Department of Economics
3105 Tydings Hall
University of Maryland
College Park, MD 20742

Phone: 301 405-3489
Fax: 301 405-3542
Email: **jrust@econ.umd.edu**
WWW: **gemini.econ.umd.edu**

May 10, 2008

Professor Andrea Prat
Department of Economics
London School of Economics
London, W23BD
England

Dear Professor Prat:

We thank you for your decision letter on June 15, 2007 that gave Sung Jin Cho and I an opportunity to revise and resubmit our paper, "The Flat Rental Puzzle". We were encouraged by your statement that "Both referees are very positive about your paper" and that

"Both referees – and myself – find what you are trying to do here quite exciting. Your paper shows that economics has made amazing progress. We now have theoretical and empirical tools to help companies make better business decisions. Obviously, these tools are far from perfect but the advantage of this style of work is that assumptions are transparent from the outset. Your paper is potentially very important and I can see this work being used as an example by many others."

We were sufficiently encouraged by this that we undertook a substantial effort to revise this paper and most importantly, to make sure it contains experimental evidence that was lacking in the previous version.

We can appreciate your concern about the overlap with our other paper, "Is Econometrics Useful for Private Policy Making? A Case Study of Replacement Policy at an Auto Rental Company" that is now accepted for publication in the *Journal of Econometrics* and should be forthcoming in the next year or so.

We have completely revised the "Flat Rental Puzzle" and we submit both papers to you and the referees to inspect to satisfy yourself that although both papers deal with the same rental car company, they really are very different papers that make very different points and the overlap is minimal. In our opinion the small overlap that remains would not be grounds for rejecting this paper for publication in the *Review of Economic Studies*, given that this revision reports the results of a field experiment that does not appear in the other paper, and this is the aspect you found so exciting about our work.

In summary, this revision of "The Flat Rental Puzzle" contains new data and new results that are not in our *Journal of Econometrics* paper. Our work convinced the company to run an experiment to test our prediction that by discounting prices of its older rental vehicles the company would increase its profits. To our surprise and to the great surprise of the executive of the rental company who authorized this experiment to take place, rental revenues actually *increased* by a significantly higher amount at the 4 "treatment locations" where the discounts were given compared to 6 "control locations" that matched the treatment locations as closely as possible but where no discounts were given (i.e. the company maintained its existing flat rental schedules at all other locations).

We are quite excited by this finding and it was the cause of the delay in resubmitting this manuscript. It was especially interesting to us that the outcome of the experiment went against the prior expectation of the executive of the rental company who expected the experiment would reduce profits and revenues at the treatment locations. The company had previously conducted an experiment that discounted rental prices of older cars and the general recollection was that the previous experiment was unsuccessful (although no data from it had been retained that would have allowed us to make an independent confirmation about what happened). But it suffices to say that the official at this company began the experiment we suggested with a great deal of skepticim that the results would be any different from their previous experiment.

Now that they have seen these results, they want to continue and extend the experiment. However we did not think it advisable to delay the resubmission for longer than we already have, especially since the first experiment already has produced fairly convincing results. The next experiment will begin in January 2009, but we were concerned that you might no longer be an editor at *ReStud* by that time.

The first experiment ended at the end of May, 2007 shortly before we received your decision letter. We had an impression that this first experiment was even terminated prematurely and that the executive in charge, having received preliminary reports from their sales offices, seem to have come to an initial conclusion that the experiment was not a success. There was additional delay in receiving the contract level data that we analyzed and included in a new section 6 of the paper, "Results from a Field Experiment". The executive told us that he was contemplating the creation of a completely new subsidiary that would be responsible for renting older used cars, and part of our delay in resubmitting is that we hoped to report results from this additional "field experiment." However after waiting for nearly a year, and realizing that the executive was distracted with other more pressing issues (including a imminent promotion to become head of the entire company), that the issue of creating a subsidiary fell to the "back burner".

We felt that it would not be a good idea to further delay the resubmission of this paper, lest it become stale and dimmer in your and the referees' memories. We did not want to treat your gracious offer to allow us to resubmit as an open-ended invitation, especially in view of the fact that your term as editor does not last forever.

So the rest of this letter provides responses summarizing how our revision dealt with the comments made by the two referees. Before I do that, let me try to succinctly summarize the difference between this version of "The Flat Rental Puzzle" and our forthcoming *Journal of Econometrics* paper so it is absolutely clear to you how we revised the paper to deal with your own main concerns, which I summarize below:

"Based on my own reading, I believe that as things stand the overlap is too serious to allow for publication. Rather than making specific suggestions as to what material should be allocated to each of the two papers, I prefer to leave this decision to you. However, in order for this paper to be publishable on the Review I would think that: (1) It should focus on the flat-rental component of the puzzle; (2) The analysis should be self-contained and the paper should be fully understandable to people who have not read your J of Etrics piece; (3) You should explain carefully what the relationship between the two papers is in the beginning, but then technical cross-references should be kept at a minimum; (4) It should report preliminary (but meaningful) results of the experiment that the car rental company is currently carrying out (I understand that they will be available shortly)."

I think you will see that we have made a conscientious effort to do what you have asked, and we put particular emphasis on your suggestions 1) to focus on the flat rental aspect of the paper (the *JE* does not), and 4) to report the results of a field experiment by the company that is not analyzed in the *JE* paper.

You will see that this version cites our *JE* paper in both the abstract and in the introduction, and this revision makes it very clear that we are using an econometric model that was developed in the *JE* paper, but the econometric details and methods are more technical and of greater interest to econometricians. Including too much discussion of these issues would have distracted from the main *economic* issues discussed in this paper. We only summarize enough of the *JE* paper to convince readers that the model is appropriate for simulating the operations of the firm at the level of individual contracts and vehicles and that this econometric model provides a very good approximation to the operations of this rental car company – at least under it *status quo* rental pricing and replacement policy.

However the vast majority of material in this paper is not in the *JE* paper including the entire discussion in section 2 that shows how theory predicts that competitive rental car prices should be a declining function of odometer or age, the dynamic programming analysis in section 5 that computes the optimal strategy for the firm under hypothetical counterfactual non-flat rental pricing strategies, and of course the analysis of the experimental results in section 6.

Thus the only sections of the paper where this is slight overlap is in section 3 where we describe the company and the data we have (since the *JE* paper analyzes data from the same rental car company), and section 4, where we summarize the main results of the econometric model that we develop and estimate in the *JE* paper. But you can see by comparing the two papers that there is a huge amount of detail about the data set, the econometric model, and the methods we used to estimate and test it that is contained in the *JE* paper but is not in the "The Flat Rental Puzzle". Instead we just refer readers who are interested in the details of how we developed and estimated our semi-Markov model of the firm's operation to the *JE* paper.

We think this is an optimal division of material between the two papers, since we think the focus on econometric details makes the other paper more appropriate for an econometric field journal (snce the *Journal of Econometrics* is considered the leading econometrics field journal, the fact that our other paper has been accepted there is evidence that our econometric model is an independent contribution). However since the *Review of Economic Studies* is the top ranked general interest journal, we have written "The Flat Rental Puzzle" to appeal to a broad audience. We think this paper combines ideas from economic theory, econometrics, dynamic programming, and experimental economics in a way that is accessible and hopefully of interest to a wide audience, including a growing literature in "behavioral economics".

In particular, "The Flat Rental Puzzle" not only describes an interesting new empirical application, but it is unique in showing how economic theory motivated real world decision makers to undertake a field experiment. We believe our paper is unique in demonstrating that economic theory and economic model building can actually affect real world decision makers, and showing how effective the interaction between model building and real world experimentation can be in improving our understanding of economic phenomena.

Given the relatively specialized nature of the focus on rental cars, we also try to make clear in both papers what the broader significance of the papers are. In the past, my own best work has been quite narrowly focused, such as my paper "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", a paper that was awarded the Frisch Medal at *Econometrica*. In this paper the editor, Angus Deaton, never required me to say anything about the broader significance of the work. But these days, it seems that editors do not like to assume that readers will find the broader implications of such narrowly focused papers to be self-evident, so we have taken the space in our longer than normal conclusion to make it explicit what we think the more general contribution is, even at the risk of annoying readers (including the referees) for being too "pedantic".

Of course, we can trim down the conclusion further and let readers draw their own conclusions about the broader significance, if any, of our work. But I do think it helps you in judging the issue of overlap to differentiate between the "broader scientific contribution" of the "Flat Rental Puzzle" and our *JE* paper.

We see the main contribution of the *JE* paper as showing how one can build an econometric model of the overall operations of a firm by adapting econometric methods that have been developed for duration and transition analysis. Our semi-Markov process model of the firm is a new contribution to econometrics because of the way the various pieces of the model fit together to provide a "microfoundations model" of the operations of the overall company. Most econometric models of firms take a "top down" approach, focusing on modeling sales revenues as an independent stochastic process, using time series methods to forecast how revenues evolve over time, etc.

Beyond the methodological innovation on how to model operations of rental car companies, the broader point that we stress in our *JE* paper is that *the value of econometric policy analysis is most effectively demonstrated in applications involving firms, since the "treatment effects" of interest for firms are "profits" and these are far easier to measure than in cases of consumers or governments, where policy changes have consequences in terms of individual welfare or social welfare, and where heterogeneous distributional effects makes it far less clear whether specific policy changes are an "improvement" of not.*

We do not view the "Flat Rental Puzzle" paper as a methodologically oriented paper in the same way our *JE* paper is. Instead we view the main contribution of "The Flat Rental Puzzle" comes from posing and solving a unique new economic puzzle that no previous researchers appear to have noticed before.

We think there are two other distinct issues of broader significance in "The Flat Rental Puzzle". First, our paper calls into question the mainstream economic view that firms are best modeled as fully rational profit-maximizers, and instead suggests that theories more closely akin to Simon's notion of *satisficing* may be a better conceptual framework for understanding behavior of many firms, especially firms such as the one we are studying that are already highly profitable.

Secondly, our paper illustrates the value of economic modeling in an environment where real world experimentation can be costly. We believe our paper shows that the use of numerical methods and stochastic simulations are enabling us to develop economic models that are increasingly realistic and increasingly *credible,* so that we are now approaching a point where they are in essence serving as *laboratories for simulating the effects of counterfactual changes in firm strategies that make experimentation far less costly for the firm.* In our case, the executives of the rental company were able to understand our model, and were sufficiently convinced by the logic of its predictions that they were willing to undertake an experiment to test these predictions, *even though the predictions of our model were contrary to their intuitive prior expectations.*

We believe there is a lot of promise for economic modeling in a variety of other, admittedly more important, situations, but the issue of *credibility* of economic models is a very critical one in our view. If real world policy makers and decision makers do not find our models understandable or credible, then they will not even be willing to undertake experiments to test whether our models are any good in the first place. This can lead to an unhealthy schism between theory and practice, with theory and modeling becoming increasingly disconnected from reality.

Our paper touches on this by mentioning some (in our view) rather far fetched explanations for the flat rental puzzle based on a hypothesis that there are multiple oligopolistic equilibria in the market, one of which can lead to an absence of any price differentiation based on observable differences in the characteristics of goods. This is Wilson's theory, and of course we have the greatest respect for him. We simply quoted a private communication from him that suggests that he too thinks this sort of highly sophisticated theoretical rationalization may be rather far fetched and that the flat rental puzzle really is a problem where it will be quite challenging for theorists to provide a compelling and convincing rational explanation for.

Thus, we would hope that this paper will stimulate a lot of additional work, both theoretical and empirical, to see if our explanation holds in other contexts, or if there are compelling alternative theoretical/rational explanations for flat rental schedules that our analysis has overlooked. We welcome the challenge that future researchers might try quite hard to show why our solution to the puzzle is wrong, or is not the best possible solution.

Response to R1

We thank the referee for these helpful comments. We agree with R1's basic assessment of the contribution of the paper, i.e. that it is primarily an empirical paper motivated by inconsistencies with what economic theory predicts, and it shows how it is possible to relax and test the hypothesis of profit maximization:

"I have to say at the outset that I am a big fan of this style of emprical work, and I would very much like to see it published in a major Economics journal 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 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)."

We do know of another paper besides Einav's: it is Levitt's paper that tests whether an MIT Economics PhD who switched careers to become a seller of donuts and bagels for a living is maximizing profits or not. Levitt concludes that the "bagel man" is not maximizing profits because he prices his products in the elastic region of the demand curve. However Levitt's approach is very different from the one taken here: Levitt does not attempt to formulate or solve an explicit model of the bagel man's decision problem, and his paper does not offer any concrete suggestion about what the bagel man might do to increase profits. In particular, since Levitt appears to know the bagel man personally, it would seem natural for Levitt to suggest experiments that would change the prices of bagels and donuts to see if profits can be increased, providing a direct test his claim that the bagel man is behaving suboptimally. If Levitt had done these things, then we would acknowledge that he would have "scooped" us and the novelty of our contribution would be lower as a result.

While we have much more admiration for the level of sophistication in economic modeling in Einav's paper, it too fails to do what we have done in our paper, i.e. specify a concrete alternative strategy and demonstrate via an actual field experiment that an alternative strategy for the timing of introduction of new movies results in higher profits. It would clearly be unrealistic to expect Einav to influence movie producers to conduct field experiments with their movie release times: this just points out the difficulty we noted above about the highly costly nature of firm experimentation.

Further, we think the "jury is still out" on whether Einav has a convincing explanation that studios are behaving suboptimally. There is another paper on this industry by Natasha Zhang Foutz and Vrinda Kadiyali "Competitive Dynamics in the Release Date Pre-announcements of Motion Pictures" and they conclude that release dates are consistent with a dynamic Markov-perfect equilibrium model of competition. They note that "Our results demonstrate that studios are forward-looking and maximize intertemporal expected payoffs during decision making; ignoring such competitive dynamics produces inaccurate inferences on studios' decision processes and strategic interactions; and accounting for competitive dynamics increases a studio's predictive power of competitor decisions and responses to pre-announced release date decisions." (p. 31). We did not have the space to get into this debate about which of these two papers is more convincing, so we opted not to cite either paper. In addition, whatever goes on in the motion picture industry seems of limited relevance for the car rental industry.

We are careful not to claim to be the first paper to find evidence that some firms may not be profit maximizers. However we do think our paper is one of the first to show how economic theory can be sufficiently realistic and credible that it could directly affect real world decision makers and influence them to undertake costly experiments to test whether a predicted alternative more profitable strategy really would be more profitable. Furthermore, I think this is one of the first papers to show the full cycle, by reporting evidence that the predictions of the theory are confirmed from the results of specially designed field experiment.

There is one paper that we think is more impressive than our's in terms of how economic theory affects real-world policymaking. It is a paper by Erin Mansur and Matthew White's "Market Organization and Efficiency in Electricity Markets." Their paper provides convincing empirical evidence that a new online auction market for electricity (which was designed in part by economists), substantially improved overall market efficiency. While there is no field experiment in this paper comparable to our's, the "before/after" comparison of outcomes they provide is fairly convincing. We think these sorts of illustrations and examples are extremely strong selling points for the value-added from economic theory and economic modeling, and points to the huge potential benefits of a tight interaction between theory and empirical work. However if you read closely, Mansur and White show that "economic theory yields ambiguous predictions" and it is not clear precisely how big a role theory played in the specific design of this auction market. This is also the case in the FCC "spectrum auctions" where in some sense the auctions are far too complex to model theoretically. This is in part why we have been content to start with smaller scale problems where we can draw a closer link between the theory and empirical work. So our paper tries to be as clear as possible on how theory and modeling motivated the experiments that the rental company undertook, although it is always realistically the case that there is unavoidable gap between the theory and the practical implementation of almost any idea.

We have already addressed R1's main "quibble" i.e. the potential overlap with our forthcoming *Journal of Econometrics* paper. The main remaining comment of R1 is one that is contrary to your recommendation, namely, R1 suggests that we place more emphasis on the puzzle of why firms replace their rental cars as quickly as they do. But for the "flat rental puzzle", R1 states that "My suggestion is to put less emphasis on that part of the puzzle" because if cars are replaced sufficiently quickly, "In the range of 0-3 year old cars, both the maintenance and (dis) utility functions are close to flat, justifying flat rates we observe in practice; even if the disuility function were not completely flat, its slope it probably too small to justify the costs of figuring out a more elaborate pricing scheme for that age range, informing consumers, etc. . . .".

Given that R1's advice and your advice conflict, we opted to follow your advice, but we did also take heed of R1's advice by describing a *joint puzzle* but we put primary emphasis on the flat rental puzzle as you suggested we do. We think the results of the experiment are not consistent with R1's view that there is small return to the firm to discounting rentals of older cars if consumers are approximately indifferent. If this were the case, the 13% average discount offered in the experiment should have caused a near 100% substitution from new cars (rented at full price) to older ones of the same make and model, and while we observed a strong substitution effect, we do provide strong evidence that new and used cars are not near perfect substitutes for all customers and the firm can design non-flat price schedules to effectively discriminate among consumer tastes for an observable characteristic that they obviously differentiate on.

We have been careful to remind readers (and R1 seems clearly aware of this) that the company we study keeps cars approximately *twice* as long at the top 4 American rental car companies. So R1's objection has less force for the company we study, although we agree that R1's point of view has more weight for the top four or five U.S. rental car companies. However even if we acknowledge that flat schedules may make sense *conditional* on the rapid replacement strategies of the top four U.S. companies, the logic for the rapidity of replacement is still quite a puzzle: as we note in the introduction, even if we do acknowledge that rental car customers are approximately indifferent between cars of different ages or odometers when these cars are sufficiently new, the used car market is *definitely not indifferent* so when it comes to these companies' replacement decisions, "we still have the puzzle of why this company insists on selling its cars so soon to suspicious buyers instead of continuing to rent them to their trusting rental customers."

So we think and hope our revision balances your desires and those of R1, and that R1 will feel his/her suggestions are reflected in this revision. R1's final main suggestion is that "it is possible that we are at a separating equilibrium in which customers who value reliability (the majority of US consumers) go to the big five, and consumers who are willing to sacrifice reliability in return for a lower rental rate go to the author's company. In this case the results of the experiment would still be interesting, but not necessarily applicable to other rental companies." We think this is a valid point and we acknowledge it in our discussion in the introduction when we discuss the role of *Rent-A-Wreck*. We do not want to claim that the findings for the particular company we study necessarily apply to other rental car companies, but we do not see any obvious reasons why our results *should not apply to them.* We do strongly feel that the biggest rental car companies in the U.S. are probably even further from optimality than the

company we are studying since the big 4 companies are replacing their cars roughly twice as fast as the company we study.

There may be a different explanation for rapid replacements of vehicles by the top 4 or 5 rental companies that are not operative for the company we study. Specifically, several top U.S. rental companies have been able to take advantage of very favorable lease and buy back arrangements offered by Detroit automakers as a way of boosting their own sales of new cars. Clearly, if there is a high pre-negotiated buy-back price with automakers, then the effective rate of price depreciation is much less than it would otherwise would be, and then it is optimal to replace cars sooner than would otherwise be the case.

We do not discuss this in the paper because the company we study does not have any such deals with automakers, and we do not know the specific terms of buy-back arrangements that the top 4 or 5 rental car companies have negotiated with automakers. So we cannot judge whether their rapid vehicle turnover policy is or is not optimal in light of this.

However we have independent concerns about the buy-back arrangement since while these incentives do increase the short term sales of American automakers (who face increasingly stringent competition from foreign car sales and declining overall market share), the buyback terms appear to be so favorable that they could result in a long term loss to the U.S. automakers. Thus we may be able to solve one puzzle, why do U.S. rental car companies replace their cars as fast as they do?, only by raising a new one, "why do Detroit automakers provide such gifts to rental companies to try to raise short run sales at the expense of long run profitability?" We feel that it would take us too far afield to get into all of this, and combined with the confidentiality concerns (to protect our continuing access to the data from the rental car company we are studying), we opted not to get into this discussion of these issues in this paper. Perhaps we can address these other possible explanations (including potential tax explanations) in a future paper if we develop a relationship with a top 4 company that has a special deal with Detroit. But so far we have had no success in communicating with and getting data from any major rental car company other than the company we are studying and *Rent-A-Wreck.*

Responses to R2

We appreciate R2's overall view that "I like this paper. I think the paper documents and analyzes an interesting phenomenon — the "flat rental puzzle" — and employs both good data and methodology in the analysis. I also believe the paper gains from its focus on policy implications, and from moving away from the assumption that firm behavior is optimal."

Below we respond to the six main comments of R2.

1. Contribution R2 was confused after reading the first version of our paper about what the primary contribution is. R2 notes that there are several potential contributions. We think the revised version of our paper is much clearer about the primary contribution and it clearly distinguishes the contribution of "The Flat Rental Puzzle" from the contribution of our *Journal of Econometrics* paper. The main contribution of "The Flat Rental Puzzle" is what we discussed above, namely, laying out the puzzle, showing how economic models can help us test the hypothesis of profit maximization and provide predictions of results of counterfactual alternative more profitable strategies, and how field experiments can help us to test these predictions. We feel that it is now clear that the econometric contribution is in the *Journal of Econometrics* paper and the contribution of "The Flat Rental Puzzle" is mainly an economic one.

2. Consumer Heterogeneity R2 would "be interested in knowing: (a) if the results of the model change when I allow for correlated consumer heterogeneity in both taste for rentals and taste for "newness"; (b) empirically, what the observed structure of the rental market is, relative to the used car market – private consumers versus firms. Even very aggregate data (the fraction of consumers that hold more than one rental at a time, say) may help allay concerns that consumer heterogeneity is driving the results." For point (a) our theory section points out that our finding does still hold when there are heterogeneous consumers: we made the assumption of homogeneous consumers to make the presentation as simple as possible, but the section does point out that at least in markets with zero transactions costs, there is a "representative consumer" result that allows us to reformulate an equilibrium with heterogeneous consumers with an observationally equivalent homogenous consumer equilibrium with an appropriately chosen representative consumer. This is in Rust's (1985) *Econometrica* paper "Stationary Equilibrium in a Market for Durable Assets" and we make this clearer now. We also discuss why transactions costs will not destroy our basic result on why rental prices should decline and cite the Konishi and Sandfort (2002) *JEDC* paper that extends the equilibrium in Rust's (1985) *Econometrica* paper to allow for transactions costs. The prove the existence of equilibrium with transactions costs and characterize its properties. Equilibrium prices are still convex, and thus rental rates in a competitive market will still be a declining function of age or odometer value.

As for R2's point (b) above, i.e. to provide more of the aggregate empirical data R2 requests, our revision does point out that the age distribution of holdings of rental cars are very different than the age distribution of consumer holdings: rental car companies hold cars that are well under 5 years old and for the top 4 or 5 U.S. rental companies they are under 2 years old, far less than the 10 year average lifetime of vehicles owned by consumers. Clearly there is evidence of heterogeneity in car holdings since richer consumers with high tastes for newness buy brand new cars and poorer consumers and those with lower taste for newness buy and hold older cars. But this has already been discussed in Rust's (1985) *Econometrica* paper and related references in that paper that use that model to show it provides a good empirical approximation to prices, the distribution of odometer values and the age distribution of cars in the U.S. (see Rust's 1985 *Transportation Research* paper, "Equilibrium Holdings Distributions in Durable Asset Markets"). We just did not feel we had the space to discuss all of these aspects in this

paper. If you agree with R2 and are willing to give us a bit more space, we are happy to include more discussion here, but we expect you will be asking us to make further cuts, not giving us more space, assuming you and the referees conclude this version of the paper is now nearly acceptable for publication in *ReStud.*

3. Unobserved Vehicle Heterogeneity This is closely related to Akerlof's "lemon's problem" and the revision does cite and discuss a number of empirical papers that do not find much convincing empirical evidence that this is driving the rapid decline in used car prices. Further, our *Journal of Econometrics* paper did test for unobserved heterogeneity, including using accidents and early replacements and abnormally low sales prices as a indicators of whether some cars are lemons. While we do observe extraordinary dispersion of used car sales prices (see figure 2 in our paper), that we cannot predict using any of the observable variables at our disposal, we do not find correlation between low sales prices an earlier dates of replacement or increased chance of accidents, for example. Instead, our *JE* finds that accidents act like "instrumental variables" prematurely terminating the lifespan of rental vehicles, and we use this instrument to show that indeed, shortening a cars lifespan even further than the company already does via its replacement policy causes profits to decrease.

In summary, we do not think that unobserved rental vehicle characteristics, even though they are certainly present, somehow invalidates our results. A good example of an unobserved characteristic would be an unsightly cigarette burn on a car seat, or a strange smell, or unstable handling resulting from abuse that renters give to rental vehicles (i.e. running into curbs, throwing wheels out of alignment, or bending stabilizer bars, etc). There would be a separate question of why the firm does not provide on the spot discounts for things like this, at least *ex ante* before the customer leaves the lot with the rented vehicle. Although unobserved to us, they should be mutually verifiable to the rental car company and the customer, and the failure of the firm to provide discounts to consumers for vehicles that have such unobserved negatives is another aspect of the flat rental puzzle, not somehow a potential solution to it.

4. Calibration R2 makes the point that instead of relying on pessimistic assumptions, "it would be preferable to calibrate those assumptions off of actual data." Now that we have the experiment results in, we do just that. That is, we are able to observe how rental price discounts affect utilization behavior of new and used cars and we re-estimated our econometric model to account for this induced behavioral change and we have re-solved the dynamic programming model taking the actual behavioral responses into account. We found that the predicted results changed very little, however (as we dicuss in section 6), since there were offsetting changes: slightly lower utilization of new cars, a lot more utilization of older cars. But factoring in the discounts for older cars, the revenue and costs effects largely cancelled out. The cancellation of costs occurred in part also because maintenance costs are flat.

5. Style R2 found the first version "somewhat didactic and repetitive". We hope the revision has addressed this satisfactorily. Given our space constraints, we worked hard to remove repetition in the paper, and we try to communicate our ideas as simply as possible. But we are willing to undertake further cuts if you and the referees think this is in order.

6. Structure R2 finds that it is problematic to refer readers to our *Journal of Econometrics* paper for details on our econometric model. However in order to minimize overlap between the two papers and due to space constraints we did not feel we have the luxury to include an appendix to go into the econometric model in more detail. We do have a short discussion of the model in section 4 and we think we describe it sufficiently well that readers can understand how the model works without getting into the nitty gritty econometric details. Instead we just refer the interested readers to our *JE* paper. We do present visual comparisons of simulations of the model versus the actual operating data to convince readers that our model is a good one, at least for the firm's behavior under the *status quo*. We use the experiment to help convince readers that the counterfactual predictions of our model are good ones, at least in the sense that the experiments support the main qualitative prediction of the model that there exist non-flat rental price functions that do not change overall revenues by much, and allow the firm to increase its profits significantly via the reduction in replacement costs from keeping its cars longer. We now include some concrete calculations in our discussion of the experiment that shows how we can ascribe nearly all of the increase in profitability to the reduction in replacement costs, and that this is a lower bound on the increase in profitability.

We do acknowledge that our econometric model is not adequate to predict utilization changes and overall changes in the volume of customers due to our lack of direct data on demand and customer choices. We hope to get better data on customers in future work with this company when it runs its next round of experiments. We are urging the company to survey customers directly to learn more about their preferences and thereby provide us information that may enable us to design even more profitable nonlinear price schedules that extract more consumer surplus using improved models of customer behavior that are based on a better understanding of customer preferences and customer heterogeneity. However we think this will take several years to acquire this data and to carry out or much more ambitious "portfolio analysis" and multidimensional revenue scheduling for the hundreds of makes and models of cars this company holds, not even to mention the millions of combinations of different vintages of vehicles and other characteristics that can be used as bases for nonlinear price discrimination. Since this is a project that could take four or five years to complete, we believe it is beyond the scope of what we can realistically do in this paper.

We hope you find a sufficient contribution in this paper, and conclude it might be worthy of publication in the *Review of Economic Studies*. As it stands, we already have invested over 3 years of research on model development and data analysis to produce what you see here. Unfortunately, economic modeling is itself a costly and time consuming task. But over time, we keep seeing the costs of building models going down and this is resulting in more realistic models that are able to make better predictions and suggest more effective policies than would be possible via haphazard trial and error experimentation. We hope that our work will contribute to further improvements in this aspect of our science, and that we will have even more convincing illustrations of how economic modeling can have tangible real-world benefits in the future.

Thank you for giving us this opportunity, and for undertaking the effort to evaluate this revision and distinguish its contributions from our already accepted *Journal of Econometrics* paper. We do hope you conclude that the "The Flat Rental Puzzle" does not lie in the space spanned by our own previous work, or anyone else's.

Sincerely,

John Rust