UNIVERSITY OF MARYLAND
DEPARTMENT OF ECONOMICS

3114 Tydings Hall
College Park, Maryland 20742-7211
301.405.3266 TEL 301.405.3542 FAX

Andrew Sweeting
Associate Professor
sweeting@econ.umd.edu
Direct tel: 301-405-8325

November 10, 2015

To Whom It May Concern:

Letter of Reference for Yichen (Christy) Zhou

I am pleased to write in support of Christy Zhou's application. Christy is an environmental economist whose job market paper uses a structural I.O. model and I have got to know Christy well in the eighteen months since I joined her dissertation committee. I feel well-placed to describe Christy's job market paper, and her attributes as a researcher. I recommend that you read Maureen Cropper's letter for a more complete overview of how Christy's research agenda fits within the field of environmental economics, and for a comparison with other environmental students that Maryland has produced. My bottom-line view is that Christy has developed the skills required to be an excellent researcher and I recommend that all environmental and policy schools and all economics departments outside of the top 15 should seriously consider hiring her.

Christy's job market paper evaluates how different policies would affect the fuel efficiency of new car sales, accounting for the fact that automakers will react to any policy, whether it affects demand or supply, by not only changing vehicle prices but also changing at least some vehicle characteristics when they update their vehicles between model-years.

This question and Christy's approach to answering it are well-motivated. My understanding of the relevant environmental economics literature is that it has focused either on measuring long-run trends in fuel efficiency or predicting very short-run responses where automakers can only change prices. But I.O. research has found pretty consistently that when one allows for non-price attributes to be endogenous one finds changes in welfare or product sales that are qualitatively or quantitatively quite different (one can think of examples coming from newspaper, radio, ice-cream and airline markets, where, for example, changes in product attributes can really matter for the welfare effects of mergers). Automakers are able to adjust quite a few engine powertrain/technologies quite quickly (possibly, for example, by including some of the features found in the more select model trims in one model-year into the base trim in the next year), so it naturally makes sense that we should see how predictions change when product characteristics are endogenized. In addition, a structural model is required to answer this question as in the US policy remained relevant constant from the mid-1980s to the late 2000s, so we cannot just measure the effects from the data.

Christy uses a two-stage static model (I will return to the good reasons for doing this in a moment). The structure of consumer demand is nested logit, where consumers' preferences depend on performance attributes of the vehicle, price and the cost of travel in dollars/mile given the fuel efficiency of the vehicle. Here Christy assumes that new car buyers expect the current price to be the fuel price when they are driving the vehicle in the future, which is common assumption in the auto-demand literature. On the supply-side firms play a two-stage game. In the second-stage they set prices conditional on vehicle attributes, while in the first stage they choose performance attributes, whether to include (adopt) a set of engine technologies that affect fuel efficiency, and spending on R&D, taking into account how these choices will affect prices. Fuel efficiency itself is not chosen directly but is determined by a frontier equation that depends on the technologies, R&D spending and performance attributes and which determines how performance and efficiency can be traded-off. A key assumption that Christy makes is that performance and technology attributes are treated as continuous variables. To do this, she uses the proportion of model-year sales that include the technology as her measure of adoption (so if the EX trim has the technology and the LX does not, and 70% of sales are of the LX trim, then the value of the adoption variable for that technology will be 0.3).

Christy estimates the demand and supply parts of the model simultaneously using GMM. Given annual, aggregate data the demand model produces a natural estimating equation, following Berry 1994. The supply-model gives a set of first-order conditions for pricing and for the choice of vehicle characteristics in the first stage. It is here that modeling the characteristics as continuous is incredibly convenient. Christy follows the work of Sofia Villas-Boas, who modeled pricing decisions in a vertical structure, to calculate the derivatives of second-stage prices with respect to first-stage choices. Given the number of characteristics that Christy allows to be endogenous this is a far from trivial exercise, and, having done this type of exercise myself, I admired Christy's persistence in making sure that this was done correctly. To estimate the model Christy also has to specify instruments for each equation. Finding valid instruments that have power in models where characteristics are endogenous is not easy, and I know of no really convincing examples in the literature (for example in my own work on dynamic radio station format choice I rely on timing assumptions; in Ying Fan's newspaper work, she exploits the structure of overlapping markets, but it is not clear that these instruments, while appealing from a validity perspective, have a great deal of power). Christy tries to exploit the fact that some of the technologies observed in vehicles were likely chosen some years before the current model-year update, and she uses these choices (which she labels `grandfathered technologies', although I find the term a little bit confusing) and the same types of choices for competitors as her instruments.

The model has a large number of parameters but the estimated coefficients are pretty sensible, with demand elasticities, for example, well within the range found in the previous literature including work that has allowed for a more flexible random coefficients structure, following BLP. On the supply side she finds that adding most fuel saving technologies to vehicles can be quite costly, and that firms' patents stocks contribute appear largely associated with process innovation, as they cause firms' implied marginal costs to be lower.

Based on the estimates, Christy considers three counterfactuals, focusing on what would happen if there was a policy shock in 2006. In each case she re-solves the equilibrium first-order conditions, adapted to take account of the change, and then computes consumer surplus and the fuel efficiency of the fleet of new cars. The first considers what would happen if there was a

$1/gallon increase in the (federal) gasoline tax. Of course, federal gas taxes are much lower than this (in the 20 cent range) but we do see variation in gas prices of more than one dollar in the data, which should be at least partially informative about how automakers would respond. The direct effect of increasing the gas tax (holding both characteristics and vehicle prices fixed) is that consumers will switch away from gas-guzzlers to more efficient vehicles. When prices are endogenous, this effect is softened as makes of fuel efficient cars respond by increasing their margins relative to makes of less efficient cars. When product characteristics are endogenized, Christy finds that the effects become even smaller. While a number of things change, the main driver of this result appears to be that makers of fuel-inefficient cars increase their investments in process innovation that lowers their marginal costs, so that they cut prices more. This in turn takes some demand away from more fuel efficient cars. While one should be concerned with the possibility of multiple equilibria in this type of model, Christy takes the sensible, practical approach of starting at the equilibrium in the data and then slowly raising the gas tax, so that each time we are looking for an equilibrium which is hopefully nearby the one previously identified.

The second counterfactual looks at the effects of an R&D subsidy. The effects of this are more straightforward, as it encourages process innovation and reduces marginal costs, but this does not have large effects on the fuel efficiency of the fleet, as prices of fuel-efficient and fuel inefficient cars can fall by roughly the same amount. The final counterfactual considers the effects of a hypothetical merger (GM/Chrysler). In general we would believe that a merger would soften competition and thus lead to higher prices and less investment on quality attributes. However, because all of the models of a firm benefit from process innovation, Christy finds that the scale effect tends to lead to the merged firm reducing its costs. This merger synergy leads to the reduction in competition being partly offset. As she moves forward, Christy is going to look at how changes in market structure could interact with increases in the gas tax: this would be useful and relevant from a policy perspective as one continually reads about the belief of automakers that the industry will become more consolidated, partly because of the cost of investing in technology.

As there are several things that one might view the model as missing, let me explain why Christy has structured the model in the way that she has. First, the framework is static – based on annual decisions for the current year – not dynamic. In practice, allowing for dynamics with a combination of both pricing and several product characteristic decisions is some way beyond the current literature. In my own research I allowed for dynamic format choices of radio stations, but I completely ignored active pricing decisions (revenues came from a reduced-form revenue equation), I dealt with common station ownership in a very imperfect way, I only considered a simple programming type choice and I was never able to convincingly deal with the possible multiplicity of equilibrium. Ex-post I believe that to answer the question of how changing music licensing fees would have affected format choices in the long-run, a static framework, where I might have addressed these issues in a much better way, would have been superior. For the most part Christy is looking at the decision to include technologies with which car-makers are likely familiar with into a particular model, and it seems plausible that car-makers have relatively short-horizons for this type of decision. Clearly, dynamics are likely to be more important for R&D and patenting choices, and I expect that Christy will end up downplaying this part of the model when the paper is submitted for publication. However, this is not really a problem as there will still be a lot of `meat' to the paper.

Second, she treats technologies as continuous, whereas they have a discrete aspect to them (you can have a 4-valve or a 6-valve engine, but not a 4.33-valve engine). Once again, this reflects the fact that we know how to estimate and (importantly given that we want to do policy counterfactuals) solve models with several continuous characteristics, where we can exploit first-order conditions, but not discrete choices unless we can reduce them to a simple in/out, Democrat/Republican flavor and only then with a very small number of competitors (and typically no way to deal with the multiple equilibria problem that arises much more strongly with discreteness). Here, I think that Christy may end up finding a way to project down the multiple technologies that she currently looks at into some single-index measuring the fuel efficiency-relevant technical sophistication of the engine and powertrain. The way the technologies enter the fuel efficiency frontier might provide a metric for doing this type of projection. With this done, the continuity assumption would be less objectionable, and solving the model and interpreting the results would become simpler, as there would be fewer margins for the firm to change.

Finally, Christy has focused on gasoline vehicles, and not hybrid, electric or diesel cars. This makes sense as gasoline cars were entirely dominant throughout the period for which she has data, but clearly this imposes some limits when she talks about R&D, for example, as much of the product innovation process is being missed. Christy has been careful to look at patents that are relevant for gasoline cars, but one should keep this in mind when reading what she concludes about the effects of R&D spending and possible R&D subsidies.

My overall assessment of the paper is that with some of the adjustments just noted, the paper is publishable in at least a top field journal, either in environmental economics or I.O, and I could also imagine the paper fitting well at a more general interest journal such as the International Economic Review, one of the AEJs or Quantitative Economics. The paper demonstrates Christy's ability to (i) assemble several distinct datasets (for example, on specific vehicle attributes from the EPA, and on patents from the OECD); (ii) write down a rich model that is estimable given the data she has assembled; (iii) sort out all of the kinks that arise when estimating structural IO model; and (iv) write up the results in way that is focused on the policy questions that motivated her to work on the project, rather than the method itself. I should also say that Christy worked tremendously hard and very independently on the paper, and has taken lots of opportunities to present the paper as on-going research at different forums, and this has partly led to her taking on some new projects with top-notch environmental economists such as Joshua Linn at RFF. She has the will to succeed, and that will serve her very well wherever she lands after the market.

In terms of teaching, Christy could obviously teach environmental economics at the Ph.D. level, but she could also, I think, teach the methods component of an I.O. course, and a full I.O. course at masters or undergraduate level. She has strong teaching experience at Maryland, including in our professional MA program, and while I have not seen her in the classroom I know that students have found her to be very approachable, and, on a personal level, I always enjoy interacting with her.

Let me compare Christy to several other environmental economists that I have advised. I view Christy as having a better paper and likely to produce better research in the near-future than Ron Chan (a UMD student who is now at the University of Manchester in the UK), and quite comparable in overall quality to Ralph Mastromonaco (a Duke student who is an Assistant Professor at the University of Oregon) and Beia Spiller (a Duke student who began at RFF, and

now has several well-published papers). She would lie between these students in terms of her technical skills.

As mentioned at the beginning, my bottom line is that Christy would be a very good fit for an economics department (including within a business school) outside the US top 15, an environmental school or a policy school that is looking for a serious and skilled researcher. I would recommend her without hesitation to any non-academic position.

If you have any questions that I would be able to help with, please do not hesitate to get in touch.

Yours faithfully,

Andrew Sweeting
sweeting@econ.umd.edu