



EIEF Working Paper 17/05
April 2017

**Short-run effects of lower productivity growth.
A twist on the secular stagnation hypothesis**

by

Olivier Blanchard

(Peterson Institute for International Economics)

Guido Lorenzoni

(Northwestern University)

Jean-Paul L'Huillier

(EIEF)

Short-run effects of lower productivity growth. A twist on the secular stagnation hypothesis

Olivier Blanchard Guido Lorenzoni

Jean Paul L'Huillier*

January 30, 2017

Abstract

Since 2010, U.S. GDP growth has been anemic, averaging 2.1% a year, and this despite interest rates very close to zero. Historically, one would have expected such low sustained rates to lead to much stronger demand. They have not. For some time, one could point to plausible culprits, from a weak financial system to fiscal consolidation. But, as time passed, the financial system strengthened, fiscal consolidation came to an end, and still growth did not pick up. We argue that this is due, in large part, not to legacies of the past but to lower optimism about the future, more specifically to downward revisions in forecast potential growth. Put simply, the anticipation of a less bright future is leading to temporarily weaker demand. If our explanation is correct, it has important implications for policy and for forecasts. It may weaken the case for secular stagnation, as it suggests that the need for very low interest rates may be partly temporary.

*Prepared for AEA meetings, January 2017. Thanks to Julien Acalin and Colombe Ladreit for excellent research assistance. Thanks to PIIIE colleagues for comments, and to Flint Brayton for help with the FRB/US model and for comments on the draft.

Since 2010, U.S. GDP growth has been anemic, averaging 2.1% a year, and this despite interest rates very close to zero. Historically, one would have expected such low sustained rates to lead to much stronger demand. They have not.

For some time, one could point to plausible culprits, from a weak financial system to fiscal consolidation. But, as time passed, the financial system strengthened, fiscal consolidation came to an end, and still growth did not pick up.

We believe that this is due, in large part, not to legacies of the past but to lower optimism about the future, more specifically to downward revisions in forecast potential growth. Put simply, the anticipation of a less bright future is leading to temporarily weaker demand.¹

If this explanation is correct, it has important implications for policy and for forecasts. It may weaken the case for secular stagnation, as it suggests that the need for very low interest rates may be partly temporary. Put another way, to the extent that investors in financial markets have not taken this undershooting into consideration, the current yield curve may underestimate the strength of future demand, and the need for higher interest rates in the future.²³

Our paper is organized in three sections:

The first section looks at the historical relation between revisions in long-run potential growth forecasts and unexpected movements in consumption and investment. It finds a surprisingly strong relation between the two. Based on data going back to 1991, revisions have typically been associated with forecast errors in consumption and investment of the same sign. Under the assumption that forecasts of potential growth far in the future are not affected by cyclical movements in output, this relation can be interpreted as causal. The effect is large, between 4 and 6% of consumption for example for a 1% downward revision in future potential growth.

¹We explored, both conceptually and empirically, the idea that short run fluctuations may be partly due to news about the future in Blanchard et al (2013).

²To be clear, our hypothesis is not an alternative to the secular stagnation hypothesis, but a twist on it. Namely, we do not question that interest rates will probably lower in the future than they were in the past. We argue that part of the current low rates may be temporary and thus, go away.

³We also believe that the productivity growth forecasts may be too pessimistic, and that growth, and by implication, interest rates, may end up stronger than currently expected. This is however is a different argument.

The second section reviews the different channels through which lower forecasts of long-run productivity growth can affect output and inflation in the short run. On the supply side, workers may be reluctant to accept a slower increase in real wage growth, leading to a temporary increase in the natural rate of unemployment. On the demand side, lower potential income growth may lead consumers to revise their assessment of permanent income; lower growth of demand, combined with temporarily lower profits may lead firms to revise their investment plans. Both may lead to an increase in actual unemployment. What happens to inflation depends on the balance of increases in the natural and the actual rate of unemployment.

The third section tries to assess the size of the effects by using the FRB/US model, which embodies most of the mechanisms described above. Under backward looking expectations, the adverse supply effects dominate: Unemployment goes up, and so does inflation. Under forward looking expectations however, the adverse demand effects dominate: The actual unemployment rate increases, and does so by more than the natural unemployment rate, leading to a decrease in inflation (all relative to baseline). The interest rate rule leads the policy rate to decline for some time, before partly recovering to its (lower) steady state value.

1 Potential growth revisions, consumption and investment

Using US data, we examine first whether revisions in long-run potential growth have historically been associated with weaker unexpected demand growth.

- For potential growth, we use forecasts from the Congressional Budget Office (CBO) (which appears to be the only source of such forecasts for a long enough period). Each January since 1991, CBO publishes forecasts of annual potential growth for the next N years. (It is reasonable to assume that, were professional forecasters to regularly publish long-run potential growth forecasts, they would be rather similar). Until 1995, N was equal to 5. Since then, N is equal to 10.

Let the annual forecast of potential growth in year $t + N$ as of January of year t be denoted by F_t , and the first difference in the series be denoted

by DF_t . This is not quite a pure forecast change, as the terminal year goes up by one from year t to year $t + 1$, but it is close.

Because of the change in the terminal year in 1996, we construct three different versions of DF_t . The first is the change in the forecast for year 5, which is available for the 24 years of the sample, which we denote $DF5$. The second is the change in the forecast for year 10, which is available only for 19 years, which we denote $DF10$. The third is the change in the forecast for the furthest available year, so 5 years until 1995 and 10 years thereafter, which we denote by DFn .

- For forecasts of consumption and investment, both non residential and residential, we use data from the Survey of Professional Forecasters (SPF). The forecasts are quarterly, so we can use forecasts of spending for year t as of various quarters of year $t - 1$. We construct forecast consumption growth in year t as the forecast of consumption for year t divided by the forecast of consumption for year $t - 1$, minus one (and similarly for investment). We denote forecast consumption growth for year t as of the i 'th quarter of year $t - 1$ by FCi_t . We define the corresponding forecast error for consumption growth in year t by $UCi_t \equiv C_t - FCi_t$, and similarly for non residential and residential investment, $UINRi_t$ and $UIRi_t$.

When looking at the relation between the forecast error for consumption and the change in the forecast of potential growth, what precise timing should one adopt? We can think of two extreme assumptions:

- One assumption is that CBO learns about potential growth throughout year $t - 1$ and publishes its forecast in the first quarter of year t , and that consumers (firms) learn about the change in potential growth forecasts only when CBO announces them and so adjust their consumption over year t .⁴
- The other assumption is that people learn roughly at the same time as CBO does, so that they adjust their consumption throughout year

⁴The usual remark about people and firms not reading the CBO publication applies. The assumption is that these forecasts find their way in the press, and in turn, in people's expectations about the future.

$t - 1$ rather than during year t , and do not learn further from the CBO announcement.⁵

Both assumptions are extreme, and this suggests the following encompassing specification. We assume that consumers learn both as future potential growth evolves and from the CBO announcement itself, in proportions a and $(1 - a)$. During year t , they learn both from the announcement in January of year t , and from the evolution of future potential growth during year t , which will eventually be reflected in the CBO announcement in January of year $t + 1$. So, in year t , they learn $(1 - a)DF_t + aDF_{t+1}$. This leads us to specify regressions of the type:

$$UC1_t = b[(1 - a)DFx_t + aDFx_{t+1}] + \epsilon_t \quad (1)$$

or equivalently:

$$UC1_t = ab(DFx_{t+1} - DFx_t) + bDFx_t + \epsilon_t \quad (2)$$

where DFx stands for either $DF5$, $DF10$, or DFn , and $UC1_t$ is the forecast error for C_t . As professional forecasters presumably access to the same information as CBO and take it into account in their forecasts, the forecasts of consumption and investment must be as of a time when information about changes in potential growth revealed in year $t - 1$ is not yet available. Ideally, we would want to use forecasts as of quarter 4 of year $t - 2$, but these do not exist. Thus we use forecasts of consumption and investment as of quarter 1 of year $t - 1$ (and use a similar approach for non-residential and residential investment).

The results are given in Table 1 for consumption, and in Tables 2 and 3 for non-residential and residential investment.

The three columns of Table 1 give the regression results for consumption based on the three versions of the revisions in forecast long-run growth, $DF5$, $DF10$, DFn . Given the specification, the coefficient on $DF_{t+1} - DF_t$ is equal to ab . The coefficient on DF_t is equal to a . The implied values of a and b are given at the bottom of the table.

⁵Here again, people and firms do not think about potential growth per se. But consumers may get signals that their labor income may not grow as fast as they thought, or firms may get signals that demand growth is likely to be lower than anticipated.

Table 1: Consumption forecast errors and changes in potential growth forecasts

	(DF5)	(DF10)	(DFn)
	UC_t	UC_t	UC_t
$DF_{t+1} - DF_t$	2.171* [1.081]	4.429*** [1.433]	2.529** [1.006]
DF_t	4.333** [1.544]	6.791*** [1.428]	4.642*** [1.345]
Constant	0.243 [0.316]	0.052 [0.313]	0.270 [0.296]
Observations	24	19	24
R-squared	0.185	0.469	0.274
estimated a	0.50	0.65	0.54
estimated b	4.33	6.79	4.64

Robust standard errors in brackets. ***: $p < 0.01$, **: $p < 0.05$, *: $p < 0.1$

All three specifications yield a value of a close to half, suggesting that people learn both during the year and following the CBO announcement. The coefficient b is large, between 4.3 and 6.8 (the larger estimate when $DF10$ is used reflects the effects of the shorter subsample rather than the effect of the longer horizon. When all three regressions are run over the shorter sample, they all give roughly the same numbers, about 0.5 for a and 6.0 for b). This implies that a 1% downward revision in long-run potential growth is associated with a 4.3 to 6.8% reduction in consumption growth, or, assuming a ratio of consumption to GDP of 68%, a reduction of GDP growth of 2.7% to 4%.

Figure 1 shows the scatter diagram corresponding to the specification in the second column of Table 1. The variable on the horizontal axis is the expression in brackets on the right hand side of equation 1, thus a weighted average of the CBO revisions in year t and $t + 1$. The variable on the vertical axis is the variable on the left hand side of the equation, the forecast error in consumption for year t . Note that most of the points are in the south west quadrant: most years have been associated with downward revisions of potential growth, and negative forecast errors for consumption. 2008 and 2009 are outliers, with large negative forecast error for consumption relative to the

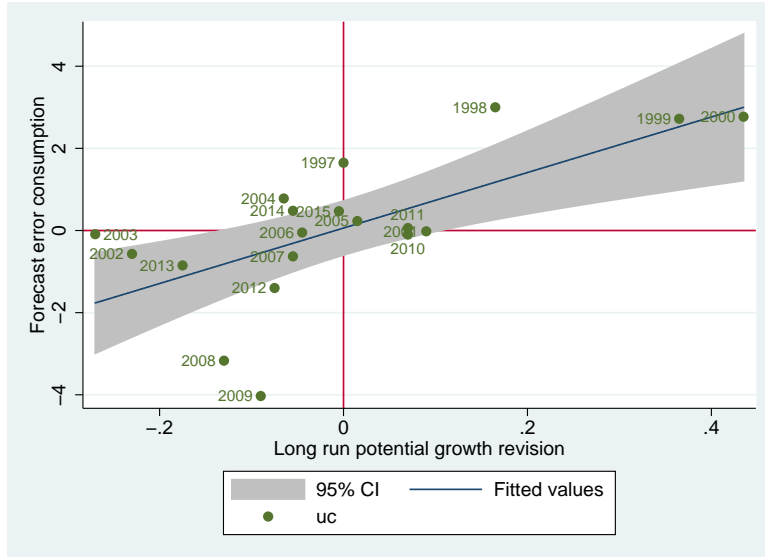


Figure 1: Consumption forecast error and Revisions of long term growth

revisions of long run growth.

An obvious issue in interpreting the results of Table 1 is that of causality. While causality from unexpected consumption to long-run potential growth revisions seems implausible, a more relevant possibility is that unexpected current income affects both unexpected consumption and the revision in long-run potential growth. The only way to settle the issue would be to find variables which affect long-run growth forecast revisions and are uncorrelated with current income. We could not think of any.⁶ The description given by CBO of the way they construct those forecasts, based on a production function approach, indicates that they try to eliminate cyclical effects.⁷ The choice we make of looking at forecasts of potential growth in year 5 or 10 also reduces the likelihood that they are strongly affected by cyclical movements. A year for which the issue seems perhaps most serious is 2009, in which sharply negative growth could have affected long-run potential growth forecasts. 10-year out potential growth was indeed revised down by 0.23%, 5-year out by 0.04%. We therefore ran the regression leaving out 2009; the results are very similar to those reported above.

⁶Demographic variables may be good instruments for the growth forecasts themselves, but not for changes in those forecasts: Their evolution from year to year is largely expected.

⁷See for example “Revisions to CBO’s projection of potential output since 2007” (2014).

Table 2: Non-residential investment forecast errors and changes in potential growth forecasts

	DF5	DF10	DFn
	$UNRI_t$	$UNRI_t$	$UNRI_t$
$DF_{t+1} - DF_t$	5.631 [3.538]	15.345** [7.156]	8.487** [3.961]
DF_t	10.401** [4.513]	19.762*** [5.519]	13.491** [4.945]
Constant	-0.829 [1.258]	-1.793 [1.288]	-0.760 [1.181]
Observations	24	19	24
R-squared	0.077	0.324	0.173
estimated a	0.54	0.78	0.63
estimated b	10.40	19.76	13.49

The three columns of Table 2 and Table 3 give the regression results respectively for non-residential investment and for residential investment, based on the three versions of the revisions in forecast long-run growth, $DF5$, $DF10$, DFn . The point estimates suggest large effects, although they are less statistically significant for non-residential investment than for consumption, and they are not significant in the case of residential investment. A 1% downward revision in potential long-run growth leads to a 10 to 20% unexpected decrease in non-residential investment; assuming a ratio of non-residential investment to GDP of 12%, this implies an unexpected decrease in GDP of 1.2 to 2.1%. A 1% downward revision in potential long-run growth leads to a 6 to 11% unexpected decrease in residential investment; assuming a ratio of residential investment to GDP of 4%, this implies an unexpected decrease in GDP of 0.3 to 0.4%.

Putting all these numbers together, what do the point estimates suggest as the contribution of downward revisions to unexpected demand growth over the past few years? Since 2012, the CBO has revised its estimates of its 5-year ahead potential growth from 2.64% to 2.10%, so a decrease of 0.54% (peak forecast was 3.23% in 2001).⁸ Adding the estimates for consumption and

⁸Over the same period, 10-year ahead potential growth has been revised down from

investment above, this suggests implied an unexpected decrease in demand of 2.3 to 3.5%, or about 0.6% to 0.9% a year for the last four years.⁹

Table 3: Residential investment forecast errors and changes in potential growth forecasts

	DF5	DF10	DFn
	URI_t	URI_t	URI_t
$DF_{t+1} - DF_t$	1.168 [6.374]	6.825 [10.709]	2.916 [6.380]
DF_t	7.784 [8.290]	11.374 [11.207]	5.868 [8.780]
Constant	-1.573 [2.091]	-2.845 [2.477]	-1.522 [2.073]
Observations	24	19	24
R-squared	0.025	0.038	0.012
estimated a	0.15	0.60	0.50
estimated b	7.78	11.37	5.87

2.37% to 1.91%, so a decrease of 0.46%.

⁹One might guess that the effect is larger when monetary policy is limited by the zero lower bound, as has been the case since 2012. There are not enough observations to estimate the equations over the crisis sample period. Estimation over the pre-crisis period (up to and including 2007) gives smaller, but not significantly smaller, coefficients. More on the potential role of the zero lower bound in Section 3.

2 Mechanisms

As demographic trends are largely predictable and capital accumulation plays a limited role in determining potential growth in the long run, revisions to long-run potential growth are mostly due to revisions to total factor productivity growth.

In the standard optimal growth model, the Euler equation for consumers implies that, in the new steady state, the decrease in productivity growth is reflected in a lower interest rate. This in turn implies a lower marginal product of capital, and thus an increase in the capital stock in efficiency units. The path of adjustment implies a smooth decrease in the marginal product of capital over time, and, typically, a smooth decrease in the interest rate over time.¹⁰

When account is taken of nominal rigidities and other distortions, a number of other relevant effects are at work and are likely to dominate.

On the supply side, while a decrease in real wage growth is eventually needed, both real and nominal wage rigidities may slowdown the adjustment, leading to a temporary increase in the natural rate of unemployment. This effect, together with the effect of higher oil prices, is seen as having been central in triggering stagflation of the 1970s (see for example Bruno and Sachs 1985).

On the demand side, for a given sequence of interest rates, aggregate demand is likely to decrease as well. Lower income growth implies a decrease in the present value of income, leading consumers to cut spending. Given lower expected demand growth, firms are also likely to cut on investment. If lower productivity growth comes from lower embodied technological progress in new capital, then two other effects are at work. The first is that, as technological obsolescence is reduced, old vintages that would have been discarded may now be kept longer, leading to a slowdown in investment for some time. The second, working in the opposite direction is that lower technological obsolescence implies a smaller decrease in the price of a given vintage of capital over time, and thus a lower user cost, leading to higher investment. Overall, the presumption

¹⁰“Typically” reflects the fact that, in the model extended to allow for habit persistence and costs of adjustment for investment, the interest rate may initially overshoot its long-run decrease. The reason is that, in order to generate an increase in investment, the shadow value of capital (Tobin’s q) may have to increase a lot, which in turn may require an interest rate far below the marginal product of capital for some time.

is that demand goes down, leading to an increase in unemployment.

Both the supply and the demand effects imply an increase in unemployment for some time. What happens to inflation depends on the balance between the two. If the increase in the underlying natural rate of unemployment exceeds the increase in the actual rate, for example if consumers and firms do not take into account right away the decrease in underlying growth so the demand effect is limited, inflation will increase. If instead, the fall in demand is sufficiently large, and the increase in the actual unemployment rate exceeds the increase in the natural rate, inflation will decrease.

This discussion of the dynamics has taken interest rates as given. Monetary policy is likely however to adjust as well. Lower interest rates can limit the effects on demand. This suggests that optimal monetary policy may take the form of a large initial decrease in interest rates, with a return to the new lower steady state in the long run.

How can we tell which of these effects dominate, whether they are consistent with the reduced form empirical evidence presented earlier, and what the path of interest rates may be under plausible interest rate rules? By simulating a model which has most of the mechanisms described above. This is what we do in the next section.

3 Simulations

The model we use for simulations is the FRB/US model developed at the Federal Reserve Board.¹¹

We use that model for two reasons. The first is that it includes the mechanisms described in the previous section—except for the vintage structure of capital. Consumers’ spending depends in part of expectations of future income and future interest rates. Firms’ investment depends in part on expectations of future output and future user costs. The wage and price equations reflect both nominal and real rigidities. The second reason is that it allows for a flexible treatment of expectations. In particular, the model can be run under “VAR” expectations, in which expectations depend on past variables, or under model consistent “MC” expectations, in which various agents, firms and

¹¹We use the November 21, 2014 package of FRB/US, downloadable from the Federal Reserve Board website. See www.federalreserve.gov/econresdata/frbus/us-models-about.htm. We benefited from a further update to the package, kindly provided to us by Flint Brayton.

workers, financial investors, consumers and firms, are assumed to have (nearly) rational, forward looking, expectations.

We present three simulations below.

In all three, the deviation from the baseline is a decrease in the rate of growth of total factor productivity of 1% a year. In order to get as close numerically to a permanent shock to the growth rate as feasible, we take the growth rate to follow an AR(1) process, with a decrease of 1% at the start of the simulation, and an AR(1) coefficient of 0.995 per quarter.

The model allows for different interest rate rules. We use the “inertial Taylor rule”, which allows for a slow adjustment of the policy rate to the deviation of inflation from target and the deviation of output from a smooth trend.¹²

The simulations differ in the treatment of expectations. The first assumes backward looking (VAR) expectations. The other two assume forward looking (MC) expectations for all agents.¹³

The first two simulations allow the central bank to freely adjust the policy rate according to the interest rate rule. This may be the right assumption for the pre-financial crisis episodes. But one important aspect of the current episode is that the Fed has been facing a zero lower bound constraint on the policy rate. While unconventional policies have allowed the Fed to decrease other rates further, leading to negative “shadow policy rates”, it has not been able to decrease even these shadow rates as much as it desired. Thus, the third simulation imposes a lower bound on the decrease in the rate. It allows the policy rate to decrease relative to the baseline by no more than 2.4%, and assumes that the constraint is potentially binding for 5 years (implicitly assuming that, after five years, the baseline interest rate is sufficiently high as to allow for the interest rate rule to be used without constraints).¹⁴ We shall refer to this simulation as the LB (for “lower bound”, but not necessarily zero) simulation. The robustness of the resulting estimates to the specific choice of

¹²The model refers to this deviation as the “output gap”, but it is not the deviation of output from the natural level of output, defined as the level of output which would prevail in the absence of nominal rigidities.

¹³The crucial assumption is that of forward looking expectations for firms and consumers. Forward looking expectations in the determination of wage and price decisions, or even in the determination of asset prices, leads to results close to the VAR assumption.

¹⁴These choices are loosely based on the evolution of the shadow rate constructed by Wu and Xia, which decreased from 0% in 2009 to between -2% and -3% in 2014, and is now roughly back to 0%.

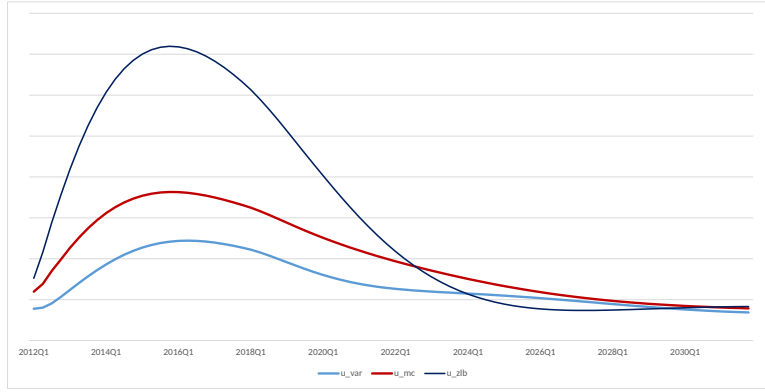


Figure 2: Unemployment rate

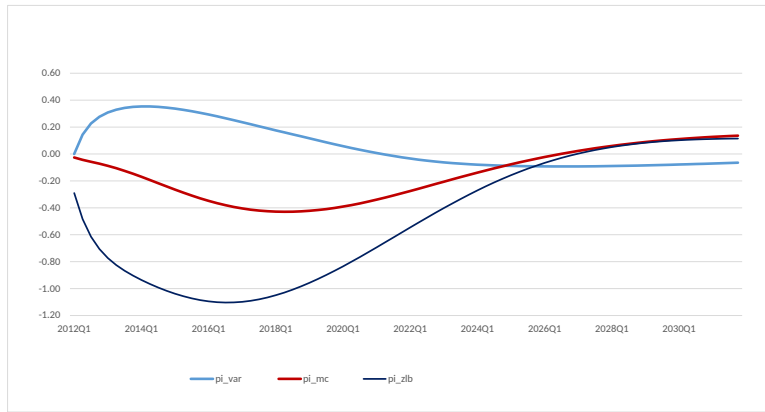


Figure 3: Core inflation rate

this lower bound is discussed below.

The results are presented in Figures 2 to 4.

Figure 2 shows the evolution of unemployment under each of the three assumptions. The maximum increase is smallest (0.7%) under the VAR assumption, larger (1.4%) under the MC assumption (as one would expect given that people and firms look forward and see a worse future), and, not surprisingly, largest (3.1%) under the LB assumption.

Figure 3 shows however that different effects are at work between the first and the other two simulations. In the first simulation, demand effects are limited as neither consumers nor firms look forward, and the increase in the natural rate exceeds the increase in the actual rate, leading to inflation pressure. Inflation increases, although only slightly. In the other two, the adverse effects dominate, leading to an unemployment gap and downward pressure on

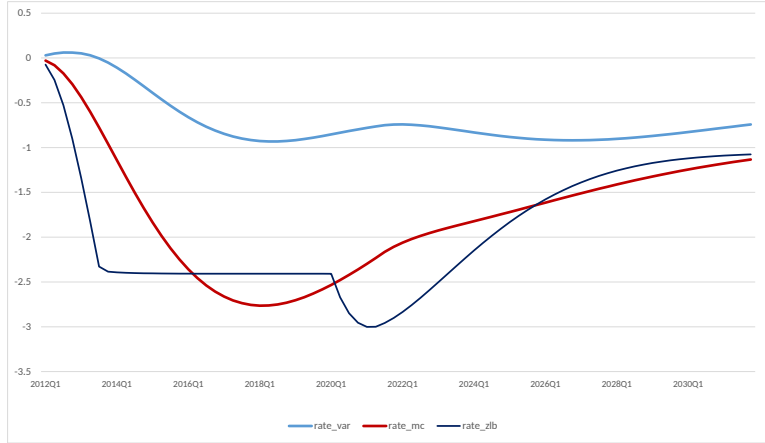


Figure 4: Policy rate

inflation. The downward pressure is stronger in the LB simulation, reflecting the largest increase in unemployment.

Figure 4 shows the behavior of the policy rate. In all three cases, the decrease in growth is associated, in the longer run, with a decrease in the policy rate.¹⁵ Not surprisingly, the initial decline in the policy rate is smallest under the VAR assumption (less than 1%). It is substantially larger under the MC assumption. The rate declines by nearly 3%, before slowly recovering to its (now lower) steady state value.

The behavior of the policy rate and the associated unemployment rate under the LB assumption is striking. Despite the fact that, relative to the unconstrained MC case, the constraint is effectively binding only for about four years, the effect on unemployment, from Figure 1, is quite dramatic. This is a reflection of the strong adverse dynamics of adverse shocks under LB constraints, in which an unemployment gap leads to lower inflation, higher real interest rates, and further increases in unemployment.¹⁶

In short, the simulations suggest that, with forward looking firms and people, a decrease in potential output growth can lead to low demand growth for some time. The effect can be very strong if monetary policy is constrained by a lower bound. They also suggest that the low interest rates, needed in

¹⁵The FRB/US model largely behaves in the long run like the standard Ramsey model, in which the Euler equation for consumers implies a tight relation between movements in the growth rate and movements in the interest rate

¹⁶Indeed, we found that if we allowed the lower bound to be potentially binding throughout (as opposed to ending after some time T), then the model typically did not solve, reflecting the fact that the economy got worse and worse, leading to no solution under rational expectations.

the short run to limit the decrease in demand, may increase substantially over time.

How consistent are the earlier reduced form results and the simulation results?

The reduced form results suggest that a 1% revision in long run potential growth in year t is perceived partly in year t and partly in year $t + 1$, and leads to an unexpected decrease in output of 4.2% to 6.5% at the end of year $t + 1$.

The simulations above imply that the same revision leads to an unexpected decrease in output at the end of year $t + 1$ of 2.9% under the VAR assumption, 3.9% under the MC assumption, and 6.7% under the LB assumption.¹⁷ The two sets of results are thus roughly consistent, although the simulations suggest a weaker effect than the reduced form when the LB is not binding, and a stronger effect than the reduced form when the LB constraint is binding.

4 Tentative conclusions

We have offered a tentative explanation for the weak US demand growth of the last four years. Namely, that, beyond crisis legacies, anticipations of a less optimistic future have led to temporarily weaker demand. Reduced form and simulation results suggest that downward revisions of productivity growth may have decreased demand by 0.5% to 1% a year since 2012. If we are right, it may well be that, as this adjustment comes to an end, this adverse effect will disappear, demand will pick up and interest rates will increase substantially. To the extent that this effect has not been taken fully into account by investors in financial markets, it may be that the current slope of the yield curve understates the increase in interest rates to come.

¹⁷The increase in the output gap is smaller, as after two years, lower productivity growth of 1% implies that potential output is at least 2% below the baseline. Using a 2% number, this implies output gaps of 0.9%, 1.9% and 4.7% respectively. Using an Okun coefficient of 2 to get the effect on unemployment, these imply an increase of 0.4, 1.0, and 2.4 percentage points, corresponding to the increases in the unemployment rate after two years reported in Figure 2.

References

- [1] Olivier Blanchard, Guido Lorenzoni, and Jean-Paul L'Huillier. News, noise, and fluctuations: An empirical exploration. *American Economic Review*, 103(7):3045–3070, December 2013.
- [2] Michael Bruno and Jeffrey Sachs. *Economics of worldwide stagflation*. Harvard University Press, Cambridge, Ma, 1985.
- [3] Jing Cynthia Wu and Fan Dora Xia. Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Chicago Booth Research Paper*, May 2015.
- [4] Congressional Budget Office. Revisions to CBO's projections of potential output since 2007. *CBO*, pages 1–21, 2014.