
Reviewer #1

This is a worthwhile paper that, on the whole, makes a positive contribution to the literature on discounting and uncertainty. Based on a convenience sample of students at roughly 100 U.S. and Canadian colleges and universities, the authors conducted a web-based survey that allowed them to measure respondents' subjective preferences using discrete-choice econometrics. The "punch line" goes something like this: On average, respondents had predictably high discount rates and a relatively low degree of risk aversion, though there is an interesting degree of heterogeneity in the data.

Response: We seem to have struck a chord with this referee in terms of our intermediate results concerning the empirical measurement of individual-specific discount rates and risk aversion papers. However, we seem to have dropped the ball on making the point that these are merely intermediate results, rather than the endpoint of the paper. Our goal is to demonstrate first that it is *possible* to use auxiliary choice information to identify distinct estimates of these two (typically unobservable) individual characteristics. Second, we demonstrate that these important characteristics of individual preferences can be used directly to explain heterogeneity in attitudes towards climate change mitigation policies. It is undeniable that time- and risk-preferences play a key role when policies involve large near-term costs and long-term benefits which are highly uncertain. Climate policy choice, unfortunately, is a "poster child" for these conditions.

As someone who works mainly on the theory of discounting and intertemporal choice, I have relatively little to say about the paper's statistical methods. The authors seem to be very much up-to-speed regarding choice modeling, survey methodology, and the interpretation of survey data using discrete-choice models.

Response: This referee's experience with theoretical aspects of discounting and intertemporal choice probably explains his/her focus on just this portion of our paper. We all tend to focus on the most-salient portions of a paper.

I do, however, have some important concerns about the paper's theoretical foundations and links to the existing literature on intertemporal choice. The authors aim to present a theoretically coherent model of intertemporal utility maximization. But they work with a setup in which: (a) the net benefits of policy interventions are measured in present-value equivalents using a particular monetary discount rate; and (b) risk aversion is defined over the distribution of present-value net benefits. Theoretically this is ad hoc and, indeed, departs from the way that intertemporal decisions under uncertainty are typically approached in both welfare economics and finance.

Response: From these comments, we realize that the paper needs to explain more clearly the distinctions between the first joint model to measure individual-specific time- and risk-preference parameters, the second "policy choice" portion of the study which conditions choices on these separately estimated preferences (in a much more ad hoc framework, dictated by the intractability of eliciting from each respondent their

subjective long-horizon time profiles for both the benefits and the costs of climate change mitigation.

Specifically, the statements are incorrect that “the net benefits of policy interventions are measured in present-value equivalents using a particular monetary discount rate.” Individual discount rates are parameterized in terms of a long list of individual characteristics, so that they vary systematically across individuals. The individual discount rates used to convert future net income amounts to present-value equivalents is estimated simultaneously with the (also systematically varying) risk aversion parameter that measures the curvature of utility with respect to present-value net income.

We now emphasize that our empirical models are “utility-theoretic” in terms of the amount of structure that motivates the choice models, and the commonality of this structure across the several different versions of the lottery winnings choices and the risky investment choices. We concede, however, that the specification remains a static choice, based on discounted net income and a specific non-linear form for the utility function. We allude to fully dynamic models of intertemporal optimization as an approach that would be preferred, in theory, but one which is difficult to implement empirically due to fairly prohibitive data requirements (which also far exceed the data actually available for this application).

One thing that we know is that, while individuals may have a unique rate of pure time preference, the discount rates they apply to future net monetary benefits generally depend on the risk characteristics of the investment or action under consideration. The key issue concerns the covariance between the net benefits of a given action and the decision-maker's future consumption stream or overall financial position. The authors rightly point out that individual economic actors may have different rates of pure time preference and risk aversion. They overlook, however, the point that financial markets provide arbitrage opportunities such that the risk-adjusted rate of return that a person demands on a given investment should be equated for each individual.

Response: We do assume that our respondents impute zero risk to the possibility that the government lottery agency may default on its commitment to make annual payments over the specified number of future years. To date, however, we know of no cases where any government lottery agency has defaulted on an annual payment commitment. If any such cases do exist, we should certainly cite them.

The problem with a story about financial market arbitrage, in an empirical setting, is that very few consumers (especially in the U.S.) have much at all in the way of savings, and those who do have savings also tend to run higher-rate balances on credit cards. Imperfectly functioning capital markets, rather than perfect capital markets, must be entertained as a reasonable maintained hypothesis in most empirical work.

One problem with the model specification assumed by the authors is that it represents the household income of each respondent at the time of the survey (Y_i) without accounting for the future time-path of the respondent's income stream.

Response: This is less a problem with the model specification than it is with our ability to elicit subjective impressions of the future time-path of the respondent's income stream. Our survey does elicit each individual's subjective future income, in brackets, in each of four future benchmark years.[working on incorporating this into the estimation]

Wealth - the present-value of future earnings including payments to labor and human capital - would be a more appropriate variable to measure in characterizing people's effective intertemporal budget constraints. Theoretically, it would be better if the authors gauged the net benefits of each prospect on a period-by-period basis given some assumption about the time path followed by the respondent's (stochastic) consumption path. Given this information, the authors would be better positioned to estimate parameters that captured each individual's degree of risk aversion and rate of pure time preference. Admittedly this kind of approach would be data- and/or assumption-laden.

Response: It would be an impressive and heroic feat to elicit from each respondent their period-by-period future wealth (including future earnings and payments to labor and human capital) so that we had a clearer picture about their subjective probability distributions on their possible effective intertemporal budget constraints. However, it is difficult enough to extract reliable information about current-period income. The referee is certainly correct in concluding that such an approach would be "data- and/or assumption laden." To be more specific, it would be exceedingly difficult and survey-time-intensive to elicit from individuals the type of data described here. The binding constraint in survey research is often the attention span of the respondent. About all we can do with respect to this objection is to point out in the paper that more-detailed time profiles of full income would be desirable and that future researchers should continue to delve into ways to extract more information of this type from individual respondents.

An alternative approach might rest on the observations that the theory of intertemporal choice suggests that:

1. Both risk aversion and time preference are captured by the (risk-adjusted) monetary discount rates that individuals employ in evaluating a given choice option.

Response: We argue that respondents assume zero risk concerning the government's ability to complete a series of annual lottery payouts if the annual payments are chosen over the lump sum payout. We concede in a footnote that respondents may have some fear that the government will collapse, or that they may die and the payments will not continue to be made to their heirs. Thus we argue that the monetary discount rate conveyed by lottery payout preferences is an essentially risk-free rate. We then constrain this risk-free discount rate to be the same discount rate that is relevant in the "investment choices" survey question, where an explicit risk is described. In this way, we believe that we are able to achieve separate identification of approximate discount rates and time-preference parameters more easily than this can be accomplished in a single choice scenario where risk aversion and discounting may indeed be confounded.

2. Given well-functioning capital markets, agents with heterogeneous underlying preferences should employ similar risk-adjusted discount rates in evaluating similar choice prospects.

Response: An assumption that there are “well-functioning capital markets” is appealing and convenient in many theoretical treatments, but it is often too strenuous in empirical contexts. Going all the way back to Hausman’s seminal empirical article about tradeoffs between capital costs and operating costs for air conditioners, it has been suggested that “imperfect capital markets” could one explanation for the remarkably high individual discount rates implied by many people’s choices.

These points support a modeling framework in which the expected net benefits of a given choice option should be discounted at a rate that reflects its perceived riskiness. Given this conjecture, the authors could presumably estimate the risk-adjusted discount rates that pertained to the various choice bundles they model. These monetary discount rates would reflect people's underlying pure rate of time preference, risk aversion, and beliefs about their future economic prospects.

Response: Is this recommendation asking for discount rates to be parameterized as a function of “perceived riskiness”? We believe that we have already purged our estimated discount rates of confounding risk preferences by combining the essentially “certain” tradeoffs embodied in the “lottery-payout” choices with the explicitly uncertain and future payouts embodied in the “risky investment” choices. This referee’s desire for “monetary discount rates that would reflect people’s underlying pure rate of time preference, risk aversion, and beliefs about their future economic prospects” seems to be moving in exactly the direction we are explicitly trying to avoid going. We have developed and implemented randomized choice scenarios designed to differentiate to the fullest extent between risk aversion and discounting—to undo the confounding of these two dimensions of preferences that has plagued previous researchers (e.g. Viscusi....).

In addition to the structural concerns raised in the preceding paragraphs, I have some comments on the exposition and presentation of the paper. For the most part, the text is crisp and well-written. On the other hand, the description of the behavioral model given on pp. 13-17 is terse and somewhat opaque despite the fact that the model itself is quite simple.

Response: The original draft of the paper presumed a basic level of familiarity with standard discrete-choice models. We have pointed this out in the revision. Random utility models (RUM) based on indirect utility-differences have been around for more than 25 years, so we elect not to use any more scarce journal pages than necessary reviewing these techniques. [what have we done to fix this?]

Finally, although the authors' discussion of their results is richly detailed, the general patterns in the data and the main conclusions that flow from the analysis are not adequately highlighted and placed front-and-center. I personally would prefer to read a shorter, more focused paper that emphasized the authors' core results on discount rates, risk aversion, and willingness-to-pay for climate mitigation policies. Instead the authors conclude their paper with a set of simulations that

seem to combine their empirical findings with contrary-to-fact scenarios (including assumptions such as risk neutrality or the imposition of specific discount rates.)

Response: We have added more description of the empirical models to the material that previously appeared in pages 13-17, while at the same time attempting to produce a “shorter more-focused” paper. This reviewer seems to want more commentary on the structural parameters of the model.

We still pursue the counterfactual scenarios because these are one of the motivations for developing these models. It has been suggested that perhaps we should not make social decisions based on higher private discount rates. Perhaps risk-loving citizens should not be permitted to dictate how we pursue climate change policy. These are relevant questions to ask—what would be demand for climate change mitigation under lower social discount rates, or in the absence of risk-loving (or even risk-avers) preferences. These question fall squarely into the realm of “libertarian paternalism”...whether we can continue to honor consumer preferences as much as possible, but at the same time, “fix” some undesirable “choice mistakes.”

In conclusion, the basic problem with this paper is that the authors have examined an explicitly intertemporal problem using a choice framework that is closely patterned after a model of utility maximization in a static setting. In my view the authors need to rethink the theoretical foundations of their model to come up with a more compelling problem formulation and specification.

Response: A fully dynamic intertemporal choice model would be, without argument, the correct way to proceed. However, the data requirements for such a model are staggering. Given the limited data we have to work with, it is indeed true that we first convert the alternative net income streams in each scenario to their present value (using an endogenously estimated individual-specific discount rate). Then we model each respondent as making a static contemporaneous choice between these different present discounted net income amounts, where utility is allowed to be non-linear in discounted net income so that the estimated curvature of the utility function allows us to impute a risk aversion parameter. This is indeed a rudimentary model, but it allows us to estimate crude measures of individual-specific time- and risk-preferences that vary systematically with observable individual characteristics.

Reviewer #2

This is an innovative and important paper that makes a number of contributions on the use of individual-specific time and risk preference parameters to account for heterogeneity in stated preference modeling of policy preferences for climate change mitigation. The unique aspect is that the authors use survey responses to a set of questions about hypothetical lottery investment choices and time preferences to jointly estimate discount rates and risk aversion measures. The fitted values from these individual characteristics results (on time and risk preferences) are then

used directly as determinants in the policy preference model (a conditional logit model on responses to a choice experiment). The fitted discount rate and risk aversion variable are shown to be highly significant, with more risk averse and patient individuals having a higher estimated willingness to pay for climate change mitigation.

Thus, the results appear to be intuitively appealing, even though the WTP estimates seem on the high side. Further a wide variety of counterfactual simulations are conducted using the modeling results to show how estimated WTP would vary under different controls for risk aversion (e.g., assuming risk neutrality), time preference and degree of respondent "informedness" about climate change, etc. However, the survey data is drawn from a web-based convenience sample of about 2000 college students (drawn from numerous universities and different classes), and the authors are clear to note that no attempt should be made to extrapolate from the results to any national estimate of household WTP for climate change mitigation. The methodological innovations on the use of time and risk preference information are certainly worthy enough to outweigh the weakness in using the convenience sampling. This is a highly sophisticated econometric paper that is also well written, well referenced and carefully crafted.

Response: This referee has successfully digested the main strategy and the main structural modeling results from the paper. The WTP results do seem to be on the "high side," but they are nevertheless similar to values estimated by Viscusi and Zeckhauser (2006) as mentioned in the paper. We are relieved to see that this referee has also picked up on the caveats that we offer, but that despite these caveats, the paper seems to have some value.

Specific Comments

Although the paper is well written, it does seem excessively long. If journal guidelines require it to be shortened then I offer a few suggestions. There is considerable amount of repetition in the paper that could be trimmed some. Footnote 3 could be dropped as it gets re-stated at the bottom of pp. 15-16. Footnote 6 could be dropped; it gets discussed later and no references are given to the experimental economics comparison. The discussion of the simple construct validity test on the subjective assessment of the policy priority seems like a minor point in the larger context and could be dropped. The last paragraph on page 17 could be dropped or considerably shortened. The discussion on page 19-20 could be considerably shortened.

Response: [Make these changed.]

Pg. 16. Eq. 1.5. Unless I missed it you haven't defined DC and IC yet in your notation. I assume these are the Domestic Cost and International Cost distributions/alternatives that vary in the policy choice scenarios?

Response: [Fix this.]

Pg. 22. The authors mention that it is possible to tease out significant sensitivity to the distribution of costs for some subgroups in the sample. Without asking the authors to change any estimation, they might tell us if these are age subgroups. That is the inference given that you are using a restricted (college age) sample, and the discussion in footnote 11. I would expect this to

be an important issue for climate change policy, so it is an important point, and it would be interesting to pursue if the absence here is simply due to the age restriction. It might also be of interest to see what the correlation is between the omitted cost distribution variables and the equity impact rating. I was wondering whether the equity rating (which shows up as highly significant in Table 5) may be picking up the effect of the cost distribution effects.

Response: [Clarify this in the paper.] Sensitivity to the distribution of costs is more evident in the tandem mail survey of a general population sample, which obviously includes a much greater proportion of working (and tax-paying) adults. In that survey, much less randomization of the presented information was feasible, so we are not able to estimate models which are as rich as those which can be entertained for the online survey fielded with mostly student samples from different universities.

The “equity” variable pertains to the individual’s subjective rating of climate change impacts under a policy of business-as-usual. Respondents were specifically asked to indicate roughly what share of the adverse impacts of climate change would be borne by the poorest 50% of the world’s population. If climate change policy could succeed in maintaining or restoring the status quo in terms of climate, then these avoided equity impacts will be among the “benefits” of the climate policy. A separate paper (by Cai, Cameron, and Gerdes) delves into the preferences over distributional outcomes in much more detail. To preview those results, differences in the initial incidence of the costs of climate change mitigation policies matter, but differently to different people, depending upon the extent to which they believe different parties should be held responsible for these costs. It is a culpability story as much as anything. It is also the basis for an entirely separate paper (and a dissertation chapter for Cai).

In terms of shortening the paper, one of my initial reactions was to simply drop the entire section C (pp. 29-32) providing the various simulations. Upon re-reading I think it is important and in the spirit of the paper’s key contributions to retain at least part of this material. To that end, I think the critical simulations are A-D; they show some of the variation in WTP, and the effects of simulating alternative assumptions about discount rates, risk aversion rates, and degree of subjective informedness. I suggest dropping the material on pages 31 and 32. The extra simulations are not necessary and the rest of the material on page 32 seems like a reiteration about sample selection issues, etc.

Response: [Take these suggestions and shorten the paper]

The paper makes the point a number of times including in the closing that that the estimated WTP values, even for this convenience sample appear to be in the range of previously estimated values in the literature for reducing global climate change and its impacts. I would like to see a little more discussion in the paper on this point, and **review of other extant estimates**. A concern I have with the particular implementation of this stated choice format (Figure A-3, p. 45) is that subjects are self-identifying their own anticipated impacts (by ranking impacts on oceans, ecosystems, health, agricultural and equity). This is not necessarily particularly problematic in itself, given that the main point of the paper is incorporating time and risk preference information in a stated preference policy choice model. However, it does raise some questions about how

valid it is to make any comparisons to other results if the change in environmental quality here is somehow heterogeneous or doesn't match well with the change evaluated in other studies.

Response: Our estimated values are within previously estimated ranges for (self-selected) student samples. They do, however, fall at the high end of estimated WTP among our tandem general-population sample of mostly adults.

Stated preference survey respondents bring to stated choice exercises their own subjective assessments of the severity of the climate change problem regardless of whether we elicit these or not. We could subsume all of the implicit benefits of climate change mitigation policies under generic alternative-specific constants for “complete mitigation.” However, it is helpful to know whether the expected scope of these benefits is systematically related to willingness to pay, or whether respondents are willing to pay the same amount regardless of what they believe will be accomplished by “preventing climate change.” In other work, we have demonstrated that willingness to pay for climate change mitigation appears to vary systematically both with expected levels of avoided impacts and the variance in each individual’s subjective distribution of impacts.