




Driving Recessions

Posted by [Stuart Staniford](#) on October 23, 2005 - 10:16am

Topic: [Economics/Finance](#)

Tags: [gas prices](#), [hubbert peak](#), [oil prices](#), [peak oil](#) [[list all tags](#)]


So I [managed to find](#) the annual data for US total vehicles mile driven for the 1970s (Table 2.1.1 in the link). It turns out we were missing some of the most interesting parts of the story... I've got incrementally more evidence that there is a causal connection between miles and GDP growth, though not an open-and-shut case, and I still can't give a plausible account of why it works so precisely.



This post is part of a never-ending series on the economic response to oil shocks, which I'm doing to gain a deeper insight into likely post-oil-peak economic occurrences. It began with a [discussion of the stability of the mix of consumer spending](#), moved on to [consider the productivity and efficiency of transportation](#), dived more deeply into what the [US economy did to become less oil intensive](#), then looked at [behavioral driving responses to oil shocks](#), and most recently [checked out some European data](#). Readers with long memories will also realize the significance of these issues to my [decline-rate based model](#) of what might happen to the US economy post peak.

It's worth noting that this line of posts has now sparked interesting discussion over at [Dead Parrot Society](#), and [Econbrowser](#).


Recall that I've been obsessing over this picture, now with the 1970s annual data and the 2003 number ([from here](#)) included.



The GDP is the real (inflation adjusted) numbers [from the Bureau of Economic Analysis](#), and the miles come from the [Transportation Energy Data Book](#), the Federal Highway Administration [Summary of National and Regional Travel Trends](#), and [Highway Statistics 2003](#). Both series have then been rescaled to make 1970 be 100%.


I have to say in passing that I thoroughly approve of the portion of my federal and state taxes that have been expended on assembling and publishing these statistics. Money well spent in my book.

So then the big question is, is the very striking agreement between these series a fluke, or is there a deep reason why it must be this way? It's definitely not the case that all things that just generally grow with the economy agree this closely with GDP. This next graph adds the time series for total US population (in green), US "economically active" population, in yellow, and number of housing units (from the Census) in purple. Clearly, these have all been generally growing, but not at the exact same rate as the GDP (though there's a very good correspondence between housing units and employed population - I guess it's hard to pay the rent if nobody in the house has a job).



We saw yesterday that there's [evidence from Europe](#) of a similar pattern over there, though the coupling is definitely somewhat looser.

Well, let us look at the year on year changes in GDP and miles (ie what was the percentage difference in the respective quantity from one year to the next) now that we can see into the Big Kahuna oil shocks in the 70s:



Ah, beautiful. We've clearly got a fair amount of correlation in the annual growth rates. For pretty much every strong, oil-shock driven, recession, the driving line takes a dive right around or slightly before the GDP line, and then recovers at the same time or slightly earlier too. Note also that now we can see the 70s, we can see all four shocks during this period:

- 1973 Arab embargo (7.8% of oil supply)
- 1978 Iranian revolution (8.9% of supply)
- 1980 Iraq-Iran war (7.2% of supply)
- 1990 Iraq invasion, and US freeing, of Kuwait (8.8% of supply)

The 78 and 80 shocks get somewhat merged. This, and the worst shock in 1973, did cause enough of a problem to actually drive miles down, albeit by only 2 1/2% in the latter case. So I guess once things get bad enough that people are shooting each other in the gas lines and the President is wearing a sweater on TV, they're bad enough to actually reduce driving a tiny bit.

However, it's also clear that not **all** features of the two curves match. There's a 1982 nasty recession that the driving data totally ignore. The driving data doesn't really follow the late nineties tech bubble and subsequent crash, which is clearly visible in the GDP growth line. There's a dual peak structure around 1985 in both lines, but they disagree about which should be the bigger peak.

So how similar are they really? There's a technical way to assess this known as R^2 , the correlation co-efficient. This measures the percentage of the fluctuations (technically variance) in one line that can be explained by a linear rescaling of the other. So if the two lines have exactly the same shape (even if size and vertical position were different), they would have an R^2 of 100%. If the two lines were completely unrelated, they would have R^2 of 0%. So these two lines have an R^2 of 35%. So, in addition to the long-term trend being the same, 35% of the fluctuations in one line about it's long term trend are explained by the fluctuations in the other line. Not too bad for a social science R^2 ...


Before anyone gets upset that autocorrelations mess this up, the 1 year lagged autocorrelation R^2 in GDP growth is only 3.8%, and the 1 year lagged autocorrelation in miles growth is only 9.6%. (The lagged autocorrelation studies how much of the fluctuations in some year were explained by where the fluctuation was up to the previous year - so they measure the degree to which the line wants to wander around smoothly rather than jerkily. These low autocorrelation R^2 values tell us these lines, in so far as they wander, wander fairly jerkily - this year's growth has very limited memory of what last year's growth was). So that maybe gives us a better sense of what this 35% correlation between the two lines means - if you want to understand this year's GDP growth, this year's mileage growth has about 10 times as much explanatory power as last year's GDP growth. And if you want to understand this year's mileage growth, this year's GDP growth has about 3 1/2 times as much explanatory power as last year's mileage growth.

I think there's a little bit of a lag here - things seem to happen a little bit earlier in the mileage line. However, it's only a few months, not a whole year, so it's a bit tricky to compute a shifted R^2 , even though it might be a shade lower. It's also noteworthy that there are no big features in the mileage line that are missing from the GDP line, but there are big features in the GDP line missing from the mileage line (such as the 1982 dip and the late nineties boom). This suggests that causation is more prone to run miles-to-GDP than the other way round (though I'm sure there's some feedback arc both ways).


Now, in addition to the local correlations, we also have the fact that the long term trends are the same. In fact, over the whole period 1970-2003, mean GDP growth is 3.1%, and the mean miles growth is 3.0%. Obviously the fact that these growth rates are so similar is the core reason the graphs line up. We can say a little more about the significance of the similarity. In particular, if we look at the population of 32 growth observations of each type, we find that the standard deviation in both is the same value: 2.0%, which (given there's not much autocorrelation so we might not be hopelessly off-base to treat them as iid and divide the standard deviation by $\sqrt{32}$), means

that the expected standard error in each rate is 0.35%. Ie, if you just looked at the pretty large growth rate fluctuations, you'd think the mean of each of them might have come out different than it did with a fluctuation size either way of 0.35%. Thus the error in the difference between their two rates is $\sqrt{2}$ larger, or 0.5%. So the two trends are quite close together compared to the fluctuations in their growth rates. But, if that difference were normally distributed, the chance of being at 0.1% different or better (ie within 0.2 standard deviations either way) would be 16%. So we can't say "there is statistically significant evidence that these lines are closer together than you'd expect just based on their average and fluctuations in growth" - we'd need a much longer trend.

I meant to repeat this with a bootstrap Monte Carlo for good measure (since the normality assumption doesn't seem too solid), but I think I need to do a bit more work. The residuals of GDP growth minus mileage growth look like this:



The autocorrelation r^2 at lag 1 in this series is only 0.6% (ie negligible), so I built a little Monte Carlo program that takes the GDP series and builds alternative mileage series by picking a random residual, bootstrap style, from the pool of residuals in the graph and multiplying last years mileage by the gdp growth plus random residual. The idea is to then build a population of such lines and see if the closeness of our actual ones is anomalous or not. Some output (two MC runs) of this is shown below. However, this process doesn't work because it produces mileage lines with too high a variance (GDP variance plus residual variance). So I need a better way to generate these random walks.



But, that will have to await a new surge of inspiration - tonight's has run out!



This work is licensed under a [Creative Commons Attribution-Share Alike 3.0 United States License](http://creativecommons.org/licenses/by-sa/3.0/us/).