$ in period $t$. $\Gamma_i$ is a shipper-specific fixed effects. To match our data, we model a period as a week. In this context, a contract in the secondary market signed in week $t$ is deemed active in week $t+1$ only. This lagged assumption is critical for studying the deliverability constraint and aligns with our data, as shown in **Figure x**, where the modal time lag between the date a contract is signed and the date it becomes active is 7 days. The goal of this section is to recover the following demand-side parameters:

$$\boldsymbol{\Theta} = \{\theta_{kj}, \gamma_{kj}, \Gamma_i, \mu_{jt}\}_{i,j,k,t}.$$

Although we write LDC regulation as $\mu_{jt}$, we assume that regulation is set at the state-year level ($st$), where states are a superset of markets $j$. To recover these demand-side parameters, we exploit the first-order conditions coming from the ex-ante markets. Let $\overline{\tau}_{jt} \geq 0$ be the (observed) average per-unit price associated with capacity released in market $j$ for contracts signed in week $t$, and let $\kappa^j_{i,t+1} \geq 0$ be the total capacity available for shipper $i$ in week $t+1$ and market $j$ coming from both the primary and secondary markets (net position). Then, for an unregulated shipper $k = 1$,

$$\overline{\tau}_{jt} = \int_{\Omega_A} \mathbf{1}\left(p_{j,t+1}(A) < P_i(\kappa^j_{i,t+1}, A)\right)\left[P_i(\kappa^j_{i,t+1}, A) - p_{j,t+1}(A)\right] dF_{j,t+1}(A),$$

where $p_{j,t+1}(A)$ is the ex-post (spot market) price of gas in market $j$ given some weather realization $A$. For a regulated shipper $k = 2$, the first-order conditions are similarly defined as follows:

$$\overline{\tau}_{jt} = \int_{\Omega_A} \Big[ \mathbf{1} \underbrace{\left\{ p_{j,t+1}(A) \geq p^c_{j,t+1}, D_i(p^c_{j,t+1}, A) > \kappa^j_{i,t+1} \right\}}_{\text{in the money, constrained}} [P_i(\kappa^j_{i,t+1}, A) - p^c_{j,t+1}]$$
$$+ \mathbf{1} \underbrace{\left\{ p_{j,t+1}(A) > p^c_{j,t+1}, D_i(p^c_{j,t+1}, A) > \kappa^j_{i,t+1} \right\}}_{\text{out of the money, constrained}} \left[ P_i(\kappa^j_{i,t+1}, A) - \tilde{p}_{j,t+1}(A, \mu_{jt}) \right] dF_{j,t+1}(A),$$

where $\tilde{p}_{j,t+1}(A, \mu_{jt}) \equiv \frac{1}{\mu_{jt}} p_{j,t+1}(A) + \left(1 - \frac{1}{\mu_{jt}}\right) p^c_{j,t+1}$. For accounting purposes, we need to construct $\kappa^j_{i,t}$, which is not directly observed in the data. Let $Q^j_{it} \geq 0$ be the capacity available for shipper $i$ from the primary market (long-term contract). Let $\tilde{\kappa}^j_{it}(s) \in (-\infty, \infty)$ be the total capacity purchased or sold (net) by shipper $i$ in the secondary market that will be available for $s \geq 1$ periods into the future[17]. Then, the following accounting identity holds:

$$\kappa^j_{it} = Q^j_{it} + \sum_{\tau < t} \sum_{s \geq t - \tau} \tilde{\kappa}^j_{i\tau}(s).$$

The difficulty for identification is that we do not observe the weather forecast in each market $F_{j,t+1}$, nor do we observe the counterfactual spot-market price realization for all possible weather realizations, so we cannot construct these FOCs directly, even if we knew demand parameters. Instead, we leverage the rational expectations assumption to construct unconditional moment conditions given the differential effects of the same cold wave across locations. Let $p_{j,t+1}(A_{j,t+1})$ be the realized spot market price in market $j$ and period $t + 1$. Then,

$$\mathbb{E}\left[ \mathbf{1}\left( p_{j,t+1}(A_{j,t+1}) < P_i(\kappa^j_{i,t+1}, A_{j,t+1}) \right) \left[ P_i(\kappa^j_{i,t+1}, A_{j,t+1}) - p_{j,t+1}(A_{j,t+1}) \right] - \overline{\tau}_{jt} \right] = 0.$$

# 7 Empirical Investigation of Cold Waves, 2014–2023

In this section, we test the model's predictions using cold waves observed between 2014 and 2023. Specifically, we examine how the secondary market responds to the arrival of a cold wave, focusing on three outcomes: changes in tariffs, the amount of capacity released, and

[17] Note that, to greatly simplify the dynamics of the model, we assume that contracts are signed for only one period. However, contracts may be signed for any number of periods in the data, and we need to keep track of all past contracts signed that are still available at period $t$.

the frequency of short-term contracts. We begin by using daily temperature data to identify all cold waves during the sample period.

**Identifying cold waves.** We identify cold waves using daily mean temperatures for all the counties in the contiguous United States from 2006 to 2023. Because unusually cold weather depends on location and season, we measure cold using standardized temperature anomalies relative to each county's historical norm for that time of year. A county is classified as cold when its temperature falls more than two standard deviations below normal. We restrict attention to November through March, when cold weather is most relevant for natural gas demand.

We calibrate the identification procedure against four well-documented cold waves: the February 2015 "Siberian Express," the January 2019 polar vortex, Winter Storm Uri in February 2021, and Winter Storm Elliott in December 2022. We select the strictest anomaly threshold that identifies all four as distinct events while capturing their principal affected states.

We first identify national cold-wave events based on the share of U.S. counties experiencing unusually cold temperatures, then determine their arrival dates in individual BEA economic areas (BEAs), which define our markets. Specifically, a national event begins when at least 20 percent of counties are cold, while a BEA is considered affected when at least 30 percent of its counties are cold. Short interruptions are allowed so that a single cold wave is not split into multiple events. This approach captures the staggered arrival of cold waves across markets, which provides the variation used in our empirical analysis.

The procedure identifies 19 cold waves during 2014–2023, of which 10 are classified as major, defined by a peak national share of cold counties of at least 40 percent. Our main analysis focuses on these major events, while Appendix G examines minor events. Appendix F provides the full identification procedure, calibration, event dates, and regional arrival patterns.

## 7.1 Tariffs, Quantities, and the Probability of a Release

We organize capacity release contracts by signing week and the BEA of their delivery point, distinguishing contracts by duration because they serve different roles in the allocation. One-month contracts account for about 80 percent of the quantity released and are typically signed in the last ten days of the preceding month, reflecting the planned transactions of the ex-ante allocation. In contrast, contracts lasting 1–7 days are typically signed one day ahead and capture short-term adjustments. For all contracts combined, and separately for one-month and 1–7-day contracts, we examine three outcomes at the BEA-week level: the

log tariff, the log quantity released, and the probability that at least one contract is signed. The first two outcomes are defined only when at least one contract of the relevant type is signed.

For each outcome $y_{bt}$ in BEA $b$ and week $t$, we estimate

$$y_{bt} = \alpha_t + \alpha_{b,m(t)} + \alpha_{b,w(t)} + \beta D_{bt} + \varepsilon_{bt}, \tag{14}$$

where $D_{bt}$ equals one during a cold wave, from its onset week through its end week in BEA $b$. The calendar-week fixed effects ($\alpha_t$) absorb shocks common to all BEAs, including the national response to cold waves. The BEA $\times$ calendar-month fixed effects ($\alpha_{b,m(t)}$) capture local seasonality, while the BEA $\times$ week-of-month fixed effects ($\alpha_{b,w(t)}$) account for local signing cycles, particularly the concentration of one-month contracts near month-end. The coefficient $\beta$ therefore measures the response in cold BEAs relative to BEAs that are not cold in the same week, conditional on these fixed effects. Standard errors are clustered at the BEA level.

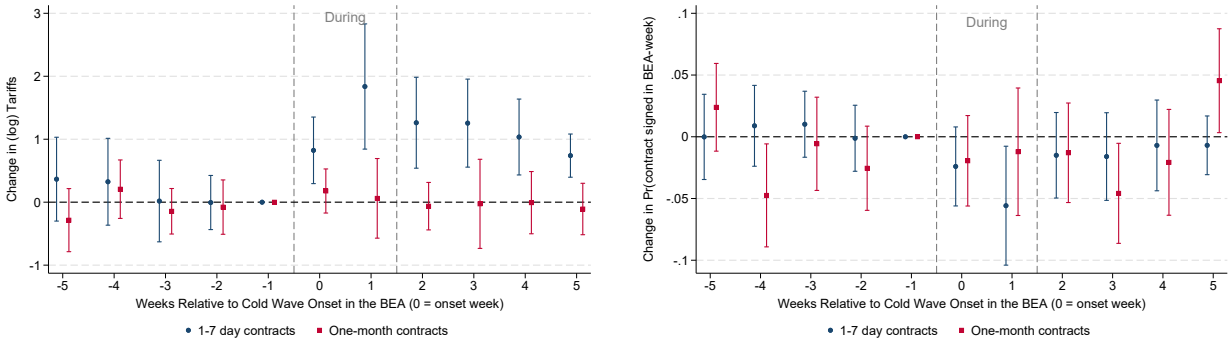
**Table 4:** Effect of Major Cold Waves on Capacity Release During the Cold Weeks

| | Log tariff | Log quantity | Pr(any signed) |
|---|---|---|---|
| All contracts | 0.194 | -0.273 | -0.015 |
| | (0.199) | (0.195) | (0.026) |
| One-month contracts | -0.072 | -0.418 | 0.004 |
| | (0.220) | (0.257) | (0.021) |
| 1–7 day contracts | 0.686** | -0.350 | -0.044** |
| | (0.287) | (0.513) | (0.019) |
| Number of BEA $\times$ week | | | |
| All contracts | 19,593 | 21,421 | 67,080 |
| One-month contracts | 13,180 | 14,368 | 67,080 |
| 1–7 day contracts | 4,869 | 5,067 | 67,080 |

*Notes:* All regressions include week, BEA $\times$ calendar-month, and BEA $\times$ week-of-month fixed effects. The probability of signing is estimated on the full BEA-week panel; log tariff and log quantity are conditional on at least one contract being signed. Standard errors clustered by BEA are in parentheses. * 10%, ** 5%, *** 1%.

Table 4 reports the estimates. Two results stand out, both for 1–7-day contracts. First, tariffs in cold BEAs rise by 0.84 log points during major cold waves, more than doubling relative to BEAs that are not cold in the same week. Second, the probability of signing a 1–

7-day contract falls by 3.2 percentage points from a baseline of 7.9 percent, a decline of about 40 percent. For all contracts combined, tariffs rise by 0.29 log points, significant only at the 10 percent level and driven by the short-term contracts; for one-month contracts, neither tariffs nor signing probabilities change significantly. Released quantities do not increase for any contract type.




(a) Log tariffs

(b) Probability that a contract is signed

**Figure 13:** Major Cold Waves: 1–7 Day versus One-Month Contracts, Week by Week

*Notes:* Coefficients on event-week indicators for the ten major cold waves of 2014–2023, relative to the week before the BEA's onset, with 95 percent confidence intervals. Week 0 is the onset week, week 1 the later cold weeks and weeks 2 to 5 the four weeks after the BEA's end; indicators for weeks −8 to −6 are included and not shown. The fixed effects and clustering are those of Table 4. Panel (a): log tariff, conditional on a contract being signed. Panel (b): probability that at least one contract of the type is signed in the BEA-week.

Figure 13 presents the event-study results. Neither series shows a pre-trend. Tariffs for 1–7-day contracts rise by 0.8 log points in the onset week and 1.8 log points in subsequent cold weeks, remaining elevated for four weeks after the cold wave ends. In contrast, one-month tariffs show no significant response. The probability of signing a 1–7-day contract falls during the cold wave and returns to baseline afterward, while the signing probability for one-month contracts shows no systematic response.

The main takeaway is that the response to cold waves is concentrated in short-term (1–7-day) contracts, both in tariffs and in the probability of signing. As a cold wave begins, fewer short-term contracts are signed, while those that are signed command substantially higher tariffs. Importantly, most cold waves last about a week, and few extend beyond two weeks. Because 1–7-day contracts commit capacity for multiple days, the allocation need not adjust day by day as the shock unfolds. Instead, capacity holders' expectations about the severity and duration of the cold wave at its onset play an important role in their release decisions. The observed response therefore reflects not only contemporaneous demand pressures but also anticipation of demand over the coming days.

As a robustness check, Appendix G repeats the event-study analysis for the nine minor cold waves in 2014–2023. We find no significant response in either tariffs or signing probabilities for 1–7-day contracts, nor any systematic response for one-month contracts. The effects documented above are therefore concentrated in major cold waves.

# 8 Conclusion

This paper studies how the U.S. natural gas pipeline network allocates transportation capacity during extreme cold waves, a question of growing importance as severe weather events become more frequent and as both residential heating and electricity generation remain heavily dependent on natural gas. We develop a theoretical framework showing that priority contracting and peak-day demand regulation can create inefficiencies in capacity allocation, contributing to high prices and excess demand precisely when reliable gas delivery matters most.

Our analysis delivers two central insights. First, when pipeline users differ in how their demand responds to weather across different realizations of a shock, ex-ante capacity allocation through long-term contracts cannot generally achieve the ex-post efficient allocation. The presence of gas peaking power plants, which activate only under extreme conditions, is a key source of such heterogeneity and illustrates why rigid priority contracting may be ill-suited to environments with substantial demand uncertainty. Second, peak-day demand regulation, while intended to protect retail gas consumers, can partly backfire: by inducing local distribution companies to over-contract pipeline capacity and under-release it in the secondary market, it grants them regulation-induced priority and crowds out other users, particularly gas-fired power plants, that also provide heating critical to households. In this way, reliability requirements designed to ensure stable supply can paradoxically increase the fragility of the broader system during the very events they are meant to address.

We test these predictions using highly granular data on the universe of U.S. interstate pipeline contracts, secondary-market capacity release transactions, daily pipeline utilization, hub-level gas prices, and weather from 2005 to 2023, leveraging the staggered rollout of major cold waves for identification. Our empirical findings confirm that the secondary market for capacity release plays an important role during demand shocks, but that contractual frictions and regulatory incentives limit its ability to reallocate capacity efficiently when it is most needed.

These results carry implications for market design. They suggest that reforms aimed at improving the flexibility of capacity reallocation during extreme events, whether through changes to scheduling priority rules, adjustments to peak-day planning requirements, or

mechanisms that facilitate faster secondary-market transactions, could meaningfully enhance the resilience of natural gas supply chains. More broadly, our findings speak to a fundamental tension in markets characterized by large sunk investments and long-term contracting: the very institutions that support infrastructure financing can impede the short-run reallocation that extreme shocks demand.

# References


**Avalos, Roger, Timothy Fitzgerald, and Randal R. Rucker**, "Measuring the effects of natural gas pipeline constraints on regional pricing and market integration," *Energy Economics*, 2016, *60*.

**Bathiany, Sebastian et al.**, "Climate Models Predict Increasing Temperature Variability in Poor Countries," *Sciences Advances*, 2018, *4*.

**Borenstein, Severin, James Bushnell, and Erin Mansur**, "The Economics of Electricity Reliability," *Journal of Economic Perspectives*, December 2023, *37* (4), 181–206.

**_ , Meghan R. Busse, and Ryan Kellogg**, "Career Concerns, Inaction and Market Inefficiency: Evidence From Utility Regulation," *The Journal of Industrial Economics*, 2012, *60* (2), 220–248.

**Brown, Stephen P.A. and Mine K. Yücel**, "Deliverability and regional pricing in U.S. natural gas markets," *Energy Economics*, 2008, *30*.

**Bushnell, James, Michaela Flagg, and Erin Mansur**, "Capacity Markets at a Crossroads," Technical Report WP 278, Energy Institute at Haas April 2017.

**Cremer, Helmuth, Farid Gasmi, and Jean-Jacques Laffont**, "Access to pipelines in competitive gas markets," *Journal of Regulatory Economics*, 2003, *24* (1), 5–33.

**Davis, Lucas W and Erich Muehlegger**, "Do Americans consume too little natural gas? An empirical test of marginal cost pricing," *The RAND Journal of Economics*, 2010, *41* (4), 791–810.

**Freeman, Gerad M., Jay Apt, and John Moura**, "What causes natural gas fuel shortages at U.S. power plants?," *Energy Policy*, 2020, *147*, 111805.

**Hoy, Darrell, Nicole Immorlica, and Brendan Lucier**, "On-Demand or Spot? Selling the Cloud to Risk-Averse Customers," in "Web and Internet Economics (WINE 2016)," Vol. 10123 of *Lecture Notes in Computer Science* Springer 2016, pp. 73–86.

**Joskow, Paul and Jean Tirole**, "Reliability and competitive electricity markets," *The RAND Journal of Economics*, 2007, *38* (1), 60–84.

**Marks, Levi, Charles F. Mason, Kristina Mohlin, and Matthew Zaragoza-Watkins**, "Vertical Market Power in Interconnected Natural Gas and Electricity Markets," Working Paper 17-27, Resources for the Future 2017. https://media.rff.org/archive/files/document/file/RFF%20WP%2017-27.pdf.

**Marmer, Vadim, Dmitry Shapiro, and Paul MacAvoy**, "Bottlenecks in regional markets for natural gas transmission services," *Energy Economics*, 2007, *29*.

**Oliver, Matthew E., Charles F. Mason, and David Finnoff**, "Pipeline congestion and basis differentials," *Journal of Regulatory Economics*, 2014, *46*.

**Safavi, Leila, Jiaxing Wu, Chiara Lo Prete, and Burçin Ünel**, "Rethinking Gas-Electric Coordination: Market Design Reforms for Electric Power Reliability During Cold Weather Events," Technical Report, Institute for Policy Integrity July 2025.

**Wang, Bin, Michela Biasutti, Michael P. Byrne, Christopher Castro, Chih-Pei Chang, Kerry Cook, Rong Fu, Alice MGrimm, Kyung Jaha, Harry Hendon et al.**, "Monsoons climate change assessment," *Bulletin of the American Meteorological Society*, 2021, *102*.

# Web Appendix for "Who Bears the Risk of Extreme Weather? Capacity Allocation in U.S. Natural Gas Pipelines"

by Yanyou Chen, Emmanuel Murray Leclair, and Adam Wyonzek

The Web Appendix is organized as follows. Appendix A documents the market structure and regulation of local distribution companies, including their concentration and service territories, rate regulation and cost pass-through, peak-day regulation by state before and after 2021, curtailment priority orders, and the evidence on gas-fired generation and heating demand during cold waves. Appendix B describes the rate regulation framework governing U.S. pipelines, including the conditions under which the Federal Energy Regulatory Commission grants market-based rate authority, and the mechanics of pipeline nomination cycles and contract-level utilization. Appendix C describes how the capacity release market works in practice, including the auction procedures for short-term and long-term releases. Appendix D documents our interviews with industry practitioners on contracting practices, the role of marketers, and the credit requirements that shape market participation. Appendix E contains proofs for the theoretical results in the main text. Appendix F describes our methodology for identifying cold waves from daily county temperature data from NOAA, the calibration of the detector on four documented cold waves, the resulting set of events with their state-level onset dates and anticipation windows, and the classification of events by how they arrive. Appendix G reports robustness checks for the empirical analysis.

# A Market Structure and Regulation of Local Distribution Companies

This appendix documents the institutional facts about local distribution companies (LDCs) that underlie the capacity obligation in Section 2.3 and the LDC's problem in Section 5.4. We focus on seven states that account for a large share of residential gas demand and that span the range of regulatory regimes in our sample: New York, New Jersey, Illinois, Colorado, Texas, Oklahoma and Kansas. Market structure data come from the U.S. Energy Information Administration's Form EIA-176 for 2023 (U.S. Energy Information Administration, 2023b). Regulatory information comes from state statutes, commission orders and utility tariffs, which we cite individually below.

## A.1 Concentration and Service Territories

Retail gas distribution is highly concentrated within each state. Table A.1 reports the share of residential and commercial volumes delivered by the three largest LDCs in each state. The top three utilities deliver between 62% and 96% of residential volumes, and the largest utility alone often delivers more than half. More importantly for our purposes, LDCs do not compete for customers. Each LDC holds an exclusive franchise over a defined service territory, so within its territory it is the only supplier of distribution service to households. Service territory maps published by state commissions confirm that the territories of major utilities do not overlap.[1] The large LDCs in every state we examine are investor-owned, for-profit corporations. Municipal and cooperative utilities exist but serve a small share of customers and are generally outside the rate jurisdiction of the state commission.

**Table A.1:** Concentration of Retail Gas Distribution by State, 2023

| State | Top-3 share, residential volume | Top-3 share, commercial volume | Largest LDC |
|---|---|---|---|
| New York | 61.6% | 66.3% | National Grid (Brooklyn Union Gas) |
| New Jersey | 90.3% | 91.6% | Public Service Electric & Gas |
| Illinois | 92.0% | 90.9% | Nicor Gas |
| Colorado | 92.3% | 89.8% | Public Service Co. of Colorado (Xcel) |
| Texas | 85.5% | 79.1% | Atmos Energy |
| Oklahoma | 96.0% | 95.4% | Oklahoma Natural Gas |
| Kansas | 93.4% | 83.7% | Kansas Gas Service |

*Notes:* Shares are computed from volumes delivered to residential and commercial customers reported on Form EIA-176 for 2023. In New York, the three largest residential distributors are Brooklyn Union Gas, Consolidated Edison and KeySpan Energy; the first and third are both National Grid subsidiaries, so the effective concentration is higher than the table suggests.

## A.2 Rate Regulation and Cost Pass-Through

Investor-owned LDCs are regulated by state public utility commissions: the Public Service Commission in New York, the Board of Public Utilities in New Jersey, the Commerce Commission in Illinois, the Public Utilities Commission in Colorado, the Railroad Commission in

[1] See, for example, the New York State gas utility service territory map, the Railroad Commission of Texas LDC service area map, and the Kansas Corporation Commission map of certified gas service areas.

Texas, and the Corporation Commissions in Oklahoma and Kansas. Three features of rate regulation are common to all seven states.

**Just and reasonable rates.** State statutes require that every rate charged by a gas utility be "just and reasonable," that rates be filed with the commission and open to public inspection, and that the utility charge neither more nor less than its filed schedule.[2] Rates are set in cost-of-service proceedings that allow the utility a reasonable opportunity to recover its operating costs and earn a return on invested capital. A utility that wants to change its rates must file notice in advance, typically 30 to 45 days, and the commission may suspend the change pending a hearing, for up to 14 months in New York and 240 days in Kansas. In these hearings the utility bears the burden of proving that the proposed rate is just and reasonable.

**Commodity pass-through.** The commodity cost of gas is not regulated and is passed through to customers through a purchased gas adjustment (PGA) or gas cost adjustment clause.[3] The clause is designed so that the utility earns zero margin on the commodity: the utility bills customers a base gas cost, reconciles actual purchase costs against billings, and refunds or surcharges the difference in a later period. The pass-through is not immediate, however. Reconciliations are annual in most states, adjustments are filed monthly or quarterly with advance notice, and commissions review purchases for prudence before allowing recovery. The Kansas Corporation Commission, for example, conducts monthly reviews of whether utilities are purchasing more gas than they need or paying more than the market price. The result is the pattern documented in Section 4: residential prices track wholesale prices with a lag and with much lower volatility, so that in a cold wave the spot price can exceed the residential price by a wide margin; the Energy Information Administration documents the same lag of residential prices behind spot prices nationally (U.S. Energy Information Administration, 2023a).

**Recovery of extraordinary costs.** When wholesale prices spike, the lag in pass-through and the prudence standard together expose the utility to the risk that extraordinary costs will not be recovered, or will be recovered only slowly. After Winter Storm Uri, utilities in the south central states incurred wholesale costs that were 10 to 100 times higher than normal. Recovery required new legislation: Texas House Bill 1520 and Senate Bill 3

[2]New York Public Service Law §§65–66; New Jersey Statutes Title 48:2-21; Illinois Public Utilities Act §§9-101, 9-102; 4 Code of Colorado Regulations 723-4; Texas Utilities Code §§104.003, 104.005; Oklahoma Administrative Code 165:45-15-2; Kansas Statutes Annotated §66-1202.

[3]16 New York Codes, Rules and Regulations §720-6.5; New Jersey Basic Gas Supply Service under N.J.S.A. 48:3-51; 83 Illinois Administrative Code Part 525; 4 CCR 723-4, Rules 4600–4603; 16 Texas Administrative Code §7.5519; Oklahoma Administrative Code Title 165, Chapter 50; Kansas Corporation Commission, Actual Cost Adjustment.

authorized customer rate relief bonds with maturities of up to 30 years (Texas Legislature, 2021); Oklahoma enacted the Regulated Utility Consumer Protection Act, which securitized extraordinary costs through state-issued bonds (Oklahoma Legislature, 2021); and Kansas enacted the Utility Financing and Securitization Act, which allows a non-bypassable surcharge for up to 32 years (Kansas Legislature, 2021). In Colorado, recovery was allowed through a two-year rider, and the legislature subsequently capped monthly gas cost recovery through Senate Bill 23-291. In Illinois, the largest utility absorbed the cost on its balance sheet and recovered only part of it in a later rate case. These episodes show that the utility, not the customer, bears the short-run cost of procuring gas at spot prices during a cold wave, and that recovery depends on regulatory and legislative discretion.

**Enforcement.** Formal penalties for failing to meet peak demand are rare. In New York and New Jersey, compliance with the design-day requirement is enforced through the commission's power to reject or modify a utility's supply plan rather than through fines. What disciplines the utility instead is the combination of prudence review, reputational cost, and political exposure that follows a curtailment. As Borenstein et al. (2012) document, utilities that are forced to curtail customers or to buy gas at extreme spot prices face a regulatory review whose severity depends on whether the utility can show it made a good-faith effort to secure supply in advance. The securitization statutes described above are the financial side of that review, and Texas subsequently introduced weatherization rules with penalties of up to $1 million per day (Railroad Commission of Texas, 2022).

## A.3 Peak-Day Regulation Across States and Over Time

Because the consequences of even a brief interruption of residential heating service are severe (American Gas Association, 2024), every state requires its utilities to plan for peak-day demand, although the form and the enforcement of the requirement vary. We now summarize how each state regulated peak-day supply before 2021 and how the regime changed afterward. Table A.2 classifies the regimes.

**New York.** From 2014 to 2018 the Public Service Commission ensured that peak-day demand was met through annual winter supply reviews, in which each utility submitted its supply and capacity plan, its strategy for meeting peak-day demand, and its treatment of interruptible customers. In 2019, physical supply constraints led Consolidated Edison and National Grid to announce moratoria on new gas connections. The Commission responded in March 2020 by opening a gas planning proceeding (New York Public Service Commission, 2020) and in May 2022 by adopting a new planning process (New York Public Service Commission, 2022). Utilities must now file long-term plans on a three-year cycle forecasting

**Table A.2:** Peak-Day Regulation of Gas Utilities by State

| State | Before 2021 | After 2021 | Enforcement |
|---|---|---|---|
| New York | Annual winter supply reviews of each utility's plan to meet peak-day demand (2014–2018); gas planning proceeding opened in 2020 | Long-term plans on a three-year cycle over a 20-year horizon, including a mandatory "no infrastructure" scenario and justification of reliance on marketers (2022) | Commission may reject or modify plans; no explicit penalties |
| New Jersey | Utilities set their own design-day standards (1-in-30 to 1-in-90 years); capacity investigation opened in 2019 | Standardized "Best Practices and Playbook for Emergencies" adopted for all gas utilities (2022) | Board evaluation; no explicit penalties |
| Illinois | Utilities report design-day demand, peak-day supply and reserve margin in rate cases (2014); energy efficiency targets | Climate and Equitable Jobs Act (2021); Future of Gas proceeding (2024) | Fines for missing efficiency plan and savings targets |
| Colorado | Traditional resource planning; utilities met peak demand through supply contracts and infrastructure | Clean Heat Law (2021); gas infrastructure plans every two years with design-day data (2023) | Cost cap on compliance; commission review of plans |
| Texas | No capacity mandate; formal policy of market-based supply determination | Senate Bill 3 weatherization mandate (2021); Rule 3.66 (2022) | Penalties up to $1 million per day |
| Oklahoma | No capacity mandate; utilities may interrupt service in emergencies | Securitization of extraordinary costs (2021) | Financial recovery focus |
| Kansas | No capacity mandate; tariff curtailment priority order | Securitization of extraordinary costs (2021) | Financial recovery focus |

*Sources:* New York Public Service Commission (2020, 2022); New Jersey Board of Public Utilities (2022); London Economics International LLC (2021); Illinois Commerce Commission (2014); Railroad Commission of Texas (2022); Texas Legislature (2021); Oklahoma Legislature (2021); Kansas Legislature (2021); Kansas Gas Service (2023). See the text for details.

a 20-year horizon, account for peak-day and peak-hour load, provide a "no infrastructure" scenario, and state explicitly how much they rely on third-party marketers during winter cold waves. The 2020 order noted that utilities had become increasingly reliant on marketers charging market prices to serve peak demand, and that marketer contracts could not be relied on for renewal. The orders contain no explicit penalties; the Commission may adopt, reject or modify a plan.

**New Jersey.** Until 2019, each gas distribution company developed its own plan to meet design-day demand, defined as the firm demand on a day of extreme cold. Two major utilities used a 1-in-30-year design day, while New Jersey Natural Gas used a 1-in-90-year standard. In 2019 the Board of Public Utilities opened an investigation into whether the state had sufficient capacity, after receiving conflicting assessments from New Jersey Natural Gas, which projected a shortfall by 2022–2023, and the Environmental Defense Fund, which argued that non-transparent design-day criteria made shortfall claims impossible to verify. The Board's consultant concluded that capacity would be sufficient in most scenarios but short on a 1-in-90-year day by 2030 (London Economics International LLC, 2021). In June 2022 the Board adopted a standardized emergency playbook requiring uniform protocols for supply shortages, procedures for calling on interruptible customers so that capacity can be reallocated to firm customers, and a planning process that prioritizes non-pipeline alternatives (New Jersey Board of Public Utilities, 2022).

**Illinois.** Since 2014, gas utilities have been required to report their forecasted design-day demand, available peak-day supply and reserve margin in general rate cases (Illinois Commerce Commission, 2014). Illinois also imposes energy efficiency targets with explicit fines: a utility that fails to file its four-year efficiency plan on time is fined $100,000 per day, and a large utility that misses its cumulative savings target must pay $600,000 to low-income assistance programs. Peak demand policy since 2021 has focused on decarbonization through the Climate and Equitable Jobs Act and the Future of Gas proceeding.

**Colorado.** Before 2021, utilities met peak demand primarily through supply contracts and infrastructure. The Clean Heat Law of 2021 imposed emission reduction targets on gas utilities, and since 2023 utilities must file gas infrastructure plans every two years that report the data and assumptions behind their design-day temperature and peak demand forecasts, and analyze non-pipeline alternatives before building new pipelines.

**Texas.** Before Winter Storm Uri, Texas imposed no capacity or weatherization mandate on gas utilities and maintained an explicit policy of market-based supply determination rather than mandatory resource planning. Regulation of gas by the Railroad Commission and of electricity by the Public Utility Commission was not coordinated. The joint FERC

and NERC inquiry into the storm attributed about three quarters of generation outages to freezing and fuel supply problems (FERC and NERC Staff, 2021). Senate Bill 3, enacted in May 2021, mandated weatherization of all gas production, transmission and pipeline facilities, and Railroad Commission Rule 3.66, effective in 2022, implements the mandate with a point-based violation system and penalties of up to $1 million per day for the most serious violations (Texas Legislature, 2021; Railroad Commission of Texas, 2022). Senate Bill 3 did not introduce a capacity requirement for gas utilities.

**Oklahoma and Kansas.** Neither state imposed a capacity mandate before 2021. Oklahoma rules allow a utility to interrupt service without notice in an emergency and to choose, in good faith, which customers to interrupt. The Kansas Gas Service tariff specifies a three-tier curtailment priority in which residential customers are in the first category and commercial and industrial customers in the second and third, and allows the utility to purchase gas from transportation customers in order to serve first-priority customers (Kansas Gas Service, 2023). After Uri, both states focused on financial recovery through securitization (Oklahoma Legislature, 2021; Kansas Legislature, 2021). During the storm the Kansas Corporation Commission ordered utilities to "do everything necessary" to maintain service, without specifying how.

## A.4 Curtailment Priority and the Position of Gas-Fired Generators

Curtailment priority orders place residential and other human-needs customers first in every jurisdiction we have examined. This ranking matters for gas-fired generators in two ways. Generators that are connected behind an LDC's city gate take service under the LDC's tariff and are subject to its curtailment order, so the LDC may curtail them, even if they hold firm service, in order to serve residential demand (MIT Energy Initiative, 2013). Generators that are connected directly to the interstate pipeline are not subject to the LDC's tariff, but they compete with the LDC for the same pipeline capacity, and because many of them hold only interruptible or short-term capacity, they are the first to be curtailed by the pipeline. The model and the empirical analysis in the main text cover only the second channel: competition for interstate capacity between LDCs and generators, or the marketers serving them, that hold their own pipeline contracts. Generators behind the city gate are exposed to both channels, so to the extent that they matter, our results understate the degree to which generators are crowded out during cold waves. Grid operators have documented the consequences of both channels: rising reliance on non-firm fuel among New York generators (New York Independent System Operator, 2023), competing demand from LDCs as a driver

of winter fuel scarcity in New England (Form Energy, 2023), and illiquid weekend gas supply for generators once LDCs and industrial customers have procured their needs (ISO New England and Midcontinent Independent System Operator and PJM Interconnection and Southwest Power Pool, 2024). The joint report of ISO New England, MISO, PJM and SPP acknowledges that LDCs have a legitimate need to reserve capacity for their core customers, but notes the lack of consistent standards for how much capacity is appropriate and for the release of unneeded capacity into the secondary market, and proposes that states revise the definition of human-needs customers to recognize that these customers also depend on electricity (ISO New England and Midcontinent Independent System Operator and PJM Interconnection and Southwest Power Pool, 2024). California has moved in this direction: in 2016 the California Public Utilities Commission approved a settlement modifying the curtailment priorities of SoCalGas and San Diego Gas & Electric to account for peak electric generator loads.

Industry participants describe the consequences of this ranking in similar terms after each major cold wave. Following Uri, generators told the Railroad Commission of Texas that curtailment rules had prioritized LDCs and left insufficient gas for power plants (Texas Alliance of Energy Producers, 2021). After Winter Storm Elliott, electric cooperatives testified to FERC that LDCs prioritized retail customers over generators, and that generators face a choice between paying for gas they may not use and paying penalties for failing to generate (National Rural Electric Cooperative Association, 2023). None of this requires strategic withholding by utilities: a 2017 study argued that New England utilities had withheld pipeline capacity to raise prices (Marks et al., 2017), but a subsequent FERC staff inquiry found no evidence of such behavior (Federal Energy Regulatory Commission, 2018). LDCs respond to a regulatory obligation in the way the obligation intends, and the allocation that results shifts the risk of extreme weather onto the users without priority capacity.

## A.5 Gas-Fired Generation and Heating Demand During Cold Waves

The deliverability constraint of Section 2 and the shape of the demands assumed in the model rest on a few documented facts about gas-fired generation. Interruptible transportation service, which is scheduled on whatever capacity remains after firm nominations, is inconsistent in normal times and is typically unavailable during periods of high demand, so that on the day a cold wave hits the pipeline allocation is essentially locked (MIT Energy Initiative, 2013). Gas-fired generators are frequently dispatched only when electricity demand peaks, and many of them cannot recover the fixed cost of firm transportation through their electricity market bids, so they rely on interruptible or short-term arrangements (MIT Energy

Initiative, 2013); the same report documents the growing interdependence of the gas and electricity systems. Gas-fired peaking units are dispatched last in the merit order and run only when electricity demand peaks (U.S. Energy Information Administration, 2012), while residential and commercial heating loads track heating degree days (U.S. Energy Information Administration, 2023c). This is why, in the model, generators' demand for gas is flat in low- and moderate-demand states of the world and spikes in extreme ones, whereas residential heating demand rises smoothly with the severity of the weather. The two demands add up in the coldest weeks: during the January 2024 cold snap, U.S. natural gas consumption set an all-time record as gas-fired generation surged on top of heating demand (U.S. Energy Information Administration, 2024).

The households served by these generators are the same ones the capacity obligation protects. Electricity supplies roughly 43 percent of residential heating services in the United States, compared with 47 percent for natural gas (U.S. Energy Information Administration, 2022), so a generator that is crowded out of pipeline capacity during a cold wave is providing heating to households through electricity, for the same health and safety reasons that place residential gas customers at the top of the curtailment order.

# B Rate Regulation in the US Pipeline Industry

Interstate pipeline rates are regulated by FERC under cost-of-service principles. The regulator determines a revenue requirement, the yearly revenue the pipeline needs to maintain service and earn a fair return, as the sum of operations and maintenance costs, taxes, depreciation, and a return on the rate base. The rate base is the pipeline's gross plant in service less accumulated depreciation. The return is computed from an assumed capital structure set by the regulator: the debt portion earns the pipeline's cost of debt, typically based on its outstanding debt, and the equity portion earns a cost of equity determined in regulatory proceedings using financial models calibrated to similarly risky assets, mainly other pipeline systems. The maximum reservation and usage charges in a pipeline's tariff are then set so that expected revenues at contracted volumes equal the revenue requirement. Pipelines may discount below the maximum rate, and, as discussed next, a small number of pipelines have been granted authority to charge market-based rates.

The United States pipeline regulatory authority, the Federal Energy Regulatory Commission (FERC), has both an established framework and legislation allowing it to approve pipelines of charging market-based rates. Much of this framework and legislation was developed in the late 80s and early 90s and today a handful of pipelines charge market-based rates, mostly oil pipelines. While the FERC grants oil and natural gas pipelines market-based rate

authority under two separate pieces of legislation generally the following must be included by a pipeline as evidence to move towards market-based rates:

- Describe proposed service.
- Define relevant product and geographic markets.
- Provide applicant's ownership and then list affiliated energy companies, services provided, and their location. If an affiliate operates in the same geographic market, the market shares of the applicant and the affiliate should be combined.
- Identify good alternatives to the proposed service - which parties provide similar service within the same geographic market. List the applicant's competitors and location.
- Include market share and Herfindahl-Hirschman Index (HHI) calculations to measure market concentration. The Commission's traditional HHI threshold is 1,800.
- Discuss other relevant competitive factors such as ease of entry and excess capacity held by competitors.
- Describe how the applicant's rates compare to the competitors.

In addition to the evidence above some pipelines have also provided econometric analysis to show that its shippers have a relatively high elasticity of demand, as a means to quantify the competitive position these pipelines face. It has also been required that pipelines provide cross-price elasticities for different products transported in their pipeline or by competitors. For example, a liquids pipeline may be constructed to ship various types of petroleum products (heavy crude, light crude, condensate, refined products, etc.) vs. another pipeline within the same area only capable of transporting a single type of petroleum product. In that case the regulator does not seem the product transported through the pipeline as homogenous and may review cross-price elasticities.

The FERC also has authority to determine if only a certain portion of a system may be granted market-based rate authority. For example, the origin market for pipelines is often considered as the production areas, which often could be considered not sufficiently concentrated and may result in the pipeline having market power in the origin market. However, many pipelines, before reaching their ultimate destination, pass through storage hubs. These storage hubs often have many pipelines entering them and are seen as a more competitive market. Therefore, a pipeline may not be granted market-based rates for its entire system but only a portion of its system, say from these storage hubs to its destination market.

As is evident the United States has a much more established framework for determining market-base rate authority. This is largely due to the more competitive nature of the pipeline industry in the United States in comparison to Canada. While this framework is established it is also important to note that this represents a relatively steep bar to cross and there are very few pipelines that are granted this authority.

**Throughput and utilization.** Most pipelines split pipeline flow into 4 nomination windows/cycles that create 3 periods of potentially unique flow conditions.

1. 7am – 12pm
2. 12pm – 4pm
3. 4pm – 7am (next day)

The scheduled capacity is what they are actually nominating to flow in any of these particular periods and the maximum capacity is essentially the size of their contract and available capacity is what they do not utilize. Note if these are firm contracts they still pay the full demand rate on the maximum capacity and pay variable charges (extremely small compared to the demand rate, essentially only fuel) on the scheduled capacity for that nomination window. However, lots of times the less sophisticated shippers will just set their scheduled capacity and leave it the same for all 4 windows, even set that for a full week or month and generally if there is no nomination the pipeline assumes it to be the previously nominated amount.

So taking the average of 4 nomination cycle scheduled capacity is a good proxy for daily flow of the contracts, with a few caveats:

1. The period of 12am – 6:59am of any given days capacity will be dependent on the intraday 2 nomination from the previous day, we could get around this by defining a day as 7am-7am (the next day).
2. When defining this utilization we want to be very specific of what we are saying here. Particularly this is utilization of an existing contract within a day/week. I think it would be too definitive of a statement to say this is the utilization of the entire pipeline because we do not have information on things like interruptible/ancillary service being sold by the pipeline after firm nominations are tallied.

The ability for pipelines to vary their capacity on a day to day basis is very limited and could be influenced by the following factors however they would likely not result in large volumes of sustained capacity:

1. Operating Pressure: Pipeline sections have licensed maximum and minimum operating pressures. You could potentially increase flow through a section of pipeline by either increasing the pressure in that section or decreasing the pressure in a downstream section. However this would still result in limited available capacity.

2. Heat Value/Ambident Temperature Assumptions: Two assumptions necessary for determining pipeline capacity is heating value (which is the assumption of the composition of the molecules that make up the natural gas and at what energy they burn at) and ambident temperature (the assumption of the temperature around the pipeline underground). Switching these assumptions could result in further capacity however they typically like to keep these assumptions very true to actual composition and heat of the pipeline to avoid any potential pipeline integrity issues. Further, pipelines have specifications for the composition of gas within a pipeline and have maximum hydrocarbon dewpoints (the measure of heat value).

3. Contracting Practices: Finally the commercial construct available for contracts may limit the ability of a pipeline to provide excess capacity on a day to day basis. Firm contracts have a minimum period and while some pipelines have short-term firm service most of the changes of the above assumptions would not likely guarantee a sustained period of stable service and so the pipelines are not likely to offer firm service by tweaking the above mentioned assumptions or pressure. They could likely create small amounts of capacity for a limited time which they would just sell as interruptible service.

# C How the Capacity Release Market Works in Practice

We first explain in more detail the capacity release programs:

1. Capacity releases can be put forward by the pipeline or requested from shippers. Both get posted on the pipeline's website.

2. There are generally two types of capacity releases biddable and non-biddable. Biddable capacity releases are subject to an open season (bidding process) where the individual who bids the highest gets the capacity. Non-biddable capacity releases are generally agreements that are entered into before posting on the regulatory website and are only eligible for releases with contract lengths of 31 days or less, or more than 1 year. This

is why we see most of the releases at 31 days in length or 1 year. Furthermore, capacity releases less than one year can be above the Tariff specified firm toll while those above a year cannot exceed the firm Tariff rate.

3. We discussed another very interesting feature the recall and reput options. Essentially if recall and reput conditions are included in a capacity release transaction then the original entitlement holder can at any time (within nomination window deadlines) "recall" that capacity and utilize it if they need it, this waives any tolls for the shipper that purchased the capacity release when it was recalled. Once the capacity is recalled the original entitlement holder can decide to "reput" that capacity if it no longer needs it and give it back to the shipper that originally purchased the capacity release. A recall/reput have to happen on separate days. From what Alan said this is a very common provision and we can see on the capacity release data from the pipeline whether these conditions are included in the contract.

4. Capacity release above the maximum tariff rate (releases less than or equal to 31 days, greater than or equal to 1 year, or bidded on).

5. The pipeline receives the same reservation charge whether it be from the releaser or the replacement shipper where they could potentially benefit is from the volumetric charge.

Below, we outilne the timelines detailing the processes for conducting open season auctions for short-term and long-term releases, respectively.

**Short Term Release (Less than 1 year)**

1. The capacity release request is posted by 9:00am on a business day (the specific timing of which will vary depending on the pipeline to meet their nomination windows, I found 9am to be the most common).

2. The open season/auction period is held between 9:00-10:00am. This is a silent auction, no parties can see what others bid, and during this period bids can be withdrawn.

3. The pipeline company begins evaluation of the bids at 10:00am, contingencies are cleared, and a determination of the best bid is made based on the process specified by the releaser/requester.

4. Both parties are notified by 11:00am and confirmed by 12:00pm. The capacity release contract is awarded within an hour.

**Long Term Release (More than 1 year)**

1. The capacity release request is posted by 9:00am on a business day.
2. The open season/auction period is then held for three consecutive business days, such that the open season process ends at 10:00am three business days after the capacity release posting.
3. The pipeline company begins evaluation of the bids at 10:00am three days later, contingencies are cleared, and a determination of the best bid is made based on the process specified by the releaser/requester.
4. Both parties are notified by 11:00am and confirmed by 12:00pm. The capacity release contract is awarded within an hour.

# D Interviews with Practitioners

Below we document our interviews with practitioners that participate in the capacity release market:

- Q: Why someone would not have an incentive to just nominate 100% of their firm contract?

  A: This is due to there also being a certain amount of variable costs associated with transportation. We discussed these briefly before Yanyou. but I was not aware it could be such a factor, for example, pipeline abandonment surcharges or fuel costs for shipping on the pipeline. If you nominate a particular level of gas for that nomination window it is assumed you use up to 100% of the gas you nominated, therefore you incur variable costs on a volumetric basis equal to your nomination. While these are small in comparison to the overall demand charges (fixed costs) there is little to no incentive for certain types of customers (industrial user or a local distribution company) to incur those costs if they do not require the gas. Whereas marketers/natural gas traders will only choose to nominate their full amount of their contract if the contract path is "in the money": price of natural gas at destination plus variable cost of transportation is greater than the price of natural gas at origin.

- Q: How often do people choose biddable agreement?

  A: In terms of biddable or non-biddable he would be surprised if more than 1% of the capacity releases were the result of a biddable process. Most are prearranged deals

where you have negotiated before posting it on the pipeline's website. In particular, even in the US many relationships are established, and many players are aware of who to call if they need a capacity release.

- Q: What is the difference between a natural gas utility, marketer, and retailer?

  A: A natural gas utility is a regulated entity that distributes natural gas, sometimes known as a local distribution company (LDC), to end use customers such as residential, commercial, and industrial customers. LDCs will often have an exclusive franchise area of a city or region and in turn their rates charged for their services are regulated by state regulatory authorities. Utilities source the transportation of their gas from interstate and intrastate pipelines which are in turn regulated by federal or state regulatory authorities, respectively. These utilities are required by their regulators to hold a certain percentage of their peak day demand requirements in long-term contracts on these pipeline systems. Utilities average day demand requirements are often much lower than their peak day requirements and therefore utilities hold excess transportation capacity on interstate and intrastate pipeline systems.

  Natural gas marketers are unregulated companies that arrange for the purchase and sale of natural gas. Most typically marketers arrange for gas supply agreements with natural gas producers or purchase gas at trading hubs. They then sell this gas to end users or at other market hubs to earn profit based on the spreads between different prices of natural gas hubs. To transport natural gas marketers must also hold capacity on natural gas pipelines but are not required to hold any typical amount of capacity and often choose to be nimble in the amount of capacity they hold on pipeline systems. When arranging the supply of gas for an end user marketers will typically charge for the cost of gas, cost of transportation, and then a service fee for arranging the transactions.

  Natural gas retailers are also unregulated companies that arrange for the purchase and sale of natural gas. However, the main distinction between them and marketers is that retailers typically hold capacity on natural gas distribution systems and sell to end users such as residential, commercial, and industrial customers. The distinction between retailers and LDCs is that the provision of natural gas distribution (responsibility of LDCs) is regulated but the retail sales of natural gas is often competitive. End use consumers can choose to purchase gas directly from the LDC, which offers a regulated rate, or from natural gas retailers which have competitive rates. Natural gas retailers typically source their natural gas from natural gas marketers or in certain circumstances are vertically integrated with a marketing company. Since both marketers and retailers are unregulated entities they do not require supply obligations and

typically hold transportation or distribution capacity that is more in line with average day demand for natural gas, in contrast to peak day demand like utilities.

- Q: Why do we see a large portion of end users of natural gas utilize marketer's services instead of contracting for their own service on interstate natural gas pipelines?

  A: A significant portion of end users of natural gas do not actually hold transportation capacity on interstate pipeline systems, or they are not connected to natural gas distribution utilities. These customers are often industrial facilities, agricultural operations, smaller power generation facilities, and natural gas retailers. Thus, a large amount of these end users of natural gas rely on natural gas marketers to meet their gas needs, both in terms of supply and transportation. The reason these end users do not hold transportation capacity on pipelines is for two primary reasons: creditworthiness and balance sheet obligations.

  The credit obligations that are necessary for a company to hold pipeline transportation service are often quite steep and that most companies could not meet. For example, a typical credit evaluation criteria for firm service on a natural gas pipeline is to provide security guarantees for three months of firm service at the maximum tariff rate for the entire volume of your contract. This requires companies to have large amounts of cash on hand (in the form of an advance deposit), a strong standing letter of credit from a financial institution, an acceptable security interest in collateral, or a guarantee from a more credit worthy parent company.

  As for the balance sheet obligations, given the take or pay nature of natural gas transportation firm service contracts financial institutions view these transportation contracts as debt obligations. If a company were to take out large amounts of transportation capacity this would result in a large liability to appear on their balance sheet which may impact their own credit metrics, impacting their ability to secure their own financing and financial obligations.

  Given these restrictive requirements many of these end users rely on the services of marketers to arrange for supply and transportation of their gas needs, since marketers require lower levels of credit requirements and the less restrictive outcomes on their balance sheet obligations. This comes at an increased cost to the end user as marketers often require a service fee or mark up for arranging supply and transportation of natural gas.

- Q: Why do we see natural gas marketers as replacers in the capacity release market?

A: Since these end users rely on marketers to provide them service when unexpected shocks in demand happen either to end use residential, commercial, and industrial demand (increased demand for natural gas retailers) or unexpected shocks in their various industries that do not impact demand for gas utilities/retailers (for example an increase in demand for steel production) they often turn to marketers to supply them with additional natural gas. While the primary market for natural gas transportation is held largely by utilities, to meet their regulatory obligations, most periods of the year they do not require the full use of their transportation contracts and would prefer releasing that capacity to a marketer that will ultimately provide the transportation services to an end user.

# E Model Appendix

This appendix collects the proofs of the paper's theoretical results and supplementary derivations. Throughout, we use the notation of the main text: $\kappa_i$ is firm $i$'s net forward position, $P_i(q, A)$ its marginal utility (inverse demand) with unconstrained demand $D_i(p, A) \equiv P_i^{-1}(p, A)$, $K$ the physical capacity of the pipeline, and $p_g(A)$ the ex-post spot price.

## E.1 Proofs

**Lemma 1** (Existence and uniqueness of the ex-post price)**.** *Given net forward positions* $\{\kappa_i\}$ *with* $\sum_i \kappa_i = K$ *and a weather realization* $A$*, the market-clearing condition* (4) *has a unique solution* $p_g(A)$*.*

*Proof.* The left-hand side of (4), inverse supply $c^{-1}(p)$, is continuous and strictly increasing in $p$, ranging over $[0, \infty)$. The right-hand side, aggregate demand $D(p, A) = \sum_i \min\{\kappa_i,\ D_i(p, A)\}$, is continuous and non-increasing in $p$, with $0 < D(p, A) \leq K$. The two sides therefore cross at most once; and they cross at least once, because at $p = c(0)$ supply is zero while demand is positive, whereas for $p$ large enough supply exceeds $K \geq D(p, A)$. Hence a solution exists and is unique. □

Next, we prove Proposition 1, which we restate for convenience:

**Proposition 1** (First-best implementation)**.** *Suppose that the capacity constraint binds with positive probability. Then, under competitive ex-post pricing, the decentralized allocation coincides with the first best in* every *state* $A$ *if and only if SIMR holds.*

*Proof.* We prove the two directions in turn.

($\Rightarrow$) *First-best implementation implies SIMR.* We argue by contraposition: if SIMR fails, the decentralized allocation differs from the first best in some state. Note first that in interior states the market is always efficient, so a SIMR violation can matter only in states where capacity binds. The assumption that for $A$ large enough, $D_i > \kappa_i$ for every firm (the capacity constraint binds with positive probability) guarantees such state occur.

Failure of SIMR means that, for some pair of holdings $(\kappa_1, \kappa_2)$ satisfying $\kappa_1 + \kappa_2 = K$, and some states $A_0$ and $A_1$, the product in Definition 1 is strictly negative. Without loss of generality, suppose that

$$P_1(\kappa_1, A_0) \geq P_2(\kappa_2, A_0) \qquad \text{and} \qquad P_1(\kappa_1, A_1) < P_2(\kappa_2, A_1).$$

Thus, the firms' marginal-utility ranking reverses between the two states. In state $A_1$, firm 2's marginal utility at $(\kappa_1, \kappa_2)$ strictly exceeds firm 1's, so a first-best allocation would shift gas from firm 1 to firm 2. By contrast, such a shift is not optimal in state $A_0$. Hence, the same allocation cannot be first-best in both states.

($\Leftarrow$) *SIMR implies first-best implementation.* Fix a reference state $A_0$. Because each $P_i(\cdot, A_0)$ is continuous and strictly decreasing, there is a unique vector $\kappa^*$ with $\sum_i \kappa_i^* = K$ that equalizes marginal utilities at capacity: $P_i(\kappa_i^*, A_0) = P_j(\kappa_j^*, A_0)$ for all $i, j$. Under SIMR, ties are preserved across states, so the *same* vector $\kappa^*$ equalizes marginal utilities in *every* state $A$ for all $i$ and all $A$, where $M(A)$ denotes the common value (note that the value of $M(A)$ can change with different realized state $A$).

We next verify that $\kappa^*$ constitutes an ex ante equilibrium. At $\kappa^*$, firm $i$'s capacity constraint binds ex post precisely when its unconstrained demand exceeds its capacity holding, that is, when $p_g(A) < P_i(\kappa_i^*, A) = M(A)$. Because this condition is the same for every firm, both the indicator and the bracketed term in (FOC-1) are identical across firms. Consequently, every firm's first-order condition is satisfied at the common capacity price

$$\tau_p = \int_{\Omega_A} \mathbf{1}\{p_g(A) < M(A)\} \left[M(A) - p_g(A)\right] f_A(A)\, dA \;\geq\; 0.$$

Moreover, the capacity market clears because $\sum_i \kappa_i^* = K$. Hence, $\kappa^*$ is an ex ante equilibrium.

We next verify that the ex-post allocation is first best in every state. Because the binding criterion is common to all firms, partial congestion never occurs at $\kappa^*$. If no cap binds, the allocation is guaranteed first-best. If every cap binds, then $q_i(A) = \kappa_i^*$, aggregate consumption equals $K$, and marginal utilities are equalized, so the market allocation coincides with the planner's congested optimum. □

Next, we prove Proposition 2, which we restate for convenience:

**Proposition 2** (Precautionary over-contracting by LDCs)**.** *Suppose SIMR holds and the ex-post market clears competitively. Then equilibrium LDC capacity $\kappa_2(\mu)$ is strictly increasing in $\mu$ and the unregulated shipper's $\kappa_1(\mu)$ is strictly decreasing; hence $\kappa_2(\mu) > \kappa_2^*$ for all $\mu > 1$.*

*Proof.* At the benchmark $(\kappa_1^*, \kappa_2^*)$ with $P_i(\kappa_i^*) = \Gamma$, the LDC's FOC integrand is

$$\mathbf{1}\{p_g(A) < \mu A\Gamma - (\mu-1)p_c\}\ [\mu A\Gamma - (\mu-1)p_c - p_g(A)],$$

while the shipper's is $\mathbf{1}\{p_g(A) < A\Gamma\}\ [A\Gamma - p_g(A)]$. For $A > \bar{A} \equiv p_c/\Gamma$ the LDC's integrand strictly dominates pointwise and through extra constrained states. By assumption, competitive pricing congests with positive probability for $A > \bar{A}$ (Section 5.2), so the LDC's FOC integral exceeds the shipper's at the benchmark; restoring equality requires raising $\kappa_2$ (lowering the LDC's valuation) and lowering $\kappa_1$ (raising $P_1$). □

## E.2 The LDC's Curtailment Choice When "In the Money"

In the main text we assume that the regulated residential price never exceeds the spot price on the domain of weather shocks we study, so the LDC is always "out of the money." For completeness, this appendix characterizes the LDC's optimal curtailment choice without that restriction. When $p_g(A) \le p_c$, the LDC is "in the money": it earns a positive margin on every unit delivered and has no incentive to curtail, although it may be forced to curtail by its capacity constraint. There are two subcases:

1. $D_2(p_c, A) \le \kappa_2$: the LDC wants to give consumers their full demand and capacity isn't binding. There is no curtailment, $q_2^*(A) = D_2(p_c, A)$ and $L = 0$.
2. $D_2(p_c, A) > \kappa_2$: the LDC wants to give consumers their full demand, but capacity is binding at $\kappa_2$. There is curtailment, $q_2^*(A) = \kappa_2 < D_2(p_c, A)$, and $L > 0$.

In both subcases the LDC would like to serve residential demand at the retail price, so its consumers' marginal utility satisfies $P_2(q_2(A), A) \ge p_c$, with equality when capacity does not bind. Combining the in-the-money and out-of-the-money regimes, the LDC's ex-post delivery rule becomes

$$q_2(A) = \min\big\{\kappa_2,\ D_2\big(\max\{p_c, \tilde{p}_g(A,\mu)\}, A\big)\big\},$$

where the max operator selects the retail price when the LDC is in the money and the effective price $\tilde{p}_g(A,\mu)$ when it is out of the money. The ex-ante first-order condition then

contains two terms, depending on the regime in which capacity binds:

$$\tau_p = \int_{\Omega_A} \Big[ \mathbf{1}\underbrace{\{p_g(A) \le p_c,\ D_2(p_c, A) > \kappa_2\}}_{\text{in the money, constrained}} [P_2(\kappa_2, A) - p_c] + \mathbf{1}\underbrace{\{p_g(A) > p_c,\ D_2(\tilde{p}_g(A,\mu), A) > \kappa_2\}}_{\text{out of the money, constrained}} [P_2(\kappa_2, A) - \tilde{p}_g(A,\mu)] \Big] f_A(A) dA.$$

Finally, the in-the-money regime adds distortions of its own relative to the first best. Even if $\mu = 1$, a fixed price $p_c$ is distortionary whenever the LDC is in the money: the lack of curtailment leads to under-consumption relative to the first best, because consumers face $p_c > p_g(A)$. Similarly, if the LDC is in the money but $\mu < 1$, then curtailment also leads to under-consumption relative to the first best.

# F Identifying Cold Waves: Data, Method and Results

## F.1 Data

We collected daily geospatial temperature data from the National Oceanic and Atmospheric Administration (NOAA).[4] The underlying series is the daily mean temperature, in degrees Celsius, of each of the county of the United States from January 1, 2006 to December 31, 2023.

We calibrate the detector on four cold waves whose timing and geography are documented by the National Weather Service, grid operators and the press: the February 2015 "Siberian Express", the January 2019 polar vortex, Winter Storm Uri in February 2021 and Winter Storm Elliott in December 2022. Table F.1 summarizes what the sources establish for each.

## F.2 The detector

The detector is designed in the following steps: we first standardize each county's daily temperature against its own climatology for that time of year, and flag the county-days that are extreme by that standard. We then define a cold wave on the national share of flagged counties: an event begins when that share crosses a high threshold and ends only once it has fallen below a lower one. Finally, we aggregate the flags to regions and attach the cold spells of individual regions to the national event, which yields for every region its own onset date and the anticipation window before it.

[4] https://psl.noaa.gov/cgi-bin/data/narr/plotday.pl/.

**Table F.1:** The four known cold waves

| Event | Span (sources) | Regions | Forecast lead time |
|---|---|---|---|
| Feb 2015 Siberian Express | Early February to mid-March 2015; arctic blasts Feb 11–16, Feb 18–20, Feb 23–24, Feb 27–28 and Mar 5–7. | East of the Rockies; the Northeast and Great Lakes most persistently; Texas in the last three pulses. | Each pulse forecast 5–7 days ahead; no long-lead forecast of the whole episode. |
| Jan 2019 polar vortex | Main pulse Jan 29–31, 2019 in the Upper Midwest; an earlier pulse Jan 21–22 and a stronger nationwide pulse Mar 4–7. | Jan 30–31: MN, WI, IA, IL, IN, MI, OH, then the Northeast. | About one week: forecasts firmed up Jan 23–25. |
| Winter Storm Uri, Feb 2021 | Cold wave Feb 6–22; northern Plains Feb 6–17; Texas below freezing Feb 10–19, coldest Feb 15–16. | Northern Plains first, then TX/OK/AR/LA, the Midwest and the South. | Models flagged the pattern by Feb 5–7; ERCOT notice on Feb 8 for Feb 11–15; 7–10 days before the Texas peak. |
| Winter Storm Elliott, Dec 2022 | Storm Dec 21–26; Texas cold Dec 22–25; deep freeze ends Dec 27. | Colorado to the East Coast and south to Florida, nearly simultaneously. | Forecasts from Dec 16–18; pipelines issued operational flow orders on Dec 20; 5–7 days. |

The parameters are chosen so that the four known cold waves are each recovered as a single event with the right geography, and the same rule is then applied to the whole 2006–2023 sample.

**County anomaly.** For county $c$ on day $t$ with day-of-year $d(t)$ we compute

$$z_{ct} = \frac{T_{ct} - \mu_{c,d(t)}}{\sigma_{c,d(t)}},$$

where $\mu_{c,d}$ and $\sigma_{c,d}$ are the mean and standard deviation of $T_{ct}$ over all days within $\pm 7$ calendar days of $d$, pooled across 2006–2023. The anomaly removes both the seasonal cycle and cross-county differences in climate, so that a county in Texas and a county in Minnesota are compared on how unusual a day is for that county at that time of year.

**County cold day.** A county is cold on day $t$ if $z_{ct} < -2$ and $t$ falls in November through March. The seasonal restriction excludes shoulder-season anomalies that matter little for heating demand.

**National event.** A cold wave is defined based on the national share of cold counties, using hysteresis: a national event begins on the first day the share reaches 20 percent and ends after the share has remained below 12.5 percent for five consecutive days.

**Region cold day.** A region (a state, or an economic area) is cold on day $t$ if at least 30 percent of its counties are cold.

**Region spell.** A region spell is a maximal run of region cold days in which interruptions of at most 4 days are allowed. Every region spell that overlaps the national event, extended by 4 days on each side, is attached to it; a region may have several spells within one event. For each region within an event we record its own onset and end, and an anticipation window consisting of the 7 days before the region's onset.

| Parameter | Value | Role |
|---|---|---|
| $z$ threshold | $-2$ | county cold if $z <$ this |
| season | Nov–Mar | months in which a county can be cold |
| entry share | 0.20 | share of all U.S. counties that are cold on the same day, which a national event must reach |
| exit share | 0.125 | share of all U.S. counties that are cold; a national event ends once the share has stayed below this for more than the gap (4 consecutive days) |
| anticipation window | 7 days | days before a region's onset |
| region share | 0.30 | region cold if this share of its counties is cold |
| gap | 4 days | interruptions allowed inside a region spell and inside a national event |

**An example: Winter Storm Uri.** The steps are easiest to follow on one event. Each county carries a daily indicator equal to one when its temperature is more than two standard

deviations below its own normal for that time of year. The share of contiguous-U.S. counties with the indicator switched on is computed for every day, and the cold wave is defined from that single national series: the national event begins on the first day the share reaches 20 percent and continues until the share has stayed below 12.5 percent for more than four days. For Uri the national share is 6 percent on February 11 and 12 percent on February 12, crosses 20 percent on February 13, 2021 (21 percent), rises to 40 percent on February 14, stays above 12.5 percent through February 19 and falls to 12 percent on February 20 and under 1 percent on February 21, so the national event is February 13–19. At this stage the event has dates but no geography.

The geography comes from the states. Each state has its own daily share of cold counties, it is cold on a day when that share is at least 30 percent, and a spell is a run of such days with interruptions of at most four days. Every state spell that overlaps the national event, extended by four days on either side, is attached to it. A state's onset and end are the first and last day of its own spell, not the national dates, and they can fall outside the national event: Montana's share crossed 30 percent on February 11 and Kansas's on February 12, two days and one day before the country as a whole, while Texas stayed cold until February 20, a day after the national event ended. Table F.2 shows a selection of the 33 states attached to the event. Texas illustrates the construction: its share of cold counties was below 5 percent through February 12 and 25 percent on February 13, then jumped to 63 percent on February 14, so its onset falls one day after the national start, three days after Montana, two after Kansas and one after Iowa, Missouri, Nebraska and South Dakota, which enter with the national start.

The anticipation window of a state is the seven days before its own onset. For Texas it is February 7–13. A shipper serving Texas was making capacity-release decisions during those days, and in the last three days of that window Montana, Kansas and the central Plains were already cold, so the cold wave was observable in other markets as well as in forecasts. Because onsets differ across states, every state has its own window, and it is this variation within an event that allows the event study to compare markets that were hit on different days by the same cold wave. Two features of the rule are worth stressing. A cold wave begins only when a fifth of all counties nationally are extreme, so a cold snap confined to one region never becomes an event. And a state is not required to be cold on the national start date; it joins the event whenever its own spell overlaps the national window, which is why the 33 states attached to Uri have onsets spread from February 11 to 16.

**Table F.2:** Winter Storm Uri: national event and selected state onsets

| State | National event | Share reaches 30% | State onset | State end | Anticipation window |
|---|---|---|---|---|---|
| Montana | Feb 13–19 | Feb 11 (0.34) | Feb 11 | Feb 15 | Feb 4–10 |
| Kansas | Feb 13–19 | Feb 12 (0.41) | Feb 12 | Feb 17 | Feb 5–11 |
| Iowa | Feb 13–19 | Feb 13 (0.43) | Feb 13 | Feb 17 | Feb 6–12 |
| Minnesota | Feb 13–19 | Feb 14 (0.46) | Feb 14 | Feb 17 | Feb 7–13 |
| Texas | Feb 13–19 | Feb 14 (0.63) | Feb 14 | Feb 20 | Feb 7–13 |
| Louisiana | Feb 13–19 | Feb 15 (0.52) | Feb 15 | Feb 17 | Feb 8–14 |
| New York | Feb 13–19 | Feb 16 (0.33) | Feb 16 | Feb 16 | Feb 9–15 |

*Notes:* The national event starts on the first day the share of all U.S. counties that are cold reaches 20 percent and ends once that share has stayed below 12.5 percent for more than four days (February 13–19, 2021); it is the same for every state. *Share reaches 30%* is the first day on which the state's own share of cold counties is at least 30 percent, with the share in parentheses. *State onset* and *State end* are the first and last day of the state's cold spell, which may fall before the national start or after the national end. The anticipation window is the seven days before the state onset.

## F.3 Calibration

We choose the parameters by requiring that each of the four known cold waves be recovered as a single event that contains the states the sources identify (New York, Minnesota and Texas for 2015; Minnesota, Illinois and New York for 2019; Texas, Oklahoma and Minnesota for Uri; Texas, New York and Illinois for Elliott). Figure F.1 summarizes a grid over the $z$ threshold ($-1.5$ to $-2.5$), the region share (0.2 to 0.4) and the gap (2 to 5 days).

Because the national event is defined on the national series, the number of events depends only on the $z$ threshold and the gap (panel b): with a 4-day gap it falls from 71 at $z < -1.5$ to 55 at $z < -1.7$, 50 at $z < -1.8$, 41 at $z < -1.9$, 34 at $z < -2$, 19 at $z < -2.25$ and 11 at $z < -2.5$. The region share governs only which regions attach to an event, and therefore whether the key states are recovered (panel a). The gap matters for the 2015 cluster, whose warm break of March 1–4 lasts four days: at thresholds of $-1.8$ or stricter a gap of 3 splits it into two events, while a gap of 4 keeps it whole. Every threshold between $-1.5$ and $-2$ recovers all four known waves as single events with a 4-day gap and a region share of 20 or 30 percent; at $-1.9$ and $-2$ a 40 percent share loses New York from the 2019 event, and at $-2.25$ the February 2015 cluster splits into two events whatever the region share, because its

warm break then lasts five days. The threshold $z < -2$ is therefore the strictest one under which all four known waves survive, and we adopt it together with a 30 percent region share and a 4-day gap. It yields 34 events rather than the 71 obtained at $z < -1.5$, and it confines the events to the days on which the cold is genuinely extreme: milder thresholds add earlier and later days to each event and a number of weaker episodes, at the cost of blurring the onset dates that the event study relies on.

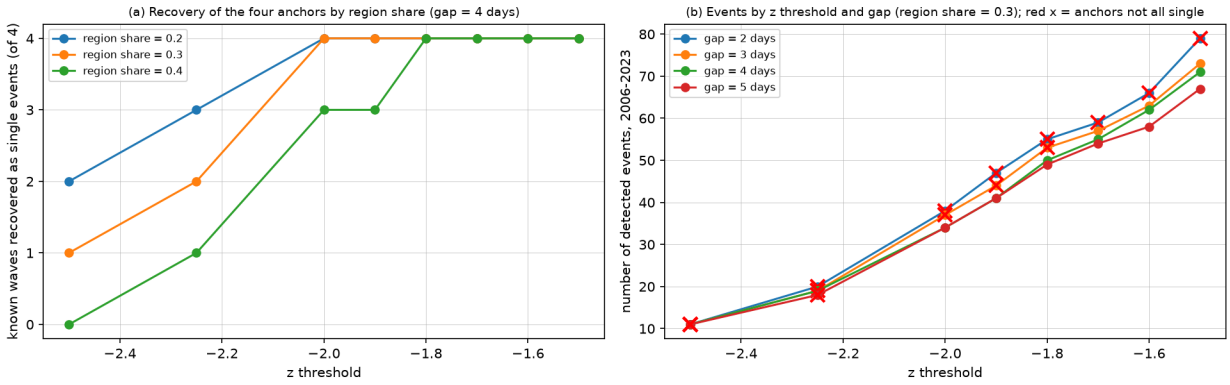



**Figure F.1:** Calibration grid. Panel (a): number of known waves recovered as single events, by region share, at the 4-day gap. Panel (b): number of detected events by $z$ threshold and gap; red crosses mark combinations under which the four known waves are not all recovered as single events.

## F.4 Results with the default parameters

The detector identifies 34 cold wave events between 2006 and 2023, about 1.9 per winter (Figure F.2). Nine of the 34 events are weak, with a peak national share between 20 and 24 percent and between 10 and 20 states attached; three of them fall in the winter of 2006–07 and one in 2007–08, two winters with no stronger event. No event is detected in the winters of 2005–06, 2011–12, 2015–16 and 2021–22, all of which were mild, and the winter of 2013–14 has the most, with six.

Figures F.3 to F.6 show the four known waves as state-by-day heat maps. Each cell is the share of a state's counties that are cold on that day, states are ordered roughly from north to south, the blue band marks the detected national event, the red lines mark the window established from the sources, and orange markers show each state's own onset and end.

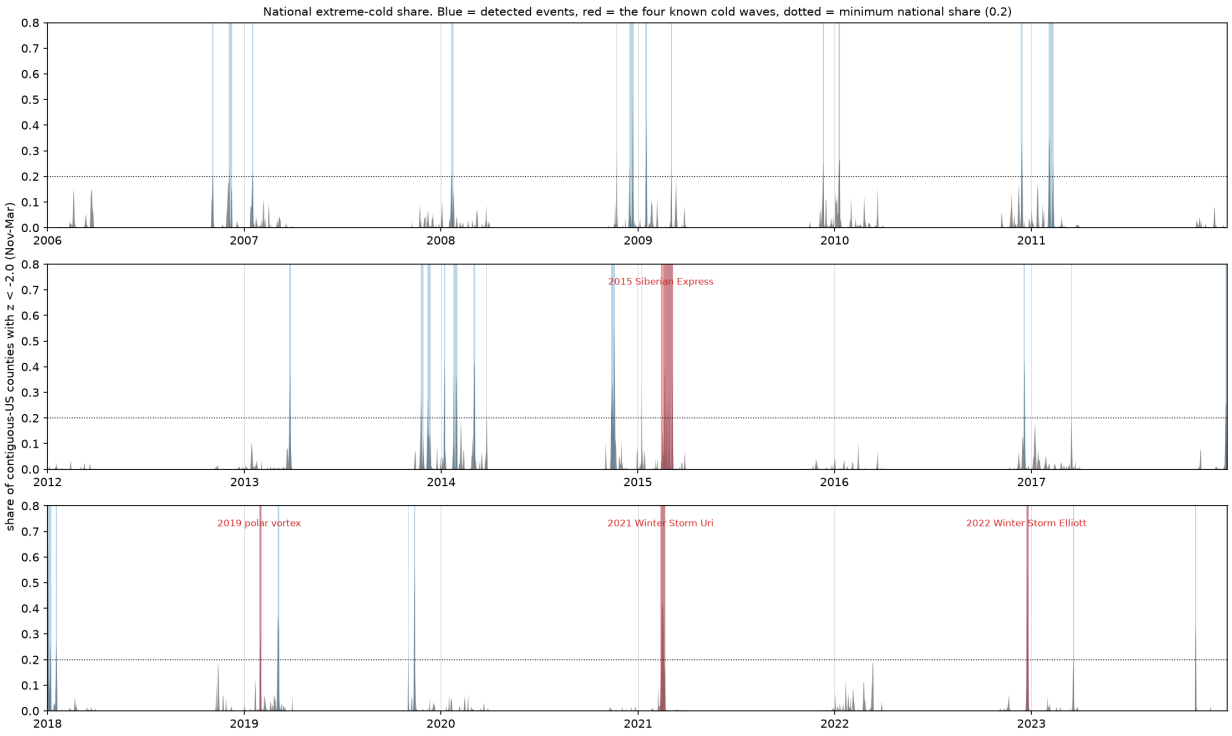



**Figure F.2:** Share of contiguous-U.S. counties with $z < -2$ (November–March), 2006–2023. Blue bands are detected national events; red bands are the documented windows of the four known waves; the dotted line is the 20 percent entry threshold.

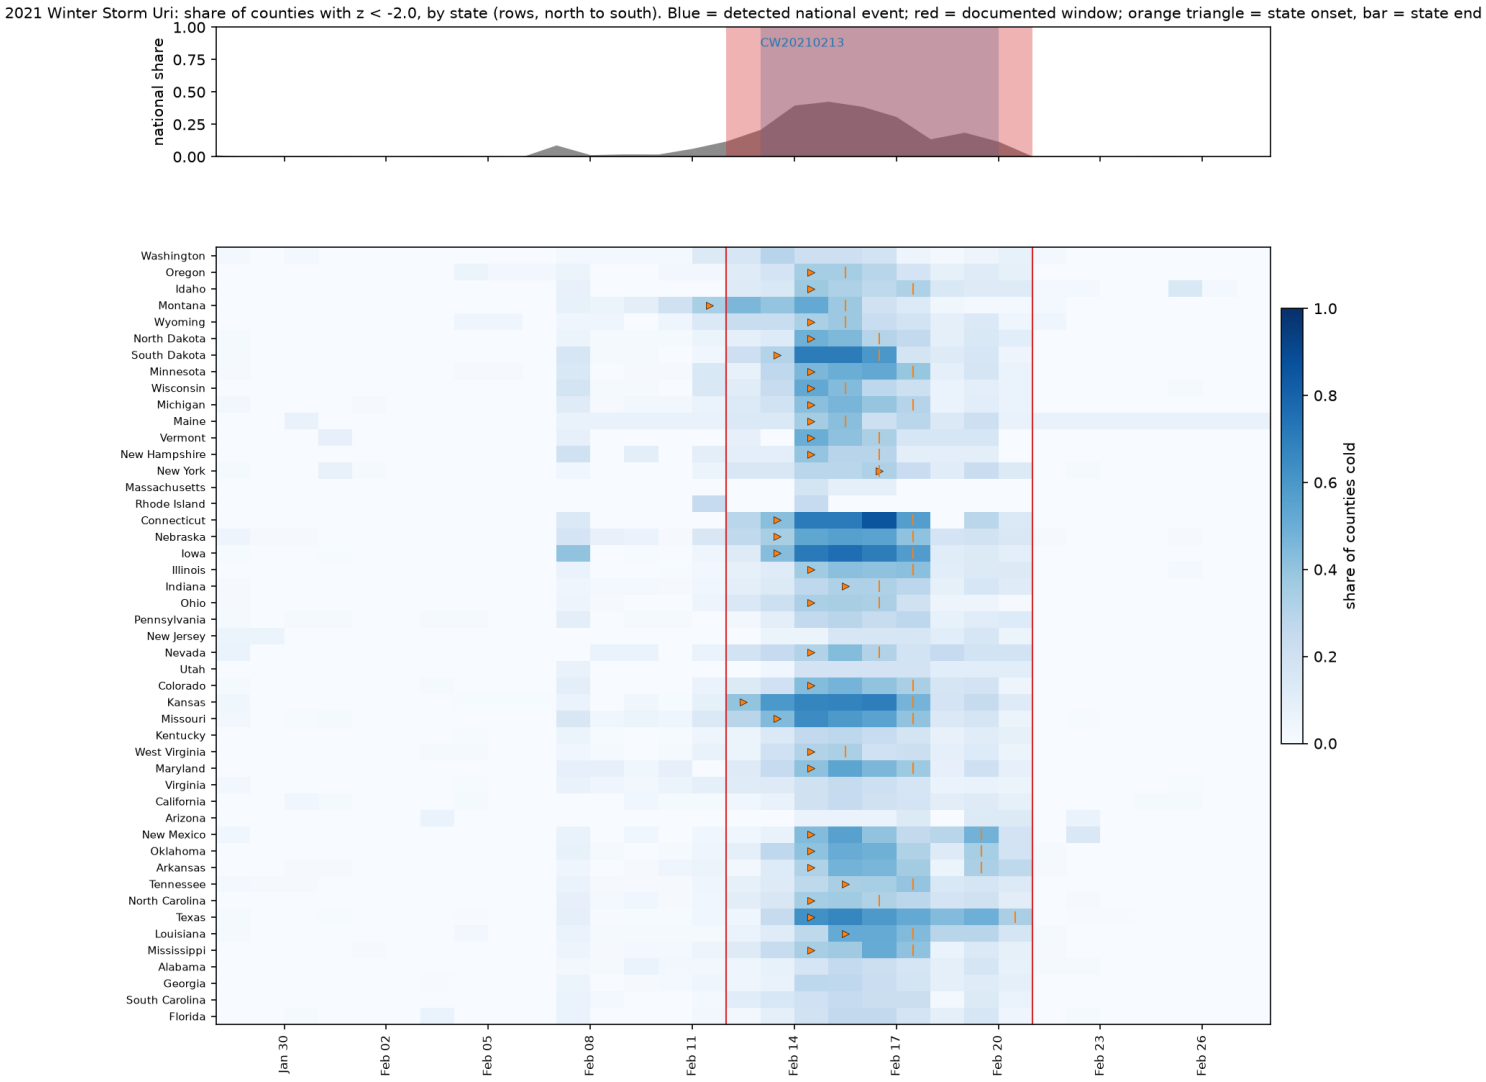



**Figure F.3:** Winter Storm Uri, February 2021. Top: national share of cold counties, with the detected national event shaded blue and the documented window shaded red. Bottom: share of cold counties by state and day; orange triangles mark state onsets and orange bars state ends.

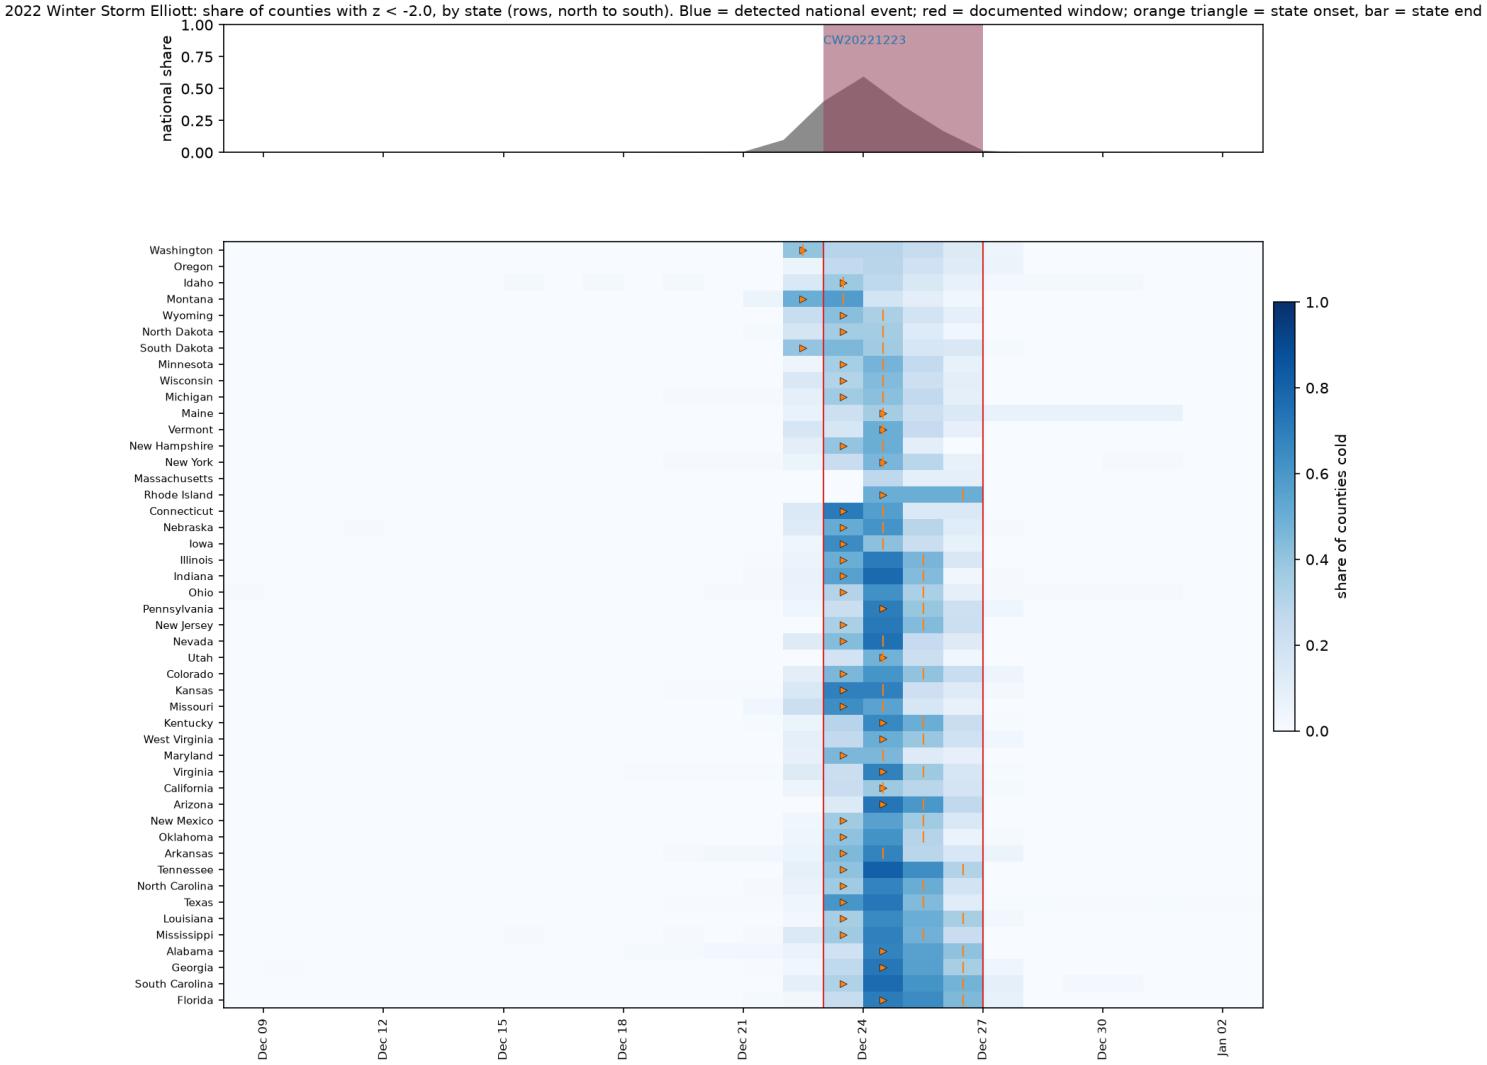



**Figure F.4:** Winter Storm Elliott, December 2022. Same layout as Figure F.3.

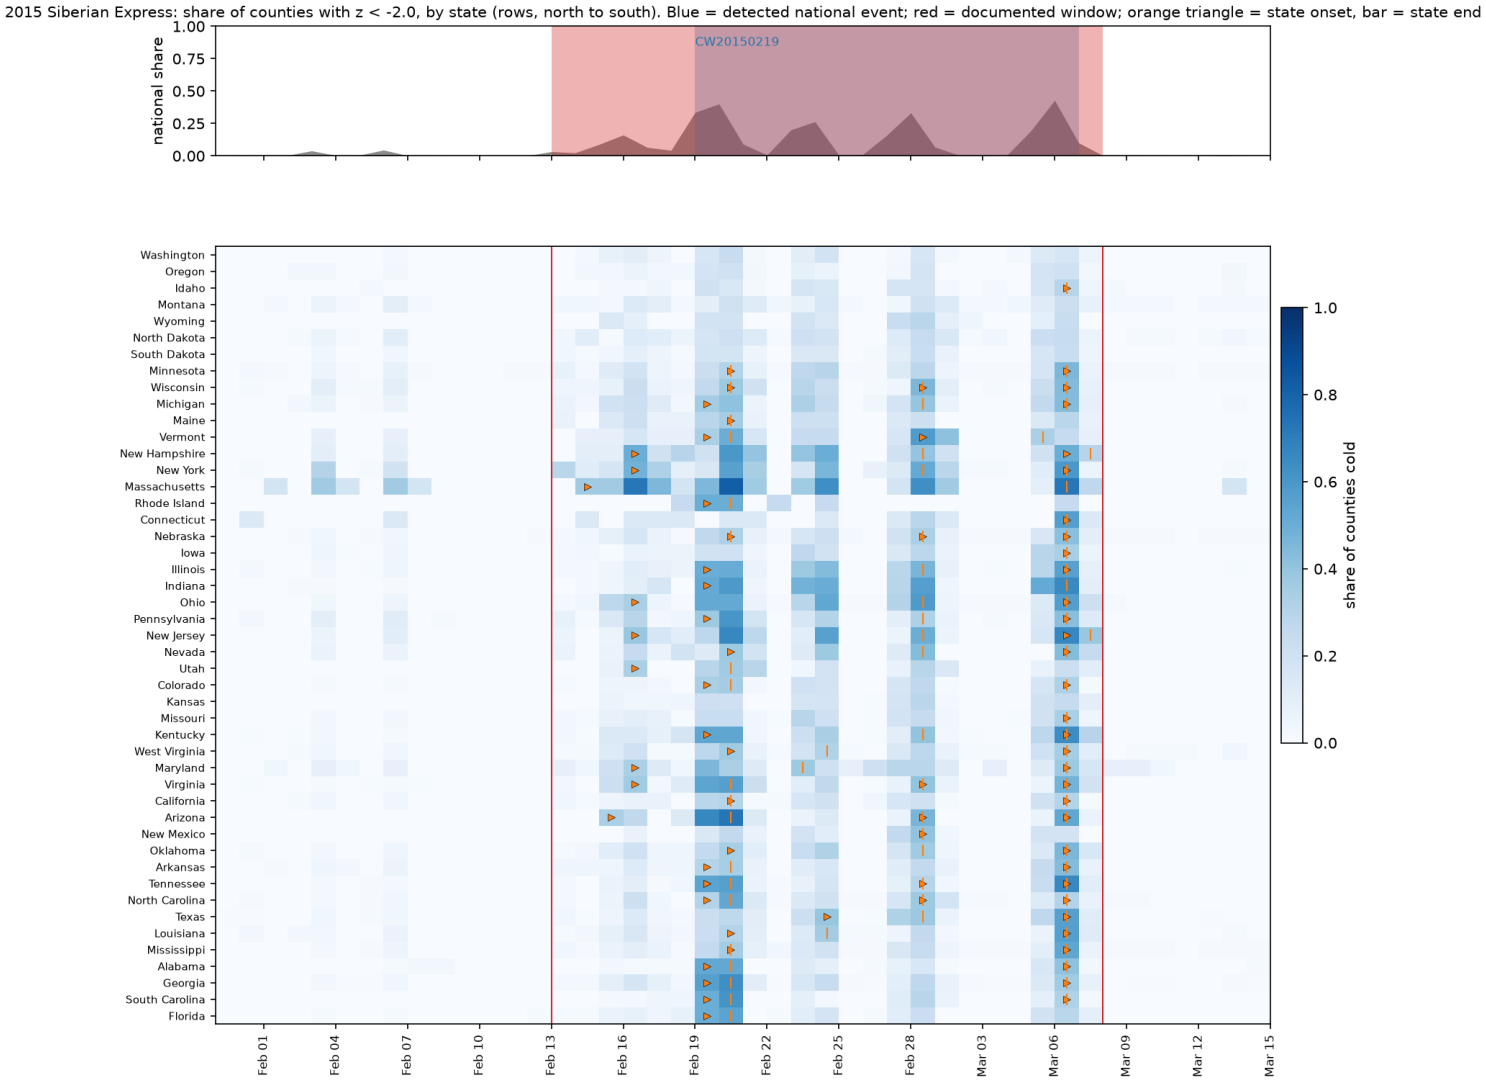



**Figure F.5:** The February 2015 Siberian Express. Same layout as Figure F.3.

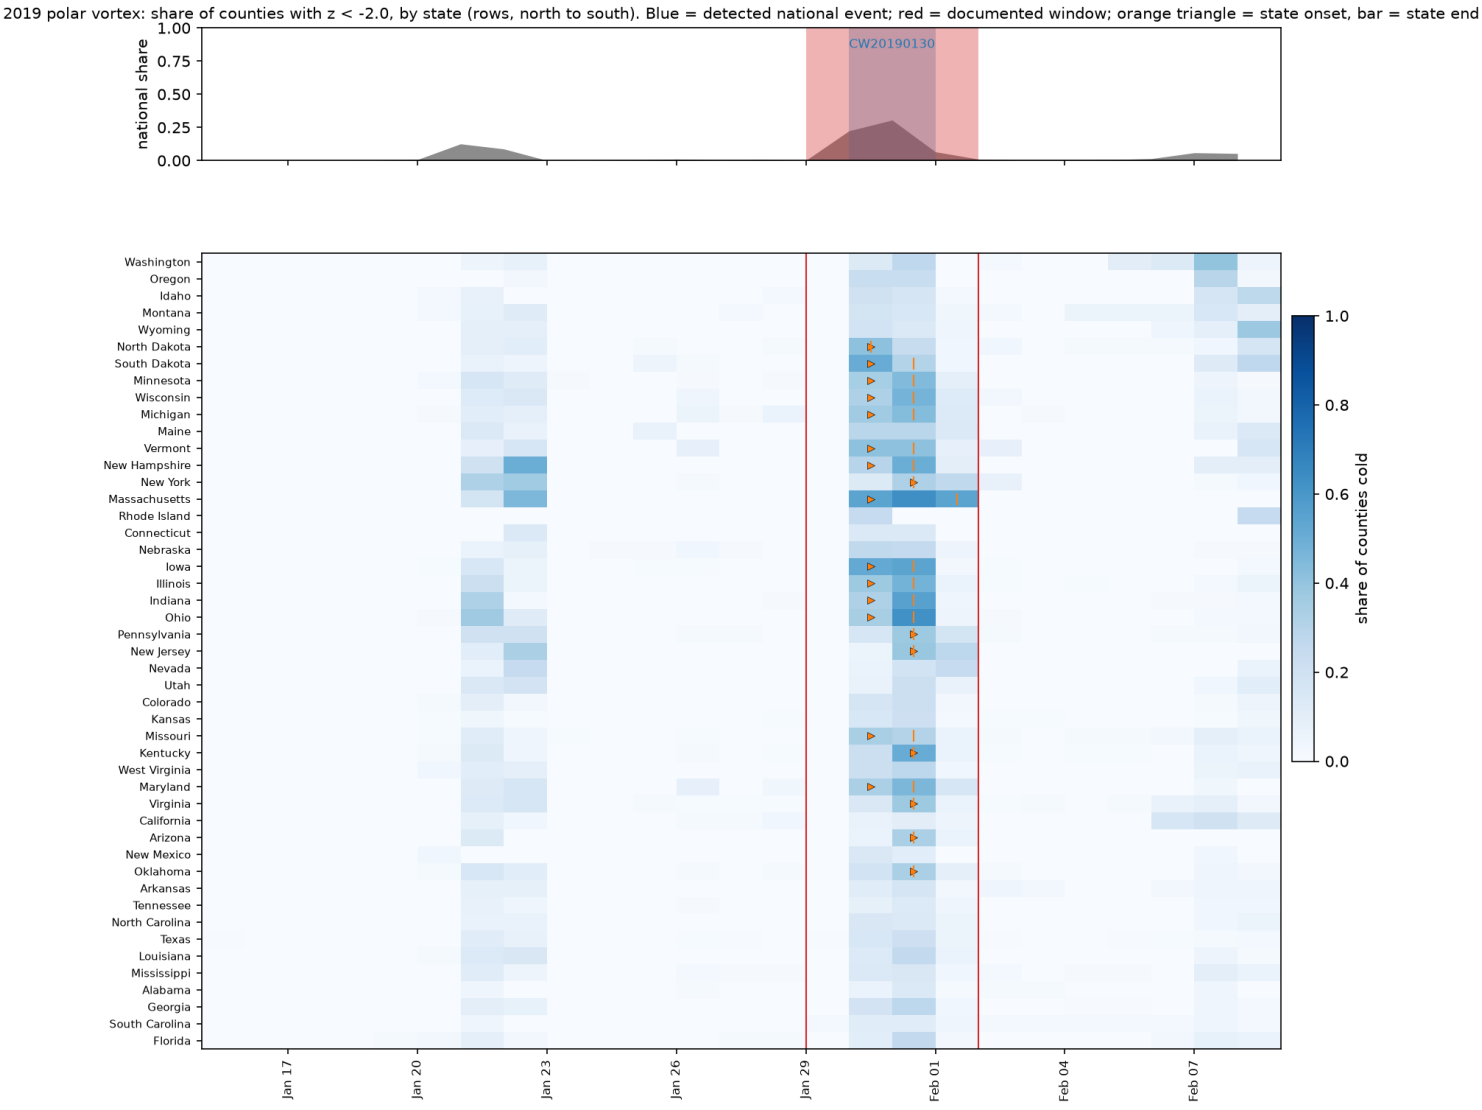



**Figure F.6:** The January 2019 polar vortex. Same layout as Figure F.3.

## F.5 Three types of cold wave events

The event-study design exploits the fact that a cold wave reaches different markets on different days, so that a market hit later can observe the cold wave in markets hit earlier. How much of this variation an event offers depends on how it arrives, and we classify every event by three statistics. The first is the number of national pulses inside the event: a pulse begins when the national share of cold counties reaches 20 percent and ends when it falls below 10 percent. The second is the spread of state onsets. A state's lag is the number of days between the national event start and the state's own onset; negative lags denote precursor cold, typically in the northern Plains or the Mountain West, that begins before the national share reaches 20 percent. The spread is the difference between the 90th and the 10th percentile of the lags across the states attached to the event, so that a single early or late state does not drive it. The third is the precursor, the number of days by which the earliest state onset precedes the national start; it captures a lead carried by only one or two states, which the spread deliberately ignores. The three types are as follows.

**Near-synchronous fronts.** One national pulse, an onset spread of at most two days and a precursor of at most one day. Fourteen of the 34 events are of this type: an arctic front crosses the country in a day or two, no state is cold more than a day before the national start, and all but a handful of states are cold within two days of it. Elliott, the January 2009 and January 2019 outbreaks, the November 2019 outbreak and the March 2019 pulse all belong to this group. For such events the anticipation window is informative only through forecasts, since no market observes the cold wave in another market first.

**Staggered rollouts.** One national pulse and either an onset spread of three days or more or a precursor of two days or more. Fifteen events are of this type, and it is these that support identification from cross-market timing, because some market is cold before others are hit. Ten of them roll across the country over several days. The November 2014 outbreak is the clearest: it enters through Montana on November 11, covers the Plains on November 12 and Texas on November 13, and reaches the Great Lakes and the East Coast only on November 15–18, so that Illinois is cold six days after the national start. The December 2006 event has the widest spread, nine days, with Kansas and Illinois cold two to three days before the national start and Louisiana five days after it. In December 2016 Massachusetts, North Dakota and Ohio are cold two to four days before the national start and Texas a day after it, and in December 2013 Illinois and Massachusetts enter five days after the start. The other five are staggered through a precursor carried by a few states, and Uri is the leading example: Montana is cold from February 11 and Kansas from February 12, two days and

one day before the national start, while 30 of the 33 attached states have onsets on February 13–15. Texas, which enters on February 14, could therefore observe the cold in the northern Plains during the last three days of its anticipation window even though the onset spread is only one day. December 2009, November 2013, the first January 2014 event and March 2014 are staggered for the same reason, with a precursor of two to three days in a few northern or northeastern states.

**Multi-pulse events.** More than one national pulse, whatever the onset spread. Five events are of this type: February 2015 (four pulses), the New Year 2018 cold wave (three), and December 2008, February 2011 and the second January 2014 event (two each). In these events a market's anticipation window may overlap an earlier pulse rather than an earlier stage of the same front. In February 2015 Massachusetts and New York are cold five and three days before the national event begins, while Texas enters five days after it because it is spared the first pulses. The New Year 2018 cold wave is also strongly staggered: it reaches the Midwest and Northeast on December 26–27 but Texas, Louisiana and the Deep South only on January 1–2, five to six days after the national start. In December 2008 a northern pulse on December 14–16, in which Montana, the Dakotas, Wyoming and the central Plains are cold, starts the event, and a second pulse covers the whole South and East at once on December 22–23.

Table F.3 lists the 19 events detected in 2014–2023, the period covered by the capacity release data, with their type, and marks the calibration anchors in the last column. We call an event *major* when its national peak share, the largest share of contiguous-U.S. counties that are cold on any day of the event, is at least 40 percent, and *minor* otherwise. Ten of the 19 events are major; they include the February 2015 Siberian Express, Winter Storm Uri and Winter Storm Elliott, whereas the January 2019 polar vortex, which was intense but confined to the Upper Midwest and the Northeast, peaks at 31 percent and is minor. The empirical analysis uses only the major events.

**Table F.3:** Detected cold wave events, 2014–2023

| Event | Start | End | Days | States | Peak day | Peak share | Severity | Type | Calibration anchor |
|---|---|---|---|---|---|---|---|---|---|
| CW20140106 | 2014-01-06 | 2014-01-07 | 2 | 42 | 2014-01-07 | 0.44 | major | staggered | |
| CW20140124 | 2014-01-24 | 2014-01-30 | 7 | 40 | 2014-01-24 | 0.40 | major | multi-pulse | |
| CW20140302 | 2014-03-02 | 2014-03-04 | 3 | 38 | 2014-03-03 | 0.44 | major | staggered | |
| CW20140326 | 2014-03-26 | 2014-03-26 | 1 | 12 | 2014-03-26 | 0.24 | minor | near-synchronous | |
| CW20141112 | 2014-11-12 | 2014-11-19 | 8 | 47 | 2014-11-18 | 0.58 | major | staggered | |
| CW20150108 | 2015-01-08 | 2015-01-08 | 1 | 22 | 2015-01-08 | 0.34 | minor | near-synchronous | |
| CW20150219 | 2015-02-19 | 2015-03-06 | 16 | 40 | 2015-03-06 | 0.43 | major | multi-pulse | Siberian Express |
| CW20161218 | 2016-12-18 | 2016-12-19 | 2 | 38 | 2016-12-19 | 0.43 | major | staggered | |
| CW20170316 | 2017-03-16 | 2017-03-16 | 1 | 14 | 2017-03-16 | 0.23 | minor | near-synchronous | |
| CW20171227 | 2017-12-27 | 2018-01-07 | 12 | 47 | 2018-01-02 | 0.67 | major | multi-pulse | |
| CW20180117 | 2018-01-17 | 2018-01-18 | 2 | 16 | 2018-01-17 | 0.28 | minor | near-synchronous | |
| CW20190130 | 2019-01-30 | 2019-01-31 | 2 | 21 | 2019-01-31 | 0.31 | minor | near-synchronous | polar vortex |
| CW20190304 | 2019-03-04 | 2019-03-06 | 3 | 37 | 2019-03-04 | 0.37 | minor | near-synchronous | |
| CW20191101 | 2019-11-01 | 2019-11-01 | 1 | 10 | 2019-11-01 | 0.24 | minor | near-synchronous | |
| CW20191112 | 2019-11-12 | 2019-11-14 | 3 | 44 | 2019-11-13 | 0.58 | major | near-synchronous | |
| CW20210213 | 2021-02-13 | 2021-02-19 | 7 | 33 | 2021-02-15 | 0.43 | major | staggered | Winter Storm Uri |
| CW20221223 | 2022-12-23 | 2022-12-26 | 4 | 45 | 2022-12-24 | 0.60 | major | near-synchronous | Winter Storm Elliott |
| CW20230319 | 2023-03-19 | 2023-03-20 | 2 | 10 | 2023-03-19 | 0.21 | minor | near-synchronous | |
| CW20231101 | 2023-11-01 | 2023-11-02 | 2 | 28 | 2023-11-01 | 0.34 | minor | near-synchronous | |

*Notes:* Events are detected with the parameters of Section F.2. Peak share is the largest national share of cold counties on any day of the event. An event is major when its peak share is at least 0.40 and minor otherwise; the empirical analysis uses the major events only.

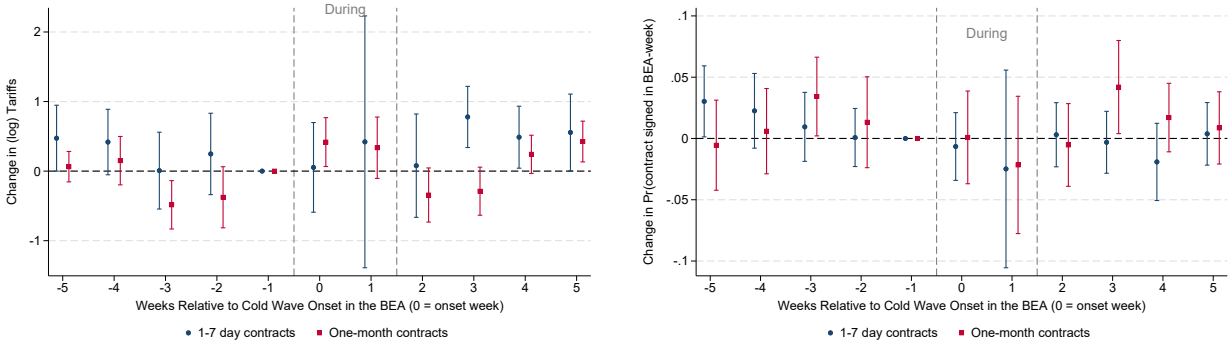
# G Robustness Analysis

## G.1 Robustness analysis using minor cold waves

Section 7 focuses on the ten major cold waves of 2014–2023, defined by a peak national share of cold counties of at least 40% (Table F.3). Here, we repeat the event-study analysis of Figure 13 for the nine minor cold waves, with peak shares between 20% and 40%. We use the same specification as in Section 7.1, retaining the major-event indicators as controls to isolate the effects of minor events. Minor cold waves cover 662 cold BEA-weeks, compared with 1,768 for major events. Because short-term contracts are infrequent, only 53 of these

cold BEA-weeks have a 1–7-day contract, limiting the precision of the tariff estimates.

Figure G.1 presents the results. We find no significant response in the tariff or signing probability of 1–7-day contracts, although the tariff estimates are imprecise, with confidence intervals spanning roughly one log point in either direction. One-month contracts also show no systematic response in tariffs or signing probabilities, despite some week-to-week fluctuations. These findings suggest that the sharp increase in short-term tariffs and decline in short-term releases documented in Section 7.1 are concentrated in major cold waves.




(a) Log tariffs

(b) Probability that a contract is signed

**Figure G.1:** Minor Cold Waves: 1–7 Day versus One-Month Contracts, Week by Week

*Notes:* As Figure 13, for the nine minor cold waves of 2014–2023. The event-week indicators of the major cold waves are included in the regression and not shown. Coefficients are relative to the week before the BEA's onset, with 95 percent confidence intervals; week 0 is the onset week, week 1 the later cold weeks and weeks 2 to 5 the four weeks after the BEA's end. All regressions include week, BEA, BEA × calendar-month and BEA × week-of-month fixed effects; standard errors are clustered by BEA.

# Appendix References


**American Gas Association**, "The Obligation to Serve: A Key Responsibility for Natural Gas Utilities," American Gas Association 2024. https://www.aga.org/the-obligation-to-serve-a-key-responsibility-for-natural-gas-utilities/.

**Borenstein, Severin, Meghan R. Busse, and Ryan Kellogg**, "Career Concerns, Inaction and Market Inefficiency: Evidence From Utility Regulation," *The Journal of Industrial Economics*, 2012, *60* (2), 220–248.

**Federal Energy Regulatory Commission**, "FERC Staff Inquiry Finds No Withholding of Pipeline Capacity in New England Markets," News release, February 27 2018. https://www.ferc.gov/news-events/news/ferc-staff-inquiry-finds-no-withholding-pipeline-capacity-new-england-markets.

**FERC and NERC Staff**, "The February 2021 Cold Weather Outages in Texas and the South Central United States," Technical Report, Federal Energy Regulatory Commission and North American Electric Reliability Corporation 2021.

**Form Energy**, "Multi-Day Energy Storage and Winter Reliability in New England," Technical Report, Form Energy September 2023. https://formenergy.com/wp-content/uploads/2023/09/Form-ISO-New-England-whitepaper-09.27.23.pdf.

**Illinois Commerce Commission**, "General Information Requirements Applicable for Gas Utilities, 83 Ill. Admin. Code §285.315," 2014. https://www.law.cornell.edu/regulations/illinois/Ill-Admin-Code-tit-83-SS-285.315.

**ISO New England and Midcontinent Independent System Operator and PJM Interconnection and Southwest Power Pool**, "Strategies for Enhanced Gas-Electric Coordination," Technical Report, Joint position paper February 2024. https://www.pjm.com/-/media/DotCom/library/reports-notices/special-reports/2024/20240221-strategies-for-enhanced-gas-electric-coordination-paper.ashx.

**Kansas Gas Service**, "General Terms and Conditions for Gas Service: Pipeline System Restrictions and Priorities, Schedule GTC11," Tariff on file with the State Corporation Commission of Kansas 2023. https://www.kansasgasservice.com/media/KGS/Tariffs/GTC11.pdf.

**Kansas Legislature**, "Utility Financing and Securitization Act, Senate Substitute for House Bill 2072," 2021. https://www.kslegislature.gov/li_2022/b2021_22/measures/documents/summary_hb_2072_2021.pdf.

**London Economics International LLC**, "New Jersey Natural Gas Capacity Study: Final Report (Public Redacted Version), Docket No. GO19070846," Technical Report, Prepared for the New Jersey Board of Public Utilities November 2021.

**Marks, Levi, Charles F. Mason, Kristina Mohlin, and Matthew Zaragoza-Watkins**, "Vertical Market Power in Interconnected Natural Gas and Electricity Markets," Working Paper 17-27, Resources for the Future 2017. https://media.rff.org/archive/files/document/file/RFF%20WP%2017-27.pdf.

**MIT Energy Initiative**, "Growing Concerns, Possible Solutions: The Interdependency of Natural Gas and Electricity Systems," Technical Report, Massachusetts Institute of Technology 2013.

**National Rural Electric Cooperative Association**, "Testimony of East Kentucky Power Cooperative, FERC Annual Reliability Technical Conference," Federal Energy Regulatory Commission, November 2023. https://www.electric.coop/wp-content/uploads/2023/11/FERC-Reliability-Conference-EKPC-Testimony-2023.pdf.

**New Jersey Board of Public Utilities**, "Decision and Order, Docket Nos. GO19070846 and GO20010033: Gas Capacity and Best Practices and Playbook for Emergencies," June 29 2022.

**New York Independent System Operator**, "2023–2024 Winter Assessment and Winter Preparedness," Presentation to the Operating Committee, November 14 2023. https://www.nyiso.com/documents/20142/41164815/04_2023-2024%20Winter%20Assessment%20and%20Winter%20Preparedness%20OC%20Presentation%20-%20Final%20111423.pdf.

**New York Public Service Commission**, "Order Instituting Proceeding, Case 20-G-0131, Proceeding on Motion of the Commission in Regard to Gas Planning Procedures," March 19 2020.

__ , "Order Adopting Gas System Planning Process, Case 20-G-0131," May 12 2022.

**Oklahoma Legislature**, "Senate Bill 1050: Regulated Utility Consumer Protection Act," 2021. https://www.oklegislature.gov/cf_pdf/2021-22%20ENR/SB/SB1050%20ENR.PDF.

**Railroad Commission of Texas**, "New 16 TAC §3.66, Relating to Weather Emergency Preparedness Standards: Adoption Preamble," August 30 2022. https://www.rrc.texas.gov/media/c5hdc4ga/rule-3-66.pdf.

**Texas Alliance of Energy Producers**, "Comments on Proposed Rules 3.65 and 3.107, Critical Designation and Curtailment Priorities," Submitted to the Railroad Commission of Texas, October 2021. https://www.rrc.texas.gov/media/radnflpe/3-65-3-107-comments-taep.pdf.

**Texas Legislature**, "Senate Bill 3, 87th Regular Session: Relating to Preparing for, Preventing, and Responding to Weather Emergencies and Power Outages; and House Bill 1520: Customer Rate Relief Bonds," 2021.

**U.S. Energy Information Administration**, "Electric Generator Dispatch Depends on System Demand and the Relative Cost of Operation," Today in Energy, August 17 2012. https://www.eia.gov/todayinenergy/detail.php?id=7590.

__ , "Residential Energy Consumption Survey (RECS), 2020: Household Energy Characteristics, Main Space Heating Fuel," 2022. https://www.eia.gov/consumption/residential/data/2020/.

__ , "Changes in U.S. Residential Natural Gas Prices Lag Spot Prices," Today in Energy 2023. https://www.eia.gov/todayinenergy/detail.php?id=60421.

__ , "Form EIA-176: Annual Report of Natural and Supplemental Gas Supply and Disposition, Natural Gas Annual Respondent Query System," 2023. https://www.eia.gov/naturalgas/ngqs/.

__ , "Units and Calculators Explained: Degree Days," 2023. https://www.eia.gov/energyexplained/units-and-calculators/degree-days.php.

__ , "Natural Gas Weekly Update, week of January 17, 2024," 2024. https://www.eia.gov/naturalgas/weekly/.