Understanding the Price Effects of the MillerCoors Joint Venture

Supplementary Materials*

Nathan H. Miller† Georgetown University

Matthew C. Weinberg‡ Drexel University

March 3, 2017

This documents summarizes a number of analyses and results that were impossible to include in the article and the published online applendix due to space constraints. We begin with Section "F" so the document can be read as a continuation of the online appendix.

F The Nesting Parameter

The nesting parameter plays an important role in the demand system because it helps determine the degree of substitution between the inside goods and the outside good. To develop intuition, consider a standard logit model (i.e., the RCNL with $\rho = \Pi = 0$). The outside good share scales with the market size assumption and, because substitution is proportional to share, so too does diversion to the outside good. The nesting parameter softens this connection and allows for more realistic substitution patterns. Table F.1 provides a number of robustness checks. First, we reestimate demand under alternative market sizes. Recall that the baseline (region-specific) market size equals 150% of the maximum observed sales. Columns (i) and (ii) show results based on scaling observed sales by 130% and 200%, respectively, to obtain the market size. The magnitude of the ρ estimate changes accordingly

*All estimates and analyses in this paper based on SymphonyIRI Group, Incorporated data are by the authors and not by SymphonyIRI Group, Inc.

†Georgetown University, McDonough School of Business, 37th and O Streets NW, Washington DC 20057. Email: nathan.miller@georgetown.edu.

‡Drexel University, Gerri C. LeBow Hall, 3220 Market Street, Philadelphia PA 19104. Email: mcw325@drexel.edu.

Table F.1: Performance of the Nested Logit Parameter

ρ Estimated:		Yes	Yes	No	No
Market Size Scaler:		1.30	2.00	1.50	1.50
Variable	Parameter	(i)	(ii)	(iii)	(iv)
		Demand Estimates			
Price	α	-0.1101	-0.0933	-0.1550	-0.1961
		(0.0166)	(0.0129)	(0.0240)	(0.0314)
Nesting Parameter	ρ	0.7752	0.8407	0.70	0.60
		(0.0489)	(0.0358)		
Income×Price	Π_1	0.0010	0.0006	0.0014	0.0019
		(0.0003)	(0.0002)	(0.0002)	(0.0003)
Income×Constant	Π_2	0.0132	0.0129	0.0055	0.0018
		(0.0063)	(0.0041)	(0.0054)	(0.0060)
Income×Calories	Π_3	0.0056	0.0030	0.0058	0.0063
		(0.0018)	(0.0014)	(0.0022)	(0.0026)
		Derived Demand Statistics			
Median Own Price Elasticity		-4.30	-5.35	-4.53	-4.15
Median Market Price Elasticity		-0.68	-0.69	-1.02	-1.25
Median Outside Diversion		16.19%	13.29%	23.24%	31.06%
$J-$Statistic		14.04	14.27	10.27	10.13

Notes: The table shows results for the RCNL-2 specification under various normalizations. There are 31,784 observations at the brand-size-region-quarter-year level. The sample excludes the quarters between June 2008 and May 2009. All regressions include product (brand×size) and time fixed effects. The elasticity and diversion numbers represent medians among all the brand-size-region-month/quarter-year observations. Standard errors are clustered by region and shown in parentheses.

such that diversion to the outside good remains relatively constant. Demand is somewhat more elastic with larger market sizes, but we have confirmed that this does not meaningfully affect inferences on the supply-side. Columns (iii) and (iv) are estimated with a normalization that the nesting parameter is 0.70 and 0.60, respectively. We exclude instruments based on the number of products and summed distance, which are no longer needed for identification. As expected, lower values of ρ correspond to greater diversion to the outside good, but otherwise the demand results are mostly unchanged. Considered together, these results indicate that the nesting parameter responds as it should to the market size, and that the results hold over a plausible range of the nesting parameter.

G Validity of the Demand-Side Instruments

We now discuss some possible concerns about the validity of the instruments. We start with the post-merger indicator for ABI and MillerCoors products. Including this in the instrument set imposes the assumption that the changes in unobserved demand for ABI/MillerCoors, before versus after the merger, are not systematically different from the changes in the unobserved demand for Modelo/Heineken. The validity of the instrument could come into question if the model understates the effect of the recession on preferences for (cheaper) domestic beer. While this is difficult to rule out entirely, the income interactions that we employ in the demand specifications allow the recession to impact demand in a sensible manner, and there is little else to be done empirically.

The number of products primarily serves to identify the nesting parameter. Variation arises because not all products are sold in each region-period. Given the structure of the data (39 regions, 39 products, 72 months) there are 109,512 possible product-region-period observations of which 94,565 (or 86.43%) are realized. Because the more popular products are sold in every region-period, variation arises due to the entry and exit of the less popular brands. This could raise concern if sales are sufficiently thin that entry/exit reflects measurement error in the data rather than true underlying variation in conditions. However, this does not seem to be the case. Consider the case of Michelob 24-packs, which are observed in 956 of the 2808 month-year combinations. Figure G.1 plots whether the product was observed in each month, by city. There are some instances in which the empirical patterns could be driven by measurement error (e.g., Boston), but more commonly the existence of sales is persistent. Analyses of other products are similar.

H Unobserved Trends

The baseline specifications incorporate time fixed effect to account for changes in cost and demand conditions that affect all products. Bias could arise to the extent that unobserved changes in the market affect products differentially. Of particular concern would be unobserved changes that affect ABI differently than Modelo/Heineken, which would undermine our key identifying restrictions. One possible source of bias is the (gradual) increase in consumption of craft beers that occurs through the sample period. It is possible that craft beer is a closer competitor to imports than to ABI/MillerCoors, on the basis that craft and imported beers tend to have higher prices and sell in smaller package sizes.

As a robustness check, we reestimate the RCNL-2 specification using the shorter time windows of 2008-2009 and 2007-2010. We still exclude a year of data after the consummation of the merger, so the 2008-2009 sample contains four quarters of data and the 2007-2010 sample contains 12 quarters. The influence of unobserved trends should be smaller with

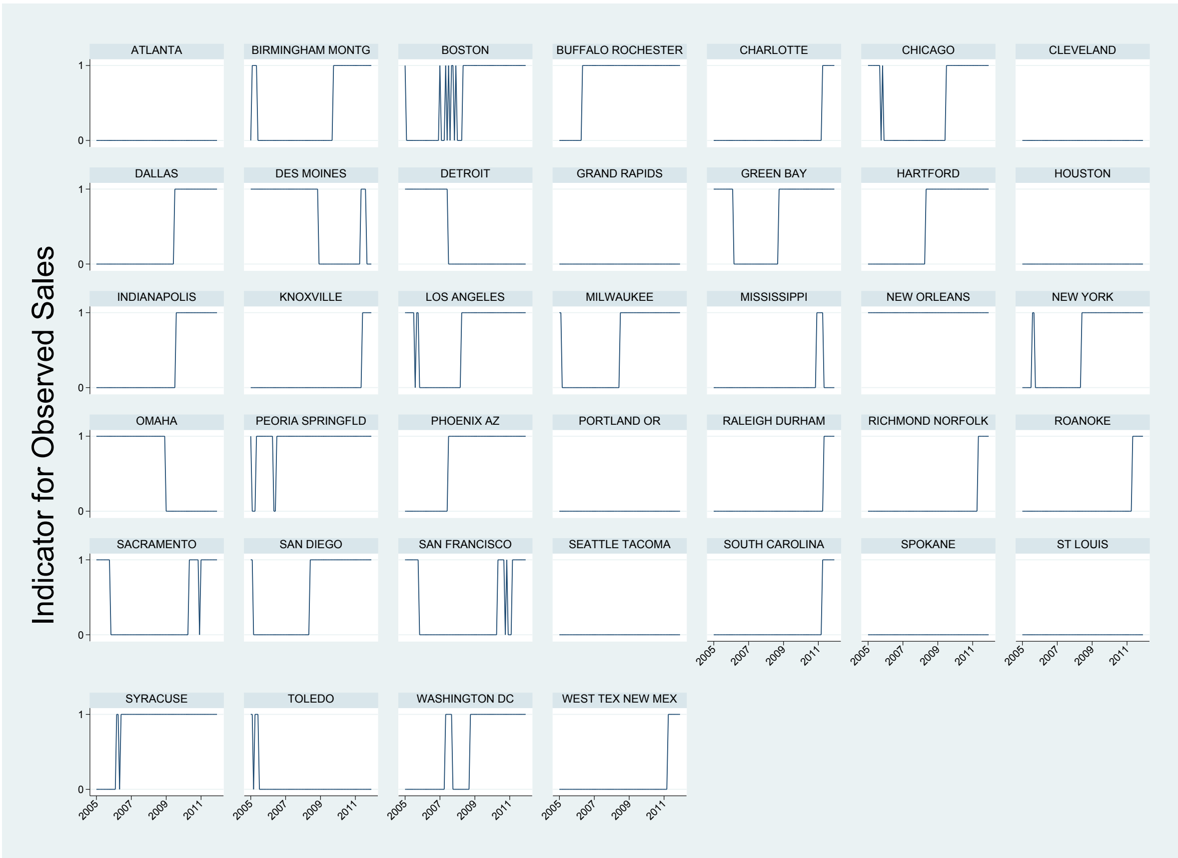

Figure G.1: Region-Periods with Observed Sales of Michelob 24-Packs

these restrictions because the window of time over which the trends operate is abbreviated. The downside is that discarding data could reduce the precision of the estimates. Table H.1 shows the results. The coefficients are similar to those obtained from full dataset, though the Income×Price parameter is somewhat smaller and no longer statistically significant. The median price elasticities increase from -4.33 with the full dataset to -5.75 and -5.64, and diversion to the outside good also is slightly lower with the shorter time periods. The supply-side estimates of the κ parameter remain positive and statistically different than zero. The somewhat smaller value of κ that is obtained with the 2008-2009 sample is consistent with results presented previously that coordination may have strengthened over the post-merger periods (e.g., recall Figure 3). The marginal cost parameters are not shown in the table, but are quite similar to those obtained from the full dataset. Overall, the analysis supports that the main econometric results are not unduly influenced by unobserved trends.

For a second check on changing demand conditions, we estimate the demand-side of the model separately with the pre-merger and the post-merger data. The post-merger ABI

Table H.1: Estimation with Short Samples

Data Sample:		2008-2009	2007-2010
Variable	Parameter	(i)	(ii)
	Demand Estimates		
Price	α	-0.0837	-0.0946
		(0.0160)	(0.0127)
Nesting Parameter	ρ	0.8536	0.8437
		(0.0373)	(0.0391)
Income×Price	Π_1	0.0003	0.0007
		(0.0003)	(0.0003)
Income×Constant	Π_2	0.0145	0.0114
		(0.0062)	(0.0055)
Income×Calories	Π_3	0.0063	0.0048
		(0.0021)	(0.0017)
	Derived Demand Statistics		
Median Own Price Elasticity		-5.75	-5.64
Median Market Price Elasticity		-0.61	-0.65
Median Outside Diversion		10.91%	11.79%
$J-$Statistic		10.95	13.30
	Selected Supply Estimates		
Post-Merger Internalization		0.1565	0.2449
of Coalition Pricing		(0.0698)	(0.1185)

Notes: The table shows results for the RCNL-2 specification using different sample windows. There are 5,409 observations in the 2008-2009 sample and 16,143 observations in the 2007-2010 sample, at the brand-size-region-quarter-year level. The samples exclude the quarters between June 2008 and May 2009. All regressions include product (brand×size) and time fixed effects. The supply-side also includes the baseline marginal cost shifters and region fixed effects. The elasticity and diversion numbers represent medians among all the brand-size-region-month/quarter-year observations. Standard errors are clustered by region and shown in parentheses.

indicator that we use as an instrument for the price coefficient is no longer applicable, so there is a loss of power in the first stage. This creates numerical problems with the RCNL. Thus, we generate results using the baseline nested logit specification, which can be estimated with 2SLS. The results are in Table H.2. The parameter estimates are similar in the two time periods. The price coefficient is larger than what is obtained with the full sample, implying more elastic demand. It also is less precisely estimated. Regressions that use the full sample but exclude the post-merger ABI indicator from the instrument set obtain similar results. While we cannot estimate the supply-side of the model with only pre-merger or post-merger data, from the numerical experiments in the published online appendix we know that more elastic demand generates larger estimates of the κ parameter, all else equal.

Table H.2: Estimation with Pre-Merger and Post-Merger Data

Data Sample:		2005-2008	2009-2011
Variable	Parameter	(i)	(ii)
	Demand Estimates		
Price	α	-0.2250	-0.2624
		(0.1986)	(0.2672)
Nesting Parameter	ρ	0.6084	0.5497
		(0.1472)	(0.1177)
	Derived Demand Statistics		
Median Own Price Elasticity		-6.06	-6.43
Median Market Price Elasticity		-1.84	-2.27
Median Outside Diversion		31.35%	36.57%

Notes: The table shows results for the NL-1 specification based on pre-merger data (column (i)) and post-merger data (column (ii)). There are 53,342 observations in the 2005-2008 sample and 41,314 observations in the 2009-2011 sample, at the brand-size-region-month-year level. All regressions include product (brand$\times$size) and time fixed effects. The elasticity and diversion numbers represent medians among all the brand-size-region-month/quarter-year observations. Standard errors are clustered by region and shown in parentheses.

Lastly, Table H.3 provides estimates of an adjusted supply-side model that incorporates bilateral nominal exchange rates (domestic-currency units per unit of foreign currency) into the marginal cost function. This allows us to determine whether exchange rate fluctuations, which affect the relative costs of ABI and Modelo/Heineken, impact estimates of κ parameter. We show results for each of the demand specifications in Table H.3. The exchange rate parameter ranges from -0.041 to 0.760 in the RCNL specifications, and the other estimates are largely unaffected. We exclude the exchange rates from the baseline specifications because structural interpretation is difficult without decomposing the costs of Modelo/Heineken into a local non-traded component due to U.S. distribution and retail and a nonlocal component due to foreign production.

I Additional Estimation Results

Figure I.1 plots the marginal cost and demand time fixed effects in the RCNL-1 specification. The demand fixed effects are seasonal, reflecting that demand for beer is strongest in the summer and weakest in the winter. There is a gradual reduction in the willingness-to-pay for beer over the sample period. Among other considerations, this could reflect the trend toward microbreweries (which are included in the outside good). The marginal cost fixed

Table H.3: Exchange Rates in the Marginal Cost Functions

Demand Model:	NL-1	RCNL-1	RCNL-2	RCNL-3	RCNL-4
Data Frequency:	monthly	monthly	quarterly	monthly	quarterly
	(i)	(ii)	(iii)	(iv)	(v)
Post-Merger Internalization	0.378	0.2654	0.249	0.289	0.345
of Coalition Pricing Externalities	(0.031)	(0.076)	(0.091)	(0.043)	(0.054)
			Marginal Cost Parameters		
MillerCoors×PostMerger	-0.608	-0.654	-0.649	-0.722	-0.526
	(0.040)	(0.050)	(0.056)	(0.042)	(0.041)
Distance	0.143	0.168	0.163	0.169	0.148
	(0.046)	(0.059)	(0.059)	(0.060)	(0.048)
Exchange Rate	0.756	0.389	-0.041	0.760	0.579
	(0.188)	(0.181)	(0.266)	(0.190)	(0.223)

Notes: The table shows the supply results with the exchange rate added to the marginal cost function. Estimation is with the method-of-moments. There are 94,656 observations at the brand-size-region-month-year level in columns (i), (iii), and (iv), and 31,784 observations at the brand-size-region-year-quarter level in column (ii). The sample excludes the months/quarters between June 2008 and May 2009. All regressions include product (brand×size) and time fixed effects. Standard errors are clustered by region and shown in parentheses.

effects decrease in the pre-merger periods, stabilize over 2009-2010, and then drop again 2011. We suspect that the observed import price decreases over 2009-2010 can be explained from changes in macroeconomic conditions, but that the import price decreases in 2011 load onto the cost fixed effects.

Table I.2 provides a histogram of the marginal cost region fixed effects in the RCNL-1 specification. The fixed effect of Atlanta is normalized to zero. As shown, most region fixed effects range over $(-1, 0)$, with a few regions having lower costs on average. It is possible that these cost fixed effects are picking up some demand-side variability because we do not incorporate region fixed effects on the demand-side (otherwise the RCNL nesting parameter is too difficult to identify with any precision).

Figure I.3 provides a histogram of the consumer-specific income coefficients obtained with the RCNL-1 specification. The histogram aggregates together all the income draws. With 39 regions, seven years, and 500 draws per region-year, this amounts to a total of 136,500 draws. We subtract the global mean from all the income draws prior to estimation. Thus, the price coefficient of -0.0887 that is reported in Table 4 represents the average effect of price on indirect utility. The 25th and 75th percentiles in the distribution of consumer-specific price coefficients are -0.1039 and -0.0772, respectively. The minimum and maximum are -0.1138 and -0.0129. The left tail of the distribution reflects that we

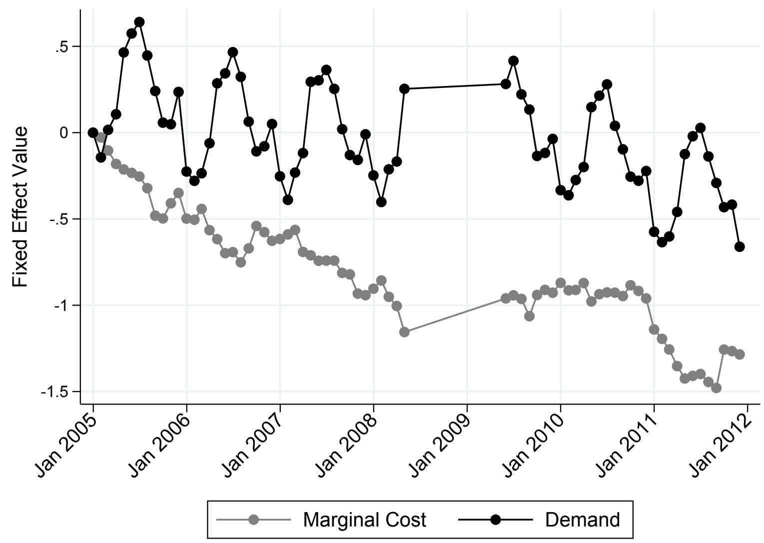

Figure I.1: Marginal Cost and Demand Time Fixed Effects in RCNL-1

Notes: The figure plots the estimated time fixed effects that affect preferences for the inside good (on the demand-side) and the marginal costs of all products (on the supply-side). We divide the demand fixed effects by the absolute value of the price coefficient, prior to plotting, so that the units are in dollars.

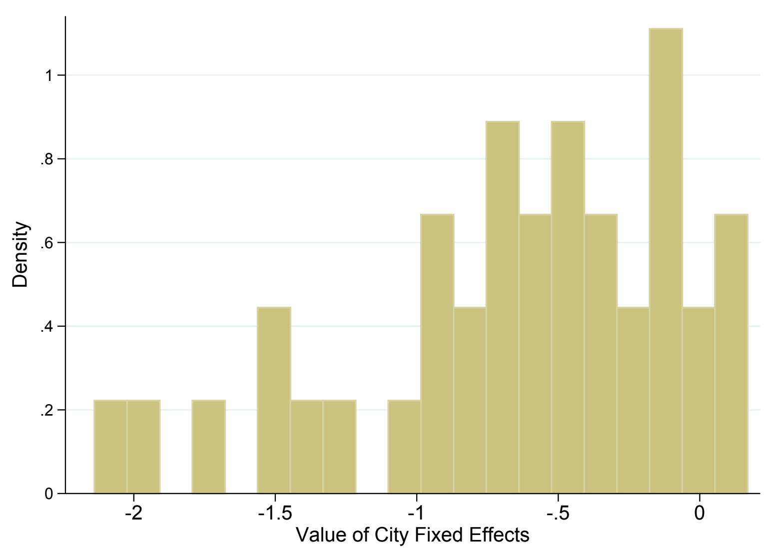

Figure I.2: Region Fixed Effects in the Marginal Cost Function in RCNL-1

drop negative incomes from the American Community Survey data.[1]

[1] In robustness checks, we have experimented with a number of different ways to sample income. These including using individual income, household income, and household income divided by the number of family members. For each of these, we have tried dropping negative incomes, dropping incomes below five thousand, and recoding negative incomes as zeros. We top-code income at the 95th percentile in each case to limit the

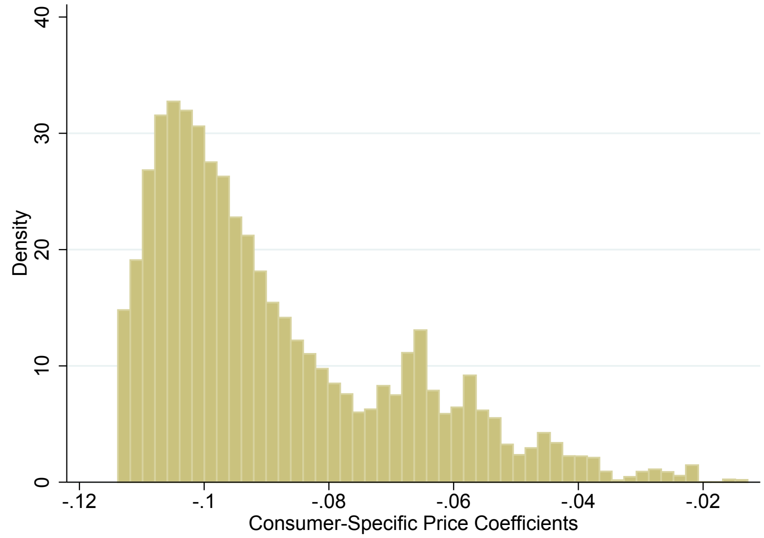

Figure I.3: Histogram of Consumer-Specific Price Coefficients in RCNL-1

Table I.1 shows the mean demand elasticities that correspond to the RCNL-3 specification of Table 4. As shown, the higher-priced imports tend to have more elastic demand. This is mechanical because the specification does not allow income to rotate the demand curve. For example, with logit demand the own price elasticity of product j is simply $\alpha p_j(1 - s_j)$. If income is allowed to rotate demand, as in the RCNL-1 specification, then we find that the higher-priced imports still have more elastic demand, but by a smaller amount. The logit restriction that consumers substitute to other products in proportion to their market shares is somewhat relaxed. Looking along column (5) for Corona Extra, there is considerable heterogeneity in the cross-elasticities, but the same is not true in column (1) for Bud Light. Thus, while the coefficient estimates of RCNL-3 introduce qualitatively similar effects on substitution, they provide less meaningful departures from nested logit demand.

influence of positive outliers. The demand-side results are robust.

Table I.1: Mean Elasticities for 12-Pack Products from RCNL-3

Brand/Category		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
								Product-Specific Own and Cross Elasticities						
(1)	Bud Light	-3.819	0.141	0.017	0.150	0.242	0.094	0.158	0.045	0.034	0.108	0.039	0.060	0.164
(2)	Budweiser	0.284	-3.958	0.017	0.146	0.281	0.104	0.187	0.049	0.033	0.105	0.038	0.062	0.162
(3)	Coors	0.281	0.144	-4.094	0.144	0.284	0.103	0.188	0.049	0.033	0.104	0.038	0.062	0.161
(4)	Coors Light	0.288	0.141	0.017	-3.969	0.233	0.092	0.151	0.044	0.034	0.109	0.039	0.059	0.164
(5)	Corona Extra	0.274	0.153	0.019	0.135	-5.349	0.133	0.275	0.062	0.033	0.096	0.036	0.069	0.158
(6)	Corona Light	0.278	0.150	0.018	0.139	0.359	-5.637	0.245	0.057	0.033	0.099	0.037	0.067	0.160
(7)	Heineken	0.272	0.154	0.019	0.133	0.416	0.137	-5.506	0.063	0.033	0.095	0.036	0.070	0.158
(8)	Heineken Light	0.220	0.112	0.014	0.112	0.273	0.100	0.187	-5.724	0.023	0.078	0.027	0.053	0.127
(9)	Michelob	0.260	0.126	0.014	0.125	0.207	0.083	0.139	0.040	-4.452	0.099	0.032	0.053	0.152
(10)	Michelob Light	0.289	0.140	0.017	0.152	0.225	0.090	0.146	0.043	0.034	-4.401	0.039	0.059	0.164
(11)	Miller Gen. Draft	0.287	0.141	0.017	0.150	0.242	0.094	0.158	0.045	0.034	0.108	-4.075	0.060	0.163
(12)	Miller High Life	0.283	0.146	0.018	0.145	0.298	0.108	0.199	0.051	0.033	0.104	0.038	-3.095	0.162
(13)	Miller Lite	0.287	0.142	0.017	0.150	0.245	0.095	0.160	0.045	0.034	0.108	0.039	0.060	-3.949
(14)	Outside Good	0.015	0.007	0.001	0.008	0.012	0.005	0.008	0.002	0.002	0.006	0.002	0.003	0.009
								Total Cross Elasticities by Category						
	6-Packs	0.279	0.147	0.018	0.141	0.336	0.116	0.227	0.054	0.033	0.101	0.037	0.065	0.160
	12-Packs	0.281	0.144	0.017	0.143	0.285	0.107	0.194	0.051	0.033	0.103	0.037	0.062	0.161
	24-Packs	0.288	0.139	0.017	0.152	0.222	0.088	0.143	0.042	0.034	0.110	0.039	0.058	0.164
	Domestic	0.287	0.140	0.017	0.151	0.234	0.092	0.153	0.044	0.034	0.109	0.039	0.059	0.163
	Imported	0.274	0.153	0.019	0.135	0.398	0.134	0.270	0.061	0.033	0.096	0.036	0.069	0.158

Notes: The table provides mean elasticities of demand for 12-packs based on the RCNL-3 specification (column (iv) of Table **??**). The cell entry in row i and column j is the percentage change in the quantity of product i with respect to the price of product j. Means are calculated across the year-month-region combinations. The category cross-elasticities are the percentage change in the combined shares of products in the category due to a 1 percent change in the price of the product in question. Letting the category be defined by the set B, the calculation is $\left(\sum_{j\in B, j\neq k} \frac{\partial s_j(p)}{\partial p_k}\right) \frac{p_k}{\sum_{j\in B, j\neq k} s_j(p)}$. The categories exclude the product in question. Thus, for example, the table shows that a 1 percent change in the price of a Bud Light 12-pack increases sales of other 12-packs by 0.281 percent.

Table I.2: Brewer Markups from RCNL-3

Brand	6-Packs Pre	6-Packs Post	12-Packs Pre	12-Packs Post	24-Packs Pre	24-Packs Post
Bud Light	4.00	4.74	4.03	4.78	4.05	4.81
Budweiser	3.97	4.70	4.00	4.74	4.03	4.78
Coors	2.73	4.56	2.75	4.62	2.77	4.66
Coors Light	2.76	4.63	2.77	4.67	2.78	4.71
Corona Extra	2.84	2.78	2.81	2.76	2.82	2.82
Corona Light	2.80	2.75	2.77	2.73	2.84	2.82
Heineken	2.64	2.63	2.63	2.61	2.74	2.72
Heineken Light	2.62	2.60	2.60	2.58	2.67	2.68
Michelob	4.00	4.85	4.06	4.88	4.04	5.04
Michelob Light	4.01	4.76	4.04	4.80	4.14	4.70
Miller Gen. Draft	3.18	4.64	3.18	4.66	3.19	4.72
Miller High Life	3.20	4.64	3.16	4.60	3.18	4.65
Miller Lite	3.17	4.61	3.18	4.66	3.19	4.70

Notes: The table provides average markups for each brand-size combination, separately for the pre-merger and post-merger periods, based on the RCNL-3 specification shown in column (iv) of Tables **??** and **??**.

Table I.2 provides the brewer markups that arise from the RCNL-3 specification. Similar before/after comparisons can be made about the effect of the merger. MillerCoors markups increase due to greater market power and lower costs; ABI markups increase due to greater market power; and Modelo/Heineken markups are relatively unaffected. The markup for products of the same firm are quite similar with the RCNL-3 specification. This is again because the demand estimates to not introduce large departures from nested logit demand (which dictates that multiproduct firms equate markups across products).

Our baseline demand specification imposes that heterogeneity in consumer tastes for product observables is due solely to income. This restriction is not necessary, and indeed many choice models allow consumer tastes to depend on random draws from a standard normal distribution (e.g., Berry, Levinsohn, and Pakes (1996); Nevo (2001)). We explore adding unobserved heterogeneity in tastes for observables using the following redefinition of the consumer-specific parameters:

$$\begin{bmatrix} \alpha_i^* \\ \beta_i^* \end{bmatrix} = \begin{bmatrix} \alpha \\ \beta \end{bmatrix} + \Pi D_i + \Sigma v_i, \qquad v_i \sim N(0, I) \tag{I.1}$$

where Σ is a diagonal scaling matrix with diagonal elements $(\nu_1, \nu_2, \dots)$ that determine the importance of the standard normal draws. Our baseline specifications implicitly set Σ equal to the zero matrix. Table I.3 summarizes the results from two augmented specifications.

In column (i) we add an interaction between price and the N(0,1) draws to the RCNL-2 specification. The coefficient on the interaction is small and has a t-statistic of about 0.02. The other coefficients are similar to their counterparts in the baseline RCNL-2 regressions. In column (ii) we add an interaction between the constant and the N(0,1) draws to the RCNL-4 specification. Again the coefficient on the interaction is small and imprecisely estimated, and the other coefficients do not change much. This pattern has held in a number of other specifications. The use of N(0,1) draws to incorporate unobserved taste heterogeneity does not affect the results, so we opt for the simpler approach in our baseline specifications.

J Additional Counterfactual Results

We develop confidence bounds for selected welfare statistics using a partial bootstrap. We start by drawing 100 values from the estimated distribution of κ. This distribution reflects the impact of demand-side uncertainty on supply-side precision. We calculate markups based on the demand estimates and each random value of κ, and reestimate the linear marginal cost parameters. Finally, we recompute equilibrium in the post-merger periods in the counterfactual scenario that the merger does not occur. We use the RCNL-2 specification with quarterly data to reduce the computation burden. The results of RCNL-2 and RCNL-1 are quite similar so the confidence bounds are relevant to the counterfactual analysis presented in the body of the paper.

Figure J.1 plots prices for Miller Light and Bud Light 12-packs. The raw data are shown in black and each of the 100 "no merger" scenarios are shown in gray. The raw data exceed all of the simulated prices. We also calculate the percentage change in producer surplus, consumer surplus, and total surplus due to the merger. We find that producer surplus increases by 19% with a 90% confidence interval of (16%,21%). Consumer surplus decreases by 3.9% with a confidence interval of (−5.6%,−2.5%). Total surplus increases by 1.5% with a confidence interval of (−0.5%,3.0%). A one-sided test rejects the null hypothesis that total surplus decreases at the ten percent level.

An interesting feature of the counterfactual "no merger" scenario is that the Bud Light prices appear to fall somewhat in the post-merger periods relative to the pre-merger periods (e.g., see Figure 5). This happens in part due to the trends down in the demand and marginal cost time fixed effects shown in Figure I.1. To explore this further, we recompute equilibrium prices under all scenarios under the following adjustments:

- Post-merger time fixed effects take the value of the May 2008 fixed effect.

Table I.3: Baseline Demand Estimates

Demand Model:		RCNL-2	RCNL-4
Data Frequency:		quarterly	quarterly
Variable	Parameter	(i)	(ii)
Price	α	-0.1033	-0.0931
		(0.0148)	(0.0156)
Nesting Parameter	ρ	0.8084	0.7949
		(0.0424)	(0.0587)
Income Interactions			
Income×Price	Π_1	0.0008	
		(0.0005)	
Income×Constant	Π_2	0.0136	0.0240
		(0.0054)	(0.0050)
Income×Calories	Π_3	0.0043	0.0043
		(0.0019)	(0.0018)
Income×Import	Π_4		0.0036
			(0.0022)
Income×Package Size	Π_5		-0.0017
			(0.007)
N(0,1) Interactions			
v_i×Price	ν_1	0.0047	
		(0.3083)	
v_i×Constant	ν_2		0.0004
			(31.0957)
Other Statistics			
Median Own Price Elasticity		-4.81	-4.87
Median Market Price Elasticity		-0.68	-0.68
Median Outside Diversion		14.52%	14.67%
$J-$Statistic		13.92	14.23

Notes: The table shows the baseline demand results. Estimation with GMM. There are 31,784 observations at the brand-size-region-year-quarter level. The samples exclude the months/quarters between June 2008 and May 2009. All regressions include product (brand×size) and period (month or quarter) fixed effects. The elasticity and diversion numbers represent medians among all the brand-size-region-quarter-year observations. Standard errors clustered by region and shown in parentheses.

- Diesel prices in the post-merger periods take the May 2008 value.
- The 2008 income draws are used in 2009, 2010, and 2011.
- The structural error terms are zero in all time periods.

Figure J.2 shows the resulting price plots for Miller Lite and Bud Light. The inter-temporal

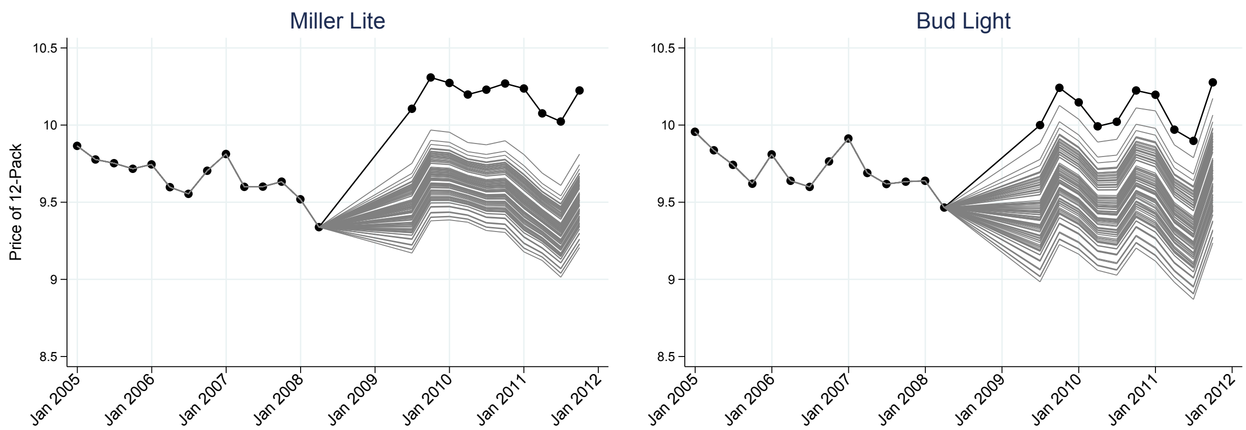

Figure J.1: Bootstrapped Counterfactual Prices

Notes: The figure plots the average prices of Miller Lite and Bud Light 12-packs in the raw data (black) and the average prices in the "no merger" counterfactual scenario under 100 random draws from the estimated distribution of supply-side parameters (gray). Results are based on the RCNL-2 specification with quarterly observations. The averages are across the 39 regions.

variability in prices that exists post-merger is nearly eliminated (some remains due to product entry/exit). The price gaps between the different scenarios are similar to what is obtained in the baseline counterfactual simulations. To the extent that the lower Bud Light prices in the baseline simulations reflect prices that are somehow too low, due to estimation error or misspecification, a "conservative" calculation of the merger welfare effects can be obtained by comparing the welfare under the raw data (e.g., as in Figures 4 and 5) to welfare in the revised "no merger" scenario shown in Figure J.2. In that comparison, we find that the merger increases producer surplus by 23%, decreases consumer welfare by 2.0%, and increases total welfare by 2.8%. These numbers are comparable to the welfare statistics shown in Table 10. Table J.1 shows welfare stats where demand and costs, the income distribution, and diesel prices are all held to May 2008 values. Again, the numbers are similar to those reported in Table 10.

K Alternative Supply-Side Model

The κ parameter in the supply-side allows for a test of post-merger Nash-Bertrand competition even if the true post-merger equilibrium is not generated by ABI and MilllerCoors internalizing their pricing externality. In that case, however, the implied brewer markups and counterfactuals could suffer from misspecification bias. Here we develop and estimate an alternative supply-side model and obtain similar results.

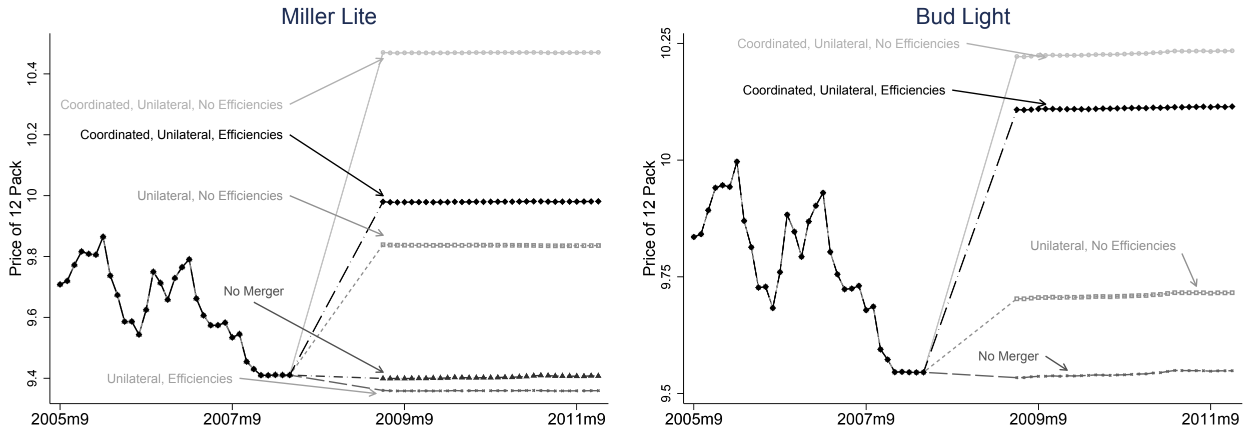

Figure J.2: Alternative Counterfactual Prices for Miller Lite and Bud Light

Notes: The figure plots the average prices of Miller Lite 12-Packs under five different counterfactual scenarios. The post-merger time fixed effects take the value of 2008:Q2 in both the demand and cost functions, demographics and gasoline prices in the post-merger periods take their 2008 values, and unobserved demand and costs are set to zero in all periods. Each dot represents the average prices across the 39 regions.

Table J.1: Welfare Effects with Fixed Demand and Costs

Coordinated Effects:	yes	yes	no	no
Unilateral Effects:	yes	yes	yes	yes
Efficiencies:	yes	no	yes	no
Producer Surplus	19.8%	17.6%	8.8%	7.5%
ABI	9.3%	18.5%	-0.3%	8.9%
Miller	36.5%	16.4%	24.3%	6.8%
Coors	50.6%	13.7%	37.1%	4.2%
Consumer Surplus	-3.1%	-4.7%	0.01%	-1.8%
Total Surplus	1.3%	-0.4%	1.8%	0.0%

Notes: The table provides the percentage changes in producer surplus, consumer surplus, and total surplus, relative to the "No Merger" scenario. All statistics are for the year 2011.

The alternative model captures a game in which all firms in the coalition increase prices above Nash-Bertrand level by the same amount. Denote the level increase $\widetilde{\kappa}$. The vector of equilibrium prices in each region-period satisfies the first order condition

$$p_t = mc_t + \widetilde{\kappa} \times \iota_t - \left[\widetilde{\Omega}_t \circ \left(\frac{\partial s_t(p_t;\theta^D)}{\partial p_t} \right)^T \right]^{-1} s_t(p_t;\theta^D) \tag{K.1}$$

where elements of the vector ι_t are indicators that equal one for ABI/MillerCoors products in

Table K.1: OLS Estimates from the Alternative Supply Model

Demand Model:		NL-1	RCNL-1	RCNL-2	RCNL-3	RCNL-4
Data Frequency:		monthly	monthly	quarterly	monthly	quarterly
Variable	Parameter	(i)	(ii)	(iii)	(iv)	(v)
Post-Merger Coalition Price Increase	$\widetilde{\kappa}$	0.798	0.664	0.600	0.779	0.749
		(0.081)	(0.081)	(0.086)	(0.084)	(0.082)
Marginal Cost Parameters						
MillerCoors×PostMerger	γ_1	-0.550	-0.613	-0.610	-0.674	-0.482
		(0.065)	(0.079)	(0.077)	(0.079)	(0.068)
Distance	γ_2	0.156	0.178	0.172	0.181	0.159
		(0.050)	(0.063)	(0.063)	(0.064)	(0.051)

Notes: The table shows the baseline supply results. Estimation is with OLS. There are 94,656 observations at the brand-size-region-month-year level in columns (i), (ii), and (iv), and 31,784 observations at the brand-size-region-year-quarter level in columns (iii) and (v). The samples exclude the months/quarters between June 2008 and May 2009. All regressions include product (brand×size), period (month or quarter), and region fixed effects. Standard errors clustered by region and shown in parentheses.

post-merger periods. The alternative ownership matrix, $\widetilde{\Omega}_t$, is standard and free of supply-side parameters. Thus markups are additively separable in the Nash-Bertrand markup and the $\widetilde{\kappa}$ parameter. We parameterize marginal costs as in equation (11). This yields an equation that can be estimated with OLS:

$$p_t + \left[\widetilde{\Omega}_t \circ \left(\frac{\partial s_t(p_t;\theta^D)}{\partial p_t}\right)^T\right]^{-1} s_t(p_t;\theta^D) = \widetilde{\kappa}\iota_t + w_{jrt}\gamma + \sigma_j^S + \tau_t^S + \mu_r^S + \eta_{jrt} \tag{K.2}$$

The empirical variation that identifies the $\widetilde{\kappa}$ parameter again relates to whether ABI price increases exceed what can be explained under Nash-Bertrand competition. Unbiasedness follows from the same identifying assumption employed in the baseline supply model: changes in the unobserved costs of ABI, before versus after the merger, are not systematically different from the changes in the unobserved costs of Modelo and Heineken.

Table K.1 shows the results for each of the demand specifications. The marginal cost parameters are close to those obtained from the baseline supply model (e.g., see Table 6). The estimates of $\widetilde{\kappa}$ range from 0.600 to 0.805. Post-merger Nash-Bertrand competition is rejected because the estimates are statistically different than zero. Interpreted strictly, the results suggest that coordination results in price increases of \$0.60-\$0.80. The magnitude of this effect is quite comparable to what arises with the baseline supply model. For example, Table 10 indicates that with the RCNL-1 specification, coordination increases the average markups of ABI from \$3.81 to \$4.45, a change of \$0.64.