Localized Competition and the Aggregation of Plant-Level Increasing Returns: Blast Furnaces, 1929-1935

Author(s): Amy L. Bertin, Timothy F. Bresnahan and Daniel M. G. Raff

Source: *Journal of Political Economy*, Apr., 1996, Vol. 104, No. 2 (Apr., 1996), pp. 241-266

Published by: The University of Chicago Press

Stable URL: https://www.jstor.org/stable/2138926

Localized Competition and the Aggregation of Plant-Level Increasing Returns: Blast Furnaces, 1929–1935

Amy L. Bertin

Economics Resource Group

Timothy F. Bresnahan

Stanford University and National Bureau of Economic Research

Daniel M. G. Raff

University of Pennsylvania and National Bureau of Economic Research

A recent empirical literature has shaken economists' confidence in the value of aggregate (industry-level) data to illuminate production relationships. But the statistical finding "you cannot aggregate," however well documented, is not an economic explanation. Plant-level relationships do aggregate in Depression-era blast furnace operations despite the presence of very substantial interplant heterogeneity, the most common economic cause of nonaggregability. The economic explanation of this lies in poor short-run substitutability of one plant's output for another's. Substitutability determines the importance of composition effects in understanding aggregate time series, constrains the potential cleansing effects of recessions, and therefore influences industry evolution quite broadly.

We would like to thank the National Science Foundation and the Sloan Foundation for financial support; seminar participants at the Federal Reserve Board, Harvard, NBER, Princeton, Stanford, and the University of British Columbia for helpful comments; and Bethlehem Steel Corp. for access to a blast furnace. Data used in the study will be archived at the Inter-University Consortium for Political and Social Research. The usual disclaimer applies.

[*Journal of Political Economy*, 1996, vol. 104, no. 2]
 0022-3808/96/0402-0001$01.50

I. Introduction

Procyclical productivity movements at the industry level are large and important. It is by now well documented statistically that industry productivity movements are not a simple or straightforward reflection of plant-level production functions (see Dunne, Roberts, and Samuelson [1989], Hamermesh [1989], Bresnahan and Raff [1991], Davis and Haltiwanger [1992], and a summary in McGuckin [1993]). Plant heterogeneity is well known to be an important part of this gap.[1] The observation that industry- and plant-level productivity movements are different, however, is not an economic explanation of either. In this paper, we examine plant-level production functions and their aggregation inside one industry—blast furnaces during the Great Depression—in search of such an explanation.

Our analysis considers three elements behind industry productivity movements over the business cycle run: plant-level production functions, heterogeneity across plants, and the industry equilibrium allocation of output across plants. We build an econometric model of the production function based on plant-level operations decisions. The hypothesis of short-run increasing returns to labor (SRIRL) has a very specific representation in our model. Some of the labor input was indivisible. Plant operation at slow production rates meant that the indivisible labor was spread over few output units. The task of the model is to define "slow production" operationally. The econometrics then serve to measure the indivisibility. By using the Alchian (1959) framework in which slow production rates are distinguished from other margins of output adjustment and by using very detailed plant-level data, our model offers a new look at the inside-the-plant phenomena behind SRIRL.

We find substantial SRIRL at both industry and plant levels. As in most industry studies, we also find wide variety in efficiency across plants. The Depression did not have much of a cleansing effect on blast furnace operations because exit did not much reduce this variety in efficiency by the 1933 trough. Industry equilibrium, then, tolerated both wasteful duplication of fixed labor costs and production at high-cost units. These two surprising features of equilibrium in the industry in this time period are closely linked by the logic of competition and aggregability. A perfectly competitive industry minimizes cost over all plants, tolerating no wasteful duplication of fixed costs and exploiting cost heterogeneity. A cluster of local monopolies in the

[1] Davis and Haltiwanger's (1992) finding that gross flows at the plant level vary dramatically and that this variation is lost in net industry flows is important evidence for heterogeneity.

same "industry" permits heterogeneity and duplication of fixed costs. As the short-run elasticity of substitution between different plants' products moves between zero and negative infinity, industry equilibrium moves from the conglomeration of local monopolies to perfect competition. The implications for short-run cost minimization move in parallel.[2] In the short run, the degree of interplant substitutability depends on the products, locations, and trading partner ties of the existing stock of plants.[3] The degree of competition mediates between plant-level phenomena such as SRIRL on the one hand and aggregate outcomes on the other. It is an empirical question whether interplant substitutability is high or low in any particular industry. As a result, it is an empirical question whether aggregate behavior will be like or unlike the behavior observed at any particular "representative" plant.

The Depression years and blast furnace operations offer an attractive opportunity to explore these matters for three distinct reasons. First, the Great Depression was a very large natural experiment. It exposed production relationships in stark relief (see Bernanke 1986; Bernanke and Parkinson 1991). Second, blast furnaces are so big and so public that it is impossible to keep knowledge of their operations secret, even when the information is commercially sensitive. Accordingly, we have been able to gather a public-source summary of operations information at the plant level. This complements the plant-level data in the recently rediscovered manuscripts of the *Census of Manufactures* conducted for the years 1929, 1931, 1933, and 1935 to form a database unusually well suited to plant-level analysis. Finally, there was very little investment or technical change in the industry and period we study. This dramatically simplifies the econometric issues.[4]

II. Decision-Making Margins in Operations

Table 1 summarizes the basic aggregate data from blast furnace operations for 1929–35. The numbers derive from the establishment-level information for blast furnace products plants in the manuscript returns to the *Census of Manufactures* and from the trade journal *Iron Age*.

[2] The existence of alternative allocation mechanisms, such as managerial discretion within the multiplant firm, has much the same implications. The extent of their influence turns equally on the degree of substitutability across plants.

[3] Transportation and transaction costs reduce interplant substitutability. Recent analyses by Williamson (1983) and Carlton (1991) emphasize the economic importance of contracting and relationship-specific knowledge between particular buyers and sellers, and these standing commercial relationships can cause low ex post substitutability.

[4] Olley and Pakes (1992) analyze production and productivity at the plant level on modern data for another industry. They make considerable progress on the very difficult problem of investment.

TABLE 1

AGGREGATE OUTPUT, LABOR, AND CAPACITY UTILIZATION, 1929–35

	1929	1931	1933	1935
Total tons (millions)	42.5	18.3	13.2	21.1
Wage earner months (thousands)	300	160	139	179
Man-hours* (millions)	72.5	33.6	22.8	30.5
Productivity:				
Tons/WEM†	141.7	113.8	95.0	117.9
Tons/MH‡	.578	.544	.575	.695
Tons/WEM, 1933 plants§	153.1	116.9	95.0	120.8
Tons/MH, 1933 plants	.622	.561	.575	.706
Number of plants	104	76	66	68
Total furnaces, active some time in year	244	141	128	141
Furnace months	2,482	1,095	948	1,358
Furnace opens plus closes per furnace month	.060	.131	.177	.085
Capacity utilization‖ (furnace months/potential furnace months)	.76	.39	.37	.55
Total days operated (average over plants)	341	285	263	314

* Man-hours is missing for 10 plants in 1929, one in 1931, and two in 1933.
† Average aggregate productivity equals total tons divided by total wage earner months.
‡ Average aggregate productivity equals total tons divided by total man-hours, where total tons is calculated for plants reporting man-hours.
§ Aggregate productivity for all four years is calculated only for the plants open during 1933.
‖ Uses sum of monthly furnace utilization divided by 12 times total furnaces at open plants. Monthly furnace series were constructed from *Iron Age*.

Output decreased 70 percent between 1929 and 1933. Depending on how labor input is measured, aggregate productivity is either strongly procyclical or flat. On a wage earner months (WEM) basis, productivity falls by about a third from 1929 to 1933, rebounding by about a fifth in the early upturn through 1935. Published census reports for 1937 (not shown) show further productivity growth as the recovery continued. On a man-hours (MH) basis, productivity can appear acyclical (compare 1929 to 1933) or even countercyclical (1931 to 1933).[5]

One possible interpretation of these numbers is that there was no SRIRL. Job sharing at the trough meant that hours, not head count, was the right measure of labor input.[6] But the next two rows undercut

[5] The published census summaries for blast furnaces do not report man-hour data. The manuscripts do. The years 1929–35 are the only interwar years from which census manuscripts survive.

[6] If job sharing occurred on a large scale at the trough, perhaps through shortening shift lengths, the actual labor input corresponding to a wage earner month would have changed over the period. Another problem could arise if the plant reported just the day shift head count, since this might overweight fixed job classifications. Secondary sources discussing the data do not address these measurement concerns. Articles in the trade press are ambiguous. Surviving archival materials (Carnegie-Illinois Steel Corp.—Duquesne Works 1936) have extremely narrow coverage. The Duquesne Works materials, for what it is worth, do not show shortening shift lengths.

this interpretation. We calculate the productivity measures for 1929–35 for only the sample of plants open in 1933. Now MH productivity also appears procyclical on a 1929 versus 1933 basis (albeit less so than WEM productivity), so that measured productivity at continuing plants is unambiguously procyclical. Clearly the aggregate data are inadequate to explore the competing economic hypotheses, and there is a puzzle to be explained in the aggregate MH versus WEM patterns.

Alchian (1959) identified a variety of operations margins for adjusting output, including changing the length of the production run and changing the production rate. Short production runs can lead to SRIRL if there are setup or shutdown tasks for laborers. If output is reduced by slowing the production rate, SRIRL can arise from lumpiness in the labor input, for example, from specialization.[7] In this section, we consider the production margins available to blast furnace operators with an eye to guiding plant-level modeling decisions and, ultimately, measurement.[8]

In this period, a blast furnace plant was a collection of one to 12 furnaces designed to convert iron ore into nearly pure iron.[9] The furnaces themselves were massive, 100 or more feet tall, and operated continuously. With minor exceptions, an individual furnace had only a single feasible operating rate, making the output choices for the managers of the plant quite simple.[10] The managers had to decide which furnaces to run at which times. Thus the production rate of a plant is captured by the number of furnaces it has operating, and the length of the production run is best defined at the furnace level. The final available adjustment margin, plant shutdown, has the obvious definition.

Once a furnace had been lit, its temperatures were slowly but steadily increased to operating levels. A temporary closure involved effort and costs at both shutdown and restart. Startup and shutdown themselves required 2 weeks and far smaller crews than regular operations. Further, the firebricks lining the furnace (to a depth of up to

[7] We ignore a third margin, shifts, since it is irrelevant to blast furnace operations.

[8] Chenery (1949) and Meyer et al. (1959) pioneered the use of engineering data in empirical studies of long-run production functions. We know of no other studies applying these methods to short-run relationships.

[9] For more detail on technical matters, see, e.g., Camp and Francis (1925, pp. 152–83) and Sweetser (1938, passim).

[10] The Depression years saw no major embodied innovations, and only one major innovation that was not embodied. A process control procedure called "fanning" allowed efficient operation at markedly lower throughput rates than had been possible up to that time. This technique appears to have emerged and diffused during the calendar year 1932. See Haven and Buell (1933), a review article in the January 1933 issue of *Blast Furnace and Steel Products*. We shall return to the implications of fanning below.

60 inches) were very durable when subjected to ongoing high temperatures but were likely to crack if the temperature changed quickly. So short production runs were characterized by high average costs, especially labor costs. Some of these costs would have been borne at the beginning and end of the production run. Other costs, as for the very considerable relining effort required when bricks approached the end of their useful life, could be scheduled freely during furnace downtime.

This production process was capital-intensive and had become considerably more so in the 20 years preceding our period (see Hogan 1971, pp. 861 ff.). By 1929, labor in day-to-day operations was primarily concerned with process control, tending the capital, and materials handling (see table 2).

An operating furnace's design determined the labor input required to feed in ore, coke, and limestone and to tend the blowing engines that blasted hot air into the furnace. Laborers to "tap" the furnace every 4–6 hours, an intermittent but dangerous and physically demanding task, were also used in proportions fixed relative to output. These classes of laborers were typically staffed in three shifts, though some smaller plants used two 12-hour shifts (the census manuscripts report shift lengths). The materials-handling tasks of labor were constant returns, not a source of SRIRL.

The services of bricklayers and others involved in the elaborate task of relining were required infrequently on any given furnace and so were shared across furnaces at multifurnace plants.[11] Large repair crews were also required on site both for emergencies and for periodic routine maintenance. These workers were specialists and could be shared across multiple furnaces. Table 2 makes clear that this could be a very substantial source of lumpiness in the labor input. The table shows nearly a quarter of the people on site in the daytime in the repairs group.[12] In sum, the engineering story suggests very specific operations measures associated with SRIRL, in particular connected with plant startup and shutdown and, through crew sharing, with the number of furnaces operating.[13]

[11] Relining probably occurred every two to three years in this period. Relining was clearly very highly skilled work, though this skill does not appear to have been furnace- or establishment-specific. The custom of having contractors carry out these tasks appears to postdate our period.

[12] Typically, workers in any particular group worked on the tasks of that group, so the lumpiness of the labor input could not easily be overcome by reallocation of workers to tasks. We thank Frederick Rorick of Bethlehem Steel and Joel Sabadasz of the National Park Service's Historic American Engineering Research Group for helpful discussions on these matters.

[13] This account has focused on features common to most blast furnace plants. There were also dimensions on which they were very different. Labor productivity varied

TABLE 2

TYPICAL BLAST FURNACE LABOR FORCE CIRCA 1929

	NUMBERS PER SHIFT				
DUTIES	Morning	Afternoon	Night	Day	TOTAL
Ore bridge operations				3	3
Bins	1		1		2
Stock house	2	2	2		6
Cast house	7	6	6		19
Engine house	5	5	5		15
Pig machine	3	3	3	1	10
Cranes	1	1	1		3
Transport	6	3	3		12
Laborers etc.	5	5	3	1	14
Repairs	6	6		7	19
Others not elsewhere classified	2	2		1	5
Grand total					108

SOURCE.—Clements (1929), p. 130.

At the bottom of table 1, we can see how the different margins were used to adjust output. All the available margins of adjustment were used, but not in the same proportions. About one-third of the plants simply closed. These plants must have been relatively small ones, since they accounted for less than a quarter of the furnaces. At the plants that remained open, production runs were varied. Total days operated shrank by about a quarter. Correspondingly, furnace openings and shutdowns per operating month roughly trebled.[14] Production rates, measured by furnace utilization, dropped by more than half. There was clearly scope for SRIRL in this.

For plausibility, the SRIRL story requires an explanation of why industry equilibrium tolerated slow production runs, since cost minimization would have economized on fixed labor costs by shutting more plants down. Our candidate explanation is simple. Most blast furnaces were associated with integrated steel works. The costs of transporting liquid iron (at 2,500° and more!) led to poor substitutability across plants, so that competition was blocked from economizing on duplication of fixed costs across plants.

---

tremendously in cross section, even at the 1929 peak. Older, smaller furnaces were associated with much larger costs. The econometric analysis documenting these assertions can be found in Bertin (1994, chap. 2).

[14] On the generous assumption of two weeks of 25 percent crew strength during shutdown and startup, the costs of short production runs trebled to about 2.2 percent of labor costs. But this is a small number. It therefore seems unlikely, at least on this engineering evidence alone, that inefficiently short production runs were an important source of low output per worker.

III. Specification of the Plant-Level Production Function

Our measurement framework is based on a short-run production function in which the important parameters measure the average physical product of labor at different production run rates. The number of furnaces operating is our observable statistic for the production run rate at blast furnace plants and therefore for the utilization rate for the fixed component of labor.[15] In our analysis, we aggregate across months to create an annual production function.[16] Our general econometric model of the production function is

$$Q_{iy} = \lambda_1 L_{iy1} + \lambda_2 L_{iy2} + \ldots + (\tau_i + \hat{\tau}_{iyf})L_{iy} + \epsilon_{iyf}. \quad (1)$$

Here Q_{iy} is the total annual output of plant i in year y, measured in tons, and L_{iyf} the labor input in months when f furnaces were operated. The parameter λ_f is the average product of labor when f furnaces are operating. In this specification, shifting from one- to two-furnace operations would raise the average product of labor by $\lambda_2 - \lambda_1$. There are three mean-zero error terms. Permanent heterogeneity in labor productivity is captured by the fixed effect τ_i. Transitory production function errors are modeled either through $\hat{\tau}_{iyf}$, independently and identically distributed shocks to productivity, or through ϵ_{iyf}, independently and identically distributed shocks to the level of the production function.[17] Note that the transitory shocks depend on the number of furnaces operating. Our preferred specification contains a correction for the endogeneity of the number of furnaces using sample selection techniques.

The *Census of Manufactures* manuscripts for the blast furnace products industry provide detailed plant-level data on inputs and outputs for 1929, 1931, 1933, and 1935. We supplement census data with information about operations decisions available from industry trade journals. Variable definitions are given in table 3, descriptive statistics in table 4, and a detailed description of data sources in Appendix A.

Labor input is measured monthly by the census in terms of both WEM and MH. The wage earner data are probably more reliable than the man-hour data.[18] Output is measured as tons of pig iron

[15] This source of productivity growth is very different from the long-run source, the exploitation of scale economies associated with larger furnaces, i.e., more capital.

[16] We have been able to build up monthly input data at the plant level, but available plant output data are annual.

[17] The choice of which of these two ways to model the transitory error turns out to matter little. In the tables in this paper, we use the production function shock, ϵ. Similar tables using the productivity shock, τ, are available from the authors.

[18] The quality of editing in the manuscripts was higher for wage earner data than for man-hour data. We simply drop 1929 man-hour data explicitly noted as incor-

TABLE 3

VARIABLE DEFINITIONS

	Production Function Data
TOTTON	Total tons of pig iron produced in the plant for the year
WEM	Annual sum of wage earners
MH	Annual sum of man-hours
WEM_0	Sum of WEM for months in which plant operates zero furnaces
WEM_1	Sum of WEM for months in which plant operates one (two, three) furnace
WEM_{2+}	Sum of WEM for months in which plant operates two or more furnaces
MH_0	Sum of man-hours for months in which plant operates zero furnaces
MH_1	Sum of man-hours for months in which plant operates one (two, three) furnace
MH_{2+}	Sum of man-hours for months in which plant operates two or more furnaces
$Months_0$	Number of months in a year in which plant operates zero furnaces
$Months_1$	Number of months in a year in which plant operates one (two, three) furnace
$Months_{2+}$	Number of months in a year in which plant operates two or more furnaces
MR	Sample selection correction to correct for endogenous furnaces
	Technology Variables
HP	1929 horsepower of prime movers and electric motors
One furnace	Equals one if plant has one furnace
	Demand Variables
Region	Annual tons (in thousands) of finished steel by region; regions defined in App. A
Prod	Annual tons (in thousands) of steel by primary downstream product category; product categories defined in App. A

SOURCE.—TOTTON, WEM, MH, and technology variables are taken from manuscripts of the *Census of Manufactures* for the blast furnace products industry. A series of monthly furnace usage by plant was constructed from *Iron Age* and combined with monthly wage earner and man-hour data from the *Census of Manufactures* to create WE_1, WE_{2+}, MH_1, and MH_{2+}. Region is taken from *AISI Annual Statistical Reports*. Prod is taken from aggregate *Census of Manufactures* data for the steel and rolling mill industry.

produced. This is available annually by plant. To conform with the annual plant-level output data, the labor input data must be aggregated to the format described in equation (1). We construct the annual variables L_{iyf}, the number of wage earners employed (or man-hours worked) at the plant in months when f furnaces were operated, and L_{iy}, the total annual WEM (MH).[19] Thus table 4 reports a 1929

rect. But there may be undetected errors in later years, since the census budgets shrank considerably.

[19] Some error in these definitions enters because the furnace data are reported as of the last day of the month, whereas wage earners are reported as of the fifteenth

TABLE 4

PLANT-LEVEL DESCRIPTIVE STATISTICS

	1929	1931	1933	1935
		All Plants		
Observations	103	75	66	68
TOTTON	409,965	243,827	200,441	310,406
	(472,133)	(214,480)	(178,856)	(284,057)
$Months_0$	.835	2.493	2.833	1.382
	(2.022)	(3.125)	(3.251)	(2.726)
$Months_1$	5.563	6.040	5.833	5.456
	(5.109)	(4.068)	(3.791)	(4.749)
$Months_{2+}$	5.611	3.467	3.273	5.162
	(5.527)	(4.494)	(4.097)	(5.430)
WEM	2,899	2,133	2,111	2,625
	(2,640)	(1,627)	(1,684)	(1,801)
WEM_0	36	148	114	63
	(112)	(840)	(205)	(133)
WEM_1	699	839	903	780
	(712)	(759)	(842)	(753)
WEM_{2+}	2,163	1,144	1,094	1,782
	(2,990)	(1,746)	(1,555)	(2,167)
HP	26,324	33,831	37,055	35,011
	(40,721)	(45,361)	(47,014)	(47,108)
One furnace	.379	.280	.257	.294
	(.487)	(.452)	(.440)	(.459)
Region	15,677	7,389	6,535	9,540
	(6,578)	(2,767)	(2,394)	(3,438)
Prod	11,322	5,014	4,059	5,830
	(5,637)	(2,490)	(2,394)	(3,502)
		Plants with Man-Hours Data		
Observations	93	74	64	68
$Months_0$	.452	2.446	2.688	1.382
	(1.247)	(3.120)	(3.182)	(2.726)
$Months_1$	5.333	6.040	5.936	5.456
	(5.184)	(4.096)	(3.833)	(4.749)
$Months_{2+}$	6.215	3.514	3.375	5.162
	(5.485)	(4.506)	(4.119)	(5.430)
MH	777,680	453,927	356,144	449,237
	(672,940)	(342,394)	(280,934)	(348,330)
MH_0	6,292	23,924	13,394	7,656
	(22,896)	(41,034)	(24,281)	(17,309)
MH_1	184,129	184,951	141,867	122,276
	(195,696)	(170,151)	(113,495)	(113,216)
MH_{2+}	587,655	245,052	200,725	319,305
	(766,924)	(375,691)	(276,059)	(408,266)

NOTE.—Entries are plant means. Standard deviations are in parentheses.

WEM_1 mean of 699; this means that the sum of employment over all months of one-furnace operation was just under 700 at the typical plant. In most of the empirical work we report here, we shall employ only three classes of plant operations: zero-furnace operations, single-furnace operations, and two or more (2+) furnace operations.[20] We also report the number of months each plant operates zero, one, or two or more furnaces as Months_0, Months_1, and Months_{2+}. Table 4 reveals that multiple-furnace operations are much less common in the early 1930s than in 1929, though still not negligible.

Even after aggregation to annual data, each plant can have up to four observations. There are a total of 312 annual observations (on average, three per plant) when we use WEM data and 299 when we use MH data.

IV. Measurement of Increasing Returns at the Plant Level

This section reports our preferred estimates of increasing returns. Section V includes additional details about the economic and econometric importance of various corrections.

Column 1 of table 5 presents estimates of scale economies from equation (1) using WEM as the measure of labor. Column 4 presents the analogous estimates using man-hour data.[21] The upper part of the table reports the actual regression coefficients. The coefficients of L_1 and L_{2+} represent the average productivity from running one furnace and two or more furnaces. We base inference on heteroscedastic-consistent standard errors since equation (1) suggests that our estimated production functions could indeed be heteroscedastic.

These preferred specifications for each labor measure have corrections both for permanent plant effects, τ_i, and for endogeneity in

and man-hours are a total. Thus if a furnace shuts down between the fifteenth and the last day, labor data will not coincide precisely with the furnace count. As a result, there is some labor incorrectly booked as in zero-furnace operations. Intramonth scale-up problems and the possibility that some of the labor on the margin may be used for bringing new furnaces on line mean that zero-furnace operations need not be unproductive.

[20] The number of three or more furnace operations in the 1930s was small. We shall return in Sec. VI to the question of whether there are short-run increasing returns beyond two furnaces. Table 4 shows clearly that zero-furnace operation was not trivial in 1931 and 1933. When we include a λ_0 term in (1), we cannot reject a value of zero for it. That the measured average physical product of labor from zero-furnace operations is zero is consistent with the view that workers at the plant at such times are doing fixed-cost setup and shutdown tasks.

[21] Plants with missing man-hours are deleted from the man-hour regression. When the two equations are run over the same sample, there are no noteworthy changes in the estimates.

TABLE 5

Panel Data Short-Run Production Function

	L = WEM			L = MH		
	(1)	(2)	(3)	(4)	(5)	(6)
L_1	98.1	91.8	67.5	.593	.509	.364
	(4.27)	(23.30)	(8.93)	(3.50)	(16.24)	(11.82)
L_{2+}	137.4	156.0	142.7	.740	.616	.637
	(5.89)	(24.09)	(18.90)	(4.30)	(33.53)	(12.61)
$L \times$ plant dummies	Yes	Yes	No	Yes	Yes	No
Sample selection correction?	Yes	No	No	Yes	No	No
Returns to scale	39.33	64.2	75.1	.147	.107	.273
$(\lambda_{2+} - \lambda_1)$	(4.25)	(6.36)	(6.14)	(3.01)	(3.49)	(4.97)
Statistics:						
Observations	312	312	312	299	299	299
R^2	.95	.93	.75	.96	.97	.77

Note.—*t*-statistics are in parentheses.

the numbers of furnaces operated. The economic logic of these two econometric points is similar. If one were not to control for permanent differences in productivity across plants, estimation could be biased toward finding scale economies since more productive plants may be systematically larger. The high-τ_i plants have high productivity and low cost. This is Friedman's "permanent plant hypothesis" (see Friedman 1955; Rosse 1970). If, plausibly, more productive plants tend to have larger-scale operations (i.e., $f_{it} > f_{jt}$, $L_{it} > L_{jt}$, and $Q_{it} > Q_{jt}$), uncorrected estimates have upward bias in estimates of $\lambda_f - \lambda_{f-1}$ and the specious implication of scale economies. To correct for this bias, plant-specific productivity fixed effects, the dummies τ_i, are included in columns 1 and 4.[22] The hypothesis that plant constants are jointly zero is decisively rejected in both specifications.

Transitory shocks to plant productivity may have the same effect as permanent ones. Again, plants could be systematically larger in the period of their higher (transitory) productivity, biasing estimates of scale economies. In our data, "transitory" means 2 years. Such a lengthy interval time leaves plenty of scope for output reallocation to efficient plants. Our preferred specification investigates reverse causation due to transitory shocks, considering the endogeneity of capital (furnaces). Our procedure for this is a modified "Heckit." The modifications, described in Appendix B, arise because we have three

[22] The mean of these effects, which are plant dummies interacted with L, is constrained to be zero.

states: zero, one, or two or more furnace operations. This correction is also highly significant statistically.

Our measure of scale economies is the increase in average product of labor when a plant moves from one to two or more furnace operations ($\lambda_{2+} - \lambda_1$). In column 1, moving a wage earner month from one-furnace operations to two-furnace operations adds about 39 tons of pig iron of output, increasing productivity nearly 40 percent. Column 4, with the man-hours measure, confirms this finding. Moving a man-hour from a one-furnace to a two-furnace operation adds about 0.14 ton for each hour. Both of these estimates are large enough to be economically significant. The labor productivity penalty for inefficiently slow production rates is 20 percent in the MH model and over 29 percent in the WEM model. Before turning to the aggregate implications of these figures, we examine measurement issues more closely in the next two sections.

V. The Labor Measures Puzzle

Section II noted very different composition effects in MH and WEM data. Much of the divergence can be resolved by examining the rest of table 5. Columns 2 and 5 contain specifications leaving out the correction for transitory shocks (the endogenous number of furnaces). Columns 3 and 6 are similar but also omit the plant-specific fixed effects. Note first that the importance of correcting for plant fixed effects is much larger in the MH data than in the WEM data. Scale economies drop by about one-seventh when WEM data are used but by more than one-half when MH data are used. The difference results from a composition effect: plant closure is much more correlated with MH inefficiency than with WEM inefficiency. For example, the plants using 12-hour shifts (high MH/WEM) all close after 1929, are very inefficient, and run one furnace. Not correcting for τ_i confuses within-plant productivity movements with between-plant movements. This explains why surviving-plant MH productivity was more procyclical than overall MH productivity in table 1.

When we add the sample selection bias correction, the effects again vary between MH and WEM models. The correction has the wrong sign in the MH model, though the magnitude (0.107 tons per man-hour vs. 0.147) is not large given the precision with which the coefficients are estimated. The WEM model performs as expected. There are two main competing hypotheses to explain this: that there is something wrong with our modified Heckit procedure or (as we thought) that the MH data are suspect.

Our modified Heckit procedure assumes normality of the errors. This might be incorrect. There are two solutions: simply instrumenting for Months_f and estimating semiparametrically both the probit and the Heckit. The conditions under which instrumenting will work are straightforward, namely $\tau_{iy1} = \tau_{iy2}$. The perfect correlation of the errors across the number of furnaces operating amounts, in terms of the economics, to the assumption that productivity shocks shift the short-run average cost curve (SRAC) vertically but never change its slope. If one instruments with the demand variables from Appendix A (in specifications not reported here), estimated scale economies actually rise dramatically. This leads us to conclude that the instrumental variables procedure is incorrect.[23] The "series" approach to the seminonparametric estimation leaves the probit and the estimated scale economies unaltered, apparently because the normal approximation works reasonably well.[24] The problem of shocks to the slope of the SRAC seems more severe, so that dummy endogenous variables methods work better. The WEM data look better, so as our initial data audit suggested, the WEM estimates seem likely to be more reliable than the MH ones.

Correcting for all the endogenous factor issues, permanent and transitory, taken together involves substantially reducing the estimated scale economies. Reverse causation from productivity to scale seems important here. Yet all the corrections do not reduce estimated scale economies to zero. It still appears that open plants are high on their SRAC. Production rates were slow, and fixed labor was not being fully exploited.

VI. Labor Hoarding

If productivity falls because labor is being hoarded, apparent scale economies are not a production function phenomenon but rather a labor market phenomenon (see Fay and Medoff 1985). Retaining the skilled labor associated with each plant may merely be dynamic cost minimization. Much of the labor used to operate blast furnaces required some skill. But the skill does not appear to have been furnace- or establishment-specific, merely industry-specific, resulting in a motive for labor hoarding. Another motive for hoarding could be political, to maintain the appearance of employment (Hoover 1952, pp. 43–44; see also Myers 1934, 1:391, 2:241). Job-sharing programs

[23] If the instrumental variables regression is correctly specified and if the reverse causation of high productivity leads to a large scale, instrumenting should lower rather than raise the estimated scale economies.

[24] We follow Olley and Pakes (1992) closely at this point. In the ordered probit estimation, we include power terms in $(Z_i)\beta$. The regression is reported as col. 2 of table 6 below.

in iron and steel were certainly widely publicized at this time (Daugherty, de Chazeau, and Stratton 1937, p. 165).[25]

For our industry in our time period, labor would need to be quite persistent for labor hoarding to explain the measured scale economies since productivity was low for a period of many years. Was it sufficiently persistent? This can all be made more operational by referring to the model of output adjustment discussed in Section II. If output is reduced by shortening production runs, labor hoarding *is* retaining laborers between runs. Slow production rates and labor hoarding involve a stock of laborers appropriate for the expected, not actual, production rate.

Graphs of employment, man-hours, and furnaces show that labor was in fact quickly responsive to the number of furnaces in operation. Figure 1 examines the within-year shifts in labor and furnace utilization at two plants. In each of the figure's two graphs, the calendar month is shown on the horizontal axis, and an index of WEM, an index of MH, and f are all on the vertical axis (the labor inputs were rescaled for readability). When the number of active furnaces changes, employment and MH adjust within a few months. This phenomenon appears to be systematic, as a simple partial adjustment model for L as a function of f (not reported here) suggests rapid adjustment. Effort intensity moved the wrong way to explain this away.[26]

To look for evidence of job sharing, we examine further short-run scale economies beyond two-furnace operations. There is some evidence of job sharing in larger plants. Political pressure in favor of job sharing would likely be most effective in the largest and most visible plants. Estimation of the production function extending the furnace categories to cover zero, one, two, three, and four or more furnace operations reveals that scale economies are consistently estimated to continue after two furnaces with WEM data whereas scale economies are largely exhausted after two furnaces with MH data (estimates are available from the authors on request). This pattern suggests that job sharing seems to have occurred principally at the larger firms and plants. But job sharing was not an economically important phenomenon since aggregate productivity movements

[25] Such programs were not, of course, unique to iron and steel.

[26] The effort intensity of most jobs in a blast furnace is determined entirely by the operations cycle of the furnace. If the furnace is running, it needs to be fed ore, coke, and limestone continuously. For jobs on the input side, then, potential intensification of effort is not an interesting problem. For workers in potentially shared crews, effort intensity clearly dropped. There were fewer furnaces to tap and less equipment in use that might break down and require repair. Fanning reduced tapping frequency and diminished the wear and tear on most equipment.

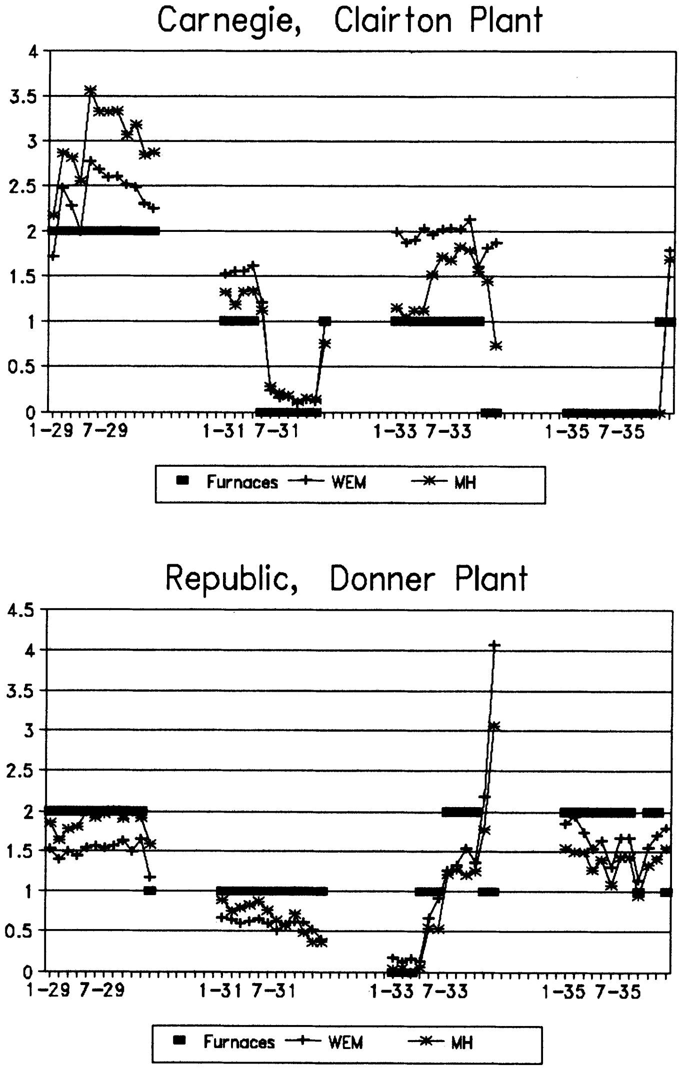

FIG. 1.—Labor and furnace usage at two plants

were driven by shifts from more than two to one-furnace operations. For this change in production rate, both labor variables lead to measured scale economies.

It is important to the argument of this paper to eliminate labor hoarding as an explanation. Our measured short-run production functions show substantial SRIRL. In what follows, we rely on the inference from this that shutting down more plants and reallocating their output to remaining facilities would raise aggregate productivity. The inference would be wrong if labor hoarding motives were the source of most of measured SRIRL. But they were not.

VII. Lack of Competition, SRIRL, and Aggregability

Lack of competition among plants would permit an equilibrium pattern of too slow production at a large number of plants rather than efficient-scale production at fewer plants. It would also block the exploitation of efficiency variations across plants. The usual evidence about competition, such as pricing or margins, is not available in this case, since blast furnaces were typically owned by their customers (steel mills) and the transfer prices for molten iron are not meaningful. But two other kinds of evidence are both salient and available: the apparent degree of geographical localization and the cost savings available from output reallocations.

Let us consider an ordered probit for the number of furnaces open. For assessment of the blockages to competition, the important variables are those we call "demand," which measure regional and product-specific demand.[27] Either an increase in demand for the steel products to which a blast furnace plant is linked or an increase in steel demand in the blast furnace's region leads to an increase in output. Table 6 shows that both coefficients are statistically significant.

To assess their economic significance, we performed simple within-sample predictions. Consider a multiple-furnace plant in Ohio associated with a plant in the heavy steel products category and having typical (sample mean) technology variables. Changing the values of Region and Prod from their 1929 to their 1933 values decreases the probability of operating two or more furnaces by more than half. Other baseline values lead to a range of similar calculations. Local demand definitely affected production rates and the extent to which scale economies were exploited.

On the other hand, there is also evidence of statistically significant cost effects, notably in the coefficient of the horsepower variable. This indicates that localized competition is not the only economically important force for this industry. There is some shifting of output in the direction of the lower-cost plants, consistent with the result above that correcting for the endogeneity of factors lowers measured scale economies. So we see a pattern of exploitation of scale economies that is intermediate between nationwide perfect competition and a cluster of local monopolies.

[27] Precise definitions for our regional and product variables are given in App. A.

TABLE 6

Reduced Forms for Factors
Dependent Variable = Months

Explanatory Variables	Estimation Method: Ordered Probit (1)	Series Ordered Probit (2)
C	−.36	−.30
	(−2.39)	(−1.13)
Demand:		
Prod	.70E-4	.63E-4
	(4.58)	(2.30)
Region	.45E-4	.43E-4
	(3.79)	(2.28)
Technology:		
Horsepower	.18E-4	.15E-4
	(6.67)	(11.9)
One furnace?	−.87	−.86
	(6.27)	(−3.75)
H	1.19	1.19
	(28.2)	(19.8)
α_2		12.1
		(.53)
α_3		−501.7
		(−.57)
α_4		3,424
		(.40)
Log(L)	−333.0	−332.8

Note.—t-statistics are in parentheses. See App. B for details.

To get more of an idea of the relative economic significance of the sources of inefficiency, we ask two questions: How much (short-run) inefficiently small scale does industry equilibrium tolerate? How much (short-run) inefficient heterogeneity does industry equilibrium tolerate? These questions are best answered with a different type of evidence.

In the counterfactual calculations described below, we assume that the output of blast furnace plants can be costlessly stored and transported. These assumptions allow us to reallocate output as a single perfectly competitive market would. This results in dramatically lower production costs than those that were actually observed. We therefore conclude that these assumptions are wrong; that is, output was not costlessly substitutable between plants.

In the first counterfactual, we imagine that each plant open during the trough operates at its efficient scale. In particular, all production rates are set at the 1933 2+ rates rather than the actual ones. Operationally, imagine that each plant runs an efficient number of furnaces for the first part of the year, costlessly storing the output. Using the estimated average physical product of labor from column 1 or 4 of table 5, we calculated aggregate counterfactual output per worker for 1931, 1933, and 1935. The figures are within 95–105 percent of the 1929 level.[28]

Our estimates of plant-level scale economies thus explain essentially all of the aggregate SRIRL. Why? Consider the peak year of 1929 and the near-trough year of 1933. Aggregate productivity fell by just under a third (table 1). The estimated gain in average physical product of labor in moving from one to two or more furnace operations is about 40 percent (table 5, col. 1). The fraction of wage earner months spent in zero- or one-furnace operations at the trough is roughly double that at the peak (table 4).[29] As a result, the plant-level average variable cost (AVC) diagram looks something like figure 2. Point *A* corresponds to 1929 behavior at the average plant. Something like *B* occurred in 1933 and 1931, because plants that operated were inefficiently high on the "average plant AVC" schedule. (The schedule is drawn with kinks at the points at which additional and less efficient equipment is brought on line.) In the counterfactual, a point like *C* is obtained: roughly 1929 efficiency at much lower output rates. While the details of the counterfactual calculation could be changed, the robust point is that the short-run average cost curve is steep, and the industry equilibrium systematically led to low production rates.

A second counterfactual explores the effect on industry productivity of reallocating output on a nationwide basis to the most productive plants. With the benchmark assumptions of substitutable output and zero transportation costs, the industry equilibrium would exploit cost differences by operating at full capacity at the most efficient plants. When the estimated fixed effects are used to order the 66 open plants, the most efficient 17 plants can produce 13.2 tons of iron operating at full capacity. In 1933, these plants made 8 million tons less than they did in 1929 and had a mean plant-level productivity of 207 tons per wage earner month. The remaining 49 plants made 8 million tons, with an average productivity of 101 tons per wage earner month. Reallocating output from the less efficient plants to the most efficient plants roughly doubles the productivity of the reallocated laborers.[30] That would leave industry productivity about 75 percent higher at the cyclical trough. Nationwide perfect competition,

[28] The same calculation for the less reliable man-hour data leads to countercyclical productivity.

[29] A slightly more complex calculation for all three years also moves some production out of zero-furnace operations, with a small savings in setup costs. A spreadsheet detailing this calculation is available from the authors.

[30] The tremendous variation in production functions found in Bertin (1994) suggests that this figure is not implausibly large.

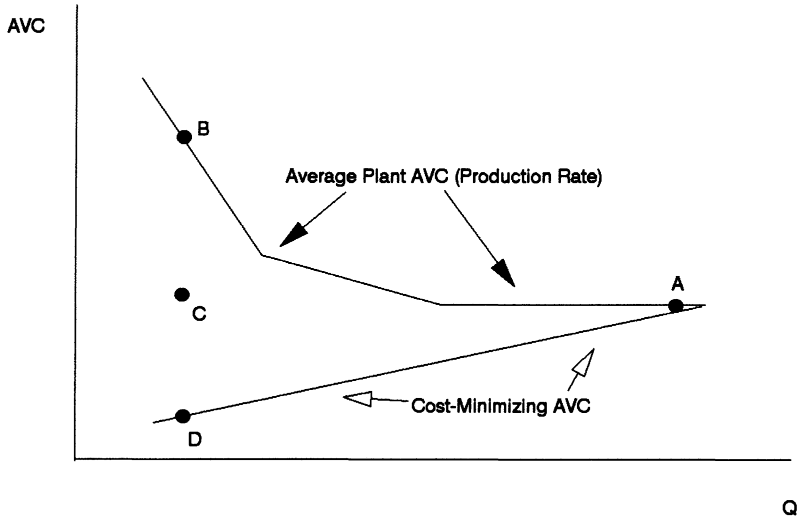

FIG. 2.—Average costs: actual and two counterfactual costs

then, would have led to a point something like D in figure 2. Since the cost savings from output reallocation were huge, transportation and storage costs must have been very large.

The lower, upward-sloping AVC in the picture is the industry AVC that comes from allocating output in a cost-minimizing way. It rises because older plants, which are utilized only as we move to the right, are less efficient than newer ones. Acyclical productivity (like C) is not necessarily an economic benchmark for an industry with heterogeneous plants. The cost-minimization benchmark involves positively countercyclical productivity at a point like D.

To rebut this conclusion, only proof of labor hoarding would do. We demonstrated in Section IV that productivity at the plant level is procyclical. There are three possible explanations of this. The first is that labor hoarding (or other similar dynamic labor behavior) could mean that the plants are not on the production function. This definitely would rebut the conclusion of this section. Unfortunately, it is refuted by the data. The second is that heterogeneity in plant efficiency could be correlated with output, so that the measured scale economies are partly specious. The third is that the measured scale economies could be correctly measured. For purposes of the economic arguments of this section, it does not matter whether the second or the third explanation is the correct one. In either case, there are substantial opportunities for reallocation of output across plants to lower production costs. If industry equilibrium did not exploit those opportunities, some other costs or rigidities prevented it.

VIII. Conclusions

Our evidence and arguments about blast furnaces in the Great Depression have three analytical themes. Together, they provide a coherent account of the industry equilibrium.

The first of these concerns SRIRL. The aggregate productivity drop in this industry did indeed reflect short-run scale economies. Inefficiencies arose through production at too low a rate (i.e., too much single-furnace operation) and through the operation of relatively less efficient plants. No important inefficiency arose from too short production runs. Operation at efficient rates at all the plants open at the trough would have shown flat productivity. If production had occurred at efficient rates at the most efficient plants, productivity would have risen significantly at the trough.

The second theme concerns the respective roles of coordination and competition. This mixture of inefficient and efficient outcomes inside blast furnace plants reflects their external environment. Each blast furnace plant coordinates with adjacent steel plants in smooth, if slow, production over time. The intermediate-output storage and shipments that would permit efficient blast furnace run rates and use only the most efficient furnaces are prohibitively costly, reflecting rigidities in physical capital. The stock of blast furnaces, steel plants, and rail lines in place was fixed. Its locational pattern would approximately minimize production plus transport costs at output levels five to 10 times those of the Depression. Given these rigidities, a general pattern of localization was likely efficient overall, though it implied substantial inefficiencies in the plants we study.

The final theme concerns aggregation and composition effects. We find that plant-level production functions exhibit SRIRL whether labor is measured by head count or man-hours. Yet aggregate productivity is flat on the man-hours measure and procyclical with the head count measure. The only important departure from localization was closure of the least efficient, small plants that ran 12-hour shifts and contributed many man-hours but few tons of output to 1929 totals. Their closure, and not job sharing, drives the movement in aggregate hours per worker over time. Job sharing, apparently more a public relations than labor relations practice, was confined to the largest and most visible plants.

The low degree of interplant substitutability is central to the economic explanation of all the phenomena we observed. With high substitutability, industry equilibrium would exploit plant-level scale economies. With high substitutability, industry equilibrium would heavily utilize the more efficient plants at the cyclical trough. Our finding that procyclical productivity at the industry level was a direct reflection of plant-level behavior is unusual. Studies based on modern data emphasize lack of aggregability (see n. 2 above). So does the Bresnahan and Raff (1993) study of motor vehicle manufacturing in the Depression era. But Depression-era motor vehicle manufacturing firms, for all the protective product differentiation that many of them deployed, had a national market. The Depression caused a shakeout of very large proportions. Since the firms that left seem to have had, on average, a production technology very different from that of the firms that stayed, the whole episode suggests the phrase "the cleansing effects of recessions" to modern analysts (see Schumpeter 1939, 1942; Caballero and Hammour 1994). The scope of plants' markets—the degree of substitutability among their output—constrains the extent of that cleansing power. Substitutability affects industry evolution.

The low level of substitutability in Depression-era blast furnaces might be anomalous. After all, the transportation and communication infrastructure to support high substitutability has been improving. Yet other modern studies, notably Carlton (1991), document important relationship-specific investments limiting interfirm substitutability. Perhaps relationship-specific investments are only a modest blockage to ex post substitutability in modern times, despite their importance in recent economic theory. Yet we do not now have a good quantitative understanding of what "low" substitutability is or even what a "large" departure from aggregability would be. Understanding the quantitative relationships between substitutability and industry outcomes such as aggregability over their typical ranges will be an important step in turning the descriptive aggregation literature into an economic one.

Appendix A

Data Sources

The blast furnace data set is an unbalanced panel of plants, constructed from the *Census of Manufactures* and from *Iron Age*. There are a total of 312 annual observations with WEM data and 299 with MH data.

Labor.—The census provides two *monthly* measures of labor input, the number of wage earners employed on the fifteenth of each month (WEM) and monthly man-hours (MH). The man-hour data are of slightly lower quality. All plants in the census sample responded to the wage earner months question, but a smaller sample of plants answered the man-hour questions. In 1929, 10 plants either did not respond or had census notations in the margins indicating that the figures were incorrect. The census edits for later years were severely limited by tight budgets. We suspect that there are more uncaught errors in the man-hours data. Even for 1929, the census did not publish aggregate man-hours statistics.

Capital.—Capital inputs, including the number of furnaces, are partially documented in censuses; each plant reports total furnaces and "furnaces active at any time during the year." We have created a more precise monthly measure of furnace utilization with information from *Iron Age*, a steel industry trade publication. Each month *Iron Age* reports furnaces that were turned on or off during each month. Using these reports and the census data, we constructed a series of furnaces in operation at month end for each plant by combining the changes data from *Iron Age* with census data about active furnaces by perpetual inventory method. The *Iron Age* data match with data from the *Census of Manufactures* reasonably well. *Iron Age* did not provide information on four plants in our sample that produced charcoal pig iron. It provided erratic information for some of the smaller southern plants. However, these plants were primarily one-furnace plants, and in these cases, the series was easily constructed with data provided on the census form.

Output.—The product of this industry, pig iron, is largely homogeneous, and we take tons produced as our measure of Q. Although some iron is foundry grade and some is high-value specialty iron (made with manganese or silicon), these products are a small percentage of output. It does not matter empirically whether they and lower-value iron are treated as identical or different (Bertin 1994). Note that the census provides plant-level data on output only annually. Aggregate output of pig iron is available monthly from *Iron Age*.

Downstream product variables.—The production of the downstream steel mills affiliated with blast furnaces was researched by examining the manuscripts from the 1929 *Census for Steel Works and Rolling Mills*. Most output is transferred to an affiliated steel mill in molten form. Each blast furnace plant was matched with its affiliated steel plant by using the names of the plants and information from the directory of the American Iron and Steel Institute (AISI). For the 55 plants that had affiliated steel works, we examined the 1929 output of the steel plant to determine the primary product and then classified blast furnace plants by the primary product of the steel mill. These groups are heavy products including plates, bars, and structural shapes; light products including sheets, strips, and bands; pipe; rails; and wire. These groupings were created so that each serves a different segment of demand. Forty plants, not affiliated with a steel mill, produced foundry-grade iron and steel-grade iron sold in ingots. These plants were classified as merchant iron producers. Eight plants produced primarily high-valued specialty iron. Annual production by category was collected from aggregated data published by the *Census of Manufactures*.

Region variables.—Ideally, regional variables would reflect logically defined regions in which it would be easy to transport the heavy-weight, low-valued steel. However, aggregate data are available only by state. Our five regional variables are defined as South if the plant is located in Alabama, Kentucky, or Tennessee; Mid-Atlantic if the plant is located in Maryland, New York, Pennsylvania, West Virginia, or Virginia; New England if the plant is located in Massachusetts; West if the plant is located in Colorado or Utah; and Central if the plant is located in Illinois, Indiana, Michigan, Minnesota, or Ohio. Regional variables were aggregated from tons of finished and rolled steel by state reported in the *AISI Annual Statistical Report.*

Technology variables.—Various operations variables characterize the technology at the plant level. Horsepower is the horsepower of motors and prime movers reported in the 1929 census. Horsepower is used to run auxiliary equipment and to generate electricity to run stoves and blowers.[31] Plants also report the total number of furnaces at the plant, and we created a dummy variable for plants with one furnace. In addition, specifications of the furnaces were collected from the AISI directory. These data included the hearth size for each furnace, its age, and its capacity. Variables including average vintage, average capacity, and average hearth size were used in some specifications of the ordered probit.

Appendix B

Heckit Details

The number of furnaces operating could be endogenous if transitionally lower costs lead to higher production rates at particular plants. This would make the variables N_0, N_1, and N_2 ($\text{Months}_0/12$, $\text{Months}_1/12$, and $\text{Months}_{2+}/12$) (which enter into eq. [1] interacted with labor variables) be correlated with the error in that equation. This problem is very similar to endogenous dummy variable problems. The difference is that we have three ordered states, 0, 1, and 2+, and that while $N_0 + N_1 + N_{2+} = 1$, all three can be nonzero fractions rather than zero-one dummies. These differences lead to minor alterations in standard dummy endogenous variable methods.

Assume that the latent variable determining the number of active furnaces is distributed joint normal with the production function errors. Under this assumption, we can undertake a two-step procedure. First, we estimate an ordered probit for N_0, N_1, and N_{2+}:

$$\Pr(\text{Months}_0) = \Phi(-z\beta),$$

$$\Pr(\text{Months}_1) = \Phi(H - z\beta) - \Phi(-z\beta),$$

$$\Pr(\text{Months}_{2+}) = 1 - \Phi(H - z\beta).$$

This will hold if a latent variable $z\beta + \epsilon_f$ crosses a threshold of H for two or more furnace operations and a threshold of zero for one or more operating furnaces. The key regressors, z, are the regional and product demand measures and two technology variables. (See table 6 for estimates of the probit and App. A for a complete description of the variables.) The main purpose of this first stage is to provide estimates of the expected value of the error

[31] Allen (1977, p. 610) provides support for the use of the horsepower variable as a proxy for capital stock. He reports that most of the cost of building a blast furnace is accounted for by the blowing engines and associated equipment. This would not have changed between his period and ours.

in the furnaces equation, ϵ_f. We then include three new regressors in the production function. They are the expected value of the error term with f furnaces operating multiplied by the number of months the plant operated f furnaces. With these variables included, and under the assumption of joint normality, production function estimates will be consistent even if the number of furnaces is endogenous.

References

Alchian, Armen A. "Costs and Outputs." In *The Allocation of Economic Resources: Essays in Honor of Bernard Francis Haley*, by Moses Abramowitz et al. Stanford, Calif.: Stanford Univ. Press, 1959.

Allen, Robert C. "The Peculiar Productivity History of American Blast Furnaces, 1840–1913." *J. Econ. Hist.* 37 (September 1977): 605–33.

American Iron and Steel Institute. *Directory of Iron and Steel Works in the United States and Canada*. New York: American Iron and Steel Inst., various issues.

Bernanke, Ben S. "Employment, Hours, and Earnings in the Depression: An Analysis of Eight Manufacturing Industries." *A.E.R.* 76 (March 1986): 82–109.

Bernanke, Ben S., and Parkinson, Martin L. "Procyclical Labor Productivity and Competing Theories of the Business Cycle: Some Evidence from Interwar U.S. Manufacturing Industries." *J.P.E.* 99 (June 1991): 439–59.

Bertin, Amy L. "Competition and Productivity in the Depression-Era Steel Industry." Ph.D. dissertation, Harvard Univ., 1994.

Bresnahan, Timothy F., and Raff, Daniel M. G. "Intra-industry Heterogeneity and the Great Depression: The American Motor Vehicles Industry, 1929–1935." *J. Econ. Hist.* 51 (June 1991): 317–31.

———. "Technological Heterogeneity, Adjustment Costs, and the Dynamics of Plant Shutdown Behavior: The American Motor Vehicles Industry in the Time of the Great Depression." First Boston Working Paper no. FB-93-09. New York: Columbia Univ., Grad. School Bus., March 1993.

Caballero, Ricardo J., and Hammour, Mohamad L. "The Cleansing Effects of Recessions." *A.E.R.* 84 (December 1994): 1350–68.

Camp, James M., and Francis, Charles B. *The Making, Shaping and Treating of Steel*. 4th ed. Pittsburgh: Carnegie Steel Co., 1925.

Carlton, Dennis W. "The Theory of Allocation and Its Implications for Marketing and Industrial Structure: Why Rationing Is Efficient." *J. Law and Econ.* 34, no. 2, pt. 1 (October 1991): 231–62.

Carnegie-Illinois Steel Corporation. Duquesne Works. "Fifth Joint Conference: Employees' and Management's Representatives." 1936. USX Corp., Duquesne Works Records. University of Pittsburgh Archives of Industrial Society Accession 87:3, box 26.

Chenery, Hollis B. "Engineering Production Functions." *Q.J.E.* 63 (November 1949): 507–31.

Clements, Fred. *Blast Furnace Practice*. Vol. 3. *Operation and Products*. London: Benn, 1929.

Daugherty, Carroll L.; de Chazeau, Melvin G.; and Stratton, Samuel S. *The Economics of the Iron and Steel Industry*. 2 vols. New York: McGraw-Hill (for Univ. Pittsburgh Bur. Bus. Res.), 1937.

Davis, Steven J., and Haltiwanger, John C. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Q.J.E.* 107 (August 1992): 819–63.

Dunne, Timothy; Roberts, Mark J.; and Samuelson, Larry. "The Growth and Failure of U.S. Manufacturing Plants." *Q.J.E.* 104 (November 1989): 671–98.

Fay, Jon A., and Medoff, James L. "Labor and Output over the Business Cycle: Some Direct Evidence." *A.E.R.* 75 (September 1985): 638–55.

Friedman, Milton. "Comment." In *Business Concentration and Price Policy.* Princeton, N.J.: Princeton Univ. Press (for NBER), 1955.

Hamermesh, Daniel S. "Labor Demand and the Structure of Adjustment Costs." *A.E.R.* 79 (September 1989): 674–89.

Haven, William A., and Buell, W. C., Jr. "Blast Furnace Activities in 1932." *Blast Furnace and Steel Products* 21 (January 1933).

Hogan, William T. *Economic History of the Iron and Steel Industry in the United States.* 5 vols. Lexington, Mass.: Lexington, 1971.

Hoover, Herbert. *The Memoirs of Herbert Hoover.* Vol. 3. *The Great Depression, 1929–1941.* New York: Macmillan, 1952.

McGuckin, Robert H. "The Importance of Establishment Data in Economic Research." Center for Economic Studies Discussion Paper no. 93-10. Washington: Census Bureau, August 1993.

Meyer, John R., et al. *The Economics of Competition in the Transportation Industries.* Cambridge, Mass.: Harvard Univ. Press, 1959.

Myers, William Starr, ed. *The State Papers and Other Public Writings of Herbert Hoover.* Garden City, N.Y.: Doubleday, Doran, 1934.

Olley, G. Steven, and Pakes, Ariel. "The Dynamics of Productivity in the Telecommunications Equipment Industry." Center for Economic Studies Discussion Paper no. 92-2. Washington: Census Bureau, February 1992.

Rosse, James N. "Estimating Cost Function Parameters without Using Cost Data: Illustrated Methodology." *Econometrica* 38 (March 1970): 256–75.

Schumpeter, Joseph A. *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process.* New York: McGraw-Hill, 1939.

———. *Capitalism, Socialism, and Democracy.* New York: Harper, 1942.

Sweetser, Ralph H. *Blast Furnace Practice.* New York: McGraw-Hill, 1938.

Williamson, Oliver E. *Markets and Hierarchies, Analysis and Antitrust Implications: A Study in the Economics of Internal Organization.* New York: Free Press, 1983.