

The Strange and Awful Path of Productivity in the U.S. Construction Sector

Austan Goolsbee*
University of Chicago
Booth School of Business and NBER

Chad Syverson
University of Chicago
Booth School of Business and NBER

February 2022

Abstract:

Aggregate data show a large and decades-long decline in construction sector productivity. This decline in such a large sector has had a material effect on secular productivity growth for the economy as a whole. Prior work has focused on the role of potential measurement problems in construction, particularly output deflators in the measurement of productivity. This paper brings some new evidence to bear on the industry's measured productivity problems and suggests that measurement error is probably not the sole source of the stagnation. First, using measures of physical productivity in housing construction, productivity is falling or, at best, stagnant over multiple decades. Second, there has been a noticeable decline over time in the efficiency with which construction firms translate materials inputs into output, and a corresponding shift toward more value-added-intensive production. Third, using state-level data, we do not find evidence of patterns of within-industry reallocation that might be expected of efficiently operating input and output markets. States with more productive construction sectors do not see growth in their shares of total U.S. construction activity; if anything, their shares fall. This may point to frictions in these markets that slow or stop what is in many other markets an important channel for productivity growth.

* Contact info: goolsbee@chicagobooth.edu and chad.syverson@chicagobooth.edu. We would like to thank Pengyu Ren, Nicole Bei Luo and Rebecca Goldgof for their outstanding research assistance and participants at the Hutchins Center at Brookings conference for comments. We thank the Initiative on Global Markets at the University of Chicago's Booth School for financial assistance.

Aggregate productivity growth is the prime determinant of long-run growth in income per capita, and an economy's productivity growth reflects the productivity growth rates of the industries and sectors within it.

Yet despite decided growth in aggregate productivity for the U.S. economy as a whole, the U.S. construction sector has diverged considerably. Indeed, for decades now, measures of labor and total factor productivity (TFP) in the sector have trended downward. To be clear, the raw BEA data suggest that the sector has become *less* productive over time. A lot less productive: value added per worker in the sector was about 40 percent lower in 2020 than it was in 1970.

Economic researchers have remarked on these troubling patterns before; see, e.g., Stokes (1981), Allen (1985), Schriver and Bowlby (1985), and Sveikauskas et al. (2016, 2018). The problem has also attracted the attention of broader analysts and audiences; see Changali et al. (2015), Economist (2017), Potter (2021), and Smith (2021) for examples. A great deal of attention has gone to the issue of whether measurement problems explain the sector's disappointing performance.

In this paper, we update some of this previous work and extend it to some new data sources and hypotheses. Together, these new approaches seem to reinforce the view that the poor performance is not just a figment of measurement error. We see similar stagnation using physical measures of productivity that are not dependent on price deflators. We also see that firms' abilities to turn materials into output has deteriorated and we document real issues with the sector's adjustment mechanisms; there is little reallocation from low productivity places to high-productivity places.

I. The Core Issue with Construction Productivity

Figure 1 shows indexes of U.S. construction sector labor productivity and TFP from 1950 to 2020. For comparison, it also plots the same indexes for the overall economy.¹

Throughout the 1950s and well into the 1960s, both measures of construction sector productivity grew steadily. Indeed, they outpaced their whole-economy counterparts during that period. By 1970, however, the construction sector's labor productivity and TFP had both begun to fall. This downturn was not temporary; the decline has continued for the past half-century.

This downturn did not mirror the economy-wide productivity pattern. Productivity in the entire economy grew throughout the period (albeit with some well-documented accelerations and decelerations). By 2020, while aggregate labor productivity and TFP were 290 percent and 230 percent higher than in 1950, both measures of construction productivity had fallen *below* their 1950 values.

This is stunningly bad productivity performance for a major sector. It is brought into special relief when compared to the over nine-fold increase in labor productivity the manufacturing sector experienced during the same period. Manufacturing, like construction, deals with the configuration and assembly of physical objects in preparation for use as either inputs into production or final consumption. Yet the two sectors experienced totally different productivity trajectories over the past 50 years.

¹ We compute productivity using the Bureau of Economic Analysis national and industry accounts data. Labor productivity is real value added divided by full-time-equivalent (FTE) employees. The change in total factor productivity (which we convert to an index level) is the growth rate in value added minus a weighted sum of the growth rates of labor and capital. The weights in a given year are the cost shares of the factor inputs, averaged across the current and prior year in the usual Divisia fashion. Labor costs are total labor compensation plus 67 percent of proprietor's income. (We base the 0.67 multiplier on labor's historical share of income being roughly two-thirds.) Capital costs are the sum of depreciation, the product of the real interest rate and the current value of installed capital, and 33 percent of proprietor's income. Analogous productivity series computed using gross output rather than value added (and in the case of TFP, subtracting the implied contribution of intermediate inputs) show similar patterns. We explore the trends in gross output, value added, and intermediate materials in the construction sector in more detail below.

Construction's poor performance might be just a curiosity if it were a trivial fraction of economic activity, but it is not. The sector's value added averaged 4.3 percent of GDP between 1950 and 2020. This share, while experiencing fluctuations, has remained fairly steady over the long-run. Construction is a sizable share of aggregate output. It is large enough that its poor productivity performance noticeably drags down aggregate productivity growth. Construction labor productivity fell at an average rate of about 1 percent per year from 1970-2020. Had it instead *grown* at the (relatively modest) rate of 1 percent per year, annual aggregate labor productivity growth would have been roughly 0.18 percent higher.² This would have resulted in current aggregate labor productivity (and plausibly, income per capita) being about 10 percent higher than it actually was.

While we focus in this paper on U.S. construction sector, the problem of laggard construction productivity growth appears more widespread. In the 29 countries for which the OECD reports construction sector value added per employee growth data over 1996-2019, 16 of the countries—as well as the EU-27 area as a whole—saw negative average labor productivity growth in their construction sectors over that 25-year period.³

Even in countries that saw positive construction productivity growth, it typically substantially lagged overall productivity growth in their economies. Average labor productivity growth across all 29 countries was 0.4 percent per year, in contrast to those countries' average overall labor productivity growth rates of 1.6 percent over the same period. The phenomena we explore at a detailed level within the U.S. may well apply more broadly. Their full international extent is worthy of future inquiry.

² This is calculated by multiplying the notional 2 percentage point increase in construction labor productivity growth by its average Domar weight (the sector's gross output as a share of GDP) of 0.090. This weight—see Domar (1961)—is the first-order approximation of the contribution of a sector or industry's productivity growth to aggregate productivity growth.

³ We use the "Productivity and ULC by Main Economic Activity" data from the OECD for these computations.

II. The Failure of Capital-Based Explanations for the Decline

First, we address perhaps the most obvious potential source of poor productivity growth, a lack of capital investment. If laggard investment led to lower capital-to-labor ratios, it would directly affect labor productivity. This could also influence TFP, if TFP were partially capital-embodied (Hulten, 1992).

At first glance, though, this does not seem consistent with data for the construction industry. Figure 2 compares changes in the construction sector's current-value capital stock to that of the entire economy. While capital in construction did not grow as steadily as capital in the wider economy, construction's total capital stock growth since 1950 has actually been a bit larger, rising 7.8-fold as opposed to the 6.5-fold increase for total capital in the economy. There was no noticeable slowdown in capital growth after 1970, when sector productivity started to fall. Moreover, capital intensity—capital stock per FTE employee—did not fall in the sector relative to the overall economy. To the contrary, it actually rose a bit faster.

A more nuanced view of capital's role in construction productivity would include intangible capital, which by definition is not contained in the capital series above. We can look, however, at what the BEA terms the capital stock of intellectual property (IP) products in the industry. The BEA defines this category of capital assets as including the capitalized value of R&D, software purchases, and—perhaps less relevant to construction—artistic originals. Following work like Brynjolfsson, Rock, and Syverson (2021), we can use IP capital stock as a proxy for intangible capital in the industry (which would include things like know-how, organizational strength, trade secrets, buyer-supplier relationships, sector-specific human capital, and so on).

The U.S. construction sector is less IP-capital-intensive than the economy overall. For instance, in 2020 IP capital accounted for 4.0 percent of the sector's total capital stock, while the same ratio for the broader economy

was 7.4 percent. On a per-employee basis, IP intensity in construction is an order of magnitude smaller than in the rest of the economy.⁴

However, the IP capital data in the construction sector are not really consistent with intangible capital driving the path of construction productivity, either. The sector had *no* recorded IP capital in the national accounts until 1970. Thus, throughout the 1950-1970 period when the sector's productivity growth kept up with or even exceeded aggregate productivity growth, there was no strong indication that the sector was putting into place large intangible investments. Only after the productivity slowdown had begun did the sector begin to invest in IP capital, and despite accumulating such capital at a rate exceeding that in the overall economy, the sector's productivity level continued to diverge from aggregate productivity growth.

III. The Traditional Confound: Measurement Problems

Because productivity is a residual—the variation in output unexplained by variation in measured inputs—mismeasurement of either output or inputs will be labeled productivity, even if unrelated to the actual efficiency of the production process. Understatements of output or overstatements of inputs would cause measured productivity to be lower than true productivity, so it is important to think about whether such issues might explain the patterns in construction.

We start from the observation in Syverson (2017) that attributing a *change* in productivity to mismeasurement requires not just establishing the presence of mismeasurement, but also a change in the *amount* of mismeasurement in the necessary direction at the same time as the measured productivity change. If we are to explain the reversal in productivity growth in construction in the late 1960s as resulting from measurement problems, we

⁴ Though part of this enormous difference is accounted for by the inclusion of the residential housing stock in economy-wide capital, which is limited in its market-activity marginal product.

need to demonstrate that some combination of growth in the understatement of output or overstatement of input occurred in the late 1960s.

We first consider labor input measurement. There are several plausible channels for labor measurement difficulties in construction, including a higher than average frequency of employees working irregular hours, contractor labor that may be misclassified by survey respondents, and, especially in more recent decades, labor supplied by undocumented workers.

Figure 3 plots three series for labor inputs in the construction sector: the sum of full-time and part-time employees (from the BEA), FTE employees (also from the BEA), and total employment (from the Current Employment Statistics of the BLS). These series capture different elements of labor inputs, such as the implied differential treatment of hours per worker when comparing summed full-time and part-time employees with FTE employees.

None of the series shows an obvious kink in the late 1960s. Moreover, they track one another closely, with no divergence at the time construction productivity started falling. The average pairwise correlations among the three series before 1970 is 0.983 and is 0.999 from 1970 on. In addition, the average annual growth rates of all three series were *lower* after 1970 than before. A mismeasurement-driven productivity slowdown would imply inputs that are growing misstatedly fast—that measured labor accelerated rather than decelerated as the data seem to show.

Given our earlier discussion about the trajectory of the construction industry's capital stock, mismeasurement of capital also seems unlikely to explain the measured productivity declines. Moreover, capital measurement problems cannot be responsible for labor productivity mismeasurements, and labor productivity in construction exhibits the same broad pattern as does TFP (and likewise for intermediate/materials inputs).

Given the lack of obvious issues with measured inputs, much of the attention in the literature has centered on problems with measuring output.

As noted above, construction's value added share of GDP exhibited no long run trend over 1950-2020, so *nominal* construction value added grew at a

similar rate to the overall economy. The components of this ratio are both nominal values, so the measured slowdown must, mechanically, be coming from differences in the output price deflators.⁵

Indeed, the construction sector's output deflator and the aggregate GDP deflator clearly start to diverge after the late 1960s. From 1950-69, the average annual growth rates of the construction and GDP deflators were almost identical—2.40 percent and 2.42 percent. From 1970 on, however, the GDP deflator averaged annual growth of 3.37 percent, while the construction value added deflator grew 5.47 percent per year. This sustained 2.1 percentage point annual difference means that even if nominal construction value added grows at a similar rate to GDP, real construction value added would grow much more slowly, and perhaps even fall.

This kink in the construction price deflator is also perhaps large enough to quantitatively explain the observed downturn in construction productivity. Figure 4 repeats the plot of Figure 1, except it computes real construction output by deflating nominal construction value added with the GDP deflator rather than the deflator for the construction sector. The difference is striking. Now construction sector productivity grows throughout the entire 1950-2020 period, nearly matching the pace of overall productivity growth.

Clearly, the construction sector price deflator is, mechanically speaking, a key source of the downturn in measured construction productivity. Computing productivity by deflating nominal construction activity with the whole-economy deflator makes the construction productivity series look like overall productivity.

Note, however, that this is not evidence that we should use the aggregate GDP deflator for the construction industry or that the construction deflator is *wrong*. It could well be right and true construction productivity has, indeed, fallen for 50 years. However, the fact that it matters so much puts onus on

⁵ The potential role of price deflators in real construction output measurement in earlier decades was taken up by Gordon (1968) and Pieper (1991).

those interested in understanding construction productivity to check the deflator's accuracy.

This is why some of the most important work on construction productivity in the existing literature, like Sveikauskas et al. (2016, 2018), focuses on subsets of the construction industry where they can build more accurate output price deflators to explore productivity growth dynamics. This work has found that in such cases, productivity declines are not as extreme in the last 30 years as suggested in the aggregate data, and in some cases productivity was in fact growing (albeit modestly). It is this idea that brings us to search for settings where we can potentially get around the output deflator issues, as we do in Section IV below.

a. Potential Sources of Growth in Construction Output Prices

Given that the construction sector output deflator is rising considerably faster than prices in the rest of the economy, it is worth exploring whether there are visible sources of this relative price increase.

A logical place to look is at construction input prices. If these are also rising faster than the overall price level after 1970, it would be a clue as to the origin of the fast-rising prices and declining real output of the sector.

We first look at construction's intermediate inputs prices, using the sector's intermediate inputs deflator from the BEA's KLEMS database. Figure 5 plots this deflator and its change since 1950, along with the construction sector output and GDP deflators for comparison. It is readily apparent that construction intermediates price growth has been on the order of price growth in the overall economy rather than the faster growth measured in construction output prices. Its time path overlaps considerably with the GDP deflator and lags the construction output deflator.⁶

⁶ In separate work, we have looked at the PPIs of 10 major construction sector inputs over the sample period. These are conceptual components of the intermediates deflator above. While there is variation in the average growth rate of these inputs' prices, their mean tracks the GDP deflator, just as the intermediates deflator does.

We next look at relative labor prices by computing the nominal implied salary per worker for the sector and comparing its trajectory to that for the overall economy. We compute the implied salary as total labor compensation plus 0.67 times sole proprietor income, divided by the sum of employees and self-employed. The construction and overall-economy series track each other quite closely throughout 1950-2020. Construction salary growth is higher, but only slightly so (a difference of 0.14 percent in average annual growth rates). Moreover, there is no obvious divergence after 1970. Given the 2.1 percentage point average annual difference in construction and GDP deflators, this difference in relative labor prices is less than a tenth of that. Thus it may explain a part of the divergence in relative output prices, but this part is very small.

Direct measures of industry-specific capital prices are not readily available, but a prominent component of capital user costs is the depreciation rate. We compute depreciation rates for both the overall economy and the construction sector using BEA data on depreciation and capital stocks. While annual depreciation rates are higher in construction than for the overall economy, averaging 3.7 percent overall and 13.6 percent for construction over 1950-2020, there is no sign that *changes* in depreciation rates imply the relative user costs of construction capital are rising. In fact, the gap between construction and overall depreciation rates fell during the period, implying that this major element of capital costs was becoming relatively cheaper in construction. For the first five years of the period, the average difference between the overall and construction depreciation rates was 13.5 percent. For the last five years, it was 9.3 percent.

In sum, there is no obvious sign that changes in the relative prices of the construction sector's major inputs—intermediates, labor, and capital—drove the observed growth of relative construction output prices.

If relative unit input prices do not account for the increase in construction's relative output price, perhaps markups rose instead. To explore this possibility, we construct an approximation for markups by comparing

sector revenues to its total measured input costs. Construction input costs include expenditures on intermediates (which we observe directly in the KLEMS data), total payments to labor (the numerator in the average salary calculation above), and payments to capital (depreciation plus the real interest rate multiplied by the industry capital stock plus 0.33 times proprietor's income). We divide revenues by the sum of these input expenditures as a proxy for markups.

The average of this markup ratio over 1950-2020 is 1.016, indicating a modest average margin of 1.6 percent in the construction sector. This margin has a slightly increasing trend of 0.024 percent per year. This corresponds to growth in the average margin of 1.7 percent over the entire 70-year period. To the extent this higher accounting margin reflects larger markups over costs, some of the increase in construction's relative output price came from growing markups.

However, there are two important caveats to this interpretation. First, this overall change is miniscule compared to the overall growth in construction prices observed during the period. It is two orders of magnitude smaller than the 2.1 percentage point annual growth in the relative price of construction output after 1970. Second, this is an increase in construction's absolute markup. If markups are also increasing in the economy overall (see De Loecker, Eeckhout, and Unger, 2020, for evidence of this, at least since 1980), then there may be no increase in construction's *relative* markup and hence relative prices.

In sum, our analysis to this point indicates that only a small part of the large increase in relative output prices of the construction sector may come from increased relative wages and (perhaps) higher markups. The vast majority of the increase in construction's price deflator relative to average overall prices cannot be explained by increasing relative prices of construction inputs or markups.

This leaves a few possibilities to explain the divergence of the construction output and GDP deflators. One is that construction sector

productivity did indeed fall, raising unit input requirements. Even though prices per unit of inputs did not rise (or rise faster than other prices outside the construction sector), the unit costs of output would increase because construction firms would need to buy more inputs than before to build a unit of output.

Another possibility is that construction output price mismeasurement became worse in the late 1960s. If, say, the quality of construction output accelerated as did prices along with it, but the deflator does not properly account for that quality difference, the increase in output quality would have to occur in a way that was not correlated with increases in construction input prices. But, increased output quality from using more expensive inputs (better materials, much more skilled labor, etc.) is not consistent with the input price evidence above that input prices have not moved together with output prices.

IV. Measuring Productivity in Physical Units

We next turn to methods that do not rely on output price measurement. Our approach is to focus on a setting where output and productivity can be measured in physical units rather than expenditure requiring a deflator. Researchers have applied this approach in the broader productivity literature, especially the part dealing with producer microdata, but of course doing so relies on having a setting where individual units are measured and there is reasonable homogeneity.⁷ In this section we consider a part of the construction industry where those conditions might hold (loosely): home construction.

Output data are available for the U.S. housing construction industry not just in value but in *number of housing units* as well. Exploring the evolution of the industry's productivity measured in houses per unit input, rather than deflated house value per unit input, lets us track a measure of the efficiency of the industry's production process that does not rely on a price index.

⁷ A few recent examples include De Loecker et al. (2016); de Roux et al. (2021), and Orr (forthcoming).

This advantage does come at an obvious cost, however. If the physical housing units change in quality over time, then a given number of housing units in one period will not be the same “real” quantity of housing as in another period. And we know that the average attributes of houses have changed over time; for instance, they have trended toward larger floor areas. Given the amount of data available on homes, however, we can also explore whether such changes materially affect the conclusions from using housing-unit-based productivity measures.

Our output data come from annual housing completions reported by the U.S. Census Bureau. These begin in 1968 and are broken out by number of units per building: 1 unit, 2-4 units, and 5+ units. We combine this with CES employment data for the housing construction industry. After 1990, the employment data are reported separately by single-family and multi-family housing construction sub-industries. An older employment series for all residential construction activity extends from 1972-2002. It unfortunately precludes a breakout into single- and multi-family construction and also includes employment in repair and remodeling that is not reflected in the Census housing construction output data. Nevertheless, it is instructive to compare the (combined single- and multi-family housing construction) results from this older series to the more precise values starting in 1990, and we do so below.

Figure 6 shows the resulting annual productivity series, measured in housing units per employee. Looking first at the two post-1990 series, while there are some notably large swings in the average number of housing units built per employee in both the single-family and multi-family housing subindustries (the largest in each being troughs during the Great Recession), there is no discernable upward trend over the three decades. Linear time trends fit to both series have statistically significant and negative slopes. For single-family housing, it is -0.023 (s.e. = 0.007) per year, and for multi-unit housing the slope is -0.081 (s.e. = 0.037) per year. Average productivity in the first five years of the single-family (multi-unit) series is 2.63 (10.6) units per

employee and in the last five years is 2.35 (10.0) units per employee.⁸ Comparing the magnitudes of the negative trend estimates to the levels of these productivity measures, housing productivity measured in units per employee is declining at a rate of roughly -1 percent per year. This is about the same as seen above in the broader construction sector and using deflated value to measure output.

The older, 1972-2002 series exhibits, if anything, a stronger downward trend than the other two series.⁹ The linear trend slope is -0.042 (s.e. = 0.005), and average units per worker in the first and last five years of the period are respectively 3.1 and 2.1. Though it is worth noting that during the 1990-2002 period of overlap, this productivity measure is basically level.

These data offer a complementary view into stagnant or declining construction productivity uninfluenced by price mismeasurement. The new residential construction industry, at least, does not seem to be becoming more efficient at building housing units.

As noted, however, the usefulness of physical-quantity-based productivity measures depends on the homogeneity of the units across settings and time. If housing units are getting better, the industry may be becoming more productive in terms of *housing quality* produced with a unit of input, even if the number of housing units per input has not increased.

One of the more obvious changes in housing units over the past several decades has been the increase in floor space per unit. Houses have been getting bigger. Data on the average square footage of completed single-family housing units is available from 1973-2020, and it rose from 1660 ft² to 2480 ft² over that period. The average size of multi-family units rose as well, but more moderately, from 1021 ft² to 1121 ft².

⁸ As seen in the figure, since the Great Recession, houses per worker recovered to roughly their pre-Great Recession levels. Time will tell if this growth extends beyond the prior level or instead merely returns to the long-run average.

⁹ Because the employment data for this series includes not just single- and multi-family housing employment but repair and remodeling employment as well, this extra employment not generally dedicated to new home construction will cause the level of this productivity series to be lower than either of the other two.

Recomputing our unit productivity series for residential construction as square footage of housing per employee, of course, exhibits more positive trends than do the housing-unit-based productivity values above (because the floor area has been trending up). The question is how much, and whether the prior significantly negative trends remain so.

The three series are in Figure 7. The significant downward trend seen in the old series remains (the series now spans 1974-2002, as no average square footage per unit for multi-family buildings was available before 1974). The linear yearly trend coefficient is -30.8 ft^2 per employee (s.e. = 9.4).

The newer series do see some changes relative to their units per employee patterns. For multi-family housing, which had a negative and significant trend in units per employee, the trend for square footage still has a negative point estimate (-42.6) but is now insignificant (s.e. = 37.9). Single-family housing now has a *positive* and significant annual trend growth of 37.4 ft^2 per employee (s.e. = 16.0).

The magnitudes are instructive, however. Average square footage built per employee in single-family housing construction over 1990-2020 is 7120 ft^2 . An annual trend growth of 37.4 ft^2 per employee off that base is 0.5 percent per year. So even in the most optimistic case, this is one quarter of the 2.0 percent annual labor productivity growth in the overall economy during the period.

V. Deteriorating Materials Productivity

Setting aside measurement problems, in this section we document a new type of productivity slowdown in the construction sector involving trends in the way which the sector converts intermediate inputs into outputs.

One summary measure of this process is the ratio of the sector's value added to gross output. The difference between these two is expenditures on intermediate inputs. Their ratio captures how much of the sector's final revenues are created through the application of value-added factors (labor and capital) as opposed to purchased from outside the sector.

Value added as a share of gross output in construction rose over the 1950-2020 period, from a level just below 0.4 in the early years to just above 0.5. The period of highest growth is circa 1985 to 2000, after which the series levels off. The sector has therefore shifted a greater portion of its final output production to activities inside the sector, done by its own labor and capital. As noted in Sveikauskas et al. (2018), specialized subcontractor labor is treated as a purchased service (a type of intermediate input) in standard KLEMS accounting. A shift toward less intensive use of subcontractors could account for some of the sector's shift toward a greater share of value added in gross output.

We break down this pattern further by looking at the separate trajectories of the revenue shares of each of the three major inputs into gross output (i.e., revenues): labor, capital, and intermediates.¹⁰ Intermediate inputs' share fell throughout 1950-2020, from an average of 60.5 percent during the first five-year period (1950-54) to 48.9 percent over the last five years. About two-thirds of this 11.6 percentage point drop was recovered by labor inputs taking a larger share (growing from 33.6 to 41.9 percent) and another third by an increase in capital's share (5.9 to 9.3 percent).

We can also look at intermediate input efficiency by computing the relative growth of real gross output and real intermediates use. This growth in “materials productivity”—output obtained per unit of intermediate input—is an analog to the more commonly measured labor productivity.

Materials productivity for both all industries and the construction sector show an interesting pattern in Figure 8. Materials productivity in construction looks much like labor productivity does, rising until the late 1960s, at which time the series turns downward and begins a half-century of decline. All industries materials productivity, on the other hand, has a slow positive trend throughout the entire 1950-2020 period.

¹⁰ Results reported here are similar if we use inputs' shares of sector costs instead of revenues, which is not surprising given the low and slow-changing markups we estimated above.

This turnaround is curious, especially when paired with the result that the sector's value added as a share of its gross output is growing. Given that the difference between these two values is intermediates purchases, this implies that over time the sector has been spending less on inputs to produce a given amount of gross output production. These productivity figures indicate this reduction in input spending is not because the sector has become more adept at converting intermediates into output (which would reduce the quantity of intermediates required). Instead, it seems to embody a substitution effect: as the industry became worse at converting intermediates into output, thereby raising the implied (output) unit costs of intermediates, the industry has shifted toward a more value-added-intensive production process.

VI. Productivity and Reallocation in Disaggregated Data

In this section we drill down below the sectoral aggregate data and investigate patterns in the disaggregated numbers. Specifically, we test whether, when there are producers with heterogeneous productivity levels, the market reallocates activity away from lower-productivity producers and toward higher-productivity producers. This would be consistent with efficiently operating input and output markets in the sector. The productivity literature has documented this reallocation mechanism as one way in which aggregate productivity growth can rise (indeed, it would facilitate aggregate productivity to grow even if no single producer's productivity rises).¹¹

We employ the state-level output and employment data from the Census of Construction conducted every five years from 1972 to 2017. Our analysis assumes that construction inputs' marginal products are positively correlated with their average products. This would be the case, for example, if the production function were Cobb-Douglas, though this property holds for many other production functions as well.

¹¹ This is sometimes referred to as the "between" mechanism (De Loecker and Syverson, 2021).

The testable empirical prediction of efficient reallocation is that states or firms with higher construction productivity today should, all else equal, see construction activity grow faster in the future, as additional resources are moved to production in that setting.¹²

Using the state-level output and employment data from the Census of Construction, we first compute the labor productivity level (real net construction value per employee) of the construction sector in each state and census year. We then regress the future change in the state's share of all US net construction value on the labor productivity level in the baseline period. With allocative efficiency, we should see a positive coefficient: states with construction sectors that are initially higher productivity see an increase in their share of total construction value.

We run this regression using data computed at two different timings. The first uses a state-by-census-year panel that regresses the five-year change in share of US construction value from year t to year $t + 5$ on the state's construction labor productivity in year t . We include year and state fixed effects in this specification. We have data for 9 inter-census changes from 1972 to 2017. The second specification is similar, but uses only a single long-difference of a state's share spanning 1972 to 2017 and regresses it on the state's 1972 construction labor productivity level.

The results are in Table 1. Column 1 shows the results from the panel; column (2) reports the outcome of the long-difference specification. Neither indicates the hypothesized positive relationship between initial productivity and future growth. Both coefficients are negative, significantly so in the panel regression.

Thus at best there is no clear relationship between the productivity of a state's construction sector and the amount of construction activity that state should expect to gain, be it the near or distant future. If anything, resources seem to move away from the more productive states.

¹² This is akin to the "dynamic allocation" test done for U.S. hospitals in Chandra et al. (2016).

This result may indicate the presence of market frictions that limit the ability of the construction sector's input and output markets to reallocate activity toward higher-productivity uses. The failure of this market mechanism blocks one of the two major channels through which productivity in an industry can grow, and may be a partial explanation of the aggregate productivity problems facing the industry.

VII. Conclusion

Measured productivity performance in the construction sector has been unusually awful for 50 years. This has caught economic researchers' attention before but has gained increasing attention in broader policy-related circles in recent years. In this paper, we have updated some of this earlier work to show that measurement problems alone likely cannot explain all of the decline and that there are some problems facing productivity in the industry that have not been documented previously.

First, from a purely accounting perspective, input-based explanations simply cannot account for the half-century decline in productivity. Most of the observed productivity decline results from the divergence between the construction sector's output price deflator from average price growth in the economy. This is summarized by the fact that if nominal construction activity is deflated by the GDP deflator instead of the sector-specific deflator, the path of implied productivity in the sector is much closer to that observed in economy-wide productivity. This importance of the deflator to the measured decline is not itself evidence that the construction sector output deflator is wrong, however.

Further exploring the deflator's growth, we find that increases in input unit prices do not seem to be driving its divergence. Either construction productivity truly is declining, raising unit input requirements, or there is some form of mismeasurement that is uncorrelated with higher input prices (i.e., ruling out quality-based explanations where the sector shifts to higher-quality and more expensive inputs).

Several pieces of evidence suggest that measurement errors in the deflator are not the sole cause of poor productivity performance in the sector. When we measure productivity performance in physical units in a key industry in the sector—residential housing construction—it also shows declining or stagnant productivity.

The sector's ability to transform intermediates into finished products has deteriorated.

We also document evidence that something keeps producers in areas where the sector is more productive from growing. Rather than construction inputs flowing to areas where they are more productive, the activity share of these areas either stagnates or even falls. This problem with allocative efficiency may be accentuating the aggregate productivity problem for the industry.

The productivity struggle is not just a figment of the data. It is real. Further research is needed to test between competing explanations and sharpen the picture of what has been happening in the sector. Certainly construction is an important enough component of total economic activity to warrant attention.

References

- Allen, Steven G. 1985. "Why Construction Industry Productivity Is Declining." *Review of Economics and Statistics*, 67(4): 661-9.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *American Economic Journal: Macroeconomics*, 13(1): 333-72.
- Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson. 2016. "Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector." *American Economic Review*, 106(8): 2110-44.
- Changali, Sriram, Azam Mohammad, and Mark van Nieuwland. 2015. "The Construction Productivity Imperative." McKinsey Productivity Sciences Center.
- De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics*, 135(2): 561-644.
- De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. 2016. "Prices, Markups, and Trade Reform." *Econometrica*, 84(2): 445-510.
- De Loecker, Jan, and Chad Syverson. 2021. "An Industrial Organization Perspective on Productivity." In *Handbook of Industrial Organization, Vol. 4*. Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri (eds.) Amsterdam: North Holland.
- de Roux, Nicolás, Marcela Eslava, Santiago Franco & Eric Verhoogen. 2021. "Estimating Production Functions in Differentiated-Product Industries with Quantity Information and External Instruments." NBER Working Paper 28323.
- Domar, Evsey D. 1961. "On the Measurement of Technological Change." *Economic Journal*, 71(284): 709-29.
- Economist. 2017. "Efficiency eludes the construction industry." *Economist*, Aug 19, 2017.

- Gordon, Robert J. 1968. "A New View of Real Investment in Structures, 1919-1966." *Review of Economics and Statistics*, 50(4): 417-28.
- Hulten, Charles R. 1992. "Growth Accounting When Technical Change is Embodied in Capital." *American Economic Review*, 82(4): 964-80.
- Pieper, Paul E. 1991. "The Measurement of Construction Prices: Retrospect and Prospect." in *Fifty Years of Economic Measurement: The Jubilee of the Conference on Research in Income and Wealth*, Ernst R. Berndt and Jack E. Triplett, eds. University of Chicago Press.
- Orr, Scott. Forthcoming. "Within-Firm Productivity Dispersion: Estimates and Implications." *Journal of Political Economy*.
- Potter, Brian. 2021. "Measuring Construction Efficiency." Oct 22, 2021. <https://constructionphysics.substack.com/p/measuring-construction-efficiency>.
- Schriver, William R. and Roger L. Bowlby. 1985. "Changes in Productivity and Composition of Output in Building Construction, 1972-1982." *Review of Economics and Statistics*, 67(2): 318-22.
- Smith, Noah. 2021. "What happened to construction productivity?" Blog post, Sep 29, 2021. <https://noahpinion.substack.com/p/what-happened-to-construction-productivity>.
- Stokes, H. Kemble Jr. 1981. "An Examination of the Productivity Decline in the Construction Industry." *Review of Economics and Statistics*, 63(4): 495-502.
- Sveikauskas, Leo, Samuel Rowe, James Mildenberger, Jennifer Price, and Arthur Young. 2016. "Productivity Growth in Construction." *Journal of Construction Engineering and Management*, 142(10): 04016045-1-8.
- Sveikauskas, Leo, Samuel Rowe, James Mildenberger, Jennifer Price, and Arthur Young. 2018. "Measuring Productivity Growth in Construction." *Monthly Labor Review*, January 2018: 1-33.

Syverson, Chad. 2017. "Challenges to Mismeasurement Explanations for the US Productivity Slowdown." *Journal of Economic Perspectives*, 31(2): 165-86.

Figure 1

Indexes of Value Added Per Worker and TFP, Overall U.S. and Construction Sector (BEA Data)




Figure 2

Capital Stocks (2012 = 100), Total Economy and Construction Sector




Figure 3

Construction Sector Labor Inputs, Three Measures

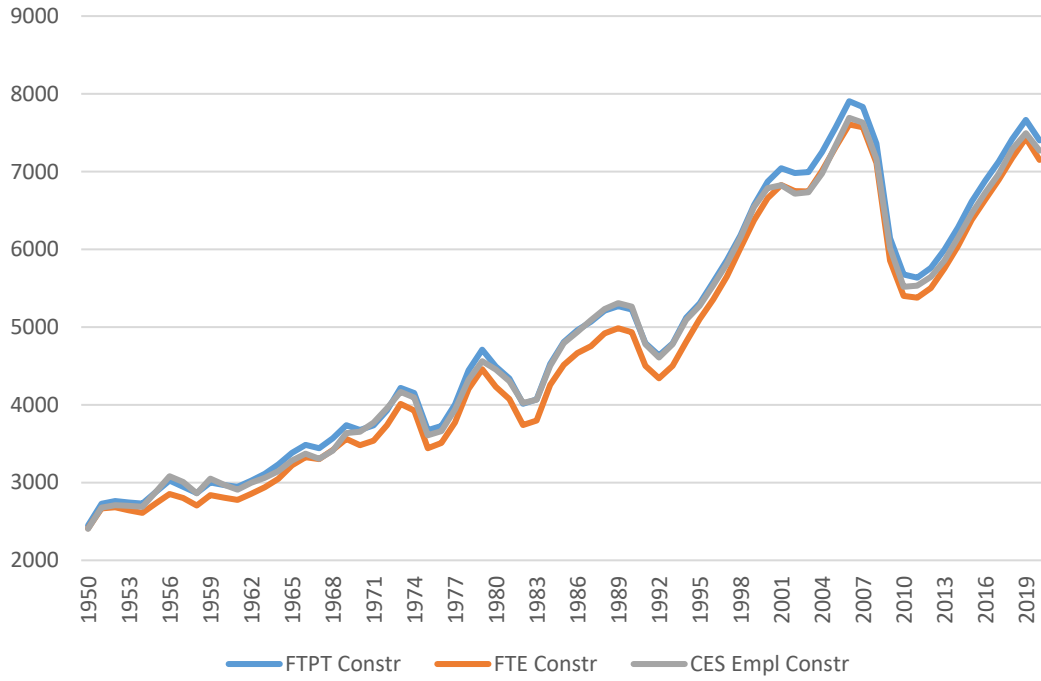


Figure 4

Indexes of Value Added Per Worker and TFP, Overall U.S. and Construction Sector (BEA Data)

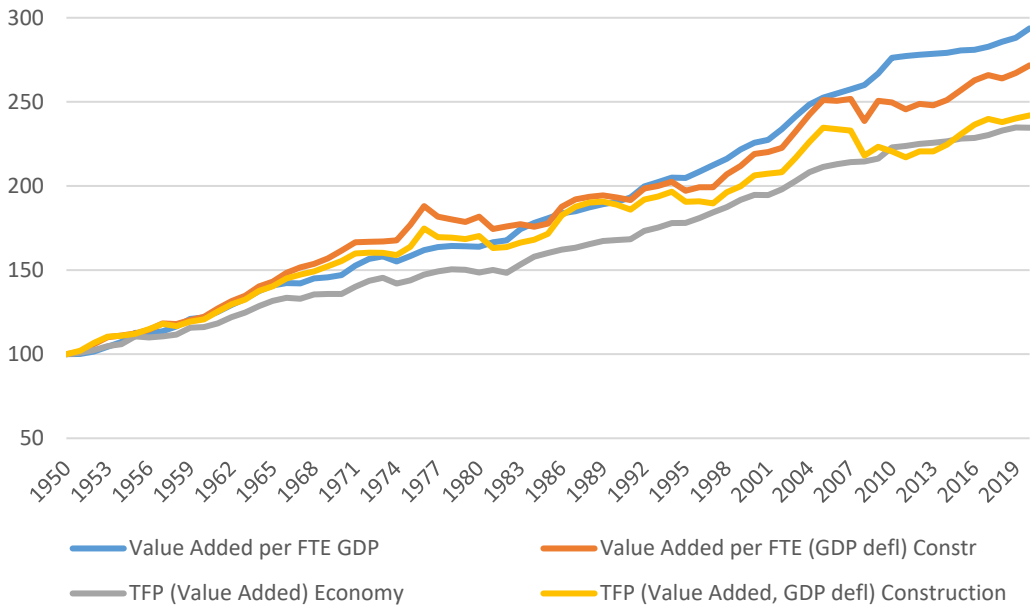


Figure 5

Construction and GDP Output Price Index Compared to Construction Intermediates Price Index




Figure 6

Housing Units per Employee, Housing Industry

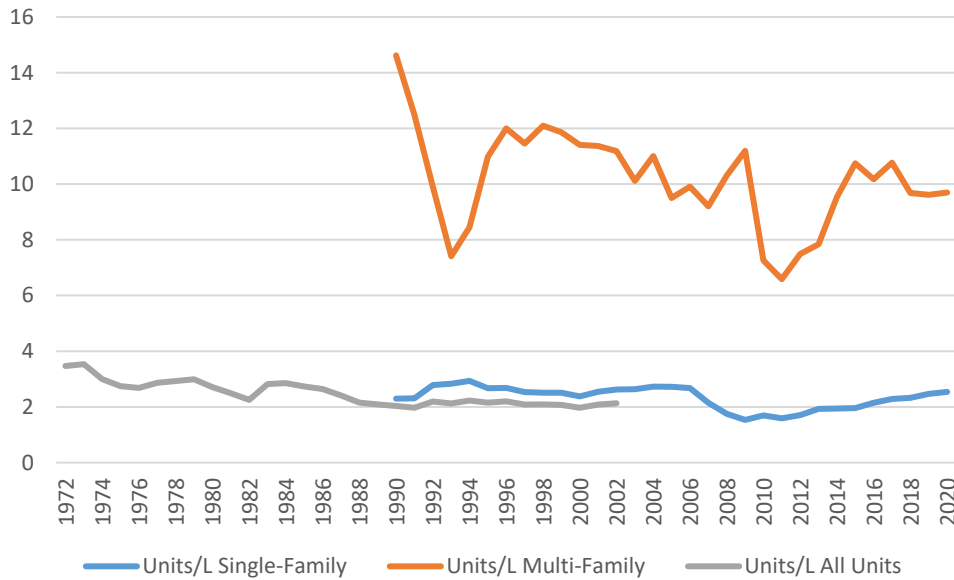


Figure 7

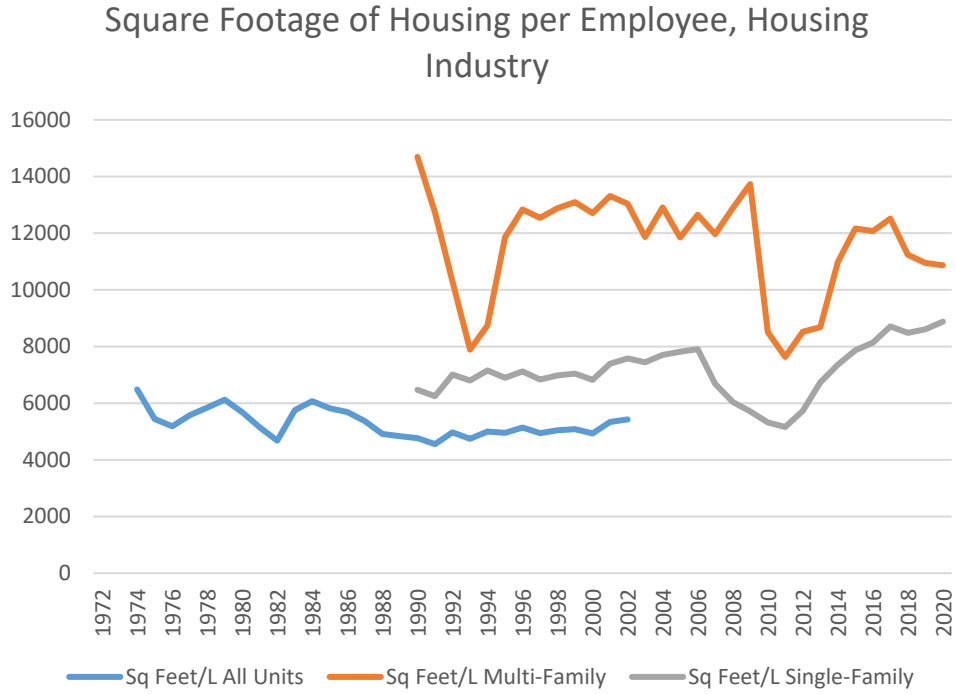


Figure 8

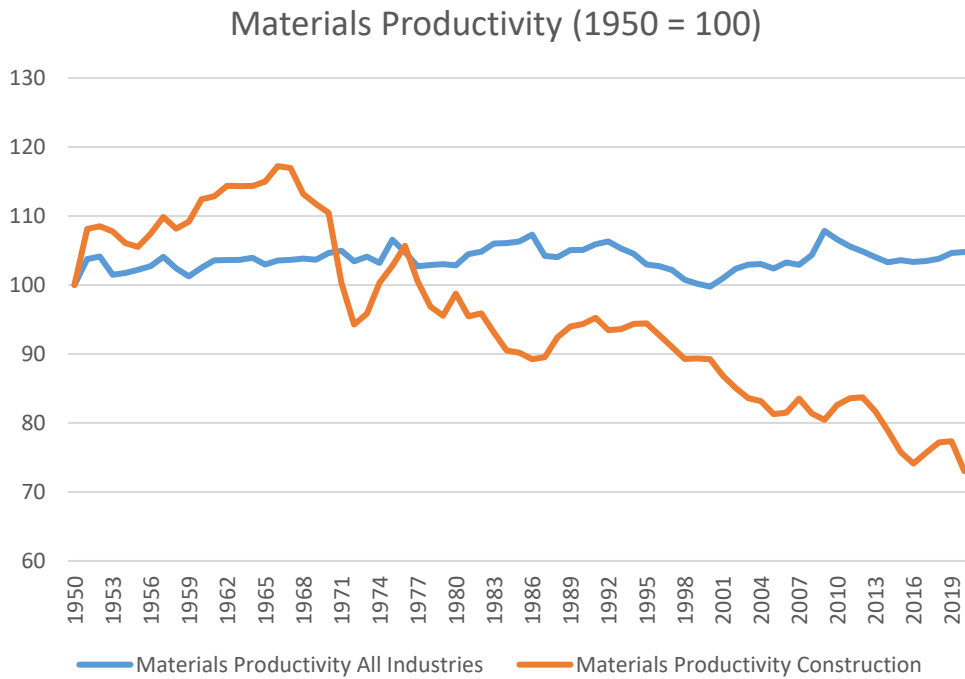


Table 1. Change in State's Share of U.S. Construction Value on Initial Labor Productivity Level

	(1) Change in state's share of U.S. net construction, year t to $t+5$ (percent)	(2) Change in state's share of U.S. net construction, 1972-2017 (percent)
$\ln(\text{state's real net construction value per employee})_t$	-2.92 (0.48)	-2.40 (3.01)
Year FE	Yes	No
State FE	Yes	No
N	459	51