

Comments on

“Essays in Econometrics of Matching Markets: Identification, Estimation and Practice”

By Marin Drlje

I enjoyed reading this thesis and congratulate the author. The comments below are given in the spirit of improving the chapters for potential publication. See the last section for my evaluation.

Essay 1: Identification of school admission effects using propensity scores based on a matching market structure

1. This paper studies the role of assignment propensities in centralized matching markets: how to compute assignment propensities, and how assignment propensities can and should be used for the estimation of treatment effects. There are two major conclusions. The first is that standard RD bandwidths include many observations with low assignment risk, which the paper argues is undesirable. The second is to suggest an alternate approach to treatment effect estimation in matching market settings: conditioning on assignment propensity, as in Rosenbaum and Rubin (1983). This approach exploits additional variation in assignment beyond what is used in a standard RD, as at least some applications that are far from any particular cutoff may nevertheless be subject to assignment risk.

2. The central goal of this paper—to understand the nature of random variation of assignments generated by centralized assignment mechanisms, and to use that variation to estimate treatment effects—is well motivated and increasingly relevant for applied work as data from centralized assignment mechanisms becomes more widely available. The paper is well written, carefully thought out, and uses excellent and (as far as I know) novel microdata. I have NN comments/suggestions.

3. My first comment is that I find the second propensity-score based argument—that we can use propensity scores to identify a larger set of individuals exposed to assignment risk than are included in traditional RD estimation, and that this allows for a broader set of cross-student comparisons and treatment effect estimates—much more compelling than the first. Specifically, though the paper stated several times that using individuals with negligible assignment risk for RD estimation was undesirable, it didn't present a convincing argument for why this was the case. A central tension in RD estimation in general is that while as-good-as-random assignment holds only in the limit, we need to use observations away from the cutoff to estimate the conditional means on either side. Even if these away-from-the-cutoff observations are not subject to assignment risk themselves, they may still be useful in RD estimation if they can help us predict these conditional means. Local linear or local polynomial regression approaches adjust for the fact that far-from-the-cutoff observations may differ from those that are arbitrarily close. The point that even restricting to narrow bandwidths does not fully balance assignment propensities across the cutoff, and that local linear regressions can address the remaining imbalance, is one made in Abdulkadiroglu et al. (2019, section 2.1, link:

<https://cowles.yale.edu/sites/default/files/files/pub/d21/d2170.pdf>). I discuss this paper in more detail immediately below. My recommendation would be to focus the first part of the paper on documenting the distribution of assignment risk and de-emphasizing the critique of existing bandwidth choices. An alternative approach would make clearer why exactly it is bad to include observations with near-zero assignment risk even if those applications are informative for estimation of conditional means.

4. My favorite part of the paper was the discussion of how to use propensity score matching to expand the analysis of random variation generated by the assignment mechanism. My main suggestion for this part is to clarify your contribution relative to two existing papers. The first of these is the roughly contemporaneous paper “Breaking Ties: Regression Discontinuity Meets Market Design” by Abulkadiroglu et al (henceforth AANP; see link above). This paper seems roughly contemporaneous to yours, so I don’t think you need to defer to it as you might to an already-published work, but it’s nevertheless important to clarify for readers what the similarities and differences are. Like this paper, AANP computes assignment propensities for schools in a centralized matching market. Their approach includes both schools that use random tiebreakers (i.e. lotteries) and schools that use non-random score-based tiebreakers (as here). Two important differences between your approach and AANP jump out to me. The first is that while AANP holds types (i.e., preferences and priorities) fixed and considers only assignment risk that emerges from the randomness of tiebreakers, your approach resamples over types as well and thus includes additional uncertainty. Your approach seems like an intuitive one to me—one way to think about it as an assumption that applicants know the type distribution from which applications are drawn and my respond to that, but cannot respond to the specific draw from that distribution—and it should generate additional uncertainty in assignment beyond the AANP approach, though how quantitatively important this additional uncertainty is is unclear to me. Second, while AANP focus on a constant effects value added model and uses the additional propensity score variation to help identify parameters in that model, this paper focuses mostly on how local estimates of particular cutoffs may change. I’d recommend exploring these differences in more detail.

The second paper I’d encourage you to engage with is Agarwal and Somaini (ECMA 2018). This paper is not about regression discontinuity estimation, but it does develop a resampling procedure similar to yours. AS use this approach to estimate rational expectations beliefs about assignment chances. These correspond to assignment probabilities, including assignment risk from changes in distribution of applicants. I’d recommend describing how your approach relates to this one, and perhaps just adopting this one.

5. My final suggestion is that the paper would benefit from an expanded empirical application. I liked the description of assignment propensities above and below the cutoff in figure 1.1. There are two additional exercises I think would be especially valuable. The first is to illustrate how controlling for assignment propensities in the expanded sample of individuals subject to random assignment produces balance on predetermined treatment characteristics. The second is to show how including these additional comparisons alters estimates of treatment effects for pretty much any outcome. I understand that data on some “signature” outcomes in college application settings, such as earnings, may not be readily available. But even some intermediate outcome, such as graduation, academic achievement, or persistence in college would be invaluable as an example of the method. Interesting angles here include a) extending estimates of treatment effects away from the threshold—how do they differ? and b) how much does the additional variation you get from resampling over applications alter treatment effect estimates relative to what would be obtained under AANP type approaches?

Essay 2: LATE estimators under costly non-compliance in student-college matching markets

1. This paper studies violations of the exclusion restriction in IV research designs that use assignment to particular college in a centralized matching system to instrument for graduation or enrollment. The focus is on “costly non-compliance,” in which individuals who want to violate an initial assignment may

be able to do so, but at the cost of a reduction in the value of the outcome of interest. The authors focus on a leading example: the time cost of re-application. “Always-taker” students who violate their initial assignment by attending (or graduating from) a preferred school even if initially rejected can only do so by waiting a year and re-applying. This will tend to reduce earnings measured in a given calendar year, because always-takers have less experience. The authors propose a nonparametric bounding approach that identifies the shares of always- and never-takers, and then selectively trims the top and bottom of the earnings distribution to produce lower and upper bounds on the LATE, in the spirit of Lee (2009).

2. I enjoyed reading this paper. My take is that at present it functions mostly as a critique of Kirkeboen, Leuven, and Mogstad (2016). To the extent the authors argue that the approach is more general, it focuses mostly on other matching market applications. Are there other examples besides the college matching market coupled with the wait-a-year-to-reapply penalty? One obvious but minor extension is that similar exclusion violations may arise in college applications without matching markets; e.g. if I really want to go to a US state university but am denied admission, I can wait a year and reapply there as well, even though there is no centralized system (see e.g. Zimmerman 2014; Goodman, Hurwitz, and Smith 2017, 2020; Kozakowski 2020). It would be nice to include a broader variety of examples like this.

3. I found key passages in the text describing the intuition behind the bounding approach very hard to understand. This is on page 44. For this kind of paper, an easy-to-read description of how the bounding procedure works is an important part of the contribution and I had to read this text four or five times before I felt as though I understood it. This is a sign it could be improved.

4. In the KLM application, a natural way to address the critique is to focus on mid-career earnings, where the age-earnings profile is flat on average and the returns to experience are approximately zero. It would be helpful to provide readers with guidance about how to choose dependent variables to mitigate exclusion-restriction violations through this channel. Intuitively, a given effect on years of experience may have very different effects on the outcome, and hence very different implications for the exclusion restriction. It would be helpful to provide readers with guidance as to what reasonable values of gamma might be, given what we know about the setting in KLM.

5. An obvious but probably not helpful suggestion is that it would be great to actually conduct an empirical analysis of earnings. I imagine those data may be difficult to obtain, though.

Essay 3: Siblings’ spillover effects on college and major choice: Evidence from Chile, Croatia, and Sweden

This paper is an important achievement in cross-national collaboration on an interesting research question. My understanding is that a revision of this paper has recently been submitted to the Quarterly Journal of Economics; I expect that the paper will be published soon. I don’t have much to add to the many comments the authors have already received and incorporated.

Overview

I enjoyed reading this thesis and learned quite a bit from it. In my view it satisfied the formal and content requirements for a PhD thesis in economics. I recommend the dissertation for a defense. The most important comments in this document are comments 3 and 4 on chapter 1. These can be addressed through changes in the text. Comment 4 on chapter two is also fairly important. I would not

make any of these issues a requirement for successful completion of the thesis but would strongly recommend the author make these changes to improve and expand the influence of the work.

Seth Zimmerman
Associate Professor of Economics
Yale SOM
November 23rd, 2020