FEDERAL RESERVE BANK OF KANSAS CITY

ECONOMIC REVIEW

Fourth Quarter 2018 *Volume 103, Number 4*

FEDERAL RESERVE BANK OF KANSAS CITY

ECONOMIC REVIEW

Fourth Quarter 2018 *Volume 103, Number 4*

The Faster Growth of Larger, Less Crowded Locations

By Jordan Rappaport

Over the past few decades, the population and employment growth of small and large locations in the United States have diverged. Many smaller cities and rural areas saw declining population and employment from 2000 to 2017 as residents and jobs migrated to larger, more prosperous locations. This migration might suggest that the benefits of size, such as business productivity and urban amenities, have become greater over time. However, the migration might also reflect other factors, such as the disproportionate specialization of smaller locations in the declining manufacturing and agriculture sectors.

Jordan Rappaport documents the faster population and employment growth of medium and large metropolitan areas compared with smaller locations and finds growth is strongly positively correlated with population. Moreover, he finds that most of this correlation is driven by size itself rather than other characteristics. However, he also finds that this relationship breaks down for the largest metro areas. His results suggest that both the benefits and costs of size have increased over the past few decades.

Machine Learning Approaches to Macroeconomic Forecasting

By Aaron Smalter Hall

Forecasting macroeconomic conditions can be challenging, requiring forecasters to make many discretionary choices about the data and methods they use. Although forecasters underpin the choices they make about models and complexity with economic intuition and judgement, these assumptions can be flawed. Machine learning approaches, on the other hand, automate as many of those choices as possible in a manner that is not subject to the discretion of the forecaster.

Aaron Smalter Hall applies machine learning techniques to find an optimal forecasting model for the unemployment rate. His results suggest that when supplied with diverse and complex data, a machine learning model can outperform simpler time-series models as well as a consensus of professional forecasters, with better performance at shorter horizons. In particular, his results show that a machine learning model can identify turning points in the unemployment rate earlier than competing methods.

The Response of U.S. Investment to Oil Price Shocks: Does the Shale Boom Matter?

By Nida Çakır Melek

After an unprecedented decline from 2014 to 2016, the real price of oil more than doubled, renewing interest in the effects of oil price fluctuations on the U.S. economy. The oil sector has become increasingly important to the U.S. economy over the past decade, and total U.S. business fixed investment appears to have followed oil investment's pattern in recent years. This positive correlation between oil prices and U.S. investment growth may be related to the surge in U.S. oil production known as the shale boom.

Nida Çakır Melek explores the effect of unexpected oil price changes (or "shocks") on U.S. investment and examines whether this effect changed after the shale boom. She finds that U.S. investment has become more responsive to demand shocks and less responsive to oil supply shocks since the shale boom. In addition, she finds that oil investment has become more responsive to oil supply and demand shocks since the boom. Her results suggest that the shale boom led to greater spillovers from the oil sector to the aggregate economy.

The Faster Growth of Larger, Less Crowded Locations

By Jordan Rappaport

Over the past few decades, the population and employment of small and large locations in the United States have been diverging. Most of the smallest locations in the United States—the approximate 1,200 counties and micropolitan areas with a population below 25,000—saw declining population and employment from 2000 to 2017 as their residents and jobs migrated to larger, more prosperous locations. Conversely, almost all medium and large metropolitan areas in the United States—those with a population of 500,000 or more—saw increasing population and employment from 2000 to 2017, many at well above the national rate.

An important question is whether this divergence between small and large locations has been driven by size itself. One possibility is that the benefits of size have become greater over time. For example, businesses may increasingly benefit from being near suppliers. Likewise, households may increasingly value access to services and amenities that are only available in larger locations. Alternatively, the divergence may be driven by characteristics that are correlated with size but not inherent to it. For example, the slower growth of smaller locations may simply reflect their disproportionate specialization in the manufacturing and agriculture sectors, which have seen declining employment.

In this article, I document the faster population and employment growth of medium and large metropolitan areas compared with smaller locations. Among these smaller locations—rural counties, micropolitan

---

*Jordan Rappaport is a senior economist at the Federal Reserve Bank of Kansas City. Anissa Khan, an intern at the bank, helped prepare the article. This article is on the bank's website at **www.KansasCityFed.org***

areas, and small metropolitan areas—growth is strongly positively correlated with population. Statistical analysis shows that most of this positive correlation is likely driven by size itself rather than location characteristics correlated with size. Among the medium and large metropolitan areas, growth is only weakly correlated with population but strongly negatively correlated with population density, a measure of crowdedness that moves closely with population. This negative correlation with density, too, is likely driven by density itself rather than correlated characteristics. Together, growth's positive correlation with population and negative correlation with density suggest that both the benefits and costs of size have increased over the past few decades.

Section I documents the relationship between population growth and size: population growth is positively correlated with size up to a population of about 500,000, uncorrelated with increases in size from 500,000 to 3 million, and negatively correlated with increases in size above 3 million. Section II lays out a framework for interpreting these correlations between growth and size: differences in locations' population and employment growth typically reflect relative changes in locations' productivity and amenities. Section III documents that the positive correlation of growth and size holds even after controlling for differences in local characteristics.

I. The Positive Relationship between Population and Employment Growth and Size

To analyze the relationship between growth and size, I look at all locations within the continental United States. Specifically, I combine the 358 metropolitan and 554 micropolitan areas delineated after the 2000 decennial census with the 1,346 remaining counties that are not part of a metropolitan or micropolitan area. The metropolitan areas, which range in population from 52,000 to 18 million, are combinations of counties surrounding a dense core of at least 50,000 residents. Most are made up of two or more counties. For descriptive purposes, I divide the metropolitan areas into three groups: small (population up to 500,000), medium (population from 500,000 to 3 million), and large (population above 3 million). The micropolitan areas, which range in population from 13,000 to 182,000, are combinations of counties surrounding a dense core of 10,000 to 50,000 residents. Most are made up of a single county. The remaining counties (henceforth, "non-core") range in population from 67 to 97,000. I measure growth rates using a constant delineation of metropolitan and micropolitan land areas. Thus, any changes in metropolitan and micropolitan area populations attributable to changes in their land area are excluded from measured growth rates.[1]

Chart 1 shows a scatter plot of the average annual growth rate of locations' population from 2000 to 2017 against the natural log of their population in 2000. I take the natural log of population to allow the horizontal axis to measure proportional rather than additive changes in population: each log point increase moving rightward along the horizontal axis represents a multiplicative increase in population by a factor of 2.7.[2]

The chart shows a clear, positive relationship between locations' population growth from 2000 to 2017 and their level of population in 2000. The blue line shows a smoothed average of this relationship, which can be interpreted as the predicted rate of growth based on initial size.[3] Predicted population growth ranges from modestly negative for locations with a low population in 2000 to moderately positive for most metropolitan areas. For example, a location with a population of 1,500 in 2000 (log population of 7.3) has a predicted population growth rate of –0.6 percent per year, leading to a 10 percent cumulative loss in population from 2000 to 2017. In contrast, a location with a population of 500,000 (log population of 13.1) has a predicted growth rate of just under 1 percent per year, leading to a 17 percent cumulative gain in population from 2000 to 2017. The predicted population growth rate is highest, 1.1 percent per year, for locations with a population close to 3 million (log population of 14.9).

Locations' predicted population growth falls off as their population in 2000 exceeds 3 million. Los Angeles and New York City, the two largest metropolitan areas with respective populations of 12.4 million and 18.3 million and log populations of 16.3 and 16.7, have predicted growth rates of 0.5 percent and 0.2 percent per year, meaningfully lower than the maximum 1.1 percent rate. But the decline from the maximum rate is slight for most large metropolitan areas. For example, Philadelphia, the fourth-largest metropolitan area in 2000 (with a population of 5.7 million and log population of 15.6) has a predicted

Chart 1
Population Growth versus Initial Population, 2000–17

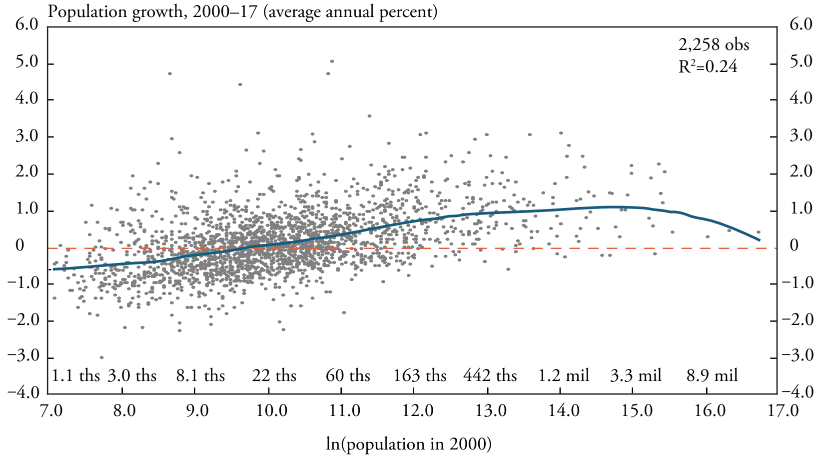

Notes: The blue line represents a prediction of locations' growth rates based on their population. The dashed orange line corresponds to a growth rate of zero. Replication code is available in an online data supplement.
Sources: U.S. Census Bureau and author's calculations.

growth rate of 1.0 percent per year, only a tick below the maximum predicted rate.

Overall, locations with population above 500,000 on average have considerably higher predicted growth than medium-sized locations, which in turn have considerably higher predicted growth than small locations. Of course, the actual growth rates of many locations differed considerably from the predicted rate based on their size. The large vertical dispersion of the scatter above and below the average line reflects that characteristics other than initial size drove most of the variation in growth rates from 2000 to 2017.

Correspondingly, small size did not preclude rapid growth. Chart 2 zooms in on the left-hand side of Chart 1, showing the same scatter and average relationship for locations with a population in 2000 below 100,000 (log population of 11.5). Many of the locations that grew fastest relative to their predicted rates—that is, those furthest above the predicted growth line—are distinguished by natural amenities such as mountains (for example, Custer County, CO; Rich County, UT; Wasatch County, UT; St. George, UT; Valley County, ID; Jackson, WY; Bozeman, MT; and Lyon County, NV) or warm winter weather (for example, Palm Coast, FL; The Villages, FL; and Walton County, FL).[4] Others are adjacent to metropolitan areas (for example, Hudspeth

Chart 2

Growth versus Initial Population, Smaller Locations

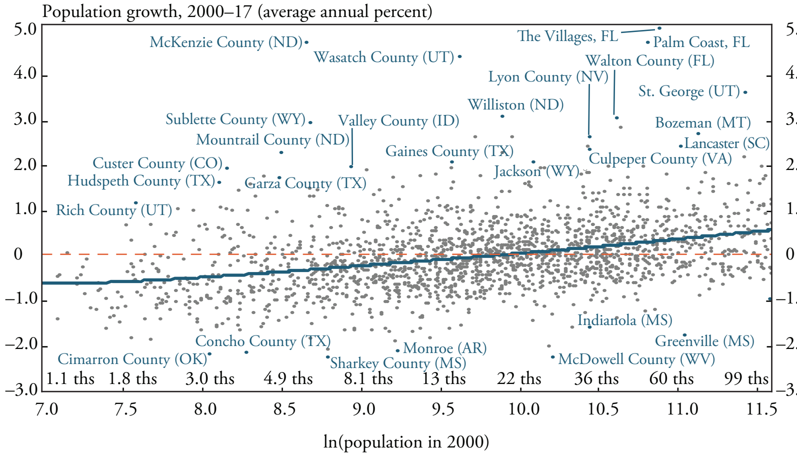

Notes: Metropolitan and micropolitan areas are labeled with the name of their largest city. The R^2 value is for all locations with a population below 100,000. The blue line represents a prediction of locations' growth rates based on their population and estimated using all 2,258 locations. The dashed orange line corresponds to a growth rate of 0. Marker for Issaquana County, MS (log population 7.7, growth rate –3.0 percent) is below the displayed area. Replication code is available in an online data supplement.
Sources: U.S. Census Bureau and author's calculations.

County, TX; Culpeper County, VA; and Lancaster, SC) or benefitted from oil and natural gas deposits that became accessible with the development of hydraulic fracturing technology (Williston, ND; Mountrail County, ND; McKenzie County, ND; Sublette County, WY; Garza County, TX; and Gaines County, TX).

Similarly, large size did not preclude population decline. Chart 3 zooms in on the right-hand side of Chart 1, showing the same scatter plot and average relationship for medium and large metropolitan areas, those with a population in 2000 of at least 500,000. In nine of these metros, population actually declined. Among these, eight are distinguished by an industrial composition skewed heavily toward manufacturing, a sector in which employment has been contracting for many decades. The disadvantages of this inherited industrial composition are likely to have offset any benefits from size.

Overall, however, declining population was relatively rare for medium and large metropolitan areas as well as for smaller metropolitan areas with a population in 2000 between 200,000 and 500,000 (Chart 4). In contrast, the majority of locations with a population

Chart 3

Growth versus Initial Population, Medium and Large Metropolitan Areas

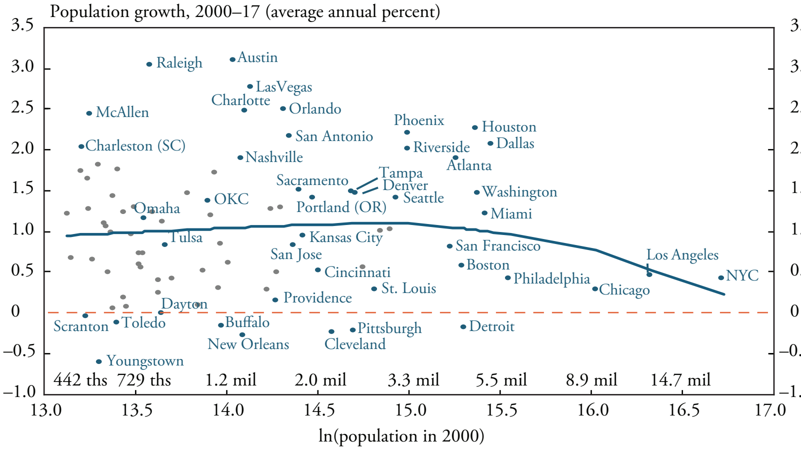

Notes: Metropolitan areas are labeled with the name of their largest city. The R^2 value is for metropolitan areas with a population of at least 500,000. The blue line represents a prediction of locations' growth rates based on their population and estimated using all 2,258 locations. The orange dashed line corresponds to a growth rate of 0. The Denver and Boulder metropolitan areas are combined. Replication code is available in an online data supplement.
Sources: U.S. Census Bureau and author's calculations.

below 25,000 contracted, as did more than 40 percent of locations with a population between 25,000 and 50,000. Across all locations, population accounts for almost a quarter of the variation in growth rates from 2000 to 2017 (as estimated by the R^2 statistic). This is a high share attributable to a single characteristic.

The relationship between population growth and size from 2000 to 2017 continued a pattern that began in the mid-twentieth century. Chart 5 shows the predicted population growth rates from 1960 to 1980 (blue line) and from 1980 to 2000 (green line) based on the corresponding initial population levels. Both predicted relationships are characterized by a positive correlation between growth and size across most locations and a negative relationship across the largest locations.[5]

However, the relationship between growth and size evolved over these periods in four important ways. First, predicted growth shifted lower over time for locations with an initial population up to about 1 million, primarily reflecting slowing national population growth. As a result, the share of small locations with predicted population decline increased over time. Second, the downward slope in the relationship for

Chart 4

Share of Locations Losing Population

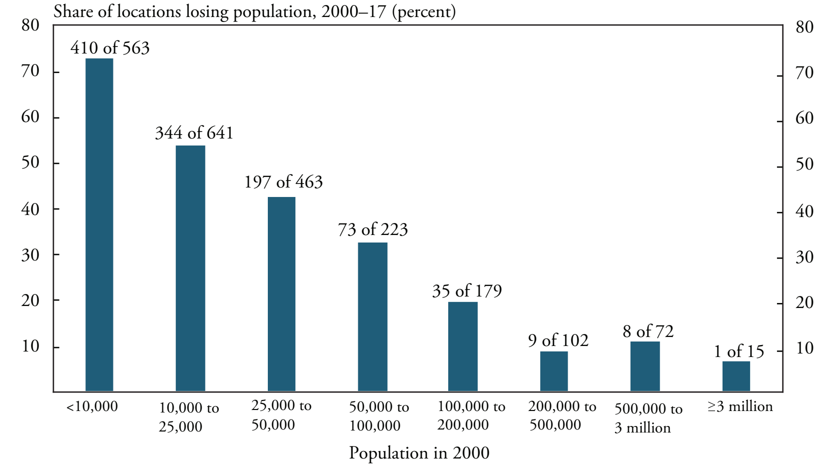

Note: Replication code is available in an online data supplement.
Sources: U.S. Census Bureau and author's calculations.

Chart 5

Historical Population Growth versus Initial Population

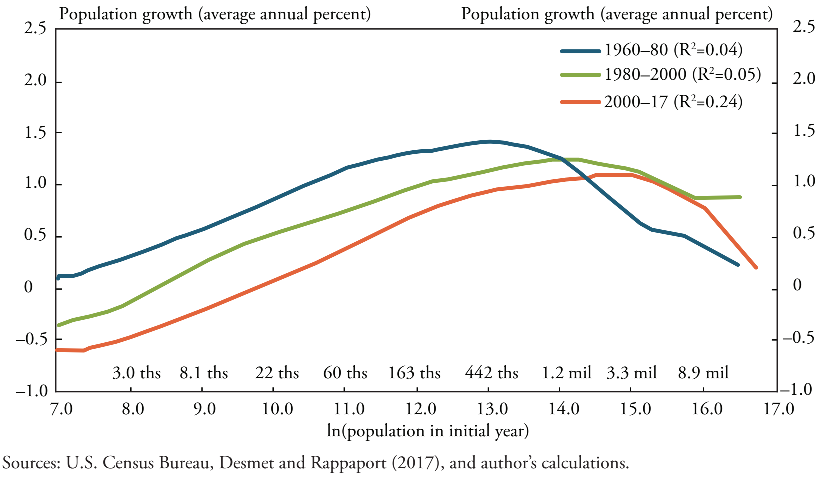

Sources: U.S. Census Bureau, Desmet and Rappaport (2017), and author's calculations.

the largest locations began at successively higher population levels over time. From 1960 to 1980, predicted growth began sloping down at a population of about 500,000. From 1980 to 2000, predicted growth began sloping down at a population of about 1.3 million. In the most recent period, it began sloping down at about 3.0 million. Third, the magnitude of the downturn for the largest locations varied over time. Predicted growth declined more in the 2000–17 period than in the 1980–2000 period but less than in the 1960–80 period. Fourth, the relationship between growth and size was strongest during the most recent period. Initial size accounted for 24 percent of the variation in growth from 2000 to 2017 compared with a maximum of 5 percent in the earlier periods. This increase in explanatory power may reflect size becoming a more important determinant of growth or other determinants of growth, such as suburbanization and the migration to the Sunbelt, becoming less important.

Growth's relationship with employment is similar to its relationship with population. Chart 6 shows a scatter of locations' employment growth from 2000 to 2014 plotted against their initial level of employment in 2000. The black line represents predicted growth.[6] The employment levels along the horizontal axis are lower than the population levels discussed previously, reflecting that the number of individuals with jobs, including both full time and part time, is less than one-half of the population in most locations. Employment growth from 2000 to 2014 was positively correlated with employment in 2000 up to a level of about 500,000 employed individuals (log employment of 13.1) and approximately uncorrelated with further increases in employment. Predicted average annual employment growth rose from about –0.8 percent for locations with fewer than 500 employed individuals in 2000 (log employment below 6.2) to about 1.0 percent for locations with more than 500,000 employed individuals.

The similar relationships between size and growth of both population and employment reflect that employment and population growth are strongly positively correlated over the long term. In particular, increases in employment tend to be matched approximately one for one by inflows of workers (Rappaport 2012). I focus my subsequent analysis on population rather than employment, as it is the better measured of the two.

Chart 6

Employment Growth versus Initial Employment, 2000–14

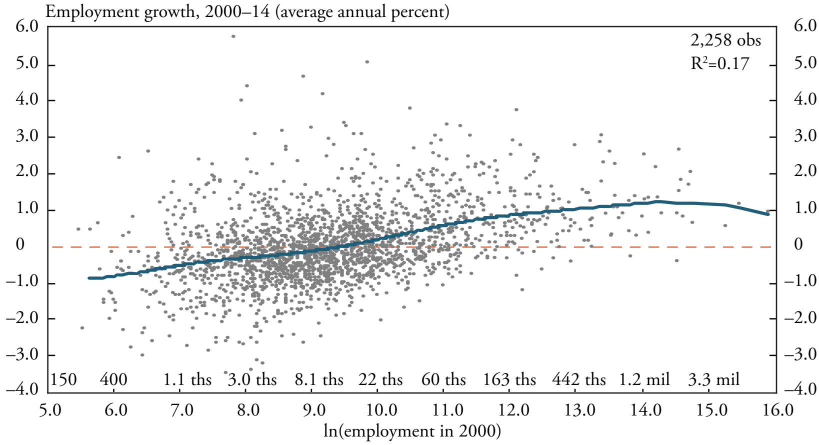

Notes: The orange dashed line corresponds to a growth rate of 0. Several locations have employment growth rates outside the displayed range. Replication code is available in an online data supplement.
Sources: U.S. Census Bureau and author's calculations.

II. Interpreting Correlations with Growth

Local population and employment growth are not inherently desirable or undesirable outcomes. Rather, they reflect changes in the underlying fundamentals determining where households and businesses choose to locate. In this section, I introduce a local growth framework that illustrates how these fundamentals are linked to population and employment. I then suggest some shifts in fundamentals that might be driving the observed relationship between growth and size.

A framework for understanding local population and employment growth

The local growth framework has three key features. First, locations have different fundamental characteristics that affect the productivity of businesses or that serve as amenities for residents.[7] Some of these characteristics are exogenous in the sense that they do not depend on local outcomes such as income and population—for example, natural resources, a natural ocean harbor, natural recreational opportunities, and nice weather. Other characteristics are endogenous in the sense that they are themselves a local outcome—such as population, employment, and income—or partly depend on a local outcome. For example, a larger population may contribute to disamenities such as traffic congestion and pollution. Likewise, the age and income distribution of local residents may affect the variety of available goods and services (Glaeser, Kolko, and Saiz 2001; Diamond 2016).[8]

Second, an economy is in a spatial equilibrium if all households and firms prefer to remain where they are rather than move elsewhere (Rosen 1979; Roback 1982). In this equilibrium, businesses cannot increase their profits by moving somewhere else because higher productivity in other locations is offset by higher wages, land prices, and other costs. Similarly, households cannot benefit from moving somewhere with higher wages or more amenities because these advantages are offset by higher housing prices, more traffic congestion, and other costs. As is intuitive, locations with characteristics that contribute to high productivity or amenities have a larger equilibrium population (Rappaport 2008a, 2008b, 2016). The larger population pushes up land prices, housing prices, traffic congestion, and other costs to the level at which businesses and households are equally willing to live in locations with lower productivity and amenities but also lower costs.

Third, a location's transition from its current level of population to its equilibrium level, driven by net flows of households and businesses, takes considerable time (Rappaport 2004; Desmet and Rappaport 2017). Locations with a current population significantly below its equilibrium level will typically grow at an above-average rate for several decades as people gradually migrate there. Locations with an initial population significantly above its equilibrium level will typically grow at a below-average rate for several decades as people gradually migrate elsewhere. These transitions are gradual for several reasons, including physical moving costs, households' ties to family and friends in origin locations, the time it takes for housing and infrastructure to deteriorate in origin locations, and the time it takes to build new housing and infrastructure in destination locations (Glaeser and Gyuorko 2005; Kennan and Walker 2011; Davis and others 2013).

Although the local growth framework emphasizes location size as a consequence of productivity and amenities, size is also an important determinant of location productivity and amenities. For example, large size contributes negatively to productivity and amenities by increasing numerous types of congestion.[9] But large size also contributes positively to productivity and amenities in numerous ways. Such agglomerative benefits typically take the form of more sharing, better matching, and increased learning. (Duranton and Puga 2004). Agglomerative sharing captures large locations' ability to spread large fixed costs, such as building an airport or sports stadium, across a broad base of customers. It also captures large locations' ability to support a wider variety of business and consumer services, especially those that are more specialized. Agglomerative matching is exemplified by a larger pool of job candidates and firms that allows for a better fit of workers' skills to firms' needs. For example, research shows that a larger pool of employers has become more important over time as the share of couples with dual careers has risen (Costa and Kahn 2000). Agglomerative learning concerns the generation and diffusion of knowledge. For example, researchers tend to discover more, as measured by patents, when working near each other (Carlino and Kerr 2015; Buzard and others 2017). Likewise, when many workers in the same occupation are concentrated in one location, they learn from each other. As Alfred Marshall observed in 1890, "the mysteries of the trade become no mysteries; but are as it were in the air." Estimates based on wages suggest that these agglomeration effects together increase a location's productivity by between 2 and 6 percent for each log point increase in population. (Combes and Gobillon 2015). No comparable estimates exist for the effect of size on amenities.

An important implication of the local growth framework is that differences in growth rates typically reflect *changes* in underlying productivity and *changes* in underlying amenities rather than levels. At any point in time, the distribution of population across locations already captures many of the differences in local productivity and amenities, which tend to persist over very long periods. Metaphorically, differences in locations' current size result from how firms and people have "voted with their feet" up until that point in time (Tiebout 1956). Differences in growth rates, on the other hand, reflect firms and people changing their "votes." Because transitions are extended, this vote changing can persist for up to several decades following a change in productivity or amenities.[10]

Interpreting the relationship between growth and size

The local growth framework suggests at least three possible interpretations of the empirical relationships between the levels of population and employment and their growth rates. One possible interpretation is that the agglomerative benefits from increases in population up to a level of 500,000 have become larger during recent decades.[11]

Following such an increase in agglomerative benefits, the equilibrium size of locations with a previous equilibrium population above 500,000 would have risen *relative* to the new equilibrium size of locations with a previous equilibrium population below 500,000.[12]

Under this first interpretation, the lack of correlation between growth and size across medium metropolitan areas suggests that the agglomerative benefits from increases in population from 500,000 to 3 million have remained approximately the same during recent decades. The negative correlation between growth and size across large metropolitan areas suggests that the agglomerative benefits from increases in population above 3 million have become smaller during recent decades.[13]

A second possible interpretation is that the agglomerative costs from increases in population up to a level of 500,000 have become smaller during recent decades. The spatial equilibrium of locations' size depends on the extent to which higher productivity and amenities are offset by higher housing prices, traffic congestion, and other agglomerative costs as more people compete for housing, road space, and other goods and services.[14] Thus, if agglomerative costs become less sensitive to population—for example, if cities relax zoning restrictions, expand public transit, or improve highway infrastructure—then larger locations will benefit proportionally more than smaller locations.[15] In essence, the lowered costs of size allow more households and businesses to crowd into locations with high productivity and amenities before rising house prices and congestion offset the gains from doing so.

Under this second interpretation, the flat and declining portions of the relationship between growth and population suggest that the agglomerative costs from increases in population from 500,000 to 3 million have remained approximately the same, while the agglomerative costs from increases in population above 3 million have become larger.[16]

A third possible interpretation is that the contributions of a location's exogenous characteristics to productivity and amenities have changed during recent decades. Such changes would likely induce correlations between growth and size because the same characteristics, through their previous contributions to productivity and amenities, helped determine the location's previous equilibrium population. For example, households during the nineteenth century were likely to have preferred, all else equal, to live where winters were less cold and summers were less hot and humid. As a result, locations with mild weather were likely to have, all else equal, a larger equilibrium population than locations that did not. During the early twentieth century, this preference for mild weather began strengthening (Rappaport 2007). The resulting faster population growth of locations with mild weather induced a positive correlation between growth and size, reflecting that locations with mild weather tended to have above-average size. However, the underlying impetus for this positive correlation was not intrinsically related to size. Similarly, the industry and occupation composition of many small locations is skewed toward agriculture and manufacturing, sectors for which employment has been declining during recent decades. The slower growth of smaller locations may partly reflect this industrial shift rather than size.

III. The Increased Benefits and Costs of Size

The local growth framework suggests three possible interpretations of the empirical correlations between growth and size. To assess which interpretation is most likely, I first run regressions of population growth on initial population and several additional location characteristics to rule out that such characteristics, rather than changes in agglomerative benefits and costs, are driving the correlations. I then run regressions of population growth on initial population density to help distinguish whether changes in agglomerative benefits or agglomerative costs are driving the correlations.

The increased benefits of size for smaller locations

Table 1 reports results from regressing average annual population growth from 2000 to 2017 on initial population in 2000 and additional characteristics such as geographic location and industry composition. I divide initial population into a “spline” of eight population ranges to allow the regression to approximate the smoothed relationship between predicted growth and population (the blue line in Chart 1). The regression coefficient on each of the population ranges estimates the slope of a linear segment corresponding to the curved line through that population range.

The results in column 1 show that regressing population growth on the spline without controlling for other characteristics approximately

Table 1

Partial Correlation of Population Growth with Population

	(1)	(2)	(3)	(4)
Partial correlations of 2000–17 population growth	Own size only	Baseline	Baseline and industry composition	Baseline, industry composition, and occupation composition
Controls:				
Metropolitan adjacency (2)		x	x	x
Weather (10)		x	x	x
Coast and river adjacency (7)		x	x	x
Hilliness (2)		x	x	x
Shale basin (6)		x	x	x
Higher education (1)		x	x	x
Industry composition (18)			x	x
Occupation composition (21)				x
ln(pop) from 5.9 to 8 (3,000)	0.06	0.16	0.36**	0.41***
128 locations	(0.12)	(0.10)	(0.15)	(0.16)
ln(pop) from 8 to 9 (8,100)	0.27**	0.35***	0.26***	0.28***
319 locations	(0.12)	(0.07)	(0.07)	(0.07)
ln(pop) from 9 to 10 (22,000)	0.30***	0.27***	0.18***	0.18***
627 locations	(0.08)	(0.07)	(0.07)	(0.06)
ln(pop) from 10 to 11 (60,000)	0.25***	0.19***	0.19***	0.20***
620 locations	(0.09)	(0.06)	(0.06)	(0.05)
ln(pop) from 11 to 12 (163,000)	0.46***	0.50***	0.39***	0.36***
292 locations	(0.08)	(0.09)	(0.09)	(0.09)
ln(pop) from 12 to 13 (440,000)	0.18	0.09	0.04	0.03
126 locations	(0.13)	(0.10)	(0.09)	(0.09)
ln(pop) from 13 to 14 (1.2 million)	0.12	0.17	0.00	–0.06
60 locations	(0.20)	(0.15)	(0.11)	(0.10)
ln(pop) from 14 to 16.7 (18.3 million)	–0.09	–0.13	–0.16	–0.16
41 locations	(0.13)	(0.12)	(0.10)	(0.10)
Observations	2,258	2,258	2,258	2.258
R^2	0.24	0.43	0.52	0.54
Adjusted R^2	0.23	0.42	0.50	0.52
Control variables		28	46	67
R^2, control variables		0.30	0.48	0.51

** Significant at the 5 percent level
*** Significant at the 1 percent level

Notes: Dependent variable is average annual population growth (percent) from 2000 to 2017. Regressions also include a constant. The smallest location has a log population of 5.9 (population of 356). Standard errors are in parentheses and adjust for spatial correlation based on Conley (1999). Italicized text reports the number of locations with a population that lies within each spline segment. Coefficients on all variables included in the baseline regression are reported in appendix Table A-1. Replication code is available in an online data supplement.

matches the smoothed relationship between predicted growth and initial size. Each log point increase in population from 8 to 12 (corresponding to a population increase from about 3,000 to 163,000) is associated with between 0.25 and 0.46 percentage point faster predicted growth per year, implying large differences in cumulative growth. Three of the corresponding coefficients statistically differ from zero at the 1 percent level, and the fourth coefficient statistically differs from zero at the 5 percent level. Predicted growth also rises modestly as log population increases from 12 to 13 and from 13 to 14 and then falls modestly as log population rises above 14. But none of the coefficients on these segments statistically differs from zero. Overall, the initial population spline accounts for 24 percent of the variation in location growth rates, the same as the share accounted for by the smoothed relationship shown in Chart 1.

The statistically significant, positive relationship between growth and size for locations with log population from 8 to 12 continues to hold after accounting for numerous other characteristics, suggesting that the correlation is indeed driven by changes in agglomerative benefits and costs rather than by changes in the contributions of exogenous characteristics to productivity and amenities. Column 2 of Table 1 shows results from a regression that controls for 28 baseline characteristics likely to affect productivity and amenities and thereby drive both growth and size. On their own, these baseline characteristics—which describe adjacency to metropolitan areas, adjacency to coasts and rivers, weather, hilliness, energy deposits, and the presence of universities and colleges—account for a considerable portion of the variation in growth rates (30 percent, reported in the bottom row) and an even larger portion of the variation in the level of population in 2000 (40 percent, not shown).[17] But they leave the coefficients on the population spline mostly unchanged, ruling out that any of the baseline characteristics is driving the correlation between growth and size.

The positive correlation between growth and size similarly continues to hold after controlling for the industry and occupation composition of locations. Column 3 reports results from a regression that includes the baseline characteristics along with variables measuring the share of aggregate employment in each of 18 industries. Column 4 reports results from a regression that includes the baseline and industry characteristics along with variables measuring the share of aggregate employment in each of 21 occupations. In both regressions, the positive coefficients on the spline segments with log population between 8 and 12 are mostly unchanged. Both regressions also estimate a large, statistically significant positive coefficient on the lowest spline segment (log population from 5.9 to 8; population from 350 to 3,000), suggesting that even across the smallest locations, the net productivity and amenity benefits of size increased. In addition, including the industry and occupation controls boosts the magnitude of the negative coefficient on the uppermost spline segment and lowers its standard error. As a result, the coefficients on the uppermost segment in columns 3 and 4 statistically differ from 0 at only slightly above the 10 percent level, suggesting that the net productivity and amenity benefits of size may have decreased for large metropolitan areas.[18]

The estimated coefficients in Table 1 imply that differences in population among smaller locations predict large differences in growth rates from 2000 to 2017. To obtain the differences in growth rates between a location with log population in 2000 of 8 (a population of 3,000) and a location with log population of 12 (a population of 163,000), I sum the coefficients of the four spline segments from 8 to 12 for each of the specifications. This simple calculation shows that a location with log population in 2000 of 12 has from 1.0 to 1.3 percentage points faster predicted annual growth from 2000 to 2017, corresponding to a larger cumulative increase in population from 19 to 25 percentage points.

Although characteristics excluded from these regressions could account for the positive correlation between growth and size, such a possibility seems unlikely. On their own, the 67 characteristics included in the column 4 regression account for more than half of the variation in population growth, a high share given the many idiosyncratic circumstances affecting local growth. Moreover, many excluded characteristics are likely to endogenously depend on size. Including such endogenous variables in a regression may help identify channels through which size affects growth but might also mask the effect of size through all channels.

The positive, statistically significant coefficients reported in Table 1 suggest that the *net* benefits of larger size—the gross agglomerative benefits of higher productivity and amenities less the gross agglomerative costs of higher housing prices and more traffic congestion—increased for non-core counties, micropolitan areas, and the smallest metropolitan areas. But the coefficients do not distinguish whether this net change arose from an increase in agglomerative benefits or a decrease in agglomerative costs. The former seems more likely, as these groups of locations have historically been characterized by relatively low home prices, minimal commuting traffic, and few other congestion costs.

Conversely, the negative coefficients on the uppermost spline segment suggest that the net benefits of larger size may have decreased for metropolitan areas with population above 1.2 million. Again, the coefficients do not distinguish whether this arose from a decrease in agglomerative productivity and amenities or an increase in agglomerative costs, both of which seem plausible.

The increased costs of size for larger locations

To distinguish whether changes in agglomerative benefits or costs are driving the relationship between growth and size, I look at the relationship between population growth and population density. While strongly positively correlated with the level of population, population density appears to be more closely related to home prices, a key agglomerative cost, than does population. In particular, population density accounts for more than twice the variation in median home prices across medium and large metropolitan areas.[19] In addition, population density is unlikely to affect businesses productivity. (Employment density, in contrast, is likely to affect businesses productivity by allowing more workers to interact with each other.)

Population density varies greatly within metropolitan areas, making "raw" density (total population divided by total land area) a poor summary measure of the density of the neighborhoods in which most residents live. For example, the raw density of the Las Vegas metropolitan area in 2000 was 174 persons per square mile. However, this measure is misleading, as 85 percent of Las Vegas residents lived in census tracts—small geographic units that typically include between 1,000 and 8,000 residents—with raw population of more than 2,100 persons per square mile.

To better reflect the density most residents actually experience, I measure mean population density, calculated as the population-weighted mean of each census tract's raw density (Glaeser and Kahn 2004; Rappaport 2008a).[20] Using this measure, the mean density of Las Vegas in 2000 was 6,500 persons per square mile. Across all locations, mean density ranged from less than 1 person per square mile for the 35 locations with the lowest value to more than 8,000 persons per square mile for the five locations with the highest value. Mean population density in the New York City metropolitan area, 32,600 persons per square mile, was almost three times that of the second most dense metropolitan area, Los Angeles.

Table 2 reports results from regressing population growth on a spline of mean population density in 2000. The results in column 1 show that before controlling for other characteristics, positive coefficients on three of the spline segments statistically differ from zero, implying that increases in mean density within each of the corresponding ranges predict faster population growth. However, controlling for the baseline characteristics and industry composition pushes down each of these positive coefficients to near zero (column 3).[21] In other words, any positive association between predicted growth and mean density may be driven by differences in the baseline characteristics and industry composition rather than by a change in agglomerative costs.

In contrast, population growth is negatively correlated with increases in mean population density within the uppermost segment, a relationship that strengthens as additional controls are added to the regression. Controlling for the baseline characteristics and industry composition, the negative coefficient on the uppermost segment statistically differs from 0 at the 1 percent level (column 3). Additionally, controlling for occupation shares leaves this coefficient essentially unchanged (column 4). Holding the baseline characteristics, industry composition, and occupation composition constant, each log point increase in mean population density from 8 to 10.4 (that is, each 2.7 multiplicative increase in mean population density from 3,000 to 33,000 persons per square mile) is associated with 0.45 percentage point slower predicted population growth per year. This implies that the New York City metropolitan area would have had 1.1 percentage point per year higher predicted population growth if it had had St. Louis' mean population density in 2000 of 3,000 persons per square mile (rather than 32,600). Correspondingly, New York City's predicted increase in population from 2000 to 2017 would have been 20 percentage points higher.

Table 2

Partial Correlation of Population Growth with Mean Population Density

	(1)	(2)	(3)	(4)
Partial correlations of 2000–17 population growth	Own size only	Baseline	Baseline and industry composition	Baseline, industry composition, and occupation composition
Controls:				
Metropolitan adjacency (2)		x	x	x
Weather (10)		x	x	x
Coast and river adjacency (7)		x	x	x
Hilliness (2)		x	x	x
Shale basin (6)		x	x	x
Higher education (1)		x	x	x
Industry composition (18)			x	x
Occupation composition (21)				x
ln(dens) up to 3 (20 persons/square mile) *398 locations*	0.07 (0.05)	0.14*** (0.04)	0.05 (0.05)	0.08** (0.04)
ln(dens) from 3 to 4 (55) *277 locations*	0.29*** (0.10)	0.12 (0.09)	0.07 (0.08)	0.02 (0.07)
ln(dens) from 4 to 5 (148) *398 locations*	−0.05 (0.10)	0.03 (0.08)	−0.03 (0.07)	−0.02 (0.06)
ln(dens) from 5 to 6 (403) *383 locations*	0.21** (0.10)	0.12 (0.08)	0.06 (0.08)	0.06 (0.07)
ln(dens) from 6 to 7 (1,100) *427 locations*	0.14 (0.11)	0.15 (0.10)	0.08 (0.08)	0.06 (0.08)
ln(dens) from 7 to 8 (3,000) *299 locations*	0.60*** (0.15)	0.39*** (0.13)	0.06 (0.13)	−0.02 (0.12)
ln(dens) from 8 to 10.4 (33,000) *56 locations*	−0.11 (0.14)	−0.20 (0.15)	−0.44*** (0.15)	−0.45*** (0.16)
Observations	2,258	2,258	2,258	2,258
R^2	0.16	0.37	0.49	0.51
Adjusted R^2	0.15	0.36	0.48	0.50
Control variables		28	46	67
R^2, control variables		0.30	0.48	0.51

** Significant at the 5 percent level
*** Significant at the 1 percent level

Notes: Dependent variable is average annual population growth (percent) from 2000 to 2017. Regressions also include a constant. Standard errors are in parentheses and adjust for spatial correlation based on Conley (1999). Italicized text reports the number of locations with mean population density that lies within each spline segment. Results for all variables included in the baseline regression are reported in appendix Table A-2. Replication code is available in an online data supplement.

Among medium and large metropolitan areas, the negative correlation between growth and density is especially strong when measured using each metro's 95th percentile density. At least 95 percent of a metro's residents live in a tract with raw density at or below its 95th percentile density and at least 5 percent live in a tract with raw density at or above its 95th percentile density. The negative relationship between growth and 95th percentile density stands out in a scatter plot of the former against the latter (Chart 7). The best-fit linear relationship between the two, shown by the black line, accounts for 12 percent of the variation in growth (as measured by an R^2 statistic; see column 1 of Table 3).[22]

Although uncorrelated with population on its own, growth across medium and large metropolitan areas is strongly positively correlated with population after controlling for 95th percentile density, illustrating the partly offsetting benefits and costs of size (columns 2 and 3 of Table 3). More specifically, the positive coefficient on population in column 3 likely captures a gross increase in agglomerative benefits, while the negative coefficient on density—which is considerably larger in magnitude than when not controlling for population (column 3 versus column 1)—likely captures a gross increase in agglomerative costs. On net, the gross increase in costs dominates the gross increase in benefits, which is reflected in the negative coefficient on density when not controlling for size (column 1).

Population density, while unlikely to benefit businesses' productivity, may contribute positively to metros' amenities. For example, high population density helps support nearby urban amenities such as pedestrian access to varied restaurants, cafes, bars, retailers, and performance venues. Consistent with this possibility, growth is positively correlated with spikes in population density, measured by the increase in log density from a metro's 95th percentile tract to its 99th percentile tract (column 4 of Table 3). This positive correlation may be closely related to the increased tendency of young professionals to live near metropolitan central business districts (Couture and Handbury 2017; Baum-Snow and Hartley 2018)

Changes in agglomerative benefits and costs, as captured by partial correlations with population and population density, account for a considerable share of the variation in population growth across medium and large metropolitan areas. Together, population and 95th percentile

Chart 7

Growth versus 95th Percentile Population Density, Medium and Large Metropolitan Areas

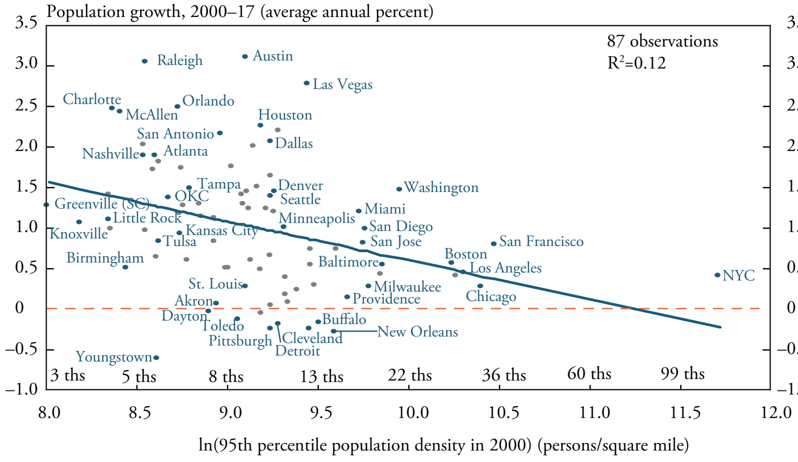

Notes: Metropolitan areas are labeled with the name of their largest city. The blue line represents the best fit based on a linear regression. The orange dashed line corresponds to a growth rate of 0. The Denver and Boulder metropolitan areas are combined. Replication code is available in an online data supplement.
Sources: U.S. Census Bureau and author's calculations.

Table 3

Partial Correlation of Population Growth with Population and Population Density

Partial correlations of 2000–17 population growth	(1)	(2)	(3)	(4)	(5)
ln(pop)		0.00	0.38***		0.34***
		(0.09)	(0.11)		(0.10)
ln(95th percentile density)	–0.48***		–0.82***		–0.79***
	(0.10)		(0.14)		(0.13)
ln(99th percentile)–ln(95th percentile)				1.13**	1.07***
				(0.47)	(0.36)
Observations	87	87	87	87	87
R^2	0.12	0.00	0.20	0.09	0.28

** Significant at the 5 percent level
*** Significant at the 1 percent level

Notes: Dependent variable is average annual population growth (percent) from 2000 to 2017. Regressions also include a constant. Standard errors are in parentheses and adjust for spatial correlation based on Conley (1999). Replication code is available in an online data supplement.

population density account for 20 percent of the variation in growth (column 3 of Table 3). Including the difference in log density between the 99th and 95th percentile boosts the share of variation accounted for to 28 percent (column 5 of Table 3). Including the 25th and 75th percentile densities further boosts the share of variation accounted for to 38 percent (not shown).[23] This ability to account for more than one-third of the variation in population growth suggests that the shifting contributions of size, as measured by both population and density, to productivity and amenities have been among the most important determinants of recent metropolitan population growth.

IV. Conclusions

The population and employment of small and large locations in the United States have been diverging for several decades. For locations with a population in 2000 up to about 500,000, population growth from 2000 to 2017 was positively correlated with initial population. For locations within this group with a population up to about 160,000, size itself is likely to have driven the positive correlation, reflecting a net increase in agglomerative productivity and amenities over the past few decades.

In contrast, population growth from 2000 to 2017 was negatively correlated with mean population density at high levels, likely reflecting a net increase in agglomerative costs such as housing prices and traffic congestion over the past few decades. Similarly, growth across medium and large metropolitan areas was strongly negatively correlated with population density measured at the 95th percentile.

This pattern of local population growth—positively correlated with population across smaller locations and negatively correlated with population density across larger locations—is likely to persist for a considerable time, as net flows of households and jobs gradually move locations toward a spatial equilibrium.

Population and employment growth's dependence on size and density has some important public policy implications. First, small locations seeking to reverse declining population and employment face a formidable challenge, as they must offset the decreasing relative productivity and amenities attributable to their small size. Small locations that have succeeded in doing so have primarily relied on exogenous characteristics, such as natural resources, nice weather, natural recreational opportunities, the presence of a university, or adjacency to a large metropolitan area. For small locations that lack such offsetting characteristics, public policy may be more effective ameliorating the negative consequences of decline than reversing it.

Second, economic development strategies that attract new jobs may benefit existing local residents and businesses if the associated agglomerative benefits exceed the associated agglomerative costs. For example, the increase in local employment may sufficiently increase the productivity of existing businesses to allow them to pay higher wages. It may also attract households and workers who sufficiently increase the local tax base to offset any associated increases in public spending. In other cases, however, successfully attracting jobs may hurt existing residents and businesses. In particular, tax incentives narrowly targeted at one or a handful of businesses may lower government services and increase the tax burden for existing residents and businesses. Rather than narrowly targeting incentives, more effective public policy might focus on policies that broadly benefit local businesses and residents, both existing and new.

Appendix

Additional Tables

Table A-1

Partial Correlation of Population Growth with Population and Baseline Controls (Regression Reported in Table 1, Column 2)

Right-hand-side variable	Coefficient	Standard error	t-statistic	p-value
Initial population spline				
ln(pop) from 5.9 to 8 (3,000)	0.16	0.10	1.51	0.130
ln(pop) from 8 to 9 (8,100)	0.35***	0.07	5.28	0.000
ln(pop) from 9 to 10 (22,000)	0.27***	0.07	3.90	0.000
ln(pop) from 10 to 11 (60,000)	0.19***	0.06	2.90	0.004
ln(pop) from 11 to 12 (163,000)	0.50***	0.09	5.58	0.000
ln(pop) from 12 to 13 (440,000)	0.09	0.10	0.95	0.343
ln(pop) from 13 to 14 (1.2 million)	0.17	0.15	1.12	0.264
ln(pop) from 14 to 16.7 (18.3 million)	–0.13	0.12	–1.11	0.267
Metropolitan adjacency (1/0 indicator)				
Micro/non-core adjacent to metro with pop. > 1 million	0.33***	0.07	4.59	0.000
Micro/non-core adjacent to metro with pop. ≤ 1 million	0.05	0.05	1.10	0.270
Weather				
Average max daily temp. in Jan. (linear)	1.92E-02***	5.89E-03	3.25	0.001
Average max daily temp. in Jan. (quadratic)	7.16E-04***	2.16E-04	3.32	0.001
Average max daily heat index in July (linear)	–3.81E-04	6.53E-03	–0.06	0.954
Average max daily heat index in July (quadratic)	–2.65E-04	3.63E-04	–0.73	0.466
Average mean daily relative humidity in July (linear)	–2.63E-02***	7.83E-03	–3.36	0.001
Average mean daily relative humidity in July (quadratic)	6.26E-05	3.08E-04	0.20	0.839
Average annual rainfall (linear)	–2.66E-03	8.19E-03	–0.32	0.745
Average annual rainfall (quadratic)	3.15E-04***	1.16E-04	2.71	0.007
Average annual number of days with rain (linear)	6.60E-03*	3.73E-03	1.77	0.077
Average annual number of days with rain (quadratic)	–1.32E-04***	3.71E-05	–3.57	0.000

Table A-1 (continued)

Right-hand-side variable	Coefficient	Standard error	t-statistic	p-value
Coast and river adjacency (1/0 indicator)				
Atlantic, Northeast census region	0.19	0.15	1.29	0.196
Atlantic, South Atlantic census division	0.53***	0.15	3.58	0.000
Gulf of Mexico	0.12	0.17	0.71	0.478
Pacific	-0.11	0.26	–0.43	0.666
Great Lakes	–0.32***	0.11	–2.95	0.003
Within 40 km of river on which nav. in 1968	–0.19**	0.08	–2.44	0.015
Within 40 km of major river	0.02	0.05	0.31	0.755
Hilliness				
Ratio of std. dev. of altitude to land area (linear)	1.43***	0.37	3.90	0.000
Ratio of std. dev. of altitude to land area (quadratic)	–0.55***	0.13	–4.36	0.000
Shale oil basins (1/0 indicator)				
Anadarko	0.35***	0.13	2.76	0.006
Bakken	0.87***	0.20	4.32	0.000
Eagle	0.56***	0.19	3.00	0.003
Haynesville	–0.13	0.11	–1.13	0.259
Niobrara	0.24**	0.11	2.27	0.023
Permian	0.08	0.15	0.50	0.617
Higher education				
Ratio of post-secondary students to pop.	2.32***	0.37	6.30	0.000

* Significant at the 10 percent level
** Significant at the 5 percent level
*** Significant at the 1 percent level

Notes: Table reports estimation results for all variables included in the regression reported in column 2 of Table 1. Standard errors adjust for spatial correlation based on Conley (1999). The p-value is the probability that the absolute value of the t-statistic would exceed its regression value under the null hypothesis that population growth is uncorrelated with the corresponding right-hand-side variable. Linear weather coefficients estimate the partial derivative of growth with respect to each of the five weather measures for a location with the mean value of that measure. Replication code is available in an online data supplement.

Table A-2

Partial Correlation of Population Growth with Mean Population Density and Baseline Controls (Regression Reported in Table 2, Column 2)

Right-hand-side variable	Coefficient	Standard error	t-statistic	p-value
Initial population spline				
ln(dens) up to 3 (20 persons/sq. mile)	0.14***	0.04	4.01	0.000
ln(dens) from 3 to 4 (55)	0.12	0.09	1.44	0.149
ln(dens) from 4 to 5 (148)	0.03	0.08	0.33	0.740
ln(dens) from 5 to 6 (403)	0.12	0.08	1.42	0.155
ln(dens) from 6 to 7 (1,100)	0.15	0.10	1.60	0.110
ln(dens) from 7 to 8 (3,000)	0.39***	0.13	3.09	0.002
ln(dens) from 8 to 10.4 (33,000)	-0.20	0.15	-1.38	1.831
Adjacency to metro area (1/0 indicator)				
Micro/non-core adjacent to metro with pop. > 1 million	0.30***	0.07	4.13	0.000
Micro/non-core adjacent to metro with pop. ≤ 1 million	0.01	0.04	0.11	0.911
Weather				
Average max daily temp in Jan. (linear)	2.74e-02***	5.94e-03	4.62	0.000
Average max daily temp in Jan. (quadratic)	8.09e-04***	2.23e-04	3.63	0.000
Average max daily heat index in July (linear)	-7.89e-03	6.81e-03	-1.16	1.753
Average max daily heat index in July (quadratic)	-4.52e-04	3.95e-04	-1.15	1.748
Average mean daily relative humidity in July (linear)	-2.59e-02	8.45e-03	-3.06	1.998
Average mean daily relative humidity in July (quadratic)	2.80e-04	3.19e-04	0.88	0.380
Average annual rainfall (linear)	7.32e-04	8.12e-03	0.09	0.928
Average annual rainfall (quadratic)	2.03e-04*	1.21e-04	1.67	0.094
Average annual number of days with rain (linear)	7.92e-03**	3.90e-03	2.03	0.043
Average annual number of days with rain (quadratic)	-1.44e-04	4.13e-05	-3.49	2.000
Coast and river adjacency (1/0 indicator)				
Atlantic, Northeast census region	0.28*	0.15	1.90	0.058
Atlantic, South Atlantic census division	0.50***	0.19	2.63	0.009
Gulf of Mexico	0.17	0.20	0.85	0.394
Pacific	0.03	0.26	0.12	0.906
Great Lakes	-0.26	0.10	-2.46	1.986
Within 40 km of river on which nav. in 1968	-0.21	0.08	-2.56	1.989
Within 40 km of major river	0.05	0.05	1.08	0.281

Table A-2 (continued)

Right-hand-side variable	Coefficient	Standard error	t-statistic	p-value
Hilliness				
Ratio of std. dev. of altitude to land area (linear)	1.06***	0.35	2.98	0.003
Ratio of std. dev. of altitude to land area (quadratic)	–0.43	0.13	–3.41	1.999
Shale oil basins (1/0 indicator)				
Anadarko	0.35***	0.10	3.32	0.001
Bakken	0.85***	0.20	4.20	0.000
Eagle	0.70***	0.21	3.37	0.001
Haynesville	–0.06	0.13	–0.46	1.353
Niobrara	0.24**	0.10	2.46	0.014
Permian	–0.07	0.17	–0.41	1.317
Higher education				
Ratio of post-secondary students to population	2.70***	0.41	6.66	0.000

* Significant at the 10 percent level
** Significant at the 5 percent level
*** Significant at the 1 percent level

Notes: Table reports estimation results for all variables included in the regression reported in column 2 of Table 2. Standard errors adjust for spatial correlation based on Conley (1999). The p-value is the probability that the absolute value of the t-statistic would exceed its regression value under the null hypothesis that population growth is uncorrelated with the corresponding right-hand-side variable. Linear weather coefficients estimate the partial derivative of growth with respect to each of the five weather measures for a location with the mean value of that measure. Replication code is available in an online data supplement.

Endnotes

[1]Holding delineations constant causes the calculated population growth rates of many fast-growing metropolitan areas to be lower than those calculated using population numbers published by the U.S. Census Bureau. The difference reflects that the official borders of metropolitan areas, delineated by the Office of Management and Budget, were redrawn after the 2010 decennial census. These border changes led to the inclusion of additional counties in many fast-growing metropolitan areas, reflecting the spread of suburbs into previously undeveloped land.

[2]Using a standard additive scale, the horizontal distance between the 1 log point increase from 3.3 million to 8.9 million (a 2.7 multiplicative increase) would be 3,000 times larger than the 1 log point increase from 1,100 to 3,000 (also a 2.7 multiplicative increase).

[3]I measure the average relationship by a linear Epanechnikov kernel with a bandwidth of 1.5 log points. Replication code is available in an online data supplement.

[4]The Palm Coast metropolitan area was merged into the Delton-Daytona Beach metropolitan area following the 2010 decennial census.

[5]Throughout the nineteenth and early twentieth centuries, population growth was negatively correlated with size across small locations. Beginning in the late nineteenth century, population growth was strongly positively correlated with size across medium and large locations. This steep positive relationship flattened in about 1960 (Desmet and Rappaport 2017).

[6]I measure employment in 2014 by values reported in the 2016 American Community Survey five-year summary file, which is based on households' responses to surveys from 2012 through 2016. I measure employment in 2000 by the number of individuals reporting they were employed the week prior to filling out their census questionnaire as disseminated in the 2000 decennial census summary files. Alternatively, measuring employment based on administrative data collected from firms, the positive relationship between growth and size is considerably weaker. This difference in the relationship between growth and the level of employment may reflect that a larger share of workers in small locations are self-employed or hold other positions for which administrative data, which are based on firms' payment of unemployment insurance taxes, are not collected.

[7]Productivity measures the efficiency with which firms transform labor and other inputs into a final output good or service. A location characteristic can be interpreted as increasing productivity if it allows businesses to pay higher wages for labor and higher prices for other inputs without hurting their profits. Low rates of taxes that fall on businesses can thus be interpreted as positively contributing to a location's productivity, reflecting that businesses care about after-tax profits. But tax incentives to lure a single business to a location typically leave the after-tax productivity of most existing businesses unaffected. A location characteristic can

be interpreted as increasing amenities if it makes households willing to pay higher house prices and accept lower wages. Low rates of taxes that fall on individuals can be interpreted as positive amenities, as they make individuals willing to accept lower pre-tax wages and pay higher housing prices compared to living elsewhere. However, low taxes may result in lower amenities in the form of public services.

[8]Many characteristics have both exogenous and endogenous components. For example, many seaports are protected by a constructed breakwater or require periodic dredging. For these ports, location along an ocean coast is clearly exogenous, while the breakwater and dredging are likely to be endogenous, based on judgments about the economic potential of the location. The industrial composition of locations similarly combines exogenous and endogenous components. In part, industrial composition depends on economic considerations from the distant past, which may no longer be relevant today. But industrial composition can also evolve over time in response to changing location productivity, amenities, and other economic circumstances.

[9]Congestion exemplifies a nonpecuniary cost, meaning that it does not take the form of an explicit monetary price. In contrast, higher land and house prices are pecuniary costs, which do not directly affect productivity and amenities.

[10]Of course, workers' skills differ as do their tastes for different consumption amenities. This is a second reason, in addition to variation in house prices, that population in a spatial equilibrium is distributed across many locations rather than clustered in a handful of locations with the highest productivity and amenities. Over the past few decades, workers have increasingly sorted into different metropolitan areas based on their skill type (Moretti 2012).

[11]This increase in agglomerative benefits can equivalently be thought of as the disadvantages of small size worsening during recent decades. Consistent with this, a measure of business dynamism has been declining in small locations relative to large locations (Brown 2018).

[12]The increase in the equilibrium population of larger locations relative to the equilibrium population of smaller locations (for locations with an initial population below 500,000) is consistent with the level of equilibrium population increasing for both types of locations. In this case, all that is required is that the proportional increase in the equilibrium population of the large locations exceed the proportional increase of the population of the small locations.

[13]This interpretation corresponds with a rise in the elasticity of productivity and amenities with respect to size for increases in population up to 500,000; an unchanged elasticity with respect to size for increases in population between 500,000 and 3 million; and a decline in the elasticity for increases in population above 3 million.

[14]Numerical results from a model of metropolitan size suggest that higher housing prices and traffic congestion contribute about equally to agglomerative costs (Rappaport 2016). If commuting speeds were to remain at their free-flow

level, metropolitan areas with the highest productivity would be an order-of-magnitude larger than they actually are.

[15]Research finds that building more highways significantly increases the number of commuters, leaving travel times mostly unchanged (Duranton and Turner 2011). Building highways can thus increase a location's equilibrium population, allowing migration from elsewhere until traffic congestion returns to its previous level.

[16]Consistent with the second interpretation, recent research suggests that the elasticity of agglomerative costs with respect to population is increasing (Combes, Duranton, and Gobillon 2015; Rappaport 2016). In other words, proportional increases in agglomerative costs due to an increase in location size are higher for larger locations. Thus, if the population of all locations proportionally increases by the same amount— for example, due to national population growth—then agglomerative costs will rise more for larger locations.

[17]The baseline controls include two indicator variables for metropolitan adjacency: the first variable takes a value of 1 for micropolitan areas and non-core counties adjacent to a metropolitan area with a population below 1 million and 0 otherwise, while the second takes a value of 1 for micropolitan and non-core counties adjacent to a metropolitan area with a population above 1 million and 0 otherwise. The 10 weather variables are linear and quadratic measures of winter temperature, the summer heat index, summer humidity, annual rainfall, and annual rainy days. The five coast variables are indicators taking the value of 1 if a location borders a coast along the Great Lakes, the Pacific Ocean, the Gulf of Mexico, the North Atlantic (Maryland north to Maine), and the South Atlantic (Virginia south to Florida). The two river variables are indicators that take a value 1 for locations that touch a major river and for locations that touch a river on which there was commercial navigation in 1968. The two hilliness variables are the linear and quadratic ratio of the standard deviation of altitude within a location, measured across 1.25-arc-minute grid cells, to the location's total land area. The six shale basin variables are indicators taking the value of 1 for locations in each of the Anadarko, Bakken, Eagle, Haynesville, Niobrara, and Permian basins. The presence of colleges and universities is measured by the share of a location's population enrolled in post-secondary classes. The appendix reports the results of the column 2 regression for all of these control variables. The variables themselves are included in the online data supplement.

[18]The respective p-values on the uppermost spline segment for the regressions reported in columns 3 and 4 are 0.116 and 0.103.

[19]The variation in log population accounts for 71 percent of the variation in the log of mean population density across all locations and 47 percent of the variation across medium and large metropolitan areas (as measured by R^2 statistics). Across medium and large metropolitan areas, log mean population density and log population account for 36 percent and 17 percent, respectively, of the variation in log median home price. Across all locations, however, log mean population

density accounts for a smaller share of the variation in log median home prices than does log population (25 percent versus 37 percent).

[20]The raw density of a metropolitan area is arithmetically equal to the land-weighted mean of the raw population density of each tract.

[21]The appendix reports the results of the column 2 regression for all control variables.

[22]The negative correlation between growth and density across medium and large metropolitan areas is much weaker for more standard benchmarks of density: mean population density can account for only 4 percent of the variation in growth, and median density cannot account for any of it. Instead, the negative correlation becomes meaningful at 75th percentile density, which accounts for 4 percent of the variation in growth. One possible explanation is that density measured at high percentiles reflects opportunities for apartment construction at sites near metropolitan centers (Rappaport 2017). The negative correlation between growth and density is also considerably weaker at the highest percentiles: measured at the 99th percentile, density accounts for only 5 percent of the variation in growth; measured at the maximum density within each metro, density accounts for only 2 percent of the variation in growth.

[23]A regression on the 25th, 75th, 95th, and 99th percentile densities in 2000 accounts for 37 percent of the variation in population growth from 2000 to 2017. An arithmetically equivalent regression has right-hand-side variables for the 25th percentile density as well as the increase in density from the 25th to 75th percentile, the increase from the 75th to the 95th percentile, and the increase from the 95th to the 99th percentile. The corresponding coefficients, each of which statistically differs from 0 at the 1 percent level, are negative on the first three variables and positive on the last variable. The negative coefficients are increasing in magnitude, implying that a 1 log point increase in 25th percentile density is associated with a smaller decrease in predicted growth than a 1 log point larger increase in density from the 25th to 75th percentiles, which in turn is associated with a smaller decrease in predicted growth than a 1 log point larger increase in density from the 75th to the 95th percentiles. Additionally, including log population boosts the share of variation accounted for to 38 percent, but the associated coefficient on log population does not statistically differ from 0. The online data supplement includes metropolitan density measured at numerous percentile benchmarks ranging from the 1st to the 99th.

References

Baum-Snow, Nathaniel, and Daniel Hartley. 2018. "Accounting for Central Neighborhood Change, 1980–2010." Working paper, July.

Blomquist, Glenn C., Mark C. Berger, and John P. Hoehn. 1988. "New Estimates of Quality of Life in Urban Areas." *American Economic Review*, vol. 96, pp. 89–107.

Brown, Jason P. 2018. "The Widening Divide in Business Turnover between Large and Small Urban Areas." *Economic Review.* Federal Reserve Bank of Kansas City, Third Quarter, pp. 5–25.

Buzard, Kristy, Gerald A. Carlino, Robert M. Hunt, Jake K. Carr, and Tony E. Smith. 2017. "Localized Knowledge Spillovers: Evidence from the Spatial Clustering of R&D Labs and Patent Citations." Federal Reserve Bank of Philadelphia, Working Paper no. 17-32. Available at https://www.philadelphiafed.org/-/media/research-and-data/publications/working-papers/2017/wp17-32.pdf

Carlino, Gerald, and William R. Kerr. 2015. "Agglomeration and Innovation." In Gilles Duranton, J. Vernon Henderson, and William Strange, eds., *Handbook of Regional and Urban Economics,* vol. 5A, pp. 349–404. Amsterdam: Elsevier. Available at https://doi.org/10.1016/B978-0-444-59517-1.00006-4

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2018. "The Costs of Agglomeration: House and Land Prices in French Cities." Working Paper. Available at http://real.wharton.upenn.edu/~duranton/research.html

Combes, Pierre-Philippe, and Laurent Gobillon. 2015. "The Empirics of Agglomeration Economies." In Gilles Duranton, J. Vernon Henderson, and William Strange, eds., *Handbook of Regional and Urban Economics,* vol. 5A, pp. 247–348. Amsterdam: Elsevier. Available at https://doi.org/10.1016/B978-0-444-59517-1.00005-2

Conley, Timothy G. 1999. "GMM Estimation with Cross-Sectional Dependence." *Journal of Econometrics*, vol. 91, no. 1, pp. 1–45. Available at https://doi.org/10.1016/S0304-4076(98)00084-0

Costa, Dora L., and Matthew E. Kahn. 2000. "Power Couples: Changes in the Locational Choice of the College Educated, 1940-1990. *The Quarterly Journal of Economics*, vol. 115, no. 4, pp. 1287–1315. Available at https://doi.org/10.1162/003355300555079

Couture, Victor, and Jessie Handbury. 2017. "Urban Revival in America, 2000 to 2010." NBER Working Paper no. 24084, November. Available at https://doi.org/10.3386/w24084

Davis, Morris A., Jonas D.M. Fisher, and Marcelo Veracierto. 2013. "Gross Migration, Housing and Urban Population Dynamics." Federal Reserve Bank of Chicago, Working Paper no. 2013-19. Available at https://www.chicagofed.org/publications/working-papers/2013/wp-19

Desmet, Klaus, and Jordan Rappaport. 2017. "The Settlement of the United States, 1800–2000: The Long Transition towards Girbrat's Law." *Journal of Urban Economics*, vol. 98, March, pp. 50–68. Available at https://doi.org/10.1016/j.jue.2015.03.004

Diamond, Rebecca. 2016. "The Determinants and Welfare Implications of U.S. Workers' Diverging Location Choices by Skill: 1980-2000." *American Economic Review*, vol. 106, no. 3, pp. 479–524. Available at https://doi.org/10.1257/aer.20131706

Duranton, Gilles, and Diego Puga. 2004. "Micro-Foundations of Urban Agglomeration Economies." In J. Vernon Henderson and Jacques-François Thisse, eds., *Handbook of Regional and Urban Economics,* vol. 4, pp. 2063–2117. Amsterdam: Elsevier. Available at https://doi.org/10.1016/s1574-0080(04)80005-1

Duranton, Gilles, and Matthew A. Turner. 2011. "The Fundamental Law of Road Congestion: Evidence from American Cities." *American Economic Review*, vol. 101, no. 6, pp. 2616–2652. Available at https://doi.org/10.1257/aer.101.6.2616

Glaeser, Edward L., and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy*, vol. 113, no. 2, pp. 345–375. Available at https://doi.org/10.1086/427465

Glaeser, Edward L., and Matthew E. Kahn. 2004. "Sprawl and Urban Growth." In J. Vernon Henderson and Jacques-François Thisse, eds., *Handbook of Regional and Urban Economics*, vol. 4, pp. 2481–2527. Amsterdam: Elsevier. Available at https://doi.org/10.1016/S1574-0080(04)80013-0

Glaeser, Edward L., Jed Kolko, and Albert Saiz. 2001. "Consumer City." *Journal of Economic Geography*, vol. 1, no. 1, pp. 27–50. Available at https://doi.org/10.1093/jeg/1.1.27

Kennan, John, and James R. Walker. 2011. "The Effect of Predicted Income on Individual Migration Decisions." *Econometrica*, vol. 79, no. 1, pp. 211–251. Available at https://doi.org/10.3982/ECTA4657

Marshall, Alfred. 1890. *Principles of Economics.* London: Macmillan.

Moretti, Enrico. 2012. *The New Geography of Jobs.* Boston and New York: Houghton Mifflin Harcourt.

Rappaport, Jordan. 2017. "Crowdedness, Centralized Employment, and Multifamily Home Construction." *Economic Review.* Federal Reserve Bank of Kansas City, vol. 102, no. 1, pp. 41–83.

______. 2016. "Productivity, Congested Commuting, and Metro Size." Federal Reserve Bank of Kansas City, Research Working Paper 16-03. Available at https://doi.org/10.18651/RWP2016-03

______. 2012. "Why Does Unemployment Differ Persistently across Metro Areas?" *Economic Review.* Federal Reserve Bank of Kansas City, Second Quarter, pp. 1–31.

______. 2008a. "A Productivity Model of City Crowdedness." *Journal of Urban Economics*, vol. 63, no. 2, pp. 715–722. Available at https://doi.org/10.1016/j.jue.2007.04.008

______. 2008b. "Consumption Amenities and City Population Density." *Regional Science and Urban Economics,* vol. 38, no. 6, pp. 533–552. Available at https://doi.org/10.1016/j.regsciurbeco.2008.02.001

______. 2007. "Moving to Nice Weather." *Regional Science and Urban Economics*, vol. 37, no. 3, pp. 375–398. Available at https://doi.org/10.1016/j.regsciurbeco.2006.11.004

______. 2004. “Why Are Population Flows So Persistent?” *Journal of Urban Economics*, vol. 56, no. 3, pp. 554–580. Available at https://doi.org/10.1016/j.jue.2004.07.002

Rappaport, Jordan, and Jeffrey D. Sachs. 2003. “The United States as a Coastal Nation.” *Journal of Economic Growth*, vol. 8, no. 1, pp. 5–46. Available at https://doi.org/10.1023/A:1022870216673

Roback, Jennifer. 1982. “Wages, Rents, and the Quality of Life.” *Journal of Political Economy,* vol. 90, no. 6, pp. 1257–1278. Available at https://doi.org/10.1086/261120

Rosen, Sherwin. 1979. “Wage-Based Indexes of Urban Quality of Life.” In Peter M. Miezkowski and Mahlon R. Straszheim, eds., *Current Issues in Urban Economics*. Baltimore: Johns Hopkins University Press.

Tiebout, Charles M. 1956. “A Pure Theory of Local Expenditures.” *Journal of Political Economy*, vol. 64, no. 5, pp. 416–424. Available at https://doi.org/10.1086/257839

The Response of U.S. Investment to Oil Price Shocks: Does the Shale Boom Matter?

By Nida Çakır Melek

After an unprecedented decline from June 2014 to March 2016, the real price of oil more than doubled, renewing interest in the effects of oil price fluctuations on the U.S. economy. An increase in oil prices can affect the economy of a net oil importing country by lowering the consumers' demand for other goods and services. At the same time, an increase in oil prices can increase the cost of production, thereby lowering profits and reducing investment.

However, the effect of higher oil prices on investment may have changed. Following the recovery in oil prices from mid-2016 to mid-2018, U.S. oil investment almost doubled. This is unsurprising, as higher oil prices make oil businesses more profitable, allowing them to increase production and investment. More surprising is that total U.S. business fixed investment appears to have mimicked the pattern of oil investment: increasing until late 2014, declining in 2015, then increasing once again since 2016. This more recent, positive correlation between oil prices and U.S. investment growth may be related to the surge in U.S. oil production known as the shale boom.

In the mid-2000s, the United States began increasing its oil and gas production through horizontal drilling and hydraulic fracturing, becoming one of the largest oil producers in the world. As a result,

---

*Nida Çakır Melek is an economist at the Federal Reserve Bank of Kansas City. Colton Tousey, a research associate at the bank, helped prepare the article. This article is on the bank's website at **www.KansasCityFed.org***

when oil prices increase, U.S. oil producers benefit more now than in the past. These benefits may also spill over to other industries, positively affecting the overall economy. Given increasing production, declining imports, and some specific properties of shale oil—high initial production rates, the continuous drilling required to maintain production, and higher sensitivity to price changes—it is plausible to expect that the United States may respond differently to oil price changes now than in the past.

In this article, I explore the effect of unexpected oil price changes—or "shocks"—on U.S. investment, a key channel through which oil price shocks affect the economy. After controlling for the source of the changes in oil prices, I investigate the effect of oil supply, oil-specific (shocks that only affect the oil market), and aggregate demand shocks on oil and non-oil components of U.S. private nonresidential fixed investment and explore how the response of U.S. investment may have changed since the shale boom. I find that oil investment has become more responsive to oil supply, oil-specific, and aggregate demand shocks since the shale boom. The changing responsiveness to fluctuations in oil prices extends beyond the more direct effects on oil investment. I find that non-oil and total investment have also become more responsive to demand shocks and less responsive to oil supply shocks since the shale boom. These greater spillovers from the oil sector are not present prior to the shale boom.

I. Oil Price Changes and the U.S. Economy

The sharp oil price decline from June 2014 to March 2016 fueled a debate about its overall effect on the U.S. economy. Oil price shocks affect the economy primarily through consumers' and firms' spending (Hamilton 2008). A decline in oil prices can increase consumption, as consumers spend less of their income on fuels and increase their demand for other goods and services. Increased demand for a firm's output can, in turn, increase business capital spending too. The lower cost of production associated with declining oil prices can also lead to increased investment. As a result, many observers expected the 2014–16 oil price decline to boost U.S. economic growth. However, that boost did not seem to materialize (Baumeister and Kilian 2016b;

International Monetary Fund 2016). The reason may be related to the recent boom in U.S. oil production.

The shale revolution—brought about by the broad application of horizontal drilling and hydraulic fracturing—ended a decades-long period of declining U.S. oil production. In 2005, the United States began increasing its oil and gas production and became a major producer in a short time. Since 2013, the United States has been the world's top producer of petroleum hydrocarbons (Energy Information Administration 2018). Therefore, although the United States is still a net importer of oil, its dependence on foreign oil has declined significantly as net oil imports have plummeted.

In line with the increased production, the share of oil investment in total U.S. investment has also increased materially. The share of oil investment in total nonresidential fixed investment in structures and equipment increased from an average of 3.4 percent in the 1986:Q1–2005:Q4 period to an average of 10.5 percent in the 2006:Q1–2014:Q2 period. As a result, when oil prices fell sharply, the oil sector was hit hard. For example, Chart 1 shows that following the substantial decline in oil prices in the second half of 2014, U.S. oil investment collapsed. Given the increased importance of this sector, it is natural to ask whether the shale boom may have changed the overall relationship between oil and the U.S. economy.

How oil price shocks are transmitted to the economy has been a central question in macroeconomics and of interest to policymakers. Academic researchers have extensively studied the effects of oil price changes on consumer spending, finding quantitatively important effects (for example, Mehra and Petersen [2005]; Edelstein and Kilian [2009]; Gelman and others [2016], Ready [2018]; Iacoviello [2016]; Alsalman and Karaki [2017]; and Baumeister, Kilian, and Zhou [2018]). However, few have studied the response of firms' investment spending to oil price changes. Notable exceptions are Edelstein and Kilian (2007) and Loria (2017). Edelstein and Kilian (2007) investigate how nonresidential fixed investment in structures and equipment responds to energy price changes. They find that while the estimated response of nonresidential fixed investment in structures and equipment excluding oil is small and statistically insignificant, the estimated response of oil investment is large and statistically significant.[1] Loria (2017) shows that the

Chart 1
Real Oil Price versus Oil Investment

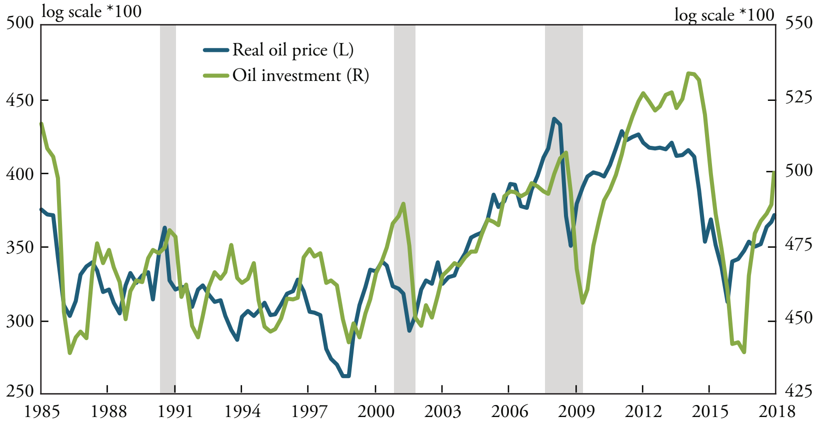

Notes: Variables are measured on a log scale and multiplied by 100 so that their swings represent percent changes. Gray bars denote National Bureau of Economic Research (NBER)-defined recessions.
Sources: BEA, BLS, EIA, NBER, and author's calculations.

size of the shock matters for the response of U.S. nonresidential fixed investment in structures and equipment. She finds that while a small oil price increase leads to a decline in investment, the effect of a large oil price increase is ambiguous, as it results in higher oil and oil-related investment but lower non-oil investment.[2]

Recent changes in the price of oil have raised questions once again about the effect of changing oil prices on the U.S. economy (see, for example, Klein [2018]; Liesman [2018]; and Yang and Sider [2018]). The real price of oil has more than doubled since its low in early 2016, bringing with it a substantial increase in oil investment (Chart 1). Motivated by the recent recovery in oil prices and the small amount of prior research, I investigate the effect of oil supply and demand shocks on U.S. private nonresidential fixed investment categories and assess whether this effect may have changed after the shale boom.

II. The Effects of Oil Price Shocks on Nonresidential Fixed Investment in Structures and Equipment

Oil prices fluctuate for several reasons. Increasing global economic activity can push up demand and increase oil prices, whereas a larger global oil supply can cause oil prices to decline. Oil prices can also move due to shifts in expectations about future oil supply or demand growth. Distinguishing between the factors driving oil price changes is important, as these factors tend to have very different effects on macroeconomic aggregates (Kilian 2009). For example, if an oil price increase is driven by an unexpected increase in global economic activity, aggregate investment will likely increase due to booming aggregate demand. However, if an oil price increase is driven by an unexpected decline in the global oil supply, aggregate investment will likely decline due to the higher cost of production. Investigating the dynamic effects of oil price shocks on U.S. investment requires a model that incorporates measures of the supply and demand shocks driving oil price changes.

Decomposing oil price shocks into oil supply and demand shocks

To disentangle the factors driving oil price movements, I use the framework from Kilian (2009). Kilian proposes a monthly three-variable structural vector autoregression (SVAR) to identify underlying demand and supply shocks in the global oil market. This framework identifies three shocks: aggregate demand shocks, oil-specific demand shocks (or precautionary demand shocks), and oil supply shocks. Aggregate demand shocks capture shifts in oil prices driven by changes in global real economic activity. These shocks reflect changes in demand for all industrial commodities. Oil-specific demand shocks, on the other hand, capture oil price changes driven by shocks specific to the crude oil market. For example, changes in expectations about future oil supply growth—such as an unexpected discovery of supply resulting in expectations of higher future supply growth—or demand growth can cause fluctuations in oil-specific demand. Finally, oil supply shocks capture shifts in oil prices driven by changes in the global oil supply. For example, a disruption to oil production would cause oil prices to increase.

I extract the monthly series of oil supply and demand shocks from the SVAR for two periods: January 1986 to December 2005—the pre-shale period—and January 1986 to December 2017—the full sample (see the appendix for details on how the shocks are recovered from observables). I end the pre-shale period in 2005 because that is the year U.S. oil and gas production began increasing (Çakır Melek 2015).

Responses of the real price of oil to the shocks extracted from the SVAR show that demand shocks and oil supply shocks have different effects (see the appendix for the responses and more details). Consistent with the findings in Kilian (2009) and Davig and others (2015), an unexpected positive aggregate demand shock leads to a persistent increase in the price of oil, while an unexpected positive oil-specific demand shock causes a sharp, very large, and persistent increase in the real price of oil that is also highly statistically significant. The increase in the price of oil due to a negative oil supply shock, on the other hand, is less persistent.

These findings reemphasize the importance of decomposing changes in real oil prices into oil supply shocks and oil-specific and aggregate demand shocks in examining the effect of oil price shocks on macroeconomic aggregates.

The response of U.S. investment to oil price shocks

After extracting the series of structural oil shocks from the model, I estimate the effects of these shocks on real private nonresidential fixed investment in structures and equipment as well as its components.

Because data on U.S. investment are only available at the quarterly frequency, I construct measures of quarterly oil price shocks by averaging the monthly structural innovations derived from the SVAR for each quarter. I then estimate the response of several U.S. investment categories to oil supply and demand shocks via ordinary least squares (OLS) according to the following equation:

$$\Delta I_{s,t} = \alpha_j + \sum_{i=1}^{4} \beta_{ji} \Delta I_{t-i} + \sum_{k=0}^{8} \gamma_{jk} OILSHOCK_{jt-k} + u_{jt,} \qquad j=1,2,3$$

where $\Delta I_{s,t}$ refers to the quarterly percent change of component s of real private nonresidential fixed investment in structures and equipment, α_j denotes the constant, u_{jt} denotes the error, and j refers to aggregate demand, oil-specific demand, and oil supply shocks.[3]

The Bureau of Economic Analysis (BEA)'s National Income and Product Accounts provide data on different components of private nonresidential fixed investment in structures and equipment. Table 1 shows the average shares of different components of private nonresidential fixed investment in structures and equipment. The BEA presents nonresidential private fixed investment in structures in five broad categories that have subcategories—or in some cases, sub-subcategories.

Table 1
Average Shares of Private Nonresidential Fixed Investment in Structures and Equipment by Category

Investment component	Share (percent)	
	1986:Q1–2005:Q4	1986:Q1–2017:Q4
Structures	31.2	32.5
Commercial and health care	13.1	11.8
Manufacturing	4.1	3.9
Mining exploration, shafts, and wells	3.0	5.0
Power and communication	4.7	5.4
Other	6.3	6.3
Equipment	68.8	67.5
Information processing	24.7	23.9
Industrial	15.8	15.0
Transportation	14.7	14.5
Other	13.2	13.2
Mining and oil field machinery	0.4	0.9
Total	100	100
Oil	3.4	5.9
Non-oil	96.6	94.1

Source: BEA.

Similarly, the BEA breaks nonresidential private fixed investment in equipment into four broad categories that themselves have subcategories or sub-subcategories.

In this article, I consider the following broad investment categories for the impulse responses: oil, non-oil, and total investment. I define oil investment as the sum of investment in mining exploration, shafts, and wells (structures) as well as investment in mining and oil field machinery (equipment). And I define non-oil investment as the sum of the rest of the components of investment in structures and equipment. Total investment is total private nonresidential fixed investment in structures and equipment. The average shares of these broad categories are presented near the bottom of Table 1.

Investment responses to aggregate demand shocks

Chart 2 shows the point estimates of the cumulative responses of each investment category to aggregate demand shocks.[4] The

Chart 2
Cumulative Response of Investment to Aggregate Demand Shocks

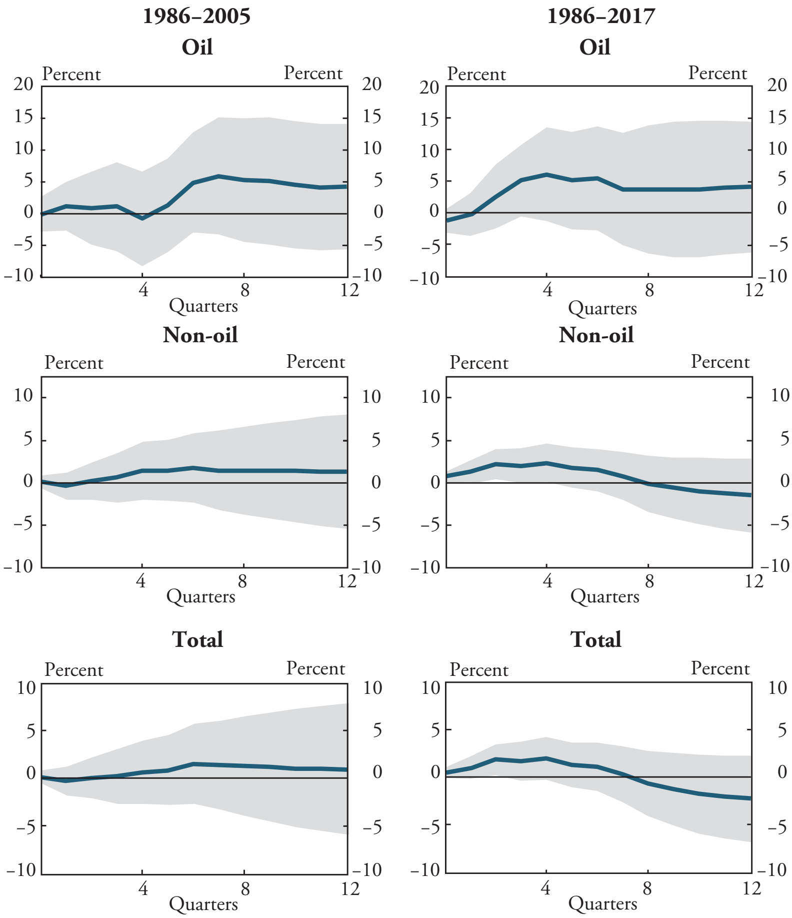

Notes: Solid lines show point estimates. Gray shaded regions represent 90 percent confidence intervals.
Sources: BEA, BLS, EIA, and author's calculations.

left-hand-side panel presents the results for the pre-shale period (1986:Q1–2005:Q4). An unexpected surge in aggregate demand drives up investment in all investment categories but with a delay of about a year in the pre-shale period. The responses of non-oil and total investment categories are flat in the first year before increasing and remaining positive. In response to an aggregate demand shock, oil investment shows a transitory increase in the first year followed by a large, sustained increase. However, the responses are not statistically significant.

The recent boom in U.S. oil production, brought about by the application of horizontal drilling and hydraulic fracturing, has been a major event in the global oil markets. In late 2017, average monthly oil production surpassed its previous peak in the 1970s. Moreover, U.S. net oil imports have declined substantially following the shale boom (Çakır Melek and Nie 2018). As the share of oil investment in U.S. aggregate investment has increased, U.S. investment patterns have changed too (Rodziewicz 2018). To assess the effect of the shale boom on the response of U.S. investment to oil price shocks, I next present impulses responses for the full sample covering 1986:Q1–2017:Q4.[5]

The right-hand-side panel in Chart 2 shows the responses of all three categories of investment in structures and equipment to aggregate demand shocks over the full sample period. An unexpected surge in aggregate demand drives up non-oil and total investment. The responses are positive on impact, statistically significant for about a year, and remain positive for almost two years before turning negative in the third year. The direct stimulating effect of higher aggregate demand on U.S. investment seems to dominate the indirect negative effect of higher oil prices in the short run. Higher costs due to higher oil (and other commodity) prices do not begin to weigh in for non-oil-related businesses until the second year. The response of oil investment, on the other hand, is positive at all horizons except in the first quarter and larger, but not statistically significant.

Investment responses to oil-specific demand shocks

Chart 3 shows the responses of each investment category to oil-specific demand shocks. The left-hand-side panel of Chart 3 presents pre-shale responses. In the pre-shale period, an unexpected increase in oil-specific demand causes oil investment to increase relatively sharply in the first year, with the effect peaking in the second year. The increase is sustained and mostly statistically significant. The responses of investment in the other two categories, however, are either flat or negative in the first year and are negative after. The responses are not statistically significant.

In the full sample period, overall, investment is even more responsive to oil-specific demand shocks, as shown by the right-hand-side panel of Chart 3. An unexpected increase in oil-specific demand causes oil investment to increase sharply in the first year. The increase is

Chart 3

Cumulative Response of Investment to Oil-Specific Demand Shocks

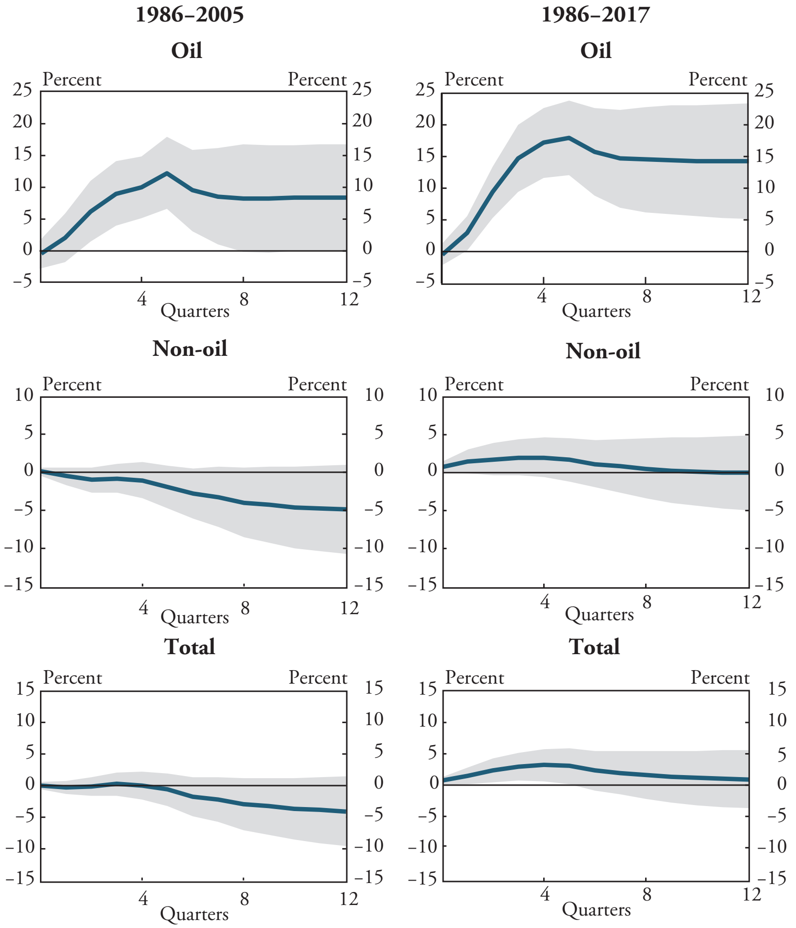

Notes: Solid lines show point estimates. Gray shaded regions represent 90 percent confidence intervals.
Sources: BEA, BLS, EIA, and author's calculations.

sustained and highly statistically significant. More striking is the positive response of non-oil investment to an unexpected increase in oil-specific demand in the first two years. Non-oil investment increases on impact, with the effect peaking at the end of the first year. The response is statistically significant for more than a quarter. These patterns carry over to the response of total investment. An oil price increase driven by an unexpected increase in oil-specific demand increases total investment on impact, with the effect peaking in the second year. The increase is persistent and statistically significant for more than one year.

The sustained increase in non-oil investment for more than a year in response to a positive oil-specific demand shock in the full sample is a notable result. Some might expect higher oil prices to *decrease* non-oil investment due to higher costs and depressed consumer demand. However, positive spillover effects from the oil sector could drive this result. Higher oil prices make oil businesses more profitable, boosting oil investment and potentially boosting investment in other sectors. To the extent that additional investment in the oil sector creates additional demand for other sectors, positive spillovers from the oil sector to the aggregate economy may lead to higher aggregate investment.

Although the oil sector represents a small share of the U.S. economy—around 1.6 percent of U.S. GDP in 2014—oil is an important production input and consumption good. As a result, oil price shocks can have important macroeconomic implications. Baqaee and Farhi (2017) find that negative shocks to crucial industries, such as "oil and gas," can have a significantly larger aggregate effect than negative shocks to larger but less crucial industries. Moreover, oil shocks are highly persistent and can thus generate significant welfare costs (Hitzemann and Yaron 2016). With potential spillover effects to other areas of the economy, the implications of oil price shocks might be amplified.

Given the emergence of the shale oil sector in the past decade, positive spillover effects are likely. In fact, several studies document such effects. For instance, Allcott and Keniston (2018); Feyrer, Mansur, and Sacerdote (2017); and Gilje, Ready, and Roussanov (2016) examine the local implications of the shale boom and find strong positive spillovers to employment and wages at the local and regional level. Çakır Melek, Plante, and Yücel (2017) and Bjørnland and Zhulanova (2018) investigate spillovers to the aggregate economy after the shale oil boom and find positive spillovers to output and investment.

In that context, the responses presented in the left and right panels of Chart 3 reveal a key result that positive spillovers were not present before the shale boom. In the pre-shale period, an oil price increase driven by shocks specific to the oil market induces almost no effect—and in the second year, a negative effect—on non-oil investment. The response of total investment is flat in the first year before turning negative. These responses contrast with the positive, significant on impact responses of non-oil and total investment in the full sample, which peak in about a year.

Investment responses to oil supply shocks

Chart 4 presents the responses of investment to an unexpected decline in the global oil supply in the pre-shale period and in the full sample. The left-hand-side panel shows that pre-shale responses are mostly negative at all horizons in all categories. The response of oil investment is flat in the first year and statistically insignificant at all horizons. The negative responses of non-oil investment and total investment are persistent and statistically significant for about a year.

The right-hand-side panel in Chart 4 shows that an unexpected decline in global oil supply causes non-oil investment and total investment to decrease and remain persistently low at all horizons in the full sample as well. Oil investment, on the other hand, turns from negative to positive in the second year and remains positive thereafter. All three responses are statistically insignificant.[6]

Key takeaways

Together, the impulse responses presented in Charts 2–4 reveal important differences in how the demand and supply shocks underlying the real price of oil affect U.S. investment—in other words, the source of the shock matters. In addition, the full sample results reveal that overall, investment is more responsive to aggregate and oil-specific demand shocks in the first two years, but not as responsive to disruptions in the global oil supply. This is consistent with results showing that the real price of oil is more responsive to aggregate and oil-specific demand shocks than oil supply shocks, and that recently, oil price fluctuations have been driven largely by aggregate and oil-specific demand shocks (see the appendix for details).

Comparing the impulse responses for the pre-shale period to those for the full sample yields two more key findings. Oil investment is less responsive to oil price shocks in the pre-shale period than in the full sample, regardless of the type of shock. In the pre-shale period, the response of oil investment to an oil price increase is delayed: either the recovery begins much later or the increase is less sharp than in the full sample (the top panels of Charts 2–4). This is in line with conventional

Chart 4
Cumulative Response of Investment to Oil Supply Shocks

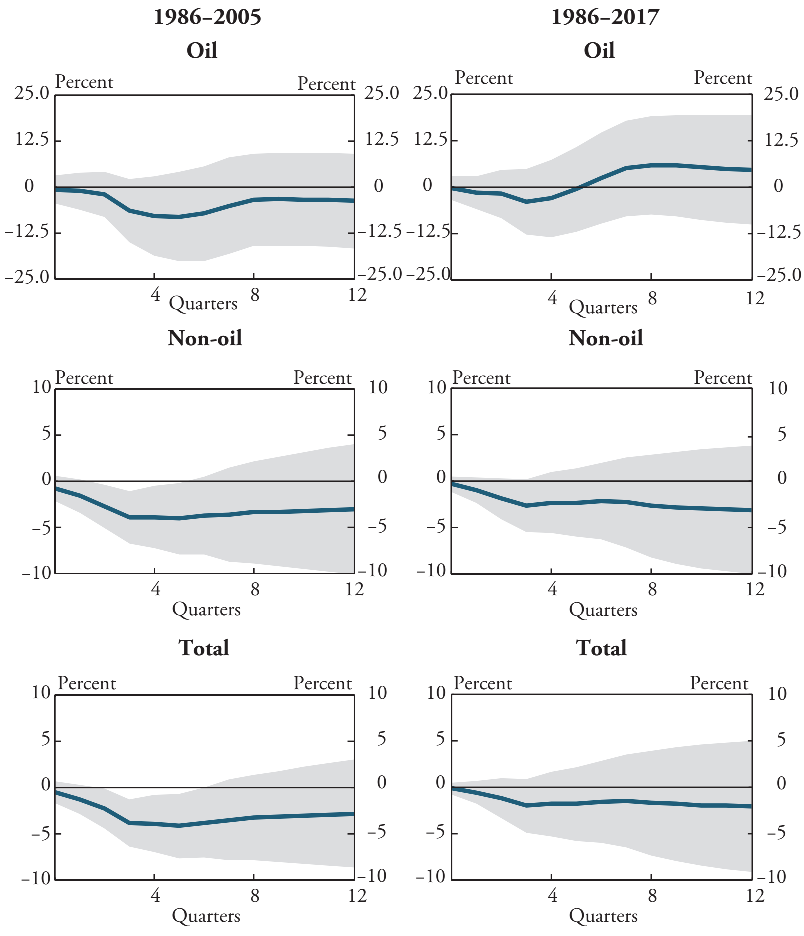

Notes: Solid lines show point estimates. Gray shaded regions represent 90 percent confidence intervals.
Sources: BEA, BLS, EIA, and author's calculations.

oil production's lower sensitivity to oil price changes than shale (unconventional) production, which relies more on investment.[7] Shale production is more capital intensive with high initial production rates, requires continuous drilling for maintaining production, and is more responsive to price changes, with shorter investment payback periods.

Finally, investment is less responsive to an aggregate demand shock but more responsive to an oil supply shock in the pre-shale period. A positive aggregate demand shock increases non-oil and total investment in the pre-shale period, too, but with a delay, and the responses are not statistically significant. An oil price increase driven by a negative oil supply shock, on the other hand, causes a larger and significant decline in non-oil and total investment for about a year in the pre-shale period compared with a more muted, insignificant response in the full sample.

In summary, my findings suggest that the U.S. shale boom changed the response of U.S. investment to oil price shocks. Oil investment has become more responsive to an oil price increase, inducing higher non-oil investment. These positive spillovers from an oil price increase to U.S. aggregate investment are not present in the pre-shale period.

III. Conclusion

Oil prices have more than doubled since their lows in early 2016, renewing interest in the effect of increasing oil prices on the U.S. economy. A primary channel through which oil price shocks affect the economy is investment. After controlling for the source of changes in oil prices, I find that oil investment has become more responsive to oil supply and demand shocks since the shale boom. Changes in the sensitivity of investment to oil price fluctuations extend beyond the more direct effects on oil investment. The response of non-oil investment has changed too. Non-oil investment has also become more responsive to demand shocks and less responsive to oil supply shocks since the shale boom, a pattern which has carried over to aggregate investment.

Together, these results suggest that the increased U.S. presence in the global oil market has led to greater spillovers from the oil sector to the aggregate economy. These findings may help explain why U.S. investment exhibited recessionary-like dynamics following the substantial decline in oil prices from June 2014 to March 2016 and then recovered as oil prices rose.

Appendix

Decomposing Changes in the Real Price of Oil into Oil Supply and Demand Shocks

Structural vector autoregressions (SVARs) are commonly used to model the global oil market and study the effect of oil price shocks on macroeconomic aggregates. The shocks recovered from an SVAR may differ depending on the variables included in the model and the identification scheme. The SVAR specification I use in this article is similar to Davig and others (2015), which is based on Kilian (2009).

I use a three-variable SVAR based on monthly data that include the percent change in global oil production, a suitable index for real global economic activity, and the real price of oil. The model is identified recursively. Oil supply does not respond contemporaneously (within a month) to changes in oil demand—that is, the short-run supply curve is vertical.[8] Additionally, changes in real oil prices driven by oil-specific shocks have no contemporaneous effect on global economic activity.

The SVAR covers the sample periods of January 1986 to December 2005 and January 1986 to December 2017 with four lags. The choice of starting date is motivated by Baumeister and Peersman (2013).[9] The lag order tends to be larger than estimates suggested by the Akaike Information Criterion conditional on an upper bound of 12 lags. Although the qualitative results are not sensitive to the lag order choice, I adopt a conservative approach due to the possibility of underfitting a VAR model (Hamilton and Herrera 2004).

Global oil production is obtained from the Energy Information Administration (EIA), and an updated time series for Kilian's (2009) index of global economic activity is obtained from his website.[10] The real price of oil is measured as the refiners' acquisition cost of imported crude oil, provided by the EIA, which is deflated by the U.S. Consumer Price Index (CPI). Both the real economic activity index and the real price of oil are expressed in logs. The model is estimated following Kilian (2009).

Charts A-1 and A-2 show the responses of the real price of oil to one standard deviation structural changes in the pre-shale period and in the full sample, respectively. Solid blue lines show point estimates with 90 percent confidence intervals. In the pre-shale period,

Chart A-1

Oil Price Responses to One-Standard-Deviation Structural Shocks, 1986–2005

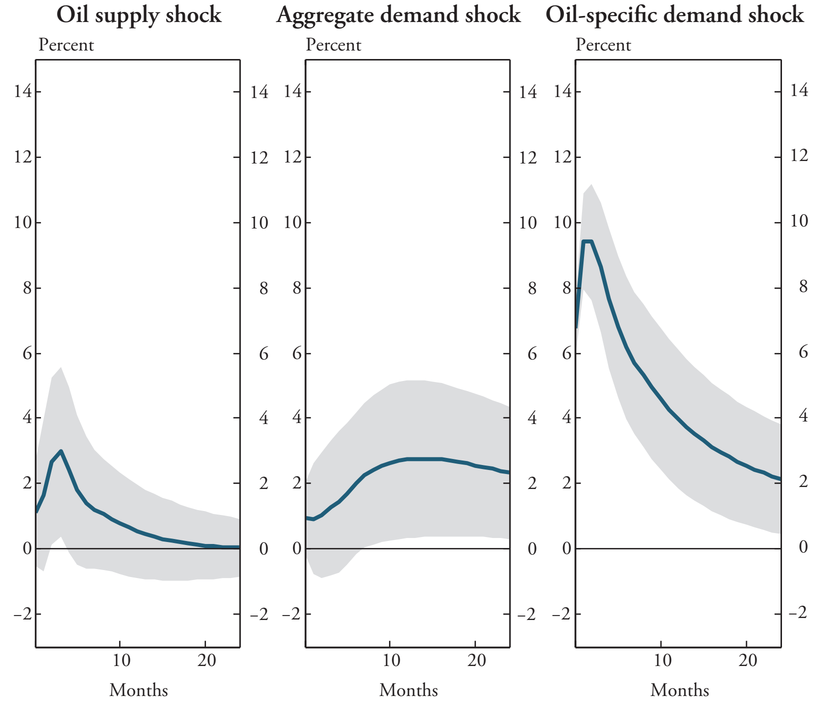

Notes: Solid lines show point estimates. Gray shaded regions represent 90 percent confidence intervals.
Sources: BLS, EIA, and author's calculations.

Chart A-2

Oil Price Responses to One-Standard-Deviation Structural Shocks, 1986–2017

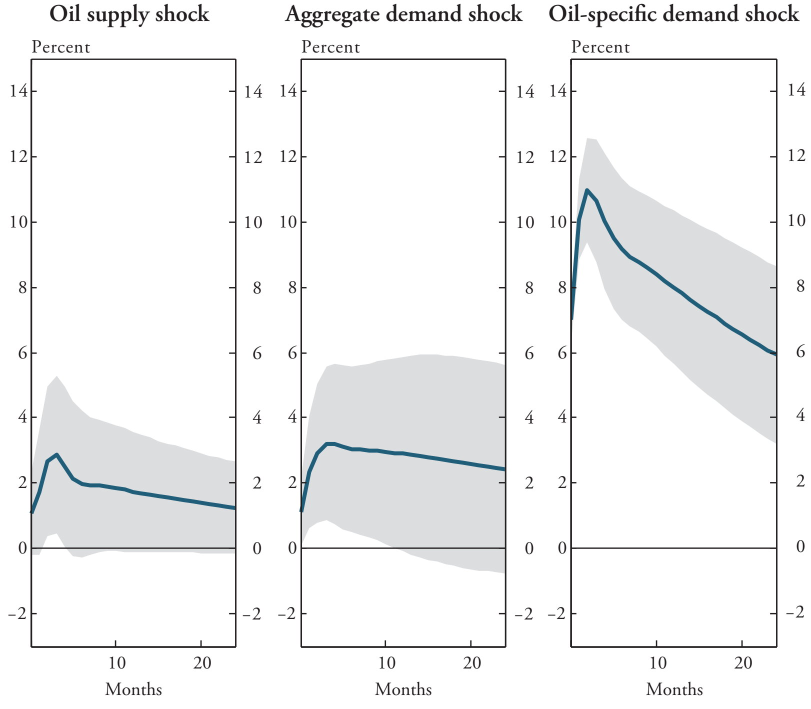

Notes: Solid lines show point estimates. Gray shaded regions represent 90 percent confidence intervals.
Sources: BLS, EIA, and author's calculations.

an unexpected negative oil supply shock causes a transitory increase in the price of oil. The effect of an unanticipated positive aggregate demand shock is similar on impact to that of a negative oil supply shock but highly persistent. A surprise increase in oil-specific demand has a more distinct effect on the price of oil. It causes a sharp, very large, and persistent increase in the real oil price, which is also highly statistically significant. The oil price responses in the full sample are similar to the pre-shale responses except a persistent increase in the real oil price in response to a negative supply shock in the full sample. These results are overall similar to estimates obtained by Davig and others (2015) using data from January 1985 to March 2015.

A historical decomposition is useful for understanding the contribution of these shocks to oil price movements. Charts A-3 and A-4 present the respective cumulative contribution of oil supply and demand shocks to the real price of oil over the two sample periods. Overall, the charts show that the real oil price is indeed driven by all shocks at all times, but their contributions differ. Historically, oil supply shocks have made smaller contributions to the real price of oil relative to demand shocks. And while aggregate demand shocks have caused long swings, oil-specific shocks are associated with fairly sharply defined swings in oil price. Charts A-3 and A-4 also suggest that demand shocks have played a larger role in recent episodes of large oil price changes, such as the global financial crisis and the 2014–16 oil price declines.[11] Oil-specific demand shocks played an especially important role in the 2014–16 decline (for more details, see Davig and others [2015]).

Chart A-3
Historical Decomposition of the Real Price of Oil, 1986–2005

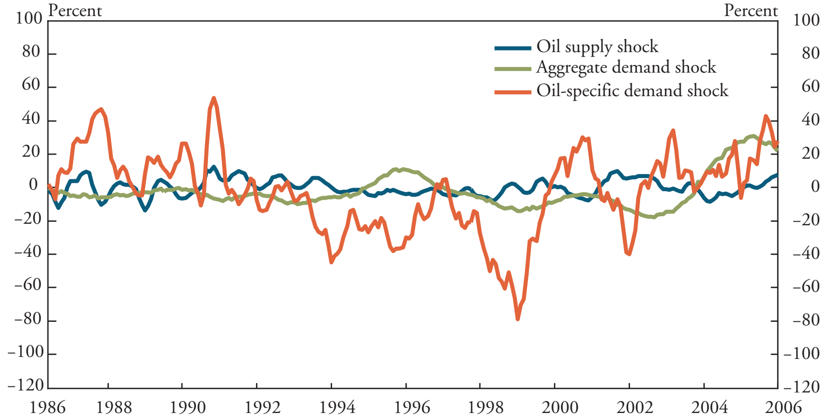

Sources: BLS, EIA, and author's calculations.

Chart A-4
Historical Decomposition of the Real Price of Oil, 1986–2017

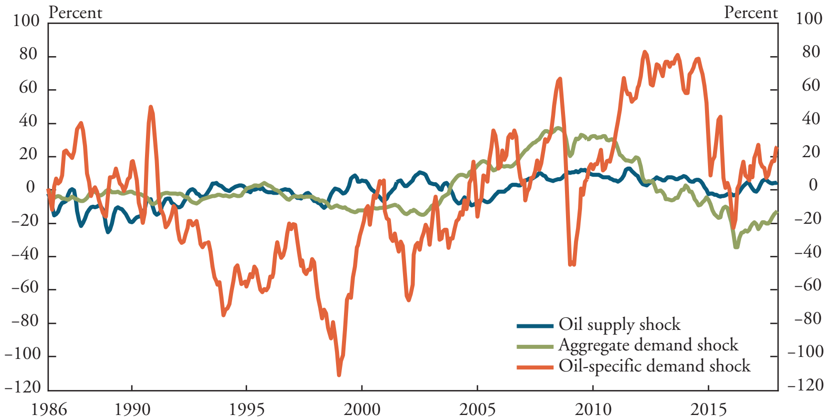

Sources: BLS, EIA, and author's calculations.

Endnotes

[1]These results rely on the assumption that investment responds symmetrically to energy price increases and decreases.

[2]She then builds a model to explain her empirical findings and shows that the oil firm's ability to cover high fixed costs in the sector, which depends on the size of the oil price shock, is important in understanding the responses.

[3]The lag structure follows existing research—see, for example, Hamilton (2003), Edelstein and Kilian (2007), Kilian (2008), and Baumeister and Kilian (2016b).

[4]I obtain the level responses for the three investment categories by cumulating the estimated impulse responses.

[5]Splitting the sample at 2005 results in too few observations for the empirical analysis in the post-2005 period.

[6]Oil investment being more responsive to demand shocks than supply shocks is consistent with Bornstein, Krusell, and Rebelo's (2017) finding that investment in the oil industry is driven mostly by demand shocks.

[7]See, for example, Dale (2016); and Bjørnland, Nordvik, and Rohrer (2017).

[8]The very low short-run oil supply elasticity estimate (annual 0.12) reported by Bornstein, Krusell, and Rebelo (2017)—along with the finding of Anderson, Kellogg, and Salant (2018) that oil production from existing wells in Texas does not respond to oil prices—supports this assumption.

[9]The authors find a considerable break in oil market dynamics in the first quarter of 1986 in a time-varying SVAR framework. Moreover, prior research frequently uses this date for splitting samples, which coincides with the collapse of the Organization of the Petroleum Exporting Countries (OPEC)'s market share and the start of the Great Moderation.

[10]The index is the cumulative average of the increase in bulk dry cargo ocean freight rates, deflated by the U.S. CPI and linearly detrended. Kilian and Zhou (2018) provide detailed information on this index and other indicators of global real economic activity.

[11]For a more detailed discussion on major oil price events, see Hamilton (2011) and Baumeister and Kilian (2016a).

References

Allcott, Hunt, and Daniel Keniston. 2018. "Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America." *The Review of Economic Studies*, vol. 85, no. 2, pp. 695–731. Available at https://doi.org/10.1093/restud/rdx042

Alsalman, Zeina N., and Mohamad B. Karaki. Forthcoming. "Oil Prices and Personal Consumption Expenditures: Does the Source of the Shock Matter?" *Oxford Bulletin of Economics and Statistics*.

Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant. 2018. "Hotelling Under Pressure." *Journal of Political Economy*, vol. 126, no. 3, pp. 984–1026. Available at https://doi.org/10.1086/697203

Baqaee, David Rezza, and Emmanuel Farhi. 2017. "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem." NBER Working Paper no. 23145, February. Available at https://doi.org/10.3386/w23145

Baumeister, Christiane, Lutz Kilian, and Xiaoqing Zhou. 2018. "Is the Discretionary Income Effect of Oil Price Shocks a Hoax?" *The Energy Journal*, vol. 39, no. 2, pp. 117–137. Available at https://doi.org/10.5547/01956574.39.SI2.cbau

Baumeister, Christiane, and Lutz Kilian. 2016a. "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us." *Journal of Economic Perspectives,* vol. 30, no. 1, pp. 139–160. Available at https://doi.org/10.1257/jep.30.1.139

———. 2016b. "Lower Oil Prices and the U.S. Economy: Is This Time Different?" *Brookings Papers on Economic Activity*, Fall, pp. 287–357. Available at https://doi.org/10.1353/eca.2016.0029

Baumeister, Christiane, and Gert Peersman. 2013. "Time-Varying Effects of Oil Supply Shocks on the U.S. Economy." *American Economic Journal: Macroeconomics*, vol. 5, no. 4, pp. 1–28. Available at https://doi.org/10.1257/mac.5.4.1

Bjørnland, Hilde C., Frode Martin Nordvik, and Maximilian Rohrer. 2017. "Supply Flexibility in the Shale Patch: Evidence from North Dakota." Norges Bank Research working paper no. 9.

Bjørnland, Hilde C., and Julia Zhulanova. 2018. "The Shale Oil Boom and The U.S. Economy: A Time-Varying Analysis." Working paper, May.

Bornstein, Gideon, Per Krusell, and Sergio Rebelo. 2017. "Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry." NBER Working Paper no. 23423, May. Available at http://doi.org/10.3386/w23423

Çakır Melek, Nida. 2015. "What Could Lower Prices Mean for U.S. Oil Production?" Federal Reserve Bank of Kansas City, *Economic Review,* vol. 100, no. 1, pp. 51–69.

Çakır Melek, Nida, and Jun Nie. 2018. "What Could Resurging U.S. Energy Production Mean for the U.S. Trade Deficit?" Federal Reserve Bank of Kansas City, *The Macro Bulletin*, March.

Çakır Melek, Nida, Michael Plante, and Mine K. Yücel. 2017. "The U.S. Shale Oil Boom, the Oil Export Ban, and the Economy: A General Equilibrium Analysis." Federal Reserve Bank of Kansas City, Research Working Paper no. 17-10, September. Available at https://doi.org/10.18651/RWP2017-10

Dale, Spencer. 2016. "New Economics of Oil." *Oil and Gas, Natural Resources, and Energy Journal*, vol. 1, no. 5, pp. 365–379.

Davig, Troy, Nida Çakır Melek, Jun Nie, A. Lee Smith, and Didem Tüzemen. 2015. "Evaluating a Year of Oil Price Volatility." Federal Reserve Bank of Kansas City, *Economic Review,* vol. 100, no. 3, pp. 5–30.

Edelstein, Paul, and Lutz Kilian. 2009. "How Sensitive Are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics*, vol. 56, no. 6, pp. 766–779. Available at https://doi.org/10.1016/j.jmoneco.2009.06.001

———. 2007. "The Response of Business Fixed Investment to Changes in Energy Prices: A Test of Some Hypotheses about the Transmission of Energy Price Shocks." *The B.E. Journal of Macroeconomics*, vol. 7, no. 1, pp. 1–41. Available at https://doi.org/10.2202/1935-1690.1607

Energy Information Administration (EIA). 2018. "United States Remains the World's Top Producer of Petroleum and Natural Gas Hydrocarbons." *Today in Energy*, May 21.

Feyrer, James, Erin T. Mansur, and Bruce Sacerdote. 2017. "Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution." *American Economic Review*, vol. 107, no. 4, pp. 1313–1334. Available at https://doi.org/10.1257/aer.20151326

Gelman, Michael, Yuriy Gorodnichenko, Shachar Kariv, Dmitri Koustas, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2017. "The Response of Consumer Spending to Changes in Gasoline Prices." NBER Working Paper no. 22969, November. Available at https://doi.org/10.3386/w22969

Gilje, Erik, Robert Ready, and Nikolai Roussanov. 2016. "Fracking, Drilling, and Asset Pricing: Estimating the Economic Benefits of the Shale Revolution." NBER Working Paper no. 22914, December. Available at https://doi.org/10.3386/w22914

Hamilton, James D. 2011. "Historical Oil Shocks." NBER Working Paper no. 16790, February. Available at https://doi.org/10.3386/w16790

———. 2008. "Oil and the Macroeconomy." *The New Palgrave Dictionary of Economics*, Second Edition.

———. 2003. "What is an Oil Shock?" *Journal of Econometrics*, vol. 113, no. 2, pp. 363–398. Available at https://doi.org/10.1016/S0304-4076(02)00207-5

Hamilton, James D., and Ana Maria Herrera. 2004. "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: A Comment." *Journal of Money, Credit, and Banking*, vol. 36, no. 2, pp. 265–286. Available at https://doi.org/10.1353/mcb.2004.0012

Hitzemann, Steffen, and Amir Yaron. 2017. "Welfare Costs of Oil Shocks." Society for Economic Dynamics, 2017 Meeting Papers no. 1381.

Iacoviello, Matteo. 2016. "Oil Prices and Consumption Across Countries and U.S. States." Working paper, November.

International Monetary Fund. 2016. "Too Slow for Too Long." *World Economic Outlook*, April.

Kemp, John. 2018. "Rising Oil Prices Boost U.S. economy: Kemp." *Reuters*, May 15.

Kilian, Lutz. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review,* vol. 99, no. 3, pp. 1053–1069. Available at https://doi.org/10.1257/aer.99.3.1053

———. 2008. "A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries." *Journal of the European Economic Association*, vol. 6, no. 1, pp. 78–121. Available at https://doi.org/10.1162/JEEA.2008.6.1.78

Kilian, Lutz, and Xiaoqing Zhou. 2018. "Modeling Fluctuations in the Global Demand for Commodities." *Journal of International Money and Finance*, vol. 88, pp. 54–78. Available at https://doi.org/10.1016/j.jimonfin.2018.07.001

Klein, Matthew C. 2018. "How Rising Oil Prices Will Affect the U.S." *Barrons*, May 11.

Liesman, Steve. 2018. "Rising Oil Prices, Once a Big Negative, May Now Be 'a Wash' for the U.S. Economy." CNBC, May 14.

Loria, Francesca. 2017. "The Effect of Oil Price Shocks on U.S. Investment: Size Matters." Job market paper, European University Institute.

Mehra, Yash P., and Jon D. Petersen. 2005. "Oil Prices and Consumer Spending." Federal Reserve Bank of Richmond, *Economic Quarterly*, vol. 91, no. 3, pp. 53–72.

Ready, Robert C. 2018. "Oil Consumption, Economic Growth, and Oil Futures: The Impact of Long-Run Oil Supply Uncertainty on Asset Prices." *Journal of Monetary Economics*, vol. 94, pp. 1–26. Available at https://doi.org/10.1016/j.jmoneco.2017.09.010

Rodziewicz, David. 2018. "Energy Investment Variability within the Macroeconomy." Federal Reserve Bank of Kansas City, *Economic Review*, vol. 103, no. 3, pp. 53–75. Available at https://doi.org/10.18651/ER/3q18Rodziewicz

Yang, Stephanie, and Alison Sider. 2018. "Oil Is Fast Approaching $70. Is the Economy Ready for It?" *The Wall Street Journal*, April 22.

Machine Learning Approaches to Macroeconomic Forecasting

By Aaron Smalter Hall

Forecasting macroeconomic conditions can be challenging. Accurate forecasts require an approach complex enough to incorporate relevant economic data but focused enough to exclude irrelevant data. Most forecasters attempt to balance focus and complexity with one of two approaches: aggregating forecasts from professional forecasters into a consensus prediction or using statistical modeling techniques typically specialized for time-series data. However, both approaches involve discretionary choices about data and methods. Although forecasters underpin the choices they make about models and complexity with economic intuition and judgement, these assumptions can be flawed.

Machine learning approaches, on the other hand, automate as many of those choices as possible in a manner that is not subject to the discretion of the forecaster. Indeed, optimizing a model without assumptions or judgement is a central issue of study in machine learning. In this article, I apply machine learning techniques to find an optimal forecasting model for the unemployment rate. After identifying a model that is not subject to discretion or assumptions, I compare the forecasts obtained from the machine learning model to forecasts obtained through traditional forecasting approaches. My experiments show that when supplied with diverse and complex data, the machine learning model

---

Aaron Smalter Hall is a senior data scientist at the Federal Reserve Bank of Kansas City. This article is on the bank's website at ***www.KansasCityFed.org***

can outperform simpler time-series models as well as a consensus of professional forecasters, with better performance at shorter horizons. In particular, experiments show that a machine learning model can identify turning points in the unemployment rate earlier than competing methods, suggesting that it can provide forecasters with more guidance about cyclical fluctuations than a consensus or autoregressive forecast.

Section I reviews the statistical and consensus approaches for macroeconomic forecasting. Section II discusses how machine learning can be used to optimize model complexity. Section III shows that one machine learning model, the Elastic Net, can outperform traditional models at all horizons, detect turning points earlier, and identify variables that can help predict unemployment from among a broad set.

I. Popular Approaches to Unemployment Forecasting

The unemployment rate is a challenging variable to forecast because it can be influenced by a wide variety of other variables. In the past, forecasters have attempted to overcome these challenges through statistical modeling or consensus forecasting. Statistical modeling benefits from strong mathematical theory that can lead to less biased forecasts and more consistent predictions. One of the most common statistical models for time-series forecasting is an autoregressive (AR) model, which assumes future observations are predicted primarily by recent observations. The unemployment rate is largely consistent with that assumption. In addition to AR models, forecasters often use a random walk model as a control or baseline. In a random walk, predictions for tomorrow's unemployment rate are based on the unemployment rate today plus an additional random shock. If a model cannot perform better than a random walk, it is unlikely it will be useful in making predictions.

Some forecasters might assume that a wide variety of variables, in addition to the past unemployment rate, may be useful in forecasting unemployment and therefore incorporate additional information in their forecasting models (Barnichon and Nekarda 2013). Still, they may overlook other important information that might influence the unemployment rate, including info on housing markets, consumer sentiment, and interest rates.

Consensus forecasting, contrasted with statistical models, leverages the wisdom of the crowd and benefits from combining as many different kinds of models as there are forecasters. These models range from sophisticated empirical systems of equations to simple gut instinct. Consensus forecasts are available for the unemployment rate from both Blue Chip Economic Indicators and the Survey of Professional Forecasters (SPF). By drawing from a variety of techniques and methods, these consensus forecasts can be more accurate than any particular professional forecast (Clemen 1989).

Both statistical modeling and consensus forecasting require a range of different choices that affect model complexity and forecasting performance. The AR model requires some very clear choices—specifically, the number of lagged variables to include. This choice can have a significant effect on AR forecasting performance. Although many researchers use some metric such as the Bayesian information criterion to determine the number of lags, they still must choose a metric and the maximum number of lags. There are not many models that do not require any choices, but the random walk model is one example; the model is purely a function of random movements based on past observations of the unemployment rate.

On the surface, the consensus forecast may not seem to involve many choices, as all the choices are being made by the forecasters themselves. This is a desirable quality, which is why the Blue Chip forecasts are commonly used as a competitive benchmark. However, researchers using these forecasts must still make choices about how to use them. Some may decide to use the median forecast instead of the average forecast. Others may choose to average only top-performing forecasters.

Rather than make restrictive modeling assumptions, researchers might alternatively consider a model that can admit any number of variables. Here, the field of choices begins to explode. For example, the FRED economic data platform, housed at the Federal Reserve Bank of St. Louis, advertises 527,000 economic series—which of these should researchers choose to include?

Machine learning techniques and models can simplify this set of choices tremendously. Models such as Elastic Net allow researchers to make a single choice about how to constrain the model's complexity rather than considering the inclusion or exclusion of all potential series. Moreover, other machine learning techniques can help researchers discover the optimal model complexity, eliminating that choice as well.

II. Introducing Machine Learning and Forecasting

The field of machine learning encompasses a wide variety of models. One way to identify the optimal forecasting model is to decompose the source of forecast errors into two parts—bias and variance. The bias of a model describes errors due to inaccurate assumptions about a problem or data. The variance of a model describes errors due to a model's sensitivity to small perturbations in the underlying data.

Figure 1 shows the effects of bias and variance on model predictions. In particular, the figure shows four bullseyes arising from the combination of two different conditions: high or low bias and high or low variance. The orange center of the bullseye represents the correct prediction, and the blue dots represent model predictions. Bias indicates how close to the center of the bullseye a model's predictions are, while variance indicates how tightly or loosely the predictions are clustered. Models with high bias and low variance yield observations that are tightly clustered but far from the correct prediction. Models with low bias and high variance, on the other hand, yield observations that are only loosely clustered around the correct prediction. Both of these sources of errors are related to the complexity of a model. Understanding the tradeoff between bias and variance can help forecasters optimize model complexity and thus minimize the total error in their predictions.

The complexity of a model can be assessed in many different ways, which are often dependent on the nature of the model itself. Sources of model complexity include the number of variables in a model, the number of parameters a model "learns," the number of parameters a user defines, or the number of variable relationships a model captures internally.

The relationship between model complexity, bias, and variance is often illustrated as a trade-off. Figure 2 shows that as the complexity of a model increases, the model typically exhibits less bias but more variance. The dashed line indicates optimum model complexity, where the total error due to both bias and variance is minimized and thus the corresponding forecast error for new observations (called the generalization error) is minimized.

To further illustrate this trade-off, consider a forecaster using a linear regression model estimated using ordinary least squares (OLS) to forecast the unemployment rate. In this case, the complexity of the model refers to the number of explanatory variables the forecaster chooses to

Figure 1

Errors of a Model Due to Bias and Variance

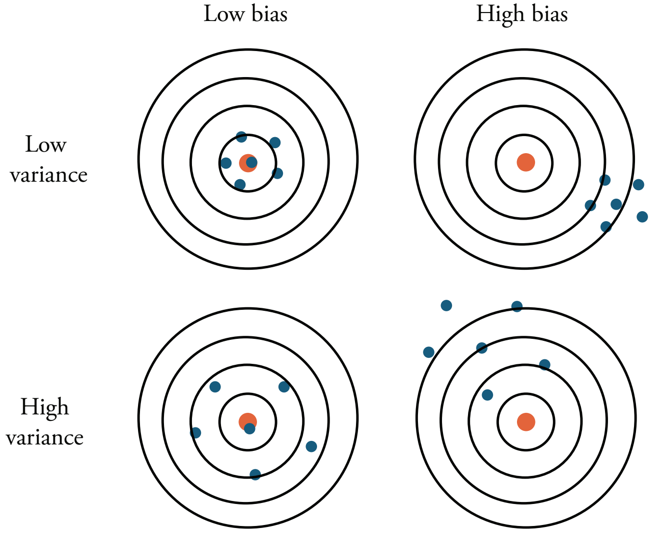

Note: The orange dot represents the true output a model is trying to capture, while the blue dots represent the learned outputs of a model under different conditions.

Figure 2

Trade-Off between Errors Due to Variance and Errors Due to Bias

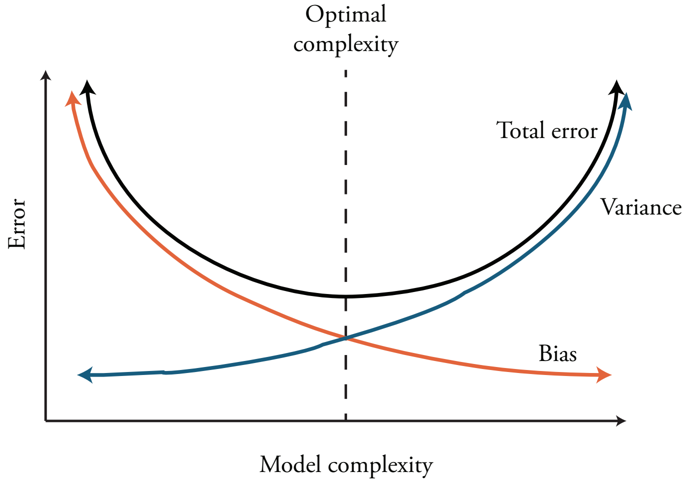

Note: Minimally complex models typically exhibit low variance but high bias, while maximally complex models exhibit the opposite.

include. If the forecaster uses a simple model with a single variable for the previous period's observed unemployment rate, then the model will have large forecast errors from bias due to the likely incorrect assumption that future unemployment is only dependent on current unemployment. However, the model will have small errors due to variance because the model is robust to small perturbations in the underlying data (that is, small changes in the observed variable translate to small changes in the forecast, proportional to the estimated model coefficient).

The forecaster can increase the complexity of their model by adding variables. In doing so, the forecast error from bias will decline because the assumption that only current unemployment determines future unemployment has been relaxed. However, the forecast error from variance will increase because small perturbations in lots of variables (including those that may be irrelevant to the unemployment rate) will now translate to disproportionally large changes in the forecast.

The Elastic Net model, which has become popular in machine learning, can precisely balance errors due to bias and variance through regularization. Regularization effectively "disciplines" a model by introducing penalties for overfitting the data. The Elastic Net model blends two different kinds of regularization by penalizing both the number of variables in the model and the extent to which any given variable contributes to the model's forecast (the magnitude of the variable's coefficient). By applying these penalties, Elastic Net learns which variables are most important, eliminating the need for researchers to make discretionary choices about which variables to include (further details on how the penalties are assessed are available in the appendix). In this way, the model is able to pare complex data down to only what is needed to capture the complexity of the problem at hand.

III. Comparing Machine Learning to Consensus and Statistical Forecasts

Machine learning provides a strong set of tools for finding the optimal complexity of a model, freeing forecasters from the need to use strong assumptions or judgement to simplify their models. To assess whether these advantages yield more accurate forecasts, I compare the results from one regularized machine learning model, the Elastic Net, to typical consensus and statistical forecasts of the U.S. civilian unemployment rate at a monthly frequency. I also assess these methods' utility in identify turning points in the business cycle as seen in the unemployment rate.

The benchmark models used for comparison are the consensus forecast from Blue Chip, an AR model, and a random walk. I use Blue Chip forecasts rather than forecasts from the SPF, as the SPF does not forecast at a monthly frequency.[1] The AR model is an integrated model that uses monthly changes in the unemployment rate to forecast the future unemployment rate. I choose the lag length using the Bayesian information criterion, with a maximum lag length of six months. By definition, the random walk and AR models only use past observations of the unemployment rate to forecast future outcomes.

In contrast, the Elastic Net model forecasts the future unemployment rate using FRED-MD, a diverse variable set of 138 macroeconomic variables drawn from a number of economic categories (McCracken and Ng 2015). Adding these variables allows the Elastic Net model more potential complexity but allows overfitting as well. The combined data set spans March 1959 through April 2017 for a total of 698 monthly observations. I use the first half of this sample (March 1959 through June 1988) to train the initial models and fit the coefficients, and the second half of the sample (July 1988 through April 2017) for the forecast comparison. I use a rolling forecast framework, in which I train a model using data observed up to current period t to fit the coefficients, make a forecast for period $t + 1$, then advance one period and repeat until the end of the forecast span. I consider forecasting horizons of up to two years ahead at three-month intervals (three, six, nine, 12, and 24 months).

Forecast accuracy

Table 1 shows that Elastic Net yields the most accurate forecasts as measured by the mean absolute error (MAE) for all horizons. Elastic Net outperforms even Blue Chip forecasting, indicating that the model is effectively learning the underlying relationships between economic series. Table 2 shows that Elastic Net outperforms Blue Chip by 0.07 percent on average. The advantage is more pronounced at shorter horizons, and the learned relationships are less useful at horizons above 12 months. The performance advantage of Elastic Net over Blue Chip is

Table 1
MAE Results

Horizon (months)	Random walk (percent)	AR (percent)	Blue Chip (percent)	Elastic Net (percent)
3	0.35	0.23	0.26	0.17
6	0.53	0.38	0.39	0.27
9	0.79	0.58	0.52	0.42
12	1.00	0.82	0.65	0.58
15	1.22	1.04	0.76	0.74
18	1.49	1.30	0.92	0.88
21	1.69	1.55	–	1.02
24	1.91	1.78	–	1.20
Average	1.12	0.96	–	0.66

Note: Forecasting results shown for all horizons with each of the four methods.
Sources: Blue Chip Economic Indicators, Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

Table 2
Differences between MAE for Elastic Net and Other Forecasts

Horizon (months)	Random walk (percent)	AR (percent)	Blue Chip (percent)
3	0.18	0.06	0.09
6	0.26	0.11	0.12
9	0.37	0.16	0.10
12	0.42	0.24	0.07
15	0.48	0.30	0.02
18	0.61	0.42	0.04
21	0.67	0.53	–
24	0.71	0.58	–
Average	0.46	0.30	0.07

Note: Positive values indicate better performance for Elastic Net.
Sources: Blue Chip Economic Indicators, Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

Table 3
P-Values for Statistical Significance

Horizon (months)	Random walk	AR	Blue Chip
3	0.008	0.000	0.000
6	0.006	0.000	0.000
9	0.171	0.000	0.000
12	0.006	0.000	0.001
15	0.000	0.000	0.021
18	0.005	0.000	0.199
21	0.003	0.000	–
24	0.043	0.000	–

Note: Significance computed for MAE with respect to Elastic Net compared with AR, random walk, and Blue Chip.
Sources: Blue Chip Economic Indicators, Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

statistically significant for most forecast horizons, as shown in Table 3. The p-values in this table indicate a strong, statistically significant difference at forecast horizons of 12 months or fewer.[2]

Forecast turning points

While comparing the errors of different models provides one way to evaluate their predictive performance, the raw errors themselves do not tell the entire story. One hallmark of the unemployment rate is its cyclical movement: the unemployment rate falls in an expansion and rises in a recession. Training a model to capture and forecast this cyclical pattern has advantages beyond statistically lowering errors. Forecasting turning points in the unemployment rate may provide useful information for policymakers.

I identify turning points in the forecast and actual unemployment rate using the following steps: first, I smooth the forecast and actual unemployment rates using a three-month centered moving window. Second, I find the months of highest and lowest unemployment for each business cycle (spanning peak to trough). Third, I calculate the difference in months between the turning point in the actual unemployment rate and the turning point in the forecast unemployment rate.

Table 4 shows that Elastic Net is also able to identify unemployment rate turning points earlier than Blue Chip and other methods. The advantage for Elastic Net is more pronounced at shorter horizons,

Table 4
Forecast Error in Predicting Turning Points (in Months)

Horizon	Random walk	AR	Blue Chip	Elastic Net
3	6.2	4.0	4.2	0.6
6	8.4	6.0	7.0	1.4
9	10.8	8.8	9.2	4.0
12	15.8	10.8	11.8	6.4
15	17.8	12.8	14.0	7.0
18	18.6	13.4	12.8	8.8
21	20.8	14.2	–	7.0
24	19.8	13.8	–	10.0
Average	14.7	10.4	–	5.65

Note: Average distances are calculated from predicted turning points to actual turning points for forecasting methods and variable sets for all horizons.
Sources: Blue Chip Economic Indicators, Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

Table 5
Difference in Turning Point Distances from Elastic Net (in Months)

Horizon	Random walk	AR	Blue Chip
3	5.6	3.4	3.6
6	7.0	4.6	5.6
9	6.8	4.8	5.2
12	9.4	4.4	5.4
15	10.8	5.8	7.0
18	9.8	4.6	4.0
21	13.8	7.2	–
24	9.8	3.8	–
Average	9.0	4.8	5.1

Note: Positive values indicate better performance for Elastic Net.
Sources: Blue Chip Economic Indicators, Bureau of Labor Statistics, Federal Reserve Bank of St. Louis Federal FRED, and author's calculations.

meaning it identifies upcoming turning points in the near term earlier than other methods. On average, Elastic Net identifies turning points 5.1 months earlier than Blue Chip, as shown in Table 5. The advantage in the number of months across horizons appears to be fairly consistent, without strong variation between short and long horizons.

Chart 1

Unemployment Forecasts at the Three-Month Horizon

Panel A: Elastic Net

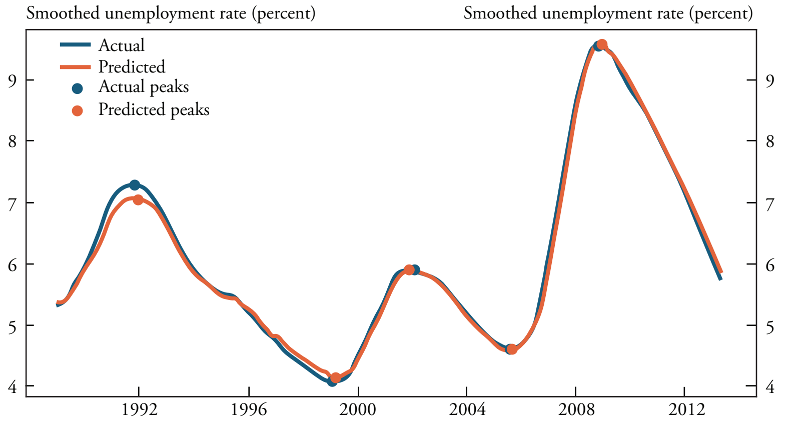

Panel B: Blue Chip Consensus

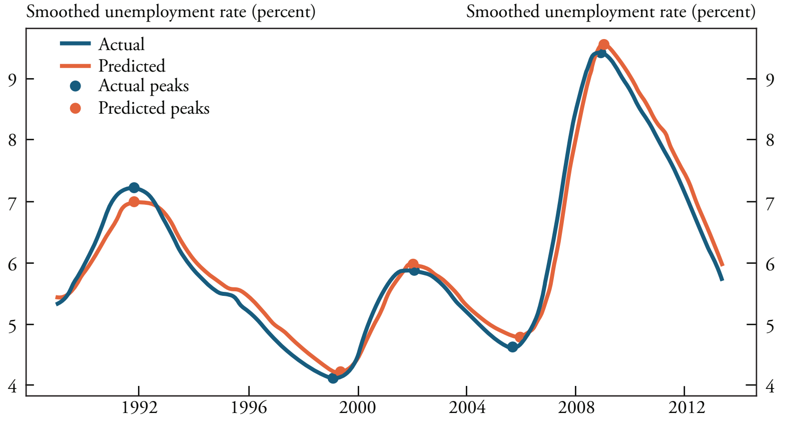

Sources: Blue Chip Economic Indicators, Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

The earlier turning points for Elastic Net can be identified visually at both short and long horizons, as shown in Charts 1 and 2. Panels A and B of Chart 1 shows that at a three-month horizon, the differences from the actual unemployment rate are relatively minor—specifically, the blue (actual) and orange (predicted) lines often overlap. However, the Elastic Net forecast appears to track (overlap) the actual rate track

Chart 2
Unemployment Forecasts at the 12-Month Horizon

Panel A: Elastic Net

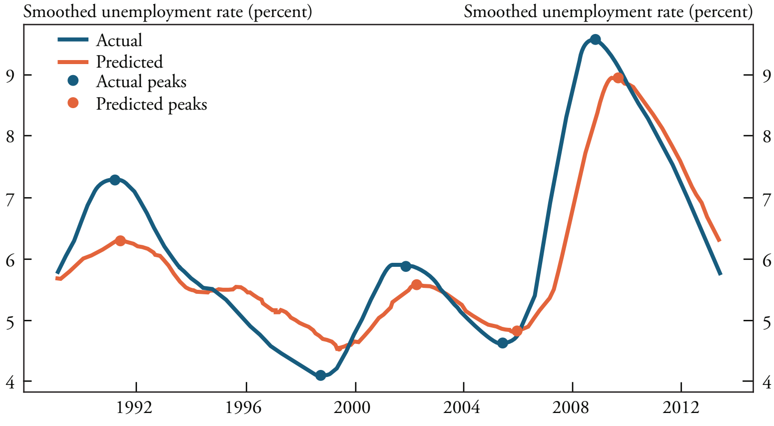

Panel B: Blue Chip Consensus

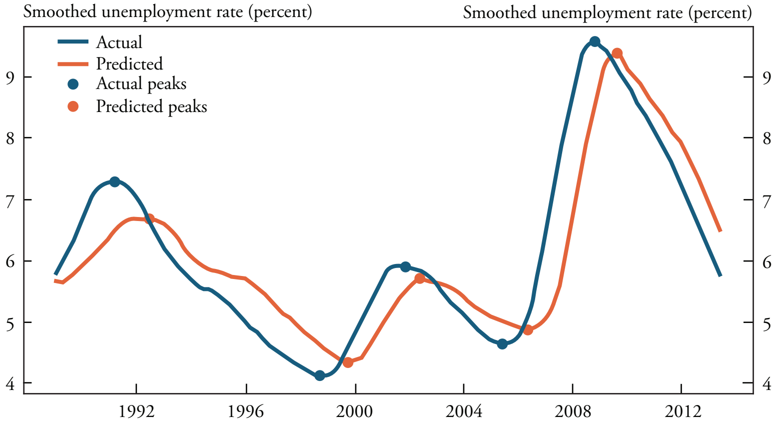

Sources: Blue Chip Economic Indicators, Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

more closely than the Blue Chip forecast. Indeed, the Blue Chip forecast appears to lag the actual rate by a consistent amount, especially during downward trends.

Panels A and B of Chart 2 show that at a 12-month horizon, the forecasts deviate more from the actual rates, as expected. The blue and orange lines in Panels A and B show very little overlap, instead crossing over in the periods between turning points. Again, however, the Blue Chip forecast appears to lag the actual rate by a consistent amount throughout the range of the forecast. The Elastic Net model, on the other hand, appears to respond to turning points earlier to change the forecast trend. These charts provide qualitative support for the differences shown in performance statistics.

Variables identified by Elastic Net

In addition to forecast improvements, Elastic Net can also identify variables potentially important to predicting the unemployment rate. By using regularization to control the size of variable coefficients in the model, Elastic Net can drive some coefficients to exactly zero and thus select a more parsimonious number of variables than an unregularized linear model. Inspecting the model to see how it is making forecasts may uncover new relationships in the data that have been previously overlooked.

The variables identified by Elastic Net cover mostly labor market and housing, with some additional variables covering consumption, output and income, and interest and exchange rates. The coefficients for these variables are shown in Tables 6 and 7, for forecast horizons of three and 12 months, respectively. At a three- and 12-month horizon, variables from housing, manufacturing, and interest rates have the largest coefficients, indicating they have the largest effect on the unemployment rate forecast. The coefficients listed in these tables are averaged over all forecast steps. Only variables with coefficients greater than two standard errors from zero are included.

Panels A and B of Figure 4 show that the strength of these coefficients can vary over time for some variables while staying constant for others. Examining the coefficients allows us to discover how the learned model changes in response to different economic conditions over time. For example, Panel A shows that while many of the coefficient series are stable at the three-month horizon (shown by the solid white bands), the housing and interest rate variables change throughout the forecast period (shown by the shifts in color over time). Prior to 1993–94, the coefficients on the housing and interest rate variables were near zero or only briefly below zero. After 1993–94, however, the coefficients were consistently strong (indicated by the more saturated orange bands).

Table 6
Coefficients for Elastic Net Model for Three-Month Horizon

Variable	Coefficient	Category
Consumer sentiment index	−0.005 (0.001)	Consumption, orders, and inventories
Help-wanted index for United States	−0.025 (0.012)	Labor market
Average duration of unemployment (weeks)	0.032 (0.012)	Labor market
Capacity utilization: manufacturing	−0.043 (0.004)	Output and income
Effective federal funds rate	−0.100 (0.033)	Interest and exchange rates
New private housing permits, Midwest	−0.110 (0.051)	Housing
New private housing permits, Northeast	−0.115 (0.044)	Housing
Three-month Treasury minus federal funds	−0.148 (0.038)	Interest and exchange rates
Average weekly hours: manufacturing	−0.164 (0.020)	Labor market
Civilian unemployment rate	0.874 (0.019)	Labor market

Note: Standard errors are in parentheses.
Sources: Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

Table 7
Coefficients for Elastic Net Model for 12-Month Horizon

Description	Coefficient	Category
Consumer sentiment index	−0.015 (0.006)	Consumption, orders, and inventories
Help-wanted index for United States	−0.036 (0.016)	Labor market
New private housing permits, Northeast (SAAR)	−0.165 (0.082)	Housing
Six-month Treasury minus federal funds	−0.205 (0.065)	Interest and exchange rates
Three-month Treasury minus federal funds	−0.216 (0.060)	Interest and exchange rates
Average weekly hours: manufacturing	−0.237 (0.073)	Labor market
Civilian unemployment rate	0.590 (0.062)	Labor market

Note: Standard errors are in parentheses.
Sources: Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

Figure 4
Coefficients over 1988–2017 for Elastic Net

Panel A: Three-Month Forecast Horizon

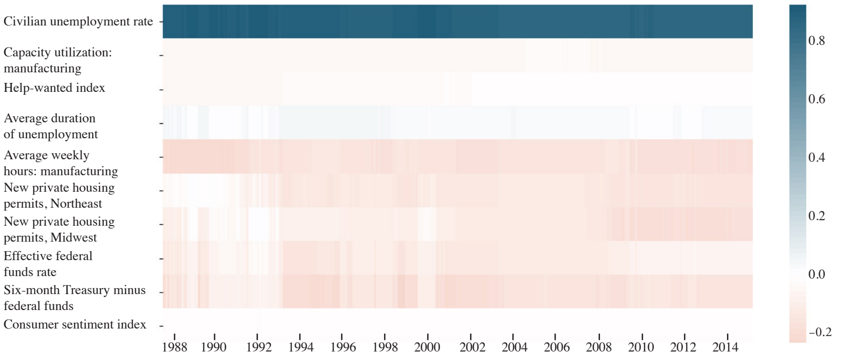

Panel B: 12-Month Forecast Horizon

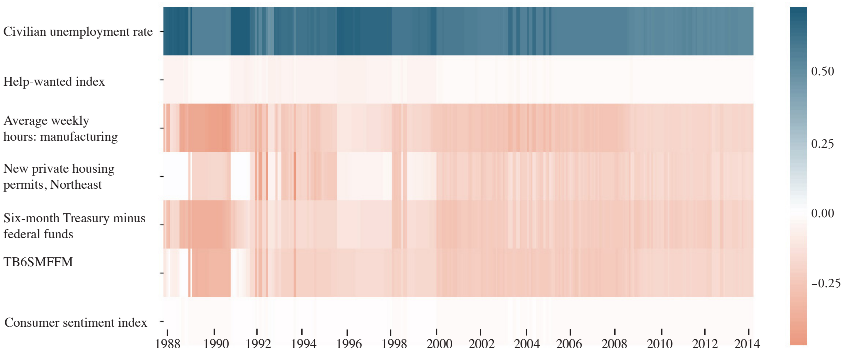

Sources: Bureau of Labor Statistics, Federal Reserve Bank of St. Louis FRED, and author's calculations.

This period of inconsistency broadly centers around the 1990–91 recession, but extends both ahead and behind. The coefficients on these variables diminished near the 2001 recession as well. Interestingly, only the coefficient on the federal funds rate variable diminished after the 2007–09 recession, indicating that this variable has diverged from the others and may no longer be relevant.

Panel B shows several other interesting patterns in variable coefficients for the 12-month forecast horizons. The conspicuous band of color change on the left side of the heat map shows that coefficients on the labor, housing, and interest variables strengthened. This band corresponds to the late-1989 to mid-1991 recessionary period of negative or slow GDP growth, suggesting the model is picking up disruptions in economic relationships associated with the recession. Another pattern of change across variables is seen in the 1996–98 period and again in 1999–2000. These years correspond to a relatively stable expansion before the 2001 recession, suggesting that the influence of manufacturing, housing, and interest rate variables strengthened after the recession. In contrast, the coefficients on these variables weakened after the 2007–09 recession; however, the coefficients on other variables weakened during this period as well, suggesting these changes may be due to broader conditions.

IV. Conclusions

Forecasting is challenging, and the wealth of new and accessible data describing economic conditions presents an opportunity to explore more complex models that can capture more of the economic data. The field of machine learning provides a number of methods to address and capitalize on this complexity, both through increasingly complex models as well as methods to control and optimize that complexity.

I compare the performance of consensus, statistical, and machine learning methods for forecasting the monthly U.S. unemployment rate. My analysis shows that a more complex model, when properly controlled and provided with enough data from which to learn, can significantly outperform consensus and simpler statistical forecasting methods. The key to this result is the control of model complexity through regularization, a machine learning technique that yields a model complex enough to avoid underfitting the data but not so complex as to overfit it.

Appendix

Details of the Elastic Net Model

This appendix discusses the technical details of the Elastic Net model and how it is estimated. For completeness, I build the Elastic Net model piecewise as I add regularization components. The objective function for estimating coefficients $\hat{\boldsymbol{\beta}} = \{\beta_1, \beta_2, \ldots, \beta_m\}$ in an ordinary least squares (OLS) model is given as:

$$\hat{\boldsymbol{\beta}}^{OLS} = \arg\min_{\boldsymbol{\beta}} \Sigma_{i=1}^{n} (\boldsymbol{x_i}\boldsymbol{\beta} - y_i)^2.$$

Here, I omit the intercept without loss of generality, and represent the set of coefficients $\boldsymbol{\beta}$ and independent variables $\boldsymbol{x}$ as vectors with the i subscript indicating the ith of n total samples, with y_i as the response for the ith sample. Next, an L_1 regularization term is added, giving the LASSO objective:

$$\hat{\boldsymbol{\beta}}^{LASSO} = \arg\min_{\boldsymbol{\beta}} \Sigma_{i=1}^{n} (x_i\boldsymbol{\beta} - y_i)^2 + \lambda \|\boldsymbol{\beta}\|_1.$$

Then, an L_2 regularization term is added as in a ridge regression, giving the Elastic Net objective:

$$\hat{\boldsymbol{\beta}}^{EN} = \arg\min_{\boldsymbol{\beta}} \Sigma_{i=1}^{n} (x_i\boldsymbol{\beta} - y_i)^2 + \lambda_1 \|\boldsymbol{\beta}\|_1 + \lambda_2 \|\boldsymbol{\beta}\|_2^2.$$

In this objective, the penalty weights λ_1 and λ_2 are chosen as positive but unbounded, so a re-parameterization is often desirable. The revised objective is:

$$\hat{\boldsymbol{\beta}}^{EN} = \arg\min_{\boldsymbol{\beta}} \Sigma_{i=1}^{n} (x_i\boldsymbol{\beta} - y_i)^2 + \alpha\rho \|\boldsymbol{\beta}\|_1 + \frac{\alpha(1-\rho)}{2} \|\boldsymbol{\beta}\|_2^2,$$

which gives two alternative parameters, α and ρ, where ρ is bounded between 0 and 1 and controls how the regularization strength is balanced between the L_1 and L_2 terms. The parameter α controls the overall strength of regularization and is positive and unbounded. The final objective function is convex and so can be minimized using any convex optimization method such as gradient or coordinate descent.

Endnotes

[1]One challenge with using Blue Chip forecasts for comparison is that some of the historical data used by models in my forecasts have been revised, while the Blue Chip forecasts were made before revised data were available.

[2]Statistical significance is computed using a two-sample Kolmogorov-Smirnov test to compute the probability that errors from two models are drawn from the same distribution.

References

Barnichon, Regis, and Christopher J. Nekarda. 2012. "The Ins and Outs of Forecasting Unemployment: Using Labor Force Flows to Forecast the Labor Market." Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series 2013-19, November.

Clemen, Robert T. 1989 "Combining Forecasts: A Review and Annotated Bibliography." *International Journal of Forecasting*, vol. 5, no. 54, pp. 559–583. Available at https://doi.org/10.1016/0169-2070(89)90012-5

Cook, Thomas R., and Aaron Smalter Hall 2017. "Macroeconomic Indicator Forecasting with Deep Neural Networks." Federal Reserve Bank of Kansas City, Research Working Paper no. 17-11, September. Available at https://doi.org/10.18651/RWP2017-11

Diebold, Francis X. 1997. "The Past, Present, and Future of Macroeconomic Forecasting." NBER Working Paper no. 6290, November. Available at https://doi.org/10.3386/w6290

Lucas, Robert E., Jr. 1976. "Econometric Policy Evaluation: A Critique." *Carnegie Rochester Conference Series on Public Policy*, vol. 1, pp. 19–46. Available at https://doi.org/10.1016/S0167-2231(76)80003-6

McCracken, Michael W., and Serena Ng. 2015. "FRED-MD: Monthly Database for Macroeconomic Research." Federal Reserve Bank of St. Louis, working paper, August. Available at https://doi.org/10.1080/07350015.2015.1086655

Pescatori, Andrea, and Saeed Zaman. 2011. "Macroeconomic Models, Forecasting, and Policymaking." Federal Reserve Bank of Cleveland, *Economic Commentary*, vol. 19, October.

Sims, Christopher A. 1980. "Macroeconomics and Reality." *Econometrica*, vol. 48, no. 1, pp. 1–48. Available at https://doi.org/10.2307/1912017

Ye, Jianming. 1998. "On Measuring and Correcting the Effects of Data Mining and Model Selection." *Journal of the American Statistical Association*, vol. 93, no. 441, pp. 120–131. Available at https://doi.org/10.2307/2669609

Zou, Hui, and Trevor Hastie. 2005. "Regularization and Variable Selection via the Elastic Net." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 2, pp. 301–320. Available at https://doi.org/10.1111/j.1467-9868.2005.00503.x

Economic Review

Fourth Quarter 2018, Volume 103, Number 4

The *Economic Review* (ISSN0161-2387) is published quarterly by the Federal Reserve Bank of Kansas City, 1 Memorial Drive, Kansas City, Missouri 64198-0001. Subscriptions and additional copies are available without charge. Send requests to the Public Affairs Department, Federal Reserve Bank of Kansas City, 1 Memorial Drive, Kansas City, Missouri 64198-0001. Periodical postage paid at Kansas City, Missouri.
POSTMASTER: Send address changes to *Economic Review*, Public Affairs Department, Federal Reserve Bank of Kansas City, 1 Memorial Drive, Kansas City, Missouri 64198-0001.

ECONOMIC REVIEW
Federal Reserve Bank of Kansas City
1 Memorial Drive
Kansas City, Missouri 64198-0001
Fourth Quarter 2018, Vol. 103, No. 4