Where the Fed Goes Wrong:

The "Productivity Gap" and Monetary Policy

(Preliminary)

David Beckworth
Assistant Professor of Economics
Texas State University
San Marcos, TX 78666
db52@txstate.edu

George Selgin
Professor of Economics
Terry College of Business
University of Georgia
Athens, GA 30602
Selgin@uga.edu

April 20, 2010

Where the Fed Goes Wrong:

The "Productivity Gap" and Monetary Policy

Introduction

Many experts claim the housing boom of 2001-2007 was to a substantial degree the unintended consequence of the Fed's easy monetary policy stance in the wake of the 2001 dot.com crash. That stance involved setting the federal funds rate target at levels that proved, in retrospect, to be too low for too long. According to this view, both the extent of the housing boom and the severity of the consequent bust would have been considerably limited had the Fed pursued a less accommodative monetary policy (e.g. Taylor 2009).

Such claims raise the question, why did the Fed behave as it did? What caused the FOMC to select a path for the federal funds rate that now appears to have helped to inflate a housing-market bubble? What beliefs caused it to conclude that, in setting its targets as it did, it was merely helping to achieve a "soft landing" from the 2001 crash, and not setting the stage for another, ultimately more serious, round of boom and bust? Have similar beliefs caused it inadvertently to contribute to past booms and busts as well?

We propose a partial answer to these questions, and by doing so hope to assist monetary policymakers in avoiding similar mistakes in the future. Our thesis, in brief, is that the Fed's occasional, unintentional exacerbation of the business cycle is largely attributable to its failure to respond appropriately to persistent changes in the growth rate of total factor productivity. In particular, when that growth rate surged following the dot.com crash, the Fed treated the surge, not as a reason for adjusting its federal funds rate target upwards, as theory suggests it ought to have done, but as instead permitting it to maintain an unusually low federal funds rate by countering the inflationary tendencies such a low rate would otherwise have entailed.

Monetary Policy and Productivity

Monetary policy in the U.S., as well as in many other nations, has long been based on interest-rate targeting, that is, on selecting a target value for a particular interest rate, and adjusting the growth rate of the monetary base in a manner calculated to make the actual rate hit the target. In the Fed's case, the interest rate that's targeted is that for overnight interbank reserve borrowings or "federal funds." When the federal funds rate rises above its targeted value, the Fed's open-market desk responds by increasing its net purchases of government securities, thereby adding to the supply of bank reserves; when the rate falls below target, it reduces its net purchases, and may even engage in net sales. Because the cost of acquiring federal funds to cover temporary reserve shortages is an important determinant of banks' own willingness to lend, raising the federal funds rate target tends to check overall credit expansion, other things being equal, while lowering it tends to encourage it.

Responsibility for setting the target rate itself falls on the Federal Open Market Committee (FOMC). In general, the FOMC's challenge can be understood as one of achieving a "neutral" monetary policy stance, meaning one that is neither expansionary nor contractionary, that is, that avoids contributing either to booms or to busts (Bernhardsen and Gerdrup 2007). The funds rate consistent with such a monetary policy stance is sometimes referred to as the "neutral" (and sometimes as the "natural") rate of interest. When the stance of monetary policy is such that the inflation rate is always zero, the "nominal" and the "real" neutral rate of interest are identical. Otherwise the "real" neutral rate is equal to the nominal neutral rate minus the prevailing rate of inflation. If, for example, policymakers aim for a 2 percent inflation target, then to avoid aggravating the business cycle they must set the federal funds target at a rate equal to the real

neutral rate plus 200 basis points, and then provide for such open-market operations as are needed to achieve that target.

As this summary suggests, were the real neutral federal funds rate directly observable, implementing a "neutral" monetary policy would be a relatively simple matter. In fact, the rate is both directly unobservable and impossible to estimate precisely. Consequently, instead of simply tracking it and adjusting the target federal funds rate accordingly, the Fed adjusts its target in response to other, directly observable variables, such as the rates of inflation and unemployment, the values of which supply a rough (and, as we'll see, sometimes very imperfect) indication of whether some prevailing rate target is too low or too high.

But although the neutral federal funds rate isn't observable, the fundamental determinants of that rate are well known and uncontroversial, as they are implied by many standard economic models. According to the Ramsey model of economic growth, for example,

$$r^n = g + n + d \qquad (1)$$

where r^n is the real neutral rate of interest, g and n are the productivity and population growth rates, respectively, and d is the household rate of time preference, that is, the degree to which the representative household prefers immediate to delayed consumption (ibid., p. 53).

Of the three fundamental determinants of the neutral rate of interest, as implied by the Ramsey model, the productivity growth rate is the one most likely to be behind and to best reflect substantial movements in neutral rates of interest in the U.S. economy. Between 1970 and 2006, the total factor productivity (TFP) growth rate ranged from as high as 4 percent to as low as -3 percent, with many relatively sharp swings, the most recent of which has been the so-called productivity "surge" that began in the mid-1990s, was interrupted by the dot.com crash, and then

Figure 1

Total Factor Productivity Growth Rate

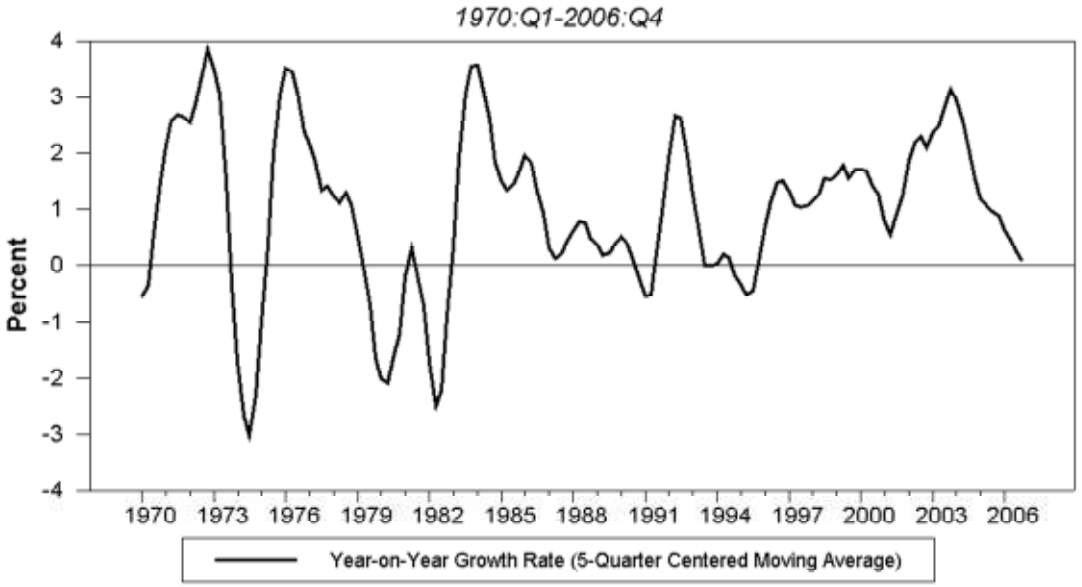

Source: Fernald (2009), Authors' Calculations

resumed with still greater vigor until 2004.[1] The U.S. population growth rate, in contrast, has been relatively stable during this time hovering around 1 percent.[2] The importance of the rate of time preference to the neutral rate is largely unknown since there are no reliable and consistent estimates of it (Frederick, Loewenstein, and O'Donoghue, 2002). However, to the extent that there are there are cyclical movements to time discounting the cyclical swings in the productivity growth rate may provide a proxy for them if not influencing them directly (e.g. a negative shock to productivity may alter time discounting by increasing uncertainty and changing preferences).

It follows from these considerations that the achievement of a neutral monetary policy stance in recent decades has depended to a considerable degree on the Fed's recognizing changes—especially persistent changes—to the U.S. productivity growth rate and adjusting its federal funds target appropriately in response to them.

[1] Gordon (2010) attributes the productivity "explosion" of 2001-2004 to a lagged effect of the information and communication technology boom of the late 1990s as well as aggressive cost cutting measures made by firms. Oliner, Sichel, and Stiroh (2007) make a similar argument.

[2] Most changes in the U.S. population growth rate have been within 28 basis points of this value.

The Productivity Gap

Although the real neutral federal funds rate cannot be measured precisely, our knowledge of the rate's fundamental determinants and also of its long-run, average value (which tends to be the same as its actual, long-run value) supplies the basis for various more-or-less crude estimates. One such crude estimate, suggested by the previous discussion, treats the real neutral rate as simply being equal to the long-run average real rate, which is generally assumed to be about 2 percent, plus the difference between the currently forecasted TFP growth rate and the mean rate of TFP growth.[3] In effect, this crude estimate treats fluctuations in the real neutral rate as depending *solely* on fluctuations in the rate of growth of productivity, as if fluctuations in other fundamental determinants of the neutral rate were of only minor importance.[4]

Among its advantages, besides that of being very simple, this proxy is available in "real" time, rather than derived ex-post, so that monetary policymakers might actually refer to it in setting their targets. Also, it serves, to highlight the crucial importance of productivity innovations, both in influencing the real neutral rate and in accounting for the Fed's past failure set its target rates correctly. Figure 2 shows a smoothed version of this real neutral federal funds rate along with the real federal funds rate for the period 1970:Q1 – 2006:Q4.[5]

In general, departures from optimal monetary policy can be understood as being measured by the spread between the real federal funds rate at any moment and the neutral real federal funds rate at that moment. Because in our simple framework the latter rate changes only owing to

---

[3] More formally, if equation (1) represents the long-run, steady-state values and current realizations of the neutral rate are denoted with time t subscripts then this approach would be stated as follows, $r_t^n = r^n + (g_t^e - g)$, where g_t^e is the expected TFP growth rate. See the data appendix for how we estimate g_t^e.

[4] Other, more sophisticated, estimates of the neutral real rate often show contradictory implications for the stance of monetary policy. For example, Laubach and Williams (2003) show an ongoing gradual decline in the U.S. neutral real rate from 2003 on while Lombardi and Sgherri (2007) show an increase. The sophisticated estimates are very sensitive to assumptions made in the estimation have caused some policymakers to question their usefulness (Ferguson, 2004). Our measure, by contrast, is simple, intuitive, and consistent with the stylized facts of the U.S. business cycle as we note in the paper.

[5] The real federal funds rate is constructed by subtracting the year-on-year inflation rate from the federal funds rate.

Figure 2
The Neutral and Actual Real Federal Funds Rate

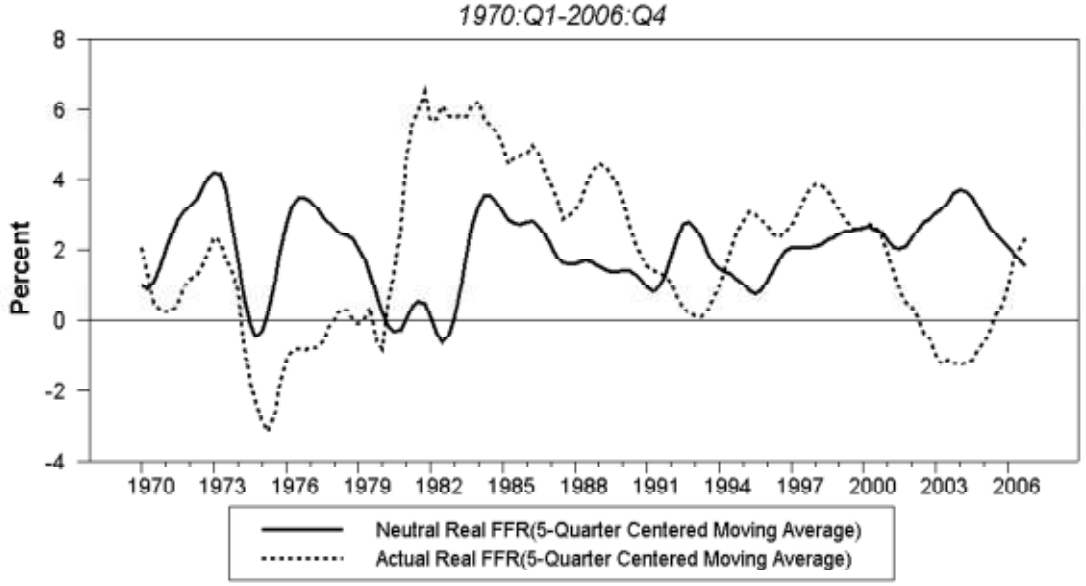

Source: Fernald (2009), FRED Database, Authors' Calculations

Figure 3
The Productivity Gap

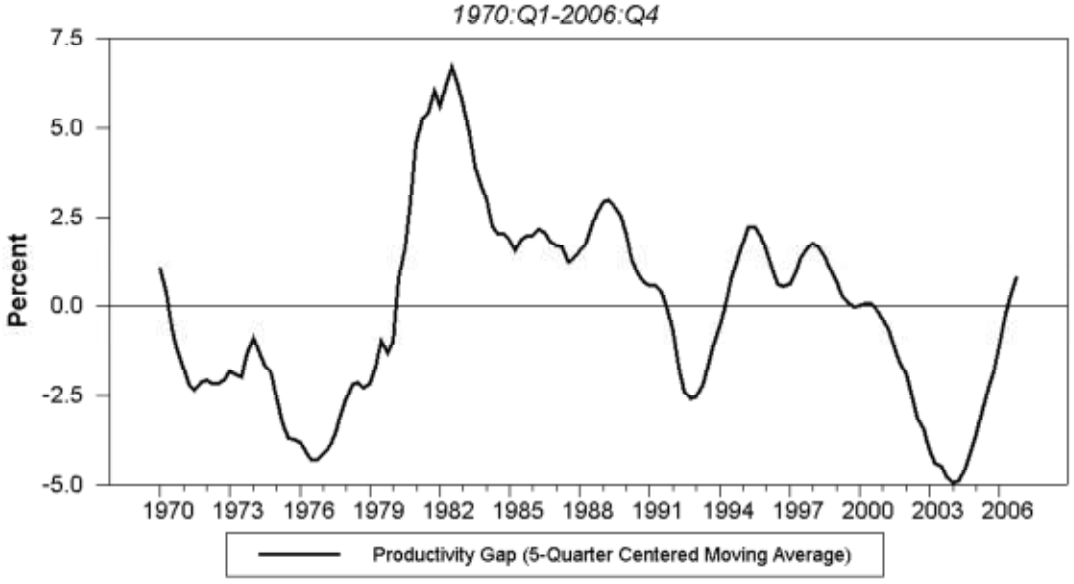

Source: Fernald (2009), FRED Database, Authors' Calculations

changes in the growth rate of productivity, we shall refer to this spread as the "Productivity Gap." Formally, we define it as follows:

$$P_t = r_t - r_t^n,$$

Where P_t is the current period Productivity Gap, r_t is the actual real federal funds rate, and r_t^n is the current period neutral real interest rate. A positive Productivity Gap means that the target real federal funds rate is too high relative to the assumed real neutral rate, so that monetary policy is too tight, whereas a negative Productivity Gap means, that monetary policy is too easy. Figure 3 shows this measure and suggests that the Fed's response to productivity innovations is occasionally inadequate, occasionally excessive, and—most significantly--occasionally perverse.

The direction and magnitude of monetary policy errors implied by the Productivity Gap are also at least roughly consistent with conventional wisdom. Thus Fed policy was, according to our measure, excessively easy throughout the 1970s, although less so in the immediate wake of the initial OPEC-induced oil supply shock than at other times, while it was excessively tight during the Volker anti-inflation campaign. During the nineties policy was at first easy and then somewhat (though not dramatically) tight. During the period immediately surrounding the tech bubble crash, monetary policy appears to have been neutral. Finally, beginning around 2002, monetary policy became increasingly easy as the Fed drove the real funds rate into negative territory despite strong productivity growth, with the Productivity Gap reaching its lowest value in the sample period at the height of the housing boom.

Figure 4

The Productivity and Output Gaps

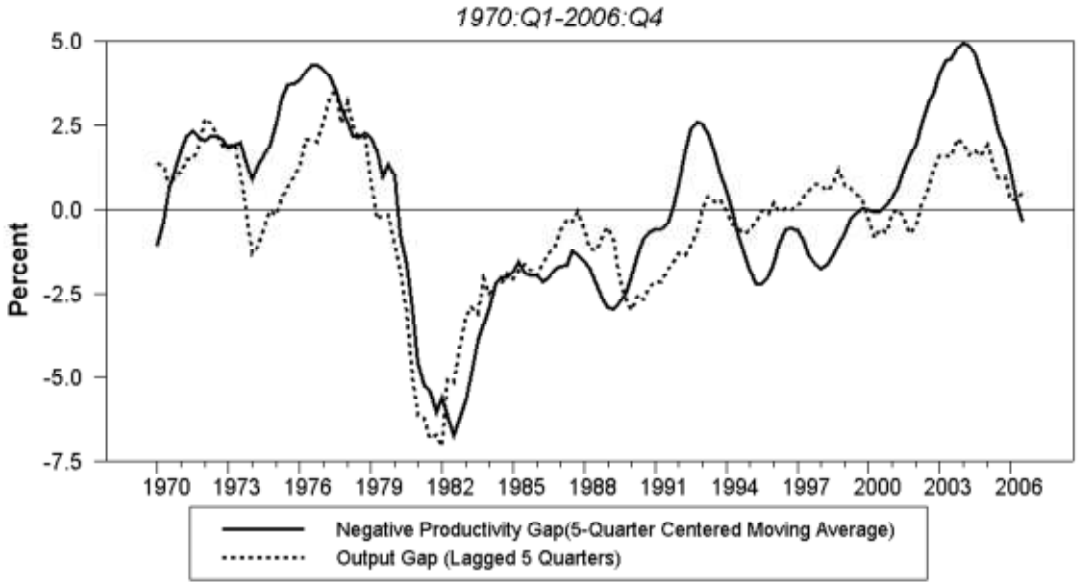

Source: Fernald (2009), FRED Database, Laubach and Williams (2003), Authors' Calculations

The Productivity Gap's merit as a rough indicator of the stance of monetary policy is further suggested by the degree to which it correlates with other such indicators, including estimates of the so-called "output gap" and measures of housing market activity. The "output gap" is defined as the difference between actual and natural or sustainable output,[6] and, as Berhardsen and Gerdrup (2007, p.54) and Williams (2003, p.1) observe, a neutral rate of interest can in fact be understood as one that serves to "close" the output gap. In contrast, the output gap will tend to be positive when interest rates are below their neutral levels, and negative when they are above. It follows that, if the Productivity Gap is a good measure of the gap between the actual and neutral federal funds rates, it should be negatively correlated with the output gap. Figure 4 tests this notion by graphing the Productivity Gap with the output gap lagged 5 quarters.[7] This figure shows, for sake of comparison, the negative of the Productivity Gap and the output gap

[6] We prefer these expressions to "potential" output, as output might be understood to be below its "potential" value whenever monetary policy might succeed in driving the economy up its short-run aggregate supply schedule, even if only temporarily.

[7] We use the Laubach and Williams (2003) output gap measure rather than the CBO's output gap since the latter's measure assumes the growth rate of output does not vary much in the short-run. See Weidener and Williams (2009) for more discussion on why the Laubach and Williams output gap is preferred over the CBO's measure.

Figure 5

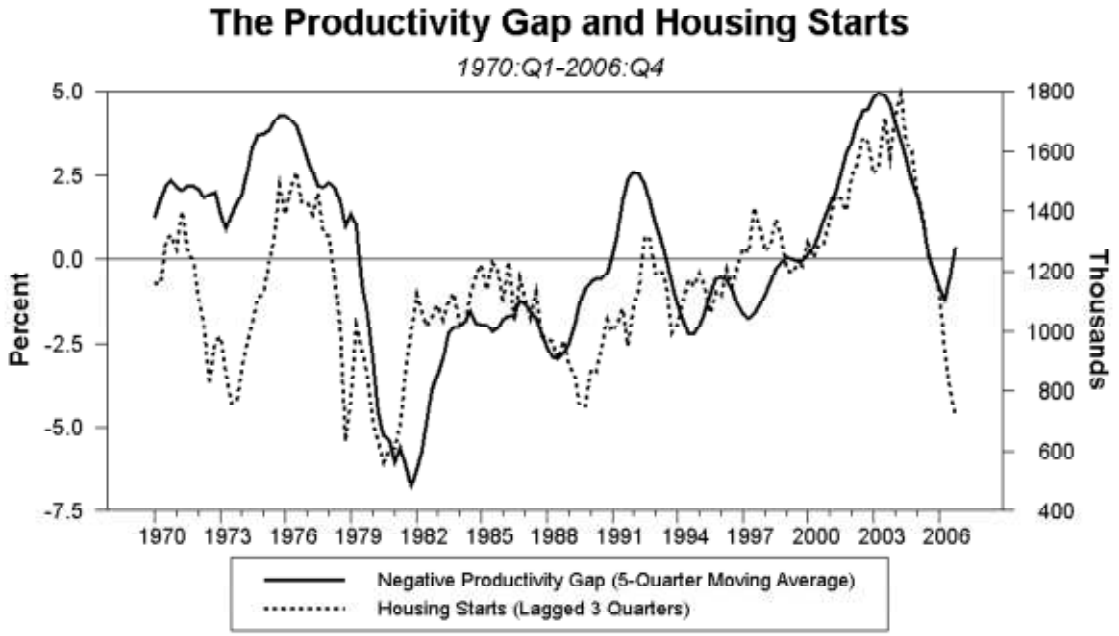

Source: Fernald (2009), FRED Database, Authors' Calculations

lagged and reveals they are in fact highly correlated. The R^2 between these two series is 66 percent, suggesting the Productivity Gap affects the output gap with a lag.

Theory also suggests that the demand for durable assets should be especially sensitive to gaps between neutral and actual rates of interest, for such gaps imply a difference between that rate at which funds can be borrowed and either the marginal product of capital or the rate at which households discount future utility flows relative to present ones. A negative gap in particular makes it generally worthwhile to acquire assets using borrowed funds, while a positive gap should discourage borrowing. If the Productivity Gap is a good proxy for the actual-real neutral rate gap, we should expect the demand for assets, and durable assets especially, to be negatively correlated with it. One measure of such demand of particular interest in light of recent events is the number of housing starts. Figure 5 show the relationship between the Productivity Gap and housing starts and suggest, once again, the former affects the latter with a lag. Here, the R^2 between these two series is 35 percent.

The Productivity Gap and the Taylor Gap

Rather than adjusting its federal funds target in response to either observed or anticipated deviations of productivity growth from trend, the Fed has tended to base its target rate adjustments mainly on observed and forecasted values of unemployment and the inflation rate. This tendency can be seen in Figure 6. Here, the cumulative responses of total factor productivity (TFP), the federal funds rate, the inflation rate, the unemployment rate, and the output gap to a typical (i.e. one-standard deviation) positive shock to the TFP growth rate are presented for the period 1970:Q1 – 2006:Q4.[8] Since this figure reports cumulative responses, it

Figure 6

Cumulative Response to a Typical TFP Shock

1970:Q1 – 2007:Q4

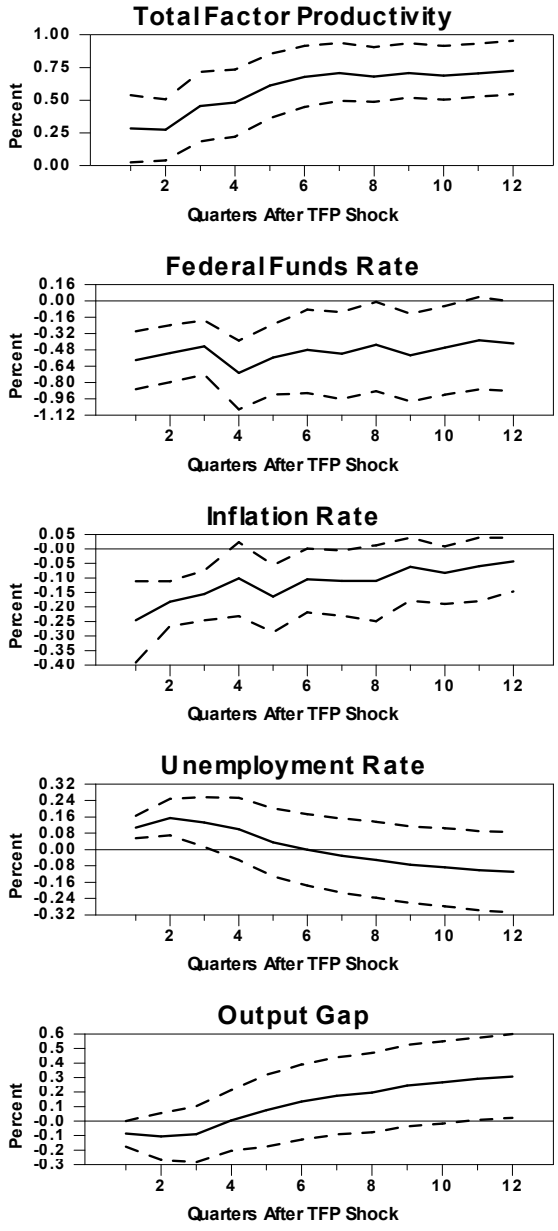

[8] These responses come from a vector autoregression (VAR) estimated over the sample period. The VAR uses 5 lags to eliminate serial correlation and identifies the TFP shock using long-run restrictions (i.e. only shocks to the TFP series can permanently affect itself). The TFP series is logged and first differenced while the other series—which are already in rate form—are just first differenced to induce stationarity. See the appendix for the data sources.

can be interpreted as showing the effect of the TFP growth rate shock on the level of these variables. The solid lines in the figures represent the average or point estimates, while the dashed lines show a simulated standard error band.[9]

Figure 6 shows that, ceteris paribus, the typical positive TFP shock during this time increased the level of TFP which, in turned, put downward pressure on the inflation rate. It also appears to have modestly increased unemployment, indicating firms may have temporarily used fewer workers given the productivity gains.[10] The Fed, however, responds by dropping the federal funds rate and holding it there so as to push the inflation and unemployment rates back to their pre-shock values. Doing so, however, means the federal funds rate is not allowed to reflect the changes in productivity—it is held below its neutral rate level. Consequently, this response pushes the economy beyond its sustainable output level as seen in the increase in the output gap.[11] This approach reflects the nature of the Fed's "dual mandate," which makes it directly responsible for both combating unemployment and stabilizing inflation, and not for maintaining a "neutral" policy stance as such, as well as the assumption, shared by most monetary economists, that a neutral target rate is also one that that will tend to preserve both a steady rate of inflation and a "natural" rate of unemployment.

In the early 1990s, John Taylor of Stanford University developed a simple formula that closely approximated the FOMC's actual target choices in the course of the preceding decade. According to this so-called Taylor Rule,

$$i^* = \pi + \tfrac{1}{2}(y-y_P) + \tfrac{1}{2}(\pi-2) + 2$$

[9] The standard error bands are calculated using standard Monte Carlo techniques.

[10] There is large debate as to whether productivity shocks in the short-run increase or decrease hours works. Our evidence suggests it leads to a decrease. See Whelan (2009) for a survey of the evidence.

[11] Unemployment also begins to drop below zero, but is not statistically significant.

where i^* is the nominal federal funds target, π is the inflation rate over the last year, and y and y_T are the observed and trend rates of real output, respectively. The Taylor rule suggests, in other words, that FOMC assigns equal importance to the "output gap" (itself a proxy for the cyclical component of unemployment) and the gap between the ongoing rate of inflation and the desired rate (here, 2 percent).

Although first intended as a description of the Fed's conduct only, rather than as a prescription, because it appeared to described the Fed's conduct during an interview of exceptional macroeconomic stability—the first part of the so-called "Great Moderation"—the Taylor Rule ultimately came to be regarded as providing a rough-and-ready guide to optimal monetary policy; and it was not long, indeed, before the FOMC began referring to the rule as a component of its own deliberations. Deviations of the Fed's actual target federal funds rate from the target consistent with the Taylor Rule therefore supply another indication of the extent to which monetary policy is or has been too easy or too tight. Taylor himself has noted that the Fed departed substantially from the Taylor Rule during the recent housing boom. He has claimed, furthermore, that the boom would have been far less substantial, and the consequent bust far less (Taylor, 2009).

Because both the difference of the federal funds rate from its Taylor Rule value—henceforth "The Taylor Gap"—and the Productivity Gap are supposed to indicate the stance of monetary policy, it is worth considering the extent to which the two measures are in agreement. That the two "gaps" (both in non-smoothed form) are in fact closely correlated can be seen by looking at Figure 7. The R^2 between these two series is 80 percent.

Figure 7

The Productivity and Taylor Rule Gaps

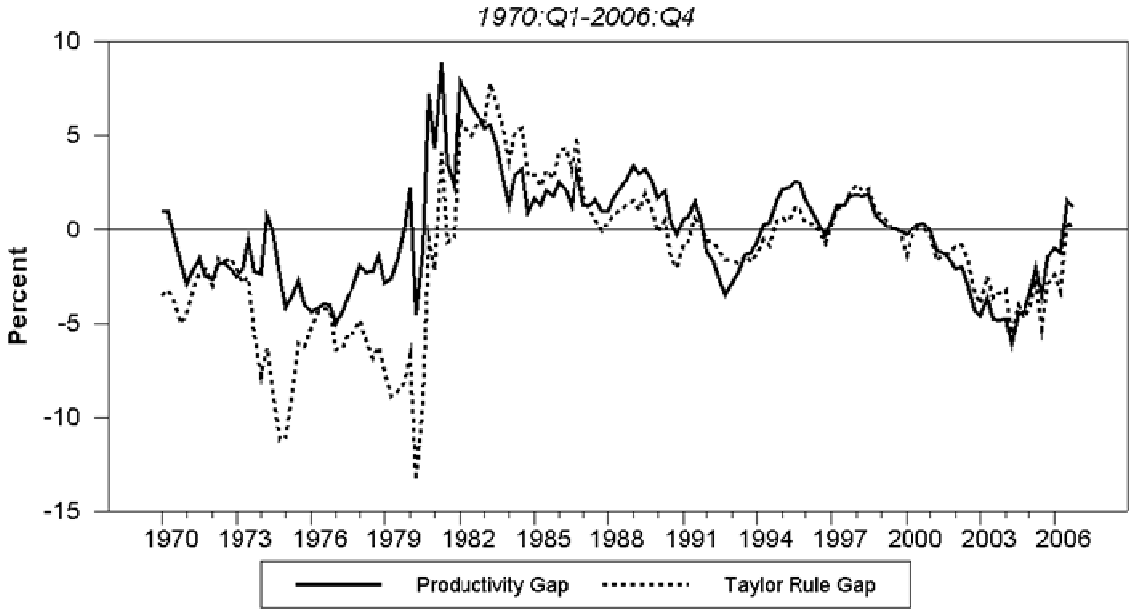

Source: Fernald (2009), Laubach and Williams (2003), FRED Database, Authors' Calculations

That the gaps move together is not surprising in light of the previously observed, negative correlation between the Productivity Gap and the output gap, the last of which forms one of the Taylor Rule's two indicators of the relative easiness of monetary policy. That the gaps tend to coincide suggests both that changes in the neutral funds rate, as well as in the output gap, are both to a very substantial degree driven by persistent changes in the economy's rate of productivity growth, and that the Taylor Rule, with its output gap component, allows reasonably well for such changes.

It follows that the Fed's departures from the Taylor Rule may be understood as being frequently due to its failure to adjust the federal funds target appropriately to changes in the economy's rate of productivity growth.

The Productivity Surge and the Subprime Boom

Both the Taylor Gap and the Productivity Gap suggest that monetary policy was excessively easy in the aftermath of the dot.com collapse, and that it was so to an extent unmatched since the

inflationary 1970s. However the Productivity Gap measure allows us to reach the further conclusion that the Fed erred specifically by failing to raise its target rate in response to the renewal of the productivity growth surge that had begun in the mid-90s, but interrupted by the dot.com crash.

A comparison of the Fed's response during the first (pre-tech bubble crash) and second (post-crash) phases of the productivity surge is revealing. During the first phase of the surge, the Fed kept the federal funds rate either at or somewhat above both its Taylor Rule level and our own rough measure of the neutral level, by raising its target rate aggressivel.[12] However, as the productivity surge continued, Fed policy became more accommodative, and especially so in the aftermath of the crash, when the Fed, instead of responding to rising productivity growth by correspondingly raising its target, responded perversely, lowering the funds rate to the point of actually driving it below zero in real terms.

This response pattern was, we believe, the result of two interacting factors. First, as Richard Anderson and Kevin Kleisen (2010) explain, the Fed only gradually came to appreciate the magnitude and enduring nature of the post-2000 revival and acceleration of productivity growth. Second, as it did so, it also came to modify its position concerning the bearing of persistent productivity growth on appropriate adjustments of the federal funds rate target. Alan Greenspan explained the nature of the modification in his January 2004 speech at the AEA meetings:

> *As a consequence of the improving trend in structural productivity growth that was apparent from 1995 forward, we at the Fed were able to be much more accommodative to the rise in economic growth than our past experiences would have deemed prudent. We were motivated, in part, by the view that the evident structural economic changes rendered*

[12] According to Timothy Fuerst (2003), the federal funds rate was also consistently above its Taylor Rule value through the latter half of the 1990s.

> *suspect, at best, the prevailing notion in the early 1990s of an elevated and reasonably stable NAIRU (non-accelerating inflation rate of unemployment). Those views were reinforced as inflation continued to fall in the context of a declining unemployment rate that by 2000 had dipped below 4 percent in the United States for the first time in three decades.*

The FOMC became increasingly inclined, in other words, to overlook the implications of accelerating productivity growth on the real neutral rate of interest, instead preferring to focus on the accelerating growth rate's implications for the rate of inflation. Because more rapid productivity growth reduced the risk of inflation, while raising that of disinflation, the FOMC reasoned that it could go further than might otherwise have appeared prudent to in encouraging recovery from the dot.com recession by refraining from rate increases. This was in contrast to the view that prevailed prior to the dot.com crash, according to which persistently higher rates of productivity growth pointed to the need for corresponding rate *increases* to avert an *inflationary* boom.

Thus in its February 2000 meeting, in response to both unexpectedly high staff productivity growth estimates and concerns that still more robust growth might be in the offing, Fed Governor Meyer opined:

> *if the acceleration in productivity leads to continued expectations of accelerating earnings per share, the only way to eliminate the wealth effects, which has to be eliminated, is for the market rate used by investors to calculate the present value of expected earnings to rise (ibid., p. 146).*

Although Meyer also worried that overly sharp rate increases might "crack the market," he nevertheless proposed a hike of 25 basis points, to which the rest of the FOMC agreed. In subsequent meetings prior to the late 2000 downturn the committee took a similar stance, further

raising the target funds rate in response to renewed staff reports of surging productivity growth, with the aim of limiting profit inflation.

Following the late 2000-2001 setback, in contrast, as productivity once again began growing at high and increasing rates, the FOMC no longer concerned itself with the possibility that a failure to raise rates would place them below their neutral levels, inflating both anticipated profits and asset prices. Instead, it concerned itself with the potential disinflationary consequences of strong productivity growth, which it sought to counter by leaving its lowered nominal target unchanged despite surging productivity.

Some members of the FOMC themselves recognized that its stance had shifted. During the committee's January 2003 meeting, Glenn Rudebusch (FOMC January 28-29 2003, p. 28) observed how, beginning in 2002, monetary policy was apparently being guided by "something other than strict Taylor Rule determinants," perhaps owing to the "collapse of the tech bubble in stock prices or to geopolitical risks" and that this something (as if it was a force other than the FOMC's own decisions) was "holding the actual funds rate below what a simple Taylor rule would say." At this time the real federal funds rate had been negative for some months, while the Federal Reserve Board's estimates of past and projected annual total factor productivity growth rates were both high and increasing, with the rate for 2002 placed at 2 ¼ percent, and those for both 2003 and 2004 placed at 2 ½ percent. Based on these estimates the FOMC also raised its estimate of the permanent component of the multifactor productivity growth rate to 1½ percent (ibid., p. 74).

Such high productivity growth rate estimates might have been expected to suggest similarly high estimates of the real neutral federal funds rate, and thus (for any given mean inflation target) to the need for correspondingly high targets for the nominal federal funds rate target. Yet instead of revising its target upwards the FOMC kept it at its low post-bust level. The committee's decision to do this was influenced by its understanding that that, because such an "accommodative"

stance would merely serve to prevent such a *decline* in inflation as would otherwise go hand-in-hand with exceptionally rapid productivity growth, instead of actually raising the rate of inflation, the stance was not excessively easy. Rapid productivity growth, in other words, was treated by the FOMC as at best offering a sort of "free lunch," whereby interest rates could be kept persistently below any reasonably reckoning of their "neutral" levels without destabilizing real economic activity, or at worst posing a "downside risk" of deflation that warranted further *easing* of monetary policy.

By the time of the FOMC's March meeting, its decision was beginning to bear fruits that might have encouraged it to reexamine its logic, including evidence of high earnings and profits generally, and of a buoyant housing market in particular. FRB Dallas President Bob McTeer (March 18, pp. 54-5) even reported hearing from one Texas authority that low mortgage rates were serving as "nicotine to the housing industry" and that "mortgage rates could rise by a percentage point or so, maybe even 2 points, from the current very low levels without having a strongly negative effect on housing demand"—an observation consistent with neutral rates having been at least that much higher than ones actually prevailing.

Yet the FOMC remained obstinately committed to avoiding a decline in the rate of "headline" (CPI) inflation at all costs, notwithstanding the fact that such a decline, had it been permitted, would have reflected declining unit production costs, and not a deceleration of the rate of growth of aggregate demand, and despite its awareness that such a policy was not in keeping with a goal of keeping interest rates at their neutral levels. Thus in the December 9th FOMC meeting Federal Reserve economist David Stockton expressed the Board's belief that the Fed might keep the funds rate below equilibrium levels until late 2006 "and still not generate any acceleration of inflation," appearing thereby to answer in the affirmative Kansas City Fed President Thomas Hoenig's query as to whether the Board intended to encourage the FOMC to boost output first

"and then worry about how far below equilibrium the short-term rate is" (FRB December 2009, pp. 21-2). Stockton elaborated:

> *[S]uppose at the end of 2005 the nominal funds rate is 2 percent and you thought that a 4 percent nominal funds rate—a 3 percent real funds rate and 1 percent inflation—was where the rate had to go. Could this Committee raise the funds rate 200 basis point in 2006? I think you certainly could...given what you've demonstrated in the past as a reasonable willingness to be aggressive (ibid., p. 22).*

Besides being aware that the funds rate was well below its 'equilibrium" level, the Fed at this date was also aware of increasing symptoms of an overly-easy policy stance, including a weakening dollar and continued housing price inflation. But all such considerations were set aside in favor of an exclusive focus on "balancing" the risks of inflation and deflation. The attitude expressed by President McTeer at the December meeting, though more cautious than that of some of his fellow committee members, was in other respects typical:

> *...I am willing to concede that rapid growth sustained by substantially negative real interest rates may prove a problem. But that problem is not imminent given the degree of productivity growth and slack in the economy .... I believe that the outlook for growth in real GDP is either balanced or biased toward more growth, and I think the outlook for inflation is now balanced.*

President McTeer apparently did not consider that exceptionally high productivity growth, far from justifying negative real interest rates, meant that those rates were likely to be even further below their neutral values than they would have been otherwise.

Many months later, Board Vice Chairman Roger Ferguson, in the course of an October 2004 talk at the University of Connecticut School of Business on the subject of the "equilibrium" real federal funds rate (his preferred term for what we are calling the "neutral" rate), observed that

the real federal funds rate had only recently moved "into positive territory for the first time in three years" (Ferguson 2004, p. 2). He then went on to observe that rapid productivity growth was "a powerful force tending to make the equilibrium real rate higher than it would otherwise be," while noting how one well-known estimate of the real neutral rate—that of Laubach and Williams (2003)—placed it at 3 percent as of mid-2002, which agrees with our own 500 basis-point estimate of the Productivity Gap for that date. Yet, rather than conclude from this that the Fed's policy stance had been excessively easy, Ferguson attempted to justify it by citing conditions supposedly warranting departures from a "neutral" monetary policy stance. By doing so Ferguson, instead of adhering to the original understanding, dating back to Wicksell, of conformity of the actual with neutral or "equilibrium" federal funds rate as a summary criterion for the correctness of monetary policy, treated it as if it were merely one of several desirable policy objectives that might profitably be traded against one another.

Assessing the Fed's Conduct

In retrospect, the FOMC's decision to forego a neutral policy stance for so long, in favor of one merely geared to avoiding inflation, proved tragically mistaken: by failing to adjust its funds rate target in a manner consistent with productivity-growth based increase in the real neutral rate of interest, the Fed contributed to a serious destabilization of asset markets, and of the housing market in particular.

Behind the Fed's failure was its tacit switch, in the wake of the tech bubble collapse, from a policy roughly in accord with the Taylor Rule to one that amounted to mere CPI inflation targeting. That switch was a mistake, because the rate of output price inflation consistent with a neutral policy stance itself varies along with growth rate of productivity, with higher than average productivity growth requiring a lower "neutral" inflation rate.

To appreciate the logic behind this conclusion is necessary is to bear in mind that, when the productivity growth rate changes, unit costs are declining. Consequently, the difference between output and input (i.e. wage rate) inflation rates must also decline. It follows that, if it is to prevent any decline in the rate of output price inflation, the stance of monetary policy must be such as will raise the rate of *input* price inflation. But since input prices are generally more sluggish to adjust or "sticky" in response to unanticipated growth in nominal spending, such an "accommodative" stance can temporarily swell both current and anticipated profits. In other words, in preventing unusually rapid productivity growth from reducing the rate of inflation, the Fed contributes to asset-price inflation no less than it does if, at a time of average productivity growth, it sets its targets at level that cause "headline" inflation to rise.[13]

The claim that a consistently neutral monetary policy must allow for a varying rate of inflation is itself not at all inconsistent with the Taylor Rule, for although in formulating his rule Taylor allowed for a 2 percent mean rate of inflation as well as for 2 percent mean real neutral funds rate, the rule also allows for the possibility that to maintain a neutral policy stance the Fed must allow the inflation rate to vary around its mean. Our "Productivity Gap" perspective suggests that the necessary variations will be opposite those to the real neutral rate itself, and hence, approximately opposite to variations in the real rate of growth in total factor productivity.[14]

---

[13]The argument is developed with formal rigor by Tambalotti (2003, pp. 30-1), who finds that "strict inflation targeting is a particularly undesirable policy to insulate the economy from the effect of trend productivity shocks," because to do so monetary must assume what is in fact an expansionary stance that "will then result in a boom in demand." In contrast, Tambalotti concludes, so long as wage rates are more "sticky" than prices of final goods, "a policy that stabilizes nominal wage inflation around its steady state growth rate...would produce better outcomes than those obtained under inflation stabilization." Tambalotti also concludes that "actual policy in the 1990s was very close to optimal according to [his] model" (ibid., p. 32).

[14] Indeed, as Donald Sutherland (2009, p. 1) has observed, "there is nothing magical about 2 percent inflation, and the [Taylor] rule can be modified for a different inflation target." Selgin (1997) argues, not just for allowing the inflation rate to vary inversely with the growth rate of productivity, but for setting the mean target rate of inflation equal to minus the mean rate of productivity growth.

Conclusion

High or accelerating productivity growth doesn't offer monetary policymakers a "free lunch." It does not allow them to hold actual rates of interest substantially below their theoretically neutral levels without contributing to the business cycle. Nor is maintaining a non-neutral rate justified by the fact that doing so is unlikely to increase the CPI or "headline" rate of inflation. Instead, to avoid triggering or exacerbating booms and busts in the face of innovations to aggregate productivity, the Fed must be prepared to allow the "headline" inflation rate to vary occasionally from its long-run target value.

The Taylor Rule itself acknowledges this truth, by calling for adjustments to the federal funds target, not merely in response to deviations of the inflation rate from its preferred level, but also in response to a positive or negative "output gap." By so doing the Taylor Rule implicitly allows for adjustments to the real neutral funds rate in response to productivity-growth based changes in the neutral rate. By drawing attention more explicitly than the Taylor Rule does to the bearing of the productivity growth rate upon the real neutral rate of interest, we hope to have supplied a better understanding both of why the Taylor Rule works reasonably well, and of the particular circumstances in which the Fed has been inclined to depart from that rule, sometimes with dire consequences.

In closing, Fed policymakers must be aware of productivity surges and the correct way to respond to them. While the tendency of such surges to reduce inflation may tempt them to set a lower than usual FFR target, theory and prudence call for them to do just the opposite, allowing the output-price inflation rate to decline, while in effect stabilizing the rate of factor price inflation.

Data Appendix

Data Sources

Total factor productivity comes from Fernald (2009) while the output gap measure comes from Laubach and Williams (2003). The rest of the variables—the federal funds rate, the CPI, the unemployment rate, and housing starts—all come from the Fred Database at the St. Louis Federal Reserve Bank.

Variable Construction

The neutral real federal funds rate is based on the equation , $r_t^n = r^n + (g_t^e - g)$, where r_t^n is the current period neutral real interest rate, r^n is the long-run, steady neutral real interest rate, g_t^e is the current expected year-on-year TFP growth rate, and g is the mean year-on-year TFP growth rate for the sample. We assume r^n=2 percent and estimate g_t^e using an exponential weighted moving average process of past year-on-year TFP growth rates with the current period weight receiving a value of 0.70. As noted in the text, the Productivity Gap, P_t , is calculated as follows: $P_t = r_t - r_t^n$, where r_t is the actual real federal funds rate. The real federal funds rate is calculated as the federal funds rate minus the year-on-year inflation rate.

References

Anderson, Richard G., and Kevin L. Kleisen. 2010. "FOMC Learning and productivity Growth (1985-2003): A Reading of the Record." Federal Reserve Bank of St. Louis *Review* (March/April): 129-53.

Bernhardsen, Tom, and Karsten Gerdrup. 2007. "The Neutral Real Interest Rate." Norges Bank *Economic Bulletin* 78 (2): 52-64.

Christiano, Motto, Rostagno

Federal Reserve System Board of Governors. FOMC minutes.

Ferguson, Roger. 2004. "Equilibrium Real Interest Rate: Theory and Application." Remarks delivered at the University of Connecticut School of Business, October 29, http://www.federalreserve.gov/boarddocs/speeches/2004/20041029/default.htm Chairman Ferguson, October 29.

Fernald, John. 2009. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Unpublished manuscript, Federal Reserve Bank of San Francisco, August 16.

Frederick, Shane; Loewenstein, George, and O'Donoghue Ted. 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40(2), 351-401.

Gordon, Robert. 2010. "Revisiting U.S. Productivity Growth Over the Past Century with a View of the Future." *NBER Working Paper* 15834.

Kahn (on direct effect of productivity on housing demand)

Laubach, T. and Williams J. 2003. "Measuring the Natural Rate of Interest" *Review of Economics and Statistics*, 85, 1063-1070.

Lombardi, Marco and Sgherri, Silvia. 2007. "(Un)Naturally Low? Sequential Monte Carlo Tracking of the US Natural Interest Rate." *ECB Working Paper* 794.

Oliner, Stephen; Sichel, Daniel, and Stiroh, Kevin. 2007. "Explaining a Productive Decade."

Brookings Papers on Economic Activity, 1:2007, 81-137.

Schmitt-Grohe and Uribe

Selgin, George. 1997. *Less Than Zero*

Sutherland, Donald. 2009. "Monetary policy and real estate bubbles." Working paper, Institute for SocioEconomic Studies, February 11.

Tambalotti, Andrea. 2003. "Optimal Monetary Policy and Productivity Growth." Unpublished ms, Princeton University, February 11.

Taylor, John. 2009. *Getting off Track*, Hoover Instituion Press, Stanford, California

Weidener, Justin. and Williams, John. 2009. "How Big is the Output Gap? "San Francisco Federal Reserve, *Economic Letters*, 2009-19, 1- 3.

Williams, John C. 2003. "The Natural Rate of Interest." Federal Reserve Bank of San Francisco *Economic Letter* 32 (October 31).

Whalen, Karl. 2009. "Technology Shocks and Hours Worked: Checking for Robust Conclusions." *Journal of Macroeconomics*, 31, 231-239.