

The Allocative Cost of Price Ceilings in the U.S. Residential Market for Natural Gas

Lucas W. Davis

University of California, Berkeley and National Bureau of Economic Research

Lutz Kilian

University of Michigan and Centre for Economic Policy Research

A direct consequence of restricting the price of a good for which secondary markets do not exist is that, in the presence of excess demand, the good will not be allocated to the buyers who value it the most. We demonstrate the empirical importance of this allocative cost for the U.S. residential market for natural gas, which was subject to price ceilings during 1954–89. Using a household-level, discrete-continuous model of natural gas demand, we estimate that the allocative cost in this market averaged \$3.6 billion annually, nearly tripling previous estimates of the net welfare loss to U.S. consumers.

I. Introduction

A large literature in economics has examined the welfare costs of price ceilings. Among the markets that have received the most attention are rental housing, telecommunications, insurance, energy, and health care.¹ In traditional welfare analysis, price ceilings reduce the quantity transacted below the competitive level, imposing deadweight losses on

Comments from two referees and the editor as well as numerous seminar participants substantially improved the paper.

¹ See, e.g., Hayek (1931), Olsen (1972), Smith and Phelps (1978), Raymon (1983), Frech and Lee (1987), Deacon and Sonstelie (1989, 1991), Suen (1989), and Frech (2000). A closely related literature studies the welfare costs of minimum wage legislation in labor markets. See, e.g., Holzer, Katz, and Krueger (1991) and Card and Krueger (1994).

[*Journal of Political Economy*, 2011, vol. 119, no. 2]

© 2011 by The University of Chicago. All rights reserved. 0022-3808/2011/11902-0002\$10.00

both buyers and sellers. In this paper we concentrate on an additional component of welfare loss that is often ignored. Notably, when there is excess demand for a good for which secondary markets do not exist, a welfare loss occurs when the good is not allocated to the buyers who value it the most. This allocative cost has been studied, for example, by Glaeser and Luttmer (2003), but its practical importance has yet to be established empirically.²

Our analysis focuses on the U.S. residential market for natural gas. Anecdotal evidence suggests that price ceilings imposed in this market during the postwar period led to severe misallocation of natural gas across households. While some households were enjoying access to cheap price-controlled natural gas, other households were locked out of the market altogether as new residential connections were unavailable in many parts of the country.³ Many of the households without access to natural gas may well have been willing to pay more than households with access to natural gas, but there was no mechanism that allowed welfare-improving reallocations to occur. The allocative cost of restricting the price of natural gas refers to the total increase in welfare that could have been obtained by reallocating natural gas to the households with the highest willingness to pay.

The natural gas market is a good candidate for an empirical study of allocative costs for several reasons. First, natural gas is a homogeneous good, eliminating concerns about differences in quality that complicate the estimation of allocative costs in other markets. Second, whereas secondary markets may act to mitigate the costs of misallocation in some markets such as rental housing, there are no resale markets for natural gas. Third, the residential market for natural gas affects millions of consumers, suggesting that allocative costs could be very large. Fourth, this market was continuously regulated between 1954 and 1989 before experiencing complete deregulation. This allows us to observe market behavior both under regulation and in the absence of regulation. Fifth, the fact that some states remained unregulated throughout this period allows us to evaluate the out-of-sample fit of our model in settings in which markets operate freely. Sixth, this market lends itself to empirical analysis, given the availability of unusually comprehensive household-

² The problem of allocative costs is aptly described in Friedman and Stigler (1946). An early theoretical treatment can be found in Weitzman (1977). The analogous problem of misallocation due to minimum wages has been discussed as early as Welch (1974) and has been studied in Luttmer (2007).

³ As described in American Gas Association (1975, 67), the gas shortage caused widespread restrictions on new residential customers and severely limited expansion into new residential customer markets by many utilities. See also MacAvoy and Pindyck (1975, 2), American Gas Association (1976, 125), and Herbert (1992, 127). Related discussion can be found in MacAvoy (1971, 1983), Breyer and MacAvoy (1973), and MacAvoy and Pindyck (1973).

level data by state and year as well as the corresponding state-level price data.

We construct estimates of the allocative costs associated with the regulation of natural gas prices by exploiting the fact that by the 1990s, the natural gas market had been completely deregulated and, in contrast to the period of regulation, all households purchasing new homes were free to choose natural gas heating systems. Our empirical strategy is to ask how much natural gas would have been consumed in 1950–2000 based on the household preferences revealed in the 1990s data. We first estimate household preferences based on heating system choices made by households that purchased new homes in the 1990s after the end of price regulation. Controlling for household demographics and housing characteristics that affect heating demand, we then compare households' actual choices under regulation with what they would have liked to choose in an unconstrained world, as implied by an economic model of consumer choice. We can therefore calculate physical shortages of natural gas and measure the allocative cost of price ceilings.

Our paper provides for the first time a detailed picture of the evolution of physical shortages in the U.S. natural gas market during the postwar period. Whereas previous studies have traditionally measured the degree of disequilibrium in the natural gas market using shortfalls in contractually obligated deliveries to pipelines, our measure of the physical shortage correctly incorporates not only demand from existing delivery contracts but the unrealized demand from prospective new customers as well.⁴ This distinction is particularly important in the residential market because shortages were accommodated by restricting access to potential new customers rather than by rationing existing users. We find that during the period 1950–2000, demand for natural gas exceeded sales of natural gas by an average of 19.4 percent, with the largest shortages during the 1970s and 1980s. In comparison to previous studies, we find that the shortages began earlier, lasted longer, and were larger in magnitude.

Physical shortages are important in describing the effect of price ceilings but do not provide a measure of their economic costs. Using a household-level, discrete-continuous model of natural gas demand following Dubin and McFadden (1984), we estimate that the allocative cost from price ceilings averaged \$3.6 billion annually in the U.S. residential

⁴ For example, Vietor (1984) reports that shortfalls in contractually obligated deliveries to pipelines increased steadily beginning in 1970, reaching approximately 3 trillion cubic feet in 1976. This is a significant amount considering that total natural gas consumption in the United States in that year was 20 trillion. As large as these curtailments were, results from our model show that they understate the true level of disequilibrium in the market because they fail to account for demand from prospective new customers.

market during 1950–2000.⁵ This figure represents, on average, 14 percent of total residential expenditure on natural gas. As a point of comparison, MacAvoy (2000) estimates that at the national level between 1968 and 1977 natural gas price ceilings transferred an average of \$6.9 billion annually from producers to consumers while causing consumers a deadweight loss of \$9.3 billion. Thus, even abstracting from allocative cost, price ceilings made consumers worse off by \$2.4 billion annually. We find that during this same period, the mean annual allocative cost in the residential market was \$4.8 billion (which is somewhat higher than the mean estimate for the full period). Incorporating the allocative cost nearly triples the estimated net welfare loss to U.S. consumers. In addition to providing estimates of the allocative cost at the national level, our household-level approach provides insights into the distributional effects of regulation that could not have been obtained using a model based on national or even regional data. In particular, we show that the allocative cost of regulation was borne disproportionately by households in the Northeast and Midwest.⁶

Our analysis has several policy implications. First, regulators need to be aware that price ceilings benefit only consumers who have access to regulated markets. When there is a shortage of a good, not all consumers will have access to the market, and those who have access will not necessarily be the consumers who value the good the most. Second, the adverse effects of price ceilings can last much longer than the regulatory policies themselves. With natural gas, since households change heating systems infrequently, households that are barred from adopting natural gas heating systems because of a price ceiling will continue to use inferior technologies for years to come. This lock-in effect helps explain the persistence and the magnitude of the allocative costs that we find and highlights the difficulty of predicting the duration of the effects of price regulation. Third, our analysis underscores the difficulty of determining in advance how the allocative cost of price regulation will be distributed geographically.

The paper is formatted as follows. Section II demonstrates the existence of an allocative cost from price ceilings in addition to the conventional deadweight welfare loss for goods for which there is no secondary market. Section III briefly reviews the history of regulation in

⁵ All dollar amounts are expressed in year 2000 dollars.

⁶ Our analysis is germane to a substantial literature that examines regulation in the U.S. natural gas industry. Early studies such as MacAvoy (1971), Breyer and MacAvoy (1973), and MacAvoy and Pindyck (1973, 1975) document gas shortages in the early 1970s and use structural dynamic simultaneous equation models to simulate hypothetical paths for prices, production, and reserves under alternative regulatory regimes. Several of these studies present estimates of the deadweight loss from natural gas price ceilings, but only Braeutigam and Hubbard (1986) and Viscusi, Harrington, and Vernon (2005) discuss the issue of allocative cost.

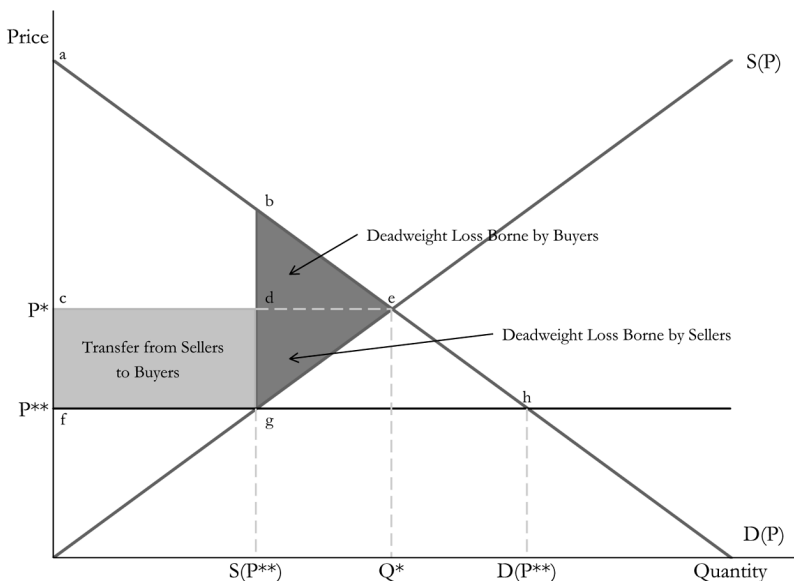


FIG. 1.—Conventional welfare analysis of the effect of price ceilings

the U.S. natural gas industry since the 1950s, emphasizing characteristics of the regulatory policies that are relevant to our analysis. Sections IV and V introduce our household-level model of demand for natural gas and discuss its empirical implementation. In Section VI, we discuss the estimates of physical shortages as well as allocative cost and the out-of-sample fit of the model. Section VII contains concluding remarks.

II. Price Ceilings and Allocative Cost

Figure 1 describes the standard problem of imposing a price ceiling. At the competitive equilibrium, the market clears at price P^* and quantity Q^* . Now consider the effect of a price ceiling P^{**} . At this price level, demand $D(P^{**})$ exceeds supply $S(P^{**})$. Compared to the competitive equilibrium, households pay $P^* - P^{**}$ less per unit, resulting in a transfer $cdgf$ from sellers to buyers. They also incur the deadweight loss given by the triangle bed because of the decrease in quantity. Firms are unambiguously worse off. They pay a transfer of $cdgf$ and incur a deadweight loss of deg . The total deadweight loss is given by the area beg .

The welfare cost of price ceilings, however, is not necessarily limited to the deadweight loss. On further inspection it becomes clear that welfare losses will be limited to the deadweight loss triangle beg if and

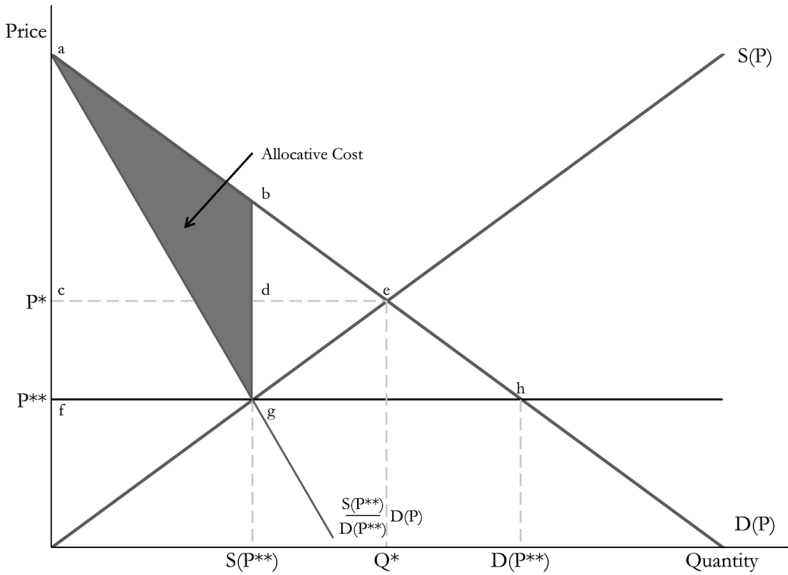


FIG. 2.—Allocative cost under random assignment

only if the good is allocated to the buyers who value it the most. In this efficient rationing case, buyers represented by the demand curve between *a* and *b* receive the good, whereas those represented by the demand curve between *b* and *h* do not. In some markets it may be reasonable to assume that a good is allocated efficiently. For example, when there is a secondary market in which goods can be resold, this secondary market ensures that buyers with the highest willingness to pay receive the good. However, in many other markets such as the market for natural gas, there is no mechanism that ensures that customers with the highest reservation price will receive the good. In these markets the welfare costs of price regulation also depend on how the good is allocated. Inefficient rationing imposes additional welfare costs.

A commonly used benchmark in illustrating these additional welfare costs is the case in which goods are allocated randomly to buyers (see fig. 2).⁷ The random allocation is inefficient because it does not allocate goods to buyers with the highest willingness to pay. At the price ceiling P^{**} , demand for the good is $D(P^{**})$, but supply is only $S(P^{**})$. If supply is allocated randomly, then only a fraction $S(P^{**})/D(P^{**})$ of buyers with a reservation price above P^{**} will be able to buy the good. This random allocation corresponds to the curve $D(P)S(P^{**})/D(P^{**})$. Now, in addi-

⁷ This analysis follows closely those of Braeutigam and Hubbard (1986), Glaeser and Luttmer (2003), and Viscusi et al. (2005).

tion to the deadweight loss, beg , there is an additional welfare loss, abg , that is the result of the loss of efficiency from not allocating the good to the consumers with the highest reservation price. This additional welfare loss represents the allocative cost of regulation in this example. In general, the allocative cost can be measured as the difference between the consumer surplus under efficient rationing and the consumer surplus under the actual rationing. If the good is allocated to the buyers at random as in figure 2, the allocative cost corresponds to the difference between the consumer surplus $abgf$ and the consumer surplus agf .

In practice, the level of the allocative cost depends not only on how the good is rationed but also on the distribution of reservation prices across households. In the market for natural gas, heterogeneity in reservation prices arises mainly for two reasons. First, there are differences across households in preferences for different types of heating systems. For example, households differ in how much they value the cleanliness and convenience of natural gas. Second, households differ in how much they value different heating systems because of technological considerations. Compared to electric heating systems, natural gas and oil heating systems are expensive to purchase but inexpensive to operate. As a result, households with high levels of demand for home heating tend to prefer natural gas or heating oil. This suggests that in estimating the allocative cost, it is important to specify a model that accounts for the heterogeneity of households.

The allocative cost depends on the location and shape of the demand curve and the observed level of price and quantity, but not on the shape of the supply curve. Accordingly, our analysis abstracts from the supply side of the natural gas market. A limitation of our approach is that without an empirical representation of the supply curve, we cannot calculate a measure of the conventional deadweight loss. For estimates of the deadweight loss, see MacAvoy and Pindyck (1975) and MacAvoy (2000). With our model we are able to simulate demand for natural gas at the prices actually observed in the market during this period and to calculate shortages, but we are not able to say what equilibrium price levels would have prevailed without price ceilings or under alternative forms of regulation. The latter question indeed is of no relevance for the measurement of the allocative cost.

III. History of Natural Gas Regulation in the United States

Most natural gas in the United States is produced in gas fields concentrated in Texas, Louisiana, New Mexico, and Oklahoma and transported by pipeline to gas consumers in the Midwest and Northeast. The traditional problem in the U.S. natural gas industry in the 1930s and 1940s had been one of overproduction. This situation changed in the 1950s.

As the domestic pipeline system expanded, supply could barely keep up with rising demand for natural gas among urban consumers in the Midwest and Northeast. Prices increased, much to the dismay of consumer advocates, and pressures arose to regulate interstate sales of natural gas, culminating in a 1954 Supreme Court decision that imposed federal price controls on natural gas sold in the interstate market.⁸ Because federal jurisdiction extended only to interstate sales of natural gas, natural gas markets in gas-producing states were left unregulated.

The price ceilings stimulated consumer demand while at the same time discouraging supply. The natural consequence was a growing shortage of natural gas. Starting in the 1950s, utility companies attempted to meet this shortage by denying access to new residential customers rather than rationing existing residential customers. The shortage became widely apparent in the 1970s, when utility companies were forced to curtail deliveries to the industrial sector. Evidence of increasing physical shortages finally brought the issue of natural gas regulation to the attention of the legislature and led to the passage of the Natural Gas Policy Act in 1978, which specified a phased deregulation of the price of natural gas. The 1978 act was a political compromise intended to reduce shortages without completely eliminating the distortions of the old pricing system. After 1978, natural gas prices temporarily spiked, supplies expanded, and curtailments were eliminated, but it was only in 1989 that all forms of price controls were officially terminated (see Bradley 1996).

Thus, residential natural gas prices in the United States were subject to price ceilings during a 35-year period from 1954 to 1989. The purpose of this paper is to quantify the allocative costs associated with these price ceilings. We exploit the fact that by the 1990s natural gas was widely available in the United States, and, in contrast to previous decades, all households purchasing new homes were free to choose natural gas heating systems. This fact is reflected in the distinct evolution of the percentage of new homes with natural gas heating in gas-importing states such as New Jersey compared with gas-exporting states such as Texas (see table 1). A priori, one would have expected a much higher share of households to demand natural gas in New Jersey with its colder climate than in Texas, yet during the height of regulation, the share of natural gas was near 40 percent in both states. With the onset of deregulation, access to natural gas in New Jersey increased and the percentage of new homes with natural gas heating increased to 63 percent,

⁸ The *Phillips Petroleum Co. v. State of Wisconsin* decision was made in spite of the fact that by all accounts U.S. natural gas production is a highly competitive industry. In 1953, the 30 largest natural gas producers controlled less than half of all reserves and accounted for only one-third of sales to interstate pipelines (Vieter 1984). See also Lindahl (1956), Cookenboo (1958), and Neuner (1960).

TABLE 1
 PERCENTAGE OF NEW HOMES WITH NATURAL GAS HEATING BY STATE

	During Height of Regulation (1969–78)	During Phased Deregulation (1979–89)	After Deregulation (1990–2000)
New Jersey	40	63	80
Texas	43	31	33

NOTE.—These percentages were calculated using data from homes constructed during the years indicated in the column headings.

whereas the corresponding share in Texas fell to 31 percent. After regulation ended, the share of new homes with natural gas heating in New Jersey further increased to 80 percent, whereas the corresponding share in Texas remained near 30 percent.

Our empirical strategy is to ask how much natural gas would have been consumed during the period of regulation on the basis of choices made by households living in homes built during the 1990s. When one controls for the covariates that affect heating demand, these choices allow the estimation of household preferences. This strategy addresses one of the central difficulties in estimating the allocative cost of price ceilings. In particular, during periods of price regulation, one observes households' behavior only under the constraints imposed by regulation, making it difficult to identify households' unconstrained preferences. Our study sidesteps this difficulty by taking advantage of the fact that in the 1990s we observe unconstrained household behavior. Comparing households' actual choices with what they would have liked to choose in an unconstrained world, as implied by the econometric model discussed in Section IV, enables us to measure the allocative cost of price ceilings. In the following sections we develop this empirical strategy in more detail.

IV. Demand Model

Calculating the allocative cost of price ceilings in the U.S. residential natural gas market requires a model of household behavior that can accurately represent households' choices of heating systems and their usage of natural gas. We use the structure of this demand model to run a counterfactual analysis of the gains from having natural gas allocated to those who value it the most. Our approach requires making explicit assumptions about households' preferences and constraints. Below we describe this parametric framework, highlighting the key modeling choices.

Joint discrete-continuous models have been the standard in modeling energy demand at the household level since Hausman (1979), Dubin

and McFadden (1984), and Dubin (1985).⁹ These models recognize that energy demand is derived from demand for household services such as heating, cooling, and refrigeration that are produced in the home using durable goods. Our interest is in modeling the residential demand for natural gas in the context of a model that allows households to choose between natural gas, electricity, and heating oil. The demand for energy (natural gas, electricity, and heating oil) and the demand for heating equipment are modeled jointly as the solution to a household production problem. Households make two choices. Households decide how much energy to purchase for heating conditional on each of the three available types of heating systems. Households also decide which heating system to purchase taking into account that whatever heating system they choose to purchase will be operated at its optimal level of utilization.

Building on Dubin and McFadden (1984), we postulate that the conditional demand function for natural gas takes the linear form

$$x_i = \beta_0 + \beta_1 p_{ig} + \beta_2' w_i + \eta_i, \quad (1)$$

where x_i denotes annual consumption of natural gas for household i (measured in British thermal units, or Btu). Demand for natural gas depends on p_{ig} , the price of natural gas. The i subscript reflects the fact that households in different states face different natural gas prices. Demand for natural gas also depends on w_i , a vector of household and housing characteristics (including heating degree days [HDD], household size, household income, number of rooms, and the number of units in the building), and η_i , which reflects unobserved differences across households in the demand for heat. The parameter β_1 measures the responsiveness of demand to changes in price, and the parameter vector β_2 describes how demand for natural gas varies across households with different characteristics. As shown by Dubin and McFadden, Roy's identity implies a partial differential equation whose solution defines the indirect utility function that corresponds to the demand function (1).

Our model postulates that households choose which heating system to purchase by evaluating the indirect utility function

$$u_{ij} = \alpha_{0j} + \alpha_{1j} p_{ij} + \alpha_{2j}' w_i + \alpha_{3j}' p_{ij} z_i + \epsilon_{ij}, \quad (2)$$

where the utility for household i of heating system $j \in \{g, e, o\}$ is a function of p_{ij} , w_i , and z_i , a vector that includes HDDs, household size, and number of rooms. The parameter α_{0j} incorporates heating system-specific factors such as purchase and installation costs as well as pref-

⁹ Examples include Bernard, Bolduc, and Belanger (1996), Goldberg (1998), West (2004), Feng, Fullerton, and Gan (2005), Mansur, Mendelsohn, and Morrison (2008), and Bento et al. (2009).

erences for particular heating systems that are common across households. The parameter α_{1j} reflects households' willingness to trade off the price per unit of heat against other heating system characteristics, and the parameter vector α_{2j} describes interactions between household characteristics and heating system alternatives. This specification allows households living in cold climates to prefer natural gas heating systems, for example. The parameter vector α_{3j} allows the importance of different energy prices to vary across households with different levels of demand for heat. For example, households living in cold climates may be more sensitive to natural gas prices. The error term, ϵ_{ij} , captures unobserved differences across households' preferences for particular heating systems.¹⁰

Given equation (2), the probability that household i selects alternative k is the probability of drawing $\{\epsilon_{i1}, \epsilon_{i2}, \dots, \epsilon_{ij}\}$ such that $u_{ik} \geq u_{ij}$ for all $j \neq k$. As is customary, we assume that ϵ_{ij} is identically and independently distributed (i.i.d) across households and heating systems with a type I extreme value distribution, $\exp[-\exp(-\epsilon)]$. Under this assumption, the probability that household i selects heating system k takes the well-known conditional logit form

$$\Pr_{ik} = \frac{\exp(\alpha_{0k} + \alpha_{1k}p_{ik} + \alpha'_{2k}w_i + \alpha'_{3k}p_{ik}z_i)}{\sum_{j=1}^J \exp(\alpha_{0j} + \alpha_{1j}p_{ij} + \alpha'_{2j}w_i + \alpha'_{3j}p_{ij}z_i)}, \quad (3)$$

and the heating system choice model can be estimated using maximum likelihood.

Following Dubin and McFadden (1984), we allow the discrete and continuous components of the model to be correlated. This correlation might arise for many reasons. For example, households that prefer warm homes may be prone to choosing natural gas heating systems as well as prone to consuming more natural gas. As a result, the distribution of η_i among households that select natural gas may not be the same as the unconditional distribution of η_i . This endogeneity problem is addressed by postulating that the expected value of η_i is a linear function of $\{\epsilon_{i1}, \epsilon_{i2}, \dots, \epsilon_{ij}\}$ and using the density of the extreme value distribution to evaluate the conditional expectation of η_i analytically. Selection correction terms $\hat{\Pr}_{ie} \ln(\hat{\Pr}_{ie}) / (1 - \hat{\Pr}_{ie}) + \ln(\hat{\Pr}_{ig})$ and $\hat{\Pr}_{io} \ln(\hat{\Pr}_{io}) / (1 - \hat{\Pr}_{io}) + \ln(\hat{\Pr}_{ig})$ can be derived that are functions of the predicted choice probabilities from the heating system choice model. When these terms

¹⁰ The indirect utility function derived using Roy's identity includes an additional scale factor that complicates estimation. This term scales the observable component of utility by a factor that varies across households as a function of price. We follow the literature in subsuming this term into the error ϵ_{ij} when estimating the model. The same approach has been adopted by Dubin (1985), Goldberg (1998), and West (2004). Mannering and Winston (1985) motivate this procedure based on a Taylor series expansion about the mean price. For an alternative derivation of this procedure, see Train (1986).

are included in the heating demand function (1), the parameters β_0 , β_1 , and β_2 can be estimated consistently.

Our specification follows previous studies in assuming that current prices are a reasonable proxy for future prices. This assumption is natural when energy prices are well approximated by a random walk and changes in energy prices are unpredictable. In most contexts this will be a reasonable assumption, although a case could be made that during the late 1970s, when deregulation was imminent, it might have been reasonable to expect natural gas prices to increase. We are also implicitly assuming that observed choices reflect the preferences of the household and not preferences of some other party. In the case in which home builders or landlords are involved, we assume that the relevant market is sufficiently competitive so that these parties act effectively as agents for the household.

In estimating the model we take the stand that variation in energy prices is not driven by unmodeled fluctuations in the residential demand for heating. There are several reasons why standard concerns about price endogeneity are likely to be much less of a problem here than in other contexts. First, a significant advantage of the census data is that they allow for an unusually rich specification of demand. By flexibly describing the responsiveness of demand to climate, household, and housing characteristics, we control for many demand determinants that might be expected to be correlated with price. Second, the identification of the parameters in equations (1) and (2) is aided by institutional characteristics of U.S. retail energy markets. For example, residential natural gas prices are largely driven by cross-sectional differences in transportation and distribution costs. Prices are consistently higher in the Northeast than in the South, where most gas is produced. Transmission costs, of course, have decreased over time, but most analysts attribute these savings to regulatory reform and technological change rather than changes in residential demand (Leitzinger and Collette 2002). In addition, the fact that there is an integrated national wholesale market for natural gas means that local changes in demand tend not to result in large local price changes. Thus, there is little scope in general for unmodeled regional demand shocks to influence prices.

Similar arguments can be made for electricity and heating oil prices. First, cross-sectional variation in electricity prices is largely driven by differences in the composition of electricity generation. States with coal-burning power plants and federally subsidized hydroelectric facilities tend to have considerably lower prices than states with natural gas and nuclear power plants. Second, time-series data for electricity prices show sustained declines in prices over almost the entire period (see fig. 2). These price reductions are due to substantial improvements in the operating efficiency of power plants rather than changes in demand (see,

e.g., Fabrizio, Rose, and Wolfram 2007). Similarly, the cross-sectional variation in heating oil prices is largely driven by differences in transportation costs across states, whereas the time-series variation mainly reflects changes in global crude oil prices, especially in the late 1970s and early 1980s, which represent supply shocks from the point of view of heating oil producers.

V. Empirical Implementation

The estimation of the model is based on a data set that we compiled from industry sources, governmental records, and the U.S. Census for 1960–2000. The census provides a 40-year history of household heating system choices, household demographics, and housing characteristics. We use 1 percent samples for 1960 and 1970 and 5 percent samples for 1980, 1990, and 2000. We also put considerable effort into constructing a 50-year panel of state-level residential prices for natural gas, electricity, and heating oil. Figure 3 shows the residential prices by region. This data set, together with the census data, makes it possible to represent formally the alternatives available to households in the United States during this period.

Table 2 reports estimates of the heating system choice model. The model is estimated using households living in homes built between 1990 and 2000 after the complete deregulation of natural gas prices. The first set of parameters corresponds to energy prices and the interactions of energy prices with household characteristics. We allow the parameters for natural gas and heating oil prices to differ from the parameters for electricity prices because electric heating systems do not require combustion and thus are much more efficient in converting energy into heat. The coefficients for price are negative and strongly statistically significant, indicating that, all else equal, households prefer heating systems with a low energy price. The price interaction terms for natural gas and heating oil are large and statistically significant. The remaining parameters correspond to household characteristics interacted with indicator variables for natural gas and electricity. All estimates are normalized relative to heating oil. For example, the negative coefficient on rooms for electric heating systems indicates that electric heating systems become less attractive relative to heating oil as the number of rooms increases. The other coefficients may be interpreted similarly. The positive constants indicate that natural gas and electricity are more attractive to households than oil heating systems, perhaps reflecting that oil systems are not as clean burning as other heating systems, require more space, and are less convenient.

Table 3 presents estimates of the parameters of the heating demand function given in equation (1). Whereas the heating system choice

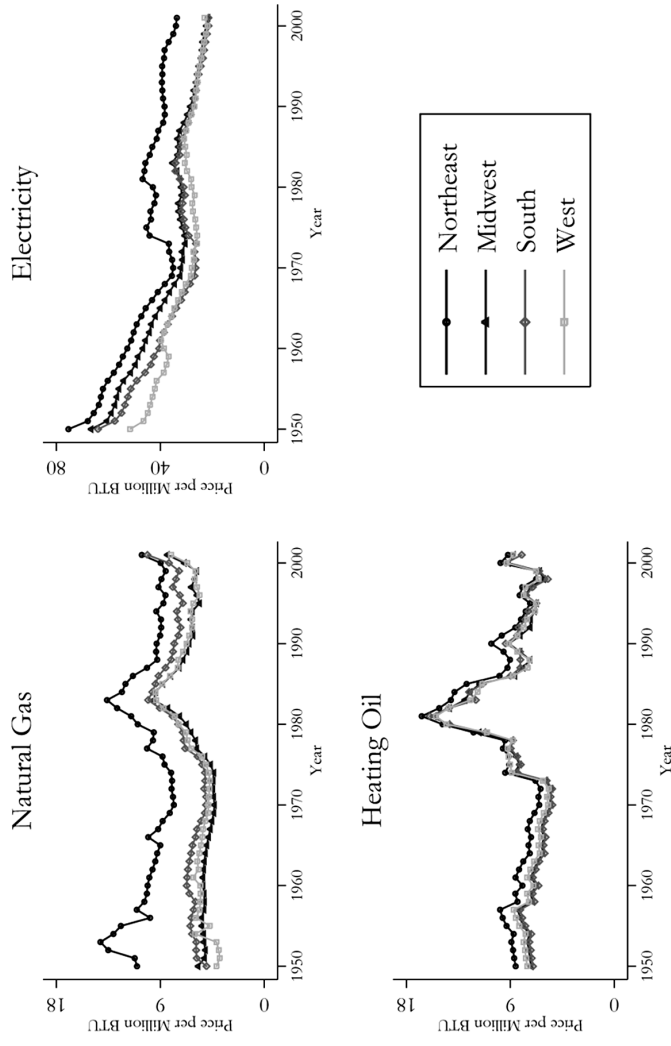


FIG. 3.—U.S. residential energy prices, by region. Sources: Natural gas, 1970–2000: Energy Information Administration; 1950–69: American Gas Association; electricity, 1970–2000: Energy Information Administration; 1950–69: Edison Electric Institute; heating oil, 1970–2000: Energy Information Administration; 1950–69: *Platt's Oil Price Handbook and Oilmanac*, various editions.

TABLE 2
ESTIMATES OF THE HEATING SYSTEM CHOICE MODEL

	Gas and Oil	Electric
Energy prices and interactions:		
Price (per Btu)	-.383 (.003)	-.119 (.0006)
Price and HDD (1,000s)	-.057 (.002)	.003 (.0004)
Price and number of rooms	-.020 (.001)	.004 (.0003)
Price and number of household members	-.020 (.002)	-.006 (.0004)
	Gas	Electric
HDD:		
HDD (1,000s)	-.99 (.03)	-1.93 (.04)
HDD ² (10,000,000s)	.37 (.03)	.50 (.03)
Demographic characteristics:		
Two household members	-.15 (.02)	-.03 (.03)
Three household members	-.32 (.02)	-.04 (.03)
Four household members	-.37 (.02)	-.08 (.04)
Five household members	-.42 (.03)	-.06 (.05)
Six or more members	-.39 (.04)	.06 (.07)
Total family income (10,000s)	.02 (.00)	-.02 (.00)
Homeowner dummy	.15 (.03)	-.13 (.03)
Housing characteristics:		
Number of rooms	.14 (.01)	-.31 (.01)
Building has 2 units	.85 (.06)	.71 (.06)
Building has 3-4 units	1.33 (.06)	1.59 (.06)
Building has 5-9 units	1.93 (.07)	2.39 (.07)
Building has 10-19 units	2.08 (.08)	2.66 (.08)
Building has 20-49 units	1.30 (.06)	1.95 (.06)
Building has 50+ units	1.05 (.06)	1.92 (.06)
Constant	5.73 (.09)	12.70 (.13)

NOTE.—Maximum likelihood estimates based on 621,478 households.

model is estimated using only households living in homes built during the 1990s, the heating demand function is estimated using all households in the 1980, 1990, and 2000 Census that use natural gas as their primary source of home heating. Our use of a longer sample exploits the fact that during the period of regulation, existing natural gas users were not rationed and could use as much natural gas as they wished. A potential concern with using this longer sample is selection bias. In estimating equation (1), we face two levels of selection. First, natural gas consumption is observed only for households that chose natural gas heating. As discussed in Section IV, we address this form of selection following Dubin and McFadden (1984) by including in these regressions a set of selection correction terms that are functions of the predicted choice probabilities from the heating system choice model. Second, not included in this sample are households that preferred natural gas but were excluded from the market because of shortages. To the extent that

TABLE 3
ESTIMATES OF THE HEATING DEMAND FUNCTION

	1980	1990	2000
Price of natural gas	-7.70 (.06)	-3.54 (.06)	-1.29 (.07)
HDD:			
HDD (1,000s)	46.5 (.24)	29.1 (.19)	26.9 (.18)
HDD ² (10,000,000s)	-31.9 (.23)	-18.2 (.20)	-21.3 (.21)
Demographic characteristics:			
Two household members	8.1 (.18)	4.9 (.15)	3.6 (.14)
Three household members	19.8 (.21)	12.7 (.19)	9.5 (.17)
Four household members	23.9 (.23)	15.4 (.20)	10.9 (.18)
Five household members	32.9 (.29)	22.4 (.27)	16.3 (.25)
Six or more members	50.2 (.36)	33.9 (.36)	23.0 (.33)
Total family income (10,000s)	.6 (.03)	1.5 (.02)	.9 (.01)
Homeowner	7.2 (.20)	2.2 (.19)	2.9 (.18)
Housing characteristics:			
Number of rooms	11.1 (.05)	8.6 (.05)	5.2 (.04)
Home built in 1940s	-10.2 (.23)	-8.7 (.23)	-6.2 (.24)
Home built in 1950s	-17.9 (.19)	-13.5 (.19)	-10.5 (.19)
Home built in 1960s	-21.6 (.19)	-15.9 (.19)	-12.5 (.20)
Home built in 1970s	-41.4 (.26)	-21.2 (.20)	-15.9 (.19)
Home built in 1980s	. . .	-26.6 (.20)	-19.2 (.20)
Home built in 1990s	-29.0 (.23)
Building has 2 units	14.7 (.32)	6.3 (.34)	7.6 (.33)
Building has 3-4 units	-10.5 (.37)	-10.5 (.38)	-3.0 (.34)
Building has 5-9 units	-27.7 (.40)	-27.2 (.37)	-14.5 (.33)
Building has 10-19 units	-32.6 (.42)	-32.5 (.36)	-18.2 (.35)
Building has 20-49 units	-41.6 (.51)	-42.4 (.43)	-28.1 (.38)
Building has 50+ units	-30.5 (.63)	-42.0 (.60)	-21.6 (.45)
Selection terms:			
Electricity selection term	-12.0 (.66)	-4.7 (.58)	40.8 (.59)
Heating oil selection term	13.7 (.89)	-1.93 (.87)	-43.8 (.80)
Constant	-7.7 (.77)	2.49 (.64)	23.2 (.83)
Number of households	1,719,743	1,882,971	2,177,998
R^2	.25	.24	.17

NOTE.—This table reports parameter estimates from three separate regressions based on U.S. Census household data, one for each decade. The dependent variable is annual natural gas consumption in millions of Btu. We instrument for the price of natural gas using regional indicator variables for each of the nine census regions. These instruments generate large first-stage F -statistics. Heteroskedasticity-robust standard errors are in parentheses.

the underlying rationing mechanism is correlated with unobserved components of heating demand, this could bias our estimates. We aim to minimize this bias by including a rich set of household and housing characteristics in the regressions.

We estimate equation (1) separately by decade. Estimating separate regressions makes it possible to capture changes in heating demand over time that are not captured by observable characteristics. It also helps us capture slow-moving changes in the energy efficiency of homes,

which is important in constructing the counterfactuals in Section VI.¹¹ The dependent variable in table 3 is annual consumption of natural gas in millions of Btu, constructed by dividing reported annual expenditure on natural gas by the average residential price of natural gas for the appropriate state and year.¹² When estimating equation (1), we instrument for the price of natural gas using regional indicator variables to address any spurious correlation between demand and price caused by measurement error in price. The parameter estimates in table 3 indicate decreasing price sensitivity over time. The point estimates imply price elasticities ranging from -0.34 in 1980 to -0.10 in 2000. As a point of comparison, Dubin and McFadden (1984) find a price elasticity for electricity demand among households with electric heating ranging from -0.20 to -0.31 . Table 3 also shows that natural gas consumption responds to household and housing characteristics with the expected signs. Temperature is one of the most important determinants of natural gas consumption, and the coefficients for HDDs are strongly statistically significant. Natural gas consumption increases with household size, the number of rooms, and the age of the home. Households in multiunit structures tend to use less natural gas than households in single-family residences, perhaps because of shared walls and other scale effects. Finally, five of the six selection terms are statistically different from zero, indicating that the unobserved determinants of heating demand and heating system choices are indeed correlated. Correcting for selection matters most in the 2000 sample, reflecting the fact that by the 1990s households faced no constraints in choosing their preferred heating system. The signs on the selection terms for 2000 are consistent with a natural ordering in which households with high levels of demand for home heating prefer heating oil to natural gas and natural gas to electricity.

Estimates of the heating system choice model and heating demand function, together with observed energy prices and household characteristics, are used to simulate heating system choices and heating demand by state and year for the period 1950–2000. Households are assumed to buy a new heating system in the year the residence is

¹¹ We would have preferred to estimate the heating demand function for the 1960 and 1970 Census as well, but households did not report energy expenditures prior to the 1980 Census. We deal with the incomplete data prior to 1980 by using the estimated parameters for 1980 in predicting heating demand. The resulting estimates are conservative because natural gas consumption prior to 1980 would have tended to be higher than consumption in 1980 because of the increasing availability of energy-efficient building materials starting in the 1970s.

¹² Little previous work has been done to assess the reliability of these self-reported measures of expenditure. Accordingly, in Sec. VI.A, we compare the self-reported measures of expenditure with residential gas sales reported by natural gas utilities. Generally, the two measures are similar, suggesting that the magnitude of the bias in the self-reported measures is small.

constructed. In addition, we assume that all households, including households living in homes built before the 1950s, replace their heating system according to annual retirement probabilities from the Energy Information Administration (EIA 2006*b*). These retirement probabilities imply that, on average, a household replaces its heating system once every 17.5 years. In the following section we compare households' actual choices with the choices implied by the model. This allows us to calculate physical shortages of natural gas and to measure the allocative cost of price ceilings.

VI. Results

A. *Physical Shortages*

This section contrasts our model's predictions of residential demand for natural gas during the period 1950–2000 with actual consumption. Although our ultimate objective is to measure the allocative cost of price regulation during this period, the measures of shortage that are discussed in this section provide a first check on the ability of the model to replicate the well-known qualitative pattern of physical shortages during this period. Moreover, these results are of independent interest in that they provide the most comprehensive assessment to date of the magnitude, timing, and geographic distribution of physical shortages.

Figure 4 describes residential demand for natural gas in the United States by year for 1950–2000. The figure plots two alternative measures of the actual level of natural gas consumption. The dashed line is residential consumption of natural gas in the United States as reported by natural gas utilities in EIA (2006*c*). The dotted line is residential consumption of natural gas from the census data. Both measures of actual consumption increase steadily between 1950 and 1970 and then level off during the later period. Although the fit is not perfect, it is reasonably close. This comparison suggests that our census data provide a reasonably accurate measure of natural gas consumption.

The solid line in figure 4 is the level of natural gas demand predicted by our model. Since our heating system choice model is estimated using choices made after natural gas deregulation, the model is able to describe the important counterfactual of what demand would have been at observed prices had all households had access to natural gas. As explained in detail in Section VI.B, our empirical methodology can be used to show how much natural gas would have been consumed at these prices during 1950–2000 based on preferences revealed in the postregulation period and controlling for household demographics and housing characteristics that affect heating demand. Simulated demand follows actual consumption reasonably closely during the 1950s and 1960s,

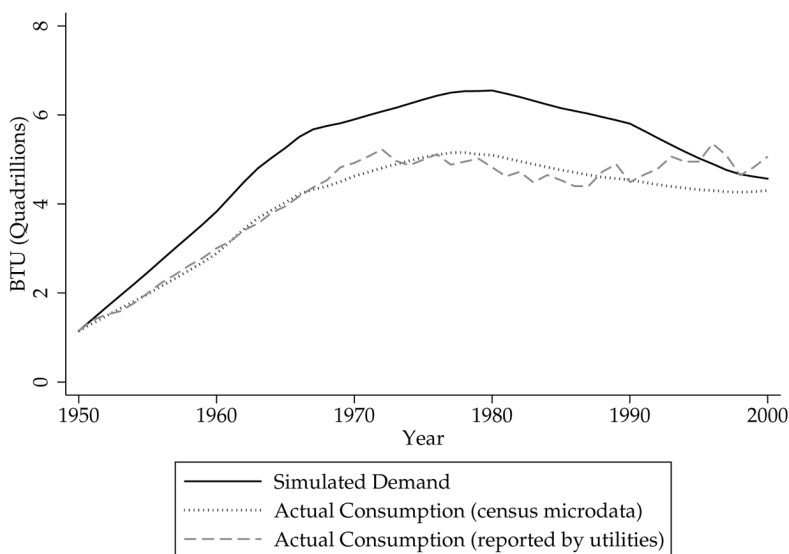


FIG. 4.—Residential demand for natural gas in the United States, 1950–2000

although even at the beginning of the sample period there is evidence of a small but growing physical shortage of natural gas. The figure indicates large differences between simulated demand and actual consumption during the 1970s and 1980s, with the gap narrowing during the 1990s. The finding of a substantial natural gas shortage during the 1970s is consistent with previous studies.¹³

Figure 5 describes residential demand by region for the same period. The pattern for the Northeast is similar to the national pattern, with large differences between simulated and actual demand throughout the period. The Midwest also experiences shortages, but they are smaller in magnitude and less persistent. In other regions there is very little shortage. In the West, in particular, there is essentially no difference between simulated demand and actual consumption. The smaller shortages in the South and West likely are due to several factors. First, these regions include the natural gas-producing states, where one would expect to see no shortage. Second, the warmer climate and greater access to cheap hydroelectric electricity in many of these states imply a lower overall level of demand for natural gas and thus less scope for shortages.

¹³ According to Vietor (1984), shortfalls in contractually obligated deliveries to pipelines began in the early 1970s and reached 15 percent of the entire market for natural gas in 1976. MacAvoy (1983, 85) reports even larger shortfalls: "Forced curtailments of committed deliveries increased from 12% of total interstate demand in 1973 to 30% in 1975. Further curtailments caused the short deliveries to exceed 40% of the total in 1978."

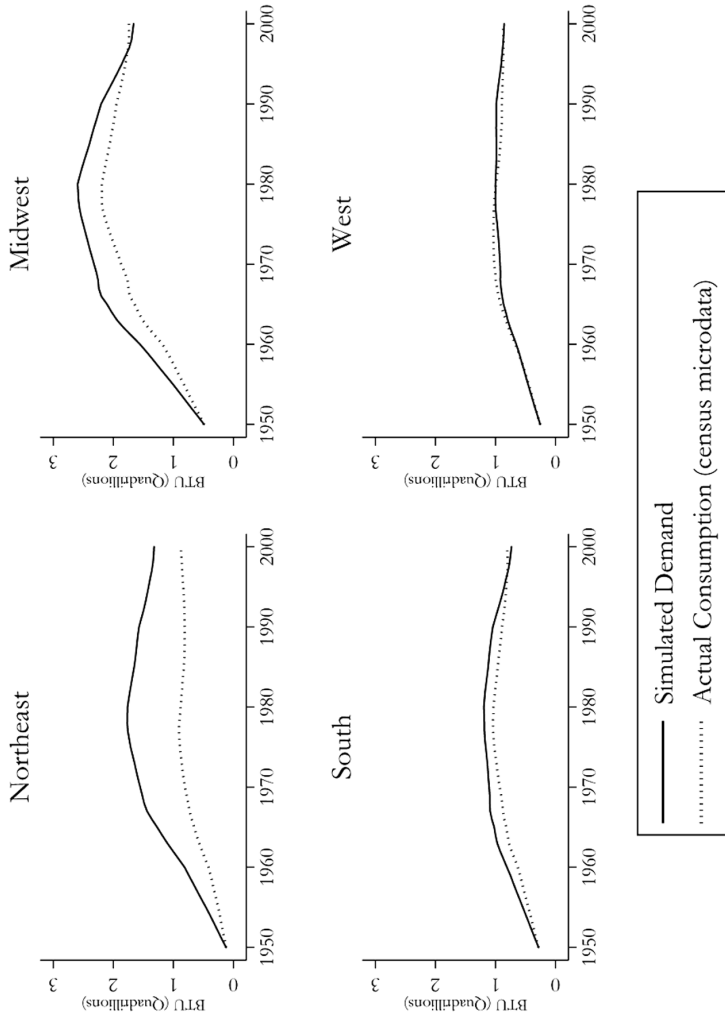


FIG. 5.—U.S. residential demand for natural gas by region, 1950–2000

Third, following deregulation, there was much more new housing construction in these regions than in the Northeast and Midwest, allowing the effects of regulation to be overcome more quickly than elsewhere.

We find that during 1950–2000, total U.S. demand for natural gas exceeded observed sales of natural gas by an average of 19.4 percent, with the largest shortages during the 1970s and 1980s. Our estimates of the physical shortage provide a more complete description of the degree of disequilibrium in the natural gas market than previously used measures such as the curtailment of gas deliveries. In particular, our measure correctly incorporates not only demand from existing delivery contracts but the unrealized demand from prospective new customers as well. The need for such a comprehensive measure of shortages has long been recognized in the literature. For example, MacAvoy (1983, 85) notes that “the excess demand of those excluded from gas markets was not listed as a ‘shortage,’ and yet substantial numbers of potential new residential and commercial customers denied service by state and federal regulations were ‘short’ by the entire amount of their potential demands.” Our estimates provide quantitative content for MacAvoy’s conjecture that curtailments understate the degree of disequilibrium in the market because they fail to incorporate demand from potential new customers who are prevented from adopting gas heating by their local utilities in an effort to preserve service to existing customers.

Finally, our analysis sheds new light on the persistence of the effects of price regulation. Previous studies assumed that the effects of regulation were quickly alleviated after the Natural Gas Policy Act in 1978. For example, many observers have pointed to the apparent “gas glut” in the early 1980s as evidence of the end of the effects of price regulation. In sharp contrast, we find that simulated demand exceeded actual consumption of natural gas throughout the 1980s and well after the complete deregulation of wellhead prices in 1989. These highly persistent legacy effects of earlier heating system choices become apparent only owing to our use of a microeconomic model, illustrating the importance of explicit modeling of household decisions. At any point in time, demand is derived from households that purchased their home heating systems many years ago. Thus, even many years after complete deregulation, a substantial fraction of households continued to be locked into suboptimal heating system choices, prolonging the effects of regulatory policies long beyond the official end of regulation. For example, under our assumptions, 71 percent of households in 2000 are living in homes with heating systems that were purchased prior to complete deregulation in 1989. This helps explain the fact that simulated demand continues to exceed actual consumption even in 2000.

B. *Allocative Costs*

In this subsection we use the demand model described in Section IV to estimate the allocative cost of price ceilings in the U.S. residential market for natural gas. As defined in Section II, the allocative cost of a price ceiling is the welfare loss that results from not allocating a good to the buyers that value it the most, measured as the difference between the consumer surplus under the actual allocation and the consumer surplus under the allocation when buyers are rationed efficiently. Under efficient rationing, the good is provided to the buyers with the highest reservation prices, and welfare cannot be improved on by reallocating the good among buyers. In an unregulated market this allocation is achieved with a national market-clearing price. In a regulated market, the actual allocation typically does not provide the good to the buyers who value it the most, and the allocative cost refers to the welfare gains that can be realized from replacing the actual allocation by the allocation under efficient rationing.

Our calculation of the allocative cost involves the following steps. In step 1, we simulate the reservation price for each household i . The reservation price, p_{ig}^* , is the natural gas price that makes household i indifferent between natural gas and the best of the alternatives:

$$p_{ig}^* = \frac{\max(u_{ie}, u_{io}) - \alpha_{0g} - \alpha'_{2g}w_i - \epsilon_{ig}}{\alpha_{1g} + \alpha'_{3g}z_i}, \tag{4}$$

where g , e , and o index natural gas, electricity, and heating oil. We evaluate u_{ie} and u_{io} according to the indirect utility function defined in equation (2) using electricity and heating oil prices, p_{ie} and p_{io} ; the observable characteristics w_i and z_i for each household from the census data; the estimates of the parameters α_{0e} , α_{1e} , α_{1o} , α_{2e} , α_{2o} , α_{3e} , and α_{3o} from table 2; and the random components ϵ_{ie} and ϵ_{io} . We compute p_{ig}^* from equation (4) using these estimates of u_{ie} and u_{io} ; the estimates of the parameters α_{0g} , α_{1g} , α_{2g} , and α_{3g} from table 2; and the random component ϵ_{ig} . All the random components are i.i.d. draws from the extreme value distribution.¹⁴

In step 2, we calculate each household’s consumer surplus for this household’s simulated reservation price. A household’s annual consumer surplus is the difference between its simulated reservation price and the actual price of natural gas for household i in the year in which the heating system was purchased, multiplied by the household’s annual demand for gas, \hat{x}_{ig} , predicted from equation (1) using the parameter

¹⁴ This procedure introduces simulation error because for each household the simulated reservation price represents one possible realization. However, because the number of households in our sample is very large, the procedure introduces little variation in national and state measures of the allocative cost. Moreover, the bootstrap standard errors reported in the paper take this simulation error into account.

estimates in table 3:

$$CS_i = (p_{ig}^* - p_{ig}) \hat{x}_{ig}. \quad (5)$$

In step 3, we compute total consumer surplus for all households under efficient rationing:

$$CS_t^{er} = \sum_{i=1}^n CS_i \times 1(er_{it}), \quad (6)$$

where $1(er_{it})$ is an indicator variable equal to one if household i receives natural gas in year t under efficient rationing. Under efficient rationing, natural gas is allocated to households in the order of their simulated reservation prices, until all available gas has been allocated. We repeat this procedure for each year, assuming that households that received natural gas in the past will be able to continue to receive natural gas, so that the allocation problem is limited to reallocating gas among potential new customers.

In step 4, we calculate θ_{st} , the fraction of households that had access to natural gas among all households in state s and year t that would have wanted to choose natural gas heating:

$$\theta_{st} = \frac{\sum_{i \in s,t} \hat{x}_{ig} \times 1(\text{census}_{it})}{\sum_{i \in s,t} \hat{x}_{ig} \times 1(p_{ig}^* > p_{ig})}, \quad (7)$$

where $1(\text{census}_{it})$ is an indicator variable equal to one if in the census data household i receives natural gas in year t , and the indicator variable $1(p_{ig}^* > p_{ig})$ is equal to one for households with a reservation price that exceeds the actual price of natural gas. Thus, $1 - \theta_{st}$ is a percentage measure of the shortage of natural gas for a given state and year.¹⁵ The numerator in equation (7) represents the actual consumption of natural gas in state s among households choosing a heating system in year t , as observed in the census data. The denominator is simulated demand, that is, the level of demand predicted by the model had all households in the market for heating systems in year t had access to natural gas at the observed market price.

In step 5, we compute the total consumer surplus under actual rationing. For this calculation we use the actual allocation of natural gas across states, and we assume that during shortages the within-state allocation of natural gas was random among households with a reservation price higher than the observed price. Thus, we randomly assign natural gas to the fraction θ_{st} of households in a particular state and year with

¹⁵ It is important to clarify that θ_{st} is calculated as a function of *new demand* rather than *total demand*. New demand refers to demand for natural gas derived from households that adopt natural gas during a particular year, either because they are purchasing a new home or because they are replacing the heating system in an existing home. Total demand for natural gas is calculated by cumulating new demand over time.

a simulated reservation price higher than the observed price. As in step 3, we postulate that households that received natural gas in the past are able to continue to receive natural gas, so that this allocation occurs only among households making a heating system choice in a particular year. Total consumer surplus under the actual allocation is then calculated by summing over all households:

$$CS_t^{ar} = \sum_{i=1}^n CS_i \times 1(ar_{it}), \quad (8)$$

where $1(ar_{it})$ is an indicator variable equal to one if household i receives natural gas in year t under actual rationing.

In step 6, we compute the allocative cost in year t as

$$AC_t = CS_t^{er} - CS_t^{ar}. \quad (9)$$

It is important to stress that the estimation of the allocative cost relies on the parametric econometric model of Section IV. Our econometric analysis builds on the framework developed in Dubin and McFadden (1984). The possibility that this parametric model may be misspecified has to be taken seriously. For example, the assumption of rational optimizing behavior embodied in the model may not reflect actual household behavior. In related work, Hausman (1979) discusses the idea that households may not be as forward looking as postulated in typical parametric models of durable goods purchases. In addition, our empirical results hinge on the premise that, controlling for covariates, households' unconstrained choices in the 1990s can be used to predict the choices that households would have made during earlier decades in the absence of regulation. To the extent that these assumptions are violated, the estimates of allocative cost will be biased. In recognition of this concern we assess the out-of-sample fit of the model in Section VI.C and show that the parametric household model appears to provide a good approximation of the aggregate data.

Apart from hinging on the model structure, our estimate of the allocative cost also hinges on the assumption that the within-state allocation of natural gas was random among households with a reservation price above the observed price. This is a reasonable approximation a priori given that natural gas retailers had no incentive to target available supplies to households with a high willingness to pay. Moreover, there was limited scope for queuing in the retail natural gas market because households had to select a heating system when they moved in or when their heating system failed. We have also examined the within-state allocation empirically and found that it was essentially random among households that, according to the model, would have preferred natural gas. Nevertheless, this assumption is only an approximation, and to the

TABLE 4
ESTIMATES OF AVERAGE ANNUAL ALLOCATIVE
COST IN MILLIONS OF YEAR 2000 DOLLARS

	Estimate
Nationwide	\$3,560 (34.1)
10 most affected states:	
New York	\$1,081 (12.4)
Pennsylvania	\$652 (10.3)
New Jersey	\$284 (4.9)
North Carolina	\$259 (5.1)
Virginia	\$253 (4.9)
Massachusetts	\$196 (4.0)
Maryland	\$162 (4.0)
Indiana	\$145 (5.2)
Connecticut	\$125 (3.0)
South Carolina	\$81 (3.0)

NOTE.—Bootstrap standard errors based on 100 replications are shown in parentheses.

extent that there are departures from the random allocation, one would expect our estimates of allocative cost to be biased.

Table 4 reports estimates of the mean annual allocative cost. During the period 1950–2000, the mean annual allocative cost in the United States was \$3.6 billion with a peak of \$5.0 billion in 1980. This amount represents, on average, 14 percent of total residential expenditure on natural gas. The table also reports bootstrap standard errors. Bootstrap replicates are constructed by reestimating all model parameters for each bootstrap sample and simulating the implied allocative cost. As one would expect given the large sample size, the standard errors are generally small. Our estimates of the allocative costs are about half as large as previous estimates in the literature for the conventional deadweight loss. MacAvoy (2000), for example, reports a mean annual deadweight loss of \$10.5 billion between 1968 and 1977. We find that during this same period, the mean annual allocative cost in the residential market was \$4.8 billion. Because this allocative cost is in addition to the conventional deadweight loss, our estimates show that total welfare losses from natural gas price regulation were considerably larger than previously believed.¹⁶

The allocative cost can be decomposed into a within-state component and an across-state component. The within-state component is the gain in consumer surplus from reallocating gas within states to households with the highest reservation prices, holding constant the allocation

¹⁶ These large allocative effects are consistent with theoretical evidence about the relative size of allocative cost and conventional deadweight loss. Glaeser and Luttmer (2003) show that allocative cost exceeds conventional deadweight loss when the demand curve is linear, and price ceilings reduce quantity supplied by less than 50 percent.

across states. The across-state component is the gain in consumer surplus from reallocating natural gas across states, assuming efficient rationing within states. All else equal, total consumer surplus increases when natural gas is reallocated toward states in which the marginal household has a high reservation price. We find that, during the period 1950–2000, the mean annual within-state allocative cost was \$2.93 billion and the mean annual across-state allocative cost was \$0.63 billion. This decomposition reveals that the bulk of the allocative cost comes from the misallocation of natural gas within states.

Most of these costs were borne by households in the Northeast, Midwest, and South, with households in the West bearing a smaller amount. Total allocative cost in the Midwest and South decreased substantially during the 1980s and 1990s, though in neither case did costs disappear by 2000. In the Northeast, costs were more persistent, with large costs remaining in 2000. The adjustment was particularly slow in the Northeast because there was less new housing construction compared to the South or West.

Table 4 reports the results by state for the 10 most affected states. It is interesting to note that price regulation was supported in Congress even by states in which households faced large allocative costs. For example, in the 1973 Buckley Amendment (S2776), most senators from states in the Northeast and Midwest voted to continue regulating prices. We have documented this pattern, confirming a statistically significant positive correlation between allocative costs and congressional support for regulation. These findings are consistent with the view that during the period of regulation, politicians focused on the benefits of price regulation for existing customers and discounted the costs to customers without access to natural gas. Although it is possible that the transfers received from gas-producing states may have outweighed the combined deadweight and allocative costs in some states, we know that this was not the case at the national level.

C. Out-of-Sample Fit of the Model

Our empirical results in Sections VI.A and VI.B hinge on the parametric specification of the model and on the premise that, conditional on observables, households' unconstrained choices in the 1990s can be used to predict the choices that households would have made during earlier decades in the absence of regulation. There are several approaches to verifying the validity of this empirical strategy. In this subsection we show that our estimation procedure yields only minimal estimates of allocative cost when applied to settings in which the market operates freely. The fact that the model performs reasonably well in

these out-of-sample contexts adds credibility to the results presented in the previous subsections.

One approach is to evaluate the out-of-sample fit of the model for Texas, Louisiana, New Mexico, and Oklahoma, the four states accounting for the bulk of U.S. natural gas production during the period of regulation.¹⁷ A key feature of the price regulation implemented during this period is that it applied only to interstate sales of natural gas. Because federal regulators did not have jurisdiction over intrastate sales, prices for gas sold within gas-producing states were unregulated. Thus, there is no reason to expect shortages in Louisiana, New Mexico, Oklahoma, or Texas. Simulated demand for these states should closely follow actual consumption. We find that our estimated model passes this test of out-of-sample fit. The model estimates imply no shortages in gas-producing states. Moreover, allocative cost within gas-producing states is effectively zero, averaging only \$5.4 million (or \$0.83 per capita) annually, compared with \$152.4 million (or \$25.59 per capita) for states in the Northeast.

An alternative approach to evaluating the out-of-sample fit of the model involves estimating the allocative cost for households that installed heating systems after deregulation of natural gas prices was completed in 1989. Since these households did not face price ceilings, their allocative cost should be negligible. We test this proposition by splitting the 2000 census data into two random samples. One of the subsamples is used to estimate household preferences conditional on covariates; the model is then used to predict allocative cost for the households in the other subsample. We find that the mean annual allocative cost among households that made heating system choices during the period of regulation was \$80.81 per household. In contrast, the mean annual allocative cost among households living in homes purchased in 2000 was \$11.02 per household. The latter amount is small compared to the allocative cost estimated for the earlier period but is not zero. If we interpret that estimate as a measure of the potential model misspecification, an alternative estimate of the mean annual allocative cost can be constructed by subtracting this baseline from the estimate of \$3.6 billion, resulting in a mean annual allocative cost of \$3.1 billion, which is smaller than the original estimate of \$3.6 billion but still of the same order of magnitude.

¹⁷ Natural gas is more expensive to transport than oil, and thus most natural gas consumed in the United States is produced in North America. Net imports of natural gas increased from less than 1 percent in 1950 to 15.2 percent in 2000, with 94 percent of the natural gas imports in 2000 coming from Canada (see EIA 2006a, table 6.3).

VII. Conclusion

Whereas the importance of allocative costs is well recognized as a theoretical matter, its quantitative importance in real-life markets has remained uncertain, owing to the difficulties of empirically quantifying such costs. Our study is the first to demonstrate how allocative costs in a market subject to price ceilings may be estimated. We focused on the U.S. residential market for natural gas. Our analysis showed that the allocative cost in this market averaged \$3.6 billion annually during the period 1950–2000. While our estimates of allocative cost are large, total allocative cost for all consumers is likely to be even larger than the magnitudes reported here, given that our analysis has been restricted to the residential market.

Our analysis illustrates the importance of careful ex ante economic analysis. Price ceilings for natural gas were supposed to help consumers, particularly consumers in northern markets, who in the 1950s were concerned about rising natural gas prices. We found that these very consumers ended up bearing a disproportionately large share of the allocative cost. While consumers with access to natural gas indeed benefited from lower prices, not all consumers had access to the market, and those who had access were not necessarily the consumers with the highest willingness to pay.

Not only is our analysis relevant for understanding the consequences of regulation in the U.S. natural gas market, but it has implications for other markets as well. Allocative costs due to price ceilings arise more generally whenever the good in question cannot be readily traded in secondary markets, as would be the case, for example, in insurance, health care, or telecommunications markets. The broader conclusion of our paper is that policy makers in conducting an ex ante economic analysis of regulatory reform ought to take careful account of the allocative cost of price regulation in addition to the conventional dead-weight loss. In particular, our analysis showed that the effects of price ceilings may be very uneven across households and states, that it is difficult to determine in advance how the cost will be distributed geographically, and that the adverse effects of price ceilings tend to last much longer than the regulatory policies themselves.

References

- American Gas Association. 1945–69. *Gas Facts: A Statistical Record of the Gas Utility Industry*. Lexington, NY: American Gas Assoc., Bur. Statis.
- Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. von Haefen. 2009. "Distributional and Efficiency Impacts of Increased US Gasoline Taxes." *A.E.R.* 99 (3): 667–99.
- Bernard, Jean Thomas, Denis Bolduc, and Donald Belanger. 1996. "Quebec

- Residential Electricity Demand: A Microeconomic Approach." *Canadian J. Econ.* 29 (1): 92–113.
- Bradley, Robert L., Jr. 1996. "Foreword." In *New Horizons in Natural Gas Deregulation*, edited by Jerry Ellig and Joseph P. Kalt. Westport, CT: Praeger.
- Braeutigam, Ronald R., and R. Glenn Hubbard. 1986. "Natural Gas: The Regulatory Transition." In *Regulatory Reform: What Actually Happened*, edited by Leonard W. Weiss and Michael W. Klass. Boston: Little, Brown.
- Breyer, Stephen, and Paul W. MacAvoy. 1973. "The Natural Gas Shortage and the Regulation of Natural Gas Producers." *Harvard Law Rev.* 86 (6): 941–87.
- Card, David, and Alan B. Krueger. 1994. "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania." *A.E.R.* 84 (4): 772–93.
- Cookerbo, Leslie J., Jr. 1958. *Competition in the Field Market for Natural Gas*. Pamphlet, vol. 44, no. 4. Houston: Rice Inst.
- Deacon, Robert T., and Jon Sonstelie. 1989. "The Welfare Costs of Rationing by Waiting." *Econ. Inquiry* 27 (2): 179–96.
- . 1991. "Price Controls and Rent Dissipation with Endogenous Transaction Costs." *A.E.R.* 81 (5): 1361–73.
- Dubin, Jeffrey A. 1985. *Consumer Durable Choice and Demand for Electricity*. Amsterdam: North-Holland.
- Dubin, Jeffrey A., and Daniel L. McFadden. 1984. "An Econometric Analysis of Residential Electric Durable Good Holdings and Consumption." *Econometrica* 52 (2): 345–62.
- Edison Electric Institute. 1945–69. *Statistical Year Book of the Electric Utility Industry*. New York: Edison Electric Inst.
- EIA (Energy Information Administration). 2006a. *Annual Energy Review 2005*. DOE/EIA-0384. Washington, DC: Energy Information Admin.
- . 2006b. *Assumptions to the Annual Energy Outlook 2006*. DOE/EIA-0445. Washington, DC: Energy Information Admin.
- . 2006c. *State Energy Consumption, Price, and Expenditure Estimates*. Washington, DC: Energy Information Admin.
- Fabrizio, Kira R., Nancy L. Rose, and Catherine D. Wolfram. 2007. "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on U.S. Electric Generation Efficiency." *A.E.R.* 97 (4): 1250–77.
- Feng, Ye, Don Fullerton, and Li Gan. 2005. "Vehicle Choices, Miles Driven and Pollution Policies." Working Paper no. 11553, NBER, Cambridge, MA.
- Frech, H. E., III. 2000. "Physician Fees and Price Controls." In *American Health Care: Government, Market Processes, and the Public Interest*, edited by Roger D. Feldman. New Brunswick, NJ: Transaction.
- Frech, H. E., III, and William C. Lee. 1987. "The Welfare Cost of Rationing-by-Queuing across Markets: Theory and Estimates from the U.S. Gasoline Crises." *Q.J.E.* 102 (1): 97–108.
- Friedman, Milton, and George J. Stigler. 1946. *Roofs or Ceilings? The Current Housing Problem*. Popular Essays on Current Problems, vol. 1, no. 2. Irvington-on-Hudson, NY: Found. Econ. Educ.
- Glaeser, Edward L., and Erzo F. P. Luttmer. 2003. "The Misallocation of Housing under Rent Control." *A.E.R.* 93 (4): 1027–46.
- Goldberg, Pinelopi K. 1998. "The Effects of the Corporate Average Fuel Efficiency Standards in the US." *J. Indus. Econ.* 46 (1): 1–33.
- Hausman, Jerry A. 1979. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *Bell J. Econ.* 10 (1): 33–54.

- Hayek, Friedrich A. von. 1931. "Wirkungen der Mietzinsbeschränkungen." *Schriften des Vereins für Socialpolitik* 182:253–70.
- Herbert, John H. 1992. *Clean Cheap Heat: The Development of Residential Markets for Natural Gas in the United States*. New York: Praeger.
- Holzer, Harry J., Lawrence F. Katz, and Alan B. Krueger. 1991. "Job Queues and Wages." *Q.J.E.* 106 (3): 739–68.
- Leitzinger, Jeffrey, and Martin Collette. 2002. "A Retrospective Look at Wholesale Gas: Industry Restructuring." *J. Regulatory Econ.* 21 (1): 79–101.
- Lindahl, Martin L. 1956. "Federal Regulation of Natural Gas Producers and Gatherers." *A.E.R. Papers and Proc.* 46 (2): 532–44.
- Luttmer, Erzo F. P. 2007. "Does the Minimum Wage Cause Inefficient Rationing?" *B.E. J. Econ. Analysis and Policy* 7, no. 1 (Contributions), art. 49. <http://www.bepress.com/bejeap/contributions>.
- MacAvoy, Paul W. 1971. "The Regulation-Induced Shortage of Natural Gas." *J. Law and Econ.* 14 (1): 167–99.
- . 1983. *Energy Policy: An Economic Analysis*. New York: Norton.
- . 2000. *The Natural Gas Market*. New Haven, CT: Yale Univ. Press.
- MacAvoy, Paul W., and Robert S. Pindyck. 1973. "Alternative Regulatory Policies for Dealing with the Natural Gas Shortage." *Bell J. Econ.* 4 (2): 454–98.
- . 1975. *The Economics of the Natural Gas Shortage (1960–1980)*. Amsterdam: North-Holland.
- Manning, Fred, and Clifford Winston. 1985. "A Dynamic Empirical Analysis of Household Vehicle Ownership and Utilization." *Rand J. Econ.* 16 (2): 215–36.
- Mansur, Erin T., Robert Mendelsohn, and Wendy Morrison. 2008. "Climate Change Adaptation: A Study of Fuel Choice and Consumption in the U.S. Energy Sector." *J. Environmental Econ. and Management* 55 (2): 175–93.
- Neuner, Edward J. 1960. *The Natural Gas Industry: Monopoly and Competition in Field Markets*. Norman: Univ. Oklahoma Press.
- Olsen, Edgar O. 1972. "An Econometric Analysis of Rent Control." *J.P.E.* 80 (6): 1081–1100.
- Raymon, Neil. 1983. "Price Ceilings in Competitive Markets with Variable Quality." *J. Public Econ.* 22:257–64.
- Smith, Rodney T., and Charles E. Phelps. 1978. "The Subtle Impact of Price Controls on Domestic Oil Production." *A.E.R. Papers and Proc.* 68 (2): 428–33.
- Suen, Wing. 1989. "Rationing and Rent Dissipation in the Presence of Heterogeneous Individuals." *J.P.E.* 97:1384–94.
- Train, Kenneth E. 1986. *Qualitative Choice Analysis: Theory, Econometrics, and an Application to Automobile Demand*. Cambridge, MA: MIT Press.
- Vietor, Richard H. K. 1984. *Energy Policy in America since 1945*. New York: Cambridge Univ. Press.
- Viscusi, W. Kip, Joseph E. Harrington Jr., and John M. Vernon. 2005. *Economics of Regulation and Antitrust*. Cambridge, MA: MIT Press.
- Weitzman, Martin L. 1977. "Is the Price System or Rationing More Effective in Getting a Commodity to Those Who Need It the Most?" *Bell J. Econ.* 8 (2): 517–24.
- Welch, Finis. 1974. "Minimum Wage Legislation in the United States." *Econ. Inquiry* 12 (3): 285–318.
- West, Sarah. 2004. "Distributional Effects of Alternative Vehicle Pollution Control Policies." *J. Public Econ.* 88:735–57.