

WORKING PAPERS



A VARIANCE SCREEN FOR COLLUSION

**Rosa M. Abrantes-Metz
Luke M. Froeb
John F. Geweke
Christopher T. Taylor**

WORKING PAPER NO. 275

March 2005

FTC Bureau of Economics working papers are preliminary materials circulated to stimulate discussion and critical comment. The analyses and conclusions set forth are those of the authors and do not necessarily reflect the views of other members of the Bureau of Economics, other Commission staff, or the Commission itself. Upon request, single copies of the paper will be provided. References in publications to FTC Bureau of Economics working papers by FTC economists (other than acknowledgment by a writer that he has access to such unpublished materials) should be cleared with the author to protect the tentative character of these papers.

**BUREAU OF ECONOMICS
FEDERAL TRADE COMMISSION
WASHINGTON, DC 20580**

A Variance Screen for Collusion

Rosa M. Abrantes-Metz*

Luke M. Froeb[♦]

John F. Geweke[^]

Christopher T. Taylor*

March 2005

Abstract

In this paper, we examine price movements over time around the collapse of a bid-rigging conspiracy. While the mean *decreased* by sixteen percent, the standard deviation *increased* by over two hundred percent. We hypothesize that conspiracies in other industries would exhibit similar characteristics and search for “pockets” of low price variation as indicators of collusion in the retail gasoline industry in Louisville. We observe no such areas around Louisville in 1996-2002.

*Consultant, NERA Economic Consulting. This study was initiated while at the Bureau of Economics, Federal Trade Commission. rosa.abrantes-metz@nera.com.

[♦]Director, Bureau of Economics, Federal Trade Commission and Vanderbilt University luke.froeb@owen.vanderbilt.edu.

[^]Professor, Departments of Economics and Statistics, University of Iowa. john-geweke@uiowa.edu.

*Deputy Assistant Director, Antitrust Division I, Bureau of Economics, Federal Trade Commission. ctaylor@ftc.gov.

We would like to thank participants at the International Industrial Organization Conference, William Dillon, Albert Metz, Timothy Muris, Frederick Warren-Boulton and our colleagues at the FTC and NERA for helpful comments and suggestions, in particular, Sumanth Addanki, Gregory Leonard and Ramsey Shehadeh. We also thank David Yans and Ning Yan for excellent assistance, and two anonymous reviewers of the FTC working paper series for their comments. The views expressed are those of the authors and do not necessarily reflect the views of the FTC Commission or any individual Commissioner, nor the views of NERA or its clients. All errors are ours. Comments solicited.

1. Introduction

Developing data screens to detect anticompetitive conspiracies has been an elusive goal of competition agencies for many years. In the 1970's, for example, the U.S. Department of Justice formed an "identical bids" unit that investigated government procurement auctions in which identical bids were submitted. In six years, the unit failed to uncover a single conspiracy.¹ Currently, the Federal Trade Commission (FTC) monitors gasoline prices to identify unusually high prices. The agency further analyzes the situation if no obvious explanation for the high prices can be found. To date, price "anomalies" have tended to be either short-lived, or found to have obvious explanations, like a pipeline break or refinery outage.²

Conspiracies are difficult to detect because they take many forms. All raise price above the competitive level, but the cost data necessary to estimate the competitive price level are rare. Consequently, use of the price level as a screen is limited to industries where the competitive level is known. Easier-to-measure characteristics of specific types of conspiracies, like identical bids, have been the focus of enforcement and academic efforts to identify collusive behavior. For example, Porter and Zona (1993) show that losing conspiracy members bid differently than non-conspiracy members. While this is useful for proving the existence of a conspiracy (Froeb and Shor, forthcoming), it is less useful as a screen unless you already know who is a member of the conspiracy. Bajari and Yi (2003) propose exchangeability tests to identify non-random patterns of bidding across auctions indicative of bid suppression, bid rotation, or the use of side payments to reward losing conspiracy members for not competing aggressively. This is potentially useful as a screen, but requires difficult-to-collect data on losing bids and bidder identities across a wide set of auctions.

We contribute to this literature by proposing a screen based on the observed differences between a conspiracy and the period immediately following its collapse. The average weekly price level *decreased* by 16%, while the standard deviation of price *increased* by 263%. We hypothesize that conspiracies in other industries would also exhibit low price variance and

¹ As told by Frederick Warren-Boulton, former Deputy Assistant Attorney General for Antitrust.

² The FTC models the price spreads between selected cities and reference cities, where collusion is unlikely due to the large number of sellers. Underlying this price spread model is an implicit assumption that big price spreads represent arbitrage opportunities which would be exploited in a competitive market. In addition the FTC has examined the historic level of the spreads for the underlying causes of regional price differences (Taylor and Fischer, 2003).

design a screen based on standard deviation of price normalized by its mean, or the coefficient of variation.

There is some theoretical justification for a variance screen if it is costly to coordinate price changes or if the cartel must solve an agency problem as in Athey, Bagwell, and Sanchiro (2004). Similarly, in Harrington and Chen (2004), to conceal the existence of a cartel from its customers, cartel members only partially pass through cost changes. This partial pass-through lowers the variance of price relative to the variance that would obtain in a competitive market. Both Athey et al. (2004) and Harrington and Chen (2004) predict that during conspiracies, price responds less to cost shocks under collusion, or that collusion reduces price variance.

We design a screen based on variance and apply it to the retail gasoline industry in Louisville in 1996-2002. A cartel the size of a city would be very costly to organize and police, but there may be a degree of market power that could make elimination of localized competition profitable. Our screen would identify a potential cartel as a group of gasoline stations located close to one another that exhibit lower price variation and higher prices relative to other stations in the city.

We estimate variance at the 263 gasoline stations in Louisville in 1996-2002 which accept “fleet” credit cards, used by sales people, or workers whose job requires driving.³ Every time a driver purchases gasoline using a fleet card, the price is recorded in an electronic data base. Prices are observed daily, unless no fleet credit cards are scanned during the twenty-four-hour period, in which case, data are recoded as missing for the day in question. In our data set there are no gasoline stations without missing observations.

To estimate variance for stations with missing observations, whose presence could be crucial in an application like ours, we use Markov chain Monte Carlo with Gibbs sampling combined with data augmentation, a Bayesian-based data imputation technique. We then look for groups of stations that exhibit unusually small price variation, as measured by the standard deviation normalized by the mean. We find no such groups of stations in and around Louisville in 1996-2002. Instead, observed pricing differences across gas stations seem to be driven more by proximity to major arteries and by brand characteristics rather than by pricing of neighboring stations.

³ These are purchases made with Wright Express fleet cards. Wright Express is the largest provider of fleet card services. Its cards are currently accepted at 90 percent of gasoline stations in the United States.

While any screen is likely to miss some conspiracies, and falsely identify others, we conclude by noting that our screen has four advantages:


- It does not require cost data to implement.
- It is easy to estimate and has a known distribution.
- It has theoretical and empirical support.
- Even if it were to become known that competition agencies were screening for low variance, it would still be costly to disguise cartel behavior if there are costs of changing or coordinating price changes.

2. The Collapse of a Bid Rigging Conspiracy

In this section we estimate price variance across the collapse of a bid-rigging conspiracy. The data were collected for a sentencing hearing following convictions of several defendants for rigging bids for frozen seafood sold to the Defense Personnel Support Center (DPSC) in Philadelphia. We focus on one of the items involved in the conspiracy, frozen perch, for which we have a good measure of cost, the price of fresh perch. A more complete description of the conspiracy as well as econometric estimation of price change is found in Froeb, Koyak and Werden (1993).

Figure 1 illustrates the collapse of the conspiracy following reports of an investigation. Prices for frozen perch dropped dramatically and remained at a much lower level. In the post-conspiracy period, price began to co-vary more closely with cost, and exhibited larger variation. We compare prices and costs in what we call the “collusive” regime (to the left of the vertical lines) to prices in the “competitive” regime (to the right of the vertical lines), and assume that the period in between the vertical lines represents a transition period from collusion to competition.

Figure 1



In the Figure above, we plot the average weekly price paid by the Philadelphia DPSC for frozen perch filets from 1987 through September 1989 (U.S. dollars per pound). The cost data are the average monthly price of fresh perch. Table 1 below summarizes our findings. In the first row of this table, we see the mean price decreased by 16% while the standard deviation of price increased by 263%. When standardized by the mean, the standard deviation of the price, or its coefficient of variation increased by 332% from collusion to competition.

The mean and standard deviation of cost are also higher under competition, but not by enough to account for the increase in price variance. We report the cost data because we have it readily available and it is necessary to rule out the alternate hypothesis that it is changes in cost that are accounting for the observed difference in the price.

Table 1
Means and Standard Deviations for Perch Price and Cost (\$/pound)

Statistics	Collusion	Competition	Differences Across Regimes
<u>PRICE</u>			
Mean	3.544	2.97	-16.2%
Std.Dev.	0.078	0.283	263%
CV = Std.Dev./Mean	0.022	0.095	332%
<u>COST</u>			
Mean	0.722	0.771	6.8%
Std.Dev.	0.114	0.173	51.8%
CV = Std.Dev./Mean	0.158	0.224	41.8%

3. The Retail Gasoline Industry in Greater Louisville

In the previous section we found that the collusive regime had lower standard deviation. We hypothesize that other conspiracies would exhibit this feature, along with a higher price, and design and apply a screen based on the coefficient of variation to retail gasoline stations in Louisville.

The size of potential conspiracies is not known, but researchers with industry experience have worked with small localized markets. Slade (1992) found fairly localized markets, and Hastings (2004) uses circular markets with radii of 0.5 and 1.5 miles. The U.S. Department of

Justice has prosecuted conspiracies ranging in size from two gasoline stations to those involving two to five jobbers and thirty to fifty stations.⁴ Due to the arguably localized nature of competition in retail gasoline, we hypothesize that a conspiracy would include stations located nearby to one another, and look for areas where stations exhibit low normalized standard deviation of price over time as an indicator of collusion.

Our retail gasoline price data comes from the Oil Price Information Service (OPIS). The data are generated from a sample of retail outlets that accept fleet cards.⁵ OPIS records the actual transaction price charged at the station on a given day. Hence, in principle, it is possible to create a panel data set consisting of specific stations' daily gasoline price. While the gasoline price data from OPIS is among the best available, a price is recorded for a specific station only if a purchase is made at that station; that is, if no one with a fleet card purchases gasoline at a station no price is recorded for that station on that day. In our data *no* single station has a complete time series of prices, and some stations have very few price quotes (e.g., fewer than one a week).⁶ Consequently, stations that sell more gasoline are more likely to be sampled on any given day. In addition, branded gasoline stations (which tend to charge higher prices) are more likely to accept fleet cards, and thus have fewer missing observations.

The OPIS data consists of the station specific price of regular grade gasoline, the brand of gasoline, and the station location. In this sample from the Louisville region, we use a total of 279 gas stations with incomplete daily retail gasoline prices from February 4, 1996 through August 2, 2002.⁷ To fill in the missing observations in our dataset we use multiple imputation, in particular, Markov chain Monte Carlo with Gibbs sampling combined with data augmentation. Details on this methodology can be found in Appendix A. Figure 2 in Appendix B provides a

⁴ This information is based on conversations with William Dillon of the U.S. Department of Justice.

⁵ Fleet cards are often used by firms whose employees drive a lot for business purposes, e.g., salesman or insurance claims adjusters. Fleet cards are often used to closely monitor what items employees charge to the firm, e.g., to ensure that an employee only bills fuel and not food when visiting a filling station.

⁶ Retail prices are reported for most weekdays with few exceptions. In 1998 and 1999 no retail prices are reported during the week of Thanksgiving (because of very small sample sizes).

⁷ According to New Image Marketing surveys there were 418 gasoline stations in the three Kentucky Counties that comprise Louisville in 1996 and 344 in 1999. Our sample of stations therefore represents between two-thirds and three-quarters of the gasoline stations.

map of the Louisville area with the main interstates and routes, as well as all the retail gasoline stations studied.

Our sample includes ten different brands of gasoline stations and another group of unbranded stations. The brands are Amoco, Ashland, BP, Chevron, Citgo, Dairy Mart, Marathon, Shell, Speedway and Super America.

When studying the price behavior of these stations, we should keep in mind that there are multiple vertical relationships between retail gasoline stations selling branded gasoline and the suppliers of branded gasoline. Some branded gasoline stations are owned and operated by a major oil company. These stations pay an unobserved wholesale transfer price. Other branded stations are owned by the major oil companies but are leased to an individual, called a “lessee dealer.” These stations pay a dealer tank wagon (DTW) price which can be station-specific or which may vary by zones. These zones are firm specific and are not publicly available. Monthly average state-wide DTW prices are available but are too aggregated to use to remove the effect of wholesale price on retail price variation. The company-owned-and-operated stations and the lessee-dealer stations are typically supplied directly by the company.

There are also branded stations that are owned by an individual who contracts with a major oil company to sell their brand of gasoline. These stations typically pay the rack price and are supplied and sometimes owned by “jobbers”. They are referred to as jobber-supplied stations. In our data, we have both direct-supplied stations that pay an internal transfer price as well as jobber-supplied stations that pay the rack price. It is thus possible that retail price variation could be driven by differing vertical relationships between brands and retail stations.

For each gasoline station, we compute the mean, the standard deviation and the coefficient of variation. Figure 3 in Appendix is a scatter plot of the standard deviation against the mean for all the stations in the data. As it is clear from this figure, stations with higher means tend to have higher variance. We are searching for “outliers” with high means and low standard deviations, which would be below and to the right of the rest of the data. As we can see from the figure, there is not much variation in the standard deviation, certainly not as much as there is between the competitive and collusive regimes in bid-rigging conspiracy described in section two.

Figure 4.A (mean), 4.B (standard deviation), and 4.C (coefficient of variation) in the Appendix present histograms for the gasoline stations in our sample. The histograms are used

to identify outliers. Four regions are chosen whose endpoints are represented by vertical lines, to identify the outliers for each statistic. “Average stations” lie in the two middle intervals at the center of the histograms where most of the mass is, while “outliers” are those on the tails of the histogram.⁸

We identify these outliers on the maps in Figures 5.A (mean), 5.B (standard deviation), and 5.C (coefficient of variation). The color red indicates a high outlier, and the color blue a low outlier. Thus, in Figure 5.A we look for red “clusters” indicating unusually high values of the mean; while in Figures 5.B and 5.C, we look for blue clusters indicating low standard deviations or low coefficients of variation.

Beginning with the coefficient of variation (Figure 5.C), we find no suspicious blue clusters of stations indicating an unusually low coefficient of variation. Likewise in Figure 5.B, there are no blue clusters of stations low standard deviation.

Some interesting patterns do emerge by looking at the geographic distribution of red stations, those with an unusually high mean. In Figure 5.A (mean), we see that with very few exceptions, the stations with the highest means are located on major roads and are without a very close competitor, like the stations marked in red on route 31, interstates 264, 64 and Shelbyville road (north and parallel to 64).

Figure 5.B maps the standard deviations for all retail stations. There is no clear pattern in this figure. Nevertheless, we notice that some of the highest standard deviations are along major roads such as interstates 64 and 65, route 31 and Shelbyville road. In the area where route 44 and interstate 65 meet, most stations have low price variation, as well as many of the stations located in Dixie Road. As previously discussed, this could be due to the fact that some stations are company owned and operated where pricing is centrally controlled (less volatile retail price), while for others pricing is decided by the local owner (more volatile retail price).

Overall, only some of the stations along Dixie Road have slightly higher means, lower standard deviations and lower coefficients of variation than most others in the area, but these differences are not very big, especially when measured against the large changes in variance estimated in the previous section. The coefficient of variation increased almost four and a half

⁸ Notice that the cutoff points chosen in Figures 4.A through 4.C are not going to influence the conclusions. This can be clearly observed from Figure 3. No matter how the intervals are divided, there are no stations in Figure 3 with distinctively high means and low standard deviations.

times from collusion to competition in the perch example. In the gasoline retail market, the highest value for the coefficient of variation is only about one and a half times higher than the lowest value. Therefore, we do not think the changes are big enough to suggest collusive behavior in Louisville in 1996-2002 .

4. Conclusion

We know far too little about how real cartels actually operate. Retrospective studies of cartel prosecutions, particularly when they result in the collapse of a cartel, allow researchers to compare a collusive regime to a competitive one to isolate the critical features of the conspiracy. In the collapse of our bid rigging conspiracy, we found a relatively small difference in price, but a huge difference in variance.

Empirical evidence like this can be used to test some of the hypotheses generated by the enormous amount of theoretical work in this area. It can also be used to guide policy. For example, there is very little evidence on how mergers affect the likelihood of collusion in a given industry. More retrospective studies might generate empirical regularities that could be used to tell enforcers both where to look for conspiracies, and which factors would hasten the collapse of a conspiracy or prevent a cartel from otherwise forming. This kind of information would also help merger enforcers determine when a merger might result in a cartel, or prevent a cartel from otherwise collapsing. Our understanding of how a merger results in “coordinated effects” is an area where economists have given enforcers very little guidance.

The failure of our screen to identify “suspicious” areas could mean only that the gas stations in Louisville are competing, but it could also indicate a failure of the screen to uncover pockets of existing collusion. Ideally, we would use a cartel collapse in the same industry to design an empirical screen for collusion. As more and better retrospective studies are done, we expect that more features of conspiracies will be uncovered that will help us design screens. Given the history of data screens to identify cartels, we remain uncertain whether screening for conspiracies is a good use of scarce enforcement resources.

References

Athey, Susan, Kyle Bagwell and Chris William Sanchiro, 2004, "Collusion and Price Rigidity," *Review of Economic Studies*, 71(2), pp. 317-349.

Bajari, Patrick, and Lixin, Ye, 2003, "Deciding Between Competition and Collusion," *Review of Economics and Statistics*, 85 (4), pp. 971-89.

Froeb, Luke and Mikhael Shor, 2000, "Auctions, Evidence, and Antitrust", forthcoming as a chapter in John Harkrider, ed., *Econometrics*, American Bar Association.

Froeb, Luke and Robert Koyak, 1994, "Measuring and Comparing smoothness in time series: The production smoothing hypothesis," *Journal of Econometrics*, 64, pp. 97-122.

Froeb, Luke M., Robert A. Koyak and Gregory J. Werden, 1993, "What is the Effect of Bid-Rigging on Prices?," *Economic Letters*, 42, pp. 419.

Geman, Donald and Geman, Stuart, 1984, "Stochastic relaxation, Gibbs distributions, and the Bayesian reconstruction of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6, pp. 721-741.

Geweke, John F., 1999, "Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication (with discussion and rejoinder)," *Econometric Reviews*, 18, pp. 1-126.

Geweke, John F., 1992, "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments," in Bernardo, J., Berger, J., Dawid, A. and Smith, A. (eds.), *Bayesian Statistics 4*, pp. 641-649. Oxford: Clarendon Press.

Geweke, John F., 1989, "Bayesian Inference in Econometric Models using Monte Carlo Integration," *Econometrica*, 57, pp. 1317-1340.

Harrington, Joseph E. and Joe Chen, 2004, "Cartel Pricing Dynamics with Cost Variability and Endogenous Buyer Detection," Working Paper, Johns Hopkins University.

Hastings, Justine, 2004, "Vertical Relationships and Competition in Retail Gasoline Markets: Empirical Evidence from Contract Changes in Southern California," *American Economic Review*, 94(1), pp. 317-328.

Koop, Gary, 2003, *Bayesian Econometrics*, Wiley Editions.

Porter, Robert H. and J. Douglas Zona, 1993, "Detection of Bid Rigging in Procurement Auctions," *Journal of Political Economy*, 101(3), pp. 518-38.

Slade, Margaret E. "Vancouver's Gasoline-Price Wars: An Empirical Exercise in Uncovering Supergame Strategies." *Review of Economic Studies*, Vol. 59 no. 2 (April 1992), pp. 257-276.

Tanner, M. A. and Wong, W. H., 1987, "The calculation of posterior distributions by data augmentation". *Journal of the American Statistical Association*, 82, pp. 528-550.

Taylor, Christopher and Jeffrey Fischer, 2003, "A Review of West Coast Pricing and the Impact of Regulations," *International Journal of the Economics of Business*, 10(2), pp. 225-243.

Appendix A

(A.1) Data Imputation, Markov Chain Monte Carlo, Gibbs Sampling and Data Augmentation

Data imputation “fills in” the missing data with predicted or simulated values. Popular types of imputation include mean substitution, simple hot deck and regression methods.

Mean substitution is appealing but corrupts the marginal distribution of the series we are trying to complete, Z , as well as its covariances and correlations with other variables. In simple hot deck, each missing value is replaced with a randomly drawn observed value. It is similar to bootstrap, preserves the marginal distribution of Z , but distorts correlations, and it is appropriate only for univariate analysis. Regression methods replace missing values with either the predicted values from a regression model, or with the predicted values plus random residuals. They become difficult in a multivariate context when more than one variable has missing values, such as in our case.

An additional problem of this type of methods is that, whenever the missing data are replaced by one set of imputed values, the analysis that follows from the use of the complete dataset does not reflect missing-data uncertainty: the sample size N is overstated, the confidence intervals are too narrow, and the type I error rates are too high. The problem becomes worse as the rate of missing observations and the number of parameters increase. For all above reasons, in our context the proper data imputation procedure to use is multiple imputation.

Multiple imputation is a simulation-based approach to the analysis of incomplete data. The procedure is to (1) create imputations (*i.e.*, replace each missing observation with $m > 1$ simulated values), (2) analyze each of the m datasets in an identical fashion and (3) combine the results. Some main advantages of multiple imputation in relation to single imputation ($m=1$) are that (i) the final inferences incorporate missing-data uncertainty, and (ii) it is highly efficient even for small m .

After creating imputations, the analysis of the imputed datasets stores m sets of point estimates and their standard errors. There are several methods to combine the results at the

end such as combining vector estimates and covariance matrices, combining p-values and combining likelihood-ratio test statistics.

Multiple imputations are created through simulated draws from a predictive distribution of the missing data. In general, multiple imputations are drawn from a Bayesian predictive distribution

$$p(z^m, \theta | z^o) = \int p(z^m | z^o, \theta) p(\theta | z^o) d\theta, \quad (1)$$

with the knowns denoted by z^o (observed) and θ (parameters), while z^m stands for the unknown or missing observations.

Predictive distributions of the missing data are usually intractable, and special computation methods are needed. One such method is Markov chain Monte Carlo, an iterative method for drawing from intractable distributions. A Markov chain is created so that it converges to the desired target. That can be accomplished through Gibbs sampling, the Metropolis-Hasting algorithm, data augmentation, or a combination of these methods. In our multiple imputation of the retail stations, the Markov chain Monte Carlo method used is Gibbs sampling combined with data augmentation.

The principle behind Gibbs sampling is simple. Let θ denote the vector of parameters, $p(z|\theta)$ is the likelihood, $p(\theta)$ is the prior distribution and $p(\theta|z)$ is the posterior distribution. We are interested in the numerical approximation of $E[g(\theta)|z]$ for any function of interest $g(\theta)$. Partition θ into various blocks as $\theta = (\theta'_{(1)}, \theta'_{(2)}, \dots, \theta'_{(B)})$, where $\theta_{(j)}$ is a scalar or vector, $j = 1, 2, \dots, B$. In many models, it is not easy to directly draw from $p(\theta|z)$, but many times it is easy to randomly draw from $p(\theta_{(1)} | z, \theta_{(2)}, \dots, \theta_{(B)})$, $p(\theta_{(2)} | z, \theta_{(1)}, \theta_{(3)}, \dots, \theta_{(B)})$, up to $p(\theta_{(B)} | z, \theta_{(1)}, \theta_{(2)}, \dots, \theta_{(B-1)})$. These distributions are known as full conditional posterior distributions, since they define a posterior for each block conditional on all the other blocks. Drawing from the full conditionals will yield a sequence $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(s)}$ which can be averaged out to produce estimates of $E[g(\theta)|z]$. It should be

noted that the state of the Gibbs sampler at draw s (i.e., $\theta^{(s)}$) depends on its state at draw $s-1$ (i.e., $\theta^{(s-1)}$), meaning that the sequence is a Markov chain.

In our approach we mixed Gibbs sampling with data augmentation. Data augmentation alternately simulates missing data and parameters, creating a Markov chain that converges to the correct predictive distribution. More specifically, given starting values it draws the missing observations at time $t+1$, z_{t+1}^m , from $p(z^m | z^o, \theta_t)$ and then draws θ_{t+1} from $p(\theta | z^o, z_{t+1}^m)$. This procedure generates the Markov chain $z_1^m, \theta_1, z_2^m, \theta_2, \dots$, which converges in distribution to $p(z^m, \theta | z^o)$. The rate of convergence depends on the fraction of missing information. Multiple runs (indexed by subscript m) of this imputation technique will collect iterates of z^m : $z_{1t}^m, z_{2t}^m, z_{3t}^m, \dots, z_{mt}^m$. In data augmentation it is important to choose t large enough, so that the successive imputations are statistically independent. This can be monitored through time series plots and autocorrelation functions of θ .

(A.2) Interpolation of Missing Retail Price Data

As previously explained, we use full Bayesian imputation to randomly assign values for the missing observations. Once this is done, analysis can proceed as if there were no missing observations. To summarize the strategy, let \mathbf{z} denote all prices for a retail station, and decompose this vector as $\mathbf{z}' = (\mathbf{z}^{o'}, \mathbf{z}^{m'})$ where o indicates observed values and m indicates missing values. Let θ denote the vector of unknown parameters in an econometric model for retail prices, to be described shortly. The model specifies the distribution of retail prices, $p(\mathbf{z} | \theta)$, as well as a prior distribution for the unknown parameters, $p(\theta)$. Then the distribution of the unknowns, θ and \mathbf{z}^m , conditional on the knowns, \mathbf{z}^o , is the predictive distribution

$$p(\mathbf{z}^m, \theta | \mathbf{z}^o) = p(\mathbf{z}^m, \mathbf{z}^o, \theta) / p(\mathbf{z}^o) \propto p(\mathbf{z}^m, \mathbf{z}^o, \theta) = p(\mathbf{z}, \theta) = p(\theta) p(\mathbf{z} | \theta). \quad (2)$$

We use Markov chain Monte Carlo methods to sample from this distribution, discarding the values of θ and keeping the values of \mathbf{z}^m . The latter are the imputed missing values.

In particular, let y_t denote an overall market price on day t ; there are T days in the sample. In practice, y_t is computed as the average price taken over all stations reporting data for day t . If no stations report data on day t , then y_t is linearly interpolated for that day.

Let y_{it} denote price at station i on day t , and let $z_{it} = y_{it} - y_t$. There are n stations in the sample. Interpolation is based on the stationary first-order autoregressive model

$$z_{it} - \mu_i = \rho_i (z_{i,t-1} - \mu_i) + \varepsilon_{it}.$$

The shocks ε_{it} are normal and independently distributed across stations and days, and identically distributed across days for each station: $\varepsilon_{it} \stackrel{iid}{\sim} N(0, \sigma_i^2)$. The stationary initial condition is $z_{it} - \mu_i \sim N[0, \sigma_i^2 / (1 - \rho_i^2)]$.

The model permits a station to have prices that tend to be higher ($\mu_i > 0$) or lower ($\mu_i < 0$) than average. On any given day there is a tendency for price at station i to return to its equilibrium value $y_t + \mu_i$. The speed of adjustment is governed by ρ_i . For the stations in our sample ρ_i averages about 0.9; some stations are lower, others higher. On any given day there is also a shock ε_{it} that perturbs prices. The standard deviation of this shock varies across stations, with a typical station having a standard deviation of about 0.015.

If we were able to fully observe all T prices for station i , then

$$p(z_{i1}, \dots, z_{iT} | \mu_i, \rho_i, \sigma_i^2) = (2\pi)^{-T/2} (\sigma_i^2)^{-T/2} (1 - \rho_i^2)^{-1/2} \cdot \exp\left\{-\left\{(1 - \rho_i^2)(z_{i1} - \mu_i)^2 + \sum_{t=2}^T [z_{it} - \mu_i - \rho_i(z_{i,t-1} - \mu_i)]^2\right\} / 2\sigma_i^2\right\}. \quad (3)$$

This constitutes the density $p(\mathbf{z}|\theta)$ in the summary of our strategy given in (2). The unknown parameters constitute $\theta' = (\mu_i, \sigma_i^2, \rho_i)$. (Note that our analysis takes place one station at a time, because conditional on $y_t (t=1, \dots, T)$ and θ , station prices are independent of one another.) The prior distributions of the three parameters are independent and uninformative: there is a flat, improper prior distribution for μ_i on the real line; a flat, improper prior distribution for σ_i^2 on the positive half-line; and a flat, proper prior distribution for ρ_i on the interval $(-1, 1)$. Thus the density from which we wish to sample is (3), but with the restriction that $\rho_i \in (-1, 1)$.

The Markov chain Monte Carlo method is Gibbs sampling. Each of the three parameters and each of the missing prices are drawn, in succession, from the conditional distribution implicit in (3). Two of the conditional posterior parameter distributions are straightforward:

$$\mu_i | [\sigma_i^2, \rho_i, z_{i1}, \dots, z_{iT}] \sim N \left[\frac{(1-\rho_i) \sum_{t=2}^T (z_{it} - \rho_i z_{i,t-1}) + (1-\rho_i^2) z_{i1}}{(T-1)(1-\rho_i)^2 + (1-\rho_i^2)}, \frac{\sigma_i^2}{(T-1)(1-\rho_i)^2 + (1-\rho_i^2)} \right], \quad (4)$$

and

$$\left\{ (1-\rho_i^2)(z_{i1} - \mu_i)^2 + \sum_{t=2}^T [z_{it} - \mu_i - \rho_i(z_{i,t-1} - \mu_i)]^2 \right\} / \sigma_i^2 \sim \chi^2(T). \quad (5)$$

From (3),

$$p(\rho_i | \mu_i, \sigma_i^2, z_{i1}, \dots, z_{iT}) \propto (1-\rho_i^2)^{1/2} \quad (6)$$

$$\cdot \exp \left\langle - \left\{ (1-\rho_i^2)(z_{i1} - \mu_i)^2 + \sum_{t=2}^T [z_{it} - \mu_i - \rho_i(z_{i,t-1} - \mu_i)]^2 \right\} / 2\sigma_i^2 \right\rangle. \quad (7)$$

The function (6)-(7) is not the kernel of a known distribution, but (7) alone is the kernel of the distribution

$$\rho_i \sim N\left[\sum_{t=2}^T (z_{i,t-1} - \mu_i)(z_{it} - \mu_i) / \sum_{t=1}^T (z_{it} - \mu_i)^2, \sigma_i^2 / \sum_{t=1}^T (z_{it} - \mu_i)^2\right]. \quad (8)$$

Thus we employ a Metropolis within Gibbs step: draw a candidate ρ_i^* from (8). If $\rho_i^* \notin (-1, 1)$ then we reject the candidate and keep the previous value of ρ_i . Otherwise, accept ρ_i^* with probability

$$\min\left\{\left[\frac{(1 - \rho_i^{*2})}{(1 - \rho_i^2)}\right]^{1/2}, 1\right\}.$$

The missing prices are drawn one at a time. The price y_{is} is missing if and only if z_{is} is missing. If $2 \leq s \leq T-1$ then from (3)

$$\begin{aligned} p\left[z_{is} \mid z_{it} (t \neq s), \mu_i, \sigma_i^2, \rho_i\right] &= p\left(z_{is} \mid z_{i,s-1} z_{i,s+1}, \mu_i, \sigma_i^2, \rho_i\right) \\ &\propto \exp\left\{-\left\{\left[z_{is} - \mu_i - \rho_i (z_{i,s-1} - \mu_i)\right]^2 + \left[z_{i,s+1} - \mu_i - \rho_i (z_{is} - \mu_i)\right]^2\right\}\right\} \end{aligned}$$

and hence

$$z_{is} \left[z_{it} (t \neq s), \mu_i, \sigma_i^2, \rho_i \right] \sim N\left[\mu_i + \frac{\rho_i (z_{i,s-1} + z_{i,s+1} - 2\mu_i)}{1 + \rho_i^2}, \frac{\sigma_i^2}{1 + \rho_i^2}\right]. \quad (9)$$

Similarly

$$z_{i1} \left[z_{it} (t \geq 2), \mu_i, \sigma_i^2, \rho_i \right] \sim N\left[\mu_i + \rho_i (z_{i2} - \mu_i), \sigma_i^2\right], \quad (10)$$

$$z_{iT} \left[z_{it} (t \leq T-1), \mu_i, \sigma_i^2, \rho_i \right] \sim N\left[\mu_i + \rho_i (z_{i,T-1} - \mu_i), \sigma_i^2\right]. \quad (11)$$

Of course, $y_{is} = y_s + z_{is}$. To sample the unknown parameters and missing data from (2), the Markov chain Monte Carlo algorithm samples, in turn, from (4), (5), (8) followed by the acceptance-rejection step, (10), (9) for $s = 2, \dots, T-2$ and (11). The algorithm is initialized

with $\mu_i = 0$, $\rho_i = 0.9$, $\sigma_i^2 = 0.015^2$, and $z_{i_s} = 0$ for all missing data. Convergence is essentially instantaneous, even for stations for which a substantial fraction of observations are missing. Our interpolation is based on ten iterations of the MCMC algorithm, which requires only a few minutes of computing time for the full sample of $n = 279$ retail stations and $T = 2,371$ days.

Appendix B

Figure 2
Louisville Gasoline Retail Stations




Figure 3

Means and Standard Deviations for all Stations

All brands, 279 stations




Figure 4.A
Means Histogram




Figure 4.B
Standard Deviations Histogram




Figure 4.C
Coefficients of Variation Histogram




Figure 5.A

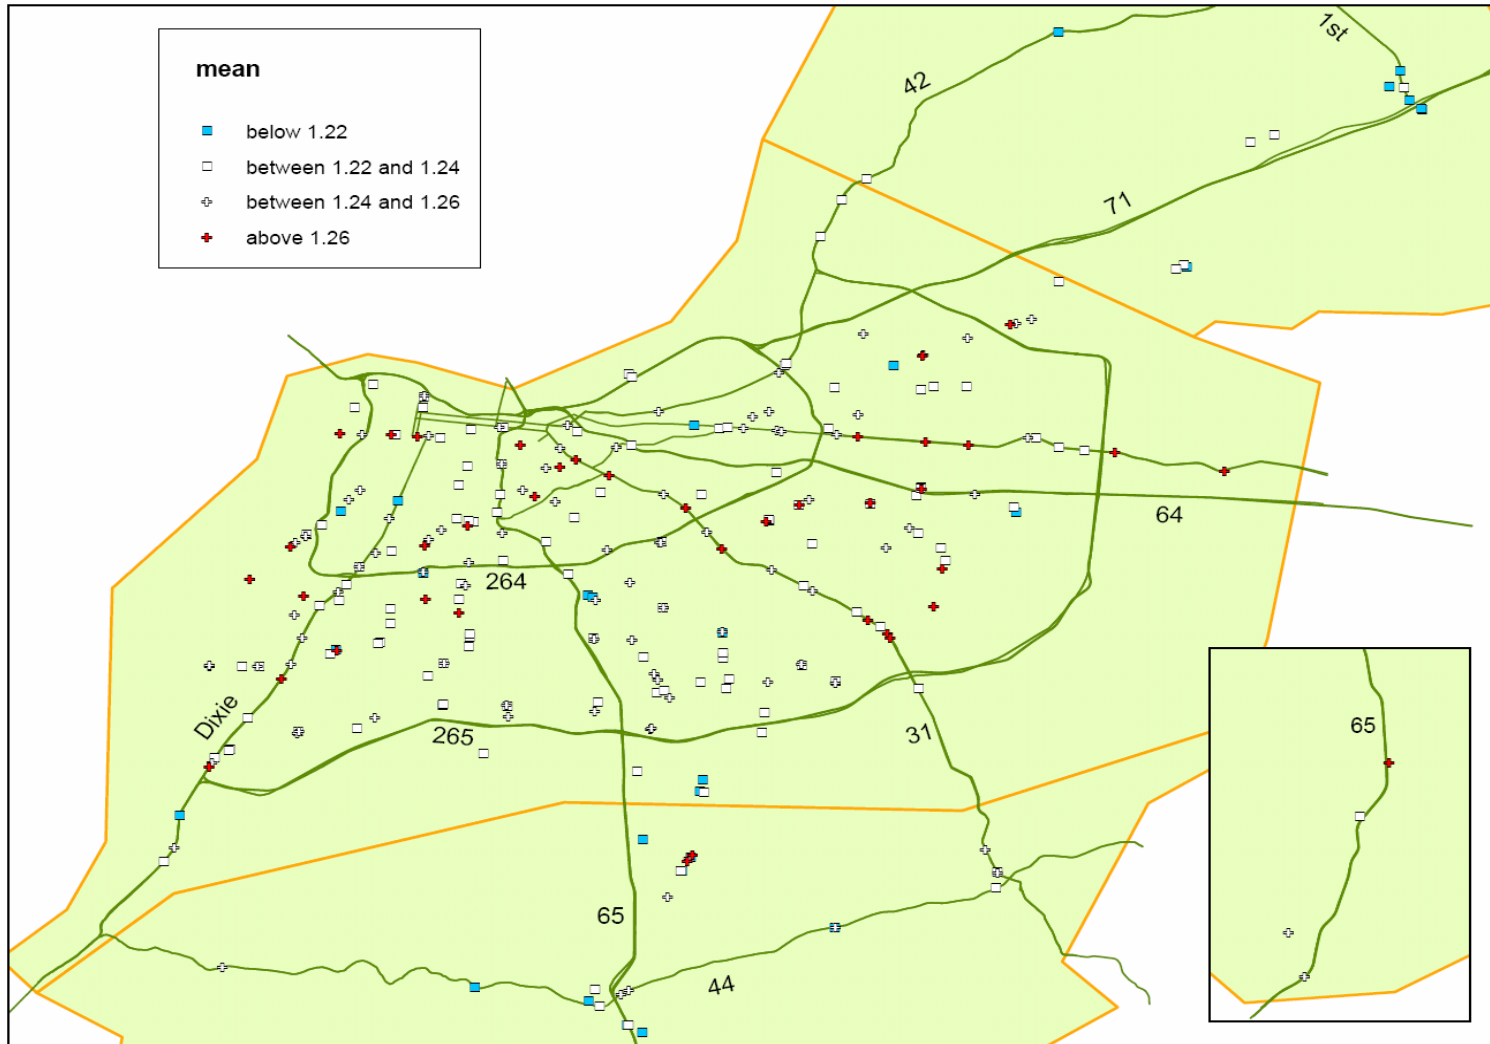


Figure 5.B

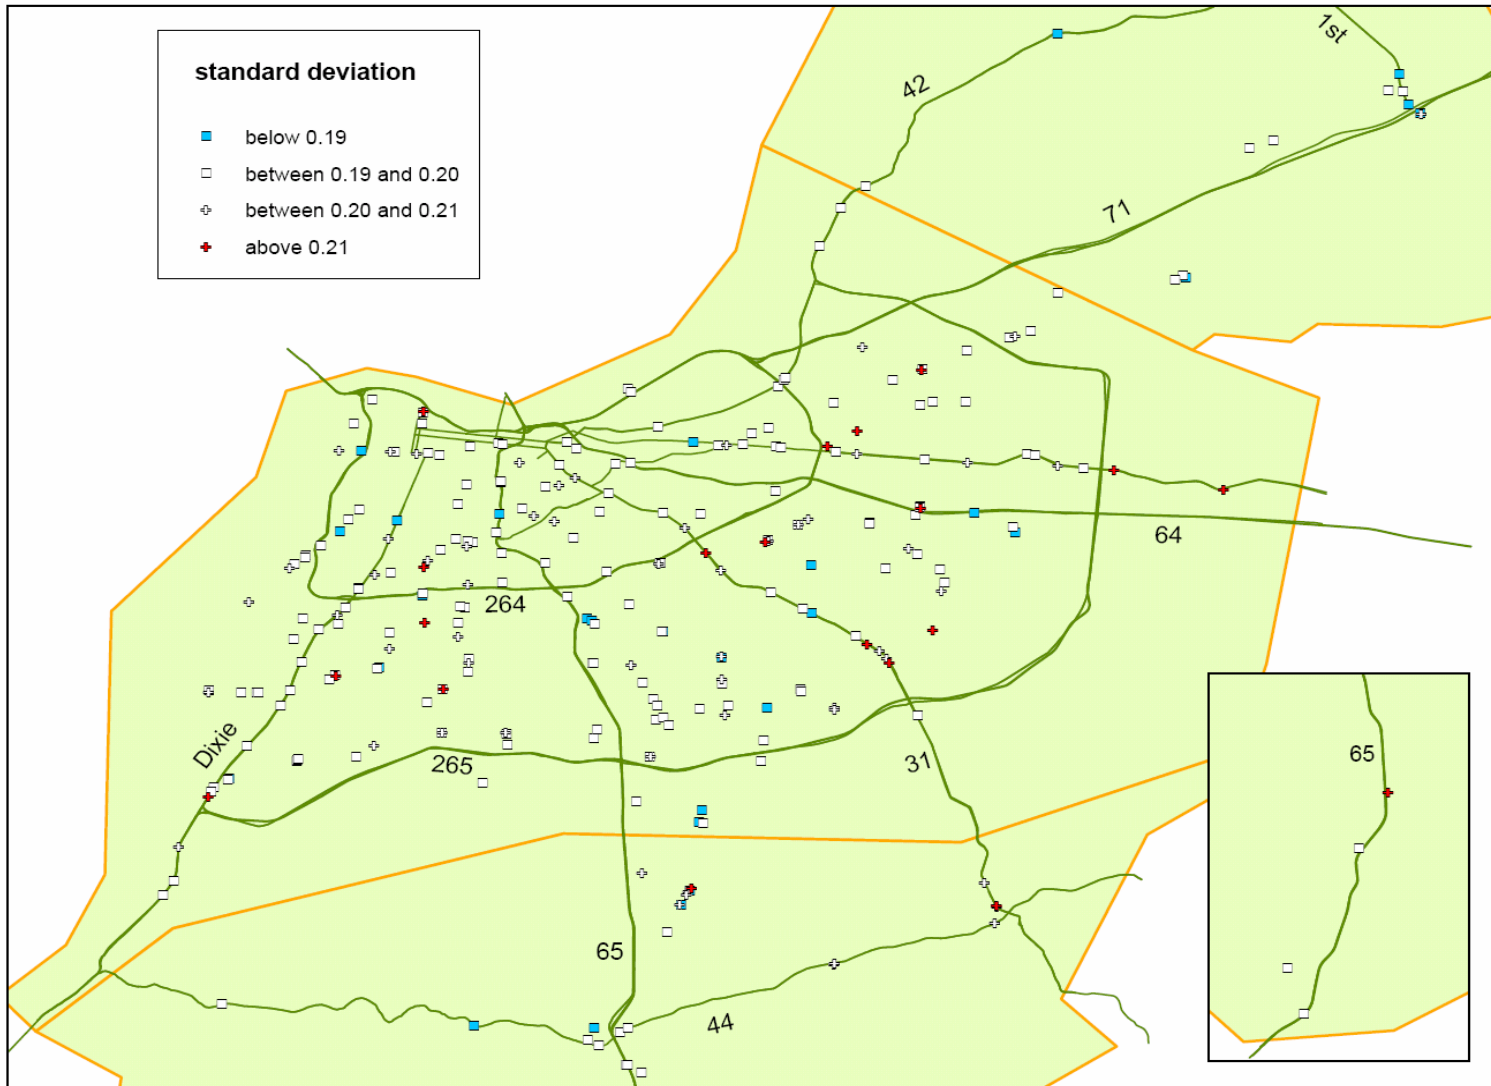


Figure 5.C

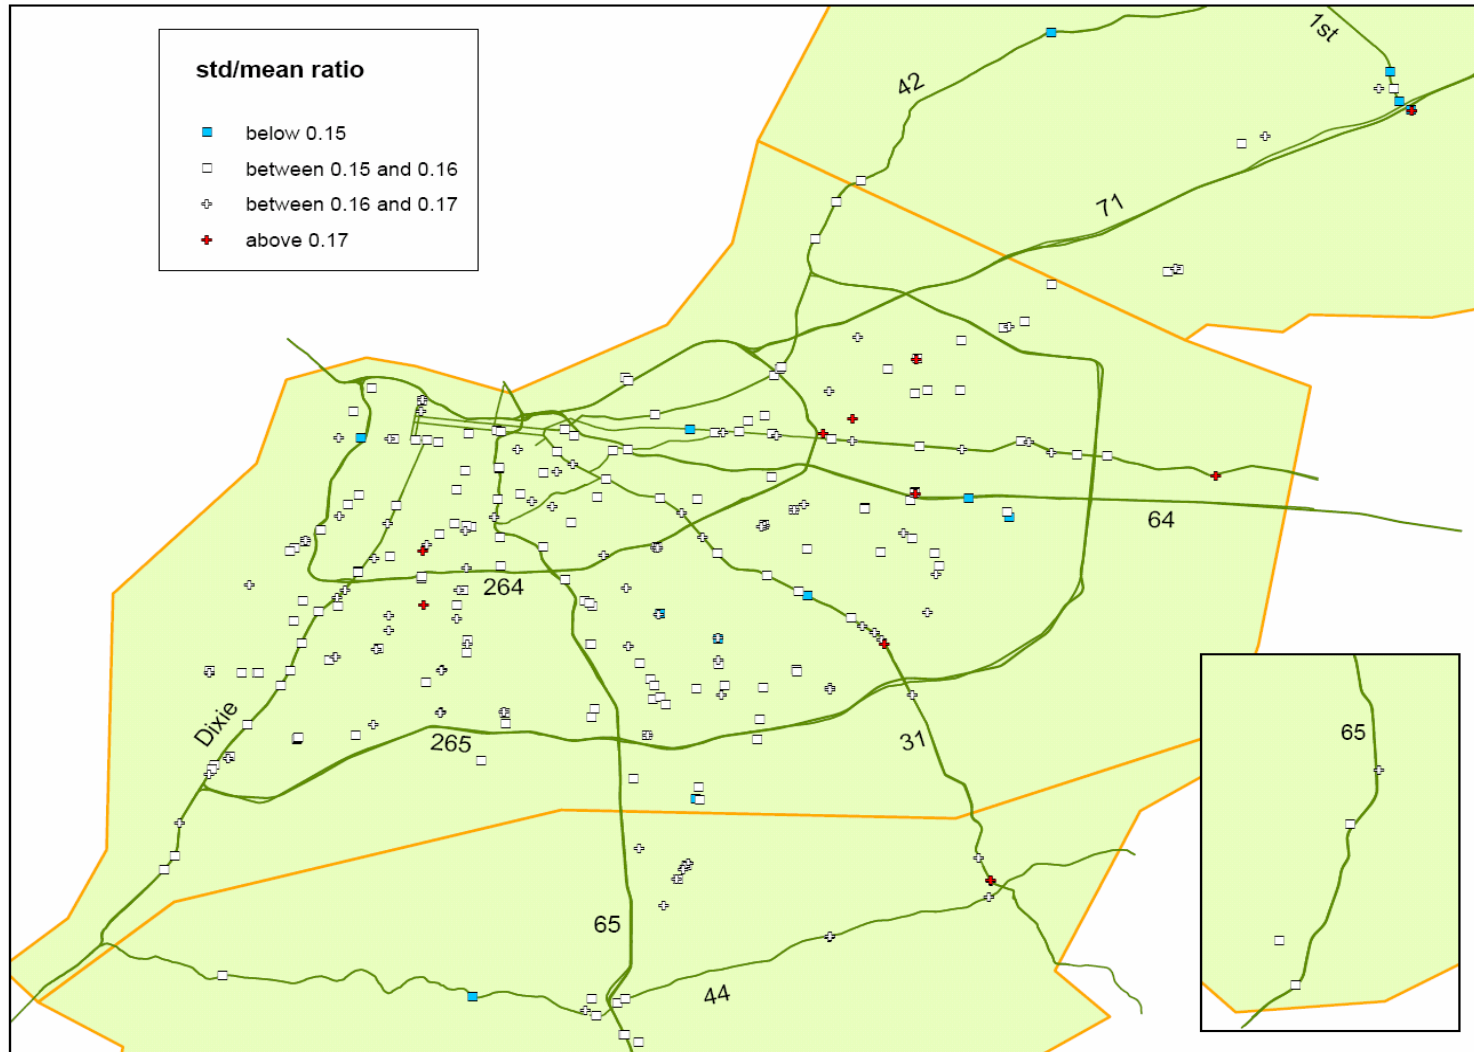


Figure 6

