FEDERAL RESERVE BANK OF KANSAS CITY

ECONOMIC REVIEW

Fourth Quarter 2022 *Volume 107, Number 4*

Hybrid Working, Commuting Time, and the Coming Long-Term Boom in Home Construction 5
By Jordan Rappaport

Can Higher Gasoline Prices Set Off an Inflationary Spiral? 41
By Nida Çakır Melek, Francis M. Dillon, and A. Lee Smith

The Evolving Role of the Fed's Balance Sheet: Effects and Challenges 59
By Chaitri Gulati and A. Lee Smith

Hybrid Working, Commuting Time, and the Coming Long-Term Boom in Home Construction

By Jordan Rappaport

The long commuting times from outer suburbs to the central business districts of large metropolitan areas have depressed single-family home construction over the last two decades. The shift to hybrid working during the COVID-19 pandemic, however, has reduced the time many people spend commuting and potentially increased their willingness to live farther from employers. Lightly settled land at the peripheries of metropolitan areas may become more desirable for development, relaxing a long-standing constraint on single-family home construction.

In this article, Jordan Rappaport estimates the statistical relationship between single-family home construction and commute duration, finding a significant negative correlation across the largest U.S. metropolitan areas. Across metropolitan areas with a population of at least 1 million in 2020, the time employees save commuting from outer suburbs to central business districts by working remotely two days a week ranges from 130 to 406 hours per year. These results suggest that in the long run, the time savings from fewer commutes could almost double single-family home construction in these metropolitan areas from its level just prior to the pandemic, an aggregate increase of 427,000 units per year. The largest metropolitan areas are likely to see an especially strong boost.

Can Higher Gasoline Prices Set Off an Inflationary Spiral?

By Nida Çakır Melek, Francis M. Dillon, and A. Lee Smith

In early 2022, with consumer price inflation already high, a spike in the price of gasoline increased public concerns that the U.S. economy could be in for a repeat of the inflationary spiral that gripped the nation in the 1970s and 1980s. During this period, energy price increases created an environment where rising inflation and rising inflation expectations reinforced one another until a deep economic contraction broke the feedback loop.

Nida Çakır Melek, Francis M. Dillon, and A. Lee Smith assess the risk of a similar spiral in the current environment by exploring whether high inflation makes consumers' inflation expectations more responsive to salient price increases—namely, higher gasoline prices. They find that in response to an increase in the national price of gasoline, individuals with higher initial inflation expectations revise up their one-year-ahead inflation forecasts by almost twice as much as those with lower initial inflation expectations. With inflation currently high and consumers' inflation expectations elevated, their results suggest that changes in salient prices could indeed have an amplified effect on inflation expectations.

The Evolving Role of the Fed's Balance Sheet: Effects and Challenges

By Chaitri Gulati and A. Lee Smith

The Federal Reserve's balance sheet more than doubled in the wake of the COVID-19 pandemic, primarily due to large-scale asset purchases. In 2022, with inflation surging and the labor market tight, the FOMC has started to withdraw policy accommodation and has set in motion a plan to significantly reduce the balance sheet. However, it may be difficult for policymakers to judge how fast and how far to unwind asset purchases, especially as it is unclear exactly how much accommodation the pandemic-era asset purchases have put in place.

In this article, Chaitri Gulati and A. Lee Smith present evidence that the Federal Reserve's expanded balance sheet, with a large portfolio of long-duration assets, has provided a significant amount of policy accommodation in recent years, depressing long-term interest rates by about 1.6 percentage points as of early 2022. They also argue that the FOMC's plan to remove this accommodation through the passive runoff of maturing securities may prove challenging. They project that the downward pressure the balance sheet is currently placing on longer-term interest rates will only gradually reverse.

Hybrid Working, Commuting Time, and the Coming Long-Term Boom in Home Construction

By Jordan Rappaport

Single-family home construction has been sluggish for most of the past two decades, from the housing price collapse in the mid-2000s to the onset of the COVID-19 pandemic in early 2020. Long commuting times between peoples' homes and businesses were an important factor weighing on construction during this period. Land available for new development in the central portion of metropolitan areas and in most older suburbs is limited; although the outermost suburbs have more land for development, many workers have been hesitant to take on the long commutes associated with living in them.

However, the shift to hybrid working since the pandemic has likely altered people's preferences. Working partly at an employer's physical workplace and partly from home hugely cuts time spent commuting, making households potentially willing to live farther from employers. This increased willingness may, in turn, open large swathes of lightly settled land at the peripheries of metropolitan areas for development, thereby relaxing a long-standing constraint on single-family construction.

In this article, I estimate the statistical relationship between single-family home construction and commute duration, finding a significant

---

Jordan Rappaport is a senior economist at the Federal Reserve Bank of Kansas City. Francis Dillon, a research associate at the bank, helped prepare the article. This article is on the bank's website at ***www.KansasCityFed.org***

negative correlation across the largest U.S. metropolitan areas. I then calculate how much time employees commuting from outer suburbs to central business districts (CBDs) save by working remotely two days a week. I find that across metropolitan areas with a population of at least 1 million in 2020, the time saved ranges from 130 to 406 hours per year.

Although home construction is contracting after recent increases in mortgage interest rates, my results suggest that in the long run, the time savings from fewer commutes could almost double single-family home construction in these metropolitan areas from its level just prior to the pandemic, an aggregate increase of 427,000 units per year. The largest metropolitan areas, where commutes have been longest, are likely to see an especially strong boost. For example, construction is predicted to more than triple in the New York, Los Angeles, Chicago, Philadelphia, and Boston metropolitan areas.

Section I describes the sluggish recovery of single-family home construction during the 2010s and documents that it was weakest in metropolitan areas where commutes from the outer suburbs to downtown office districts were longest. Section II calculates the commuting time saved from hybrid working—due both to fewer trips and less traffic congestion—and argues that the reduction will boost single-family home construction most in the outer suburbs of the largest metropolitan areas.

I. Weak Home Construction and Long Commutes

Single-family home construction was weak throughout the 2010s despite strong underlying demand for homes and vigorous growth in numerous other measures of economic activity. The negative correlation between home construction and commuting time from the outer suburbs to the CBDs of large metropolitan areas suggests that households' reluctance to move ever farther from their place of employment was an important cause.

The incomplete recovery of home construction following the Great Recession

Panel A of Chart 1 illustrates the incomplete recovery of home construction following the Great Recession with an index that plots the level of single-family starts relative to its level in 2000, which I use as a benchmark.[1] Although starts moved upward during most of the 2010s, they remained almost 20 percent below their 2000 level at the

Chart 1

Single-Family Home Starts Relative to 2000

Panel A: Single-Family Home Starts, 1989–2022

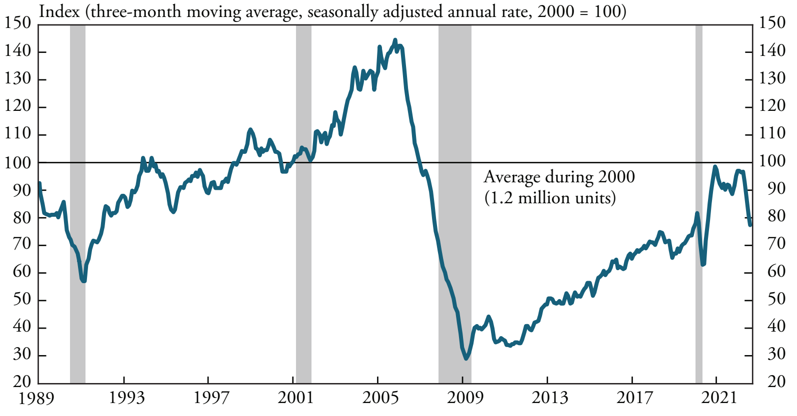

Panel B: Single-Family Home Starts as a Share of U.S. Households, 1959–2022

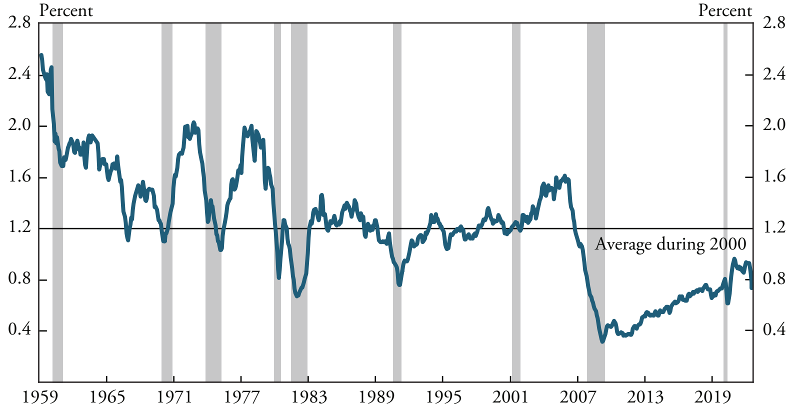

Notes: Data are shown through August 2022. Gray bars denote National Bureau of Economic Research (NBER)-defined recessions.
Sources: U.S. Census Bureau (Haver Analytics), NBER, and author's calculations.

beginning of 2020.[2] During the pandemic, starts surged, briefly touching their benchmark level, but then swung sharply down during the first half of 2022 as mortgage interest rates rose.

Single-family home construction during the 2010s looks even more sluggish considering the increase in the U.S. population during this period. Indeed, Panel B of Chart 1 shows that the ratio of starts to households during the 2010s remained well below its level from the late 1950s until the Great Recession, excepting only lows during the 1981–82 and 1990–91 recessions.

The incomplete recovery of single-family construction during the 2010s contrasted with other measures of economic activity. By the end of 2019, real GDP had risen more than 45 percent above its 2000 level, while real disposable income had risen more than 55 percent above its 2000 level. As is intuitive, the increase in income was accompanied by increases in underlying consumption demand for most goods and services, including single-family homes. However, this increase in underlying demand—the number of homes that households wish to purchase at a fixed level of prices—did not boost single-family home construction.

Instead, increased underlying demand, coupled with the weak recovery in construction, contributed to soaring home prices and subsequently dampened the *actual* number of homes demanded as prices exceeded many households' budgets. Chart 2 shows the sales and rental prices of single-family housing, adjusted for broad-based inflation, since the late 1980s. Real sales prices increased by 40 percentage points during the 2010s, while real rent increased by almost 15 percentage points, leaving both well above their levels in 2000, just prior to the start of the home construction boom and eventual house price bubble that precipitated the Great Recession.

Rising prices and rents considerably depressed the number of households who could afford to live in single-family homes. The black line in Chart 3 represents a demographic-based projection of the number of single-family homes that households would demand if home prices, home rents, incomes, and other non-demographic factors remained close to their levels in 2000.[3] The blue line represents the actual number of occupied single-family homes. During the construction boom in the early 2000s, the number of occupied single-family homes moved modestly above the projection. But beginning in 2010, the number of

Chart 2

Real Price of Single-Family Housing, 1988–2022

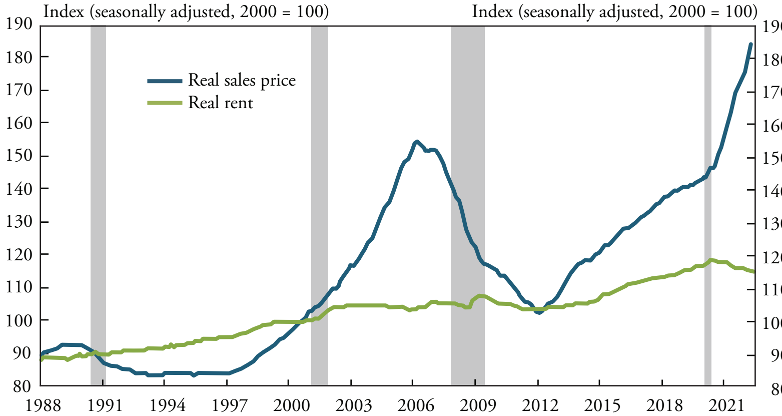

Notes: Data shown are through June 2022. Nominal prices and rents are converted to real using the PCE price index. Gray bars denote NBER-defined recessions.
Sources: Standard and Poor's (Haver Analytics), U.S. Bureau of Labor Statistics (Haver Analytics), U.S. Bureau of Economic Analysis (Haver Analytics), NBER, and author's calculations.

Chart 3

Occupied Single-Family Homes in the 2000s and 2010s, Projected and Actual

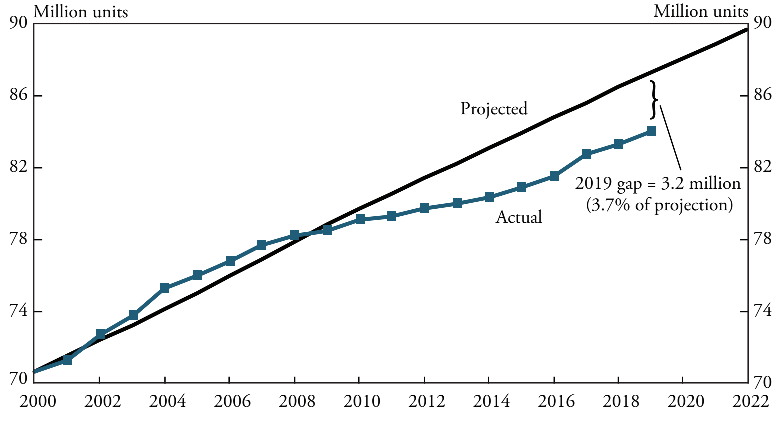

Sources: U.S. Census Bureau (Haver Analytics) and U.S. Bureau of Labor Statistics (Haver Analytics).

occupied single-family homes fell increasingly below it, with the gap eventually widening to 3.2 million units in 2019.

The increasing burden of commuting

Increased commuting times, reflecting worsening traffic congestion and the outward expansion of suburbs, were an important factor preventing builders in the 2010s from responding more vigorously to rising prices and rents. Commuting delays due to traffic congestion have been steadily increasing since the 1980s, especially in the largest metropolitan areas. Chart 4 shows the average annual hours of congestion-related delays experienced by commuters in metropolitan areas of four size categories (Schrank and others 2021).[4] Specifically, the chart measures the hours that traffic congestion adds to the duration of car travel above what would be required if traffic moved at its free-flow speed. The blue line shows that delays became especially punishing in metropolitan areas with a population above 3 million, where the extra time attributable to congestion more than doubled, from 37 hours per year in 1982 to 84 hours per year in 2019.

A different measure of commuting burdens, which captures both the speed of commutes as well as their distance, focuses on the longest commutes rather than average delays. It, too, is positively related to metropolitan population. Chart 5 plots the 95th-percentile commuting time per one-way trip of workers driving alone to the CBD of a metropolitan Core-Based Statistical Area (CBSA) during 2012–16 against the population of that CBSA in 2010.[5] Metropolitan CBSAs, the most commonly used delineation of metropolitan areas, are constructed as combinations of whole U.S. counties based on commuting flows. However, these delineations miss the significant share of workers employed in a CBSA who commute from residences outside it (Humann and Rappaport 2022). To account for this discrepancy, I measure commuting times based on workers driving to a CBD regardless of whether they live in its CBSA. For each CBD, the 95th-percentile commuting time is longer than 95 percent of driving commutes and shorter than 5 percent of driving commutes. As is intuitive, the 95th percentile primarily reflects the driving time of workers living in outer suburbs. Maps in the appendix illustrate the origin location of commutes to the CBDs of the five large metropolitan areas with the longest 95th-percentile commuting times (New York;

Chart 4

Average Delay per Commuter by Metropolitan Area Population, 1982–2019

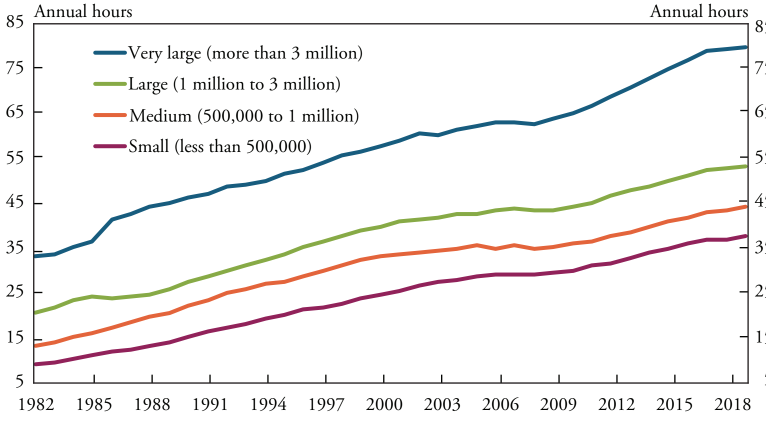

Notes: Metropolitan areas are proxied by Urbanized Areas as delineated by the U.S. Census Bureau following the 1980, 1990, 2000, and 2010 decennial censuses. Metropolitan areas are assigned to a size category based on their population in the contemporary year.
Source: Schrank and others (2021).

Chart 5

95th-Percentile Commuting Time to CBD versus Metropolitan Population

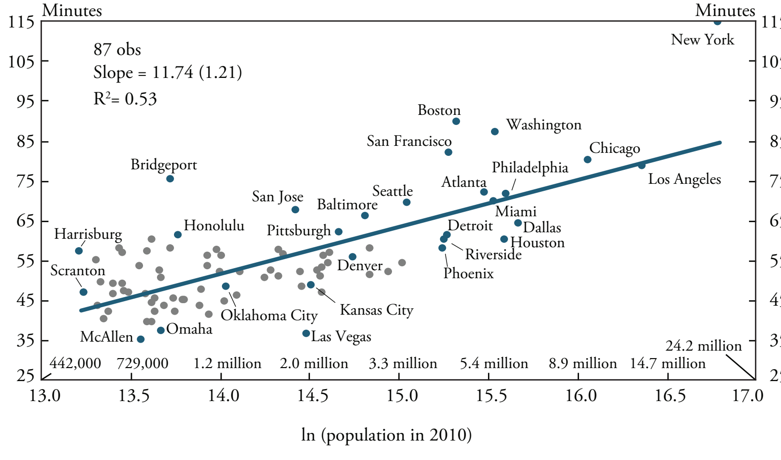

Notes: Travel time is based on workers commuting to a census tract in the CBD, regardless of whether they live in the same CBSA. Population is for medium and large CBSAs with population of at least 500,000 in 2010.
Sources: U.S. Census Bureau and author's calculations.

Boston; Washington, DC; San Francisco; and Chicago). Maps of additional large metropolitan areas with long commuting times are available in a separate online supplement.[6]

The 95th-percentile times vary considerably across metropolitan areas. Across medium CBSAs—those with population between 500,000 and 1 million in 2010—the 95th-percentile commuting time ranged from 35 minutes (for trips to the CBD of McAllen, TX) up to 76 minutes (for trips to the CBD of Bridgeport-Stamford-Norwalk, CT). Across large metropolitan areas—those with population in 2010 above 1 million—the 95th-percentile commuting time ranged from 37 minutes (for trips to the CBD of Las Vegas) up to 115 minutes (for trips to the CBD of New York).

The negative effect of commuting time on home construction

For a large share of metropolitan residents, long commutes contribute to making the outer suburbs a less desirable place to live than places closer to the metropolitan center. Although only a moderate share of jobs may be located in or near the CBD itself, far more workers commute toward the CBD than away from it. For example, across large CBSAs in 2010, the median share of employment located within 10 miles of the CBD was 51 percent but the median share of residents located within 10 miles of the CBD was just 39 percent. Moreover, numerous civic and cultural amenities—such as museums, zoos, live performance venues, and sports teams—are located in the interior of metropolitan areas, further decreasing the desirability of living in outer suburbs. This centralization of employment and amenities constrains the extent to which metropolitan areas can expand outward.[7]

Commuting's negative effect on home construction reflects that constructing single-family homes is typically less expensive in the outer suburbs, where commuting times are longest. In other words, although developers can sell homes at more affordable prices in the outer suburbs than in locations closer to the metropolitan center, relatively few households take advantage of this opportunity due to long commuting times.

Developing homes in the outer suburbs is less costly for several reasons. First, the price of land is less expensive at the periphery of metropolitan areas. Second, zoning near the outermost suburbs is typically less restrictive than in more interior locations. Third, outer suburbs are more likely to have large tracts of open land, which lets builders develop new subdivisions that take advantage of economies of scale, constructing numerous homes in close proximity. Most metropolitan areas, including those with the largest population, have considerable peripheral land that is lightly settled. For example, in all but one of the 56 CBSAs with population above 1 million in 2020, at least half the land area has a population density below 500 persons per square mile, the threshold below which the U.S. Census Bureau classifies land as rural. Even more open land is available just outside the borders of most CBSAs.[8]

Empirically, home construction is negatively correlated with various measures of car commuting time when duration is high but positively correlated with car commuting time when duration is low. Chart 6 plots the average single-family permitting rate from 2015–19 in large metropolitan CBSAs against the 95th-percentile commuting time to their CBD during 2012–16.[9] The blue line shows the estimated statistical relationship between the two. The line has a gradual upward slope until commuting time reaches 58 minutes, indicating that construction and 95th-percentile commuting time are positively related up to this point. As commuting time exceeds 58 minutes, however, the slope of the line turns negative and steepens as the minutes increase further. In other words, the higher the 95th-percentile commuting time to a CBD grows above 58 minutes, the lower that CBSA's predicted single-family permitting rate will be. For example, the chart predicts that the Boston metropolitan area, which has a 95th-percentile commuting time of 90 minutes, will have a single-family permitting rate more than 0.8 percentage points lower than Seattle, which has a 95th-percentile commuting time of 70 minutes.

As long commutes are undesirable, the positively sloped portion of Chart 6 may seem counterintuitive. However, the statistical relationship reflects co-movement (that is, correlation) rather than causality. The undesirability of long commutes, all else equal, contributes to negative co-movement. But all else is unlikely to be equal: residents tend to be willing to endure longer commutes and pay higher house prices in metropolitan areas with more amenities or where higher productivity allows businesses to pay higher wages, both of which are likely to strengthen home construction. Essentially, desirable metropolitan characteristics contribute both to stronger single-family construction

Chart 6
Single-Family Permitting Rate versus 95th-Percentile Commuting Time

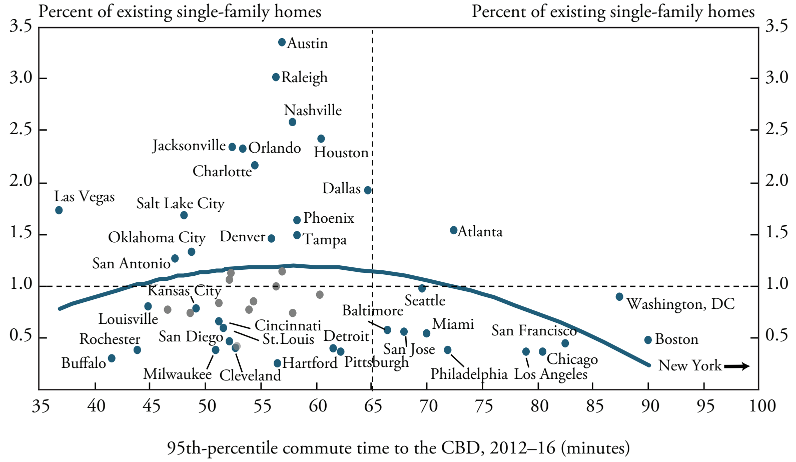

Notes: The blue line corresponds to the estimated statistical relationship between the single-family permitting rate in large metropolitan CBSAs and the 95th-percentile commute duration to their CBD. The single-family permitting rate is measured as the average annual number of construction permits for single-family houses during 2015 through 2019 divided by the average number of single-family houses during 2010 through 2014. The New York CBSA, which is excluded from the estimation, has a 95th-percentile commute duration and single-family permitting rate of 115 minutes and 0.33 percent, respectively.
Sources: U.S. Census Bureau and author's calculations.

and longer commutes, leading them to move together up to a certain threshold. In addition, the increase in population that typically accompanies strong single-family construction often worsens traffic congestion, further contributing to the positive co-movement.

Taking these factors into account, the *causal* relationship of increases in commuting time on single-family construction is likely to be negative at all durations. For example, an increase in a metropolitan area's 95th-percentile commuting time from 40 minutes to 60 minutes, all else equal, will cause single-family permitting to decrease, not increase. Moreover, the causal relationship is likely to be more strongly negative than the co-movement relationship when the co-movement relationship in Chart 6 is negative.[10] For example, an increase in 95th-percentile commuting time from 60 minutes to 90 minutes (again, all else equal) will cause single-family permitting to decrease by more than the decline implied by the co-movement relationship.

More qualitatively, Chart 6 also suggests that long commuting times impede single-family construction. The vertical dashed line separates metropolitan areas with 95th-percentile commuting times above or below 65 minutes, while the horizontal dashed line separates metropolitan areas with permitting rates above or below 1 percent. Among metropolitan areas with a commuting time above 65 minutes, only one has a single-family permitting rate above 1 percent. Among metropolitan areas that have both a 95th-percentile commuting time below 65 minutes and a single-family permitting rate below 1 percent, most have large manufacturing sectors that have been declining since the early 1980s. The low permitting in these areas likely reflects the weak population growth that accompanied industrial decline, both lowering demand for new homes and shortening commuting times. In other words, permitting was lower in these metropolitan areas notwithstanding shorter commutes.

More broadly, population growth is an important mechanism mediating the negative effect of long commutes on housing construction. Just as long 95th-percentile commuting times lower the desirability of living in outer suburbs, long commuting times from residential locations throughout a metropolitan area make living in that metropolitan area less desirable, discouraging people from moving there and so lowering demand for new homes. Consistent with this, Rappaport (2018) documents that the population growth from 2000–17 of the largest metropolitan areas was negatively correlated with their population.

II. Hybrid Working and Future Home Construction

Although long commutes appear to have constrained single-family home construction in the 2010s, the COVID-19 pandemic and associated rise in hybrid working may change this relationship going forward. Since the start of the pandemic, an increasingly large number of businesses have announced plans to permanently adopt hybrid working—that is, letting employees work some of their days in the office and some remotely. The associated time saved commuting is likely to increase workers' willingness to live farther from their workplaces, spurring single-family construction in and beyond the current outer suburbs of large metropolitan areas. Some businesses have also announced plans to let a portion of their employees work remotely full time, potentially freeing them to move anywhere and thus spurring home construction in locations throughout the United States. But if these workers must come into the office even a handful of times per year, or come into the office on short notice, they will still probably need to live no more than a few hours' drive away. In this sense, hybrid working loosens the tether between where someone lives and works but does not break it.[11]

Time saved commuting

One of the largest benefits of hybrid working is reduced time spent commuting, a function of both fewer weekly trips and faster driving speed due to reduced traffic congestion. The amount of time saved will depend on the number of days worked remotely, which will likely vary considerably across employees. For illustrative purposes, Chart 7 reports predictions of the annual time saved from working remotely two days per week for a worker with a pre-pandemic commuting time equal to the 95th-percentile duration to the CBD in 2012–16. The 10 large metropolitan areas shown are those with the longest 95th-percentile durations. The predictions assume that the share of employees working remotely is about the same on each day of the workweek, causing a long-run increase in speed on trips to the CBD in each metropolitan area equal to one-half the increase in speed on all trips in the same metropolitan area from 2019 to 2020 (when the pandemic significantly reduced traffic volumes).[12] The annual commuting time saved from hybrid working ranges from 247 hours for workers driving to the Philadelphia CBD to 406 hours for workers driving to the New York CBD, with most of the time savings coming from making fewer trips rather than making trips at a faster speed.[13]

More generally, the predicted time saved is positively correlated with metropolitan population. The average predicted time saved each year is 254 hours in CBSAs with a population of at least 3 million in 2020, 180 hours in CBSAs with a population between 1 and 3 million, and 161 hours in CBSAs with a population between 500,000 and 1 million.

Three factors contribute to this positive correlation between predicted time saved and metropolitan population. First, commuting duration is positively correlated with population, implying that each skipped commute saves more time in a large metropolitan area than a small one. Second, the distance of long-duration commutes to the CBD tends to be greater in large metropolitan areas, implying that identical speed increases will save more time in large metropolitan areas than small ones.[14] Third, decreases in driving volume increase speed by

Chart 7
Commuting Time Saved Working Remotely Two Days per Week

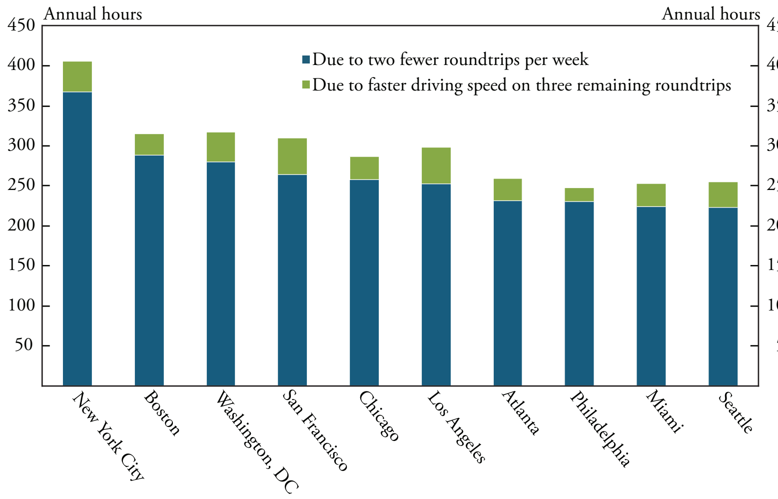

Note: Chart shows annual time saved from working two days per week from home for a worker with pre-pandemic commuting duration equal to the 95th-percentile time to the CBD of the enumerated metropolitan area.
Source: Author's calculations.

proportionally more where congestion is worst, which tends to be in large metropolitan areas. For example, a 10 percent reduction in traffic volume on an interstate highway segment where traffic had previously slowed speed to 45 miles per hour would cut driving time by about 20 percent. The same 10 percent reduction in traffic volume on an interstate segment where traffic had previously slowed speed to 25 miles per hour would cut driving time by 40 percent.[15]

As is intuitive, the relaxed burden from commuting is likely to make workers willing to live farther from their place of employment. Consistent with this hypothesis, Delventhal and Parkhomenko (2021) estimate that the 90th-percentile distance of commutes by full-time workers increases from 23 miles for those who do not regularly work from home, to 26 miles for those who work from home one day per week, to up to 32 miles for those who work from home two days per week.[16]

The disproportionate boost to single-family home construction in large metropolitan areas

To estimate how much the predicted time savings from commuting will boost single-family home construction, I make three benchmark assumptions. First, I assume that the effect of commuting time on single-family permitting in a CBSA is equal to the magnitude of the negative correlation between permitting and the 95th-percentile commuting time to its CBD when the commute duration is 75 minutes.[17] Under this assumption, a 10-minute increase in the 95th-percentile commuting time to the CBD causes the single-family permitting rate to decline by 0.32 percentage points. This assumption is arguably conservative given the factors that offset negative co-movement of the permits rate and commuting time.

Second, I assume that the annual hours an individual saves from less frequent and faster commutes have the same effect as an equivalent decrease in the duration of their one-way commute. For example, I assume that the predicted 406 annual hours saved for workers in the New York CBSA who make two fewer roundtrip commutes per week will have the same effect on permitting as a 51-minute decrease in individuals' one-way commuting time.

Third, I assume that a significant share of employees will be allowed and will choose to work remotely two days per week. Of course, the frequency of working remotely will vary considerably across occupations, companies, and individual employees. But only a moderate share of households will need to increase their willingness to live farther from their employer for single-family construction to boom (see endnote 10).

Given these benchmark assumptions, I predict that reduced commuting times will eventually boost aggregate single-family permits in the 56 CBSAs with a population of at least 1 million in 2020 by 427,000 per year, increasing single-family construction in these CBSAs by 92 percent above its level in 2019 and increasing national single-family construction by 49 percent above its level in 2019. Based on this benchmark, national single-family permits will eventually rise to a long-term annual rate of 1.4 million.[18] For some metropolitan areas, the boost may take place primarily outside the formally delineated CBSA.

Hybrid working is also likely to increase single-family construction for reasons only indirectly related to decreased commuting, further boosting permitting's long-term annual rate. For example, hybrid working is likely to increase demand for larger residences to accommodate home offices. Although some households may be able to accommodate this need by renovating their current home, many others will move to new homes, further driving up single-family construction.[19] In addition, many households are likely to purchase second homes, where hybrid working will allow them to spend more time than previously.

The predicted long-term increase in single-family construction is skewed toward the largest metropolitan areas. Table 1 shows that permits are predicted to increase by more than 25,000 per year in the New York, Los Angeles, and Chicago CBSAs, and by at least 9,500 per year in each of the remaining CBSAs with a population of at least 4 million in 2020. In contrast, permits are predicted to increase by no more than 3,100 per year in each of the CBSAs with a population between 1 and 1.5 million. One reason for this positive skew is that the predicted increases in the permitting rate are assumed to be proportional to time saved commuting, which is highest in the largest metropolitan areas. In essence, the switch to hybrid working boosts the desirability of living in metropolitan areas the most where one-way commuting times are longest.[20] Another reason is that for a given increase in the predicted permitting *rate*, the predicted increase in permits is directly proportional to the number of existing homes, which of course is higher in the largest metropolitan areas.

Partly offsetting the positive skew, employees who have the option of working remotely full time, without any geographic tether to their employer's office, may choose to move to smaller locations to take advantage of lower home prices there. Other employees who can work remotely full time may choose to move closer to family or to places with more amenities such as mountains and nice weather (Rappaport 2018). In these cases, full-time remote working will shift some home construction from the largest metropolitan areas to smaller locations. But these shifts are unlikely to affect the overall strength of the construction boom, reflecting that relocating workers will increase demand for new homes in the locations to which they move.

In addition, the predicted increases do not account for geographic constraints on home construction. For example, the predicted large increases in permits in the Los Angeles, San Francisco, Seattle, and Miami metropolitan areas will likely be tempered by extensive mountainous terrain and wetlands. Limited water resources may also temper home construction, especially in the West and Southwest.

More recently, the sharp increase in mortgage interest rates since the start of 2022 has pushed single-family home construction into a

Table 1

Predicted Increases in Single-Family Permits due to Less Time Spent Commuting

Population rank	Metropolitan area with population in 2020 of at least 1 million	2020 population	2015–19 annual permits	Predicted long-term increase in annual permits
1	New York-Newark-Jersey City, NY-NJ-PA	20,838,000	11,200	54,900
2	Los Angeles-Long Beach-Anaheim, CA	13,201,000	9,600	30,800
3	Chicago-Naperville-Elgin, IL-IN-WI	9,619,000	8,100	25,700
4	Dallas-Fort Worth-Arlington, TX	7,708,000	33,100	15,200
5	Houston-The Woodlands-Sugar Land, TX	7,122,000	37,700	13,200
6	Washington-Arlington-Alexandria, DC-VA-MD-WV	6,371,000	13,400	18,900
7	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	6,245,000	6,900	17,600
8	Miami-Fort Lauderdale-West Palm Beach, FL	6,138,000	7,000	12,900
9	Atlanta-Sandy Springs-Roswell, GA	6,090,000	24,200	16,200
10	Boston-Cambridge-Newton, MA-NH	4,942,000	4,900	12,700
11	Phoenix-Mesa-Scottsdale, AZ	4,846,000	20,800	10,600
12	San Francisco-Oakland-Hayward, CA	4,749,000	4,500	12,700
13	Riverside-San Bernardino-Ontario, CA	4,600,000	10,100	9,800
14	Detroit-Warren-Dearborn, MI	4,392,000	5,800	12,100
15	Seattle-Tacoma-Bellevue, WA	4,019,000	9,200	9,500
16	Minneapolis-St. Paul-Bloomington, MN-WI	3,705,000	8,400	7,600
17	San Diego-Carlsbad, CA	3,299,000	3,200	5,400
18	Tampa-St. Petersburg-Clearwater, FL	3,175,000	12,400	6,800
19	Denver-Aurora-Lakewood, CO	2,964,000	10,700	5,900
20	Baltimore-Columbia-Towson, MD	2,845,000	4,900	8,000
:				
50	Buffalo-Cheektowaga-Niagara Falls, NY	1,167,000	1,000	1,800
51	Rochester, NY	1,090,000	1,200	1,900
52	Grand Rapids-Wyoming, MI	1,083,000	2,700	1,800
53	Tucson, AZ	1,043,000	2,900	1,800
55	Tulsa, OK	1,015,000	3,100	1,700
56	Fresno, CA	1,009,000	2,200	1,300
All 56 CBSAs with population in 2020 of at least 1 million		189,976,000	435,000	427,000

Notes: Predicted increases are based on the predicted annual time savings from working from home two days per week for a worker with a pre-pandemic commute duration equal to the 95th-percentile duration of workers commuting to the CBD of the enumerated metropolitan area during 2012–16. Metropolitan areas are the CBSAs delineated by the Office of Management and Budget in 2013 with population of at least 1 million in 2020. Additional statistics for CBSAs not shown above can be found in the appendix.
Sources: U.S. Census Bureau and author's calculations.

cyclical downturn. In particular, the increase in interest rates paired with the sharp run-up in home prices over the past few years is considerably straining the affordability of monthly mortgage payments on home purchases. To the extent that mortgage interest rates remain near or above their level in mid-2022, builders may need to adjust their development strategies—for example, by shifting their product mix toward homes that are less expensive to construct or by employing new production techniques, such as modular construction, which can significantly cut costs (Bertram and others 2019). More generally, home construction has always fluctuated cyclically and is likely to continue to do so.

When single-family construction begins to rebound, supply constraints are likely to slow its climb to its predicted long-term rate. Shortages of workers, construction materials, and ready-to-build lots are all likely to constrain the growth of single-family construction in the short term. Proportionately scaling up employment to match the predicted increase in single-family construction to 1.4 million units per year would require developers to hire 1 million more construction workers than were employed in mid-2022. The production of lumber and other construction materials would also need to increase significantly. Moreover, developing new subdivisions is a drawn-out process, typically taking several years from the conception of a project to breaking ground on new homes. Ramping up will be even more challenging in many metropolitan areas. For example, annual single-family permits are predicted to more than triple in the New York, Chicago, Philadelphia, and Boston metropolitan areas.

Once single-family home construction ramps up, it is likely to remain high for many years. A widespread change in where people want to live can be accommodated only gradually. For example, the shifts in preferences toward living in suburbs and in the Sunbelt induced transitions that played out over decades (Rappaport 2004, 2007). The coming accumulation of booming home construction and the associated outward expansion of metropolitan settlement will require numerous complementary processes. Roads and utility capacity will need to be expanded, and new schools built and staffed. Similarly, growing neighborhoods at the periphery of metropolitan areas will need to attract new retailers, restaurants, and local service providers. These processes, which each take considerable time on their own, must proceed in rough concurrence. Only so much home construction can accumulate before the associated congestion dampens demand.

Conclusions

Hybrid working dramatically cuts time spent commuting and so is likely to make workers willing to live farther from their place of employment. This greater willingness, in turn, is likely to boost single-family construction, especially in the outer suburbs of the largest metropolitan areas, where the time saved from reduced commuting is highest. Based on an estimate of the relationship between commuting time and single-family home permitting and some benchmark assumptions, I predict that hybrid work and the associated reduction in commuting time will eventually boost aggregate single-family permits in large CBSAs by a total of 427,000 per year—a 92 percent increase from its aggregate level in large CBSAs in 2019 and a 49 percent increase above the national rate of single-family permitting just prior to the COVID-19 pandemic. Several additional considerations are likely to boost national single-family construction by even more, including increased demand for homes in locations just outside large CBSAs, increased demand for space to accommodate home offices, and increased demand for second homes.

The skew in the predicted home construction boom toward the largest metropolitan areas may be partly offset by a rise in full-time remote working. Households that are no longer tethered to employer offices can choose to move anywhere, boosting home construction in locations throughout the United States. Their choices are likely to depend on three considerations: the presence of family, friends, and social networks; proximity to amenities; and the affordability of homes. It may take some time before trends emerge on the specific locations that attract these newly mobile households.

Appendix

Table A-1
Selected Statistics

Population rank	Metropolitan statistical area	Population in 2020	Driving time to CBD: 95th percentile, 2012–16 (minutes)	Time saved working remotely two days per week (hours)	Annual permits, 2015–19 (average)	Predicted annual permits	Predicted annual permit increase
1	New York-Newark-Jersey City, NY-NJ-PA	20,838,000	114.8	406	11,200	66,100	54,900
2	Los Angeles-Long Beach-Anaheim, CA	13,201,000	79.0	298	9,600	40,400	30,800
3	Chicago-Naperville-Elgin, IL-IN-WI	9,619,000	80.5	286	8,100	33,700	25,700
4	Dallas-Fort Worth-Arlington, TX	7,708,000	64.7	223	33,100	48,300	15,200
5	Houston-The Woodlands-Sugar Land, TX	7,122,000	60.5	214	37,700	50,900	13,200
6	Washington-Arlington-Alexandria, DC-VA-MD-WV	6,371,000	87.5	317	13,400	32,200	18,900
7	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD	6,245,000	72.0	247	6,900	24,600	17,600
8	Miami-Fort Lauderdale-West Palm Beach, FL	6,138,000	70.0	253	7,000	19,800	12,900
9	Atlanta-Sandy Springs-Roswell, GA	6,090,000	72.5	259	24,200	40,400	16,200
10	Boston-Cambridge-Newton, MA-NH	4,942,000	90.0	315	4,900	17,600	12,700
11	Phoenix-Mesa-Scottsdale, AZ	4,846,000	58.3	209	20,800	31,400	10,600
12	San Francisco-Oakland-Hayward, CA	4,749,000	82.5	310	4,500	17,200	12,700
13	Riverside-San Bernardino-Ontario, CA	4,600,000	60.5	221	10,100	19,900	9,800
14	Detroit-Warren-Dearborn, MI	4,392,000	61.6	210	5,800	17,900	12,100
15	Seattle-Tacoma-Bellevue, WA	4,019,000	69.7	255	9,200	18,700	9,500
16	Minneapolis-St. Paul-Bloomington, MN-WI	3,705,000	54.5	190	8,400	16,000	7,600
17	San Diego-Carlsbad, CA	3,299,000	52.3	190	3,200	8,600	5,400
18	Tampa-St. Petersburg-Clearwater, FL	3,175,000	58.3	206	12,400	19,200	6,800

Table A-1 (continued)

Population rank	Metropolitan statistical area	Population in 2020	Driving time to CBD: 95th percentile, 2012–16 (minutes)	Time saved working remotely two days per week (hours)	Annual permits, 2015–19 (average)	Predicted annual permits	Predicted annual permit increase
19	Denver-Aurora-Lakewood, CO	2,964,000	56.0	203	10,700	16,600	5,900
20	Baltimore-Columbia-Towson, MD	2,845,000	66.5	237	4,900	12,900	8,000
21	St. Louis, MO-IL	2,820,000	51.7	172	5,300	11,500	6,200
22	Orlando-Kissimmee-Sanford, FL	2,673,000	53.4	188	14,400	19,100	4,600
23	Charlotte-Concord-Gastonia, NC-SC	2,638,000	54.5	192	14,900	20,100	5,200
24	San Antonio-New Braunfels, TX	2,558,000	47.3	162	7,500	11,400	3,800
25	Portland-Vancouver-Hillsboro, OR-WA	2,513,000	57.1	208	7,100	12,200	5,200
26	Sacramento-Roseville-Arden-Arcade, CA	2,397,000	56.5	198	6,300	11,400	5,000
27	Pittsburgh, PA	2,371,000	62.3	212	3,000	10,000	7,000
28	Austin-Round Rock, TX	2,283,000	57.0	205	15,400	19,100	3,700
29	Las Vegas-Henderson-Paradise, NV	2,265,000	36.8	130	9,200	12,000	2,800
30	Cincinnati, OH-KY-IN	2,234,000	51.3	176	4,100	8,600	4,500
31	Kansas City, MO-KS	2,192,000	49.2	164	5,200	9,600	4,300
32	Columbus, OH	2,139,000	48.8	166	4,200	8,000	3,800
33	Indianapolis-Carmel-Anderson, IN	2,111,000	52.3	180	6,400	10,800	4,400
34	Cleveland-Elyria, OH	2,088,000	52.8	176	2,600	7,300	4,700
35	Nashville-Davidson-Murfreesboro-Franklin, TN	2,014,000	58.0	205	13,200	17,300	4,100
36	San Jose-Sunnyvale-Santa Clara, CA	2,000,000	68.0	254	2,300	6,600	4,300
37	Virginia Beach-Norfolk-Newport News, VA-NC	1,763,000	51.3	175	4,200	7,700	3,500
38	Providence-Warwick, RI-MA	1,677,000	53.0	173	1,600	4,300	2,700
39	Jacksonville, FL	1,606,000	52.5	184	9,600	12,600	3,000
40	Milwaukee-Waukesha-West Allis, WI	1,575,000	51.0	173	1,600	4,400	2,800
41	Oklahoma City, OK	1,426,000	48.8	164	5,400	8,100	2,600

Table A-1 (continued)

Population rank	Metropolitan statistical area	Population in 2020	Driving time to CBD: 95th percentile, 2012–16 (minutes)	Time saved working remotely two days per week (hours)	Annual permits, 2015–19 (average)	Predicted annual permits	Predicted annual permit increase
42	Raleigh, NC	1,414,000	56.4	194	10,200	12,900	2,600
43	Memphis, TN-MS-AR	1,345,000	46.7	159	3,100	5,600	2,600
44	Richmond, VA	1,339,000	52.5	174	4,300	6,900	2,700
45	Louisville/Jefferson County, KY-IN	1,318,000	44.9	155	3,100	5,500	2,400
46	New Orleans-Metairie, LA	1,272,000	58.0	209	2,700	5,800	3,100
47	Salt Lake City, UT	1,258,000	48.1	165	4,600	6,400	1,800
48	Hartford-West Hartford-East Hartford, CT	1,214,000	56.6	193	800	3,300	2,500
49	Birmingham-Hoover, AL	1,181,000	54.1	186	2,700	5,400	2,700
50	Buffalo-Cheektowaga-Niagara Falls, NY	1,167,000	41.6	140	1,000	2,800	1,800
51	Rochester, NY	1,090,000	44.0	147	1,200	3,200	1,900
52	Grand Rapids-Wyoming, MI	1,083,000	45.3	150	2,700	4,500	1,800
53	Tucson, AZ	1,043,000	45.4	157	2,900	4,700	1,800
54	Urban Honolulu, HI	1,017,000	61.7	230	1,000	2,700	1,700
55	Tulsa, OK	1,015,000	42.4	140	3,100	4,800	1,700
56	Fresno, CA	1,009,000	45.8	149	2,200	3,500	1,300
57	Worcester, MA-CT	979,000	59	191	1,300	3,100	1,800
58	Omaha-Council Bluffs, NE-IA	968,000	38	131	2,900	4,300	1,400
59	Bridgeport-Stamford-Norwalk, CT	957,000	76	264	700	3,200	2,400
60	Greenville-Anderson-Mauldin, SC	928,000	45	149	4,600	6,100	1,500
61	Albuquerque, NM	917,000	44	151	2,000	3,600	1,600
62	Bakersfield, CA	909,000	43	145	2,200	3,400	1,200
63	Knoxville, TN	903,000	46	155	3,100	4,800	1,700
64	Albany-Schenectady-Troy, NY	899,000	51	167	1,200	2,800	1,600
65	McAllen-Edinburg-Mission, TX	871,000	35	117	3,000	3,800	800
66	El Paso, TX	869,000	40	130	2,300	3,300	1,000
67	New Haven-Milford, CT	865,000	53	175	400	1,900	1,500

Table A-1 (continued)

Population rank	Metropolitan statistical area	Population in 2020	Driving time to CBD: 95th percentile, 2012–16 (minutes)	Time saved working remotely two days per week (hours)	Annual permits, 2015–19 (average)	Predicted annual permits	Predicted annual permit increase
68	Allentown-Bethlehem-Easton, PA-NJ	862,000	61	202	1,000	3,100	2,100
69	Baton Rouge, LA	850,000	58	203	3,500	5,300	1,800
70	Oxnard-Thousand Oaks-Ventura, CA	844,000	40	136	700	1,900	1,100
71	North Port-Sarasota-Bradenton, FL	834,000	50	171	5,900	7,700	1,700
72	Columbia, SC	829,000	54	184	4,100	5,700	1,600
73	Dayton, OH	814,000	47	154	1,000	2,700	1,700
74	Charleston-North Charleston, SC	800,000	50	175	4,700	6,100	1,400
75	Stockton-Lodi, CA	779,000	58	195	2,200	3,700	1,400
76	Greensboro-High Point, NC	777,000	47	154	1,900	3,200	1,400
77	Boise City, ID	765,000	40	140	6,100	7,200	1,100
78	Cape Coral-Fort Myers, FL	761,000	48	166	4,800	6,400	1,500
79	Colorado Springs, CO	755,000	43	143	3,800	4,900	1,100
80	Little Rock-North Little Rock-Conway, AR	748,000	57	188	1,800	3,400	1,600
81	Lakeland-Winter Haven, FL	725,000	55	181	4,500	5,800	1,300
82	Akron, OH	702,000	48	157	900	2,400	1,500
83	Ogden-Clearfield, UT	695,000	35	116	2,600	3,300	800
84	Madison, WI	681,000	56	190	1,600	2,900	1,300
85	Winston-Salem, NC	676,000	45	150	2,600	3,800	1,200
86	Des Moines-West Des Moines, IA	672,000	42	138	3,600	4,600	1,000
87	Provo-Orem, UT	671,000	42	142	4,700	5,400	700
88	Deltona-Daytona Beach-Ormond Beach, FL	669,000	47	154	2,800	4,200	1,300
89	Syracuse, NY	662,000	47	151	600	1,700	1,200
90	Wichita, KS	655,000	41	134	1,300	2,400	1,100
91	Springfield, MA	628,000	50	164	500	1,500	1,000
92	Augusta-Richmond County, GA-SC	611,000	46	150	2,600	3,700	1,000
93	Palm Bay-Melbourne-Titusville, FL	607,000	46	151	2,100	3,300	1,100

Table A-1 (continued)

Population rank	Metropolitan statistical area	Population in 2020	Driving time to CBD: 95th percentile, 2012–16 (minutes)	Time saved working remotely two days per week (hours)	Annual permits, 2015–19 (average)	Predicted annual permits	Predicted annual permit increase
94	Toledo, OH	606,000	44	146	700	1,900	1,100
95	Spokane-Spokane Valley, WA	599,000	38	129	1,800	2,600	800
96	Harrisburg-Carlisle, PA	592,000	58	188	1,200	2,600	1,300
97	Durham-Chapel Hill, NC	589,000	54	182	3,100	4,200	1,100
98	Jackson, MS	575,000	51	169	1,500	2,600	1,100
99	Fayetteville-Springdale-Rogers, AR-MO	570,000	46	156	3,600	4,500	900
100	Scranton-Wilkes-Barre-Hazleton, PA	568,000	47	154	600	1,800	1,200
101	Chattanooga, TN-GA	563,000	49	163	1,800	2,900	1,100
102	Lancaster, PA	553,000	49	156	1,000	1,900	1,000
103	Modesto, CA	553,000	48	158	600	1,500	900
104	Portland-South Portland, ME	552,000	55	176	1,800	3,100	1,300
105	Youngstown-Warren-Boardman, OH-PA	541,000	47	154	300	1,600	1,200
106	Lexington-Fayette, KY	517,000	51	174	1,400	2,500	1,000
107	Pensacola-Ferry Pass-Brent, FL	510,000	45	153	2,400	3,300	900

Sources: U.S. Census Bureau and author's calculations.

Map A-1

Commuting Time to Central Business District: New York

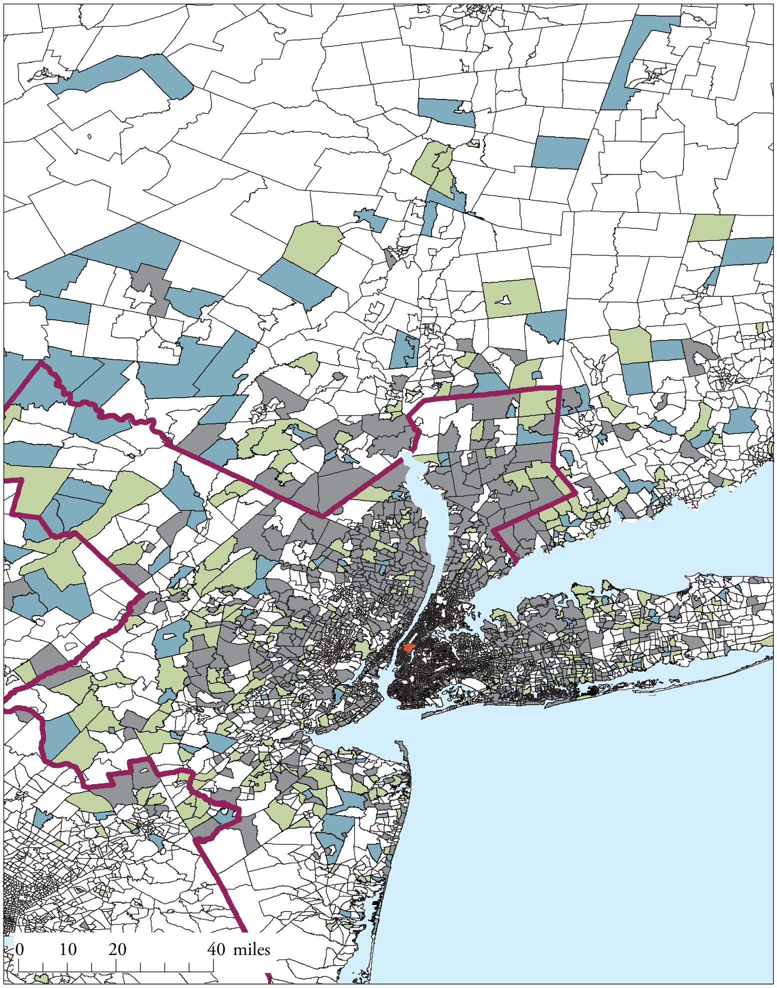

New York-Northern New Jersey-Long Island CBSA

Central business district (CBD)

Insufficient data

Duration <80th percentile (76.3 minutes)

Duration ≥80th, <95th percentile (114.8 minutes)

Duration ≥95th percentile (99th = 169.3 minutes)

Notes: Map shows percentile commuting times during 2012–16 of workers driving to the CBD of the New York metropolitan area. Census tracts with large land area have low population density.
Sources: U.S. Census Bureau and author's calculations.

Map A-2

Commuting Time to Central Business District: Boston

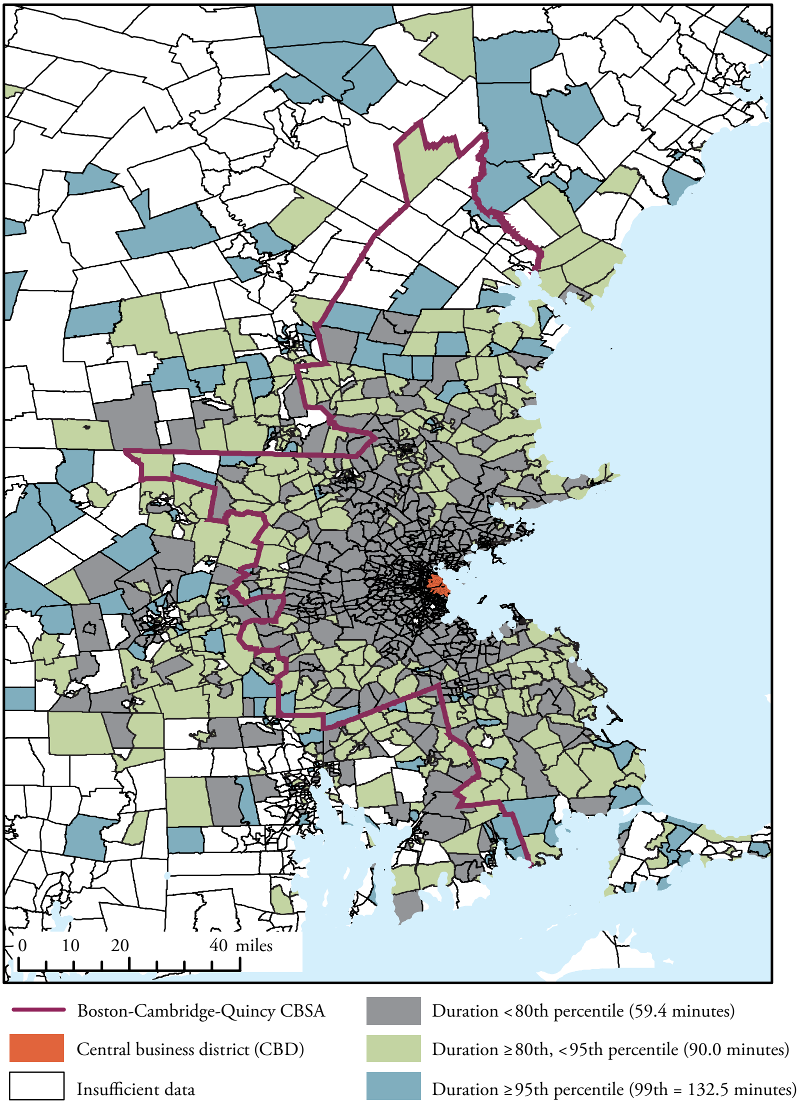

Notes: Map shows percentile commuting times during 2012–16 of workers driving to the CBD of the Boston metropolitan area. Census tracts with large land area have low population density.
Sources: U.S. Census Bureau and author's calculations.

Map A-3

Commuting Time to Central Business District: Washington, DC

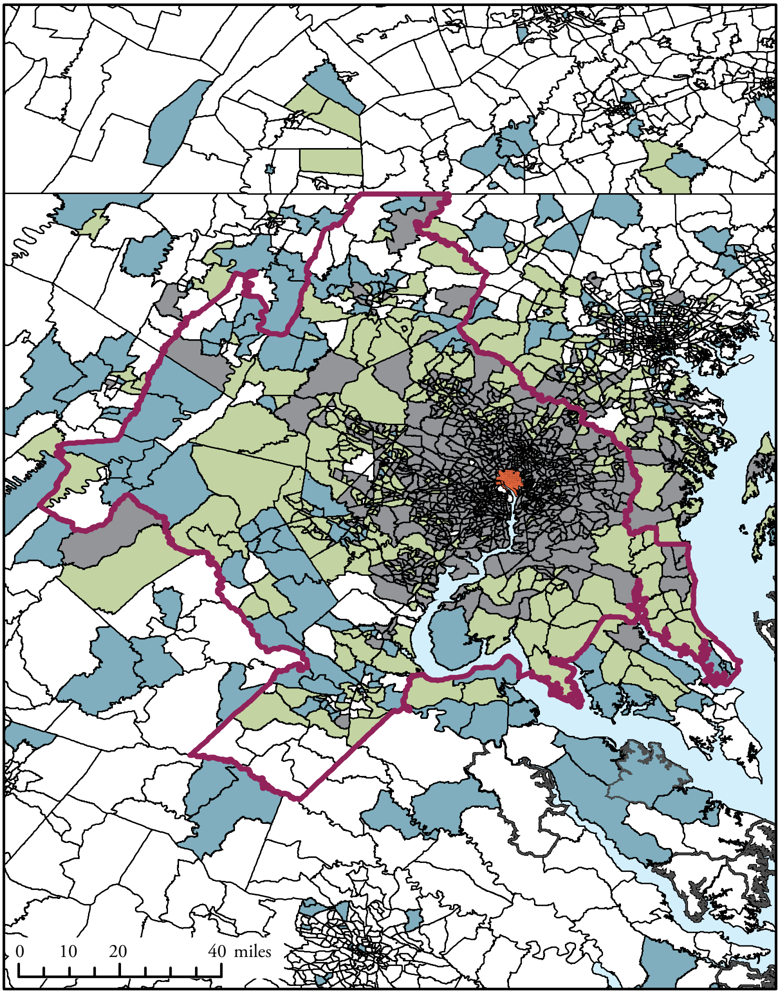

Washington-Arlington-Alexandria CBSA
Central business district (CBD)
Insufficient data
Duration < 80th percentile (56.3 minutes)
Duration ≥ 80th, < 95th percentile (87.5 minutes)
Duration ≥ 95th percentile (99th = 128.3 minutes)

Notes: Map shows percentile commuting times during 2012–16 of workers driving to the CBD of the Washington, DC metropolitan area. Census tracts with large land area have low population density.
Sources: U.S. Census Bureau and author's calculations.

Map A-4
Commuting Time to Central Business District: San Francisco

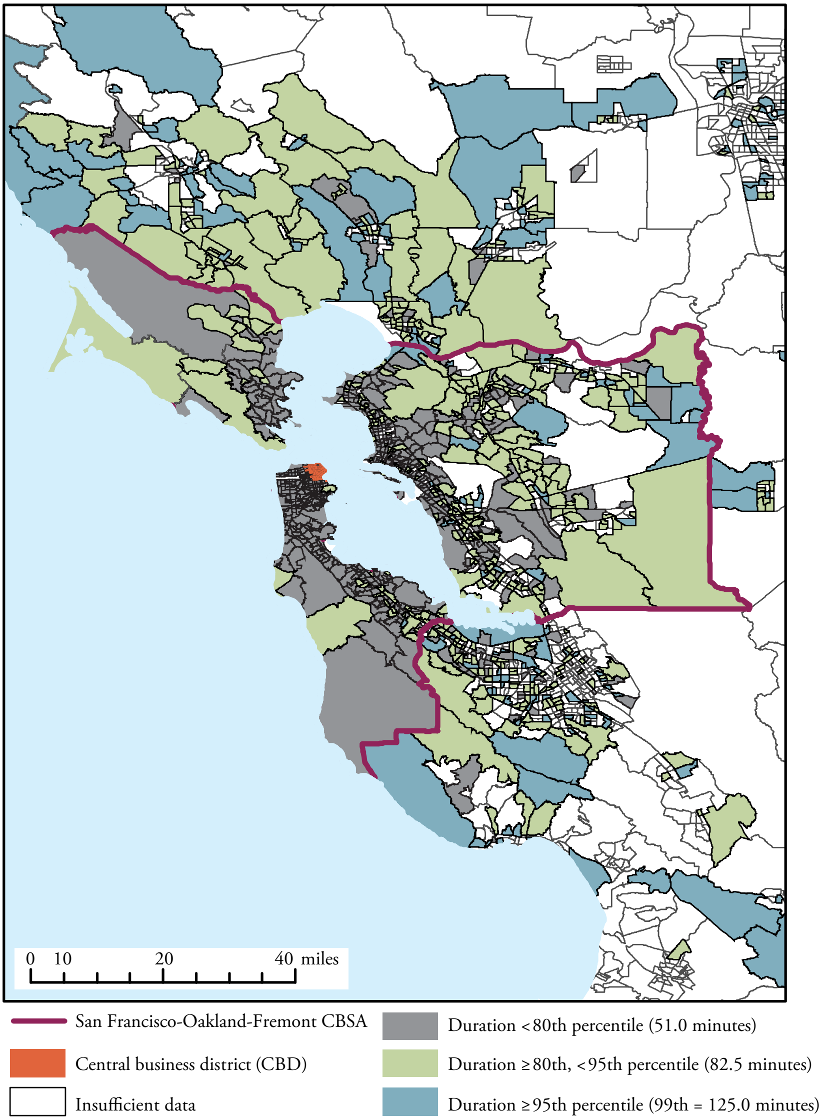

Notes: Map shows percentile commuting times during 2012–16 of workers driving to the CBD of the San Francisco metropolitan area. Census tracts with large land area have low population density.
Sources: U.S. Census Bureau and author's calculations.

Map A-5

Commuting Time to Central Business District: Chicago

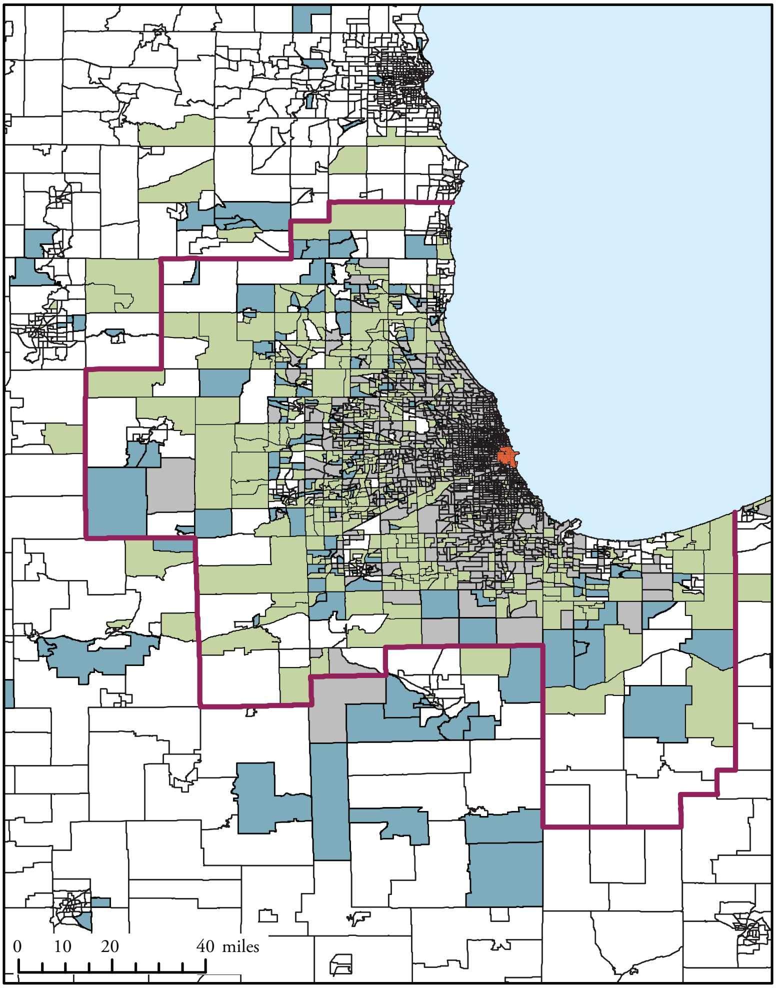

Chicago-Naperville-Joliet CBSA

Central business district (CBD)

Insufficient data

Duration < 80th percentile (53.0 minutes)

Duration ≥ 80th, < 95th percentile (80.5 minutes)

Duration ≥ 95th percentile (99th = 122.5 minutes)

Notes: Map shows percentile commuting times during 2012–16 of workers driving to the CBD of the Chicago metropolitan area. Census tracts with large land area have low population density.
Sources: U.S. Census Bureau and author's calculations.

Endnotes

[1]My choice to use 2000 as a benchmark primarily reflects that it was a year just prior to a recession and subsequent boom in home construction, which proved unsustainable. In addition, single-family home starts stayed near their level in 2000 for an extended period. Measured as a 12-month moving average, single family starts remained continuously within 10 percentage points of their 2000 benchmark from May 1998 to November 2002. And as illustrated in Panel B of Chart 1, the ratio of single-family home starts to the number of U.S. households stayed near its level in 2000 throughout most of the 1980s and 1990s.

[2]Although an overhang of unsold, newly constructed homes contributed to the crash in construction that started in 2006, this excess inventory had been more than worked off by the end of 2009, at which point the inventory of unsold new homes dropped to its lowest level since 1971.

[3]The projection uses the share of each five-year age range—20–24, 25–29, …, 85–89, 90 and older—that lived in a single-family home in 2000 to calculate the number of occupied single-family homes there would be if individuals in the same age range made the same housing choices in subsequent years as the individuals in that age range made in 2000 (Rappaport 2013). For example, 13.5 percent of women age 25–29 in 2000 lived with one other person in a single-family home in 2000. Thus, the projection assumes that this 2000 trend continued for women age 25–29 in subsequent years. As described in the main text, the projection also assumes that the relative price of single-family housing services remained at its 2000 level in subsequent years. Another critical assumption is that the relative purchasing power of a household's disposable income remained at its 2000 level. Additional assumptions include that people's preferences remained unchanged and that the stock of single-family homes in any year had similar characteristics, including the commuting time associated with location, as in 2000.

[4]Schrank and others (2021) report delays for Urbanized Areas (UAs), which are the densely settled core of metropolitan Core-Based Statistical Areas (CBSAs), a more commonly used delineation of metropolitan areas. I classify the size of each UA in each year by its population in that year. A similar chart in Schrank and others (2021) classifies UAs in all years by their population in 2020.

[5]Except for New York, I designate a tract as belonging to the CBD of a metropolitan area if it is located within two miles of the tract with the highest employment density in a metropolitan CBSA, using the CBSA delineations promulgated in Office of Management and Budget (2013) and employment reported in the Longitudinal Employer-Household Dynamics Origin-Destination Employment Statistics (LODES) for 2012. For the New York metropolitan area, the LODES data locates most employees of the New York City public school system in the tract of its Brooklyn headquarters, falsely making that tract appear to have the highest employment density in the metropolitan area. Instead, I anchor the

New York CBD by Grand Central Station and consider all tracts within two miles of it to belong to the CBD. I calculate the distribution of commuting times using the number of drivers and mean driving time for each origin-destination pair of tracts with a destination place of employment in the CBD and an origin place of residence within 150 miles of it, excluding trips with passengers. The number of drivers and mean travel time for each pairwise flow are for 2012–16 as reported in the Census Transportation Planning Product for those years.

[6]To access supplemental maps, visit https://doi.org/10.18651/ER/v107n-4Rappaport.

[7]Humann and Rappaport (2022) find that the elasticity of settled metropolitan land area with respect to metropolitan population is increasingly below as metropolitan population rises, suggesting centripetal forces limit metropolitan expansion. These centripetal forces would likely weaken if employment could expand outward in proportion to population expanding outward. One reason employers choose not to move outward is the need to attract workers from throughout a metropolitan area. Another reason is that businesses in some industries and occupations are more productive in proximity to each other, discouraging individual companies from moving away from an existing cluster. Consistent with this, employment in some occupations—including law, finance, information technology, science, and media—is considerably more centralized than employment in other occupations (Brown and others 2017). In addition, urban amenities—such as restaurants, cafés, live performance venues, and pedestrian retail districts—are a key force pulling young college graduates to live near metropolitan centers, incentivizing employers to locate nearby (Smart Growth America 2015; Couture and Handbury 2020).

[8]To be sure, in some metropolitan areas, a considerable portion of lightly settled land may be unsuitable for development—for example, due to mountainous terrain and wetlands (Saiz 2010). And especially in the West and Southwest limited water resources may preclude developing land.

[9]I henceforth measure single-family construction by permits rather than starts, as only data on the former are available for metropolitan areas. The statistical relationship is estimated by a fitted regression of the single-family permitting rate on linear and quadratic terms for the 95th-percentile commuting time to the CBD. I exclude the New York City metropolitan area from the regression because its 95th-percentile commuting time, 115 minutes, is considerably above the next highest duration, 90 minutes, causing New York to have disproportionate leverage in estimating coefficients.

[10]Consistent with this causal relationship being more steeply negative, economic theory suggests that the causal relationship is likely to curve downward (that is, to become more steeply negative as commuting time increases). Increasingly long commutes leave increasingly less time for household responsibilities and leisure. For this reason, a reduction in commuting time by a few minutes is likely to be valued more by households with a long commute than those with al-

ready short commutes. Rappaport (2016) shows that workers' marginal valuation of time saved commuting increases with the duration of their commute if they are not able to choose their weekly hours (for example, if they are required to work 40 hours per week). Under a baseline calibration, the marginal valuation of a minute saved commuting approximately doubles as one-way commuting time rises from 20 to 40 minutes and then approximately doubles again as one-way commuting time rises from 40 to 60 minutes.

[11]An ongoing monthly survey finds that 28 percent of full-time paid employees worked remotely one to four days per week in July 2022 and 14 percent worked five days per week remotely (Barrero, Bloom, and Davis 2021). These shares of hybrid and full-time remote work are likely to change over time as businesses and workers experiment with different setups. But even if the extent of remote working settles at shares well below the survey results, the number of households able to move farther from on-site workplaces is likely to be sufficient to drive a long-term boom in single-family home construction. For example, each increase of 0.1 percentage point in the share of existing households in 2022 moving into a new single-family home would require 128,000 newly constructed units. Correspondingly, hybrid working would only need to spur an additional 0.4 percent of households per year to move into a newly constructed home to exceed my baseline prediction of the single-family construction boom.

[12]The annual values are based on working 48 weeks per year. Schrank and others (2021) report annual travel time indices (TTIs) for several hundred metropolitan areas, which are calculated as the ratio of the actual duration of all vehicle miles driven in a metropolitan area during a calendar year to the duration for the same vehicle miles if they had been driven at free-flow speed. As a benchmark for time saved due to faster speed, I use the ratio of a metropolitan area's TTI in 2020 to its TTI in 2019. This ratio is likely to considerably understate the time savings in 2020 due to remote working's effect on speed for commutes to and from the CBD during rush hours. One reason is that the TTI for 2020 includes the 10 weeks prior to the mid-March lockdown, during which speeds were similar to those in 2019. A second reason is that the ratio of TTIs captures speed improvements averaged over all vehicle miles driven in a metropolitan area, which are likely to have been considerably smaller than the speed improvements experienced by workers commuting to the CBD during rush hour. CBDs are disproportionately occupied by large office buildings, which in turn are disproportionately occupied by workers in occupations that can be done from home (Dingel and Neiman 2020). Consistent with this, traffic volume near CBDs during 2020 fell by far more than total metropolitan traffic volume. For example, according to Kastle Systems, daily attendance at large office buildings remained extremely low through the end of 2020: average attendance in 10 large metropolitan areas initially fell 85 percent and remained 75 percent below its pre-pandemic level at the end of the year. The total volume of traffic in large metropolitan areas in 2020 fell by considerably less, dropping 40 percent initially and remaining 20 percent below its pre-pandemic level throughout the second half of the year (Schrank and others 2021). In addition, even if the percentage decrease in volume had been the same throughout

a metropolitan area, the percentage reduction in travel time would have been greater near CBDs to the extent that pre-pandemic traffic congestion there was worse (Small and Verhoef 2007). Of course, the eventual reduction in rush-hour driving time on routes to the CBD may be less than the reduction in travel time on all trips within a metropolitan area during 2020. For this reason, I assume that the long-term reduction in travel time on trips to the CBD in each metropolitan area will be one-half of the reduction in travel time on all trips in the same metropolitan area in 2020 compared with 2019. For example, the travel time on trips in the New York metropolitan area decreased by an average of 14 percent in 2020 relative to 2019. I assume that hybrid working will cause travel time on trips to the New York CBD to decrease by 7 percent from their duration in 2019. One source of uncertainty on time saved is that some workers may switch from taking public transit to driving, thereby offsetting some of the decrease in driving volume from making fewer commutes.

[13]In addition to saving time, hybrid working cuts commuting expenses considerably. The 95th-percentile distances for the 10 metropolitan areas shown in Chart 7 range from 24.6 miles for workers driving to the Miami CBD to 51.7 miles for workers driving to the New York CBD. Using the Internal Revenue Service's 2019 reimbursement rate for business driving, $0.58 per mile, the corresponding annual savings from making two fewer round trips per week for each of 48 weeks per year ranges from $2,740 to $5,760.

[14]Across the 51 CBSAs with population of at least 1 million in 2010, regressing the mean distance of flows with long commuting times—those with commuting times to the CBD between the 90th and 95th percentiles—on population gives an R-squared value of 0.12. The associated coefficient implies that a doubling of a CBSA's population is associated with an increase of 1.5 miles in the distance of long commutes. Regressing the 95th-percentile distance of commute to the CBD, regardless of their duration, on population gives an R-squared value of 0.22; the associated coefficient implies that a doubling of a CBSA's population is associated with an increase of 2.6 miles in the 95th-percentile commuting distance to the CBD.

[15]I calculate the increase in driving speed and associated reduction in travel time using the updated Bureau of Public Roads function and recommended parameters for freeways reported in Small and Verhoef (2007). An additional factor that may contribute to saving more time in large metropolitan areas is that commuting time there tends to be less reliable than in smaller metropolitan areas, in the sense that travel time can vary considerably day to day. The flexibility made possible by hybrid working may allow employees to strategically choose to work remotely on days when congestion is especially high (Kahn 2022).

[16]Also consistent with an increased willingness to live farther from work, Ramani and Bloom (2021) document that population inflows and home price growth during 2020 and 2021 were greatest in the low-density suburbs of the largest metropolitan areas, which they term a "donut effect."

[17]In other words, I assume the causal relationship is linear with a slope equal-to that of a line that is tangent to the co-movement curve in Chart 7 at a 95th percentile duration of 75 minutes. As described in the main text, the causal relationship of commuting duration on permitting is likely to be negative at all durations with a steeper negative slope than the co-movement relationship when the latter is negative. Allowing for the likely concavity of the causal relationship implies that the difference in the increase in permits in metropolitan areas with long commutes compared with the increase in metropolitan areas with short commutes will be greater than the benchmark prediction, reinforcing the conclusion that hybrid working will disproportionately boost home construction in large metropolitan areas.

[18]I calculate the 1.4 million long-term average by summing the benchmark 427,000 per-year boost in single-family permits in CBSAs with population above 1 million in 2020, a corresponding 66,000 per-year boost in single-family permits in smaller CBSAs, and the 957,000 average annual rate of national single-family permitting during December 2019 through February 2020.

[19]The need for home office space is also likely to increase apartment construction. Although many households currently living in apartments may seek to move into single-family homes, many others are likely to seek to move into larger apartments. In addition, many individuals living with roommates or family members are likely to choose to move out on their own, and renting an apartment is typically the most affordable option.

[20]Commuting and home prices are two key congestion mechanisms equating the utility of living in different metropolitan areas. In a quantitative framework, lowering the frequency of commutes increases the population of metropolitan areas with high productivity and amenities relative to those with low productivity and amenities (Rappaport 2016).

References

Barrero, Jose Maria, Nicholas Bloom, and Steven J. Davis. 2021. "Why Working from Home Will Stick." National Bureau of Economic Research, working paper no. 28731, April. Available at https://doi.org/10.3386/w28731

Bertram, Nick, Steffen Fuchs, Jan Mischke, Robert Palter, Gernot Strube, and Jonathan Woetzel. 2019. *Modular Construction: From Projects to Products*. McKinsey & Company, June.

Brown, Jason P., Maeve Maloney, Jordan Rappaport, and Aaron Smalter Hall. 2017. "How Centralized Is U.S. Metropolitan Employment?" Federal Reserve Bank of Kansas City, Research Working Paper no. 17-16, November. Available at https://doi.org/10.18651/RWP2017-16

Couture, Victor, and Jessie Handbury. 2020. "Urban Revival in America." *Journal of Urban Economics*, vol. 119, September. Available at https://doi.org/10.1016/j.jue.2020.103267

Delventhal, Matthew J., and Andrii Parkhomenko. 2022. "Spatial Implications of Telecommuting." SSRN working paper, February 23.

Dingel, Jonathan I., and Brent Neiman. 2020. "How Many Jobs Can Be Done at Home?" *Journal of Public Economics*, vol. 189, September. Available at https://doi.org/10.1016/j.jpubeco.2020.104235

Humann, McKenzie, and Jordan Rappaport. 2022. "The Size of U.S. Metropolitan Areas." Federal Reserve Bank of Kansas City, Research Working Paper no. 21-02, June. Available at https://doi.org/10.18651/RWP2021-02

Kahn, Matthew E. 2022. *Going Remote.* Oakland: University of California Press.

Office of Management and Budget. 2013. "Revised Delineations of Metropolitan Statistical Areas, Micropolitan Statistical Areas, and Combined Statistical Areas, and Guidance on Uses of the Delineations of These Areas." OMB Bulletin no. 13-01, February 28.

Ramani, Arjun, and Nicholas Bloom. 2021. "The Donut Effect of Covid-19 on Cities." National Bureau of Economic Research, working paper no. 28876, May. Available at https://doi.org/10.3386/w28876

Rappaport, Jordan. 2018. "The Faster Growth of Larger, Less Crowded Locations." Federal Reserve Bank of Kansas City, *Economic Review*, vol. 103, no. 4, pp. 5–38. Available at https://doi.org/10.18651/ER/4q18Rappaport

———. 2016. "Productivity, Congested Commuting, and Metro Size." Federal Reserve Bank of Kansas City, Research Working Paper no. 16-03, January. Available at https://doi.org/10.18651/RWP2016-03

———. 2013. "The Demographic Shift from Single-Family to Multifamily Housing." Federal Reserve Bank of Kanas City, *Economic Review*, vol. 98, no. 4, pp. 29–58.

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

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

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics,* vol. 125, no. 3, pp. 1253–1296. Available at https://doi.org/10.1162/qjec.2010.125.3.1253

Schrank, David, Luke Albert, Bill Eisele, and Tim Lomax. 2021. *2021 Urban Mobility Report*. Texas A&M Transportation Institute.

Small, Kenneth A., and Erik T. Verhoef. 2007. *The Economics of Urban Transportation*. London: Routledge.

Smart Growth America. 2015. *Core Values: Why American Companies Are Moving Downtown*. Washington, DC: Smart Growth America.

Can Higher Gasoline Prices Set Off an Inflationary Spiral?

By Nida Çakır Melek, Francis M. Dillon, and A. Lee Smith

Imbalances between supply and demand following the COVID-19 pandemic pushed consumer price inflation, as measured by the consumer price index (CPI), to highs last witnessed in the early 1980s. Then in early 2022, Russia's invasion of Ukraine triggered further increases in energy prices, bringing retail gasoline prices to all-time highs and further lifting inflation. The timing of this recent spike in the price of gasoline—a salient good that is historically important in shaping consumers' inflation expectations—has increased public concerns that the U.S. economy could be in for a repeat of the inflationary spiral that gripped the nation in the 1970s and 1980s. During this period, a sequence of energy price increases helped to set off an inflationary spiral, as rising inflation and rising inflation expectations reinforced one another until a deep economic contraction broke the feedback loop.

In this article, we assess the risk of a similar spiral in the current environment by exploring whether high inflation makes consumers' inflation expectations more responsive to salient price increases—namely, higher gasoline prices. We explore this risk using differences across individual responses in survey-based microdata in the pre-COVID-19 era. We find that in response to an increase in the national price of gasoline, individuals with higher initial inflation expectations revise up their one-year-ahead inflation forecasts by almost twice as much as those

Nida Çakır Melek is a senior economist at the Federal Reserve Bank of Kansas City. Francis M. Dillon is a research associate at the bank. A. Lee Smith is a vice president and economist at the bank. This article is on the bank's website at ***www.KansasCityFed.org***

with lower initial inflation expectations. We interpret our findings as evidence that high rates of prevailing inflation can make consumers' inflation expectations more sensitive to salient price increases. In further evidence of this interpretation, we find a similar pattern when looking across time, as the sensitivity of one-year-ahead inflation expectations to gasoline prices also increases following periods when CPI inflation has been high. With current inflation high and consumers' inflation expectations elevated, our results suggest that changes in salient prices could indeed have an amplified effect on inflation expectations in the current environment.

Section I reviews research on the economic importance and determinants of inflation expectations. Section II presents our empirical analysis of how gasoline price increases affect inflation expectations across individuals and time. Section III discusses the implications of our findings for inflation expectations in the current environment.

I. The Link between Households' Inflation Expectations and Salient Prices

Inflation expectations are a key determinant of actual inflation. For instance, if firms anticipate higher future inflation and hence rising input costs, then they are likely to set higher prices today to protect their profit margins. Similarly, if households anticipate higher future inflation, then they are likely to negotiate for higher wages to preserve their spending power. In addition, inflation expectations can influence a range of forward-looking decisions households and firms make, including savings and consumption decisions as well as investment choices, thereby indirectly influencing inflation.[1] For example, if households expect higher inflation in the future, they may choose to make big-ticket purchases such as household appliances today; this pulling forward of future demand, in turn, drives up actual inflation. It is therefore not surprising that central banks—who are tasked with stabilizing inflation—pay close attention to inflation expectations.

One way to gauge inflation expectations is through survey-based measures such as the inflation expectations gathered from the Michigan Survey of Consumers (MSC). The MSC is one of the longest-running household surveys, conducted by the Survey Research Center at the University of Michigan, with data available monthly since 1978. The

MSC has around 500 participants each month and is weighted to be representative of the U.S. population. The survey also has a rotating panel component, so that each month, about 60 percent of interviewees are first-time respondents, and the remaining 40 percent were interviewed six months prior.[2]

The long history and rotating panel structure make the MSC especially useful for studying consumer inflation expectations. And Coibion and Gorodnichenko (2015) provide convincing evidence that the MSC measure of inflation expectations, compared with other survey measures of inflation expectations, is more closely linked with realized inflation outcomes.

The blue line in Chart 1 shows that average one-year-ahead inflation expectations from the MSC have shot up since the pandemic-induced recession in 2020, underscoring concerns about a feedback loop between inflation and inflation expectations. The role that these inflation expectations appear to play in shaping realized inflation prompts an important question: what determines households' inflation expectations?

Individuals' inflation expectations appear to be highly susceptible to changes in prices of the goods that they purchase most often. For example, Nobel laureate economist Robert E. Lucas has postulated that individuals form their expectations about aggregate inflation based on their source of information on the current state of the economy (Lucas 1972, 1973, 1975). D'Acunto and others (2021) bring Lucas's assertion to the data and find that the price changes of goods consumers purchase frequently in grocery stores influence their inflation expectations, while price changes of goods that they do not purchase in the same stores do not. Their evidence suggests that consumers interpret price changes in their individual consumption bundles as signals about aggregate price changes. Similarly, Brachinger (2008) shows that the aggregate perceived inflation of German consumers can be explained by overweighting the inflation rates of goods consumers purchase frequently. Overall, existing research suggests that frequent price changes in the goods consumers purchase most often—henceforth referred to as salient goods—play an important role in shaping inflation expectations.[3]

Gasoline is one such salient good, the price of which consumers frequently observe at retail gasoline stations. Indeed, several studies have documented a strong relationship between the price of gasoline and

Chart 1
Households' Short-Term Inflation Expectations and the Price of Gasoline

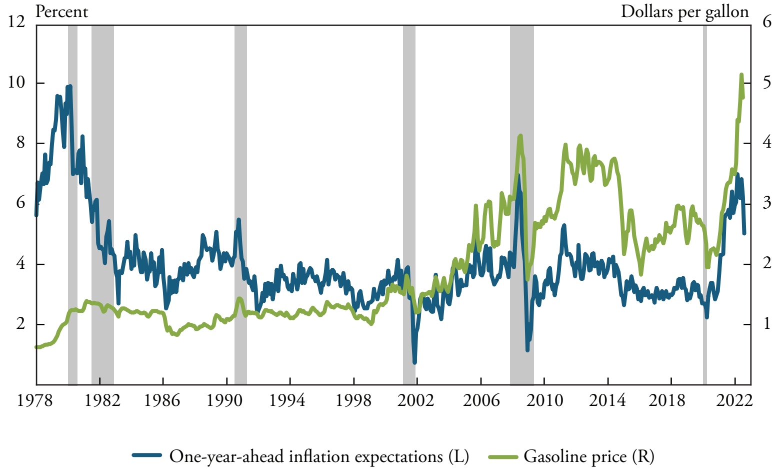

Note: Gray bars denote National Bureau of Economic Research (NBER)-defined recessions.
Sources: U.S. Energy Information Administration (EIA), University of Michigan Survey Research Center, and NBER (Haver Analytics).

households' inflation expectations (Hastings and Shapiro 2013; Georganas, Healy, and Li 2014; Coibion and Gorodnichenko 2015; Binder 2018). The volatile nature of gasoline prices contributes to this correlation, as consumers are more likely to remember extreme movements in prices and use them to form expectations about the future (Bruine de Bruin, van der Klaauw, and Topa 2011). Chart 1 presents visible evidence of the positive association between gasoline prices and one-year ahead inflation expectations from the MSC. For example, inflation expectations (blue line) rose in the mid-2000s, fell in 2008, and then rose again from 2009 to 2011, in line with swings in gasoline prices (green line).

Critically, this connection between gasoline prices and inflation expectations also appears to affect inflation dynamics. For example, Coibion and Gorodnichenko (2015) show that the link between oil prices and inflation expectations helps explain why inflation did not fall in the Great Recession, a puzzle dubbed "the missing disinflation." Indeed, they attribute the lack of disinflation from 2009 to 2011, a period of severe economic contraction, to the simultaneous increase in households' short-run inflation expectations. They then assign the increase in inflation expectations during this period to the remarkable sensitivity of inflation expectations to oil prices. The findings from Coibion and Gorodnichenko (2015) demonstrate that energy prices—operating through consumers' inflation expectations—can influence economy-wide inflation dynamics.

II. The Changing Sensitivity of Household Inflation Expectations When Inflation Is High

The linkages between gasoline prices, inflation expectations, and inflation demonstrated by Coibion and Gorodnichenko (2015) highlight the risk that surging gasoline prices can pose to the inflation outlook. We build on this result and investigate whether a surge in gasoline prices *in an already high-inflation environment* puts inflation expectations at a heightened risk of destabilizing and driving inflation even higher.

To formally assess whether consumer inflation expectations, and hence inflation, are more sensitive to salient price increases in a high-inflation environment, we analyze MSC microdata. We use individual consumer responses that have not been aggregated to measure changes in one-year-ahead inflation expectations during the pre-COVID-19 era, from 1981 to 2019.[4] We match each MSC survey with the price of gasoline in the corresponding month. This allows us to observe how changes in the national price of gasoline over a six-month period relate to concurrent changes in inflation expectations for the same individual.[5]

In our ideal dataset, we would be able to link the inflation rates experienced by individuals with their inflation expectations. This linkage would enable us to directly test whether the inflation expectations of individuals experiencing higher inflation are more sensitive to increases in gasoline prices than the expectations of individuals experiencing lower inflation. Unfortunately, experienced (or perceived) inflation is not regularly reported in the MSC. However, in a special installment, economists at the Board of Governors of the Federal Reserve System surveyed MSC respondents on how much prices changed over the past year to study the relationship between inflation perceptions and inflation expectations (Detmeister, Lebow, and Peneva 2016). They found a strong positive relationship between the inflation rates individuals perceive and the inflation rates that individuals expect in the future, suggesting that individuals reporting higher expected inflation are also likely to be currently experiencing high rates of inflation.

Leveraging these results, as a first step, we treat individuals' initial inflation expectations as a proxy for the inflation rates they experience and examine how these individuals revise their inflation expectations in response to an increase in gasoline prices over a six-month period. We split our sample into three subgroups: individuals with initial inflation expectations equal to 0, 1, or 2 percent; individuals with initial inflation expectations equal to 3 or 4 percent; and individuals with inflation expectations greater than or equal to 5 percent.[6] The blue bars in Chart 2 illustrate the estimated results for each of the three subgroups.[7] More specifically, the bars show the upward revision in inflation expectations of each group in response to a 1 percent increase in the price of gasoline.

Two key results are visible in the chart. First, an increase in the price of gasoline is associated with a statistically significant upward revision in inflation expectations for all groups, in line with the results reported in Coibion and Gorodnichenko (2015).[8] Second, and novel, the sensitivity of inflation expectations to changes in gasoline prices is monotonic—that is, the higher an individual's initial inflation expectations, the more sensitive their one-year-ahead inflation expectations are to an increase in gas prices. In fact, individuals with higher inflation expectations are almost twice as sensitive to an increase in the price of gasoline as those with lower inflation expectations. Specifically, respondents with initial short-term inflation expectations of 5 percent or higher revise their inflation forecasts up by 0.029 percent in response to a 1 percent increase in the price of gasoline, while respondents with initial inflation expectations of 0, 1 or 2 percent revise up their forecasts by only 0.017 percent in response to the same price increase.[9]

Although these forecast revisions may seem small in isolation, they can have a substantial influence on inflation expectations. Our estimates are based on a 1 percent increase in the price of gasoline; however, gasoline prices are highly volatile and can post large fluctuations (up to 60 percent) over a six-month period. As a result, even a 0.017 percent increase in expected inflation in response to a 1 percent increase in the price of gasoline—as estimated for the least sensitive group—can have an economically meaningful effect. For example, our results suggest a downward revision ranging between 1.02 and 1.74 percent in inflation expectations in response to a 60 percent decline in gasoline prices, a decline similar in size to that experienced during the global financial

Chart 2
Inflation Expectation Revisions

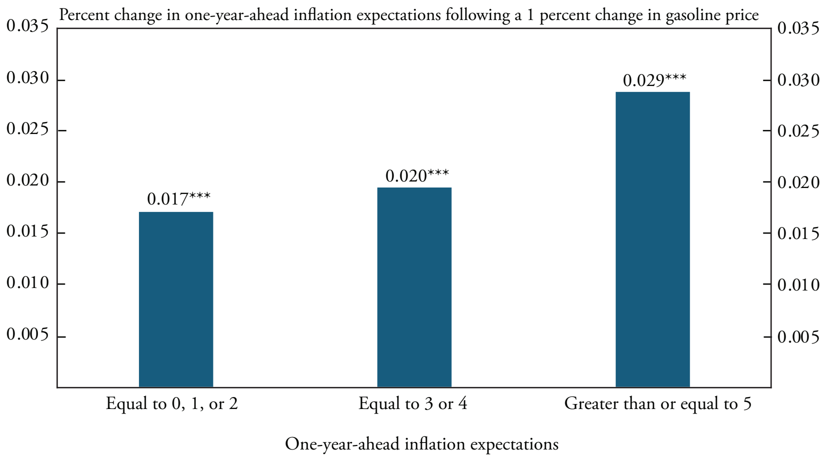

*** Significant at the 1 percent level

Notes: Individuals surveyed after December 2019 or who revised their inflation expectations by more than 15 percentage points are excluded from the sample. Observations are weighted using MSC microdata weights.
Sources: EIA, University of Michigan Survey Research Center, and authors' calculations.

crisis. Therefore, our evidence suggests that high inflation can make individuals' inflation expectations more sensitive to salient price increases.

Our analysis has thus far assumed that inflation expectations proxy for experienced (or perceived) inflation. We now relax this assumption at the cost of reduced data granularity. Instead of dividing our sample into subgroups based on inflation expectations and estimating each subgroup's sensitivity to gasoline price changes, we now estimate the sensitivity of inflation expectations—aggregated across all survey respondents—to changes in gasoline prices on two-year rolling samples. In other words, we estimate the sensitivity of inflation expectations to gasoline price changes over two-year intervals that move forward one month at a time. For example, the first estimate is for January 1981, which represents the effect of a 1 percent change in the price of gasoline on one-year-ahead inflation expectations estimated over the January 1981–December 1982 period. The second estimate is for February 1981, representing the sensitivity of inflation expectations to a 1 percent change in gasoline prices estimated over the February 1981–January 1983 period. We do this rolling estimation using data through the end of 2019, and hence lay out how the relationship

Chart 3
Changes in Sensitivity and Inflation

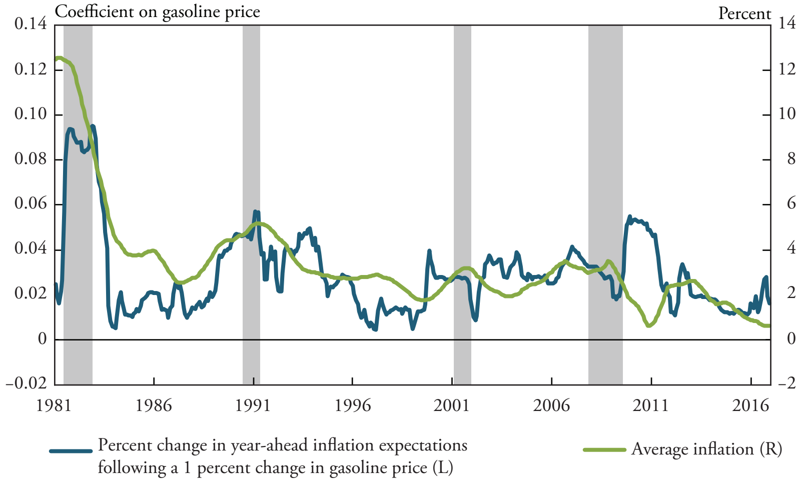

Notes: Individuals surveyed who revised their inflation expectations by more than 15 percentage points or who had negative inflation expectations are excluded from the sample. Average inflation is average headline CPI over the previous two years.
Sources: U.S. Bureau of Labor Statistics (Haver Analytics), EIA, University of Michigan Survey Research Center, NBER (Haver Analytics), and authors' calculations.

between revisions to inflation expectations and gasoline prices might have varied over different time periods.

The blue line in Chart 3 presents these time-varying estimates and shows that the sensitivity of inflation expectations to gasoline prices has in fact varied meaningfully across time. Moreover, the time variation in the relationship between revisions to inflation expectations and gasoline prices correlates positively with official inflation measures in previous years. The green line in Chart 3 shows realized average inflation over the two years preceding the survey as measured by the CPI. The chart shows a clear, positive relationship between the estimated sensitivity of inflation expectations to changes in gasoline prices and recent inflation. Indeed, the correlation between the two lines is 0.58, a strong positive correlation that suggests that following periods of high inflation, consumers make larger revisions to their inflation forecasts in response to changes in the price of gasoline compared with the same change in gasoline prices following a period of low inflation. Overall, the results from both the subgroup and aggregate analyses provide evidence that a high-inflation environment can make individuals' inflation expectations more sensitive to changes in salient prices.

III. Implications for the Current High-Inflation Environment

Given that the United States is experiencing its highest bout of inflation in decades, our findings from the previous section underscore that inflation expectations are at risk of moving even higher when consumers absorb a large increase in gasoline prices in a high-inflation environment. Inflation has become an everyday topic for many consumers, possibly strengthening the linkages between salient prices, inflation, and inflation expectations.

As households are already navigating large increases in the cost of living after a long period of price stability, the public has become much more attuned to inflation. For example, in August 2022, Google searches for "inflation" hit their highest level since the company began tracking searches in 2004. The run-up in inflation and the attention it has garnered has led individuals to increase their expectations for near-term inflation. The orange shaded area in Chart 4 shows the share of MSC respondents since 1978 who expect prices to go up by 5 percent or more in the next year. This share has increased from around 30 percent in 2019 to around 60 percent in early 2022, the highest since the early 1980s. Gasoline prices have also increased recently. Monthly gasoline prices increased by almost 50 percent in the first half of this year, partly due to shortages driven by the Russia-Ukraine war. Our pre-COVID-19 estimate presented in Section II suggests that, all else equal, this 50 percent jump in gasoline prices could lead individuals who already expect short-term inflation to be higher—individuals who now make up the majority of the survey respondents—to revise their expectations up by 1.45 percent (0.029×50).[10] Somewhat mechanically, our results suggest that the increase in the share of MSC respondents with high inflation expectations will likely lead to a stronger link between aggregate consumer inflation expectations and gasoline prices.

Our results can be interpreted in the context of recent research on rational inattention, a concept that policymakers have recently highlighted to emphasize the risk that high inflation can pose to inflation expectations. For example, in his speech at the 2022 Jackson Hole Economic Symposium Chair Jerome Powell said: "one useful insight into how actual inflation may affect expectations about its future path is based in the concept of 'rational inattention.' When inflation is persistently high, households and businesses must pay close attention and incorporate

Chart 4
One-Year-Ahead Inflation Expectations Since 1978

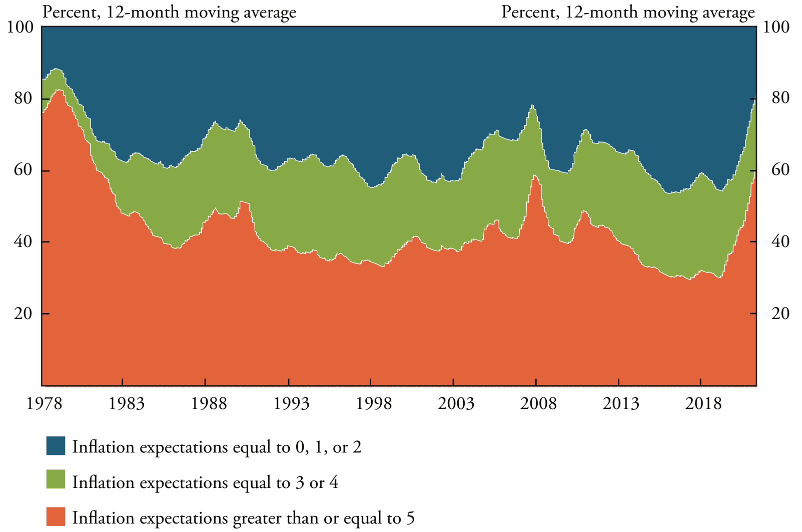

Note: Inflation expectations less than zero are excluded from calculations.
Sources: University of Michigan Survey Research Center and authors' calculations.

inflation into their economic decisions. When inflation is low and stable, they are freer to focus their attention elsewhere."

As highlighted by Chair Powell, the theory of rational inattention assumes that economic agents cannot process all available information but instead choose which fraction of available information to process based on their economic environment. For example, in a low-inflation environment, such as much of the past 30 years in the United States, inflation is likely to be less prominent in the minds of consumers, as errors in understanding inflation are likely to come at a low cost. In contrast, in a high and volatile inflation environment, inaccurate information on inflation can be costly, so rational inattention theory predicts that people will allocate more attention to inflation signals (Mankiw, Reis, and Wolfers 2003).

Cavallo, Cruces, and Perez-Truglia (2017) provide evidence consistent with this hypothesis based on survey experiments. They document that households in Argentina, an environment of high and volatile inflation at the time, were well-informed about recent inflation and monetary policy. In a parallel analysis for the United States, an environment of low and stable inflation at the time, they find that households were inattentive to inflation and monetary policy developments.

According to this previous research, households appear to adapt how they allocate their attention to the economic environment. With inflation high and the majority of individuals expecting higher inflation now than in the past, increased attentiveness to inflationary developments—such as higher salient prices—supports concern over further increases in inflation expectations. As more households raise their inflation expectations and hence become focused on inflationary developments, the effects of increasing gasoline prices on inflation and inflation expectations could be amplified. In a similar vein, a decrease in gasoline prices can have an amplified effect on lowering inflation expectations when inflation is high.

Conclusion

For much of the past decade, the risk of rising inflation and inflation expectations appeared a distant memory. For instance, in 2019, Federal Reserve Chair Jerome Powell characterized persistently low inflation and inflation expectations as the more pressing economic challenge of our time (Powell 2019). However, the re-emergence of high inflation during the recovery from the COVID-19 pandemic has reignited concerns about the possible feedback loop between rising inflation and inflation expectations that damaged the economy more than 40 years ago. The spike in gasoline prices in the summer of 2022 accentuated the similarities of inflation today to the 1970s and 1980s. Although much has changed since then, including the importance that monetary policymakers place on managing inflation expectations, our results underscore that inflation expectations remain susceptible to salient price shocks when inflation is already high.

Appendix

Methodology and Additional Results

We use microdata from the Michigan Survey of Consumers (MSC), which has been conducted monthly since January 1978. The survey asks individuals "by about what percent do you expect prices to go up/down on the average, during the next 12 months?" Respondents are probed if their answer is greater than 5 percent to make sure they understand the question, and responses in the microdata are capped at an absolute value of 25 percent.

Beginning in January 1981, the MSC began resurveying individuals who responded to their survey six months prior, allowing us to observe revisions in individual inflation expectations across six-month periods while also controlling for individual fixed effects.

We regress changes in individuals' one-year-ahead inflation expectations against changes in the price of gasoline over the same six-month periods, building on Coibion and Gorodnichenko (2015). We cluster standard errors by year and month. The model is as follows:

$$E_t^i \pi_{t,t+12} - E_{t-6}^i \pi_{t-6,t+6} = \alpha + \beta \times \log\left(\frac{GasP_t}{GasP_{t-6}}\right) \times 100 + error_{i,t}$$

where i and t represent individual i at time t, and $E_t^i \pi_{t,t+12}$ is the one-year-ahead inflation expectation during the second time the individual was surveyed. $GasP_t$ is the price of gasoline (as measured by the U.S. city average retail price of all grades of gasoline, in dollars per gallon, including taxes) during the second time an individual was surveyed.

Similar to Coibion and Gorodnichenko (2015), we remove outliers that could bias the data, so we exclude individuals whose absolute value of change in one-year-ahead inflation expectations $\left(|\, E_t^i \pi_{t,t+12} - E_{t-6}^i \pi_{t-6,\, t+6} \,|\right)$ was greater than or equal to 15. We also choose to exclude individuals who had inflation expectations of less than zero when they were initially surveyed. This group made up only 3.3 percent of the sample but oftentimes had large upward revisions to their inflation expectations (bringing them much closer to actual inflation) when they were reinterviewed. Of this group, over 80 percent had non-negative inflation expectations six months later. We view this group as an outlier, and thus exclude them from our analysis.

Table A-1
Effects of Gas and Food Price Changes, with and without Controls

	Gas			Food		
	Coefficient	Observation	R^2	Coefficient	Observation	R^2
Inflation expectations						
Equal to 0, 1, or 2	0.017***	30,525	0.006	0.002**	28,590	0.001
Equal to 3 or 4	0.020***	20,097	0.011	0.001**	19,142	0.001
Greater than or equal to 5	0.029***	28,867	0.010	0.003***	27,636	0.001
Inflation expectations including controls						
Equal to 0, 1, or 2	0.019***	28,622	0.017	0.002***	26,842	0.010
Equal to 3 or 4	0.019***	18,984	0.019	0.002***	18,085	0.009
Greater than or equal to 5	0.029***	27,114	0.012	0.003***	25,962	0.002

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

Notes: Regressions exclude individuals who had changes in inflation expectations greater than 15 or who were sampled after December 2019. Observations are weighted using MSC microdata weights. Standard errors are robust and clustered by year and month.

As a robustness check, we also regress changes in individuals' one-year-ahead inflation expectations against changes in food prices (as measured by the change in consumer price index food inflation) over the same six-month periods. In addition, we run both gasoline and food price regressions with a host of demographic variables to control for income, sex, age, education, and geographic region. Table A-1 presents the results.

Endnotes

[1]Among others, recent research documenting the role of inflation expectations in shaping economic decisions, such as consumption or financing, are Burke and Ozdagli 2021; Binder and Brunet 2020; Ichiue and Nishiguchi 2015; D'Acunto, Hoang, and Weber 2018, 2022; Dräger and Nghiem 2021; and Botsch and Malmendier 2021.

[2]The survey elicits inflation expectations in two steps. First, the survey asks respondents whether "prices in general" will increase, decrease, or stay the same over the next 12 months. Second, the survey asks those who answered "increase" or "decrease" by about what percentage they expect prices to go up or down, on average.

[3]Other factors such as lifetime experiences also affect the formation of consumers' beliefs and expectations about inflation. Malmendier and Nagel (2016) argue that individuals place extra weight on information about past inflation levels that they personally experienced when forming expectations. Relatedly, Binder and Makridis (2022) find that consumer sentiment becomes more pessimistic as gasoline prices rise, with the strongest effect for consumers who lived through the 1970s. Personal experiences matter not only for the expectations of households, but also for policymakers. Malmendier, Nagel, and Yan (2021) document that members of the FOMC who personally experience higher inflation during their lifetime are more likely to indicate higher inflation expectations in their semiannual Monetary Policy Reports to Congress. Recent research also highlights the role of cognitive abilities in driving inflation expectations, as forming expectations and making economic decisions require the use of cognitive resources. For example, using rich Finnish micro data, D'Acunto and others (forthcoming) find that cognitive abilities (measured by IQ) predict individuals' inflation expectations beyond the direct effects of income, education levels, wealth, and other proxies for economic sophistication.

[4]As mentioned in the previous section, the MSC has been conducted monthly since January 1978. Since January 1981, the MSC has resurveyed a fraction of participants from the prior six months (thus, in January 1981, the MSC resurveyed certain individuals who participated in the July 1980 survey).

[5]By looking at these changes over time for the same individual, we can implicitly control for individual fixed effects, or attribute beyond initial inflation expectations that lead some individuals to revise their expectations differently from others.

[6]Individuals with inflation expectations of less than zero are rare, and most individuals in this group have large upward revisions to their expectations between surveys; thus, we exclude them from our sample. In addition, following Coibion and Gorodnichenko (2015), we exclude individuals who revised their inflation expectations by more than 15 percentage points. The three subgroups are then chosen to have sample sizes sufficiently large and similar in size to one another.

[7]Our analysis tests whether coefficient estimates are different from zero in a statistically significant sense. In testing if the responses are statistically different from one another, we obtain overall similar results. Given the current period of high inflation and the recent increase in gasoline prices, Chart 2 illustrates revisions to inflation expectations following an *increase* in gasoline prices. However, our regression model is linear, so the plotted responses would have the opposite sign in response to a *decrease* in gasoline prices.

[8]Although still elevated, gasoline prices have dropped in recent months. Following this decline, and in accordance with our results, we have recently observed declines in one-year-ahead inflation expectations from their recent highs.

[9]We obtain similar results when we consider food prices instead of gasoline prices. Moreover, our results are robust to including different controls such as income, age, gender, region, and education.

[10]The 1.45 percent upward revision estimate is likely conservative given that it is based on pre-COVID-19 data.

References

Binder, Carola Conces. 2018. "Inflation Expectations and the Price at the Pump." *Journal of Macroeconomics*, vol. 58, pp. 1–18. Available at https://doi.org/10.1016/j.jmacro.2018.08.006

Binder, Carola, and Gillian Brunet. 2020. "Inflation Expectations and Consumption: Evidence from 1951." SSRN, April 15. Available at https://doi.org/10.2139/ssrn.3594809

Binder, Carola, and Christos Makridis. 2022. "Stuck in the Seventies: Gas Prices and Consumer Sentiment." *Review of Economics and Statistics*, vol. 104, no. 2, pp. 293–305. Available at https://doi.org/10.1162/rest_a_00944

Botsch, Matthew J., and Ulrike Malmendier. 2021. "The Long Shadows of the Great Inflation: Evidence from Residential Mortgages." SSRN, July 14. Available at https://doi.org/10.2139/ssrn.3888762

Brachinger, Hans Wolfgang. 2008. "A New Index of Perceived Inflation: Assumptions, Method, and Application to Germany." *Journal of Economic Psychology*, vol. 29, no. 4, pp. 433–457. Available at https://doi.org/10.1016/j.joep.2008.04.004

Bruine de Bruin, Wändi, Wilbert van der Klaauw, and Giorgio Topa. 2011. "Expectations of Inflation: The Biasing Effect of Thoughts about Specific Prices." *Journal of Economic Psychology*, vol. 32, no. 5, pp. 834–45. Available at https://doi.org/10.1016/j.joep.2011.07.002

Burke, Mary A., and Ali K. Ozdagli. 2021. "Household Inflation Expectations and Consumer Spending: Evidence from Panel Data." Federal Reserve Bank of Boston, Research Department Working Papers, no. 20-15. Available at https://doi.org/10.29412/res.wp.2020.15

Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices." *American Economic Journal: Macroeconomics*, vol. 9, no. 3, pp. 1–35. Available at https://doi.org/10.1257/mac.20150147

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics*, vol. 7, no. 1, pp. 197–232. Available at https://doi.org/10.1257/mac.20130306

D'Acunto, Francesco, Daniel Hoang, and Michael Weber. 2022. "Managing Households' Expectations with Unconventional Policies." *Review of Financial Studies*, vol. 35, no. 4, pp. 1597–1642. Available at https://doi.org/10.1093/rfs/hhab083

———. 2018. "Unconventional Fiscal Policy." National Bureau of Economic Research, working paper no. 24244, January. Available at https://doi.org/10.3386/w24244

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber. Forthcoming. "IQ, Expectations, and Choice." *Review of Economic Studies*.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber. 2021. "Exposure to Grocery Prices and Inflation Expectations." *Journal of Political Economy*, vol. 129, no. 5, pp. 1615–1639. Available at https://doi.org/10.1086/713192

Detmeister, Alan, David Lebow, and Ekaterina Peneva. 2016. "Inflation Perceptions and Inflation Expectations." Board of Governors of the Federal Reserve System, FEDS Notes, December 5. Available at https://doi.org/10.17016/2380-7172.1882

Dräger, Lena, and Giang Nghiem. 2021. "Are Consumers' Spending Decisions in Line with a Euler Equation?" *Review of Economics and Statistics*, vol. 103, no. 3, pp. 580–596. Available at https://doi.org/10.1162/rest_a_00909

Georganas, Sotiris, Paul J. Healy, and Nan Li. 2014. "Frequency Bias in Consumers' Perceptions of Inflation: An Experimental Study." *European Economic Review*, vol. 67, pp. 144–158. Available at https://doi.org/10.1016/j.euroecorev.2014.01.014

Hastings, Justine S., and Jesse M. Shapiro. 2013. "Fungibility and Consumer Choice: Evidence from Commodity Price Shocks." *Quarterly Journal of Economics*, vol. 128, no. 4, pp. 1449–1498. Available at https://doi.org/10.1093/qje/qjt018

Ichiue, Hibiki, and Shusaku Nishiguchi. 2015. "Inflation Expectations and Consumer Spending at the Zero Bound: Micro Evidence." *Economic Inquiry*, vol. 53, no. 2, pp. 1086–1107. Available at https://doi.org/10.1111/ecin.12176

Lucas, Robert E., Jr. 1975. "An Equilibrium Model of the Business Cycle." J*ournal of Political Economy*, vol. 83, no. 6, pp. 1113–1144. Available at https://doi.org/10.1086/260386

———. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review*, vol. 63, no. 3, pp. 326–334.

———. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory*, vol. 4, no. 2, pp. 103–124. Available at https://doi.org/10.1016/0022-0531(72)90142-1

Malmendier, Ulrike, and Stefan Nagel. 2016. "Learning from Inflation Experiences." *Quarterly Journal of Economics*, vol. 131, no. 1, pp. 53–87. Available at https://doi.org/10.1093/qje/qjv037

Malmendier, Ulrike, Stefan Nagel, and Zhen Yan. 2021. "The Making of Hawks and Doves." *Journal of Monetary Economics*, vol. 117, pp. 19–42. Available at https://doi.org/10.1016/j.jmoneco.2020.04.002

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2003. "Disagreement about Inflation Expectations." *NBER Macroeconomics Annual*, vol. 18, pp. 209–270. Available at https://doi.org/10.1086/ma.18.3585256

Powell, Jerome. 2022. "Monetary Policy and Price Stability." Speech at the Federal Reserve Bank of Kansas City Economic Policy Symposium, Jackson Hole, WY, August 26.

———. 2019. "Transcript of Chair Powell's Press Conference." Board of Governors of the Federal Reserve System, March 20.

The Evolving Role of the Fed's Balance Sheet: Effects and Challenges

By Chaitri Gulati and A. Lee Smith

The Federal Reserve's balance sheet more than doubled to nearly $9 trillion in the wake of the COVID-19 pandemic, primarily due to large-scale asset purchases (LSAPs). Although the Federal Open Market Committee (FOMC) had previously employed LSAPs during the global financial crisis, the pace of asset purchases in March 2020 was unprecedented, as the FOMC sought to relieve severe strains in financial markets that threatened to halt the flow of credit to households and businesses. The Federal Reserve continued to purchase assets at a reduced pace through March 2022 to support the economic recovery.

More recently, with inflation surging and the labor market tight, the FOMC has started to withdraw policy accommodation and has set in motion a plan to significantly reduce the balance sheet. However, the process of balance sheet reduction is likely to be challenging, as policymakers have much less experience with adjusting the balance sheet compared with their primary policy instrument, the federal funds rate. Indeed, they have engaged in quantitative tightening (QT), or balance sheet reduction, only once before. Moreover, it is not clear exactly how much accommodation the pandemic-era asset purchases have put in place, making it difficult for policymakers to judge how fast and how far to unwind them.

---

*Chaitri Gulati is an assistant economist at the Federal Reserve Bank of Kansas City. A. Lee Smith is a vice president and economist at the bank. This article is on the bank's website at **www.KansasCityFed.org***

In this article, we attempt to quantify the accommodation stemming from the expansion of the Federal Reserve's balance sheet from 2020 to 2022 and discuss the challenges policymakers may face in removing it. We present evidence that the Federal Reserve's expanded balance sheet, with a large portfolio of long-duration assets, has provided a significant amount of policy accommodation in recent years, depressing long-term interest rates by about 1.6 percentage points as of early 2022. We also argue that the FOMC's plan to remove this accommodation through the passive runoff of maturing securities, rather than outright asset sales, may prove limiting. Based on the 2017 episode of QT, as only soon-to-mature Treasury securities run off the balance sheet and prepayments of the Federal Reserve's mortgage holdings slow, we project that the downward pressure the balance sheet is currently placing on longer-term interest rates will only gradually reverse. Our results highlight an inherent asymmetry between the pace at which central banks can expand and contract their balance sheets. For asset purchases to provide effective policy accommodation, central banks must credibly commit to holding assets for a sustained period; unwinding these purchases sooner than expected could undermine the effect of future LSAPs.

Section I reviews the evolution of the size and composition of the Federal Reserve's balance sheet since the FOMC first deployed LSAPs in 2008. Section II analyzes previous research to estimate the amount of accommodation arising from the 2020–22 LSAPs. Section III reviews the FOMC's recently initiated plan for significantly reducing the balance sheet and draws on the similarly structured 2017–19 balance sheet runoff to project the effects of the current QT policy.

I. The Recent Evolution of the Federal Reserve's Balance Sheet

Prior to the 2007–09 global financial crisis, the Federal Reserve's balance sheet played only a supporting role in implementing monetary policy. In this era, the Federal Reserve's asset holdings were not the primary consideration when adjusting the balance sheet. Instead, adjustments to the balance sheet were made largely to achieve the target the FOMC set for the federal funds rate, the overnight interest rate at which banks borrow and lend reserves to each other. Each day, the Open Market Trading Desk at the Federal Reserve Bank of New York, in conjunction with the Federal Reserve Board of Governors, would forecast the amount of reserves the financial system needed to maintain the target funds rate. This projected level of reserves was achieved in practice by conducting open market operations (OMOs)—regularly adjusting the amount of reserves in the banking system by buying or selling Treasury securities in the System Open Market Account (SOMA) portfolio. The focus of these operations was on achieving the desired level of reserves, a liability on the Federal Reserve's balance sheet; consequently, these OMOs were typically small and concentrated in relatively short-maturity Treasury securities. For example, at the end of 2006, the Federal Reserve's Treasury holdings amounted to less than 9 percent of the Treasury market, and about half of these holdings matured in less than one year. Therefore, before the financial crisis, the balance sheet primarily affected longer-term interest rates through persistent changes in the overnight federal funds rate, achieved through adjustments in reserves (Ihrig, Meade, and Weinbach 2015).

The economic and financial fallout from the global financial crisis led to a radical shift in the way the FOMC used the balance sheet to conduct monetary policy. In December 2008, the Committee set the target range for the federal funds at zero to 0.25 of a percentage point, an all-time low. With economic and financial conditions continuing to deteriorate—and with no appetite to push rates into negative territory—the FOMC turned to LSAPs to reduce longer-term interest rates more directly and stimulate the economy. These LSAPs targeted mortgage securities and longer-term Treasury securities, drastically affecting not just the size but also the composition and maturity profile of the Federal Reserve's asset holdings relative to the pre-crisis norm.

Chart 1 shows the evolution of the Federal Reserve's bond holdings in recent decades, with shading to denote specific asset purchase programs, often referred to as rounds of quantitative easing, or QE. The blue region denotes par values of Treasury holdings—that is, the amount the bondholder receives when the security matures—while the green region denotes par values of federal agency debt and federal agency mortgage-backed security (MBS) holdings.

Although the Federal Reserve's balance sheet should naturally increase over time with the size of the economy to meet growth in the demand for currency and reserves, the increase in the balance sheet since

Chart 1

Evolution of the Federal Reserve's Bond Portfolio

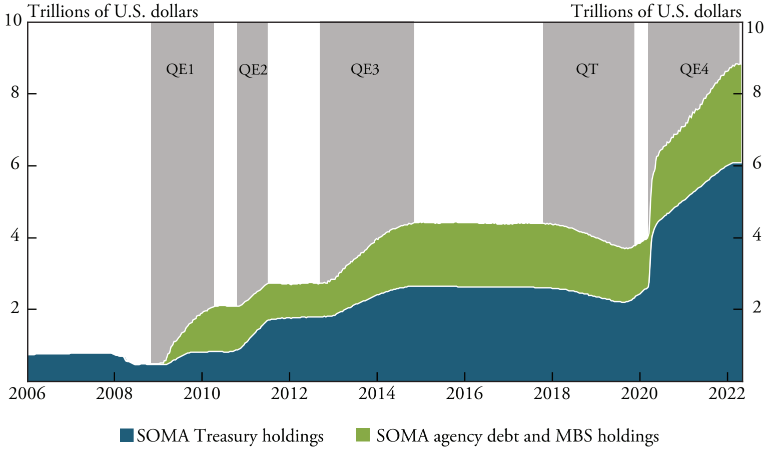

Source: Board of Governors of the Federal Reserve System (Haver Analytics).

the start of LSAPs has far outpaced economic growth. From 2008 to 2014, the size of the Federal Reserve's balance sheet increased nearly five-fold—from roughly $0.9 trillion to around $4.5 trillion—after three rounds of asset purchases dubbed QE1, QE2, and QE3. Then, from 2017 through 2019, the balance sheet shrank modestly during a first-of-its-kind process dubbed QT. In the spring of 2020, the FOMC turned once more to LSAPs, as the COVID-19 pandemic again pushed the federal funds rate to the zero lower bound. However, the scale of asset purchases in response to the pandemic was unprecedented, with purchases in April 2020 alone exceeding $1 trillion. Cumulative purchases since 2020, which we have denoted as "QE4" in Chart 1, exceeded $4.6 trillion, more than all three previous QE programs combined.

In addition to the size, the composition of asset holdings has also changed substantially since 2008. Chart 1 shows that prior to 2008, the balance sheet comprised almost entirely Treasury securities (blue region). However, during the financial crisis, the Federal Reserve entered the then-turbulent agency debt and MBS markets to contain upward pressure on mortgage rates. In particular, the FOMC began QE1 by purchasing debt issued by the three federal agencies—the Federal National Mortgage Association (Fannie Mae), the Federal Home Loan Mortgage Corporation (Freddie Mac), and the Government National

Mortgage Association (Ginnie Mae)—and the MBS guaranteed by them.[1] By the end of QE1, the Federal Reserve owned more than 20 percent of the agency MBS market. The Federal Reserve's presence in the MBS market has fluctuated since but never fully receded. At the beginning of 2020, the Federal Reserve again turned to agency MBS purchases to combat the COVID-19 crisis. By March 2022, the Federal Reserve's MBS holdings stood at more than $2.7 trillion, which amounted to nearly 30 percent of outstanding agency MBS securities.

The maturity profile of the Federal Reserve's asset holdings has also increased since 2008, due to increased holdings of both MBS and longer-term Treasury securities. The MBS purchases not only expanded the scope of assets held by the Federal Reserve, but also helped to increase the maturity of the Federal Reserve's asset holdings: most mortgages in the United States are paid off more slowly than the bills and other short-term Treasury securities that once comprised the majority of the Fed's asset holdings. Accompanying this increase in maturity from MBS purchases, the maturity of the Federal Reserve's Treasury portfolio has also significantly lengthened in recent decades. In 2009, the FOMC expanded the QE1 purchase program to include sizeable purchases of longer-term Treasury securities; subsequent asset purchase programs have further increased the Federal Reserve's holdings of longer-term Treasury securities. Chart 2 shows that the share of Treasury securities in the Fed's Treasury portfolio maturing more than five years into the future doubled from less than 20 percent at the end of 2006 to more than 40 percent at the end of March 2022.

The average maturity of the Federal Reserve's Treasury holdings peaked after the completion of the 2011–12 maturity extension program (MEP) and has since remained elevated. The MEP differed from other asset purchase programs in that it did not aim to increase the size of the Federal Reserve's balance sheet. Instead, the MEP aimed to increase the maturity profile of the Federal Reserve's Treasury holdings while keeping the overall size of the balance sheet constant, a goal achieved by selling shorter-maturity Treasury securities and using the proceeds to purchase longer-maturity Treasury securities. Chart 3 shows that the average maturity of the Federal Reserve's Treasury holdings increased from roughly three years at the end of 2006 to a peak of more than 10 years at the end of 2012. More recently, following the completion of

Chart 2
Maturity Profile of the Federal Reserve's SOMA Treasury Holdings

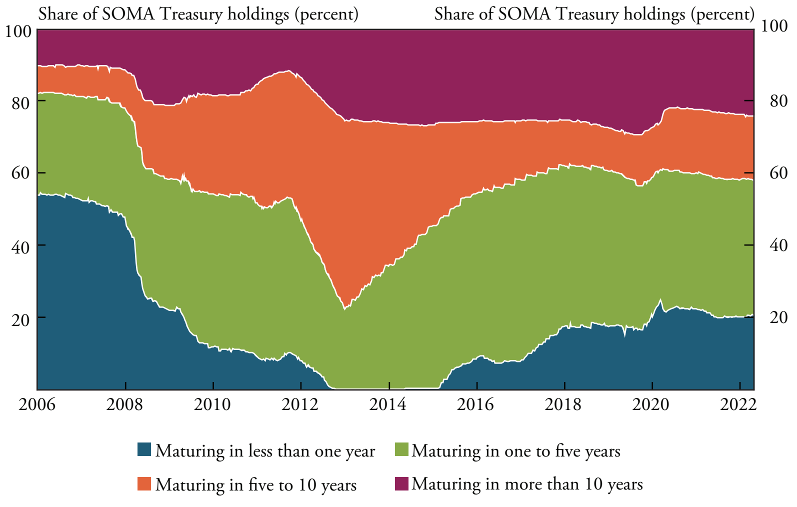

Source: Board of Governors of the Federal Reserve System (Haver Analytics).

Chart 3
Weighted-Average Maturity of the Federal Reserve's SOMA Treasury Holdings

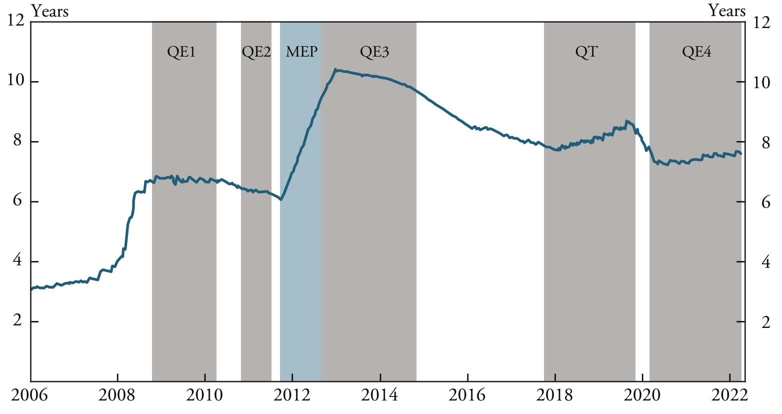

Note: Average maturity is calculated based on the par value weighted-average maturity.
Sources: Federal Reserve Bank of New York and authors' calculations.

the pandemic-era purchases in March 2022, the average maturity of Treasury holdings stood at 7.6 years.

A useful way to summarize changes to both the size and maturity composition of the Federal Reserve's balance sheet is to convert the Federal Reserve's bond portfolio into 10-year equivalents. Just as it sounds, this involves scaling the (par) value of each security by its maturity relative to the maturity of a 10-year Treasury note. For example, a Treasury security maturing in five years would receive approximately half its value in 10-year equivalents, whereas a Treasury security maturing in 20 years would receive approximately double its value in terms of 10-year equivalents. Rescaling each security on the Federal Reserve's balance sheet in this way and then summing the value of the 10-year equivalent holdings provides a single measure that can simultaneously account for changes in the size and maturity structure of the balance sheet. More formally, we compute the duration of each bond—the number of years it will take a bond holder to be repaid the bond's price by way of discounted future payments—compared with the duration of a 10-year Treasury security (see Appendix A for details on this calculation).

Chart 4 compares the size of the Federal Reserve's asset holdings in par value, as is typically reported, as well as the duration-adjusted size reported in 10-year equivalents. Although each round of QE is clearly visible in both measures, the 2011–12 MEP is only noticeable after the balance sheet is converted into 10-year equivalents. Therefore, measuring the size of the Federal Reserve's balance sheet in 10-year equivalents allows us to better capture the full range of balance sheet policies and their effects on longer-term interest rates.

The duration-adjusted size of the Federal Reserve's balance sheet summarizes a central mechanism through which central bank balance sheet policies are thought to operate; namely, by changing the supply of long-duration assets in the hands of investors.[2] The Federal Reserve can either remove long-duration assets from the market by purchasing them and thus expanding its balance sheet, or, as the MEP neatly illustrates, by adjusting the composition of its balance sheet toward longer-duration securities without necessarily changing its size. In either case, when the Federal Reserve reduces the supply of duration in the hands of investors by purchasing long-duration assets, the price of duration is effectively bid higher. As the pricing of duration is

Chart 4
Size and Duration-Adjusted Size of the Federal Reserve's SOMA Holdings

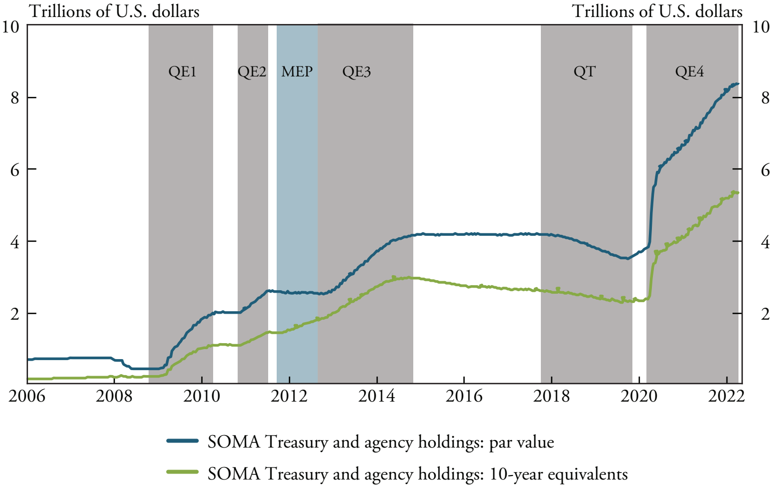

Notes: The 10-year equivalents are calculated based on the Treasury yield curve as of July 2014. See Appendix A for more details.
Sources: Federal Reserve Bank of New York, Board of Governors of the Federal Reserve System (Haver Analytics), Bloomberg LP, and authors' calculations.

embedded in all longer-term assets, the Federal Reserve's asset purchases have the potential to increase prices across a range of assets. And since bond prices and yields move inversely, increases in the duration-adjusted size of the Fed's bond portfolio should place downward pressure on longer-term interest rates.

QT acts to reverse these effects by removing long-term assets from the Fed's balance sheet, thereby increasing the supply of long-duration assets in the hands of investors. Although the qualitative effects of these adjustments on longer-term interest rates are straightforward, the quantitative effects remain an open question.

II. Quantifying the Accommodation from the Balance Sheet

To quantify the effects of balance sheet adjustments on longer-term interest rates, we review previous studies on the relationship between the supply of duration and longer-term interest rates. Many research papers estimate this relationship using event studies, which measure how market interest rates respond to announcements of asset purchases by the Federal Reserve. However, event studies often focus on announcements of asset purchases made during crises; as a result, these studies may capture channels that are not operative in normal times and that do not apply symmetrically to balance sheet expansions and reductions.[3] For example, during periods of severe financial distress, market functioning may be impaired, leading asset purchase announcements to have outsized effects, such as during QE1 (D'Amico and King 2013). Furthermore, during crises, asset purchases may generate important liquidity effects not present in normal times. Vissing-Jorgensen (2021) argues that liquidity provision was a key channel through which the FOMC's asset purchases transmitted to Treasury and MBS markets in the spring of 2020. However, this liquidity channel of asset purchases may not be present in periods of normal market functioning when the balance sheet is reduced and that liquidity is withdrawn.

For these reasons, we restrict our attention to research that meets three criteria: (i) the research has been peer-reviewed or formally subjected to comments from other researchers, (ii) the research does not rely solely on estimates from the QE1 period nor the spring of 2020, and (iii) the analysis enables us to convert estimated effects in terms of duration-adjusted quantities (10-year equivalents) on the 10-year Treasury yield or comparable long-term Treasury rates. We impose the first criterion to ensure the estimates are of sound quality, the second criterion to ensure the estimates do not reflect the effects of asset purchases during periods of severe market dislocations, and the third criterion to ensure the estimates can be applied to the duration-adjusted size of the balance sheet.

Table 1 summarizes the eight studies that meet our criteria and accordingly inform our estimates of the effect of the balance sheet on longer-term rates. These eight studies cover a range of estimation samples and statistical techniques, ensuring our estimate of the balance sheet effect on long-term rates is not driven by any one event or estimation strategy. To facilitate comparison, we apply the same thought experiment to each study—specifically, we consider how much the 10-year Treasury yield would fall if investors were asked to hold $100 billion less in 10-year equivalent U.S. government debt (for example, because the Federal Reserve removed duration from the market through

Table 1
Summary of Estimates of Asset Purchases on 10-Year Treasury Yield

Study	Estimation sample	Event analyzed	Par value purchases (billions of dollars)	Duration-adjusted purchases (billions of dollars, 10-year equivalents)	Estimated effect on 10-year yield (basis points / $100 billion, 10-year equivalents)
Greenwood and Vayanos (2014)	1952–2007	QE1, QE2	$900	$569	–2.3
Hamilton and Wu (2012)	1990–2007	MEP	$400	$400	–3.5
Swanson (2011)	1961	QE2	$600	$400	–3.8
Krishnamurthy and Vissing-Jorgensen (2011)	2010	QE2	$600	$400	–4.5
Gagnon and others (2011)	1985–2008	QE1	$1,750	$750	–4.5
Hanson (2014)	1989–2014	QE1	$1,750	$750	–7.2
Li and Wei (2013)	1994–2007	QE1, QE2, MEP	$2,750	$1,550	–9.7
D'Amico and others (2012)	2002–2008	QE1, QE2	$900	$569	–14.1
Median estimate (25–75 percentile range)					–4.5 (–3.7, –7.8)

Notes: Greenwood and Vayanos (2014) and D'Amico and others (2012) study only the Treasury component of QE1 purchases. The 10-year-equivalent amounts are from Hanson (2014).

asset purchases). Answering this question for each study requires some conversions of the estimates in the respective papers (details on how we convert these estimates to the values shown in Table 1 are provided in Appendix B). The far-right column of Table 1 shows the resulting estimates (in basis points) for each study, along with the median estimate across studies.

The median estimate from our meta-analysis suggests that every $100 billion in 10-year equivalents purchased by the Federal Reserve reduces the 10-year Treasury yield by 4.5 basis points. To arrive at a measure of the accommodation stemming from the Federal Reserve's balance sheet, we apply this median estimate to the size of the duration-adjusted SOMA portfolio scaled relative to the size of the economy. By this measure, the Federal Reserve's balance sheet is providing significant accommodation, depressing the 10-year Treasury by roughly 160 basis points as of the first quarter of 2022 (more details on this calculation are provided in Appendix B). Perhaps as expected, given the doubling

Box

Estimates of the Effects of Asset Purchases on Longer-Term Interest Rates

The estimates from our meta-analysis in Table 1 largely align with the estimates from other studies synthesizing existing research on the effects of asset purchases on the 10-year Treasury yield.

Two widely cited reviews are from Williams (2013) and Gagnon (2016), both of which compile estimates of the effects of LSAPs on the 10-year Treasury yield across a range of studies. However, unlike our analysis, the surveys in Williams (2013) and Gagnon (2016) do not account for the duration of the Fed's asset purchases. Williams (2013) instead argues that a $600 billion (par value) asset purchase program targeting medium- and longer-term Treasury securities reduces the 10-year Treasury yield by 15 to 25 basis points. If we report the purchase amounts in the studies cited in Table 1 by par value rather than 10-year equivalents, the median estimate across the eight studies we analyze would imply that $600 billion in purchases reduces the 10-year Treasury yield by about 18 basis points, near the midpoint of the range reported in Williams (2013).

Like Williams (2013), Gagnon (2016) does not adjust asset purchases for duration but does normalize par value purchase amounts as a share of GDP. The median estimate from Gagnon (2016) suggests that an asset purchase program amounting to 10 percent of nominal GDP reduces the 10-year Treasury yield by 50 basis points. In 2011, when QE2 and the MEP took place, two events that served as the policy scenarios that many of the estimates in Table 1 were based on, 10 percent of nominal GDP amounted to roughly $1,580 billion in par value purchases. Therefore, the median estimate in Table 1 implies that an asset purchase program amounting to 10 percent of nominal GDP reduces the 10-year Treasury yield by 48 basis points (= 18 basis points × $1,580 billion / $600 billion), remarkably close to the median estimate of 50 basis points reported in Gagnon (2016).

The proximity of the estimates from Table 1 with those from prior research synthesizing the effects of asset purchases suggests that our criteria for selecting which studies to include in our analysis did not meaningfully bias the estimates one way or another.

Chart 5
Estimate of Accommodation from the Federal Reserve's Balance Sheet

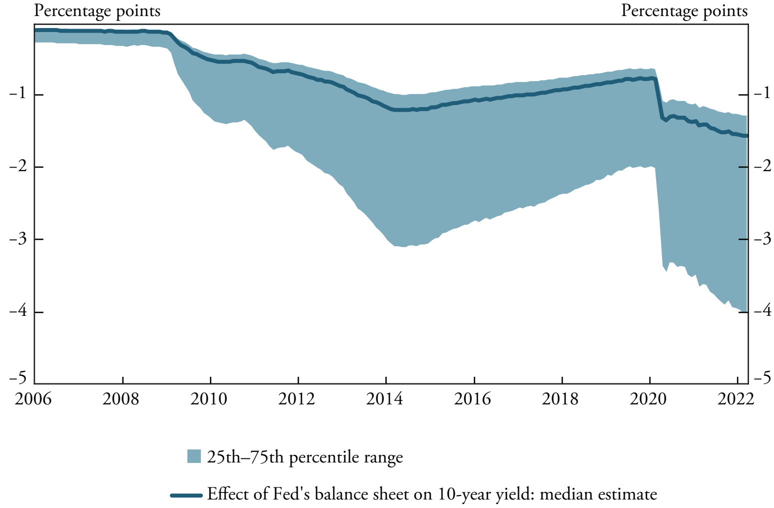

Note: See Appendix B for more details.
Sources: Federal Reserve Bank of New York, Board of Governors of the Federal Reserve System (Haver Analytics), Bloomberg LP, U.S. Bureau of Economic Analysis (Haver Analytics), and authors' calculations.

of the Federal Reserve's balance, about half of the 160 basis points of downward pressure on longer-term interest rates stems from asset purchases since the spring of 2020.

However, wide uncertainty surrounds this estimate. For example, Greenwood and Vayanos (2014) estimate much more modest effects from asset purchases on longer-term interest rates, while D'Amico and others (2012) estimate much larger effects. Differences across estimates reflect numerous factors, including different samples and methodology, as well as more fundamental differences in the channels through which asset purchases can influence longer-term interest rates.[4] We attempt to acknowledge the uncertainty around our estimate by also reporting the 25th–75th percentile range of estimates from Table 1. This range suggests that the Federal Reserve's balance sheet was depressing longer-term rates by anywhere from 130 to 270 basis points as of the first quarter of 2022.

Chart 5 shows the evolution of the accommodation emanating from the balance sheet. The solid line and surrounding shaded area show the median estimate and the 25th–75th percentile range, respectively, of the effect of the Fed's asset holdings on the 10-year Treasury yield. Since late 2008, the Federal Reserve's balance sheet has proven to be an effective means of providing policy accommodation. Given that the federal funds rate was at or near the zero lower bound for much of the last 15 years, the stimulatory effects from balance sheet expansions have played an important role in supporting the U.S. economy at times when reductions in the federal funds rate were unavailable. However, this accommodation has proven to be persistent, never fully reversing. Particularly remarkable is the fact that despite a concerted effort to shrink the balance sheet from 2017 to 2019 through QT, the accommodation emanating from past asset purchases is estimated to have only partially receded. This underscores a potential challenge in relying on the balance sheet for policy accommodation: the difficulty in removing accommodation through the balance sheet once the economy no longer requires added monetary policy support.

III. Challenges in Removing Policy Accommodation through the Balance Sheet

As the FOMC embarks on another episode of QT, it faces new and unprecedented challenges in unwinding the balance sheet. During the recovery from the global financial crisis, inflation remained low, and employment gains were frustratingly slow; therefore, the minimal withdrawal of policy accommodation through the balance sheet did not impede the pursuit of the FOMC's economic objectives. However, in the pandemic recovery, the economy is in a starkly different position. Inflation is reaching multi-decade highs, and by some measures, the labor market is historically tight. Against this backdrop, the FOMC is pursuing a more rapid withdrawal of policy accommodation through the balance sheet. Whereas the FOMC waited nearly three years after the end of QE3 in 2014 to begin QT in 2017, the FOMC only waited three months between ending net asset purchases in March 2022 and beginning to reduce the balance sheet in June 2022. Moreover, the pace of balance sheet reduction is nearly double the pace from the 2017–19 QT.

Nevertheless, the FOMC's current plan for reducing the balance sheet follows the same core strategy employed during the 2017–19 QT. In particular, the Committee is reducing the balance sheet by no longer reinvesting the proceeds from a portion of maturing securities rather than selling assets outright. Although the FOMC signaled that it would consider MBS sales once balance sheet runoff is well underway, no plans have been announced (Board of Governors 2022a).

Some may argue that the current economic backdrop of low unemployment and high inflation calls for a more deliberate approach to QT—one that includes asset sales. Indeed, the decision to not sell assets limits the speed at which the balance sheet can shrink. However, this decision also preserves the power of asset purchases in the event of a future crisis.[5] For asset purchases to be effective in moving long-term interest rates, market participants must believe the central bank will hold the assets it purchases for a significant period of time (Potter 2016). For example, if the central bank announces its intention to buy a large amount of long-duration assets and at the same time announces its plans to sell said assets a short period later, the effects on longer-term interest rates would likely be nil. While extreme, this example illustrates that a quicker-than-expected reduction in the balance sheet may diminish the effectiveness of future asset purchases. This logic imparts an inherent asymmetry in the pace at which the central bank's balance sheet can grow and shrink while retaining the effectiveness of LSAPs and underscores a fundamental challenge with removing accommodation from past asset purchases.

Given that the FOMC is not currently pursuing asset sales, we leverage the similarities between the FOMC's current plan for reducing the balance sheet and the 2017–19 QT episode to predict how the current balance sheet runoff will evolve. The FOMC's stated objective of QT in 2017 was to arrive at a smaller, Treasury-only portfolio (Board of Governors 2022b). To achieve this objective, the FOMC allowed principal payments from maturing securities to run off the balance sheet rather than fully reinvesting the proceeds. The FOMC also capped the pace at which these payments were allowed to run off to reduce the balance sheet "in a gradual and predictable manner." These caps, once fully phased-in, ensured that no more than $30 billion in Treasury securities and $20 billion in agency MBS ran off the balance sheet each month.

The FOMC's passive approach to balance sheet runoff failed to significantly decrease the duration-adjusted size of its asset holdings. Therefore, the previous QT episode led to an incomplete withdrawal of the accommodation stemming from past asset purchases. From October 2017 through September 2019, the par value of the balance sheet shrank by only about $700 billion—a modest amount given the scale of the preceding asset purchase programs, which totaled more than $3.5 trillion. Moreover, by solely relying on principal paydowns of maturing securities from 2017 to 2019, only soon-to-mature Treasury securities rolled off the balance sheet while longer-dated Treasury securities remained. As Chart 3 shows, this approach skewed the weighted-average maturity of the Treasury portfolio higher. As a result, we calculate that the duration-adjusted decline in the balance sheet from October 2017 through September 2019 totaled just $290 billion in 10-year-equivalents. According to the median estimate from Table 1, the effect of this decrease in the duration of the SOMA portfolio would translate to an average increase in the 10-year Treasury yield of only 10 basis points from 2017 to 2019 (with an interquartile range of 8 to 17 basis points). Smith and Valcarcel (2022) arrive at a similar estimate using an entirely different empirical strategy; they estimate that QT pushed up the 10-year Treasury yield by just 8 basis points on average over this same 2017–19 period.

Moreover, the comparatively gradual runoff of MBS holdings limited the FOMC's progress toward achieving a portfolio consisting primarily of Treasury securities. Most mortgages in the United States are issued for 15- or 30-year terms; accordingly, none of the MBS purchased since 2008 were maturing from 2017 through 2019. Although MBS are frequently repaid before maturity, as when a house is sold or a mortgage is refinanced, these activities tend to slow in a rising-interest-rate environment, leading actual runoff of MBS to fall below the runoff cap once it was fully phased-in in late 2018. As a result, the comparatively slower rate of MBS runoff led to an increase in the share of agency MBS in the SOMA portfolio during balance sheet runoff, moving further away from the Treasury-only portfolio the FOMC had envisioned.

In 2022, the FOMC is pursuing a similar strategy for reducing the balance sheet, albeit with higher runoff caps. The fully phased-in runoff caps are roughly double the 2017–19 pace, increasing from $50 billion per month to $95 billion per month. However, as in 2017–19, only the shortest-duration Treasury securities will mature, and reinvestments above the caps will be spread across new issuance. As a result, the average maturity of the Treasury portfolio will likely increase once again, dampening the duration-implied effects of the runoff. Moreover, projections based on current mortgage rates suggest agency MBS holdings will rarely reach the fully phased-in cap (Logan 2022). Both factors suggest that the pace of the actual runoff—both in par value and in terms of 10-year equivalents—will be slower than the notional cap of $95 billion per month and that progress toward a primarily Treasury portfolio will be limited.

Based on the current runoff strategy, we anticipate that a significant amount of the accommodation from past asset purchases is likely to remain in place in the coming years. Recent projections from the Federal Reserve Bank of New York suggest that the Federal Reserve's bond portfolio will decline by about $2.5 trillion in par value from June 2022 through the end of 2025.[6] As in the previous QT episode, only soon-to-mature Treasury securities will run off the balance sheet in coming years, and the decline in terms of 10-year equivalents will therefore be significantly smaller. Based on the relative decline in 10-year equivalents versus par value from 2017 to 2019, we project that the decline of $2.5 trillion in par value will only reduce the balance sheet by about $1 trillion in 10-year equivalents from 2022 through 2025.[7] Relying on the median estimate from Table 1, absent other developments, this decline in the balance sheet would be expected to increase the 10-year Treasury yield by roughly 30 basis points by the end of 2025.[8] Given that we estimate that the balance sheet is depressing the 10-year Treasury yield by roughly 160 basis points, we anticipate that the downward pressure on longer-term interest rates stemming from the balance sheet will only partly reverse in coming years. More fully removing this accommodation will likely require a prolonged period of reinvestments from principal payments into shorter-maturity Treasury securities, thereby shrinking the duration-adjusted size of the balance sheet as a share of GDP closer to its pre-2008 share.

Conclusion

Recently, policymakers have come to rely on expansions of the balance sheet through LSAPs to provide further policy accommodation when the federal funds rate is constrained by the zero lower bound. Although LSAPs have been shown to be effective at delivering desired accommodation, the Federal Reserve's limited experience with shrinking the balance sheet suggests that reducing this accommodation can be a slow and challenging process.

We argue that the challenges associated with balance sheet reduction are inherent to the use of the balance sheet as a policy tool. In particular, retaining the full effectiveness of future balance sheet expansions likely requires somewhat gradual reductions. In light of this intrinsic asymmetry between the pace at which the balance sheet can grow and shrink, policymakers may need to weigh future balance sheet expansions against the potential costs of putting in place persistent policy accommodation.

Appendix A

Converting the SOMA Portfolio into 10-Year Equivalents

We use Committee on Uniform Security Identification Procedures (CUSIP)-level data of the Federal Reserve's SOMA portfolio to convert the SOMA portfolio from par or face value into 10-year equivalents. These data are available Wednesday of each week from the Federal Reserve Bank of New York from July 2003 to the present.

Over this sample, the SOMA portfolio has consisted of Treasury bills, notes, and bonds; Floating Rate Notes (FRNs); Treasury Inflation Protected Securities (TIPS); inflation compensation on TIPS holdings; agency debts; agency mortgage-backed securities; and commercial mortgage-backed securities.

We convert each security from par value into 10-year equivalents based on the Treasury yield curve as of July 30, 2014. This follows Greenwood and others (2016) and prevents shifts in the Treasury yield curve from affecting the measure of 10-year equivalents week to week due solely to changes in interest rates that, in turn, affect the way future coupon payments are discounted to the present. On this day, the duration of a 10-year Treasury note is measured to be 8.9 years. Although the duration of Treasury securities can be directly calculated with the yield curve data and the detailed security-level data from the Federal Reserve Bank of New York, calculating MBS duration is considerably more complicated, as it depends on the probability that a mortgage is refinanced or paid-off early given current interest rates. We therefore follow Hanson (2014) and use a measure of MBS duration obtained from Bloomberg LP. Consistent with our Treasury duration calculations, we hold this duration measure fixed as of July 30, 2014. For each type of security, we use the following methods to convert the par or face value of the security into 10-year equivalents.

- *Treasury bills, notes, bonds, TIPS, and agency debts*. Duration is measured using the Macaulay modified duration formula based on par value, issue and maturity dates, and coupon rates with assumed biannual coupon payments (where applicable).
- *FRNs*. Duration is set to six days unless it matures before then, as the rates on these notes are reset weekly (on Tuesdays).

- *Agency mortgage-backed securities*: Bloomberg U.S. MBS Modified Duration is set to 5.1 years, its value on July 30, 2014.
- Inflation compensation on TIPS holdings and commercial MBS are excluded from the calculation.

We then scale the par or face value of each security by the ratio of its measured duration relative to 8.9 years, the duration of a 10-year Treasury note on July 30, 2014. We then sum over all securities to arrive at a measure of the SOMA portfolio in 10-year equivalents.

Appendix B

Converting Reported Duration or Asset Purchase Effects on the 10-Year Treasury Yield to Effects per $100 Billion of 10-Year Equivalents

The meta-analysis in Section II draws on estimates of changes in the supply of long-duration assets or asset purchases on the 10-year Treasury yield from eight previously published studies. Here we detail how the previously published estimates are converted to effects on the 10-year Treasury yield per a $100 billion reduction in the supply of 10-year equivalents, shown in the last column of Table 1.

To summarize the estimated effects of the balance sheet on the 10-year yield shown in the last column of Table 1, we first identify the event analyzed (QE1, QE2, or MEP) of each study in our sample and gather the published estimate of how the analyzed event(s) affected the 10-year Treasury yield in basis points. We then obtain the dollar amounts for these events in terms of 10-year equivalents from Hanson (2014). We take the ratio of these two items: the effect on the 10-year Treasury yield (in basis points) divided by the total amount of purchases for the analyzed event(s) in terms of 10-year equivalents (expressed in billions of U.S. dollars). Multiplying this ratio by 100 provides the implied estimate of the basis point effect on the 10-year Treasury yield per $100 billion in 10-year equivalents. The exact inputs for these calculations are detailed in Table B-1.

We now walk through this conversion for each of the eight studies shown in Table B-1.

- *Greenwood and Vayanos (2014).* This study analyzes the effect of changes in the supply of maturity-weighted debt relative to nominal GDP on the 10-year Treasury yield over a 1952–2007 sample. Based on the regression estimates, the authors note that a one-unit decrease in maturity-weighted debt to GDP lowers "long-term" Treasury yields by 40 basis points (p. 685, Table 2). The authors impute that the Treasury component of QE1 and QE2 purchases reduced the maturity weighted debt-to-GDP ratio by 0.32, leading us to infer a total effect of these purchases on long-term yields of about 13 basis points (= 0.32 × 40).

- *Hamilton and Wu (2012).* This study contemplates a MEP-style operation of buying $400 billion in long-term yields funded by creating reserves which, at the zero lower bound, is roughly equivalent to selling $400 billion in short-dated Treasury securities and using the proceeds to buy $400 billion in longer-term Treasury securities. We assume the operation is structured to remove $400 billion in 10-year equivalents, as was the case with the actual MEP according to Hanson (2014). Based on the pre-crisis estimates of their term structure model, Hamilton and Wu estimate this operation would lower the 10-year Treasury yield by 14 basis points (p. 32, Table 5).
- *Swanson (2011).* This study analyzes the effects of "Operation Twist" in the early 1960s, whereby the Federal Reserve purchased longer-term Treasury securities while the Treasury shifted issuance towards shorter-term securities. Swanson argues that the size of Operation Twist purchases ($8.8 billion) is comparable in size to QE2 ($600 billion), relative to the size of the economy and the Treasury market at the time. Swanson estimates that Operation Twist—which can be seen as roughly similar in size to QE2 —reduced the 10-year Treasury yield by roughly 15 basis points (pp. 174–175, Table 3).
- *Krishnamurthy and Vissing-Jorgensen (2011).* This study analyzes the change in Treasury yields and other asset prices in a one-day window around the announcements by the Federal Reserve related to QE2. The authors estimate that the cumulative change around the two most relevant announcement dates related to QE2 lowered the 10-year Treasury yield by 18 basis points (p. 248, Table 5).
- *Gagnon, Raskin, Remache, and Sack (2011).* This study analyzes the effect of changes in the net supply of 10-year equivalents relative to nominal GDP on the term-premium component of the 10-year Treasury yield over a 1985–2008 sample. Based on the regression estimates, the authors predict that the QE1 purchases would lower the 10-year Treasury yield by 38 basis points (p. 30). The authors also predict that QE 1 reduced the supply of 10-year equivalents by $850 billion, rather than

the $750 billion cited in Hanson (2014). To make the various estimates comparable, we use the $750 billion figure and therefore scale down the estimated effect from 38 basis points to 34 basis points (= 38 basis points × $750 / $850).

- *Hanson (2014).* This study regresses changes in MBS duration on 10-year Treasury yields and finds that a 1 standard deviation decline in MBS duration, which amounts to a $503 billion decline in 10-year equivalents, lowers the 10-year Treasury yield by 36 basis points (p. 286, Table 4). Hanson further computes that QE1 lowered the supply of 10-year equivalents by $750 billion, implying a total effect on the 10-year yield of 54 basis points (= 36 basis points × $750 / $503).
- *Li and Wei (2013).* This study analyzes the effects of changes in the supply of long-duration assets relative to nominal GDP on Treasury yields over a pre-crisis sample of 1994–2007 using a term structure model. They then use the estimated model to infer the effects of QE1, QE2, and the MEP. The authors note that the combined effects of QE1, QE2, and the MEP removed an amount of duration that would be predicted to lower the 10-year Treasury yield by 150 basis points (pp. 28–29, Table 6).
- *D'Amico, English, Lopez-Salido, and Nelson (2012).* This study analyzes changes in the supply of long-duration assets using detailed, security-level data to estimate a broader range of channels of how LSAPs could affect longer-term yields. They estimate that through both the supply of long-duration assets (12 basis points) and local scarcity channels (23 basis points), QE1 reduced the 10-year Treasury yield by 35 basis points (p. 441). Similar calculations for QE2 imply an estimated reduction in the 10-year Treasury yield of 45 basis points (pp. 441–442; 10 basis points from duration, 35 basis points from scarcity). The sum implies an 80-basis point reduction in long-term Treasury yields from the Treasury components of QE1 and QE2.

Given that the size of the Federal Reserve's balance sheet should naturally increase with the size of the economy, many of the estimates in Table B-1 are based on the supply of duration relative to nominal

Table B-1:
Breakdown of Estimates of Asset Purchases on 10-Year Treasury Yield

Study	Event analyzed	Duration-adjusted purchases (billions of dollars, 10-year equivalents)	Estimated effect on 10-year yield (basis points)	Estimated effect on 10-year yield (basis points / billions of dollars, 10-year equivalents)	Calculation of estimated effect ×100 (basis points / $100 billion, 10-year equivalents)
Greenwood and Vayanos (2014)	QE1, QE 2	$569	–13	–13/569	–2.3
Hamilton and Wu (2012)	MEP	$400	–14	–14/400	–3.5
Swanson (2011)	QE2	$400	–15	–15/400	–3.8
Krishnamurthy and Vissing-Jorgensen (2011)	QE2	$400	–18	–18/400	–4.5
Gagnon and others (2011)	QE1	$750	–34	–34/750	–4.5
Hanson (2014)	QE1	$750	–54	–54/750	–7.2
Li and Wei (2013)	QE1, QE2, MEP	$1,550	-150	-150/1550	–9.7
D'Amico and others (2012)	QE1, QE2	$569	–80	–80/569	–14.1
Median estimate					–4.5

GDP. To arrive at a measure of the effect of the duration-adjusted size of the Federal Reserve's balance on the 10-year Treasury yield, as shown in Chart 5, we therefore scale our measure of the duration-adjusted size of the Federal Reserve's balance sheet by the size of the economy before applying these estimates. Intuitively, scaling the Federal Reserve's duration-adjusted balance sheet by nominal GDP is sensible: some of the studies use samples dating as far back as the 1950s, and today's economy is considerably larger. In particular, we benchmark these estimates to 2011:Q4 and then scale down the effects of $100 billion in 10-year equivalents in future years by the ratio of nominal GDP in 2011 to nominal GDP in more recent years. We choose 2011:Q4 as a reference period because the above estimates are generally linked to the 2011 timeframe when QE2 and the MEP were ongoing.

Endnotes

[1]The debt and MBS issued by these agencies became officially backstopped by the federal government once they went into conservatorship in September 2008 (Rappaport 2020).

[2]Vayanos and Vila (2021) and Doh (2010) provide theoretical foundations for this channel of asset purchases.

[3]One channel that these event studies capture that is likely relevant in normal times is the way FOMC announcements influence expectations about future adjustments in the supply of long-duration assets. Although we do not capture this channel in our approach, our estimates should eventually capture the duration effects of announced purchases once the purchases are completed.

[4]For example, D'Amico and others (2012) study not just duration effects, but also local scarcity effects of asset purchases. The scarcity channel emphasizes that investors may have difficulty replacing the particular assets purchased by the Federal Reserve, which would increase their price and, for bonds, lower their yield.

[5]Sengupta and Smith (2022) provide another rationale for moving gradually to shrink the balance sheet at least initially: the comparatively unsettled state of financial markets in 2022 relative to 2017.

[6]These are approximate values taken from the Federal Reserve Bank of New York's report prepared for the FOMC on "Open Market Operations during 2021" (Federal Reserve Bank of New York 2022).

[7]As discussed in the text, from 2017 through 2019, the par value of the Federal Reserve's asset holdings declined by about $700 billion. However, we estimate that the balance sheet only declined by about $290 billion in terms of 10-year equivalents. We apply this ratio of 0.41 (= $290 / $700) to the projections for a par value decline of $2.5 trillion from 2022 through 2025 to arrive at a projected decline in terms of 10-year equivalents of roughly $1 trillion.

[8]This estimate of the effects of the projected balance sheet runoff on the 10-year Treasury yield is similar though a bit smaller than those from the Crawley and others (2022).

References

Board of Governors of the Federal Reserve System. 2022a . "Minutes of the Meeting of the Federal Open Market Committee, March 15–16, 2022."

———. 2022b. "History of the FOMC's Policy Normalization Discussions and Communications." Last updated May 4.

Crawley, Edmund, Etienne Gagnon, James Hebden, and James Trevino. 2022. "Substitutability between Balance Sheet Reductions and Policy Rate Hikes: Some Illustrations and a Discussion." Board of Governors of the Federal Reserve System, FEDS Notes, June 3. Available at https://doi.org/10.17016/2380-7172.3147.

D'Amico, Stefania, William B. English, David Lopez-Salido, and Edward Nelson. 2012. "The Federal Reserve's Large-Scale Asset Purchase Programs: Rationale and Effects." *Economic Journal*, vol. 122, no. 564, pp. 415–446. Available at https://doi.org/10.1111/j.1468-0297.2012.02550.x

D'Amico, Stefania, and Thomas B. King. 2013. "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply." *Journal of Financial Economics*, vol. 108, no. 2, pp. 425–448. Available at https://doi.org/10.1016/j.jfineco.2012.11.007

Doh, Taeyoung. 2010. "The Efficacy of Large-Scale Asset Purchases at the Zero Lower Bound." Federal Reserve Bank of Kansas City, *Economic Review*, vol. 95, no. 2, pp. 5–34.

Federal Reserve Bank of New York. 2022. "Open Market Operations during 2021." May.

Gagnon, Joseph E. 2016. "Quantitative Easing: An Underappreciated Success." Peterson Institute for International Economics, Policy Brief no. 16-4, April.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian P. Sack. 2011. "The Financial Market Effects of the Federal Reserve's Large-Scale Asset Purchases." *International Journal of Central Banking*, vol. 7, no. 1, pp. 3–43.

Greenwood, Robin, Samuel G. Hanson, Joshua S. Rudolph, and Lawrence H. Summers. 2016. "Debt Management Conflicts between the U.S. Treasury and the Federal Reserve," in David Wessel, ed., *The $13 Trillion Question: How America Manages Its Debt*. Brookings Institution Press.

Greenwood, Robin, and Dimitri Vayanos. 2014. "Bond Supply and Excess Bond Returns." *Review of Financial Studies*, vol. 27, no. 3, pp. 663–713. Available at https://doi.org/10.1093/rfs/hht133

Hamilton, James D., and Jing Cynthia Wu. 2012. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." *Journal of Money, Credit, and Banking*, vol. 44, pp. 3–46. Available at https://doi.org/10.1111/j.1538-4616.2011.00477.x

Hanson, Samuel G. 2014. "Mortgage Convexity." *Journal of Financial Economics*, vol. 113, no. 2, pp. 270–299. Available at https://doi.org/10.1016/j.jfineco.2014.05.002

Ihrig, Jane E., Ellen E. Meade, and Gretchen C. Weinbach. 2015. "Rewriting Monetary Policy 101: What's the Fed's Preferred Post-Crisis Approach to Raising Interest Rates?" *Journal of Economic Perspectives*, vol. 29, no. 4, pp. 177–198. Available at https://doi.org/10.1257/jep.29.4.177

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity*, vol. 2011, no. 2, pp. 215–287. Available at https://doi.org/10.1353/eca.2011.0019

Li, Canlin, and Min Wei. 2013. "Term Structure Modelling with Supply Factors and the Federal Reserve's Large Scale Asset Purchase Programs." *International Journal of Central Banking*, vol. 9, no. 1, pp. 3–39. Available at https://doi.org/10.2139/ssrn.2191189

Logan, Lorie K. 2022. "Federal Reserve Asset Purchases: The Pandemic Response and Considerations Ahead." Remarks at New York University's Stern School of Business (delivered via videoconference), March 2.

Potter, Simon. 2016. "Discussion of 'Evaluating Monetary Policy Operational Frameworks' by Ulrich Bindseil." Remarks at the 2016 Economic Policy Symposium at Jackson Hole, WY, August 26.

Rappaport, Jordan. 2020. "What to Do about Fannie and Freddie: A Primer on Housing Finance Reform." Federal Reserve Bank of Kansas City, *Economic Review*, vol. 105, no. 2, pp. 5–29. Available at https://doi.org/10.18651/ER/v105n2Rappaport

Sengupta, Rajdeep, and A. Lee Smith. 2022. "Assessing Market Conditions ahead of Quantitative Tightening." Federal Reserve Bank of Kansas City, *Economic Bulletin*, July 11.

Smith, A. Lee, and Victor J. Valcarcel. 2022. "The Financial Market Effects of Unwinding the Federal Reserve's Balance Sheet." Federal Reserve Bank of Kansas City, Research Working Paper no. 20-23, January. Available at https://doi.org/10.18651/RWP2020-23

Swanson, Eric T. 2011. "Let's Twist Again: A High-Frequency Event-Study Analysis of Operation Twist and Its Implications for QE2." *Brookings Papers on Economic Activity*, vol. 2011, no. 1, pp. 151–188. Available at https://doi.org/10.1353/eca.2011.0006

Vayanos, Dimitri, and Jean-Luc Vila. 2021. "A Preferred-Habitat Model of the Term Structure of Interest Rates." *Econometrica*, vol. 89, no. 1, pp. 77–112. Available at https://doi.org/10.3982/ECTA17440

Vissing-Jorgensen, Annette. 2021. "The Treasury Market in Spring 2020 and the Response of the Federal Reserve." *Journal of Monetary Economics*, vol. 124, pp. 19–47. Available at https://doi.org/10.1016/j.jmoneco.2021.10.007

Williams, John C. 2013. "Lessons from the Financial Crisis for Unconventional Monetary Policy." Panel discussion at the NBER Conference, Boston, MA, October 18.

Economic Review

Volume 107, Number 4

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