

Faust, Jon

Working Paper

Comments on Piazzesi and Schneider's "Bond positions, expectations, and the yield curve"

Working Paper, No. 2008-4

Provided in Cooperation with:

Federal Reserve Bank of Atlanta

Suggested Citation: Faust, Jon (2008) : Comments on Piazzesi and Schneider's "Bond positions, expectations, and the yield curve", Working Paper, No. 2008-4, Federal Reserve Bank of Atlanta, Atlanta, GA

This Version is available at:

<https://hdl.handle.net/10419/70740>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Comments on Piazzesi and Schneider's
"Bond Positions, Expectations, and the Yield Curve"

Jon Faust

Working Paper 2008-4
February 2008

Comments on Piazzesi and Schneider's "Bond Positions, Expectations, and the Yield Curve"

Jon Faust

Working Paper 2008-4
January 2008

Abstract: This working paper comments on Monika Piazzesi and Martin Schneider's "Bond Positions, Expectations, and the Yield Curve," delivered at the Fiscal Policy and Monetary/Fiscal Policy Interactions conference held at the Atlanta Fed on April 19–20, 2007.

Key words: expectations, surveys, interest rates, portfolio choice, asset positions, term structure, yield curve

The views expressed here are the author's and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the author's responsibility.

Please address questions regarding content to Jon Faust, Director, Center for Financial Economics and Economics Department, Johns Hopkins University, 3400 N. Charles Street, Baltimore, MD 21218, 410-516-7614, faustj@jhu.edu.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's Web site at www.frbatlanta.org. Click "Publications" and then "Working Papers." Use the WebScriber Service (at www.frbatlanta.org) to receive e-mail notifications about new papers.

The allocation of research resources in the profession between fiscal and monetary policy has long been quite skewed toward monetary policy. This conference on monetary and fiscal interactions reminds us once again that the overemphasis of monetary policy may not make sense even from the narrow standpoint of understanding monetary policy. It is very fitting that the Federal Reserve Bank of Atlanta is promoting a more integrated view of policy analysis.

Another area would also benefit from a more integrated approach: macroeconomics and finance. There is an immense macroeconomic literature on how monetary policy is transmitted through the yield curve and a similarly immense finance literature on the yield-curve. Only recently have we begun to integrate these literatures. A proper integrated treatment of monetary and fiscal interactions cries out for an integrated view of finance and macro. Understanding policy interactions requires resolving important questions about yield curve dynamics, their relation to policy expectations, and to the magnitude and composition of outstanding debt.

Given this background, I was very pleased to be asked to discuss a paper entitled “Bond Supply, Expectations, and the Yield Curve,” authored by Monika Piazzesi and Martin Schneider, two of the pioneers in the fledgling enterprise of integrating modern finance and macro. This is an excellent piece of research and I learned a great deal in reading it.

In these comments, I’ll try to provide some perspective on key choices faced in producing this ambitious work. In doing so, I’ll highlight some ways in which I like the paper more, perhaps, than the authors do. I’ll also discuss some directions for further progress.

Dynamic, Stochastic Models in Finance and Macro

One central project in finance is to answer the question: Can we rationalize observed data on asset prices and portfolio shares? It turns out that this question is too challenging, and so we often simplify by ignoring quantities, except to assume that the representative, or marginal, agent holds a bit of all assets. That is, we assume the portfolio choice problem is at an interior solution. This allows us to examine many questions regarding asset prices by studying Euler equations arising in optimization problems, and we can get a very long way without understanding, or even specifying, most of the model. The ultimate goal remains to understand prices *and* quantities, but this generally requires a more complete model.

Until recently, there were very few results on the difficult problem of solving a dynamic, stochastic, general equilibrium (DSGE) model of portfolio choice in the realistic setting in which agents face a rich array of assets. The last 10 years, however, have seen an explosion of finance work in dynamic, stochastic models. We see brute-force numerical solution of models, complete solutions of special cases, and clever approximations around the special solved cases. Monika and Martin's paper falls in this latter category, building on the important work of Campbell, Chan, and Viceira (2003), who derive an approximate solution to a simple dynamic, stochastic portfolio problem.

Of course, since the work of Kydland and Prescott, macroeconomics has also made great strides in analyzing DSGE models. The macro models present agents with a rich set of choices about real stuff--labor supply, consumption, physical investment, etc--while including a trivial set of intertemporal claims (assets). For example, there may be

only direct ownership of capital and one-period bonds. Alternatively, there may be complete contingent claims markets, and this trivializes the problem in another way. In contrast, in the finance DSGE work, agents face a rich portfolio problem, but the macroeconomic side of the model is trivial: the real outcomes on which agents write claims might, for example, be simply a random endowment stream.

Of course, we would like to see DSGE models with both rich macroeconomics and rich portfolio choice problems. Ultimate resolution of many problems including those regarding monetary and fiscal interactions will almost certainly require a model that is rich in both dimensions.

One main roadblock we face on the pathway to integration is that it is very hard to reconcile asset price dynamics with economic fundamentals. Thus, an important aspect of current integration work involves documenting puzzles and anomalies in asset price behavior from the standpoint of existing models. Where do we stand on the project of resolving these anomalies? I believe we are not at the end of this process, or even at the beginning of the end. Perhaps work such as the current paper brings us to the end of the beginning.

Given the current state of understanding, anyone building an integrated model faces a challenging problem of how in the face of all the anomalies to use data to improve our current models. In both macro and finance, there are four important challenges we face in bringing DSGE models to data:

1. *Pick a subset of all possible implications of the model to take seriously.*

We cannot solve all problems at once, so every current DSGE will have *problem areas*.

We ignore these in order to focus on certain *focal* implications. For example, we might emphasize the consumption or investment Euler equations.

2. *Pick a small subset of all potentially relevant data to use in drawing inferences about the model.* Our models simply cannot account for the joint dynamics of a very large set of variables. Thus, in macro, we may focus on aggregate GDP, consumption, and inflation, along with, say, a single rate of return. In finance, we may leave out most, if not all, macro variables and focus on a larger (but still smallish) set of returns.

3. *Pick a treatment of expectations.* Here there are many choices, but the problems begin to multiply. We may use “rational” expectations, expectations that are derived to be consistent with the full set of implications of the model. Of course, in steps 1 and 2 we admitted that we are not really comfortable with the full set of implications of the model. Thus, the case for rational expectations, which to many economists might otherwise be compelling, is substantially weakened. An alternative is to do inference on the model using instrumental variables (IV), which leaves the treatment of expectations largely implicit. This allows the researcher to sweep under the rug some troubling questions about expectations, but in my view this is not a virtue of the approach.

A little-used approach is to estimate expectations using an *ad hoc* time series model (*ad hoc* in the sense of not imposing any model-based restrictions). Of course, this approach is a close cousin of the IV approach; in both approaches expectations are inferred from reduced-form time series correlations in the data. The principle difference

between the IV and *ad hoc* time series approaches, in my view, is that the *ad hoc* time series approach is more explicit and open to discussion and criticism. I see this as a virtue.

4. *Choose a metric for success of the model.* Here, of course, we have formal methods, Bayesian (e.g., posterior odds ratios) and classical (e.g., J-tests in the GMM case). Given the fact, however, that we know our current models have gross deficiencies, formal inference is sometimes argued to be uninteresting: it is fairly straightforward to make current models look about as bad as we like using formal inference. In light of this, we often find ourselves using informal methods, which amount to discussing whether or not the model “looks pretty good” in certain dimensions. This is surely an intermediate step on the path of research, but has proven a useful technique.

The Four Choices in the Current Paper

The basic plan in the paper starts with measuring at each point in time the portfolio allocation of U.S. households. For this purpose, assets are gathered into broad categories to form a small number of synthetic assets and measures are created of the evolution of portfolio shares through time. Next, the authors build a time series model of returns on those synthetic assets. We take the predicted values from this model as a proxy for the expectations of households. Given asset holdings, price expectations, and a model of optimal behavior, we then invert the decision process of the agents to infer what asset prices would be implied. These model-deduced prices are compared to actual prices to determine if the model is “pretty good.” Let me make some comments organized around the four choices discussed above.

Choice 1. This paper begins with a Campbell, Chan, Viceira-style (CCV) solution of a dynamic asset allocation problem. Using a clever approximation around a solved case, CCV derive fairly simple expressions for the joint behavior of three items: asset prices, portfolio allocations, and expectations. Given any two of these items, we can make inferences about the third. CCV derive implications for asset allocations based on a price data and a model for expectations. The current paper also takes seriously CCV-style implications for the joint behavior of asset prices, portfolio allocations, and expectations.

As a small aside, CCV and the current paper perform two of the three natural permutations with the framework. The final exercise is to take price and quantity data and to derive implications for expectations: given the price and quantity data, what must people have been expecting? (Or as my Mother used to say when faced with *my* choices, “What in God’s name were you thinking!?”) In many ways, this third exercise is the most natural. We have excellent data on prices, decent data on quantities, but essentially no data on expectations. It probably makes sense to use the better-measured items to learn about the less well measured.

Choice 2. Regarding data, the current paper is nearly unique in focusing on portfolio allocations. The authors construct a good deal of the portfolio share data required for the paper. Constructing these data required a great deal of work and I think this alone is a major contribution of the research project. Consistent with most work coming from the finance side, macroeconomics is treated in a fairly trivial way; for example, the only macro variable involved in the forecasting model is inflation.

Choice 3. When it comes to expectations, the authors sound an apologetic note in taking the *ad hoc* time series approach. I do not think such apologies are needed. As

noted above, the rational expectations approach is questionable in the current context and the IV approach has little to recommend it over the *ad hoc* time series approach.

Choice 4. The authors take the informal approach to evaluating the merits of the model. For example, the paper presents plots of actual asset prices and prices inferred using the model; the authors then discuss whether the model-based prices “look pretty good.” We would ultimately prefer more formal criteria, but I think this paper is another excellent demonstration of the fact that these informal comparisons may be quite informative, especially at this early stage in the work synthesizing macro and finance.

On the whole then, I think the authors have dealt with the four choices in ways that are astute at best, and natural at worst. Let me raise two other issues.

First, I would like to see better motivation for the auxiliary time series model used to generate expectations. The main point of this exercise may be boiled down to asking whether observed dynamics of portfolio shares are associated in the proper way with agent’s predictions regarding returns. From this standpoint, it is essential that the time series model reflect the predictions agents may have made. Is a low-dimensional VAR likely to do this? We know that VARs generally produce very poor predictions, and they are essentially never used in practical forecasting. At minimum, I would like to see documentation that the predictions from the model are somewhere near the frontier of best practice in forecasting excess returns.

Further, the authors impose the restrictions of a particular arbitrage-free term structure model on the time series model. It is becoming common to do this, but I have yet to see an argument as to why it is coherent to do so. Informally, I put the issue this way. If the restrictions hold with probability one in the data, then there is zero (precisely

zero) benefit from imposing them. If they do not hold in the data, then we need a theory of why and how they do not hold before we know whether it is coherent to impose the restrictions. Greg Duffee (2007) has made this point much more thoroughly in a recent paper.

Thus, while I have no difficulty with using an auxiliary time series model for expectations, I would like to see a much stronger case made that the model used here is appropriate.

Second, as the authors are aware, the treatment of the residual asset is currently very unsatisfactory. The assets in the economy essentially divide into a set of bonds and the residual asset, which as the name implies is everything else: equities, real estate, human capital, etc. From the standpoint of integrating finance and macro, I think it is fair to say that the vast majority of macro is residing in the residual asset. In the current version of the paper, the return on this asset is taken to be independent and identically distributed and independent of the return on all other assets.

From the standpoint of integrating realistic macro and finance models, this model is squarely in the tradition of treating fundamentals as an exogenous process with no interesting dynamics. Further, there is no sensible view of economics or finance in which the returns on a firm's bonds and equities are related in this way. Indeed, this surely violates accounting identities.

The authors call this assumption about the return on the residual asset a *benchmark*. Because it cannot be made consistent with any reasonable view of economics or finance, I think that this is a highly problematic benchmark and the results

under this benchmark are of little interest.

Bottom line

Work integrating finance and macro is at an early stage and the initial forays into it will have readily apparent areas for future improvement. I see these not as problems, but as opportunities. This is an outstanding paper pointing out a direction in which we should be moving. From the standpoint of the issue in the title of this conference, I think this paper shows we have a long way to go. Understanding monetary and fiscal interactions will require a model with a rich finance side, perhaps like that in this paper, but also a much richer treatment of the economic fundamentals.

References

Campbell, J., Y. Chan, and Viceira, L, 2003. "A multivariate model of strategic asset allocation." *Journal of Financial Economics*, 67, 41-80.

Duffee, Gregory, 2007. "Forecasting with the term structure: the role of no-arbitrage." manuscript, Berkeley.