Comments on the thesis "Essays on macroeconomic models with heterogeneous agents" by Ivan Sútoris

Michael Reiter, Institute for Advanced Studies, Vienna

June 29, 2018

This is a very good thesis with three substantive papers. In general, I find the technical level of the analysis quite impressive. The first chapter offers a new and sophisticated solution method for HA models, or at least a modificiation of existing methods. The second chapter deals with an asset pricing problems in HA models and applies a variety of technical tools, providing both analytical and numerical insights. In the third paper, a difficult HA model is solved numerically. All three papers are competently executed. The combination of analytical and computational methods is fruitful, and shows a wide skill set of the candidate.

With the emphasis on theory and methods, it is understandable that the empirical side is not very extensive. The results of Chapters 2 and 3 are of a relatively abstract and theoretical nature, illustrating certain mechanisms, which are not easy to identify in the data. These mechanisms are shown to be quantitatively relevant, but the effects are not strong enough to be a "game changer", that would, for example, solve the empirical asset pricing puzzles. To further improve the papers, the most promising route is therefore to make a stronger connection to the data and look at more specific empirical applications.

1 Chapter 1

Chapter 1 presents a new way of computing a second-order perturbation approximation to the equilibrium of a HA model. The use of second-order perturbations is well motivated. So far they have computed using a low-dimensional parameterization of the cross-sectional distribution (the Winberry toolkit), which assumes a very smooth distribution. The author uses a non-parametric approach with a finite number of point masses. The transition law for the distribution uses a natural approximation, taken from Young (2010). The advantage of this approach is flexibility, the disadvantage is a lack of precision when the number of mass points is small. In the concrete application, the number of mass points is 30, which is a rather coarse approximation. It is computationally expensive to use more points, because the number of coefficients of each approximated function grows quadratically in the number of states. So it is an interesting experiment to see whether one can get a reasonably accurate solution with this approach.

One technical remark: it is crucial to obtain precise first and second derivatives of the model equation. It is stated in the paper that they have been derivied by numerical differentiation, which I assume means some finite difference scheme. It would be better to use automatic differentiation, but there is probably no suitable tool around for doing this in Matlab.

I find the results in Table 1.1 encouraging. They show the level effect of aggregate volatility on saving, measured by the mean of capital in long simulations. Linearity in aggregate variables ingores by construction the effect of uncertainty on the saving function (precautionary saving), while a second-order approximation can capture it. The effect of uncertainty on aggregate capital is therefore much stronger in the quadratic than in the linear approximation.

When checking for accuracy, the author is right in focussing on the differences between the path of the cross-sectional distribution obtained from the model (with the coarse grid) and from a panel of individuals, since this is likely to be the most problematic part. The results are shown in Figure 1.2. I understand that results look better with high volatility, because the variation generated by the model dominates approximation error from the discrete approximation. What I do not understand is why first order approximation is so much worse than second order, because the error we measure here should come mostly from the approximation of the grid, not from the approximation of the decision functions. This should be clarified. The results in Table 1.2 conform to expectations, except for the extreme p-values of the distribution test, which should be checked.

2 Chapter 2

Chapter 2 tackles a rather hard problem, namely asset pricing and asset choice in a production economy with Epstein-Zin preferences. Again, the author first derives a number of analytical insights from a simple model, before going to a full-fledged model analyzied numerically.

This chapter derives a rich theoretical apparatus, but the economic motivation (in contrast to the technical motivation) is a bit thin at the moment. The introduction speaks about the short-comings of existing approaches, but there should be a more focussed discussion of what are the important empirical regularities that the model should address. This should also give a clearer yardstick on whether the model is successful or not.

Section 2.2 presents a lot of derivations, but the discussion of the result is not clear and focussed enough, although because he says the model is "too simplified". One should nevertheless say clearly what the take-away is from this model.

Table 2.2 shows that the heterogeneous agent model features a significantly higher price of risk than the RA model, measured by the Sharpe ratio. Should we interpret this as a success? The price of risk is higher by about one third, but the model still misses the data by a factor of about 3, so one might consider the result interesting but a bit disappointing. A bit more interpretation of the results would be welcome here.

3 Chapter 3

Chapter 3 deals again with an interesting economic example. It is well structured between presenting some new empirical evidence, illustrating the economic mechanism with a simple model, and then presenting numerical results on a full dynamic model.

The empirical results of Section 3.2.2 are in line with the theoretical predictions. It is well known by know that uncertainty measures are robustly counter-cyclical. Since these measures partly reflect an endogenous response of the economy, not just an exogenous shock, it is not obvious what the parameter estimates mean. Do they reflect a causal effect of uncertainty shocks? Do the uncertainy measures just stand in for some unobserved characteristics of the economy? A discussion of these issues would be welcome.

This model is not easy to solve. The individual state space has one exogenous variable (firm productivity), which can be modeled as discrete, and two endogenous continuous variables, capital and bonds. The choice between capital and bonds is a subtle one (using bonds to buy capital this period or next period), so that the firm value function has little curvature along the line of equal net wealth. This makes it hard to solve, also because one has to make sure that the value function approximation does not introduce artificial non-concavity. It is therefore important to describe the algorithm in detail, and to provide accuracy checks, both for the steady state solution of the firm problem and for the general equilibrium dynamics, similar to the analysis in Chapter 1. From the remarks in Section 3.3 I get an idea of what the solution method is, and the appendix adds some of it, but a more detailed and organized exposition is necessary.

As concrete steps, I propose the following.

1. The model features investment irreversibility, but no financing frictions, fixed costs etc. In the paper, there seems to be some confusion about the role of convexity, For example in Section 3.1, there is the sentence "In the presence of nonconvex adjustment costs (such as irreversibilities or fixed costs)...". Fixed costs make the problem non-convex, but irreversibilities per se do not: the feasible set becomes smaller but is still convex. It seems to me that the firm problem of this chapter is a convex problem, so that the first order conditions are still sufficient. This should be clarified and made explicit in the paper. The difference between convex and non-convex optimization problems is crucial. Non-convex problems need thorough optimization algorithms that make sure that the global maximum is found.

2. Plot firm policy functions (investment as a function of capital and bonds, keeping productivity fixed) in steady state, to make sure they conform to theoretical predictions.

   It would also be nice to see some cross-sectional distributions in steady state (or long-run averages in simulations). Are they smooth?

3. Show statistics or plots of firm accuracy measures (Bellman equation residuals) in steady state.

4. Given that the method uses a low-dimensional aggregate state space, show accuracy measures for aggregate fluctuations, for example the R2's of price forecasts based on the aggregate states, for one period and several periods. A thorough accuracy check, similar to what is provided in Chapter 1, is needed here.

5. For the calibration, it would be good to have a systematic comparison between targets (moments in the data) and moments generated by the model in steady state.

There is one modeling decision in this chapter that I find overly restrictive. All non-entrpreneurs are hand-to-mouth consumers. This eliminates one state variable that is important in many other models, namely the aggregate financial net wealth of the entrepreneurs. In a future revision of the paper, I would reconsider this modeling decision.

This is a very rich model, whose potential is not yet fully exploited in the paper. The model replicates qualitatively the empirical finding on how the uncertainty effect interacts with financial development, but no effort is made for a quantitative exercise. Are there other microeconomic data that the model can adress? Can it replicate investment dynamics at the firm level? The analysis in the paper and in the conclusions highlight some general mechanisms, but too little is done to argue for the empirical importance of these mechanisms.

4 Summary

This thesis satisfies both in form and in content the requirements for a PhD thesis in economics.

I therefore recommend this thesis for the defense.