FEDERAL RESERVE BANK OF KANSAS CITY

ECONOMIC REVIEW

Second Quarter 2016 *Volume 101, Number 2*

FEDERAL RESERVE BANK OF KANSAS CITY

ECONOMIC REVIEW

Second Quarter 2016 *Volume 101, Number 2*

Global Uncertainty and U.S. Exports

By Nicholas Sly

In recent years, demand for U.S. exports has been soft, dragging down U.S. economic growth. Some of the reduced demand for U.S. exports can be attributed to changes in foreign income levels and the value of the dollar. But another, less obvious factor influencing the demand for U.S. goods is uncertainty about foreign growth and financial volatility.

Nicholas Sly constructs a measure of uncertainty for U.S. trading partners and estimates how changes in foreign uncertainty influence foreign demand for U.S. exports. He finds that periods of greater uncertainty and financial volatility are associated with substantially lower demand for U.S. goods. Moreover, he finds that changes in uncertainty and volatility have been relatively more important determinants of U.S. exports in recent years. The results suggest that if foreign growth expectations stabilize—even if they remain relatively weak—U.S. export activity will likely increase.

Consumption Growth Regimes and the Post-Financial Crisis Recovery

By Andrew Foerster and Jason Choi

The financial crisis and recession of 2007–09 hit household balance sheets hard. Even as the economy began to recover, diminished income, a stagnant labor market, and tight credit conditions made it difficult for households to increase their consumption as rapidly as they had a few years earlier. Indeed, consumption has grown more slowly after the Great Recession than in recoveries from previous recessions, suggesting a fundamental shift in the economy.

Andrew Foerster and Jason Choi compare consumption growth's historical behavior with its behavior during the most recent recovery. The authors find that the recent period of slow consumption growth was due not to new or transitory factors but rather the persistent influence of factors unusual to see outside recessions. They find that durables and nondurables consumption behaved much as they did during previous recoveries; total and services consumption, however, grew more slowly than usual throughout the expansion.

Has the Relationship between Bank Size and Profitability Changed?

By Kristen Regehr and Rajdeep Sengupta

In recent years, community bankers and industry analysts have become concerned that small banks may need to grow larger to be successful. New electronic banking platforms—along with new regulations introduced after the 2007–09 financial crisis and recession—have increased fixed costs for all banks, which could place smaller banks at a competitive disadvantage relative to their larger competitors.

Kristen Regehr and Rajdeep Sengupta examine whether the relationship between bank size and profitability has changed since the financial crisis. They find that the relationship has remained stable over time: both before and after the crisis, profitability increased with bank size but at a decreasing rate. Moreover, they find that banks need not grow larger to be successful: in achieving higher profitability, small differences in bank- and market-specific factors are equivalent to large differences in size.

Global Uncertainty and U.S. Exports

By Nicholas Sly

Exports of goods and services account for a substantial share of total U.S. economic activity, with a total value upward of 13 percent of GDP since the year 2000. With so much production, investment, and employment concentrated in the export sector, changes in foreign demand for U.S. goods have important implications for domestic growth. For example, in the early years of the current recovery, exports were a key driver of economic growth; in recent years, however, declining net exports have been a drag on economic growth. Recognizing the importance of the export sector to the U.S. economy, policymakers pay close attention to global factors that influence the demand for goods produced domestically.

In recent years, key factors such as foreign income levels and the value of the dollar have changed dramatically with clear consequences for the demand for U.S. goods. Another, less obvious factor that influences demand for U.S. exports is uncertainty about global growth and related financial volatility. In 2015, economic growth slowed in several emerging markets, with spillovers to their trading partners that are difficult to forecast. The fog does not seem to have cleared much in the beginning of 2016. Movement in oil prices, volatility in equity and bond markets, and changes in monetary policy environments across

Nicholas Sly is a senior economist at the Federal Reserve Bank of Kansas City. William Xu, a research associate at the bank, helped prepare the article. This article is on the bank's website at ***www.KansasCityFed.org.***

countries have all contributed to uncertainty about future economic growth. Regardless of the total size or income of foreign economies, greater uncertainty about their expected growth path may deter resident consumers and firms from ordering goods and components produced in the United States. Likewise, greater certainty about future economic conditions may boost demand for U.S. goods even if foreign incomes and exchange rate levels remain unchanged.

In this article, I estimate how changes in global uncertainty influence foreign demand for U.S. exports. Evidence from 2002 to 2015 across the overwhelming majority of U.S. trading partners suggests periods of greater uncertainty are associated with substantially lower foreign demand for U.S. goods. Specifically, I find that a 1 percentage point increase in the spread between reported high and low foreign GDP growth forecasts, my preferred measure of economic uncertainty, is associated with 2.8 percent lower U.S. export activity on an annualized basis. Volatility in financial conditions within foreign countries, which often portends future volatility in real economic conditions, is also associated with substantially lower demand for U.S. exports, distinct from the role of global uncertainty. The evidence suggests that changes in global uncertainty and financial volatility have been relatively important determinants of U.S. exports in recent years.

Section I presents the empirical model I use to study the role of uncertainty in determining demand for U.S. exports. Section II explains how I measure economic uncertainty across countries and provides information about the trade and income data used in the analysis. Section III presents estimates of both the negative effect of global uncertainty on the demand for U.S. exports and the drag from heightened periods of financial volatility among U.S. trading partners.

I. A Model of U.S. Export Demand

Much like the forces of gravity, the economic forces that determine global trade flows correspond to size and distance, in this case the size of trading nations' economies and the distance between their borders. Just as large physical bodies attract one another, large economies attract substantial trading activity from one another. In addition, faraway nations tend to attract fewer exports from one another: higher shipping costs make the goods of distant countries more expensive than goods of nearby countries. While seemingly simplistic, the gravity model of international commerce has widespread empirical success in explaining cross-border trade flows. An additional benefit of using the gravity model to study international trade flows is that it allows potential determinants of demand other than size and distance—such as measures of global economic uncertainty—to be included.

Furthermore, the gravity model is consistent with theories of consumer behavior and firm production. Several standard models of international trade imply that demand for U.S. exports within another country has a simple (log) linear relationship with the country's national income and the relative prices of goods from the United States compared with other potential exporters.[1]

As a first step toward estimating the demand for U.S. exports, I take the benchmark empirical gravity model given by

$$Exports_{it}^{US} = \alpha + \beta GDP_{it} + \gamma p_{it}^{US} + \varepsilon_{it},$$

where $Exports_{it}^{US}$ denotes purchases of U.S. goods by country i observed in period t, GDP_{it} is the importing nation's GDP at time t, ρ_{it}^{US} captures the relative price difference importer i must pay to purchase U.S. goods at time t, and ε_{it} is variation in importing activity due to other factors not correlated with incomes and relative prices. The term α is a constant capturing the average level of exports observed across countries due to other factors, and β is an estimated parameter reflecting the effect of greater foreign GDP on demand for U.S. goods. Conventional wisdom holds that large economies attract trade from one another and that distance between countries reduces trade. Hence, the estimate of β is expected to be positive, while the estimate of γ is expected to be negative, reflecting that higher relative prices of U.S. goods will lead to lower export activity.

The next step in estimating demand for U.S. exports is incorporating measures of global economic uncertainty about future economic conditions into the model. Leibovici and Waugh provide a simple trade model that accounts for the fact that exporting is time intensive and thus incorporates the potential role of uncertainty about future economic conditions into contemporaneous export decisions.[2] They argue that the current delivery of imports depends on the importers' national income (GDP) from the previous period (when orders for the delivery were made) as well as their willingness to substitute directly between domestic purchases, which can be ordered and delivered immediately, and imported goods that arrive later from the United States.

Several potential factors may affect the trade-off between current and future consumption. As uncertainty about global economic conditions has heightened in the last few years, I focus on uncertainty about future growth expectations as one such potential factor. Regardless of the expected level of future income, risk-averse importers may respond to changes in uncertainty about their own future levels of consumption. Specifically, greater uncertainty about future economic growth reduces the expected benefit of future consumption, making consumers less willing to sacrifice resources today for U.S. exports that will arrive tomorrow. Given the time intensiveness in international trading activity, consumers' responses to more uncertain environments may manifest as lower demand for U.S. exports that will arrive in a later period.

To empirically evaluate this prediction, I build on Leibovici and Waugh by incorporating measures of uncertainty about future GDP growth into the simple gravity equation and estimate the following:

$$Exports_{it}^{US} = \alpha + \gamma Uncertainty_{it-1} + \beta GDP_{it-1} + \gamma p_{it-1}^{US} + \varepsilon_{it}.$$

The time subscripts $t-1$ on the right-hand side explicitly highlight that the delivery of exports at time t is determined by factors that affect demand at the time orders are placed.[3]

II. Measuring Export Activity and Global Economic Uncertainty

The sample used in the empirical analysis is an unbalanced panel covering 26 countries with quarterly observations from each country over the period 2002:Q1 to 2015:Q4. Together, the sample of countries accounts for approximately 85 percent of total U.S. export activity.

A key data requirement is constructing a measure of uncertainty about aggregate economic growth. My approach to measuring global economic uncertainty is to use information derived from a range of forecasts for annual GDP growth reported each month by Consensus Economics. These reports include several independent forecasts of the current and next calendar year's GDP growth across countries. I measure uncertainty about GDP growth as the difference between the highest and lowest forecast reported for each country within each quarter. Given that Consensus Economics reports monthly observations, while trade and income data are available only quarterly, I take the quarter average of the highest and lowest GDP forecasts before calculating the difference. For example, observed uncertainty for Argentina in 2005:Q3 is calculated as 1.33=7.63–6.5, where 7.63 is the three-month average of the highest 2005:Q3 GDP forecast over the quarter, and 6.5 is the three-month average of the lowest 2005:Q3 GDP forecast over the quarter. This measure of economic uncertainty increases as high and low forecasts diverge.

One issue with using spreads between forecasts of annual growth to measure uncertainty is the annualized horizon for each forecast may not correspond exactly to the planning horizon of the foreign firms and consumers ordering U.S. goods. Early in the year, forecasts for annualized growth look ahead several periods and are more likely to reflect the perspective of those demanding exports that will arrive several periods later. Hence, the gap between high and low forecasts of the current calendar year should have a larger effect on the demand for U.S. exports early in the year. In contrast, later in the year, the gap between high and low forecasts of the next calendar year better reflects uncertainty about the future and thus influences demand for exports. To account for these effects, I estimate regression models that allow the role of uncertainty in determining export demand to vary across quarters within a given year.

As many countries have more volatile GDP series on average than others, I include country-specific fixed effects to account for differences in the average level of uncertainty within countries over time. Including country-specific fixed effects is consistent with gravity models of export demand, as it accounts for average differences in relative prices between countries due to the costs of transporting goods across fixed distances. While theoretically consistent with gravity models, and empirically justified by differences in average levels of uncertainty, including country-specific fixed effects implies that identifying the effect of uncertainty relies on variation in the spread between high and low forecasts within specific countries over time.

Global financial conditions are often a harbinger of future economic conditions. As a result, volatility in financial conditions within U.S. trading partners may reflect an alternative source of global uncertainty affecting demand for U.S. goods. To investigate this alternative channel, I calculate the standard deviation of interday yields of 10-year sovereign bonds within each quarter for each country. Greater variation in day-to-day bond yields within a quarter indicates greater financial volatility, which may then reflect greater uncertainty about real economic conditions within each country. Unlike other common measures of financial volatility, such as the VIX, the prices of government debt issuances are available for a wide set of countries (although this measure includes fewer countries than the previous measure).[4] These data are taken from Bloomberg.

Chart 1 plots the time series of measured uncertainty and volatility for each country in the sample. Panel A illustrates variation in uncertainty and bond prices for G-7 export destinations, which make up the bulk of U.S. export activity, while Panel B plots uncertainty and financial volatility for all other countries in the sample. Although the panels display clear differences in the average levels of economic uncertainty across countries, no significant trends within countries are apparent over time, alleviating concerns about spurious trends driving results.

The U.S. Bureau of Economic Analysis reports quarterly exports to each county in millions of U.S. dollars, with separate series reported for exports in goods only and exports of both goods and services.[5] My preferred measure of export activity is trade in goods, both because more countries are available in the sample, and because uncertainty has a more ambiguous effect on services trade due to the variable time it takes to deliver specific services to foreign consumers.[6] I present estimates using detrended series of quarterly exports, taken from a Hodrick-Prescott (HP) filter that accounts for secular growth in trade flows over time.

The final data requirements are national income levels and exchange rates, which independently affect export demand. I take quarterly GDP levels in billions of seasonally adjusted U.S. dollars from Haver Analytics. I report results obtained using detrended GDP series taken from an HP filter, which correspond to the measure of exports used throughout the analysis. I take exchange rate data from *The Wall Street Journal* and report values for the number of local currency units per U.S. dollar. Hence, higher values of the variable *ForEx* reflect a higher cost to purchase U.S. goods. Summary statistics for the sample are reported in Table 1.

Table 1
Summary Statistics

Variable	Mean	Number of observations	Standard deviation	Maximum	Minimum
Detrended ln(Exports)	0	1,375	0.094	0.347	–0.511
Uncertainty (current year)	1.385	1,375	0.971	8.567	0.133
Financial volatility	0.166	1,138	0.135	2.181	0.013
ln(ForEx)	0.055	1,375	0.196	0.000	1.381
Detrended ln(GDP)	0	1,375	0.016	0.053	–0.098

Chart 1
Uncertainty and Financial Volatility over Time

Panel A: G-7 Countries

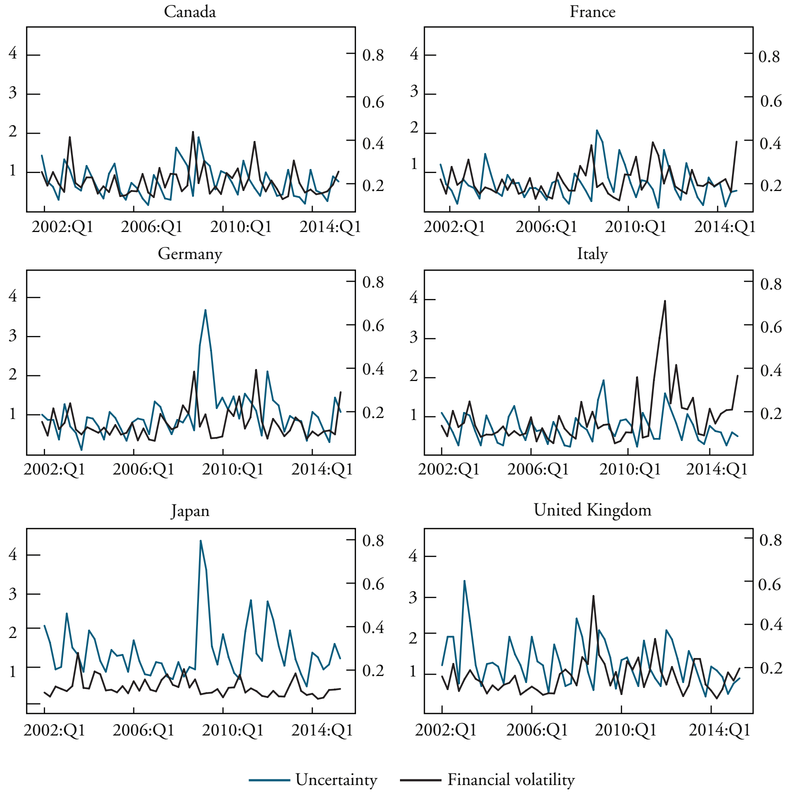

Notes: Left scale measures uncertainty in GDP growth for each country in percentage point differences in forecasts observed each quarter. Right scale measures financial volatility as the standard deviation in interday bond prices (yields) over each quarter.
Sources: Bloomberg, Consensus Economics, and author's calculations.

Panel B: Non-G-7 Countries

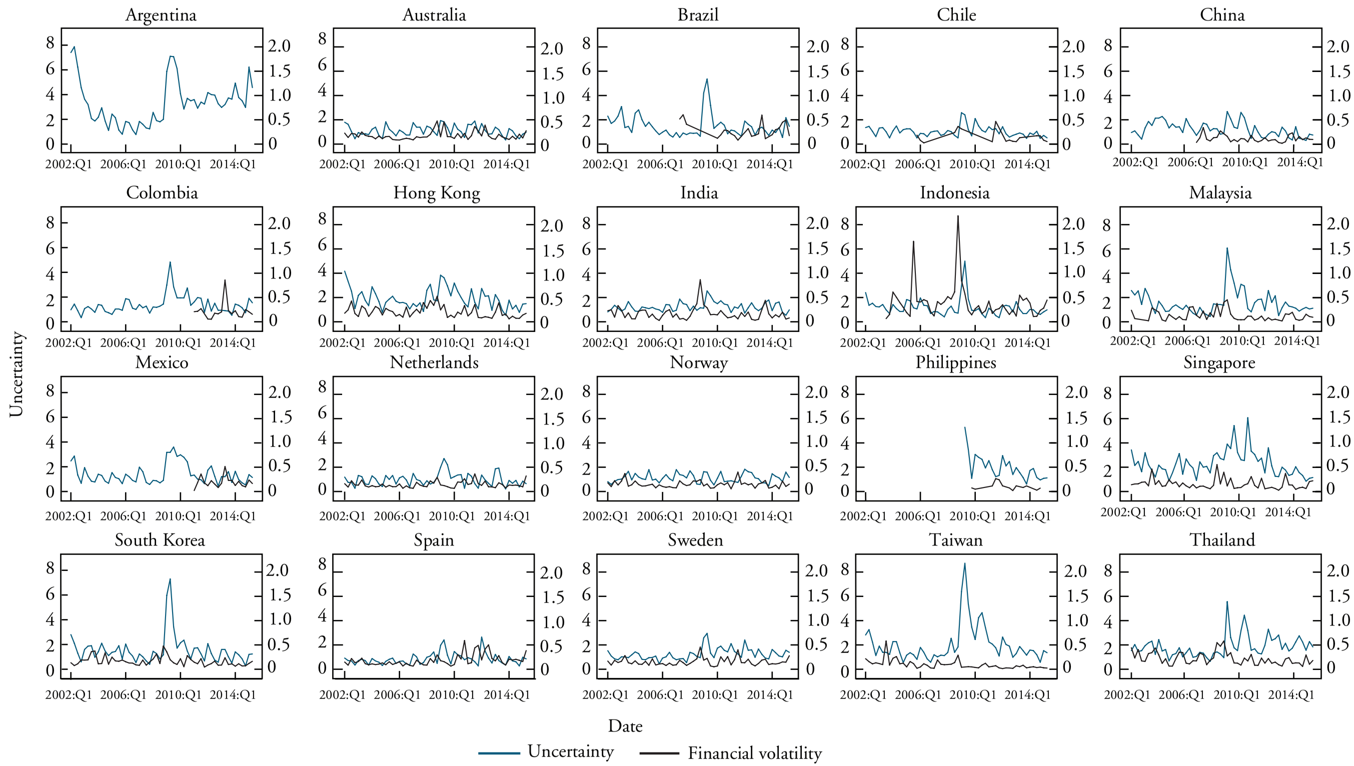

Notes: Left scale measures uncertainty in GDP growth for each country in percentage point differences in forecasts observed each quarter. Right scale measures financial volatility as the standard deviation in interday bond prices (yields) over each quarter.

Sources: Bloomberg, Consensus Economics, and author's calculations.

III. The Relationship between Uncertainty and Export Activity

Before turning to the results of the regression exercise, I examine the relationship between uncertainty and exports without accounting for other factors. Chart 2 plots the preferred measure of uncertainty, the gap between high and low forecasts of GDP in the current quarter, against detrended quarterly exports of U.S. goods to each country. Even without considering any other potential determinants of foreign demand for U.S. goods, the negative correlation illustrated in Chart 2 indicates that higher levels of uncertainty about global growth are indeed associated with lower demand for U.S. exports.

While the pattern in Chart 2 is clear and in line with expectations, the simple negative correlation fails to account for the timing of export activity as well as differences in income and the relative cost to deliver U.S. goods across countries. As these are known to be important determinants of demand for U.S. exports, I turn next to the regression analysis, which takes such factors into account.

The effect of foreign economic uncertainty on U.S. exports

Looking across several models of the demand for U.S. goods, I find that heightened uncertainty about growth in foreign countries exerts a substantial drag on U.S. export activity. Table 2 reports results from the baseline gravity specification of export demand that includes lagged measures of GDP, relative prices as measured by the value of the dollar relative to countries' local currencies, and measures of uncertainty about future economic growth. Column 1 reports estimates from a regression of U.S. exports to each country on measured uncertainty, (lagged) GDP level in logs, the (lagged) foreign exchange value of the dollar, and country-specific fixed effects. In line with expectations, the coefficient on measured uncertainty about future aggregate growth indicates that higher economic uncertainty within the economies of U.S. trading partners is a drag on U.S. export demand. Specifically, the coefficient on uncertainty of –0.020 implies that a 1 percentage point increase in the gap between the highest and lowest forecasts results in approximately a 2 percent reduction in demand for U.S. goods. Put simply, the effect of uncertainty on import demand appears substantial in economic magnitude and is significant at high degrees of statistical confidence.

Chart 2

Correlation between U.S. Exports and Foreign Growth Uncertainty across Countries and Time

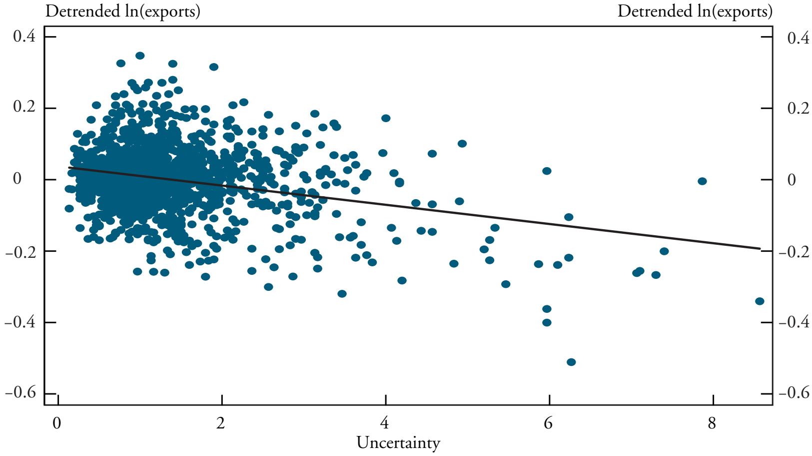

Note: Each dot represents the value of U.S. exports to a specific country at a particular quarter in the sample. These values of U.S. exports are plotted against measured uncertainty about the respective foreign country's GDP growth. The line illustrates a fitted linear regression across the whole sample.
Sources: U.S. Bureau of Economic Analysis, Consensus Economics, and author's calculations.

As in prior analyses of export activity, higher levels of national income (GDP) are associated with a higher demand for U.S. goods. Given that the model is estimated in logs, the coefficient on ln(GDP) can be interpreted as the observed income elasticity of demand. Hence, the point estimate of 2.761 on GDP in column 1 implies that a 1 percent increase in aggregate income results in a 2.761 percent increase in demand for imports from the United States, although these estimates fail to take into account time-specific effects across years or quarters. Not surprisingly, increases in the value of the dollar relative to local currencies reduce the demand for U.S. exports.

The results in column 1 use the preferred measure of uncertainty, which considers spreads between forecasts for current year GDP growth among importers of U.S. goods. However, as the year progresses and new data become available, the typical spread between forecasts will naturally fall. If trade flows also exhibit systematic variation within a year, spurious correlations may contaminate the estimates in column 1. Moreover, common global factors that vary year to year may also affect each country's individual demand for U.S. exports. To account for such issues, column 2 introduces quarter- and year-fixed effects into the analysis.

Table 2

Effect of Foreign Economic Uncertainty on Demand for U.S. Exports

Variables	(1) ln(Exports)	(2) ln(Exports)	(3) ln(Exports)	(4) ln(Exports)
Current-year uncertainty	–0.020*** (0.003)	–0.016*** (0.003)	–0.024*** (0.007)	–0.031*** (0.008)
Current-year uncertainty×Q2			0.010 (0.006)	0.016* (0.009)
Current-year uncertainty×Q3			0.011 (0.010)	0.021** (0.010)
Current-year uncertainty×Q4			0.018 (0.012)	0.049*** (0.017)
Next-year global uncertainty				0.006 (0.006)
Next-year global uncertainty×Q2				–0.008 (0.007)
Next-year global uncertainty×Q3				–0.013 (0.009)
Next-year global uncertainty×Q4				–0.035*** (0.010)
Lagged ln(GDP)	2.761*** (0.210)	1.517*** (0.239)	1.529*** (0.239)	1.454*** (0.226)
Lagged ln(ForEx)	–0.090** (0.035)	–0.047 (0.028)	–0.046 (0.029)	–0.041 (0.032)
Constant	0.244*** (0.085)	0.151* (0.076)	0.163* (0.081)	0.152* (0.088)
Observations	1,349	1,349	1,349	1,349
R^2	0.362	0.449	0.452	0.464
Country fixed effects	Yes	Yes	Yes	Yes
Quarter fixed effects	No	Yes	Yes	Yes
Year fixed effects	No	Yes	Yes	Yes

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.
Notes: Standard errors are in parentheses.

In column 2, the point estimate on uncertainty, –0.016, is only slightly different from the –0.020 point estimate obtained in column 1. However, accounting for year- and quarter-specific factors appears important to the estimates of the role of income fluctuations. Including time-specific effects, I find the coefficient on GDP growth, at 1.5, is more in line with standard estimates and remains both economically and statistically significant.

One issue with measuring uncertainty using spreads in current year growth forecasts is that Consensus Economics does not report growth forecasts for the same horizon across quarters, nor do these horizons correspond exactly to the time it takes to deliver exported goods. To account for these facts, the specifications in columns 3 and 4 investigate how the effect of uncertainty on export activity evolves within a calendar year. Measured uncertainty—the gap between high and low forecasts of the current calendar year—is expected to have a relatively larger effect on export activity within the first few months of a year; in contrast, in later months of a year, the gap between high and low forecasts of the next calendar year should have a larger influence on export activity.

The preferred specification in column 4 includes the forecasts for both the current and next calendar years as well as estimates of their differential effects across quarters. Consistent with expectations, the coefficient of –0.031 on uncertainty indicates higher uncertainty about a country's economic growth in the first quarter of the year reduces U.S. exports to that country by approximately 3 percent. In line with expectations, the positive coefficient on uncertainty for the second quarter of the year, 0.016, suggests that the drag on U.S. exports is smaller, though the difference is only marginally statistically significant. The even larger positive estimate of 0.021 for the third quarter suggests that uncertainty about the current year's growth exerts even less drag on U.S. exports. A statistical test confirms the estimated negative net effect of uncertainty on U.S. exports in the third quarter remains statistically significant.

By the fourth quarter, however, uncertainty about the current year's economic growth is no longer a drag on a country's demand for U.S. exports. The estimated effect of uncertainty about the current year within the fourth quarter (–0.31+0.5=0.2) is statistically indistinguishable from zero. Instead, uncertainty about economic growth in the next calendar year affects decisions to purchase U.S. exports. The coefficient on uncertainty about next year's growth in the fourth quarter is approximately –0.035, which is statistically significant at high degrees of confidence. On an annualized basis, the estimates in column 4 imply a 1 percentage point increase in the spread between reported high and low foreign GDP growth forecasts is associated with 2.8 percent lower demand for U.S. exports. In line with expectations, I find that uncertainty about next year's growth has no statistically discernable effect on demand for U.S. exports in the first three quarters of each year.

Table 3

The Effect of Global Uncertainty and Financial Volatility on U.S. Exports

Variables	(1) ln(Exports)	(2) ln(Exports)	(3) ln(Exports)	(4) ln(Exports)
Lagged financial volatility	–0.078*** (0.012)	–0.071*** (0.010)	–0.064*** (0.012)	–0.059*** (0.013)
Current-year uncertainty	–0.020*** (0.004)	–0.016*** (0.004)	–0.027*** (0.008)	–0.031*** (0.008)
Current-year uncertainty×Q2			0.013 (0.008)	0.015 (0.010)
Current-year uncertainty×Q3			0.020** (0.009)	0.023** (0.010)
Current-year uncertainty×Q4			0.026* (0.014)	0.057*** (0.019)
Next-year uncertainty				0.003 (0.008)
Next-year uncertainty×Q2				–0.004 (0.008)
Next-year uncertainty×Q3				–0.006 (0.007)
Next-year uncertainty×Q4				–0.033*** (0.011)
Lagged ln(GDP)	2.631*** (0.212)	1.548*** (0.269)	1.562*** (0.275)	1.494*** (0.258)
Lagged ln(ForEx)	–0.120*** (0.019)	–0.068*** (0.024)	–0.069*** (0.023)	–0.070*** (0.025)
Constant	0.292*** (0.043)	0.198*** (0.057)	0.217*** (0.056)	0.221*** (0.063)
Observations	1,114	1,114	1,114	1,114
R^2	0.357	0.435	0.440	0.453
Country fixed effects	Yes	Yes	Yes	Yes
Quarter fixed effects	No	Yes	Yes	Yes
Year fixed effects	No	Yes	Yes	Yes

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.
Notes: Standard errors are in parentheses.

The effect of foreign financial volatility on U.S. exports

Volatility in financial markets may also inhibit foreign consumers from ordering exports of U.S. goods, as real economic strife often follows bouts of financial stress. Table 3 shows the additional influence that variation in nations' sovereign bond prices, a measure of financial volatility, may have on their demand for exports of U.S. goods. The specifications in Table 3 are identical to those reported in Table 2, with the addition of measures of financial volatility. As data on financial availability are not available for a small number of countries, the number of observations differs between Tables 2 and 3.

The estimates in column 1 reveal that uncertainty stemming from foreign financial volatility is also a significant drag on U.S. export activity. The estimated coefficient of –0.078 implies that a one standard deviation increase in financial volatility within a foreign country decreases its demand for U.S. exports by nearly one standard deviation. Put simply, the effect of financial volatility is economically substantial. The effect of financial volatility is stable across specifications, dipping only to 0.059 in column 4—though this difference is not statistically different from the point estimate in column 1.

The effect of foreign financial volatility appears largely independent of the effect of uncertainty in economic growth forecasts. Both are significant at high degrees of confidence, and the estimates on uncertainty are unchanged when financial volatility measures are included. This suggests that financial volatility and uncertainty about growth within our trading partners represent distinct risks to U.S. export demand.

The estimates that include measures of foreign financial volatility continue to show a correlation between higher foreign GDP and higher demand for U.S. goods. In addition, higher values of the U.S. dollar relative to foreign currencies appear to deter foreign demand for U.S. exports. The coefficient of 0.078 in column 1 indicates that a 1 percent increase in the value of the dollar against foreign currencies reduces demand for U.S. exports by 0.07 percent on average.

IV. Factors Affecting Demand for U.S. Exports in Recent Years

The evidence in the previous section confirms that foreign economic and financial phenomena affect demand for U.S. exports in addition to conventional factors such as exchange rates and global GDP growth. More precisely, the evidence in the last section demonstrates that such factors tend to affect demand for U.S. exports on average. In this section I take a closer look to see which factors have been most important in explaining recent fluctuations in U.S. exports.

Chart 3 illustrates the estimated contributions of changes in global uncertainty, international financial volatility, foreign income levels, and exchange rates to changes in U.S. exports over recent time horizons. Panel A shows these contributions over the last decade, while Panel B focuses on the last five years. I calculate the contributions using four-quarter moving averages of each factor and their respective estimated effects from column 4 of Table 3.

The primary determinants of export demand have varied over time. The dark blue bars in Panel A, which chart foreign GDP growth, show that from 2006 to 2008, foreign growth spurred increases in demand for U.S. goods. Then, at the onset of the global financial crisis and subsequent recession in 2007–09, reductions in foreign GDP growth lowered demand for U.S. exports. The light blue bars, which represent financial volatility, show that in the first quarter of 2009, the global financial crisis also pulled down U.S. export activity. And the black bars, which represent uncertainty, show that while some of the global financial stress abated in the second quarter of 2009, uncertainty about foreign economic conditions kept demand for U.S. exports low. However, all three bars climbed during the early parts of the recovery in late 2009 and early 2010 as increases in foreign incomes combined with decreased financial volatility and foreign economic uncertainty to boost demand for U.S. exports.

Over the last few years, the value of the dollar and foreign uncertainty have played more prominent roles in U.S. export activity. Panel B of Chart 3 shows relatively large contributions from foreign uncertainty (black bars) and the dollar (gray bars) compared with foreign GDP (dark blue bars) in influencing foreign demand for U.S. goods in 2013. During the last half of 2015, the dollar's rapid rise markedly increased U.S. exports, leading to substantial drag on export demand from 2015:Q2–Q4. The recent episode of heightened global financial volatility associated with China's devaluation of the yuan late in 2015:Q3 also appears to have lowered demand for U.S. goods at year-end. The contrast between Panels A and B suggests global uncertainty and foreign financial volatility have had a larger effect on U.S. exports in recent years, primarily because other factors, particularly foreign GDP growth, have been less volatile.

V. Conclusion

Sluggish export activity has been a drag on U.S. growth recently. In addition to slowing foreign growth and a high relative value of the dollar, uncertainty in the foreign growth outlook has caused demand

Chart 3

Foreign Factors Affecting Demand for U.S. Exports over Time

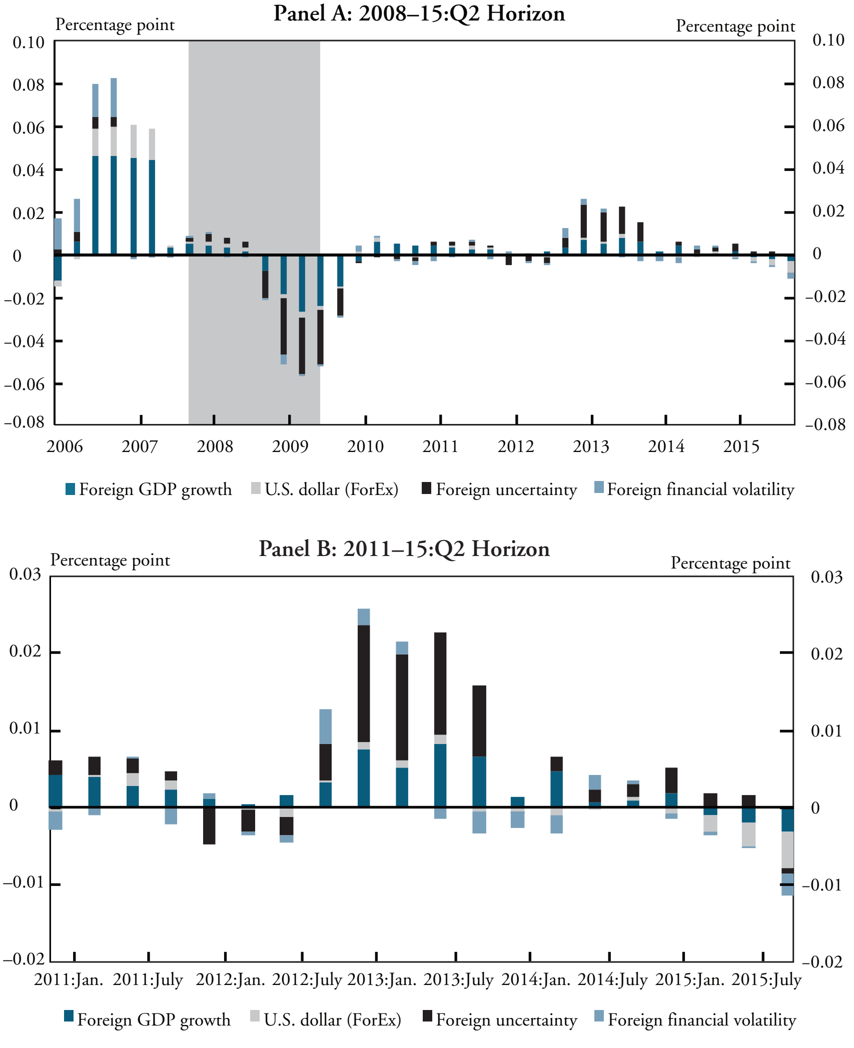

Notes: Gray shaded region denotes NBER-defined recession. The bars show the expected contributions of each variable to changes in U.S. export demand around its long-run trend.
Source: author's calculations.

for U.S. exports to wane. Stress in foreign financial conditions has further contributed to the declining demand for U.S. goods. The propensity for uncertainty to diminish orders of U.S. goods suggests that export activity would likely pick up if foreign growth expectations were to stabilize, even if the expectations for growth remain relatively weak. In the first few months of 2016, global uncertainty and financial volatility surged, potentially dampening export demand. As these forces abate, more certainty and stability in foreign economic and financial conditions will likely contribute to U.S. export growth.

Endnotes

[1]See, for example, Krugman; Anderson and van Wincoop; Eaton and Kortum; and Melitz.

[2]See Hummels and Schaur for evidence about the time intensiveness of international trading activity.

[3]An alternative specification of the gravity model would include the expected level of future GDP to account for the time intensiveness of international trading activity. I investigate this option and find that the role of measured uncertainty remains qualitatively robust.

[4]One concern is that variation in sovereign security prices is driven by trend movements within a quarter, which would spuriously measure high financial market volatility for relatively stable price movements along a trend path. I investigate measures of financial market volatility that use interday price changes to account for such concerns and find similar results.

[5]These data series are reported without seasonal adjustment. To concord with other data used in the analysis, I seasonally adjust the reported series.

[6]Regardless, I show that the results are quantitatively robust using either measure of trade.

References

Anderson, James E., and Eric van Wincoop. 2003. "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, vol. 93, no. 1, pp. 170–192.

Bloomberg L.P. 2002–15. Historical 10-year sovereign bond end-of-day day yields based on bid prices in local currencies for select countries, January 2, 2002 –December 31, 2015. *Bloomberg Finance L.P,* accessed March 30, 2016.

Consensus Economics Inc. 2002–15a. *Consensus Forecasts,* January 2002–December 2015.

_____. 2002–15b. *Asia Pacific Consensus Forecasts*, January 2002–December 2015.

_____. 2002–15c. *Latin American Consensus Forecasts*, January 2002–December 2015.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade," *Econometrica*, vol. 70, no. 5, pp. 1741–1779.

Hummels, David L., and Georg Schaur. 2013. "Time as a Trade Barrier," *American Economic Review*, vol. 103, no. 7, pp. 2935–2959.

Krugman, Paul. 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, vol. 70, no. 5, pp. 950–959.

Leibovici, Fernando, and Michael E. Waugh. 2014. "International Trade and Intertemporal Substitution," NBER working paper no. 20498, September.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, vol. 71, no. 6, pp. 1695–1725.

Consumption Growth Regimes and the Post-Financial Crisis Recovery

By Andrew Foerster and Jason Choi

The financial crisis and recession of 2007–09 hit household balance sheets hard and resulted in large numbers of job losses. Diminished wealth and income, high unemployment, and a stagnant labor market—in combination with tight borrowing and lending conditions—made it difficult for households to increase consumption as rapidly as they had just a few years earlier. After the Great Recession, consumption has grown more slowly than in the recoveries from previous recessions, suggesting a fundamental shift in the economy.

Consumption growth reflects a variety of both persistent and transitory factors. Shifts in underlying factors such as labor markets or financial conditions can persistently change the speed and volatility of consumption growth; other determinants of consumption such as weather or temporary tax changes can have transitory effects. Characterizing consumption growth during the recovery as being due to either persistent or transitory factors can help determine exactly how the recovery differed from previous ones. If the factors driving consumption growth are fundamentally different now from the past, previous recoveries may no longer indicate how the economy might rebound from recessions. But if the factors driving consumption growth are not too different, previous recessions still may provide insight into how consumption growth may evolve.

Andrew Foerster is a senior economist at the Federal Reserve Bank of Kansas City. Jason Choi is a research associate at the bank. This article is on the bank's website at ***www.KansasCityFed.org.***

In this article, we compare consumption growth's historical behavior with its behavior during the recovery from the Great Recession. We conclude that the slow growth was due not to a shift to previously unseen behavior, but rather the continued influence of persistent factors that are unusual to see outside recessions. While durables and nondurables consumption behaved much as they did during previous recoveries, both total and services consumption saw an atypical continuation of recessionary behavior during the recovery. If the recessionary behavior had not continued, the United States would have had higher total and services consumption throughout the expansion. Section I presents a graphical analysis of consumption growth after recessions. Section II presents a statistical model demonstrating that growth of total consumption and its components did not behave differently during the recovery but merely returned to previously seen behavior. Section III uses the statistical model to highlight that while factors driving consumption growth in the previous recovery mimicked those in history, their behavior in the cases of total and services consumption was unusual for periods immediately after recessions.

I. A Graphical Perspective on Consumption Growth after Recessions

One method for gauging the speed of recoveries involves graphing and comparing normalized consumption series across several recessions. This method indicates the relative speeds of consumption growth and its components after recessions and suggests that the slow recovery in total consumption from the 2007 recession was primarily due to sluggish services growth.

A graphical perspective of consumption growth

Graphs comparing the paths of total consumption across different recessions can illustrate, by historical standards, how normal or abnormal growth was after the 2007 recession. Chart 1 displays real consumption series for the last four business cycles.[1] Panel A shows these series normalized to equal 1 in the quarter of the previous expansion's peak, which represents the start of each recession. Panel B shows a similar graph with the series normalized to equal 1 in the quarter of each recession's trough, which represents the end of each recession.

Chart 1

Consumption over the Business Cycle

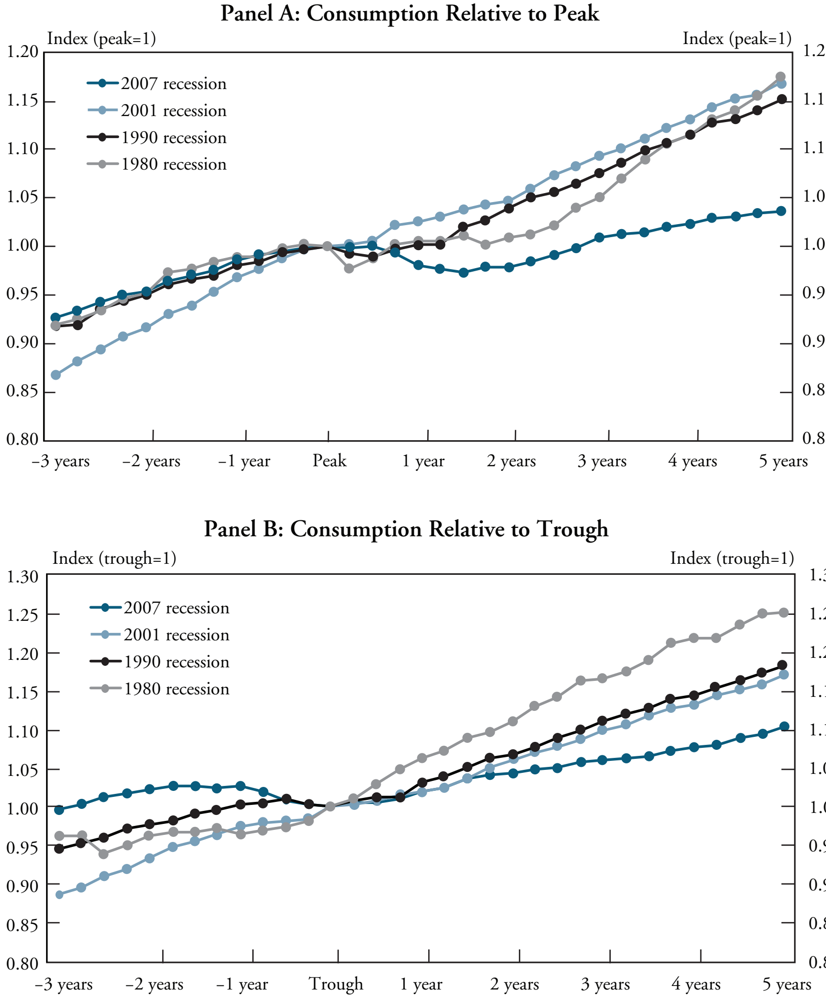

Sources: Haver Analytics, Bureau of Economic Analysis and authors' calculations.

Normalizing the series at two different points in the business cycle allows the panels to show different perspectives on how consumption behaved during and after the recessions. In both panels, the slope of the lines indicates the growth rates over each business cycle, with steeper slopes indicating faster growth rates. The level of the line accumulates these growth rates to show the relative size of consumption.

Regardless of the normalization, the recovery from the 2007 recession looks quite weak compared with previous recessions. Panel A (consumption normalized across peaks) shows that the relative level of consumption after the 2007 recession was well below its relative level after previous recessions. In addition, consumption continued to decline for more than a year after the recession's peak; in previous recessions, consumption resumed its upward climb almost immediately. The sharp and persistent drop in consumption after the peak reflects that the 2007 recession was much deeper and drawn out than previous recessions. Panel B (consumption normalized across troughs) reinforces this conclusion. Although consumption grew steadily after the 2007 recession's trough, its growth clearly lagged behind that of previous recoveries; two years after the trough, growth was especially slow. This graphical analysis highlights that the 2007 recession had not only a longer-lasting and larger decline than previous recessions, but also languishing growth in the years that followed.

Although normalized charts effectively illustrate the decline in consumption growth after the 2007 recession, they are less useful in explaining the source of this decline. Furthermore, analyses based on total consumption may mask differences among the subcomponents of consumption, which differ in how they behave during and after recessions.

A graphical perspective of consumption growth by component

Consumption is divided into three components: durables, nondurables, and services. Durables account for approximately 12 percent of total consumption and include items that are typically purchased infrequently such as vehicles, furnishings, and household appliances. Nondurables make up around 22 percent of consumption and include more regularly purchased items such as restaurant meals, clothing, and gasoline. Services make up the remaining 66 percent of consumption and include expenditures on housing, utilities, health care, and financial services.[2]

The three components of consumption—durables, nondurables, and services—behave differently over the business cycle. Similar to Chart 1, Charts 2, 3, and 4 compare these components of total consumption relative to the peaks and troughs of past business cycles.

Chart 2

Durables Consumption over the Business Cycle

Panel A: Durables Consumption Relative to Peak

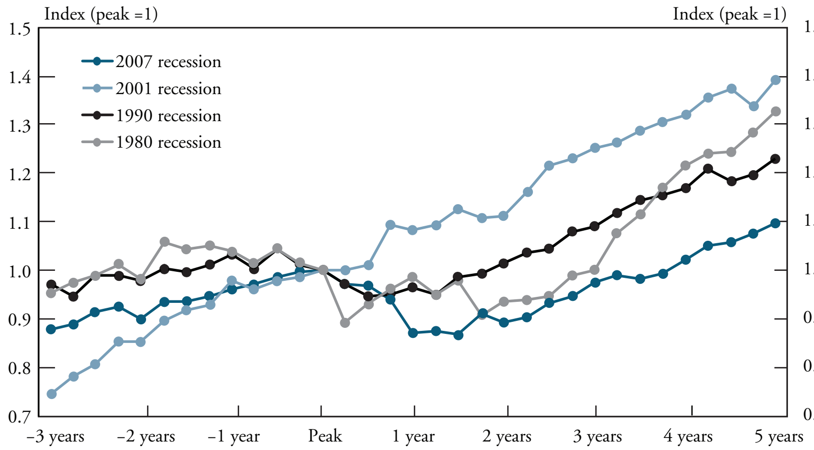

Panel B: Durables Consumption Relative to Trough

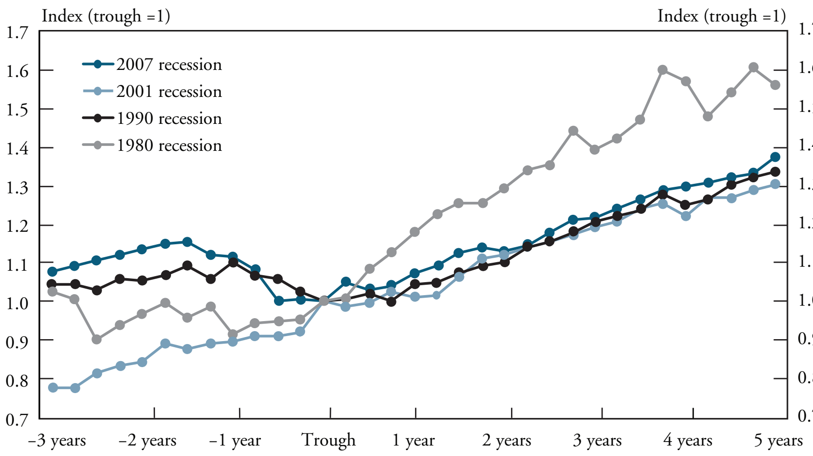

Sources: Haver Analytics, Bureau of Economic Analysis and authors' calculations.

Panel A of Chart 2 plots durables consumption across past business cycles relative to their peaks. Under this normalization, the recovery after the 2007 recession looks slow by historical standards. However, in Panel B, which plots consumption relative to business cycles' troughs, the recovery looks normal. The overall trends are largely consistent with those shown for total consumption in Chart 1. However, durables fluctuate more than total consumption. These sharp movements imply significant volatility in the growth rates of durables consumption. Both the prolonged downturn and the higher volatility of durables may result from consumers delaying purchases of durable goods when economic conditions warrant caution. When financial conditions are tight, consumers may find it necessary to postpone purchases of cars or large household items rather than cut back on small, recurring household purchases.

As with durables, nondurables consumption behaved somewhat differently than total consumption during and after the 2007 recession. Panels A and B of Chart 3 show the normalized paths for nondurables consumption. Relative to the business cycle's peak, nondurable consumption growth looks weak in the most recent recovery, although this weakness partially reflects the longer-than-usual decline in nondurables consumption during the recession. Relative to the business cycle's trough, nondurables consumption looks in line with previous recoveries for the first two years; thereafter, growth stalls and nondurables flatten for about two years before resuming their previous growth. Consequently, five years after the 2007 business cycle's trough, the level of nondurables consumption lags its level in previous recoveries. However, this may be due to a temporary slowdown in growth two to four years after the recession ended rather than persistently weak growth throughout the entire recovery. The volatility of nondurables consumption growth again appears to be higher than for total consumption, although not quite as high as durables consumption. Postponing purchases of many of the goods making up nondurables, such as clothing, may be relatively difficult, leading to somewhat less volatility than durables consumption.

Services consumption, as plotted in Chart 4, behaved quite differently after the 2007 recession than after previous recessions. When compared across peaks and troughs, the growth of services consumption looks weak relative to previous recoveries. In fact, in the first year after the business cycle's peak, the path of services consumption mimics those in previous episodes, implying that differences among the recoveries, not the recessions, are primarily responsible for the change. In addition, services consumption contrasts with both durables and nondurables in that the paths look extremely

Chart 3
Nondurables Consumption over the Business Cycle

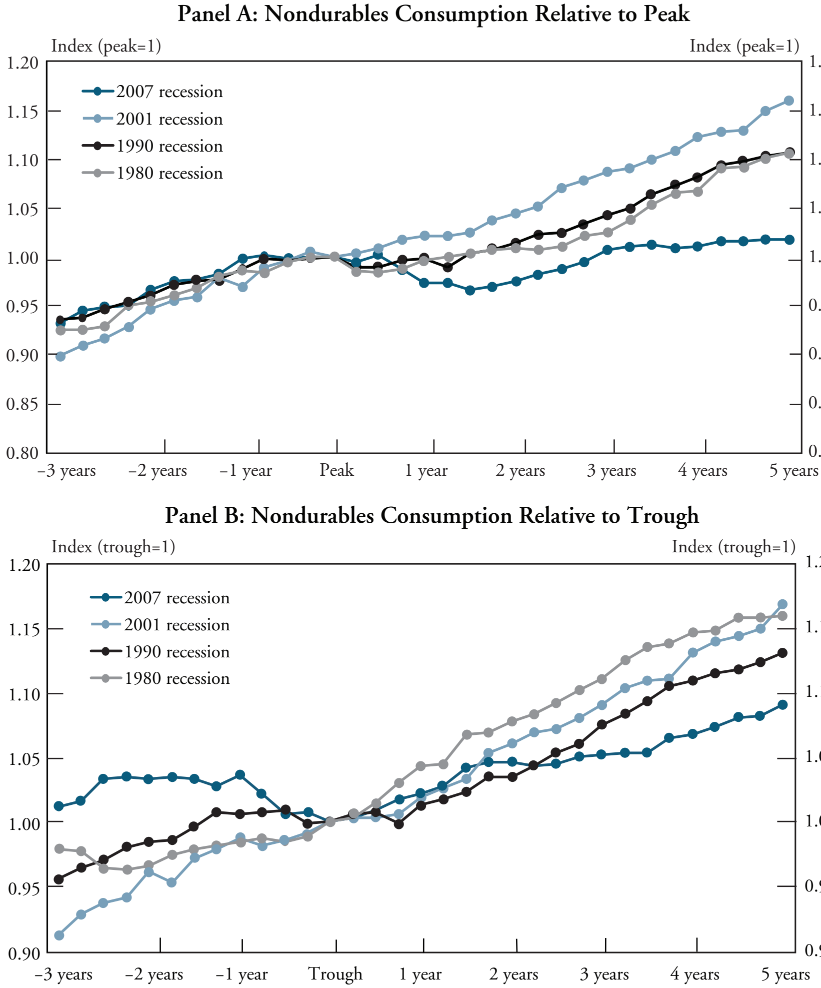

Sources: Haver Analytics, Bureau of Economic Analysis and authors' calculations.

smooth, implying much less variation in growth rates. Consumers essentially avoid postponing services such as health care or housing, leading to very stable growth in services consumption.

The components of consumption paint a more nuanced picture of consumption during recoveries than the total series does. While total consumption growth after the 2007 recession looks weaker than in previous recoveries, its components suggest tepid growth in services, which

Chart 4

Services Consumption over the Business Cycle

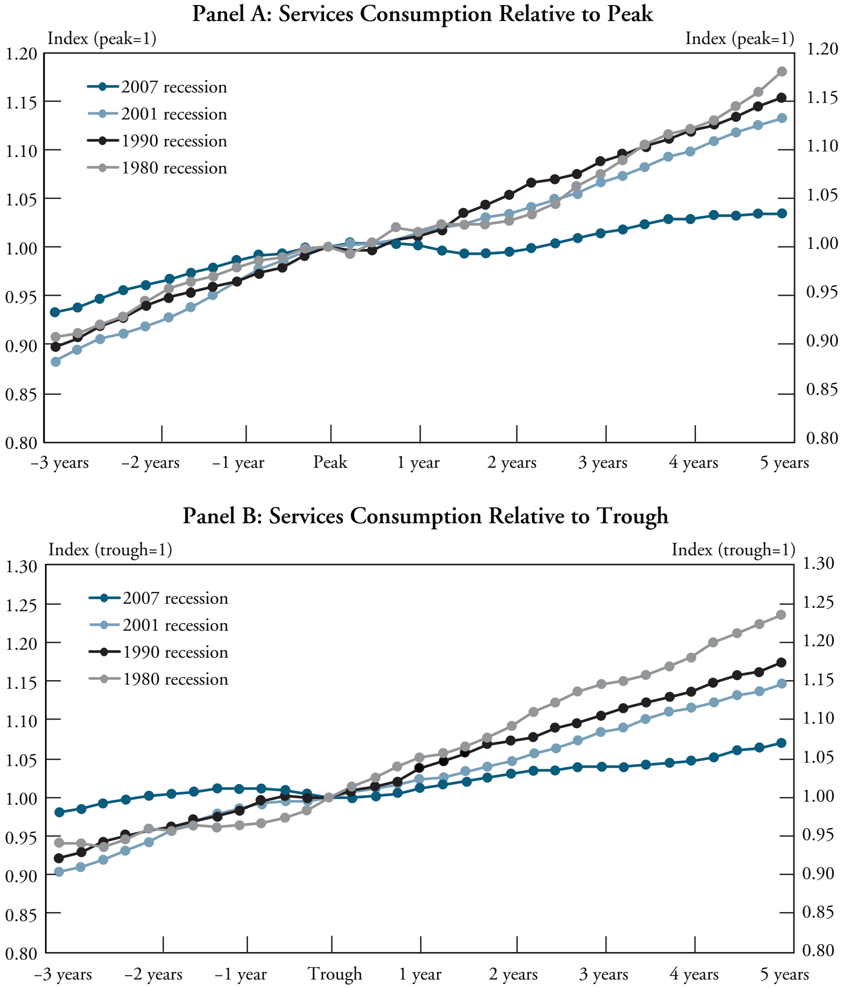

Sources: Haver Analytics, Bureau of Economic Analysis and authors' calculations.

accounts for around two-thirds of total consumption, is primarily to blame. However, this graphical analysis is limited in certain respects. First, by focusing exclusively on performance relative to peaks or troughs, the plots may obscure longer-term trends. Second, the graphical analysis does not clearly indicate whether growth in consumption and its components reflects transitory or persistent factors, nor does it explain how these factors may have differed from those underlying previous recessions. To account for these limitations, we turn to a statistical model that does not rely on comparisons across peaks and troughs to help decompose consumption growth into persistent versus transitory elements.

II. Consumption Growth Regimes

Graphical analysis indicates that total consumption growth during the recovery from the 2007 recession was low relative to historical standards, primarily due to slow growth in services. To evaluate whether the 2007 recession fundamentally changed the behavior of consumption growth, we construct a statistical model for both total consumption growth and its three components that allows growth to evolve differently across time and to depend on both persistent and transitory factors.

To allow for persistent shifts in the behavior of consumption and to distinguish those shifts from transitory movements, our statistical model allows consumption growth to depend on different "regimes" that dictate the average level or volatility of growth. These regimes capture distinct shifts in the behavior of consumption which can happen suddenly—such as during the onset or end of a recession—instead of gradually. By modeling consumption as dependent on separate growth and volatility regimes, we can capture a wide range of possible factors that may affect the average level or volatility of growth independently rather than together. Changes in fiscal policy, for example, might alter the average growth rate of consumption while leaving its volatility unchanged; likewise, foreign shocks might lead to large fluctuations in consumption without changing its average growth rate. Separating average growth and volatility regimes allows us to capture a more nuanced view of consumption growth at various points in history.

Estimates from the model suggest consumption and its components do not grow in a stable manner over time. Total consumption growth and its services component oscillate between two regimes with different average growth rates; both total consumption and services stayed in the low average growth regime during the most recent recovery. In contrast, both durables and nondurables consumption have only one regime with a constant average growth rate, implying they did not deviate from their historical behavior. Total, durables, and nondurables consumption all exhibit high- and low-volatility regimes; after briefly entering the high-volatility regime during the recession, all three components of consumption returned to the low-volatility regime. In contrast, services consumption has only one volatility regime over time.

A statistical model for consumption growth

To assess whether persistent or temporary factors drove consumption growth in the most recent recession, we use a statistical model known as a Markov-switching model, introduced by Hamilton, which allows us to relate the quarterly growth rate of total consumption or one of its three components to a level and a volatility term. This model allows the level and volatility to vary over time between regimes. The model is as follows:

$$\Delta C_t = \mu(S_t) + \sigma(V_t)\,\varepsilon_t.$$

The variable ΔC_t represents the quarterly percent change in either total consumption or one of its components—durables, nondurables, or services. The variable $\mu(S_t)$ denotes the average level of consumption growth, which varies according to the regime S_t. The variable $\sigma(V_t)$ denotes the volatility of consumption growth, which varies according to the regime V_t. The shock ε_t accounts for differences in consumption growth from the average level and is scaled by the volatility term.[3]

The variables S_t and V_t allow the average level of consumption growth and its volatility to take one of several values at each point in time—that is, they tie consumption growth and volatility to a value dictated by the regime. These regimes offer a reduced-form way to capture a variety of factors such as wage growth, financial conditions, household expectations, or policy that affect consumption growth. Decomposing growth into these different regimes helps assess whether persistent factors—indicated by shifts in the average level of growth, $\mu(S_t)$—or temporary factors—indicated by the size of the composite error term, $\sigma(V_t)\varepsilon_t$—played a larger role. For example, in explaining low consumption growth in a given quarter, if the statistical model shows a low average level of growth with a small error term, then persistent factors may be at play, and growth may be low in the future. On the other hand, if the model shows a high average level of growth and a low realization of the error term, then more transitory factors are to blame, and growth should be higher in the future.

The number of regimes used to characterize average growth and volatility can dramatically alter the conclusions. To determine the correct number of regimes, we review specifications with one, two, or three regimes for each of S_t and V_t and and pick the version of the model that best fits the data. For example, if the model allows average growth, S_t, to have three regimes, then the economy can have three different average growth rates: low, medium, and high. If the model allows S_t to have two regimes, then the economy can have two growth rates: low or high. And if the model allows S_t to have only one regime—low for simplicity—then the average growth rate remains unchanged. The regimes available for volatility, V_t, are similar.

After determining the preferred specification and analyzing how similar or different the regimes are from one another, the model can attribute them to different time periods and identify how consumption growth behaved in those periods. This method allows for a previously unseen regime—in other words, one that was not in place during previous recoveries—to arise and dictate consumption growth during the current recovery, which would suggest markedly new consumption behavior. But the method also allows for the regime in place during the recovery to simply be a repeat of a previous regime, which would suggest consumption growth similar to previous recoveries.

The evolution of the regimes follows a Markov process, which implies that the regime in a given period depends probabilistically on the regime in the previous period. For example, the probability P^{μ}_{ll} denotes the probability that the economy will be in the low average growth regime in one period if it was in the low average growth regime the previous period. Similarly, the probability P^{μ}_{lh} denotes the probability of switching from the low average growth regime in one period to the high average growth regime in the next. Corresponding probabilities exist for the volatility variable, $\sigma(V_t)$, and for each possible combination of regimes. The regimes for average growth and volatility evolve according to independent processes (see Kim and Nelson, McConnell and Perez-Quiros, and Lettau and others).

Consumption growth regimes

To characterize the behavior of total consumption growth, we first need to determine the number of regimes that avereage growth and volatility can enter. This determination matters significantly for consumption dynamics, since one regime for the average growth rate implies that consumption growth behaves stably, whereas three regimes implies that consumption switches between high, medium, and low average growth regimes. Likewise, multiple volatility regimes might signal that consumption growth switches between a low-volatility regime, in which it fluctuates only slightly from its average level, and a high-volatility regime, in which it fluctuates significantly from its average level.

We compare different specifications for the number of regimes using goodness-of-fit statistics and find consumption growth is best characterized using two level and two volatility regimes. The values of the Schwarz-Bayesian Information Criterion (SBIC) measure, shown in the first column of Table 1, assess how well the model with various numbers of regimes fits the data while penalizing the inclusion of additional regimes. Lower values of the SBIC imply more favorable model specifications, with the lowest value achieved in the "2 Average, 2 Volatility" specification.[4] As a result, we use this specification in the model, which suggests consumption growth tends to fluctuate between low and high average growth regimes and low- and high-volatility regimes. The economy switches between four regimes in total, with the following combinations of the average growth and volatility regimes: low average growth and low volatility, low average growth and high volatility, high average growth and low volatility, and high average growth and high volatility. As the average growth and volatility regimes are independent from one another, the probabilities of switching between each of these four regime combinations simply depend on the probabilities of switching between each part of the combination. For example, the probability of staying in the low average growth and low-volatility regime is given by $P^{\mu}_{ll}P^{\sigma}_{ll}$, the probability of switching to the low average growth and high-volatility regime is $P^{\mu}_{ll}P^{\sigma}_{lh}$, and so on.

Table 2 shows that the estimated values for this model differ greatly in the low and high average growth regimes as well as the low- and high-volatility regimes. In the low average growth regime, quarterly total consumption growth averages (0.45 percent) around 1.81 percent at an annualized rate. In the high average growth regime, quarterly total consumption growth (1.03 percent) more than doubles, averaging 4.17 percent at an annualized rate. Similarly, the high-volatility regime is more than twice as volatile as the low-volatility regime.

Table 1
Schwarz-Bayesian Information Criterion (SBIC) for the Markov-Switching Models

	SBIC			
Number and type of regime	Total consumption	Durables consumption	Nondurables consumption	Services consumption
1 average, 1 volatility	553.94	1,304.58	566.35	354.17
1 average, 2 volatility	537.66	1,279.23*	558.36*	359.48
1 average, 3 volatility	544.30	1,295.70	578.99	382.96
2 average, 1 volatility	538.65	1,307.40	565.30	305.35*
2 average, 2 volatility	530.10*	1,289.86	565.48	308.60
2 average, 3 volatility	533.91	1,289.06	570.61	323.44
3 average, 1 volatility	546.24	1,307.63	584.30	310.44
3 average, 2 volatility	537.82	1,302.56	579.21	312.09
3 average, 3 volatility	543.73	1,303.07	581.62	315.42

* Denotes the preferred model with the lowest SBIC

The estimates for the probabilities show that both the average growth and volatility regimes are likely to persist for several quarters. The low average growth regime has an expected duration of around eight quarters, while the high average growth regime has an expected duration of around 15 quarters.[5] The volatility regimes last even longer, with expected durations for the low- and high-volatility regimes of around 36 quarters and 64 quarters, respectively.

Consumption growth regimes by component

To gain a more nuanced view of the number and behavior of regimes, we repeat the statistical analysis on the individual components of consumption. We use similar goodness-of-fit statistics to pick the preferred number of regimes for each component. The preferred number of regimes varies across components, and the estimated coefficients indicate the components' behavior differs across regimes.

Table 2

Average Growth Rates and Volatilities across Consumption Regimes

Panel A: Average Growth Rates

Average growth regime	Consumption	Durables	Nondurables	Services
Low	0.45	1.61	0.61	0.32
	(6.18)	(12.08)	(15.21)	(5.87)
High	1.03	--	--	1.03
	(20.14)			(32.03)
Prob. low -> low	0.88	--	--	0.86
	(2.59)			(3.25)
Prob. high -> high	0.93	--	--	0.95
	(3.78)			(6.07)

Panel B: Volatilities

Volatility regime	Consumption	Durables	Nondurables	Services
Low	0.31	1.00	0.43	0.37
	(8.10)	(7.20)	(8.47)	(16.14)
High	0.75	3.77	0.95	--
	(18.34)	(23.41)	(11.00)	
Prob. low -> low	0.97	0.84	0.89	--
	(3.84)	(3.52)	(3.38)	
Prob. high -> high	0.98	0.94	0.88	--
	(5.11)	(7.20)	(2.92)	

Note: t-statistics in parentheses.

The goodness-of-fit statistics for the models of durables, nondurables, and services consumption imply the preferred models differ from that for total consumption (Table 1). For example, the preferred model specification for durables and nondurables consumption has a single average growth regime and two volatility regimes. As a result, unlike total consumption, which switches between high and low average growth regimes, durables and nondurables consumption do not switch average growth regimes over time. However, these components do switch between low- and high-volatility regimes, suggesting that growth in durables and nondurables fluctuates more in some periods than in others. In contrast, the preferred model for services consumption has one volatility regime and two average growth regimes. These regimes imply that services consumption experiences periods of both high and low average growth but that its volatility does not change around those average levels.

The estimated coefficients for the preferred durables and nondurables models show that both their average growth rates and volatility differ from total consumption. The single average quarterly growth regime estimates for the durables and nondurables model (1.61 percent and 0.61 percent, respectively) are 6.6 percent and 2.5 percent at an annualized rate (Table 2). The volatility estimates show standard deviations in the high-volatility regime are much higher than in the low-volatility regime—by more than a factor of two for nondurables and by nearly a factor of four for durables. The particularly large deviations in durables growth may reflect that consumers can more easily postpone purchases of durables than nondurables, leading to more variation in durables growth over time. The estimates of the probabilities show growth for both durables and nondurables has a lower chance of remaining in the low- or high-volatility regimes than total consumption does. For durables, the low-volatility regime has an expected duration of slightly over six quarters, while the high-volatility regime has an expected duration of nearly 16 quarters. For nondurables, both volatility regimes have an expected duration of between eight quarters and nine quarters. These estimates suggest durables and nondurables switch more frequently between volatility regimes than total consumption.

The estimates for the services consumption model mimic the average growth results for total consumption. In the low average growth regime, quarterly services consumption (0.32 percent) grows around 1.31 percent at an annualized rate. In the high average growth regime, services consumption (1.03 percent) grows around three times faster, at approximately 4.17 percent at an annualized rate. The probabilities show that the high average growth regime tends to persist longer than the low average growth regime; their expected durations are around 21 quarters and seven quarters, respectively. In addition, in the single volatility regime, the standard deviation for services is lower than in either volatility regime for durables and nondurables, highlighting that consumption growth for services tends to be much more stable than for other components.

III. Consumption Growth Regimes across History

The previous section modeled consumption growth under different regimes and demonstrated that growth behaves differently across them. In this section, we apply the estimates from the statistical model to historical fluctuations in consumption to infer which regimes held during which times. Our results show that total and services consumption often enter the low average growth regime during recessions and switch to the high average growth regime during recoveries; however, in the recovery after the most recent financial crisis, total and services consumption remained in the low average growth regime. A counterfactual exercise suggests that if consumption had returned to the high average growth regime during the recovery, total and services consumption would have followed more traditional paths.

Historical decomposition of consumption regimes

Although total consumption grew slowly by historical standards after the most recent crisis, the slow growth resulted not from a fundamental change in consumption behavior—in other words, not from the appearance of a new, third regime—but from an extension of the low average growth regime not typically seen during recoveries. Since the regime in place is unobserved, the statistical model places relative probabilities on being in each regime at given points in time. For each quarter in the sample, the model takes into account all of the data and places a probability—called the smoothed probability—on whether the economy was more likely to be in the low or high average growth or volatility regime at any point in time.[6]

The smoothed probabilities for total consumption growth regimes show an unusual combination of regimes during the recovery after the 2007 recession. Chart 5 shows the smoothed probabilities for total consumption growth: Panel A shows the probability the model places on consumption being in the low average growth regime, while Panel B shows the probability of being in the low-volatility regime. Panel A conditions on being in the low average growth regime but allows either volatility regime to be in place. Likewise, Panel B conditions on being in the low-volatility regime but allows either average growth regime to be in place. The blue line in Panel A spikes upward in most of the shaded regions, suggesting consumption was more likely to be in the

Chart 5
Consumption Growth and Volatility Regimes

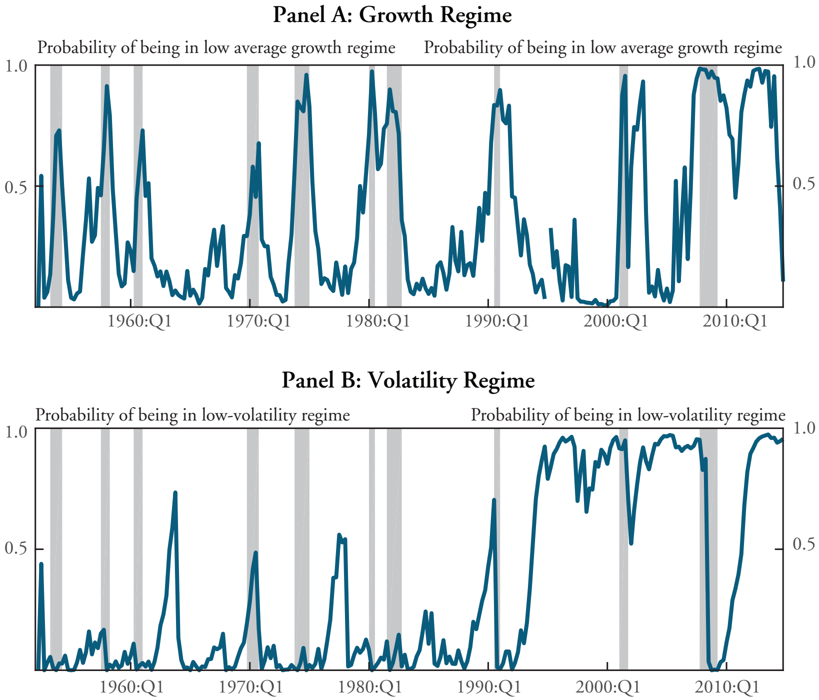

Note: Gray bars denote NBER-defined recessions.

low average growth regime during recessions, when consumption tends to fall or exhibit very weak growth. Panel B suggests the low-volatility regime became much more likely starting in the early 1990s. However, the most recent recession and recovery saw important departures from these first two trends. In particular, the high-volatility regime reappeared during the downturn and persisted well into the recovery. The low average growth regime also dominated for most of the recovery, in contrast to the usual quick shift after recessions back to the high average growth regime. As a consequence, total consumption growth tended to have a lower average and a higher volatility during the post-financial crisis recovery than typical recoveries.

Historical decomposition of consumption regimes by component

While total consumption had two possible average growth regimes during the crisis, both durables and nondurables had only one; in addition, the smoothed probabilities for the two volatility regimes show the high-volatility regime tended to hold during recessions. Charts 6 and

7 depict the smoothed probabilities for durables and nondurables consumption. Since the preferred model has only one average growth regime, the plots show only the probability of being in the low-volatility regime. Unlike total consumption, the implications of the probabilities for durables and nondurables are not as clear-cut. The probabilities do not readily match with recessions or recoveries, nor do they show clear shifts at any one time. However, both durables and nondurables entered the high-volatility regime during the recent financial crisis and recession and switched to the low-volatility regime during the current recovery. This result implies that durables and nondurables consumption growth tended to be relatively close to their average values during the recovery after exhibiting much bigger swings during the crisis.

The probabilities for services consumption are similar to those for total consumption. Chart 8 displays the probabilities of services consumption being in the low average growth regime, since the preferred model has two average growth regimes but only one volatility regime. Similar to the results for total consumption, the low average growth regime for services tends to hold during recessions. Again, in contrast to typical recoveries, the low average growth regime dominated in the post-financial crisis recovery. While past recessions saw rapid shifts back to the high average growth regime, the low average growth regime persisted after the 2007 financial crisis.

Counterfactual regimes in the post-financial-crisis recovery

The statistical model places relatively high probability on the economy being in the low average growth regime for total and services consumption even after the most recent recession ended. A counterfactual exercise can assess what the path of consumption would have looked like during this time if the economy had switched to the high average growth regime, as was typical after most recessions.

The model plays a central role in developing this alternative scenario. Since the model attributes consumption growth to changes in the average level or to transitory deviations from that average level, a scenario that considers different average levels of consumption growth allows us to identify the effects of potentially persistent changes in the growth rate. As our preferred model for durables and nondurables con-

Chart 6

Durables Consumption Volatility Regime

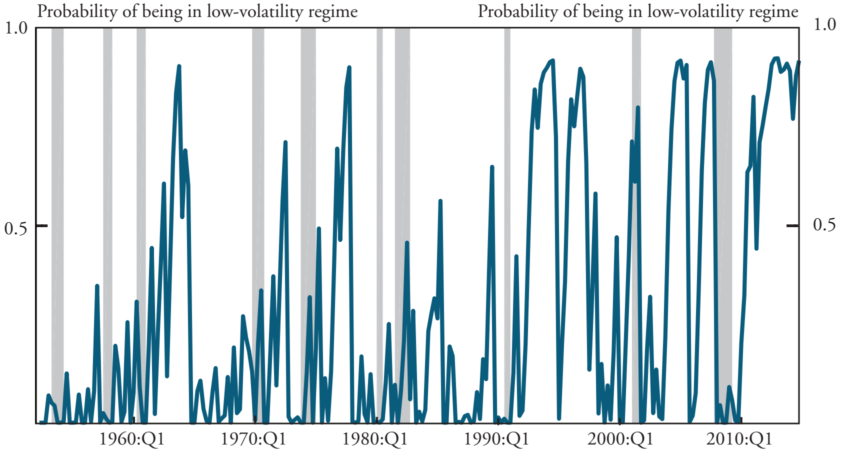

Note: Gray bars denote NBER-defined recessions.

Chart 7

Nondurables Consumption Volatility Regime

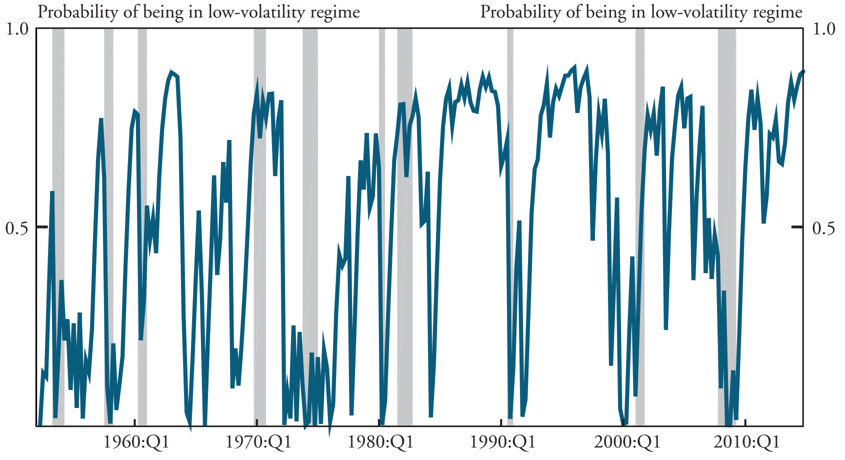

Note: Gray bars denote NBER-defined recessions.

Chart 8
Services Consumption Growth Regime

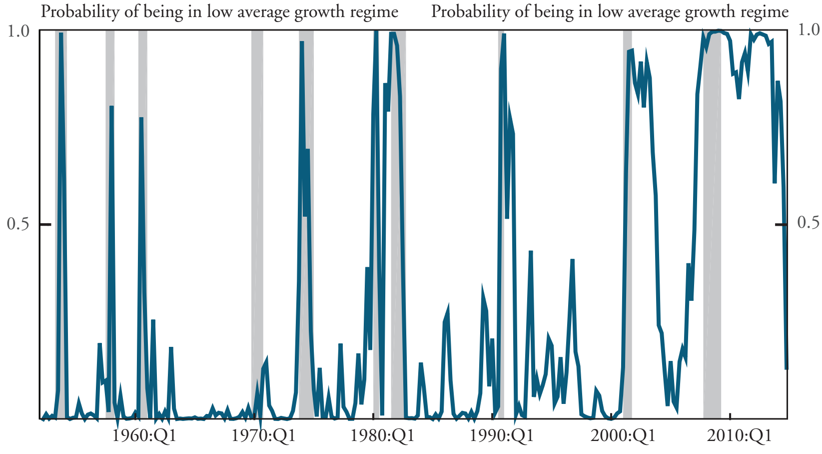

Note: Gray bars denote NBER-defined recessions.

sumption growth has only one growth regime, any deviations from the average growth rate are completely transitory. Our preferred models for total and services consumption, however, have two average growth regimes, allowing deviations from the average growth rate to persist. In this way, a counterfactual that changes the average growth rate regime implicitly considers an alternate history that changes factors that persistently alter consumption dynamics—such as financial constraints, fiscal policy, or changes in productivity—rather than those that only transitorily alter consumption growth, such as weather.

The counterfactual series shows that total and services consumption would have been much higher if the economy had switched back to the high average growth regime after the trough. Chart 9 shows the actual and counterfactual series for total and services consumption for the latest business cycle, normalized to equal 1 in the second quarter of 2009 (the recession's trough). The counterfactual shows a more rapid increase in total and services consumption relative to the actual series. By the middle of 2014, five years after the end of the recession, actual total consumption was only 10.4 percent higher than it was when the recession ended; in contrast, in the counterfactual series, total consumption is around 21.5 percent higher. Similarly, actual services consumption was only 7.1 percent higher by the middle of 2014 than

Chart 9

Counterfactual Consumption after the Great Recession

Panel A: Consumption Relative to 2007 Recession Trough

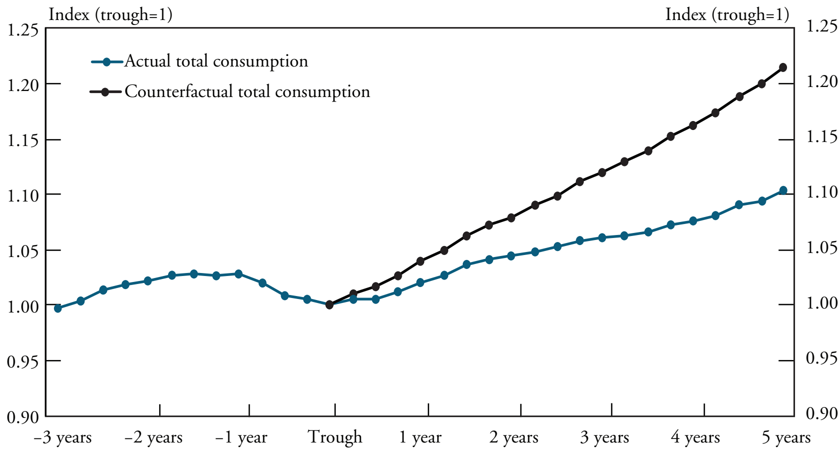

Panel B: Services Consumption Relative to 2007 Recession Trough

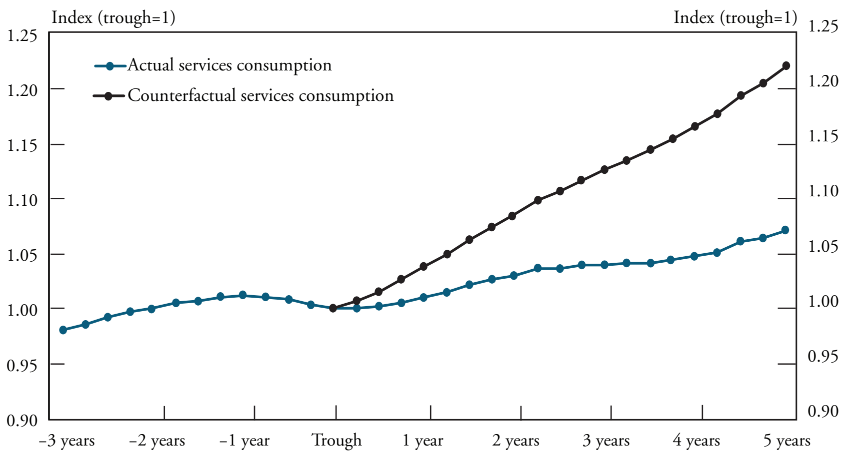

Sources: Haver Analytics, Bureau of Economic Analysis and authors' calculations.

it was at the end of the recession; in the counterfactual series, services consumption was 21.9 percent higher.

These counterfactual exercises help quantify the effect of persistent, rather than transitory, factors on total and services consumption. They suggest that if the recovery after the financial crisis had been ruled by the high average growth regime, by the second quarter of 2014, total consumption would have been about $1.1 trillion higher and services consumption would have been about $1 trillion higher.

IV. Conclusion

In this article, we use a regime-switching model to show that slow growth in total and services consumption after the Great Recession was due not to transitory factors or fundamental changes in consumption behavior but to the unusual persistence of a low average growth regime during the recovery. Low average growth regimes typically reflect the influence of persistent factors such as a slow labor market recovery, restrictive financial conditions, or weak productivity growth. Thus, policies that eliminated or alleviated these headwinds might have helped strengthen growth, leading to significantly higher consumption in the recovery.

However, one caveat to our results is that by modeling consumption growth with a regime-switching framework, we have only captured consumption dynamics in a reduced form. Although our analysis shows consumption growth remained in the low average growth regime after the crisis, it does not explain why. An analysis of the economic fundamentals driving growth would better indicate why growth after the recession was slow and which policies could have affected it. In addition, this analysis has focused on consumption growth in isolation, and ignored possibly important spillover effects into other parts of the economy, such as investment and labor markets, which may further affect consumption growth.

Endnotes

[1]This graphical analysis combines the "double dip" recessions of 1980 and 1981 with the peak occurring in 1980:Q1 and the trough in 1982:Q4.

[2]These shares vary over time, both in terms of a trend and over the business cycle. The shares reported here are based on recent data; in the 1950s and 1960s, for example, durables made up a smaller share of consumption (around 5 percent), and nondurables had a larger share (around 31 percent). The services share has remained relatively stable.

[3]The specification does not permit an autoregressive component to consumption growth, and therefore consumption follows a pure random walk with possibly time-varying drift. Taking account of persistence may be important for other measures of the business cycle (Davig), the dynamic responses of the macroeconomy to monetary policy shocks (Sims and Zha), or many other applications.

[4]Model selection for Markov-switching models can be problematic, since regularity conditions needed for likelihood ratio tests break down (Smith and others), and other information criteria such as Akaike's Information Criterion tend to over-fit and select too many regimes (Fruhwirth-Schnatter). The SBIC here provides a more accurate criterion and has been shown to be useful in larger structural models (Liu and others).

[5]Given the Markov transition probabilities, the expected duration of the low average growth regime is $1/(1-P_{ll}^{\mu})$, and similar for the high average growth and both volatility regimes. Note that small differences in the probability of switching regimes can have relatively large differences in the regimes' expected durations.

[6]This probability is "smoothed" in the sense that it is both backward and forward looking; in other words, the calculation for a given quarter uses information from data both before and after that quarter to infer the regime. This method contrasts with a backward-looking probability that would only use data up to the quarter in question.

References

Davig, Troy. 2008. "Detecting Recessions in the Great Moderation: A Real-Time Analysis," Federal Reserve Bank of Kansas City, *Economic Review*, vol. 93, no. 4, pp. 5–33.

Fruhwirth-Schnatter, Sylvia. 2006. *Finite Mixture and Markov Switching Models.* New York: Springer.

Hamilton, James. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, vol. 57, no. 2, pp. 357–384.

Lettau, Martin, Sydney Ludvigson, and Jessica Watcher. 2008. "The Declining Equity Premium: What Role Does Macroeconomic Risk Play?" *Review of Financial Studies*, vol. 21, no. 4, pp. 1653–1687.

Liu, Zheng, Daniel Waggoner, and Tao Zha. 2011. "Sources of Macroeconomic Fluctuations: A Regime-Switching DSGE Approach," *Quantitative Economics*, vol. 2, no. 2, pp. 251–301.

Kim, Chang-Jin, and Charles Nelson. 1999. "Has the Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle," *Review of Economics and Statistics,* vol. 81, no. 4, pp. 608–616.

McConnell, Margaret, and Gabriel Perez-Quiros. 2000. "Output Fluctuations in the United States: What has Changed Since the Early 1980s?" *American Economic Review,* vol. 90, no. 5, pp. 1464–1476.

Sims, Chris, and Tao Zha. 2006. "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review*, vol. 96, no. 1, pp. 54–81.

Has the Relationship between Bank Size and Profitability Changed?

By Kristen Regehr and Rajdeep Sengupta

In recent years, community bankers and industry analysts have raised concerns that smaller community banks need to grow larger to be successful. Today, banks face new and higher costs to both implement complex new regulations, especially those introduced after the 2007–09 financial crisis and recession, and transition to new electronic banking platforms. For small banks, these higher fixed costs are spread over a smaller asset base, which may put them at a competitive disadvantage relative to larger competitors. In addition, if the competitive disadvantage threatens the profitability and long-run viability of smaller banks, smaller communities and rural areas not large enough to support viable banks may lose access to their local banking services. Even if these communities do not lose banking services, a reduction in the number of banks can reduce competition, which may then lead to higher loan rates and lower deposit rates.

However, size is not the only factor that affects a bank's long-run profitability. In fact, profitability depends on the characteristics of both individual banks and the markets in which they operate. For example, bank-specific factors such as business strategies, reflected in the composition of banks' assets and liabilities, can affect profitability. Likewise, market-specific factors, such as growth in the markets in which banks

Kristen Regehr is an assistant economist at the Federal Reserve Bank of Kansas City. Rajdeep Sengupta is an economist at the bank. This article is on the bank's website at ***www.KansasCityFed.org.***

operate, can affect banks' long-run profitability. Any analysis that thoroughly examines the relationship between bank profitability and bank size must account for such bank-specific and market-specific factors.

In this article, we analyze how bank profitability changes with bank asset size after accounting for other factors that affect bank profitability. More specifically, we examine whether the size-profitability relationship has made smaller community banks less competitive in the post-crisis recovery. We find that profitability, measured by banks' return on assets, increases with bank size but at a decreasing rate. However, we also find no statistically significant difference in the size-profitability relationship before and after the crisis, suggesting the relationship has not changed in recent years in ways that disadvantage community banks relative to their larger competitors.

Section I describes the factors that affect bank profitability and the size-profitability relationship. Section II conducts a statistical analysis of the relationship between bank size and bank profitability after controlling for other factors. Section III examines whether the size-profitability relationship has remained stable over time by comparing the relationship during the periods before, during, and after the financial crisis.

I. Bank Profitability and Size

The banking industry has undergone significant restructuring over the last three decades. Since the mid-1980s, the number of commercial banks has declined, while the average assets of banks have continued to increase. These changes appear to have had a disparate effect on small banks. From 1984 to 2011, the number of banks with assets less than $100 million declined by over 11,000, largely due to the consolidation of bank charters. And while banks' average assets increased over the same period, most of the growth can be attributed to banks with more than $10 billion in assets (FDIC).

Banks have good reasons to believe profitability and size are related. Increasing bank size can increase bank profitability by allowing banks to realize economies of scale. For example, increasing size allows banks to spread fixed costs over a greater asset base, thereby reducing their average costs. Increasing banks' asset size can also reduce risk by diversifying operations across product lines, sectors, and regions (Mester 2010). Lower risk can promote profitability directly by reducing losses or indirectly by making liability holders willing to accept lower returns, thereby reducing banks' funding costs. Furthermore, as the scale of operations increases, banks may be able to better use specialized inputs such as loan officers with expertise in commercial and industrial business lines, resulting in greater efficiency. Realizing economies of scale may lead to a healthier banking system by eliminating inefficiencies and reducing risks.

However, scale economies are not the only way size can affect profitability. Small banks may be able to form stronger relationships with local businesses and customers than large banks, allowing them access to proprietary information useful in setting contract terms and making better credit underwriting decisions (Berger and others). Indeed, these informational and pricing advantages may fully offset any loss of scale economies. To determine how size affects bank performance, then, it is important to use a measure such as profitability that summarizes the various costs and benefits of size.

A simple comparison across the various size groups suggests that, on average, larger banks have higher returns. Table 1 shows return on average assets (ROAA) for different bank size groups from 2001 to 2014. The second column of Table 1 shows that average returns are highest for the more than $10 billion group at 1.09 percent and smallest for the less than $1 billion group at 0.77 percent. However, this relationship differs across the three subperiods. Larger banks saw higher returns in the pre-crisis and post-crisis periods, but smaller banks saw higher returns during the crisis. One problem with drawing conclusions about the effect of size on returns from Table 1 is the comparisons do not take other factors that can affect bank returns into account. An analysis that controls for these factors can better determine the relationship between size and profitability and whether this relationship has changed over time.

To determine how bank size affects bank profitability, we develop a simple model where a bank's profitability is a function of its size and characteristics as well as the characteristics of the markets in which it operates. Bank-specific factors include business strategies and other bank characteristics such as organizational structure. Market-specific factors include market competition and local economic conditions. Controlling for the influence of bank- and market-specific factors allows us to isolate the relationship between bank size and profitability.

Table 1
ROAA and Size Group

Size group	Sample		Pre-crisis expansion		Crisis		Post-crisis expansion	
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
All banks	0.78	1.38	1.04	1.31	0.47	1.58	0.61	1.25
Less than $1 billion	0.77	1.35	1.01	1.31	0.50	1.50	0.59	1.24
$1–$10 billion	0.86	1.69	1.34	1.35	0.15	2.27	0.78	1.38
More than $10 billion	1.09	1.38	1.58	1.11	0.14	1.94	0.93	0.89

Bank-specific factors

Banks' business strategies can affect profitability. We assess banks' business strategies by examining the strategic decisions that affect the composition of banks' balance sheets—the level of earning assets, the proportion of assets allocated to loans and securities, and the proportion of funding generated through core deposits and wholesale liabilities. Banks that focus on loans (as opposed to securities), for example, tend to generate higher interest income but entail higher expense and risk. In addition, we distinguish between funding strategies that rely on core deposits, a safer and more liquid source of funding, and those that rely on brokered deposits, which are more easily obtained but also less stable. Each strategy has advantages and disadvantages; we do not know in advance which strategy leads to higher profitability.

Other bank-specific characteristics, such as organizational structure, can also affect bank profitability. To account for these factors, we first analyze differences between single-market and multimarket banks. Multimarket banks may derive benefits, such as lower funding costs, from diversifying across different markets. In contrast, single-market banks are significantly smaller and therefore more likely to benefit from the advantages of small banks. For example, the geographically undiversified nature of single-market bank loan portfolios encourages banks to increase profitability by building up local lending relationships with a loyal customer base over time.

We also examine differences between banks that file taxes under Subchapter S (S-Corp banks) and banks that file under Subchapter C (C-Corp banks) of the U.S. tax code.[1] S-Corp banks have fewer owners due to restrictions on the number of shareholders allowed. Such concentrated ownership relative to C-Corp banks may reduce agency problems and subsequently improve shareholder control over management and risk management practices, leading to higher profitability.

Market-specific factors

Market competition can directly affect a bank's profitability. Banks in less competitive markets, for example, tend to offer lower deposit rates and charge higher loan rates, leading to higher returns. In more competitive markets, however, banks may realize lower returns as they bid for funds with higher deposit rates and try to attract borrowers with lower loan rates.

Other market-specific factors affecting profitability include market size and economic conditions. Large markets, as measured by population, may provide banks more opportunities to increase returns but may also be more competitive. Markets with stronger economic conditions, as reflected in lower unemployment rates, tend to raise bank profitability.

The size-profitability relationship

To evaluate the size-profitability relationship appropriately, we must account for other factors that affect profitability. Figure 1 shows a hypothetical bank size-profitability relationship for a given set of bank-specific and market-specific factors. The curve shows that bank profitability increases with bank size but at a decreasing rate. Consider, for example, two banks with the same characteristics operating in the same market that differ only in size—one bank has $300 million in assets, and the other bank has $500 million in assets. Figure 1 shows higher returns for the $500 million bank than the $300 million bank, though the slope of the curve suggests these effects diminish as the bank's asset size increases. Still, when all bank-specific and market-specific factors are accounted for, greater size is associated with higher profitability.

However, changes in bank-specific or market-specific factors can raise or lower profitability for banks of all sizes. Figure 2 shows how the size-profitability curve shifts in response to these changes. The lower curve represents the size-profitability relationship for banks that use the average funding strategy, while the higher curve represents the relationship for banks that use a more profitable core funding strategy. The

Figure 1
The Size-Profitability Relationship

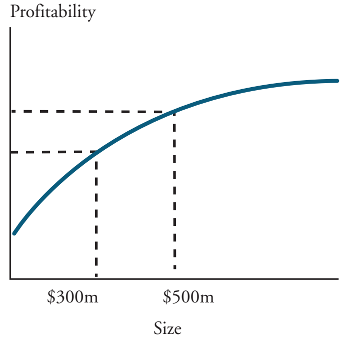

Figure 2
Shift in the Size-Profitability Relationship

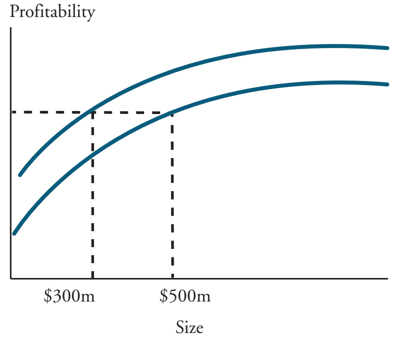

difference between the curves illustrates how a $300 million bank with a more profitable funding strategy can generate the same returns as a $500 million bank with the average funding strategy. While greater size is still associated with higher profitability within each funding strategy, Figure 2 illustrates how other factors can enable smaller banks to compete effectively with larger banks.

II. How Does Bank Size Affect Profitability?

The hypothetical exercise in Figure 2 illustrates the importance of controlling for other factors in determining the relationship between size and profitability. We use simple regression techniques on bank-level data to estimate a size-profitability model that controls for bank-specific and market-specific factors. Our results show that size is an important determinant of bank profitability and that its effect increases but at a diminishing rate.

Our sample comprises an unbalanced panel of annual observations for 8,315 community and regional banks with assets less than $100 billion (valued at 2014 U.S. dollars) from 2001 to 2014.[2] We choose the 2001–14 sample period to allow for the pre-crisis and post-crisis comparison in our statistical analysis. We divide the years in the sample into three sub-periods: the pre-crisis expansion from 2001 to 2006 (7,451 banks), the crisis period during 2007–09 (6,510 banks), and the post-crisis recovery period from 2010 to 2014 (6,326 banks).[3]

We measure bank size as the natural logarithm of total assets (valued at 2014 U.S. dollars).[4] We measure bank profitability as the tax-adjusted ROAA, which is a bank's tax-adjusted net income divided by its average total assets over the past year. We also include the square of the logarithm of total assets in the model to capture changes in the size-profitability relationship as bank size changes—specifically, to capture changes in the rate at which profitability increases or decreases as bank size increases. We calculate ROAA, bank size, and other bank financial ratios using annual bank-level data on U.S. commercial banks from the Consolidated Reports of Condition and Income for a bank, popularly known as the Call Reports (see Appendix for a description of the data).

We measure the level of competition within markets using the market Herfindahl Index (HHI). Higher HHIs indicate more concentration and less competition. The HHI for a market is calculated using the market shares of deposits from the FDIC's Summary of Deposits. More specifically, the HHI is calculated using deposit shares at banks belonging to the same bank holding company (BHC). Calculating HHIs at the BHC level rather than the bank level is reasonable since two banks in the same market belonging to the same holding company are unlikely to compete with each other.

We also use data from the Summary of Deposits to distinguish between single-market and multimarket banks.[5] For multimarket banks, we define the market variables to include all areas in which the bank or its branches are located. Accordingly, the market population for a multimarket bank is the sum of the population in every market area in which the bank has a branch. The HHI and the unemployment rate for multimarket banks are weighted averages for the banks' market areas. The HHI is weighted by the relative size of the population, while the unemployment rate is weighted by the relative size of the labor force.

Base model

The size-profitability model regresses banks' tax-adjusted ROAA on their asset size and the square of the size variable. The regression estimates how bank size affects bank asset returns while controlling for variations in bank-specific and market-specific factors. Bank-specific factors include balance sheet composition variables such as loan to asset ratio, securities to assets ratio, core deposits to total deposits ratio, and binaries for single-market banks and S-Corp banks. Market-specific factors include a measure for bank competition (HHI), population size, and the unemployment rate. The regressions also use other explanatory variables such as bank age, risk, and a binary for rural banks that control for potential variations in profitability. Finally, we use bank-specific, time-invariant binaries (fixed effects) to control for bank-level heterogeneity; the annual GDP growth rate to control for macroeconomic factors; and binary variables for each of the three periods before, during, and after the financial crisis to control for other variations over time (see the Appendix for a complete variable list).

The bank size variables are our principal interest. The estimated coefficients on the size variable capture the change in profitability associated with a 1 percent increase in real assets holding all other factors constant. A positive coefficient indicates that profitability increases with size, whereas a negative coefficient indicates profitability decreases with size. The squared term captures the rate of acceleration or deceleration in profitability associated with a percentage change in real assets. Accordingly, a positive effect on the squared term indicates an increasing rate of change, whereas a negative effect indicates a decreasing rate of change.

The coefficients of the size and size-squared variables, as shown in the second column of estimates in Table 2, are positive and negative, respectively, and statistically significant at the 1 percent level (see Appendix for full regression results).[6] The coefficients indicate that percentage increases in size are associated with increasing conditional ROAA but at a decreasing rate (the ROAA estimated here is conditional on holding bank-specific

Table 2

ROAA Regression Results

Variables	Size-profitability model	Post-crisis break model
Size	2.915***	2.064***
Size^2	−0.071***	−0.048***
Loan to asset ratio (one-year lag)	1.112***	1.105***
Security to asset ratio (one-year lag)	1.199***	1.168***
Core deposit to deposit ratio (one-year lag)	0.498***	0.496***
Single-market bank	0.118***	0.118***
Age	−0.401***	−0.408***
Risk	−0.094***	−0.093***
Subchapter S bank	0.066***	0.056**
Rural bank	−0.026	−0.016
Population level	−0.052***	−0.052***
Unemployment rate	−0.121***	−0.121***
HHI	0.269	0.264
Real GDP growth rate	0.050***	0.049***
Crisis binary variable	−0.206***	−7.255***
Size × crisis binary variable		0.862***
Size^2 × crisis binary variable		−0.026***
Post-crisis expansion binary variable	−0.026	−1.707
Size × post-crisis expansion binary variable		0.203
Size^2 × post-crisis expansion binary variable		−0.006
Observations	86,706	86,706
Number of banks	8,315	8,315
Adjusted R^2	0.085	0.089
F-stat	179.73	151.13

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.
Notes: Standard errors are robust and clustered at the bank level. See Appendix for full regression results.

Table 3
Conditional ROAA and Asset Size

Asset size (millions)	Size-profitability model (basis points)	Post-crisis break model		
		Pre-crisis expansion (basis points)	Crisis (basis points)	Post-crisis expansion (basis points)
$100	16.40	19.06	11.70	17.43
$200	6.41	9.02	3.57	7.80
$300	3.13	5.46	1.07	4.48
$400	1.61	3.69	–0.01	2.85
$500	0.79	2.66	–0.55	1.93
$600	0.30	2.01	–0.85	1.37
$700	–0.01	1.55	–1.01	0.97
$800	–0.22	1.23	–1.11	0.70
$1,400	–0.66	0.37	–1.19	0.01
$1,500	–0.67	0.30	–1.17	–0.03
$1,600	–0.69	0.25	–1.17	–0.07
$2,400	–0.70	0.01	–1.03	–0.22
$2,500	–0.69	–0.02	–1.00	–0.23

Note: The values show the change in ROAA given a $100 million increase in asset size for a bank with the initial asset size shown in the first column.

and market-specific variables constant at their mean values). In particular, a 1 percent increase in a bank's real assets is associated with an increase in conditional ROAA of 2.9 basis points minus twice the bank's initial size (measured in logarithm of real assets) times 0.07 basis point.[7] In other words, the change in conditional ROAA associated with a given change in bank real assets varies with the initial size of the bank.

Next, we use the coefficients from the size-profitability model to calculate the change in conditional ROAA associated with a $100 million increase in bank assets. The first column of estimates in Table 3 shows the results for different bank sizes. The estimates suggest the smallest banks experience large increases in conditional ROAA as they grow. Specifically, an increase in bank size from $200 million in assets to $300 million in assets is associated with an increase in conditional ROAA of 6.4 basis points. However, the size of this increase diminishes as bank size increases. Banks with larger asset sizes experience much smaller increases in conditional ROAA.

Conditional returns are maximized at $755 million under the size-profitability model. Increases in size beyond this level are associated with decreases in conditional ROAA. Returns on assets continue to be positive beyond this size but are lower for larger banks. The returns decline slowly for each $100 million increase in size. As shown in Table 3, increasing size from a $1.5 billion bank to a $1.6 billion dollar bank is associated with a decrease in conditional ROAA of 0.7 basis point.

III. How Has the Bank Size-Profitability Relationship Changed in Recent Years?

Changes in technology and regulation have the potential to affect both bank profitability and the size-profitability relationship. Changes in factors that affect the profitability of all banks will shift the size-profitability curve up or down. However, changes in factors that have disparate effects on banks of different sizes will change the size-profitability relationship. Figure 3 illustrates a change in the size-profitability relationship from a technological change that favors larger banks. The lower curve shows the size-profitability relationship before the technological change, and the higher curve shows the size profitability relationship after the technological change. As Figure 3 shows, larger banks experience a greater increase in profitability after the change.

Often, technological and regulatory changes favor larger banks. The Interstate Banking and Branching Efficiency Act (IBBEA) of 1994, for example, paved the way for consolidation, thereby allowing banks to exploit scale economies.[8] Recent studies of scale economies claim that the efficient scale of commercial banking has risen over the past 20 years (Wheelock; Feldman, Mester, and DeYoung). Improvements in information technology have also increased productivity and scale economies in processing electronic payments (Berger). However, technological changes can benefit small banks as well: for example, small banks may be able to benefit from the services of third-party providers without having to develop new banking platforms on their own.

After the financial crisis and recession of 2007–09, the banking industry underwent significant technological and regulatory changes. Banks introduced new technological innovations in mobile and online banking in the post-crisis period. In addition, the Dodd-Frank Act of 2010 introduced new financial regulations to reduce risks to the banking sector and to enhance overall financial stability. We want to determine whether these changes have significantly altered the size-profitability relationship in the post-crisis period.

Figure 3
Change in the Size-Profitability Relationship

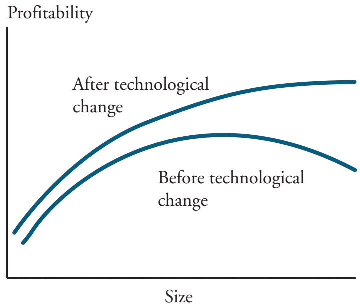

We develop a post-crisis break (PCB) model to examine the relationship between bank size and profitability over three periods: the pre-crisis expansion, the crisis, and the post-crisis recovery. The PCB model allows the coefficients on the size variables in the size-profitability model to vary across the subperiods, allowing us to evaluate whether the size-profitability relationships changed in these periods. The PCB model includes a crisis binary variable indicating whether the sample observation belongs to the crisis period of 2007–09. The model also includes a post-crisis binary variable indicating whether the sample observation belongs to the post-crisis expansion.

The PCB model finds a statistically and economically significant change in the size-profitability relationship during the crisis. The second column of estimates in Table 2 shows the coefficients for the post-crisis binary variables and for the binary variables interacted with the size variables. The estimated coefficients on the interaction terms are statistically significant at the 1 percent level and indicate that while conditional ROAA still increased with size during the crisis, these returns diminished at a faster rate than in the pre- and post-crisis expansions. As a result, the bank asset size associated with the maximum conditional ROAA was significantly smaller during the crisis.

The coefficients on the post-crisis binary and its interactions with size and size-squared variables allow us to compare the size-profitability relationship during the pre- and post-crisis expansions as well. The estimated coefficients on the post-crisis expansion variables again indicate that increasing bank size was associated with higher conditional ROAA, but returns in this period diminished at a slower rate than during the crisis. However, none of the estimated coefficients on the post-crisis binary variable are statistically significant, suggesting the size-profitability relationship in the two periods was not significantly different.[9] The observed differences in ROAA in the pre-crisis and post-crisis expansions shown in Table 1 are due instead to changes in economic and competitive conditions that shift conditional ROAA downward without changing its sensitivity to size. For example, a high unemployment rate could lower profitability during the post-crisis recovery. The mean post-crisis unemployment rate across all banking markets, 7.3 percent, was significantly higher than the pre-crisis mean of 5.1 percent (see Appendix for summary statistics on the regression variables).

The PCB model columns in Table 3 show how returns change with increases in size under the PCB model. For a benchmark bank—one with bank-specific and market-specific factors at their mean values—conditional asset returns in the pre-crisis expansion are maximized at $2.5 billion.[10] We derive the increases in the post-crisis column from the estimates on the post-crisis binary variable and interacted size variables in Table 2. While the size of the benchmark bank in Table 3 is smaller in the post-crisis period, we calculate the size using estimated coefficients that are not statistically different from zero. For small banks, the increases in returns for a $100 million increase in size are comparable in the pre-crisis and post-crisis periods. For larger banks, the difference in returns between the two periods is less than one basis point.

Chart 1 shows the estimated relationship between size and the benchmark conditional ROAA for the pre-crisis and post-crisis periods. In the pre-crisis period, the benchmark conditional ROAA is higher than in the post-crisis period for most bank sizes. However, the two curves are similar in that the relationship between profitability and size remains fairly stable during the pre-crisis and post-crisis expansions. Our statistical analysis confirms that while ROAA is lower on average for all bank sizes in the post-crisis period than in the pre-crisis period, the reduction in ROAA cannot be attributed to a diminished effect of size on ROAA.

Our analysis reveals that the difference in the size-profitability relationship between the PCB model and the baseline size-profitability model is statistically significant, largely due to a change in the size-profitability relationship during the crisis period.[11] The PCB model allows

Chart 1

Predicted Conditional ROAA and Bank Size (Pre-Crisis and Post-Crisis Expansions)

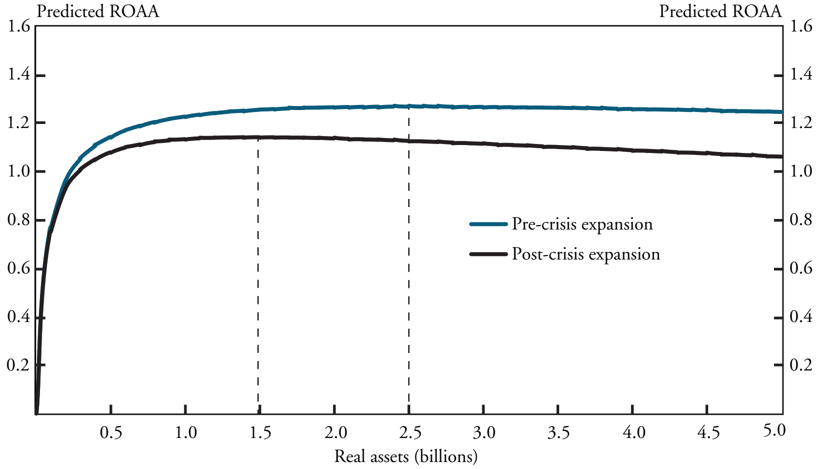

us to isolate the influence of the crisis period and estimate the equilibrium relationship during the pre-crisis and post-crisis expansions. However, the PCB model shows no statistically significant difference between the pre-crisis and the post-crisis periods. We therefore use this equilibrium relationship that remains unaltered over the pre-crisis and post-crisis expansions to examine the effects of bank-specific and market-specific factors on profitability.

The effect of factor variables

We use the PCB model to study the effect of bank-specific and market-specific factors on bank profitability. Specifically, we explore how a 10 percent change from the mean value of select bank-specific and market-specific variables affects profitability. As shown in Table 4, these 10 percent changes are small in terms of our sample—less than the standard deviations of the sample variables. Table 4 estimates the minimum size bank that can achieve the maximum conditional ROAA for the $2.5 billion benchmark bank (benchmark conditional ROAA) given a 10 percent change in a single factor. We focus on two bank-specific factors—the core deposit ratio and whether the bank is a single-market or multimarket bank—and one market-specific factor, the unemployment rate.

Table 4
Benchmark Conditional ROAA and Asset Size

Variables	Mean	Standard deviation	Value after 10 percent change from mean*	Change as a percentage of standard deviation	Real assets (millions)
Loan to asset ratio	0.63	0.16	0.69	40	759.8
Security to asset ratio	0.23	0.15	0.26	15	1,191.2
Core deposit to deposit ratio	0.83	0.10	0.91	79	1,001.6
Population level	5.61	2.27	5.05	25	1,160.7
Unemployment rate	6.01	2.20	5.41	27	736.0
Single-market bank**	-	-	-	-	525.2

*Assumes change in variable increases ROAA
**Asset size needed to achieve maximum estimated ROAA if bank is a single-market bank
Note: Table shows the minimum size bank that can achieve the maximum conditional ROAA for the $2.5 billion benchmark bank (benchmark conditional ROAA) given a 10 percent change in a single factor.

The PCB model column in Table 2 shows that the core deposit to total deposit ratio has a positive and statistically significant effect on conditional ROAA. This implies that increasing core deposits increases bank returns an upward shift of the curve in Figure 2. We quantify this effect by considering a 10 percent increase in the core deposit ratio from 83 percent (the sample mean) to 91 percent. Table 4 shows that this change is less than 1 standard deviation of the variable but would allow a $1 billion bank to achieve the same ROAA as the $2.5 billion benchmark bank.

We obtain similar estimates for market-specific factor variables. Table 2 shows that the estimated coefficient for the unemployment rate is negative and significant under the PCB model. Again, we would expect banks in market areas with a lower unemployment rate to show higher returns. Thus, we consider a 10 percent decrease in the unemployment rate from 6 percent (the sample mean) to 5.4 percent in a given market area. Table 4 shows that this change, which is smaller than one-third of the variable's standard deviation, would allow a $736 million bank to achieve the same ROAA as the $2.5 billion benchmark bank.

Single-market banks are smaller in size and scope, and they lack the advantages of diversification that multimarket banks typically accrue. In the PCB model, the coefficient of the indicator variable for single-market banks is economically and statistically significant (Table 2). The estimated coefficient shows that, after controlling for size, single-market banks have higher returns than multimarket banks. Table 4 shows that a $525 million single-market bank can achieve the same ROAA as the $2.5 billion benchmark bank.

The comparisons in Table 4 show that favorable market conditions and changes in bank-specific characteristics other than size also increase conditional ROAA. Small changes in bank-specific and market-specific factors can be equivalent to large changes in size in achieving higher ROAA.

IV. Conclusion

Our results support industry analysts' view that there are significant scale economies in banking, especially for the smallest community banks. However, this is not merely a post-crisis phenomenon. Throughout our sample period, small community banks have exhibited significant scale economies. While the smallest banks can benefit significantly from growth, the advantages of growth become progressively smaller until they are exhausted. For most midsized community banks, the increase in returns relative to size is modest; these banks would need large increases in size to realize significantly higher returns. The relationship between size and profitability remains unchanged between the pre-crisis and post-crisis expansions. In other words, we find the post-crisis economic and regulatory environment has not disproportionately affected the size-profitability relationship for small community banks.

An important caveat is that our results are not causal: higher returns are associated with larger banks, but increases in size do not necessarily *cause* increases in returns. Indeed, banks with higher returns may simply be better positioned to grow.

Regardless, our analysis suggests the competitive disadvantage of community banks in the post-crisis period may be overstated. The decline in profitability during the post-crisis recovery cannot be attributed to size or any change in the size-profitability relationship. Rather, changes in economic and competitive conditions lowered post-crisis profitability without affecting its sensitivity to size. In particular, our analysis shows that factors other than size, such as large differences in banking market unemployment rates between the pre-crisis and post-crisis expansions, can account for the lower profitability of community banks in the post-crisis recovery.

Finally, our results show that favorable market outcomes and changes in other bank-specific characteristics also increase returns. In achieving higher profitability, small changes in bank-specific and market-specific factors are equivalent to large changes in size. Therefore, banks need not grow larger to be successful: business strategies and local economic growth are no less important in determining bank profitability than size.

Appendix

Data Description and Variable Definitions

The primary data sources are the FDIC's Summary of Deposits (SOD) and the FFIEC Call Report (031/041). The SOD data are as of second quarter (the FDIC conducts the survey annually at the end of the second quarter), while the Call Report data are as of fourth quarter to match annual bank profitability to the annual macroeconomic data available for the geographic regions. The regression data are an unbalanced panel of annual observations from 2001 to 2014. The sample excludes banks with real assets of $100 billion or more, de novo banks (defined as banks less than five years of age), and other nontraditional banks, such as credit card banks and banks that do not take deposits or make loans.

The variables are divided into six different categories and are defined as follows:

Dependent variable

ROAA: Annual net income divided by average total assets over the year. For S-Corp banks net income is adjusted to account for differences in tax treatment.

Bank-specific variables

Size: Natural logarithm of real assets measured in 2014 dollars.

Size-squared: Square of the natural logarithm of real assets measured in 2014 dollars.

Age: Number of years that the bank has been operating.

Risk: Volatility of bank earnings measured by the standard deviation of quarterly ROAA for prior three years.

The following variables are ratios and do not require any inflation adjustment. In the regression analysis, lagged values of the variables are used.

Loan to asset ratio: Total loans divided by total assets

Security to asset ratio: Total securities divided by total assets.

Core deposit to deposit ratio: Sum of transactions accounts, money market deposit accounts, time deposits of less than $100,000, and other non-transaction savings deposits divided by total deposits.

Competition variable

HHI: Sum of squared bank deposit shares in a market. For multimarket banks, HHI is weighted by the relative size of the population.

Market variables

Population: Natural logarithm of annual market population from the Census Bureau. For multimarket banks, population is the sum of the population in every market area in which the bank has a branch. (Source: Census Bureau)

Unemployment Rate: Annual market unemployment rate from the Bureau of Labor Statistics. For multimarket banks, unemployment rate is weighted by the relative size of the labor force.

Macroeconomic variable

Real GDP Growth: Annual growth rate of real GDP from the Bureau of Economic Analysis.

Binary variables

S-Corp bank: Bank that has elected to be taxed under subchapter S of the tax code.

Single-market bank: Bank that has at least 99 percent of its deposits in a single market.

Rural bank: Bank that has at least 90 percent of its deposits in counties located outside of metropolitan or micropolitan statistical areas.

Table A-1
Regression Results

Variables	Size-profitability model	Post-crisis break model
Size	2.915*** 0.616	2.064*** 0.678
Size^2	−0.071*** 0.016	−0.048*** 0.018
Loan to asset ratio (one-year lag)	1.112*** 0.284	1.105*** 0.285
Security to asset ratio (one-year lag)	1.199*** 0.218	1.168*** 0.218
Core deposit to deposit ratio (one-year lag)	0.498*** 0.163	0.496*** 0.162
Single-market bank	0.118*** 0.031	0.118*** 0.031
Age	−0.401*** 0.079	−0.408*** 0.078
Risk	−0.094*** 0.023	−0.093*** 0.023
Subchapter S bank	0.066*** 0.023	0.056** 0.023
Rural bank	−0.026 0.040	−0.016 0.040
Population level	−0.052*** 0.013	−0.052*** 0.013
Unemployment rate	−0.121*** 0.005	−0.121*** 0.005
HHI	0.269 0.169	0.264 0.169
Real GDP growth rate	0.050*** 0.004	0.049*** 0.004
Size × crisis binary variable		0.862*** 0.248
Size^2 × crisis binary variable		−0.026*** 0.006
Size × post-crisis expansion binary variable		0.203 0.206
Size^2 × post-crisis expansion binary variable		−0.006 0.005
Crisis binary variable	−0.206*** 0.016	−7.255*** 2.407
Post-crisis expansion binary variable	−0.026 0.018	−1.707 2.001
Constant	−27.590*** 5.979	−19.960*** 6.510

Table A-1 Continued

Observations	86,706	86,706
Number of banks	8,315	8,315
Adjusted R^2	0.085	0.089
F-statistic	179.73	151.13

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.
Notes: The dependent variable is the annual return on average assets for U.S. commercial banks during 2001–14. All regressions include bank- and year-fixed effects. Standard errors are clustered by bank.

Table A-2

Summary Statistics of Regression Variables

	Sample		Pre-crisis expansion		Crisis		Post-crisis expansion	
Variable	Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation
Adjusted ROAA	0.78	1.38	1.04	1.31	0.47	1.58	0.61	1.25
Real assets (millions)	589	3,300	608	3,730	570	3,020	573	2,790
Loan to asset ratio	0.63	0.16	0.63	0.16	0.66	0.16	0.60	0.16
Security to asset ratio	0.23	0.15	0.24	0.15	0.21	0.14	0.23	0.16
Core deposit ratio	0.83	0.10	0.84	0.10	0.81	0.10	0.82	0.11
Single-market bank	0.61	0.49	0.65	0.48	0.59	0.49	0.56	0.50
Age	73	40	72	38	73	40	75	42
Risk	0.62	1.10	0.49	0.95	0.57	1.19	0.85	1.20
Subchapter S bank	0.33	0.47	0.28	0.45	0.37	0.48	0.38	0.49
Rural bank	0.23	0.42	0.25	0.43	0.23	0.42	0.21	0.41
Population (thousands)	1,998	4,744	1,790	4,164	2,016	4,607	2,279	5,521
Unemployment rate	6.01	2.20	5.10	1.36	6.07	2.45	7.26	2.36
HHI	0.18	0.12	0.18	0.12	0.18	0.12	0.19	0.12
Real GDP growth rate	1.76	1.59	2.56	0.94	-0.39	1.86	2.05	0.43
Number of banks	8,315		7,451		6,510		6,326	

Endnotes

[1]S-Corp banks are able to pass their federal corporate income tax obligations through to their shareholders. The profitability measure used in this analysis is the ROAA adjusted for tax effects of S-Corp status and tax-advantaged investments.

[2]We also conduct the statistical analysis on a sample of banks with assets less than $50 billion, but the results are not materially different.

[3]The sample is affected by survivorship bias, as banks that fail during the period drop out of the sample. This tends to bias the results toward higher returns post-crisis because poorly performing banks are no longer in the sample.

[4]The distribution of bank assets are positively skewed (long right tail). A logarithmic transformation gives us a symmetric distribution more suitable for regression analysis.

[5]We define banks holding at least 99 percent of their deposits in a single market area as single-market banks. We define banks that do not meet this criterion as multimarket banks.

[6]The size-profitability relationship could be better described by a higher degree polynomial in size. To test this hypothesis, we include size-cubed as an explanatory variable. However, its estimated coefficient is not statistically significant.

[7]The F-test of the estimated coefficient on size plus twice the estimated coefficient on size-squared is statistically significantly different from zero. We apply this F-test throughout the analysis to test whether size has a significant influence on returns in any period.

[8]IBBEA permitted banks and BHCs to expand across state lines. However, individual states were granted powers to restrict entry by out-of-state banks using different means, such as restricting de novo interstate branching (Strahan and Rice). States took advantage of these powers and in some cases have progressively deregulated entry and competition in banking even in recent years.

[9]We also estimate a fully interacted model in which all variables in the size-profitability model are interacted with the crisis and post-crisis binary variables. The fully interacted model yields the same result: the size-profitability relationship is not significantly different in the pre-crisis and post-crisis periods.

[10]The large difference in sizes at which conditional returns are maximized under the size-profitability model and the PCB model can be attributed to the effect of the crisis. If we estimate the size-profitability model for a subsample that includes only the pre-crisis and post-crisis periods, conditional returns are maximized at a bank size of $1.9 billion.

[11]An F-test rejects the null hypothesis that the interaction terms are all zero. The F-statistic gives a p-value equal to zero, so we can reject the null hypothesis at the 1 percent level.

References

Berger, Allen. 2003. "The Economic Effects of Technological Progress: Evidence from the Banking Industry," *Journal of Money, Credit and Banking*, vol. 35, no. 2, pp. 141–176.

Berger, Allen N., Nathan H. Miller, Mitchell A. Petersen, Raghuram G. Rajan, and Jeremy C. Stein. 2005. "Does Function Follow Organizational Form? Evidence from the Lending Practices of Large and Small Banks," *Journal of Financial Economics*, Elsevier, vol. 76, no. 2, pp. 237–269, May.

Federal Deposit Insurance Corporation (FDIC). 2012. "FDIC Community Banking Study," December.

Mester, Loretta J. 2010. "Scale Economies in Banking and Financial Regulatory Reform," *The Region*, Federal Reserve Bank of Minneapolis, September, pp. 10–13

———. 2008. "Optimal Industrial Structure in Banking," in *Handbook of Financial Intermediation*, Arnoud Boot and Anjan Thakor, eds. Amsterdam: North-Holland, pp. 133–162.

Wheelock, David C. 2012. "Too Big to Fail: the Pros and Cons of Breaking up Big Banks," *The Regional Economist*, Federal Reserve Bank of St. Louis, October, pp. 10–11.

Economic Review

Second Quarter 2016, Volume 101, Number 2

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
Second Quarter 2016, Vol. 101, No. 2