Report on "The Flat Rental Puzzle", by S. Cho and J. Rust

This paper is motivated by a puzzling observation: Rental prices of cars offered by the major rental companies appear to be flat as a function of age and odometer values; moreover, the companies seem to turn over used cars too fast – the average holding period is 1.3 to 3 years, while a casual observer would presume that cars' durability today extends well beyond this time frame and that companies could make money by renting older cars to consumers at discounted rates. Rather than justifying these facts within a particular model of optimizing behavior (as is traditionally done in Economics), the authors offer an alternative, rather unconventional, explanation: the documented rental patterns are not necessarily the result of optimization. To make this point, the authors proceed in two steps. First, they present an econometric model that fits the data they have acquired from one specific rental company quite well. Next, they solve the dynamic optimization problem for this company, and compute the prices and holding patterns implied by optimizing behavior under fairly pessimistic assumptions regarding maintenance costs and rental rates for older cars. They show that optimization would result in substantially longer holding periods for used cars, and substantially larger profits for the company. This is a case where it seems Econometrics could assist and improve private decision making.

I have to say at the outset that I am a big fan of this style of empirical work, and I would very much like to see it published in major Economics journals such as the *Review of Economic Studies.* While I had some initial reservations as to how interesting the particular topic would be to a wide audience, I believe that the paper's contribution lies in its ability to convince readers that high quality econometric work can be of value not only to public policy makers, but also to private agents. This is a novel element in the literature, as the traditional approach in structural empirical work is to interpret the data patterns within theoretical frameworks that assume that economic agents already optimize (the only exception to this pattern I am aware of is Liran Einav's work on the release times of movie pictures, which he shows are not consistent with optimizing behavior). As such, this traditional approach is of limited value to private agents - given that they already optimize, what is there to do better? The current line of work opens up a whole new array of uses for Economics and Econometrics. More importantly, the authors have managed to convince the company to engage in an experiment, in which the company will implement the authors' recommendations and examine whether its profits truly increase. The results of this experiment will be interesting no matter what. If the authors are right, then their argument is proven right, and yes, private agents can benefit from listening to economists; if profits do not increase, then it will be interesting to examine where the authors went wrong – this will be a useful lesson for future modelers. At any rate, the experiment makes this exercise true science (in Popper's sense of the word), as it will give the authors' model and predictions a chance of being falsified – a test that empirical work in Economics is rarely subjected to.

Regarding the specifics of the paper, the empirical analysis is well crafted and transparent, and the paper is polished. I have few concrete comments for improvement as the authors are upfront about the assumptions they need to make and the potential

caveats. So my report will focus mostly on alternative explanations for the authors' observations and some expositional issues. But before I get to that, just one quibble: The paper is very similar to an earlier paper by the same authors titled "Is Econometrics Useful for *Private* Policy Making? A Case Study of Replacement Policy at an Auto Rental Company". The acknowledgements in this other paper seem to imply that it is coming out elsewhere. If it is, the authors should make sure to minimize overlap between the two papers. While it would be a pity to alter the exposition in the current paper in a significant way (as stated earlier the current version is very clear and transparent), it is best to be safe…

Coming to the substance of the paper, what the authors label as "the flat rental puzzle" actually consists of two "puzzles": (a) the fact that rental rates are flat as a function of car age and odometer value for *the range of car ages observed in the sample*; and (b) the fact that companies do not offer older used cars at discounted rates. I believe these two "puzzles" are distinct, even though closely related – a point the authors acknowledge in the concluding section. Regarding the first part, I would argue that the companies' policy of offering flat rates seems hardly puzzling *conditional on the age of the cars they offer* and can easily be justified by appropriate functional forms for maintenance costs and (dis)utility for age/odometer values. Specifically, to make my point in the language of the model, equation (7) of the paper that specifies the equilibrium rental rate R as a function of odometer value x, shows that the shape of the R(x) function will depend on the shape of the maintenance cost function m(x) and the expected depreciation. The subsequent empirical section demonstrates convincingly that m is flat as a function of x, which is hardly surprising given that these are fairly new cars. I would argue that the *perceived* depreciation of the car from the point of view of the consumer is also flat for the *range of odometer values that are observed in the sample.* Cars that are relatively new (less than 3 years old) hardly break, and maintenance does not increase linearly with age in that range. Consumers know that, and hardly care if a car is 1 or 2 or 3 years old. Then comes the point when everything breaks… Maintenance costs jump up all at once and consumers require big discounts to rent cars with ages above a certain threshold value. In short, I am arguing that both maintenance costs and disutility for older cars are highly non-linear. In the range of 0-3 year old cars, both the maintenance and (dis)utility functions are close to flat, justifying the flat rates we observe in practice; even if the disutility function were not completely flat, its slope is probably too small to justify the costs of figuring out a more elaborate pricing scheme for that age range, informing consumers, etc… The authors justify the upward sloping q(x) function by drawing from observations in the second-hand car market. But I do not think that insights from *purchases* in the used car market are particularly useful here. The "lemons" problem can for example justify why a new car loses a substantial fraction of its value the moment it is driven off the lot; however, it cannot justify why someone who is renting a car for a week and is thus not concerned about long-term effects or resale value would require a discount in order to rent a 5-month old versus a new car. My suggestion is to put less emphasis on that part of the puzzle.

The second part of the "puzzle", the companies' reluctance to hold cars for longer periods, seems more intriguing, especially given the reported results from a previous

experiment in which the company offered a 20% discount for vehicles that were older than 2 years and found out that most customers would rent these older vehicles and not newer ones. The hypothesis that rental companies are missing something here is plausible; but it is equally plausible that the managers of these companies know something that we (as outside economists) do not – a point that the authors again candidly acknowledge. As stated earlier, the experiment the company is currently undertaking will hopefully shed light on this question one day. At this point though I would speculate that to the extent that the authors are missing something, this is related to the demand side of the market. The companies' current behavior is consistent with a model of consumer demand in which there is a very high penalty for a car that breaks down. This is a-priori plausible; when consumers rent a car, they presumably value reliability substantially more than when they buy one – if I have a one-week vacation, I do not want to spend this week dealing with car problems… Or being stranded in the middle of nowhere waiting for AAA to come… I would also presume that the demographics of people who rent cars are different from those of the general population, with people who rent having higher incomes on average, access to credit cards, and lower price sensitivity. The interaction between people who rent cars and car rental companies could be described as a repeated game, in which the consumers keep coming back to the same company as long as they have a good experience, but never go back if they experience a problem with the car they rented, thus severely penalizing the company for an unreliable car. In such a setup it is easy to see why companies would place a high premium on reliability; given that the best predictor of reliability is age, or odometer value, it is sensible to keep the fleet as young as possible.

Still, one would argue that there should be *some* consumers out there who would be willing to rent older cars at discounted rates. Again, the results of the experiment will speak directly to this claim one day. Even with the experiment in place though, I have one final reservation regarding the general applicability of the conclusions. The data and experiment are specific to one company. This company is hardly representative – as the authors point out, it is a company that already keeps its cars longer than the major rental establishments in the U.S... By focusing on one company only, the authors ignore the cross-sectional dimension of the market that is quite interesting. For example, the fact that this company – despite offering older cars at discounted rates – has a very small market share in the U.S. market tells us something about the nature of demand in this market. To come back to my earlier point, this small market share could be interpreted as evidence that even though there are *some* consumers who would be willing to rent older cars at lower rates, there are simply not enough such consumers to justify a change of policy by the major companies. Or, to put it in different terms, it is possible that we are at a separating equilibrium in which consumers who value reliability (the majority of U.S. consumers) go to the big five, and consumers who are willing to sacrifice reliability in return for a lower rental rate go to the authors' company. In this case, the results of the experiment would still be interesting, but not necessarily applicable to other rental companies. Obviously, there is little the authors can do about this – after all, the other companies refused to share their data with them… But this is something to keep in mind when they present their results and make general recommendations.