Appendix

A.1. Defining Our Products

As background to how we define our products, it is first necessary to understand the product coding scheme that Census uses. There are three types of codes that we highlight. First, Census codes flags products from administrative records (AR) sources. We exclude all of these AR products from our analysis. (Including in our measures of PPSR since it is obviously not possible to assign these AR products to a single 7-digit code.) Second, Census uses balancing codes to correct cases in which the sum of the total value of shipments of reported individual products does not sum to the reported total value of shipments. Census identified these balancing codes using special suffixes for the product codes in every census year except in 1987. Where balancing codes are identified, they have been deleted. Finally, Census collects data on receipts for contract work, miscellaneous receipts, and resales of products. These products are excluded from our calculations of PPSR (again, because it is obviously not possible to assign these AR products to a single 7-digit code). As a final exclusion, we did not include any products in that have a negative value since these are presumably balancing codes. The precise definitions of our eleven products are listed below (with 7-digit product codes in parentheses).

Boxes is defined as the sum of boxes classified by their end use and boxes classified by their materials. Boxes classified by end use are: food and beverages (2653012), paper and allied products (2653013), carryout boxes for retail food (2653014 category starts in 1987) glass, clay, and stone products (2653015), metal products, machinery, equipment, and supplies except electrical (2653016), electrical machinery, equipment, supplies, and appliances (2653018), chemicals and drugs, including paints, varnishes, cosmetics, and soap (2653021), lumber and wood products, including furniture (2653029), and all other ends uses not specified above (2653029 in 1977 and 1982, 2653030 in 1987). Boxes classified by their materials are: solid fiber (2653051), corrugated paperboard in sheets and rolls, lined and unlined (2653067), and corrugated and solid fiber pallets, pads and partitions (2653068). The physical data for boxes is measured in short tons.

Bread is defined as one 7-digit product, white pan bread (2051111), until 1992 when it was split into two products white pan bread, except frozen (2051121) and frozen white pan bread (2051122). The physical data for bread is measured in thousands of pounds.

Carbon Black is defined as one 7-digit product, carbon black (2895011 in 1977, 2895000 thereafter). The physical data for carbon black is measured in thousands of pounds.

Coffee is the sum of whole bean (2095111), ground and extended yield (2095117 and 2095118 in 1982 and 2095115 thereafter), and ground coffee mixtures (2095121). The physical data for coffee is measured in thousands of pounds.

Concrete is defined as one 7-digit product, ready-mix concrete (3273000), over our entire sample. Some of the products coded as 3237300 in 1987 were in fact census balancing codes and thus were deleted from our sample. The physical data for concrete is measured in thousands of cubic yards.

Flooring is defined as one 7-digit product, hardwood oak flooring (2426111), over our entire sample. The physical data for flooring is measured in thousands of board feet.

Gasoline is defined as one 7-digit product, motor gasoline (2911131), over our entire sample. The physical data for gasoline is measured in thousands of barrels.

Block Ice is defined as one 7-digit product, can or block ice (2097011), over our entire sample. The physical data for block ice is measured in short tons.

Processed Ice is defined as one 7-digit product, cubed, crushed, or other processed ice (2097051), over our entire sample. The physical data for processed ice is measured in short tons.

Plywood is defined as one 7-digit product, hardwood plywood (2435100), over 1977-1987. Starting in 1992, plywood is the sum of veneer core (2435101), particleboard core (2435105), medium density fiberboard core (2435107), and other core (2435147). The physical data for plywood is measured in thousands of square feet surface measure.

Sugar is defined as one 7-digit product, raw cane sugar (2061011), over our entire sample. The physical data for sugar is measured in short tons.

A.2. Measurement of input levels and input elasticities in the TFP indexes.

This section reports details on the measurement of input levels and elasticities in the TFP measures described in Section II.

Labor inputs are measured as plants' reported production-worker hours adjusted using the method of Baily, Hulten and Campbell (1992). This involves multiplying the production-worker hours by the ratio of total payroll to payroll for production workers. Prior work has shown this measure to be highly correlated with Steven J. Davis and Haltiwanger's (1991) more direct imputation of nonproduction workers, which multiplies a plant's number of nonproduction workers by the average annual hours for nonproduction workers in the corresponding two-digit industry calculated from the CPS. Capital inputs are plants' reported book values for their structure and equipment capital stocks deflated to 1987 levels using sector-specific deflators from the Bureau of Economic Analysis. The method is detailed in Foster, Haltiwanger and Krizan (2001). Materials and energy inputs are simply plants' reported expenditures on each deflated using the corresponding input price indices from the NBER Productivity Database.

To compute the industry-level cost shares that we use to measure the input elasticities α_j, we use the materials and energy expenditures along with payments to labor to measure the costs of these three inputs. We construct the cost of capital by multiplying real capital stock value by the capital rental rates for the plant's respective two-digit industry. These rental rates are from unpublished data constructed and used by the Bureau of Labor Statistics in computing their Multifactor Productivity series. Formulas, related methodology, and data sources are described in U.S. Bureau of Labor Statistics (1983) and Michael Harper, Ernst R. Berndt, and David O. Wood (1989).

A.3. Rules for Inclusion in the Sample

We detail here our sample inclusion criterion discussed in Section II.A.

Census reports physical product data for only a subset of the 11,000 products reported in the Census of Manufactures. While we use only products for which physical output is reported, the collection of this data has changed over time for two of our products. (See Table A.1.) Census did not collect physical output for ready-mix concrete in 1997, and the unit of measurement for boxes changed over our sample period in a way that makes the 1992 and 1997 data incomparable to the earlier periods. Additionally, there are recording flaws in the 1992 quantity data for processed ice that make using it unfeasible.

The Census Bureau relies on administrative record data for very small establishments (typically with less than five employees). In these cases all production data except total revenues and the number of employees are imputed, and production operations are classified only up to the four-digit industry level. Since our unit of analysis is more detailed than the four-digit industry, we cannot determine whether a particular administrative record establishment actually produces the product of interest. For these reasons, we exclude administrative records cases from our sample. While about one-third of CM establishments are administrative records, their output and employment shares are much less because they are such small plants.

We also exclude establishments whose data appear to be imputed or suffer from reporting or recording errors. The Census Bureau imputes physical quantities when product-level data are not fully reported. Unfortunately, imputed data are not explicitly identified. To distinguish and remove imputed product-level data from the sample, we use techniques similar to those employed by Roberts and Supina (1996, 2000). To minimize the influence of reporting and recording errors, we also remove a small number of plants reporting physical quantities that imply prices greater than ten times or less than one-tenth the median price in a given year. In order to maintain the same sample over all exercises, we delete observations that are missing any one of the main regression variables. We also delete observations when the plant's labor or materials cost share is less than one-tenth of the corresponding industry's average cost share for that year, or when the cost share is more than one. Finally, we still find a relatively small number of obvious outliers in physical quantity measures, so we trim the one-percent tails of the physical productivity (TFPQ) distribution.

Our product specialization criterion requires that plants obtain at least 50 percent of their revenue from our product of interest. The text discusses the measurement reasons for imposing this restriction as well as describing a robustness check with respect to this product specialization cutoff.

A.4. Characteristics of Establishments by Product

In this section we briefly characterize some of the relevant properties of the establishments that produce our products. Table A.1 shows characteristics of the sample by product. The first five columns show the number of establishments in our sample by year for each product. The second to last column shows the real revenue shares of each product. Real revenue is the weight used in our weighted regressions. Concrete clearly dominates our sample in terms of the number of establishments while gasoline dominates in terms of the revenue share. The table's last column shows mean logged income (income is taken from Census reports for the county in which the plant is located) for each product in our sample. Concrete has the highest mean log income while carbon black has the lowest.

Table A.2 shows the entry and exit rates by product for the data pooled over all available years. Entry rates range from a low of 3.9 for sugar to a high of 26.6 for concrete, while exit rates range from a low of 9.0 for gasoline and to a high of 27.7 for processed ice. Some products appear to be in a period of retrenchment or consolidation. Sugar for example, has a very low entry rate (3.9) but a high exit rate (17.0). The number of plants in the sugar and confectionary products industry (SIC 2061) has fallen from 66 in 1977 to 39 in 1997. Other products appear to simply have a high degree of churning. For example, concrete and both types of ice products all have entry rates and exit rates that exceed 20 percent. The number of establishments in ready-mixed concrete (SIC 3273) industry increases over our sample period, while the number of establishments in the block and processed ice industry (SIC 2097) falls somewhat over our sample, from 675 establishments in 1977 to 582 establishments in 1997.

A.5. Robustness of Demand Estimates

There are two potential problems with using TFPQ as an instrument. We have explored both and found our results to be robust. We briefly describe these potential issues here.

The first regards how selection on profitability impacts the assumption that TFPQ is uncorrelated with demand shocks at the plant level. This assumption, which is necessary for TFPQ to be a valid instrument in the demand estimation, strikes us as quite reasonable for *innovations* to the dynamic plant-level TFPQ and demand processes. Hence for entering plants, which get their initial idiosyncratic productivity and demand levels as draws from the respective innovation processes, the assumption of no correlation is likely to hold. However, TFPQ and demand levels may be correlated for *continuing* plants (those operating in both the current and previous periods) because of selection. Here is why. A producer continues operation if it is profitable, which depends on idiosyncratic productivity and demand. Since a producer with a higher TFPQ draw can tolerate lower demand draws (and vice versa) while remaining profitable, those producers that chose in the previous period to continue operations into the current period—i.e., the continuing plants—will tend to have negatively correlated lagged TFPQ and lagged demand levels. The correlation arises in the lagged values because the plants' decisions to continue operating into the current period were conditioned on that information. But because both of these processes are persistent, their current values could be correlated as well, making the orthogonality assumption necessary for consistent demand elasticity estimation questionable.

To explore the sensitivity of our results to this issue, we estimated product demand curves (10) using an alternative instrument for price that is based only upon innovations to TFPQ. For entrants, this was simply the observed level of TFPQ (after taking out product-year fixed effects). For continuers, on the other hand, we regressed the plants' current TFPQ levels on their lagged values (again including product-year fixed effects), and used the residual from this to instrument for price. In this way, we use the innovation to producers' productivity levels—information that was plausibly not included in their decision in the previous period to continue operations, but that should affect current costs—to gain exogenous variation in current prices. (The fact that we only observe plants every five years makes necessary some assumptions of timing here. The first is that the period when plants have received their idiosyncratic profitability draws but have yet to decide whether to exit, which is instantaneous in our two-stage model, can in reality actually correspond to a period of a few years. The entering plants in our sample are assumed to be in this stage. The second is that the TFPQ innovation does not impact producers' exit decisions in intercensal years.)

When we use this alternative TFPQ-innovation-based instrument, we find that the pattern of elasticities and even more importantly the pattern of demand shocks to be very robust. We also find the results that depend upon demand shocks in Tables 3-6 to be virtually unaffected. These results are available upon request.

A second potential problem is measurement error. We compute prices by dividing reported revenue by quantity and any measurement error in physical quantities will overstate the negative correlation between prices and TFPQ. Since the first stage of the IV estimation regresses plants' prices on their TFP levels, measurement error would yield biased estimates of the fitted prices used in the second stage, possibly leading in turn to biased price elasticities and idiosyncratic demand measures. To address this concern, we estimated a specification that should be robust to measurement error, implemented as follows. Rather than using current plant TFP directly as an instrument, we use a "fitted" value constructed as a projection of current TFP on several values expected to be correlated with plant productivity but orthogonal to measurement error. These include the plant's lagged TFP, lead TFP, the average TFP of the *other* plants in the same industry owned by the plant's firm (these three variables are used as available, depending on the plant), a set of birth-cohort-by-industry dummies, and a set of "survival-cohort"-by-industry dummies denoting how many more CMs a plant exists before exiting (or if it survives until the end of our sample).

All of these variables should have predictive power over a plant's current TFP level but be orthogonal to classical measurement error in quantity. Lag and lead TFP values are relevant to current TFP because of productivity persistence, the average TFP of other plants owned by the same firm because of management-driven productivity spillovers or assortative matching of plants, and the birth and survival cohort indicators because of age-specific productivity evolutions and because future survival is determined by current productivity. Second, besides this fitted TFP value, we also include the local average wage as an instrument in the first stage. This factor price is of course an additional cost shifter for the plant.

As with the other robustness check above, the results of this alternative demand estimation procedure yield demand estimates that closely track those obtained in our benchmark specification. Moreover, the subsequent results using the estimated demand shocks from this alternative estimation are very similar to those reported in the paper. Again, we do not report these here for space reasons.

In considering these two robustness checks, we note that they employ essentially opposite strategies: one uses the transitory variation in TFPQ while the other uses the permanent component. We are reassured by the fact that the results are robust to using either component of TFPQ variation. If either potential problem was a major driver of our results, we would expect the demand estimates to be quite sensitive to these distinctions.

Table A.1: Characteristics of the Sample by Product

Product	Number of Observations 1977	1982	1987	1992	1997	Real Revenue Share (percent)	Mean (log) Income
Boxes	936	905	1045	NA	NA	7.9	17.4
Bread	195	142	110	92	92	2.4	17.0
Carbon Black	31	23	22	21	18	0.7	16.2
Coffee	61	84	79	77	77	4.7	18.0
Concrete	2184	3316	3236	3427	NA	7.0	17.1
Hardwood Flooring	8	10	16	25	24	0.2	16.7
Gasoline	99	99	94	95	71	75.0	17.1
Block Ice	40	43	26	23	10	0.0	16.9
Processed Ice	87	155	144	NA	NA	0.1	16.8
Plywood	71	68	42	42	37	0.6	16.5
Sugar	40	36	30	35	26	1.3	16.6

Note: This table shows the number of establishments in our sample by product and year, as well as each product's share of total real revenue in the sample (pooled across all years). Mean log income, used in our demand estimation procedure, is for plants' corresponding Economic Areas (see text for details) based on data pooled over all years.

Table A.2: Entry and Exit Rates by Product

Products	Entry Rates	Exit Rates
All Products	22.3	19.6
By Product:		
Boxes	12.4	12.2
Bread	7.6	18.9
Carbon Black	4.8	13.4
Coffee	9.1	15.6
Concrete	26.6	21.8
Hardwood Flooring	18.7	11.9
Gasoline	4.2	9.0
Block Ice	24.5	26.5
Processed Ice	23.1	27.7
Plywood	7.4	10.3
Sugar	3.9	17.0

Note: This table shows the plant entry and exit rates (averaged across all years in the sample). Entry (exit) is determined by plants' first (last) appearance in a CM. See text for details. Entry rates for gasoline reflect a small number (less than five) births in the 1987, 1992, and 1997 CMs, even though the Energy Information Administration reports no new refineries were built during that period. It is not clear why there is a discrepancy in the Census data, but for the sake of consistency and given the low entry rate in the industry overall, we defined these new plants as entrants.