# Even the Representative Agent Must Die: Using Demographics to Inform Long-Term Social Discount Rates

Ethan T. Addicott, Eli P. Fenichel, Matthew J. Kotchen

**Abstract:** We develop a demographic approach for estimating the utility discount rate (UDR) portion of the Ramsey rule. We show how age-specific mortality rates and life expectancies imply a natural, long-term UDR for individuals at each age, and these can be aggregated into a population-level social UDR. We provide estimates for nearly all countries and the world. Our estimates fall within the range of those currently employed in the literature, and the empirical basis of our methodology provides a useful point of comparison for alternative assumptions about the UDR. We use our results to derive heterogeneous social discount rates across countries and explore the consequences for an integrated assessment model of climate change. We find that introducing regional heterogeneity of UDRs into the RICE model has a small effect on the business-as-usual trajectory of global emissions, yet a more substantial effect on optimal emissions and the distributional burden of emission reductions.

**JEL Codes:** H43, Q21, Q54

**Keywords:** social discounting, demography, Ramsey rule, IAM

FEW TOPICS IN ECONOMICS are as fundamental and generate as much controversy as discounting. Benefit-cost analyses of long-lived public projects—such as those related to environmental protection, infrastructure, education, and health—rely heavily on social discounting. While most regulatory agencies and subfields within economics have established procedures for discounting future benefits and costs, economists continue

Ethan T. Addicott (corresponding author) is at Yale University (ethan.addicott@yale.edu). Eli P. Fenichel is at Yale University (eli.fenichel@yale.edu). Matthew J. Kotchen is at Yale University and the National Bureau of Economic Research (matthew.kotchen@yale.edu). We are grateful to the Knobloch Family Foundation for financial support, Whitney Boone for valuable research assistance, and Lint Barrage, Glenn-Marie Lange, and Antony Millner for helpful comments and discussion. We also acknowledge helpful comments on a earlier version of the paper by Simon Deitz and three anonymous referees.
*Dataverse data:* https://doi.org/10.7910/DVN/DJGIHW

Received July 25, 2018; Accepted October 9, 2019; Published online February 5, 2020.

JAERE, volume 7, number 2.  2333-5955/2020/0702-00XX$10.00 https://doi.org/10.1086/706885

to debate what constitutes an appropriate discount rate and, even more fundamentally, how discounting should be applied. The potential consequences of the debate are significant because small changes in the discount rate and procedures can have a substantial influence on present-value calculations of long-lived projects. Discounting therefore plays a critical role in policy evaluation.

Economic analysis of climate change has brought many of the important issues about long-term discounting to the fore, showing that whether a more or less aggressive climate policy passes a benefit-cost test depends critically on the social discount rate. This is the central insight of the highly influential and contrasting contributions of Nicholas Stern (2007) and William Nordhaus (2007).[1] Both employ the Ramsey-Cass-Koopmans framework with a constant utility discount rate (UDR) within integrated assessment models (IAMs) of climate change. In this context, the UDR represents a social planner's rate of pure time preference between generations irrespective of differences in consumption.[2] While Stern uses a very low rate that supports more aggressive climate policy compared to Nordhaus, the differences in how they justify their assumptions are central to the motivation of the present paper. Stern follows classical economists and argues that the choice should be based on ethical considerations, whereas Nordhaus argues that discounting should be consistent with behavior reflected in observable market interest rates.

The Stern-Nordhaus exchange rekindled a long-standing debate about "prescriptive" versus "descriptive" approaches to discounting (see, e.g., Arrow et al. 1996), and subsequent papers have sought to further clarify the role of normative and positive assumptions implicit in the economics of climate change and discounting more generally.[3] While the existing literature furthers the understanding of conceptual issues surrounding the choice of social discount rates, there remains relatively little empirical guidance on how to choose the underlying parameters, especially with regard to the UDR in the standard Ramsey equation.[4] As a result, researchers and policy makers are typically left to choose parameters based on one of two approaches. The first is a reliance on some ethical or normative criteria, many of which push in opposite directions. The second is to back out values after calibrating to observable parameters, including a priori assumptions about what the overall social discount rate should be.

In this paper, we develop an alternative, demographic approach for estimating the UDR that serves as a useful benchmark. We show how age-specific mortality rates

1. Stern and Nordhaus have written several papers on the topic, but here we reference the two where the distinction and focus on the discount rate first arose.

2. The UDR is a parameter of the Ramsey discounting equation that is sometimes referred to as the social rate of pure time preference, in addition to other variants in the literature.

3. Examples include Weitzman (2007), Dasgupta (2008), Arrow et al. (2012), Goulder and Williams (2012), Schneider et al. (2012), and Gollier (2013).

4. The other parameters implicit in Ramsey discounting, which we will discuss in more detail later in the paper, are the elasticity of the marginal utility of consumption and the rate of growth in consumption.

and life expectancies imply a natural, long-term UDR for individuals at each age in a population, and these can be aggregated into a population-level social UDR. In particular, we show how the procedure gives rise to a UDR that a social planner might choose under the assumptions that discounting is based on mortality risk and all generations, existing and future, are equally weighted.[5] While the approach oversimplifies how individuals may, or may not, have preferences for the future, arguments can be made in support of either over- or underestimation. One's own concern for the future would tend to lower the UDR, whereas general impatience would tend to increase it. Unlike other approaches, however, the methodology that we describe has an empirical and observable basis, thereby making it a useful point of comparison for existing assumptions about the UDR in the literature. Indeed, our approach shares the same basic structure as the macroeconomic literature on life-cycle models that weight future utility based on survival probabilities.[6] In effect, the demographic basis of our approach is "descriptive" at the individual level, and the aggregation (within and between generations) is "prescriptive," where the latter is necessary to arrive at an overall social discount rate. We estimate and discuss two alternative aggregation rules that we consider based on the mean or median of a population, both of which are consistent with common social choice approaches.

We then empirically estimate UDRs for nearly all countries of the world with detailed demographic data from the World Health Organization. A striking feature of our results across countries is that they fall within the range economists generally employ and consider reasonable. Overall, the approach yields global estimates of the UDR at 2.1% and 1.3% for the mean and median aggregation, respectively. When comparing the results across countries, we emphasize the offsetting role of two demographic effects: age and life expectancy. A younger age structure and longer life expectancy tend to decrease a country's UDR, because more years to live will cause individuals to discount the future less. But, of course, a country's age distribution and life expectancies are closely related. Countries that are younger tend to have shorter life expectancies (i.e., many developing countries), whereas countries that are older tend to have longer life expectancies (i.e., many developed countries). The result, as we will show, is that countries with very different demographic profiles can have similar estimates of the UDR.

Our primary contribution is the demographically based approach for estimating UDRs, but we also place our empirical results in the context of deriving the overall social discount rate, often referred to as the consumption discount rate. Specifically, we combine our estimates of the UDR with estimates of the other parameters in the Ramsey

5. Note that mortality risk here differs from the risk of catastrophe at the population level, which has been discussed in the literature as it relates to the UDR (e.g., Dasgupta 2007; Weitzman 2009; Goulder and Williams 2012).

6. Yaari (1964) provides a seminal example, and recent papers use the same approach to explain patterns of economic growth and observed interest rates (e.g., Carvalho et al. 2016; Gagnon et al. 2016; Eggertsson et al. 2017).

equation. The analysis is primarily illustrative, and we use a combination of expert opinion and country-specific forecasts of economic growth. For example, based on Drupp et al.'s (2018) survey of experts on their views about consumption growth per capita (1.7%) and the elasticity of marginal utility (1.35), our estimates of the world's UDR imply social discount rates of 4.4% and 3.6% for the mean and median aggregation, respectively.[7] In what follows, we also estimate heterogeneous social discount rates for all countries based on country-level UDRs and forecasts of economic growth.

Our approach provides a new methodology for deriving regionally specific UDRs, with potential application to IAMs of climate change. We provide one such application using the Regional Integrated Climate Economy (RICE) model (Nordhaus and Yang 1996; Nordhaus 2010). We compare the results of two alternative calibrations: one where all 12 regions of the model have the same global estimate of the UDR, and one where each region has its own estimate of the UDR based on our methodology. We find that introducing UDR heterogeneity has little affect on the business-as-usual trajectory of emissions; however, it does meaningfully affect the efficient trajectory of emissions. We find that adding the UDR heterogeneity results in an efficient carbon tax that is 28% greater by the end of the century. Underlying the aggregate effects is a shift among countries such that those with lower UDRs reduce emissions more. Not only does a lower UDR impose greater concern for future climate damages, as is often noted in the literature, it also increases a country's emissions trajectory as a result of greater savings, capital accumulation, and output.

The remainder of the paper proceeds as follows. The next section provides a motivating background for the use of demographics as the basis for a social UDR. Section 2 develops our conceptual framework for deriving UDRs and contrasts it to previous approaches in the literature. Section 3 describes our data and reports the main estimates: country-specific and global results, along with an analysis of how demographic trends over decades affect the estimates. Section 4 places our results in the context of overall social (consumption) discount rates. Section 5 reports the results of our RICE analysis. Finally, section 6 concludes with a summary of our main results and a discussion of broader implications.

## 1. MOTIVATING BACKGROUND

We begin with the standard motivation for long-term discounting of social welfare in the Ramsey-Cass-Koopmans framework. There is an additively separable, utilitarian social welfare function of the form

$$W = \sum_{t=0}^{\infty} \left(\frac{1}{1+\delta}\right)^t U(C_t), \tag{1}$$

7. As we will discuss later, Drupp et al. (2018) also find that the average estimate of the UDR among economists is 1.1%.

where $\delta$ is the UDR, and $U(\cdot)$ is assumed to be time invariant. The convention in the literature is to assume a utility function that has a constant elasticity of the marginal utility of consumption, denoted $\eta$.[8] Then, on the optimal growth path, where consumption grows at a constant rate $g$, differentiating and rearranging (1) yields the well-known Ramsey equation:

$$r = \delta + \eta g, \tag{2}$$

where the long-term social discount rate $r$ is the sum of two terms. The first is the UDR, often interpreted in this context as reflecting the pure rate of time preference. The second is the product of the elasticity of marginal utility and the consumption growth rate. Notice that we have chosen to motivate the Ramsey equation such that all of the parameters in (2) are time invariant. This is the simplest setup and the one most commonly used, yet more general formulations allow for time-varying parameters and therefore a time-varying discount rate.[9] While our primary focus is on deriving a time-invariant estimate of the UDR, $\delta$, we will have more to say about time-varying discount rates, along with $\eta$ and $g$, later in the paper.

Calibration of the Ramsey equation is fundamental to IAMs of climate change, as well as to more general long-term discounting in benefit-cost analysis. As discussed previously, a simplified assessment of the literature is that there are two camps for choosing the UDR. One side interprets $\delta$ as an ethical parameter capturing concern for future generations. This perspective typically leads to values at or very close to zero (e.g., Ramsey 1928; Stern et al. 2007). The other side emphasizes that $r$ represents the real return on capital, and $\delta$ should be calibrated by matching real-world interest and savings rates (e.g., Nordhaus 2007). This perspective typically favors higher values of the UDR after accounting simultaneously for $\eta$ and $g$.

Our approach differs in that we seek to derive a mortality-based estimate of the UDR. Specifically, we consider the UDR that a social planner might choose under the assumption that long-term discounting is based solely on mortality risk, yet all generations, existing and future, are weighted equally. Note that we are not considering the risk to a population based on the chances of a catastrophe, which is sometimes invoked as a justification for using a nonzero UDR (e.g., Dasgupta 2007; Weitzman 2009; Goulder and Williams 2012). Instead, we consider individual risks based on mortality rates and life expectancy. In doing so, we exploit the fact that individuals of different ages are alive at any point in time and therefore face different mortality risks. Furthermore, as

8. Specifically, the assumed functional form of the utility function is $U(C_t) = C_t^{1-\eta}/(1-\eta)$, where $\eta$ in this setting is often interpreted as the degree of inequality aversion between generations.

9. See, e.g., Dasgupta et al. (1999), who discuss how changes in the consumption growth rate affect the discount rate over time. Gollier (2013) also provides numerous extensions to the Ramsey rule that provide not only for time variance but also for other concerns such as uncertainty and risk aversion.

time goes on, generations replace themselves, and the UDR can reflect an aggregation over how each generation cares about its own long-term future utility. We thus aim to provide an empirically based estimate that serves as a useful benchmark for evaluating normative judgments and highlights the potential importance of heterogeneity across populations (e.g., countries). While the approach abstracts from other possible motives that are likely to affect long-term discounting, some of which we discuss later in the paper, we argue that it provides a useful benchmark to a question in economics that generates much controversy with little direction on how the debate should be resolved.

The general idea of discounting based on mortality risk is not new. It has long been recognized that uncertainty about survival affects the way individuals discount the future and therefore make intertemporal choices. Frederick et al. (2002) provide a quotation from Rae (1834) that goes back almost a century before Ramsey: "When engaged in safe occupations, and living in healthy countries, men are much more apt to be frugal, than in unhealthy, or hazardous occupations, and in climates pernicious to human life. Sailors and soldiers are prodigals. In the West Indies, New Orleans, the East Indies, the expenditure of the inhabitants is profuse. The same people, coming to reside in the healthy parts of Europe, and not getting into the vortex of extravagant fashion, live economically. War and pestilence have always waste and luxury, among the other evils that follow in their train" (Rae 1834, 57). More recently, Fisher (1930) is well known for showing how survival probabilities are important to understand trade-offs between present and future consumption, and Samuelson (1958) uses mortality as a foundation in his seminal life-cycle model of lending markets. Eckstein (1961) describes a way that one can quantitatively infer a mortality-based UDR, writing, "The utility to be enjoyed at each future moment must be multiplied by the probability of being alive at that time, and since this probability falls with the remoteness of the period, a kind of pure discount factor emerges" (Eckstein 1961, 456). Eckstein (1961) goes on to describe that while this approach assumes individuals place no bequest value on wealth they leave behind upon death, it does suggest a pure rate of time preference that can be used for comparative purposes across countries where mortality rates may differ substantially. This is precisely the idea that underlies our approach in this paper.[10]

## 2. CONCEPTUAL FRAMEWORK

Consistent with the motivating background and the Ramsey equation in (2), we consider a social planner that seeks a constant UDR for use in long-term discounting.[11] We

10. Beyond the general idea, however, our approach differs from what Eckstein (1961) actually does in his application to the United States and India. We describe how the approaches differ after establishing our framework.

11. Even though some argue in favor of a declining social (consumption) discount rate $r$ over time, this does not create an inconsistency with a constant UDR over time. See, e.g., the discussion in Arrow et al. (2014).

assume further that the planner places equal weight on all current and future generations, including those not yet born. In this section, we show how demographic information can be used to derive an estimate of the social UDR, and despite an equal weighting across current and future generations, the estimate is greater than zero, reflecting forward-looking views about mortality risk that apply equally across generations.

Consider a population of individuals alive at time $t = 0$ that is partitioned into age cohorts indexed by $K$. For example, all individuals in cohort $K = 10$ are 10 years old and so forth for all ages $K = 0, 1, \dots, \bar{K}$, where $\bar{K}$ is the maximum age in the population. The distribution of cohorts is given by a normalized vector $\phi$ that identifies the proportion of individuals at each age, denoted $\phi_K$, where $\Sigma_{K=0}^{\bar{K}} \phi_K = 1$. As individuals age through time, they advance from one cohort to the next.[12] We assume a stable distribution among cohorts over time so that $\phi$ remains time invariant, though we address the impact of this assumption later in the paper.

Two further sources of data are used to characterize a cohort. The first is life expectancy at age, $L_K$, which represents the expected number of years a member of cohort $K$ has left to live. The second is the age-specific mortality probability, $\gamma_K$, which represents the probability of death before reaching age $K + 1$ for an individual in cohort $K$. Note that the assumption of a stable age distribution means that as time advances and individuals age, their future life expectancies and mortality probabilities will satisfy $L_K(t) = L_{K+t}$ and $\gamma_K(t) = \gamma_{K+t}$. That is, the life expectancies and mortality probabilities at each age are constant through time, and aging is equivalent to moving from one cohort to the next.

## 2.1. A Representative Individual of Age $K$

The basic idea underlying mortality-based discounting is to weight future utility by the probability of survival, and hence age-specific mortality rates imply a demographically based discount factor. Let $\beta_K(t)$ denote the discount factor that an individual in cohort $K$ would apply $t$ years into the future, so that

$$\beta_K(t) = \prod_{\tau=0}^{t} (1 - \gamma_K(\tau)) = \prod_{\tau=0}^{t} s_K(\tau), \tag{3}$$

where $s_K(t)$ denotes, in more compact notation, the age-specific survival probability after $t$ years for an individual initially in cohort $K$. It is useful to note that as long as $t < L_K$, the discount factor in (3) will tend to decrease with $t$ for two reasons. The first is the standard compounding, which would occur even if the individual faced a constant mortality risk across periods. The second is demographic, owing to rising mortality risk (i.e.,

12. The cohort partition can further delineate gender-by-age groups or any other classification for which sufficient demographic information is available. We develop the theoretical model based on an age delineation only, but a generalization is straightforward and only requires additional indexing. In our empirical application, however, we will consider gender-by-age cohorts.

decreasing survival probabilities) as the individual ages. Clearly, a direct application of (3) poses limits for long-term discounting from an individual's perspective. As the time horizon lengthens, an individual's survival probability will eventually equal zero, so that $\beta_K(t)$ will equal zero for sufficiently large $t$. This would imply an infinite long-term discount rate, as the individual does not care at all about future utility after death.

The problem is nevertheless different from the social planner's perspective. Because the planner weights current and future generations equally, an individual of age $K$ is taken to represent oneself, along with future individuals of the same age. The aim is therefore to consider a long-term discount factor that can be applied not only for individuals alive now, but also for future individuals that will eventually be of the same age, including those not yet born. We therefore consider the discount factor that applies in the most distant period that an individual expects to live. Specifically, for an individual in cohort $K$, we set $t = L_K$ and use equation (3) to derive $\beta_K(L_K)$. This represents the discount factor that an individual of age $K$ would apply $L_K$ years into the future from the perspective of period $t = 0$. Note that this discount factor is a multiplicative function of all future survival probabilities that the individual faces from age $K$ to age $K + L_K$.

It is useful to recognize that the discount factor $\beta_K(L_K)$ for all $K$ can be translated into a constant period-by-period discount parameter that yields the same discount factor in $L_K$ years:

$$\hat{\beta}_K = \left[\prod_{\tau=0}^{L_K} s_K(\tau)\right]^{\frac{1}{L_K}}. \tag{4}$$

This expression is simply the geometric mean of the survival probabilities over the expected life years remaining for an individual in cohort $K$. The key feature of equation (4) is that $\hat{\beta}_K^{L_K} = \beta_K(L_K)$, which follows directly from equation (3). This means that the mortality-based discount factor out to one's life expectancy is equivalent to applying the constant period-by-period discount parameter defined in equation (4) for $L_K$ periods.[13] It follows immediately that equation (3) further defines cohort-specific UDRs, denoted $\delta_K$ for all $K$. In discrete time, the UDR is defined by $\hat{\beta}_K = (1 + \delta_K)^{-1}$, implying that $\delta_K = \hat{\beta}_K^{-1} - 1$. Assuming mortality risks are increasing as individuals age, those in older cohorts will face a future with lower survival probabilities and therefore higher UDRs. The simple intuition is that when older individuals look into the future they will tend to have fewer expected years left to live, and those years are the ones with relatively high mortality rates.

13. It is important to underscore what is not being asserted here. Application of the discount factors $\hat{\beta}_K^t$ and $\beta_K(t)$ are not equivalent for all $t$. In general, they will imply quite different discount factors for all $t < L_K$, with the former being lower than the latter because of high survival rates at younger ages. The equivalence occurs precisely at $t = L_K$, which is of particular interest here because that is the number of years into the future that an individual of age $K$ expects to live.

Before turning next to aggregation into an overall social UDR, we remark on how the discount parameter in (4) accords with two desirable characteristics from a social planner's perspective. The first is that estimation of $\hat{\beta}_K$ is a forward-looking process. For each individual, and at each age, the planner looks as far into the future as the individual is expected to live, taking account of the full stream of expected survival probabilities. The estimate of $\hat{\beta}_K$ is therefore a summary statistic of the mortality-based discounting process that applies to individuals from each age $K$ out to the end of their life expectancy. In other words, based on survival probabilities, it identifies an exponential discounting process that is consistent with what an individual would prefer looking out to the expected end of her life from the current period. The second desirable characteristic is time invariance due to the assumption of a stable age distribution. If survival probabilities and life expectancies remain constant, the estimate of $\hat{\beta}_K$ will remain constant through time for individuals at age $K$. With each passing year, the discounting process for an individual advances from $\hat{\beta}_K$ to $\hat{\beta}_{K+1}$, but the estimate for each age does not vary with time and applies equally to those currently alive and those not yet born.

## 2.2. Social Aggregation

Having established how to characterize mortality-based UDRs at the individual level, we now consider the social planner's problem of aggregating across individuals to arrive at a long-term, population-level, social UDR. The question of aggregating discount rates has received considerable attention in the literature. The seminal approach of Weitzman (1998, 2001) takes the expectation of heterogeneous discount factors through time, which results in discount rates that decline into the future and converge to the lowest of the rates being considered. One concern with this approach is that it conflates heterogeneity in the UDR with other components of the overall consumption discount rate in the Ramsey equation (Heal and Millner 2014b; Freeman and Groom 2015). Other studies focus more directly on the UDR. Gollier and Zeckhauser (2005) and Jouini et al. (2010) consider an exchange economy where intertemporal efficiency is shown to imply a declining UDR over time for a representative agent, based on a weighted mean across individuals. Heal and Millner (2014a) adapt a similar approach in the context of evaluating climate policy, and Millner and Heal (2018) compare the two approaches of utilitarian aggregation and majoritarian voting.[14]

Here we employ two aggregation rules with a social choice basis rather than a focus on efficiency. Specifically, we consider the mean and median of equation (4) across individuals in the population. Our use of these aggregates helps to maintain our focus on the demographic basis for UDRs, which can ultimately support any number of aggregation

14. In another paper, Millner (2018) considers a "nondogmatic" approach to resolving differences of opinion among planners about the long-run social discount rate. The insight is that if all planners are uncertain about their own views, all will agree on the same value of the social discount rate to approve in the long run.

rules.[15] The unweighted arithmetic mean and median among individuals of a population are simple, transparent, and make for useful comparisons. One justification for the mean is that it represents the demographic information of the representative individual in the population, whereas the median is consistent with the democratic principle of majoritarian voting, as considered by Millner and Heal (2018).[16]

In both cases, we first aggregate over the discount parameters in equation (4) such that

$$\beta^{\text{mean}} = \sum_{K=0}^{\bar{K}} \phi_K \hat{\beta}_K, \tag{5}$$

and

$$\beta^{\text{median}} = \left\{ \hat{\beta}_{\tilde{K}} : \sum_{\hat{\beta}_K < \hat{\beta}_{\tilde{K}}} \phi_K < \frac{1}{2} \text{ and } \sum_{\hat{\beta}_K > \hat{\beta}_{\tilde{K}}} \phi_K \leq \frac{1}{2} \right\}, \tag{6}$$

where the two equations are the cohort-weighted mean and median, respectively. We then derive the corresponding mean and median UDRs with the identity $\delta = \beta^{-1} - 1$. Notice that the order of operations—that is, aggregating over the estimates in (4) and deriving the UDR versus deriving the UDR for each cohort and then aggregating—has no effect on the median because the conversion from the discount parameter $\beta$ to the discount rate $\delta$ is a monotonic transformation. While the order does affect the mean UDR, we first aggregate over the $\hat{\beta}_K s$ because they directly represent the summary statistic of future survival probabilities and are more equally representative for each individual. It can be shown that aggregating over UDRs with the arithmetic mean is equivalent to taking the geometric (vs. arithmetic) mean of the $\hat{\beta}_K s$.[17]

15. Nevertheless, implementation of more complex weighting mechanisms that account for intertemporal efficiency, such as that in Gollier and Zeckhauser (2005), would require a broader modeling framework that accounts for optimal saving behavior and intergenerational risk sharing and smoothing.

16. A recent contribution by Emmerling et al. (2017) emphasizes how a focus on the median individual is also useful for taking account of inequality in the derivation of social discount rates.

17. To see this, consider the weighted average of the cohort-specific UDRs, $\sum_K \phi_K \delta_K$. Letting $N$ denote the total number of individuals in the population and $n_K$ the number of individuals aged $K$, the weighted average can be rewritten as

$$\frac{1}{N}\sum_K n_K \delta_K = -\ln\left[\prod_K (\hat{\beta}_K)^{n_K}\right]^{\frac{1}{N}},$$

where the equality follows because $\delta = -\ln(\beta)$ in continuous time. Notice that the right-hand side is simply the natural log of the geometric mean of $\hat{\beta}_K$ across all individuals. Hence, because the arithmetic mean of positive numbers is always greater than the geometric mean, the UDR that follows from eq. (5) will always be less than the alternative approach of taking the weighted average over UDRs.

The aggregations in equations (5) and (6) implicitly define the notion of long term, with implications for interpretation of the corresponding UDRs. For each cohort, $\hat{\beta}_K$ is defined based on the survival probabilities over all years of the remaining life expectancy. This means that for every individual the long term must begin further in the future than one's own remaining life expectancy. Implicit in our analysis therefore is a notion of the long term that begins at least $T$ years into the future, where we define $T$ such that it satisfies $T \geq \max L_K$. That is, the long term is defined as beginning after the longest life expectancy and continuing indefinitely thereafter, and the estimate of the social UDR, with either aggregation, does not change over time, because of the assumption of a stable population distribution. The set of weights in $\phi$ do not change, and nor does the estimate of $\hat{\beta}_K$ for all $K$.[18] The overall results are estimates of a constant and long-term social UDR that is (i) forward looking to the expected end of life and demographically based for all individuals and (ii) treats all individuals equally, including those of all ages currently living and those not yet born.

### 2.3. Relation to Other Approaches

Throughout the remainder of the paper we will focus primarily on the mean aggregation and make direct comparisons to the median. The mean aggregation in equation (5) implies weights among heterogeneous estimates of the discount parameter. These population-based weights reflect a normative stance to treat all individuals equally, but other weighting schemes are possible. For example, Weitzman's (1998, 2001) approach to deriving the long-term discount rate would effectively count only $\max \hat{\beta}_K$ and therefore yield a long-term UDR of $\delta = [\max \hat{\beta}_K]^{-1} - 1$, which is the lowest UDR among all cohorts. Nevertheless, it is worth clarifying that the parameter $\hat{\beta}_K$, defined here, does not give rise to the same future discount factor as in Weitzman's approach. Ours, in contrast, is based on a forward-looking process reflecting survival probabilities, applies only in the period of life expectancy for an individual, and changes from the starting point of different ages.

As mentioned previously, our approach is most closely related to Eckstein's (1961) estimates of the social UDR for select countries based on the aggregation of individual mortality rates. However, his approach, along with a subsequent application by Kula (1985), differs from ours primarily in the way individual discount parameters are computed. Specifically, they estimate a representative individual's UDR as the geometric mean of expected mortality rates over the life cycle, rather than the geometric mean of the implied discount factor, that is, the survival probability. In our view, there is no consistent basis for taking the geometric mean of the mortality rates, whereas there is one for the discount factor, as it can well approximate the survival probability at life expectancy.

A mortality rate in any period will also have a different interpretation than a discount rate, yet Eckstein (1961) and Kula (1985) assume that the discount rate at age, $\delta_K(t)$, is

18. Note that technically the stable age distribution requires that the weights remain the same across cohorts, along with the survival probabilities at each age, but the total number of individuals in the population need not.

equal to the mortality rate, $\gamma_K(t)$, which is also inconsistent with their premise that survival rates, $s_K(t) = 1 - \gamma_K(t)$, provide a natural discount factor, $\beta_K(t)$. To see this, the following identities must hold by definition $\delta = (1/\beta) - 1 = [1/(1 - \gamma)] - 1 = \gamma/(1 - \gamma)$, and this implies that $\gamma < \delta$ rather than $\gamma = \delta$. This explains, in part, why empirical implementation of our approach results in slightly greater estimates of the UDR for the select countries that Eckstein (1961) and Kula (1985) consider (the United States, India, and the United Kingdom, respectively) when using the same data.[19] Our approach of using survival probabilities to define discount factors (rather than mortality rates to define discount rates) is consistent with that taken in the recent macroeconomics literature that seeks to account for survival probabilities in life-cycle models (Carvalho et al. 2016; Gagnon et al. 2016; Eggertsson et al. 2017).

## 3. ESTIMATION OF SOCIAL UTILITY DISCOUNT RATES

We now turn to empirical estimation of the mortality-based UDRs for nearly all countries of the world and for the world as a whole. We obtain demographic data from two sources. The first is country-specific life tables from the Global Health Observatory data repository of the World Health Organization (WHO).[20] These tables report age-specific mortality rates by gender and binned in the following age classes: < 1 year, 1–4 years, 5–9 years, . . . , 90–94 years, 95–99 years, and 100+ years. They also include life expectancies (i.e., the expectation of additional years to live) for each bin and gender. The life tables are available for 194 countries. We use the 2012 tables for our primary estimates, but also use tables for 1990 and 2000 to examine how the results have changed over time.

The second source of data is country-specific population estimates from the United Nations Population Division.[21] These include the number of people in each country by 5-year age groups and by gender. The estimates are available every 5 years. We associate these data for the year 2010 with the 2012 WHO life tables. When conducting the same analysis back to 1990 and 2000, we associate the UN and WHO data for that same year.[22]

---

19. Note that while $\gamma \approx \delta$ for very small $\gamma$, the condition does not hold at greater ages where mortality rates begin to climb. Furthermore, the difference affects the estimates of the period-by-period discount factor for all ages (i.e., eq. [4] here and the analog for Eckstein [1961] and Kula [1985]) because the geometric mean is taken from the current age out to life expectancy.

20. We obtained the data with queries through the web service, Athena.Links. Instructions on how to make data queries are available at http://apps.who.int/gho/data/node.resources.api.

21. The data are available at https://esa.un.org/unpd/wpp/Download/Standard/Population/.

22. When merging the two data sets, one adjustment is needed in the first age bin. The WHO data on mortality rates and life expectancy disaggregates the first age bin (0–4 years) into two, providing data for individuals <1 year old and individuals 1–4 years old. In the UN data, however, the population statistics are for the whole bin of 0–4 years old. We can nevertheless decompose the population statistics in the UN data using the mortality rates in the WHO data. After doing so, our analysis is based on bins from <1, 1–4, and 5-year increments all the way up to 100+.

These two sources of data are sufficient to estimate a population-level, social UDR as described in the previous section for nearly all countries of the world. To see the precise steps involved, first consider how to derive the discount parameter for an individual of a particular age and gender within a specific country. Take the example of a 52-year-old female in the United States. Based on the 2012 WHO life table for the United States, she has a remaining life expectancy of 33.15 years; that is, she has a full life expectancy of 85.15 years conditional on having reached age 52. For each of her expected remaining years, we also have an estimate of her survival probability from the WHO data, although the estimate is the same for each year within an age bin. We can therefore take the geometric mean as described in equation (4) with the following numerical values:

$$\begin{aligned}\hat{\beta}^{f}_{52} &= [(0.994)^3(0.991)^5(0.987)^5(0.987)^5(0.980)^5(0.968)^5(0.946)^5(0.871)^{.15}]^{\frac{1}{33.15}}\\ &= 0.978,\end{aligned}$$

where $f$ denotes female, and the numbers in parentheses are one minus the mortality rate for the corresponding year. Note that the last year is raised to the power 0.15, reflecting how the remaining life expectancy is 33.15 years. Our procedure for each individual in a population follows the same steps for the median age in each bin and by gender. We use the median age for a bin and the bin's remaining life expectancy because we do not know the distribution of ages within each bin. This produces an estimate of $\hat{\beta}_K$ for each age group and gender. For example, the illustrative case of a woman at age 52 would apply to all women in the age bin from 50 to 54, because 52 in the median age of the bin.[23]

Population estimates of $\beta^{\text{mean}}$ and $\beta^{\text{median}}$ are then derived using the aggregation rules in either equation (5) or (6), with weights based on the UN population data. For each aggregation, the corresponding population-level, social UDR is based on the simple conversion to a rate by $\delta = \beta^{-1} - 1$.

## 3.1. Country-Specific Results

We estimate social UDRs for 182 countries, all of those for which complete data are available. We report the mean and median results in table A1 for the most recent year, 2012. We report the results for 1990 and 2000 in table S1 (available online), and make explicit comparisons with the 1990 results below. We begin our analysis with a focus on the mean aggregation, before making comparisons to the median.

Figure 1*A* shows the distribution of the mean estimates across countries that range from 1.44 (Cambodia) to 3.50 (Bulgaria).[24] The distribution is left skewed, and the

23. For the first two and last categories we do not use the median, but rather ages 0, 2, and 100, respectively.

24. Countries with the 10 lowest estimates, in ascending order, are Cambodia, Guatemala, Rwanda, Honduras, Gambia, Kenya, Ethiopia, Tanzania, Syria, and Qatar. Countries with the 10 highest estimates, in descending order, are Bulgaria, Ukraine, Latvia, Russia, Serbia, Belarus, Hungary, Romania, Croatia, and Lithuania.

A

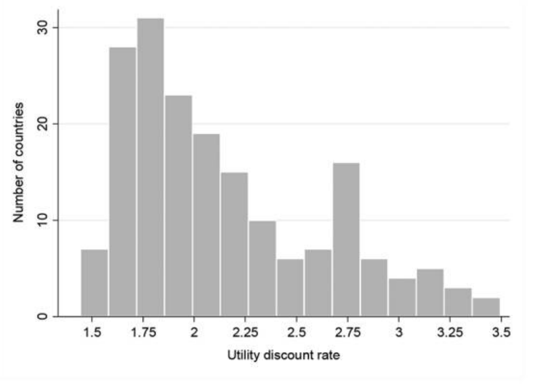


B

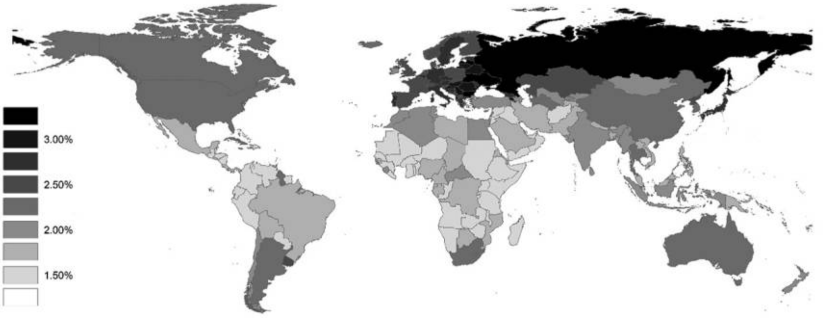


Figure 1. Histogram (*A*) and map (*B*) illustrating the distribution of the 2012 estimates of the utility discount rates across countries, using the mean aggregation. Color version available as an online enhancement.

median across the mean estimates is 2.01%. Figure 1*B* illustrates the geographic heterogeneity. There is a clear pattern of higher rates in the northern latitudes, especially in the former Soviet Union countries. The lower rates tend to be clustered in Africa, Central America, and the Middle East. Underlying these results are two offsetting yet related effects. The first is that countries with longer life expectancies will have lower UDRs, all else equal. This follows because a longer life expectancy means more years over which individuals experience high survival rates. Empirically, these same countries also tend to have older populations, and this is the second effect that pushes in the opposite direction. Older populations will have higher rates, all else equal, because there are fewer years over which to live, and these years include those with low survival rates.

It follows that two countries can have the same mortality-based UDR for very different reasons. The United States and Lesotho provide an example. We find that both countries have mean UDRs of 2.4%, yet they differ substantially in their demographics.

Life expectancies at birth in the two countries are 78.6 and 48.8 years, respectively, while the median ages are 37.2 and 20.1. The US population is older and tends to live longer, and these differences offset each other in comparison to Lesotho's younger population with a shorter life expectancy at birth or any given age.

Figure 2 illustrates the relationship between life expectancy, median age, and our estimates of the UDR across countries. As one would expect, median age increases with life expectancy. Longer life expectancy tends to decrease the UDR, holding median age constant, while an older median age tends to increase the UDR, holding life expectancy constant. As both are reduced (or increased), the figure shows how the UDRs change less, and something close to level sets within the figure's UDR bins emerge. Note, however, that median age is only one possible indicator of a population's age distribution, and the estimated UDRs are a function of life expectancies at each age, rather than at birth. Nevertheless, the figure illustrates how a population's demographic structure affects the estimate. In particular, it explains why differences do not necessarily align with a country's level of development. For the same UDR, one country may be developed with a longer life expectancy and an older population, while another may be developing with a shorter life expectancy and younger population. A second example of this phenomenon is Australia and South Africa with UDRs of 2.2%, and respective life expectancies at birth of 81.8 and 59.3, and median ages of 37.0 and 24.3.

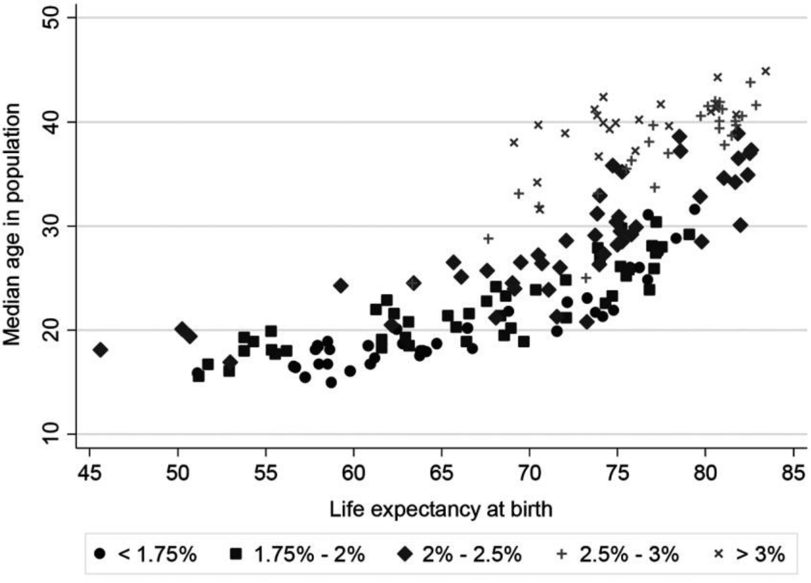


Figure 2. The relationship between life expectancy, median age, and the 2012 estimates of the utility discount rates across countries (bins denoted in the legend), using the mean aggregation. Color version available as an online enhancement.

We have thus far focused on the mean rate for the population aggregation rule in equation (5), but we also consider the median rate for a population as described in equation (6). To show how the estimates compare, figure 3*A* plots the median estimate against the mean for all countries. In all cases the median is less than the mean, reflecting the general leftward skew of population distributions toward more younger people. Hence we find that using the median aggregation rule uniformly provides lower estimates of the mortality-based UDR. The percentage point difference is greater among countries with higher estimates of the mean UDR. Figure 3*B* shows the distribution across countries of the percentage change in the estimate when converting from the mean to median aggregation rule. Many of the smaller declines occur in the Middle East,

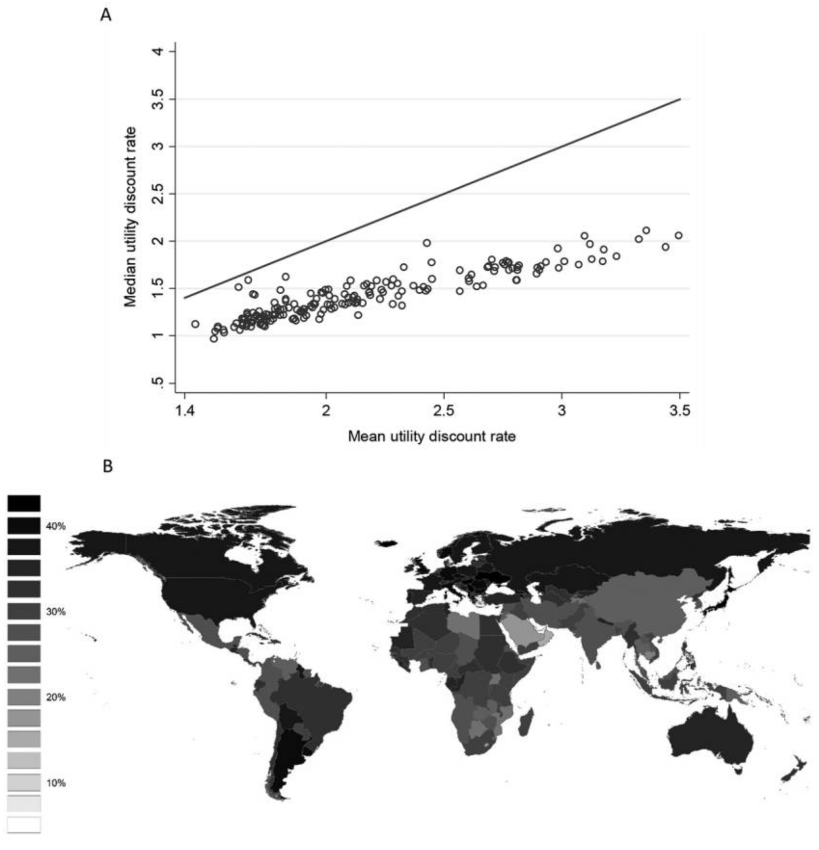


Figure 3. *A*, Scatter plot of each country's median estimate of the 2012 mortality-based utility discount rates against its mean, along with the 45-degree line. *B*, Geographical illustration of the percentage decrease when moving to the median from the mean estimate for each country, that is, $-100[\tilde{\delta} - \bar{\delta}]/\bar{\delta}$. Color version available as an online enhancement.

whereas many of the larger declines occur in Eastern Europe.[25] The percentage decline tends to be greater in countries where the population is older and life expectancy is shorter. On average across countries, the UDR decreases 33%, and for 95% of the countries the decrease is greater than 25%. These differences illustrate how the aggregation rule can have important consequences, and we return to this topic later in our discussion.

### 3.2. Global Results

We now turn to a global estimate of the mortality-based UDR. We follow the same procedure but now treat the world as a single population using the world data from the WHO and UN data sets. Specifically, we combine the demographic information for all countries into a single distribution upon which we derive an estimate of the world UDR. Using the mean aggregation, we estimate a global UDR of 2.13%. Using the median aggregation, we estimate a rate of 1.33%. Interestingly, these estimates are surprisingly close to those employed in the literature when calibrating discount rates to the Ramsey rule with applications to climate change. As discussed previously, some use rates at or near zero (e.g., Cline et al. 1992; Arrow 1999; Stern et al. 2007), but others employ rates between 1% and 3% (e.g., Nordhaus and Yang 1996; Nordhaus 2007, 2008; Weitzman 2007). Moreover, a recent survey of expert opinion among economists finds a mean estimate of 1.1% for the UDR, with a standard deviation of 1.47 and a range between 0% and 8% (Drupp et al. 2018).

While these estimates are consistent with treating the world as a single population, we note that other aggregation rules may be considered reasonable among countries. One could, for example, use the median across countries of the mean or median country-specific estimates, yielding 2.01% and 1.37%, respectively. Alternatively, one could take some weighted average of the country-specific estimates, based on population or some other normative criteria. We find, for example, that a country populated weighted mean of the mean and median aggregation yields estimates of 2.25% and 1.53%, respectively.

As mentioned previously, our aim here is not to advocate one aggregation rule over another, either between countries or within countries. Instead, we employ common aggregation rules (means and medians) as benchmarks to compare against estimates currently used in the literature. In doing so, our finding is that the choice of the UDR used for calibration of the Ramsey rule—within several IAMs for climate change and more generally—is very much in line with the estimates here based on demographic mortality risks.

### 3.3. Changes in Population Age Structures over Time

Our estimates of the social UDR are based on a population's existing age structure and life expectancies at each age. The aim is to develop a procedure for estimating long-term

25. Countries with the 10 lowest percentage declines, in ascending order, are United Arab Emirates, Qatar, Bahrain, Kuwait, Oman, Saudi Arabia, Lesotho, Cambodia, Libya, and Brunei Darussalam. Countries with the 10 highest percentage declines, in descending order, are Hungary, Ukraine, Japan, Albania, Iceland, Serbia, Armenia, Jamaica, Uruguay, and Czech Republic.

UDRs that can serve as useful benchmarks. As discussed previously, if the population age structure is constant over time, our procedure generates estimates that are time invariant. It is nevertheless reasonable to question how the estimates might change over time. We consider this question on both conceptual and empirical levels.

The notion of a stable age distribution is a foundational starting point in demography (Lotka 1922), yet we know that population age structures have undergone long-term shifts. Preston et al. (1989) show how the mean age of a population changes over time based on birth and death rates, the mean age of people living and dying, and age-specific growth rates. Our setting is similar in that the objects of interest (i.e., eqs. [5] and [6]) are functions of the age distribution rather than the age distribution itself.[26] One implication is that the results are less sensitive to demographic shifts than one might expect. Consider an increase in life expectancy due to improved or expanded medical interventions. This will have relatively little effect on the $\hat{\beta}_K$ for young individuals, because their estimate of this parameter accounts for survival probabilities over many years, and not just those most affected near the end of life. It will, however, have a greater effect on older individuals, but there are relatively fewer of them in the population. Another important demographic variable is infant mortality, but in many countries the magnitude of future gains will be relatively small because significant reductions have already occurred.[27] In other places where significant gains remain possible, they are likely to be offset to some degree by changes in fertility (Preston 1978). This leaves more general changes in fertility, immigration, and mortality (due perhaps to war or disease) as variables of potential importance, yet the empirical evidence on changes over time when data are available indicate relatively small differences.

Following the same methodology described previously, we estimate decadal UDRs for the United States over a century from 1910 to 2010. Data on mortality rates and life expectancies are from the US Social Security Administration, and data on the population distribution by age and gender are from the US Census Bureau. Because of reporting differences over time on the age bins of the population distribution, we aggregate data to the finest resolution that is common across the full time series.[28] Figure 4 shows

26. Specifically, while Preston et al. (1989) examine the mean age as a summary statistic for aging, we examine the mean and median of a function for each individual that is itself a function of population parameters.

27. Global child mortality fell from 32.1% in 1920, to 18.2% in 1960, and 4.3% in 2015. No country today has child mortality above 10%, whereas in 1970 some (e.g., Yemen) had mortality rates above 20% (Roser 2019).

28. This requires the inclusion of all individuals greater than 74 years of age into a 75+ top age bin. While this means that the resulting estimates are comparable across time, it also implies a downward bias in the estimate of the UDRs. This follows because individuals older than 75 will be ascribed the higher survival probability of a 75-year-old, which means a higher $\hat{\beta}_K$ and therefore a lower UDR. This explains the difference between these estimates and those for the United States using the UN and WHO data sets. We provide further details on the data sources and inconsistencies over time in the supplementary material (available online).

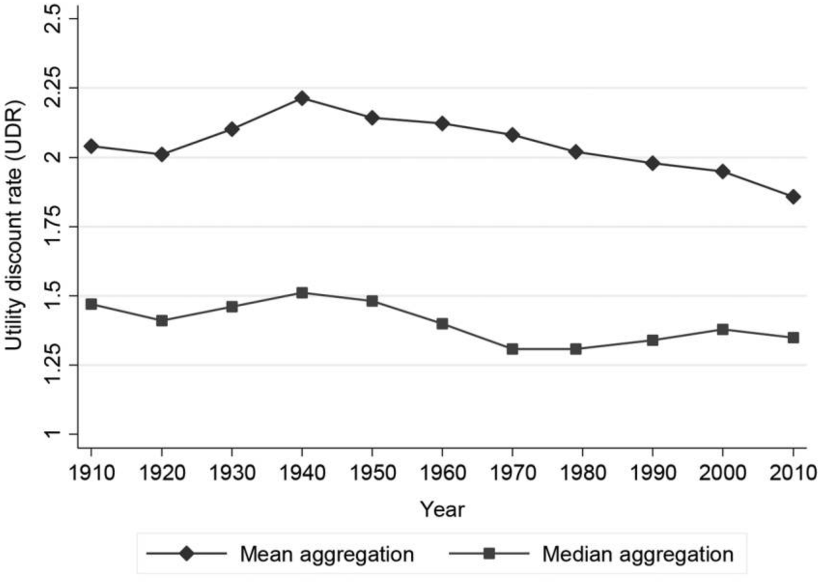


Figure 4. Trend in the mortality-based utility discount rates for the United States over a century using data from the US Social Security Administration and the US Census, showing the mean and median aggregation rules. Color version available as an online enhancement.

the full time series of the UDRs based on the mean and median aggregations, which have mean estimates of 2.0 and 1.4. The key observation is that the estimates only change a modest amount, reaching a peak in the 1940s, after which the increase in fertility associated with the baby-boom generation brings down the population UDRs. Overall, the mean estimate ranges between 1.9 and 2.2, and the median estimate between 1.3 and 1.5. We believe it is reasonable to interpret these estimates as quite stable, given the much broader range of estimates regularly employed in the literature.

For all other countries, including the United States, we examine stability of the estimates over time using the UN and WHO data for the years 1990 and 2000. The year 1990 is the earliest that comparable data is available for the broad sample of countries. We report the results for all countries, both years, and aggregation rules in the table S1. Here we focus discussion on the largest time span from 1990 to 2012. Figure 5 plots the 2012 estimate against the 1990 estimate for all countries, showing the mean and median aggregation separately. The figure shows how the estimates both increase and decrease across countries with nearly the same variance across the range of 1990 UDRs. Over the 22-year period, the mean change across countries for the mean aggregation is an increase of 2.4% with a standard deviation of 11.1.[29] For the median aggregation, the mean

29. Note that, in contrast to the figure, this is a percentage change and not a percentage point change.

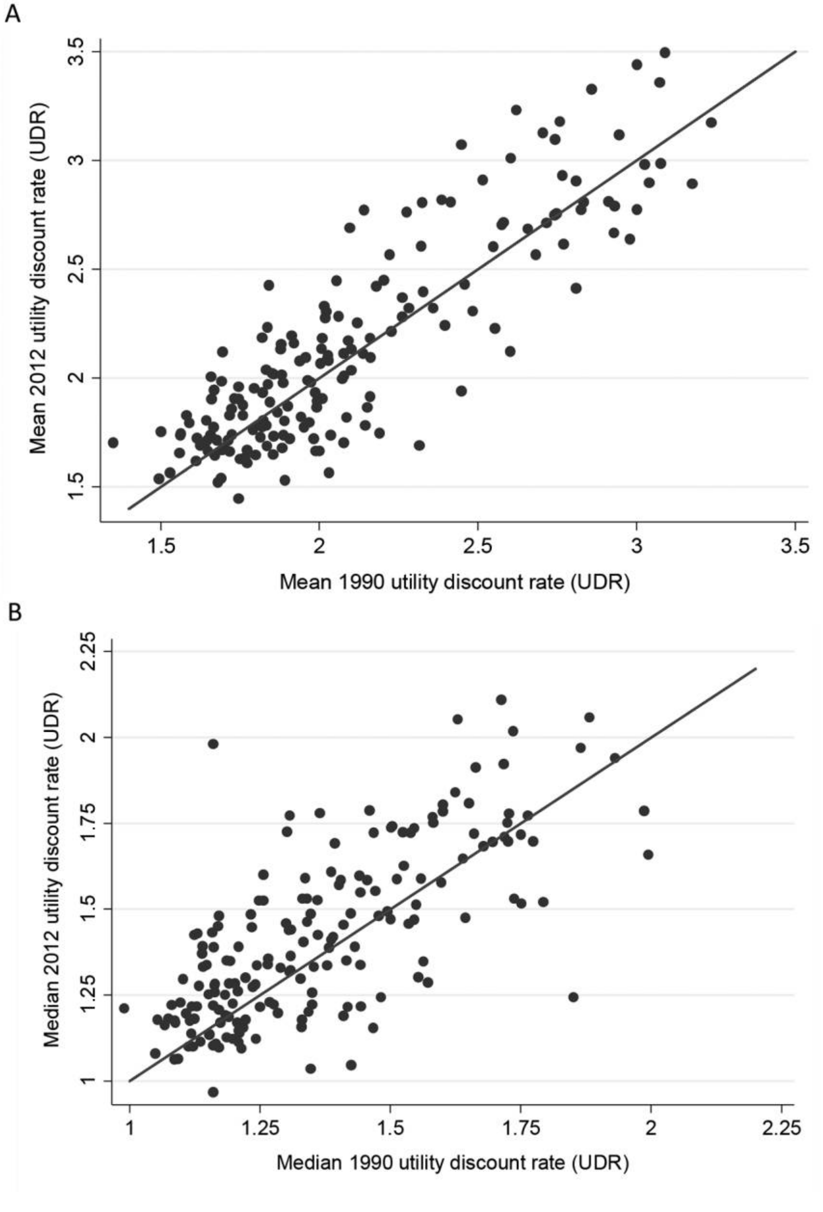


Figure 5. Scatter plot of the 2012 mortality-based utility discount rates against the 1990 estimates for all countries, showing the mean aggregation rule (*A*) and the median (*B*). Color version available as an online enhancement.

difference is an increase of 4.2% with a standard deviation of 12.8. While we interpret these results as showing a reasonable degree of stability for long-term discounting, the differences across countries are a function of how demographics have changed: a shift toward a younger (older) population with longer (shorter) life expectancies will decrease (increase) the estimated UDR.[30] Other factors that may influence the differences are changes in data availability within countries and adjustments to the methods of estimating parameters of life tables.[31]

When considering global estimates of the UDR, there is relatively little change over time. For all three estimates in 1990, 2000, and 2012, the mean aggregation remains nearly constant at 2.04%, 2.10%, and 2.13%, respectively. The same holds for the median aggregation with UDRs of 1.27%, 1.39%, and 1.33%. The stability of these estimates builds further support for using the demographic characteristics of the world's existing population to inform long-term discount rates.

## 4. RELATION TO THE SOCIAL (CONSUMPTION) DISCOUNT RATE

The primary focus of our analysis is on deriving demographically based estimates of the UDR. But this is only one component of the overall social (consumption) discount rate, defined in equation (2), that is used for economic evaluation of long-term projects. In this section, we derive overall social discount rates using our estimates of the UDR as an input.[32] The intent is to place our results in context for exploring heterogeneous social discount rates across countries. In doing so, we combine estimates of the different parameters in the Ramsey equation to produce an estimate of the overall social discount rate. In the next section, however, we show how our estimates of the UDR can be used in a calibration exercise to match observed interest rates.

Because of the importance of using estimates of the UDR as the basis for deriving a consumption discount rate in (2), we first remark on the conceptual basis of building up the Ramsey equation from its component parts. While most empirical studies build

30. Countries that make the "bottom" 10 for a decrease in the UDR for both the mean and median aggregation are Zambia, Ethiopia, Antigua and Barbuda, Liberia, Rwanda, Sierra Leone, Angola, Eritrea, Cambodia, and Uganda. With respect to the "top" 10 for an increase in the UDR, there is greater heterogeneity across the aggregation rules. The top 10 for the mean aggregation are Lesotho, Montenegro, North Korea, Nicaragua, Armenia, Algeria, Serbia, Uzbekistan, Philippines, and Macedonia. Those for the median aggregation are Lesotho, China, Thailand, Montenegro, Cuba, Algeria, Saudi Arabia, Morocco, Lithuania, and Swaziland.

31. See, for example, a description of the WHO methods and data sources for generating life tables 1990–2015 (World Health Organization 2016).

32. We are of course following convention here and throughout the paper by assuming a constant social discount rate. For recent discussions about the potential use of declining discount rates over time, see Arrow et al. (2012) and Cropper et al. (2014).

toward a consumption discount rate by separately identifying its constituent parameters (e.g., Nordhaus 2007; Stern et al. 2007), the theoretical basis for doing so is more subtle. Jouini et al. (2010) derive two conditions under which the Ramsey discount rate can be constructed based on the sum of an aggregate UDR and beliefs about future consumption. The first condition is that individual UDRs are certain. The second is that the elasticity of the marginal utility of consumption, $\eta$, is independent of the UDR. We find no reason why these assumptions cannot reasonably apply in our setting. In effect, we assume that individuals are certain about their expected mortality, and because the UDRs depend solely on demographic factors, we find little cause for concern about potential correlation with $\eta$. Furthermore, we are unaware of previous applied studies that employ discount rates with correlated values for $\delta$ and $\eta$.

We use the following procedure to obtain estimates of the other parameters in the Ramsey equation. The elasticity of the marginal utility of consumption $\eta$ that we use is the mean value of 1.35 from an expert opinion survey conducted by Drupp et al. (2018).[33] For the growth rate of per capita real consumption $g$, we use forecasts about real gross domestic product (GDP) per capita. Although consumption itself is the correct measure rather than GDP, such data are only available for a small set of countries. Following Gollier (2013, 2015), we use the International Macroeconomic Data Set constructed by the Economic Research Service of the US Department of Agriculture.[34] The data set provides historical and projected real GDP per capita for 189 countries that account for more than 99% of the world economy. For each country and the world as a whole, we take the arithmetic average of growth rates for all years 2013 through 2030. This gives us an estimate of each country's future annual growth rate based on the series of annual estimates.[35] We then use these estimates along with those for the UDRs and $\eta$ to examine heterogeneity in the overall social discount rates across countries. We use the mean aggregation of the UDR for illustrative purposes and to be consistent with the implicit averaging in GDP per capita.[36]

Figure 6*A* shows the distribution of the most recent estimates across countries. We also report the specific estimate for each country in table A1. The mean social discount

33. The survey responses range from 0 to 5, and it is straightforward to run alternate scenarios with different values. An alternative approach is to use the progressive tax schedule to infer $\eta$ (Groom and Maddison 2018).

34. These data are available online at https://www.ers.usda.govdata-productsinternational-macroeconomic-data-set.

35. In contrast, taking the geometric average would give an estimated growth that matches the cumulative growth rate of all the estimates.

36. It is nevertheless a simple exercise to generate comparable results using the median aggregation of the UDR. In this case, however, one might also want to use a measure of median income growth for consistency. This would imply in the United States, for example, a social discount rate approximately equal to the median UDR because median income growth has been close to zero.

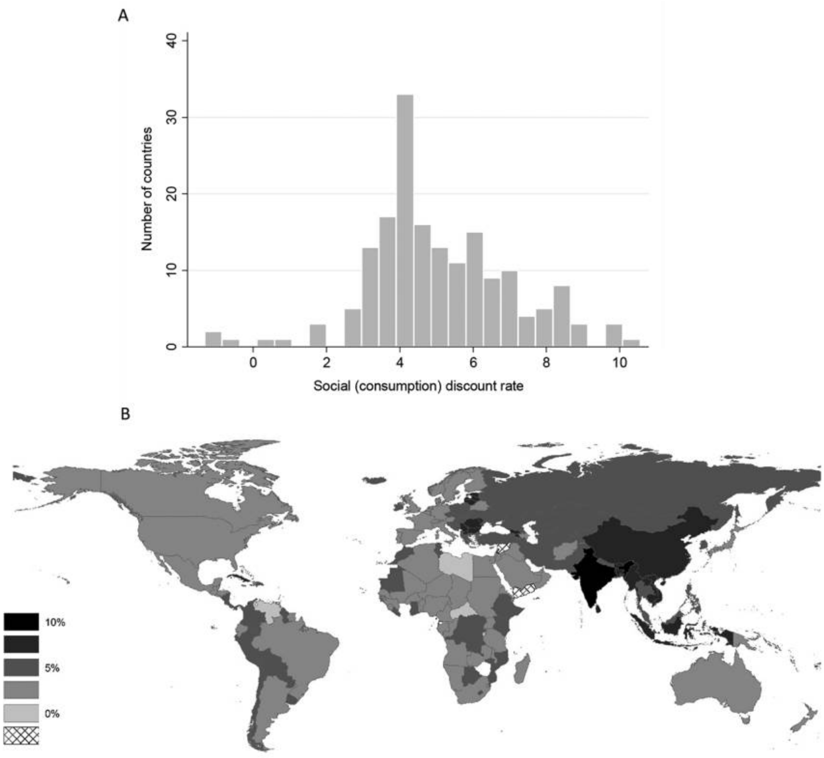


Figure 6. Histogram (*A*) and map (*B*) illustrating the distribution of Ramsey social (consumption) discount rates across countries, using the mean aggregation for the utility discount rates, an elasticity of marginal utility of 1.31, and forecasted GDP growth per capita through 2030 for the consumption growth rate. Color version available as an online enhancement.

rate is 5.07%, with a range from −1.32% to 10.56%.[37] The few negative rates arise because of sufficiently large forecasts of negative GDP growth through 2030. Here again we find that, in general, the results are well within the range of those used in the academic literature and in practice. As a point of reference the US Office of Management and Budget recommends using a social discount rate between 3% and 7% for regulatory impact analysis (Office of Management and Budget 2003). Using our procedure to estimate the social discount rate for the United States, we find a rate of 4.32%. This is also very close to the overall world estimate of 4.61%.

37. Countries with the 10 lowest estimates, in ascending order, are Syria, Equatorial Guinea, Yemen, Venezuela, Burundi, Libya, Brunei Darussalam, Central African Republic, Belize, and Cyprus. Countries with the 10 highest estimates, in descending order, are India, China, Myanmar, Latvia, Vietnam, Bangladesh, Moldova, Laos, Cuba, and Cambodia.

Panel *B* of figure 6 illustrates geographic heterogeneity of the social discount rates. The highest rates tend to be clustered, and China, India, and Southeast Asia, where forecasts of GDP growth per capita, tend to be large. More generally, however, we find that the social discount rates are approximately equally driven by the forecasts of GDP growth and the UDRs. Both contribute somewhat equally because they have somewhat similar magnitudes, and $\eta$ at 1.35 provides a roughly equal weighting between the two. Finally, the estimates of $\delta$ and $g$ are relatively uncorrelated, with a pairwise correlation coefficient of only 0.15.

## 5. APPLICATION TO THE RICE MODEL

We have previously discussed how much of the renewed debate about the choice of discount rates stems from the importance of discounting in IAMs of climate change. In this section, we consider how our heterogeneous estimates of the UDR across countries affect the results of one such model. Our illustration requires a regionally specific model, and we use the archetype and publicly available Regional Integrated Climate Economy (RICE) model (Nordhaus and Yang 1996; Nordhaus 2010). In contrast to the previous section where we derive social discount rates from parameter estimates, the RICE application is a calibration to observed interest rates by region subject to our estimates of the region-specific UDRs.

The RICE model is a dynamic, optimal growth model based on the Ramsey framework. The model's objective is to maximize present discounted utility of consumption through time. Population and technology grow exogenously, but capital accumulation is endogenous, based on the rate at which forgone consumption today is traded off against increased consumption in the future. The buildup of greenhouse gases in the atmosphere causes economically harmful climate damages. Investments in emission reductions operate in much the same way as capital. Emission reductions are costly, they lower consumption today, but increase future consumption possibilities by avoiding future climate damages. A key feature of the model for our purposes is its parameterization to 12 different regions. In the original formulation, Nordhaus and Yang (1996) compare a market scenario, where there is no correction for the climate-change externality, to a globally optimal scenario, where emissions are controlled efficiently across time and regions.

We follow the same approach here, but instead compare the results of two different calibrations. The first is one where all 12 regions of the model have the same global estimate of the UDR, which is the case in Nordhaus and Yang (1996) and Nordhaus (2010). Specifically, we assume a UDR of 2.13% for all regions, corresponding with our 2012 global estimate. This scenario falls between previously published versions of the RICE model, where Nordhaus and Yang (1996) use 3% in their original analysis and Nordhaus (2010) subsequently uses 1.5%. The second calibration that we employ uses heterogeneous estimates of the UDR for each of the 12 regions. To derive these estimates, we apply our procedure described previously to the pooled population

Table 1. Utility Discount Rates Used in the RICE Model for Each Region

| | Utility Discount Rate ($\delta$) | | |
|---|---|---|---|
| Region | Homogeneous | Heterogeneous | Difference |
| Africa | 2.13 | 1.81 | −0.32 |
| Middle East | 2.13 | 1.82 | −0.31 |
| Latin America | 2.13 | 1.91 | −0.22 |
| Other | 2.13 | 2.00 | −0.13 |
| India | 2.13 | 2.17 | 0.04 |
| Other high income (OHI) | 2.13 | 2.31 | 0.19 |
| United States (US) | 2.13 | 2.43 | 0.31 |
| China | 2.13 | 2.45 | 0.33 |
| European Union (EU) | 2.13 | 2.69 | 0.57 |
| Japan | 2.13 | 2.81 | 0.69 |
| Eurasia | 2.13 | 2.89 | 0.77 |
| Russia | 2.13 | 3.33 | 1.21 |

Note. Reported rates are percentages and based on the mean aggregation within the corresponding region or country.

demographics for all countries within a region.[38] Table 1 lists the 12 regions and the UDR for each used in our homogeneous and heterogeneous calibrations, along with the difference. The heterogeneous estimates range from a low of 1.81 for Africa to a high of 3.33 for Russia. Compared to the homogeneous calibration, the heterogeneous UDR is lower in four regions, Africa, Middle East, Latin America, and Other.

When adjusting the UDRs in the RICE model, globally and regionally, the calibration also requires adjustments to the elasticity of the marginal utility of consumption $\eta$. This is done to match the same initial conditions of the model for the overall discount rate and rate of consumption growth. Nordhaus (2007) follows a similar procedure to match observed market conditions when calibrating the related DICE model using different values of $\delta$ in response to Stern et al. (2007).[39]

Figure 7*A* shows the simulated annual emissions corresponding with the market and optimal scenarios for both the homogeneous and heterogeneous calibrations. We find that introducing heterogeneity in the UDRs has very little affect on the annual emissions

38. Note that this is not a weighted average across countries, but rather a weighted average across age bins and life expectancies for all individuals in the population defined by the countries that comprise a RICE region. In other words, a RICE region is treated as a population in the language of our model.

39. In the supplementary material, we provide further details about how to implement the simulations described here using the Excel formulation of RICE-2010. We also report the calibrated values of $\eta$ for both simulations.

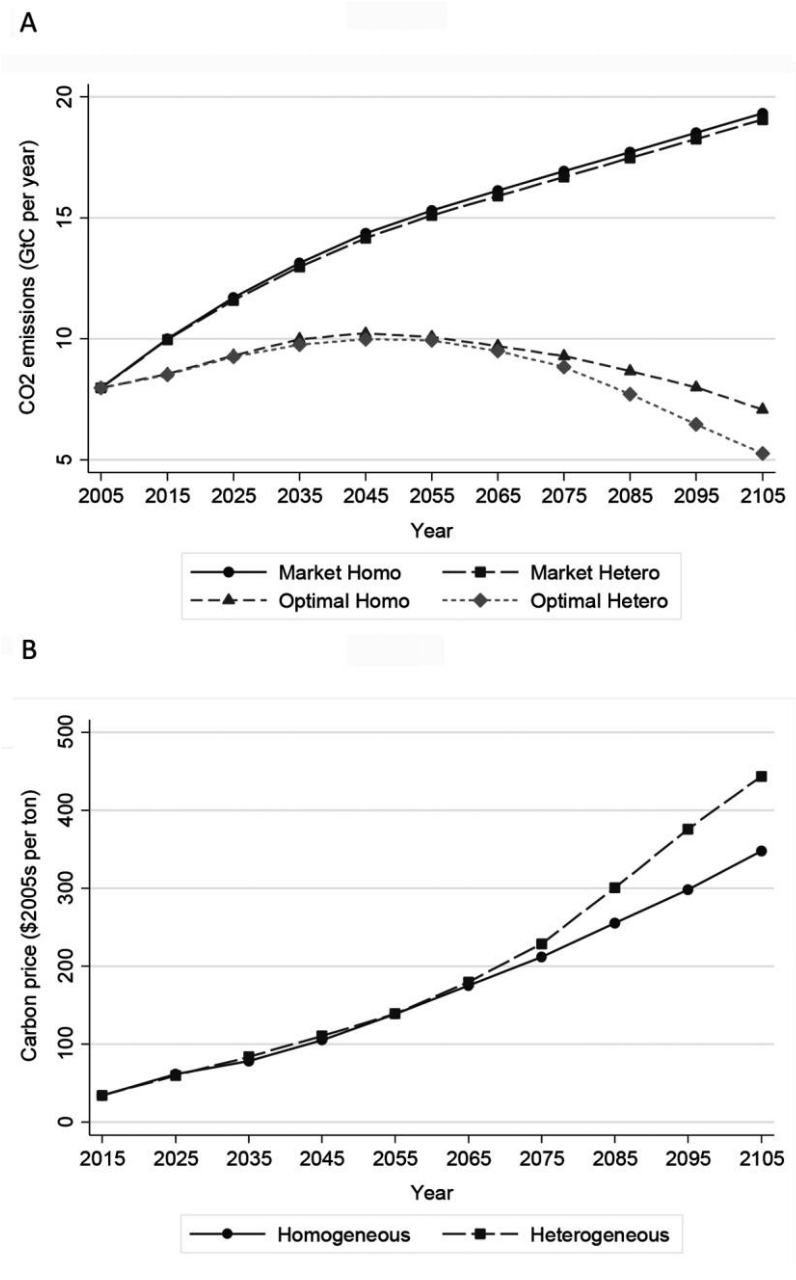


Figure 7. Global annual carbon dioxide emissions, market versus optimal scenarios, homogeneous and heterogeneous utility discount rates (*A*); the optimal carbon price for homogeneous and heterogeneous utility discount rates (*B*). Color version available as an online enhancement.

for the market scenario. The emissions profile is also very similar for the optimal scenario until about 2065, when emissions start to diverge and are lower with heterogeneous UDRs. Panel *B* of the figure shows the uniform, optimal carbon tax for both calibrations. Consistent with the differences in optimal emissions, the carbon tax begins to be greater with the heterogeneous UDRs around 2065. The difference is 28% higher by the end of the century.

Because aggregate emissions may mask some important heterogeneous effects, we now look at region-specific results. Changes in the discount rate affect both the market and optimal scenarios, and it is helpful to look at how each region's emission profile changes even without climate policy. Figure 8 shows the percentage change in annual emissions for each region when moving from the homogeneous to heterogeneous calibrations of the market scenario. A clear pattern emerges whereby regions with lower UDRs increase emissions and those with higher UDRs decrease emissions. The reason stems from the fact that a greater UDR, for example, means less concern for the future compared to the present. This implies a lower savings rate, less capital accumulation, and a decrease in future in output. Then, because emissions are proportional to output in the RICE model, we see a decrease (increase) in emissions for regions that shift to a higher (lower) UDR. See, for example, how after an initial adjustment Africa and the Middle

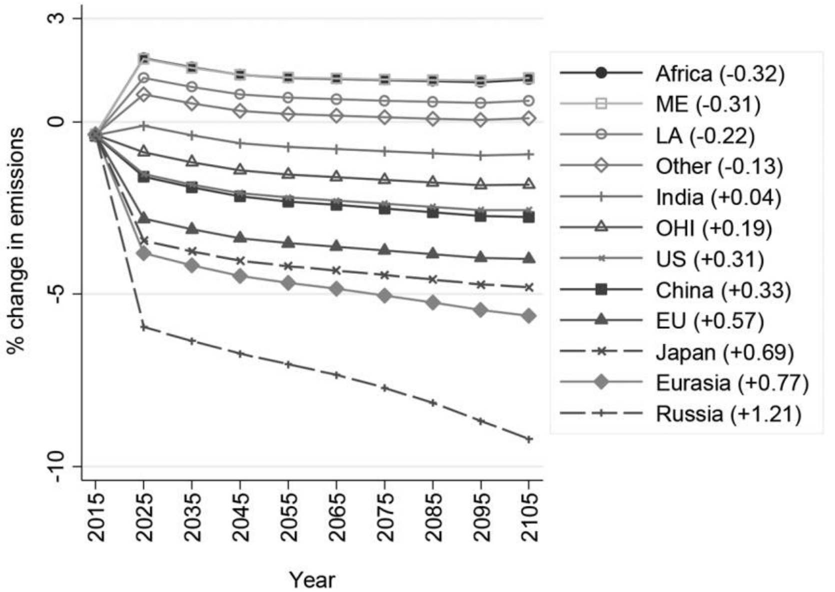


Figure 8. Percentage change in annual emissions by region when moving from the homogeneous to heterogeneous utility discount rates for the market scenario; percentage point change in the UDR is given in parentheses. Color version available as an online enhancement.

East regions increase business-as-usual emissions by about 1.5%, and Russia decreases its business-as-usual emissions more than 9% by the end of the century.

These results underscore the importance of macroeconomic adjustments to changes in the parameters of the discount rate. While researchers focused on climate change often emphasize how a lower UDR supports greater concern for future climate damages, these results also illustrate how the lower UDR also changes the baseline to favor future consumption and therefore future emissions.

What do these differences in the baseline emissions mean for optimal climate policy across regions? Figure 9 shows the percentage change in optimal abatement across regions and over time when moving from the homogeneous to heterogeneous UDRs. The panel for each region is essentially a difference-in-differences calculation because abatement is the difference in emissions between the market and optimal scenarios. Here again there is a clear pattern. Regions with lower UDRs, and therefore higher baseline emissions, abate more in all periods after an initial adjustment, with the increase in abatement reaching about 15% by century's end. Regions with higher UDRs, which have lower baseline emissions, undertake less abatement until just before 2080. The regions with relatively little change in the UDR—that is, India, Other High Income, United States, and China—have very little change in abatement until it increases just before 2080. The differences across regions offset each other such that there is relatively

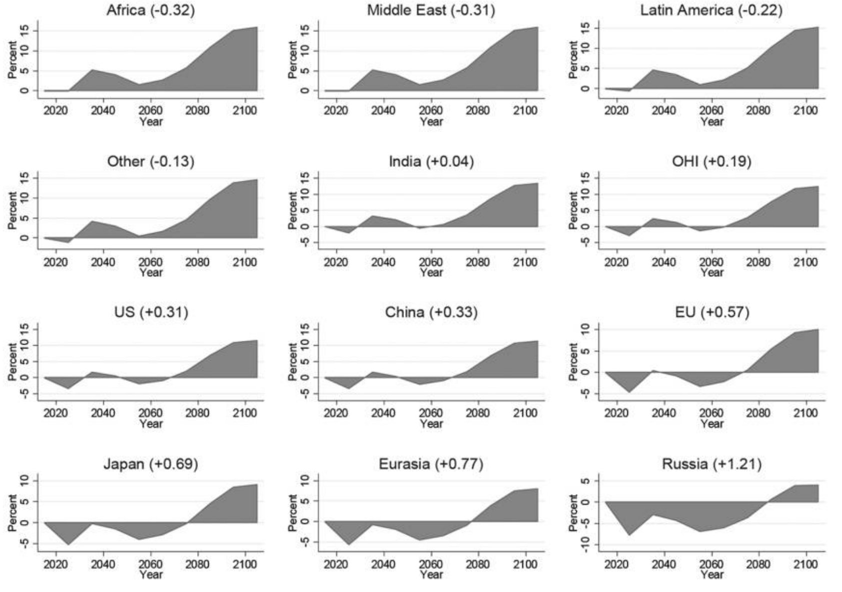


Figure 9. Percentage change in optimal abatement across regions when shifting from the homogeneous to heterogeneous utility discount rates; percentage point change in the UDR is given in parentheses

little change in aggregate emissions between the homogeneous and heterogeneous calibrations for the optimal scenarios prior to 2075 (see fig. 7*A*). Yet aggregate emissions begin to diverge and are lower for the heterogeneous case later in the century, when the large emitting countries, including the United States and China, begin increasing their abatement.

## 6. CONCLUDING REMARKS

The Ramsey discounting rule is a cornerstone of macroeconomic modeling. It is also fundamental to the economics of climate change through its use in IAMs. As originally conceived, the Ramsey model is based on an infinitely lived representative agent. Its interpretation in the climate change literature is generally still that of a representative agent, and the time steps are taken to represent a sequence of nonoverlapping generations. While the Ramsey model provides an organizing framework for pulling together the different components of a long-term discount rate, the difficulty is that the framework itself provides little guidance about how to calibrate the Ramsey rule in practice.

The question of how to choose a social discount rate has long been a subject of debate among economists and philosophers. Comparisons between the two primary schools of thought—a prescriptive or descriptive approach—identify the fundamental difference as explicit or implicit assumptions about compensation among generations (Arrow et al. 1996). This narrows the focus somewhat on the choice of the social UDR. Should the UDR be chosen as an ethical parameter or be backed out of a calibration exercise to match observable market interest rates? The current state of affairs in economics remains one where applied researchers must choose, or remain agnostic and undertake sensitivity analysis that covers a wide range of values. One respect in which the options are particularly unsatisfactory is that the choice between a prescriptive versus descriptive approach is typically inseparable from whether one thinks the discount rate should be low or high, respectively.

This paper presents an alternative, demographic-based approach that can serve as a useful point of comparison. The starting point of our analysis is recognition that even a representative agent must die. Specifically, we show how age-specific mortality rates and life expectancy imply a natural, long-term UDR for individuals at each age in a population. A social planner can then aggregate these preferences into a population-level, social UDR. Moreover, if equal weight is placed on current and future generations, including those not yet born, a stable age distribution implies that the social UDR is constant through time and can be used for long-term program evaluation. A striking part of the analysis, based on the social choice aggregations of using the simple mean or median, is that the results across countries and for the world as a whole fall within the range of UDRs employed in the literature. The results might therefore be interpreted as a reference point for how current practice aligns with the implications of a basic demographic approach.

We are aware that a leading criticism of our approach is that focusing only on mortality risk oversimplifies the basis for individual pure rates of time preference. The criticism is warranted, yet arguments can be made that the approach provides either over- or underestimates. That individuals are well known to be impatient even without mortality risk could be used to argue for a higher UDR. In contrast, the presence of bequest motives for future generations would imply a lower UDR. While many other arguments are certainly possible, the advantage of a purely mortality-based approach is transparency, an empirical basis, and broad data availability—that is, the mortality-based approach gives a nonarbitrary starting point for further adjustments. Another concern that might arise, and is a topic for further research, is the possibility for policies under consideration to substantially affect life expectancies or population distributions. This would imply that a demographically based UDR is endogenous to the policy under consideration. Indeed, Jones (2016) focuses on such endogeneity with respect to the evaluation of policies that are designed to affect life expectancy, and one could reasonably argue that long-term considerations of climate change might have similar effects.

We also show how our approach can provide the basis for exploring heterogeneity of overall social discount rates across countries, with application to IAMs for climate change. We find that introducing regional heterogeneity of UDRs into the RICE model has a small effect on the business-as-usual trajectory of global emissions. However, introducing heterogeneity has a more substantial effect on the trajectory of optimal emissions and the corresponding carbon tax. With heterogeneity, regions with lower UDRs (e.g., Africa, the Middle East, and Latin America) undertake greater emission reductions, whereas those with higher UDRs (e.g., Japan, Eurasia, and Russia) reduce their emissions by less. Although our intent here is primarily illustrative, such results may prove useful in future research for understanding positive and normative aspects of burden sharing across countries and their willingness to participate in international climate agreements.

## APPENDIX

Table A1. Mean and Median Estimates of the 2012 Utility Discount Rate for Each Country

| Country | Mean UDR | Median UDR | Ramsey *r* |
|---|---|---|---|
| Afghanistan | 1.65 | 1.18 | 3.41 |
| Albania | 2.81 | 1.59 | 6.95 |
| Algeria | 2.12 | 1.42 | 4.18 |
| Angola | 1.70 | 1.22 | 3.24 |
| Antigua and Barbuda | 1.94 | 1.30 | 4.20 |
| Argentina | 2.28 | 1.34 | 4.30 |
| Armenia | 3.07 | 1.75 | 8.36 |
| Australia | 2.32 | 1.47 | 4.47 |
| Austria | 2.79 | 1.71 | 4.24 |
| Azerbaijan | 2.18 | 1.45 | 4.30 |
| Bahamas | 2.04 | 1.39 | 3.41 |
| Bahrain | 1.83 | 1.63 | 4.44 |

Table A1 (*Continued*)

| Country | Mean UDR | Median UDR | Ramsey *r* |
|---|---|---|---|
| Bangladesh | 1.94 | 1.34 | 9.04 |
| Barbados | 2.37 | 1.53 | 4.35 |
| Belarus | 3.18 | 1.91 | 4.72 |
| Belgium | 2.81 | 1.70 | 3.92 |
| Belize | 1.64 | 1.18 | 2.53 |
| Benin | 1.73 | 1.11 | 3.71 |
| Bhutan | 1.80 | 1.26 | 8.54 |
| Bolivia | 1.97 | 1.18 | 5.77 |
| Bosnia and Herzegovina | 2.82 | 1.74 | 7.46 |
| Botswana | 1.83 | 1.37 | 5.74 |
| Brazil | 2.02 | 1.33 | 3.45 |
| Brunei Darussalam | 1.78 | 1.35 | 1.72 |
| Bulgaria | 3.50 | 2.06 | 8.30 |
| Burkina Faso | 1.68 | 1.15 | 3.74 |
| Burundi | 1.74 | 1.20 | .99 |
| Cabo Verde | 1.89 | 1.25 | 4.35 |
| Cambodia | 1.44 | 1.12 | 8.61 |
| Cameroon | 1.78 | 1.22 | 4.26 |
| Canada | 2.40 | 1.48 | 4.18 |
| Central African Republic | 2.09 | 1.41 | 1.81 |
| Chad | 1.86 | 1.34 | 4.56 |
| Chile | 2.11 | 1.41 | 5.36 |
| China | 2.45 | 1.77 | 9.78 |
| Colombia | 1.75 | 1.23 | 5.08 |
| Comoros | 1.78 | 1.18 | 3.98 |
| Congo | 1.73 | 1.11 | 3.08 |
| Costa Rica | 1.90 | 1.25 | 5.63 |
| Croatia | 3.12 | 1.97 | 5.93 |
| Cuba | 2.45 | 1.60 | 8.63 |
| Cyprus | 2.31 | 1.42 | 2.69 |
| Czech Republic | 2.89 | 1.66 | 6.01 |
| D. R. Congo | 1.80 | 1.22 | 6.82 |
| Denmark | 2.78 | 1.77 | 4.39 |
| Djibouti | 1.90 | 1.28 | 4.52 |
| Dominican Republic | 1.72 | 1.14 | 6.43 |
| Ecuador | 1.74 | 1.17 | 3.75 |
| Egypt | 2.09 | 1.35 | 4.01 |
| El Salvador | 1.86 | 1.18 | 4.31 |
| Equatorial Guinea | 1.91 | 1.21 | −1.17 |
| Eritrea | 1.78 | 1.29 | 3.46 |

Table A1 (*Continued*)

| Country | Mean UDR | Median UDR | Ramsey *r* |
|---|---|---|---|
| Estonia | 2.98 | 1.92 | 8.17 |
| Ethiopia | 1.57 | 1.04 | 5.87 |
| Fiji | 2.16 | 1.53 | 6.02 |
| Finland | 2.78 | 1.70 | 4.20 |
| France | 2.60 | 1.61 | 3.97 |
| Gabon | 1.86 | 1.18 | 4.50 |
| Gambia | 1.54 | 1.10 | 4.98 |
| Georgia | 3.01 | 1.78 | 7.89 |
| Germany | 2.99 | 1.72 | 4.87 |
| Ghana | 1.72 | 1.20 | 5.58 |
| Greece | 2.93 | 1.78 | 5.33 |
| Grenada | 2.11 | 1.36 | 4.40 |
| Guatemala | 1.52 | .97 | 4.39 |
| Guinea | 1.72 | 1.13 | 3.20 |
| Guinea-Bissau | 1.91 | 1.19 | 4.18 |
| Guyana | 2.42 | 1.48 | 5.91 |
| Haiti | 1.89 | 1.26 | 3.74 |
| Honduras | 1.54 | 1.08 | 3.96 |
| Hungary | 3.17 | 1.79 | 6.31 |
| Iceland | 2.32 | 1.32 | 5.02 |
| India | 2.17 | 1.55 | 10.56 |
| Indonesia | 2.08 | 1.44 | 7.90 |
| Iran | 1.96 | 1.39 | 5.61 |
| Iraq | 1.71 | 1.17 | 3.31 |
| Ireland | 2.23 | 1.39 | 5.21 |
| Israel | 2.13 | 1.39 | 4.80 |
| Italy | 2.91 | 1.70 | 3.68 |
| Jamaica | 2.13 | 1.22 | 3.80 |
| Japan | 2.81 | 1.59 | 4.23 |
| Jordan | 1.72 | 1.26 | 3.45 |
| Kazakhstan | 2.61 | 1.58 | 5.14 |
| Kenya | 1.56 | 1.06 | 6.52 |
| Kiribati | 1.84 | 1.30 | . . . |
| Kuwait | 1.69 | 1.44 | 2.71 |
| Kyrgyzstan | 2.08 | 1.33 | 5.31 |
| Laos | 1.82 | 1.28 | 8.65 |
| Latvia | 3.36 | 2.11 | 9.64 |
| Lebanon | 2.04 | 1.30 | 4.14 |
| Lesotho | 2.43 | 1.98 | 7.17 |
| Liberia | 1.75 | 1.15 | 5.96 |

Table A1 (*Continued*)

| Country | Mean UDR | Median UDR | Ramsey *r* |
|---|---|---|---|
| Libya | 1.83 | 1.39 | 1.57 |
| Lithuania | 3.10 | 2.05 | 8.02 |
| Luxembourg | 2.41 | 1.52 | 3.69 |
| Madagascar | 1.69 | 1.14 | 2.98 |
| Malawi | 1.72 | 1.19 | 3.39 |
| Malaysia | 2.01 | 1.42 | 6.56 |
| Maldives | 1.98 | 1.46 | 7.33 |
| Mali | 1.66 | 1.12 | 3.65 |
| Malta | 2.70 | 1.80 | 5.98 |
| Mauritania | 1.74 | 1.10 | 5.95 |
| Mauritius | 2.25 | 1.57 | 6.64 |
| Mexico | 1.79 | 1.30 | 4.14 |
| Micronesia | 1.91 | 1.28 | 3.16 |
| Mongolia | 2.09 | 1.53 | 7.07 |
| Montenegro | 2.77 | 1.78 | . . . |
| Morocco | 2.19 | 1.43 | 6.01 |
| Mozambique | 2.00 | 1.49 | 6.98 |
| Myanmar | 2.18 | 1.46 | 9.71 |
| Namibia | 1.67 | 1.21 | 4.12 |
| Nepal | 1.99 | 1.27 | 6.17 |
| Netherlands | 2.68 | 1.72 | 4.38 |
| New Zealand | 2.24 | 1.46 | 4.55 |
| Nicaragua | 1.70 | 1.21 | 6.06 |
| Niger | 1.65 | 1.11 | 4.17 |
| Nigeria | 1.77 | 1.22 | 3.59 |
| North Korea | 2.69 | 1.74 | . . . |
| Norway | 2.67 | 1.53 | 3.93 |
| Oman | 1.70 | 1.43 | 3.55 |
| Pakistan | 1.88 | 1.28 | 7.08 |
| Panama | 1.78 | 1.22 | 6.02 |
| Papua New Guinea | 2.01 | 1.49 | 4.84 |
| Paraguay | 1.74 | 1.18 | 5.14 |
| Peru | 1.75 | 1.22 | 5.61 |
| Philippines | 2.01 | 1.33 | 6.84 |
| Poland | 2.71 | 1.68 | 6.69 |
| Portugal | 2.81 | 1.72 | 4.00 |
| Qatar | 1.63 | 1.51 | . . . |
| Republic of Korea | 2.28 | 1.60 | 5.13 |
| Republic of Macedonia | 2.76 | 1.79 | 7.56 |
| Republic of Moldova | 2.91 | 1.74 | 8.98 |

Table A1 (*Continued*)

| Country | Mean UDR | Median UDR | Ramsey $r$ |
|---|---|---|---|
| Romania | 3.13 | 1.81 | 8.42 |
| Russian Federation | 3.33 | 2.02 | 5.13 |
| Rwanda | 1.53 | 1.05 | 5.41 |
| Saint Lucia | 2.15 | 1.35 | 4.25 |
| St. Vincent and Grenadines | 1.98 | 1.23 | 5.64 |
| Samoa | 1.93 | 1.32 | 3.07 |
| Sao Tome and Principe | 1.63 | 1.06 | 6.40 |
| Saudi Arabia | 1.80 | 1.48 | 4.76 |
| Senegal | 1.68 | 1.10 | 4.32 |
| Serbia | 3.23 | 1.84 | 7.91 |
| Seychelles | 2.21 | 1.59 | 5.95 |
| Sierra Leone | 2.12 | 1.35 | 3.99 |
| Singapore | 2.10 | 1.59 | 3.71 |
| Slovakia | 2.76 | 1.75 | 6.39 |
| Slovenia | 2.75 | 1.77 | 6.78 |
| Solomon Islands | 1.76 | 1.16 | 3.44 |
| Somalia | 1.74 | 1.20 | . . . |
| South Africa | 2.30 | 1.55 | 3.66 |
| South Sudan | 1.82 | 1.22 | . . . |
| Spain | 2.72 | 1.72 | 4.23 |
| Sri Lanka | 2.19 | 1.52 | 8.47 |
| Sudan | 1.66 | 1.10 | 2.81 |
| Suriname | 1.68 | 1.19 | 4.31 |
| Swaziland | 1.99 | 1.45 | 3.40 |
| Sweden | 2.90 | 1.72 | 4.44 |
| Switzerland | 2.62 | 1.65 | 3.77 |
| Syrian Arab Republic | 1.62 | 1.13 | −1.32 |
| Tajikistan | 1.93 | 1.45 | 5.51 |
| Tanzania | 1.61 | 1.09 | 4.86 |
| Thailand | 2.33 | 1.73 | 5.94 |
| Timor-Leste | 1.87 | 1.16 | . . . |
| Togo | 1.65 | 1.14 | 4.34 |
| Tonga | 2.07 | 1.34 | 4.52 |
| Trinidad and Tobago | 2.57 | 1.69 | 4.96 |
| Tunisia | 2.13 | 1.36 | 6.31 |
| Turkey | 2.02 | 1.28 | 5.43 |
| Turkmenistan | 2.28 | 1.53 | 7.26 |
| Uganda | 1.66 | 1.24 | 4.04 |
| Ukraine | 3.44 | 1.94 | 6.90 |
| United Arab Emirates | 1.67 | 1.59 | 3.31 |
| United Kingdom | 2.64 | 1.52 | 4.85 |

Table A1 (*Continued*)

| Country | Mean UDR | Median UDR | Ramsey *r* |
|---|---|---|---|
| United States | 2.43 | 1.49 | 4.32 |
| Uruguay | 2.57 | 1.47 | 6.26 |
| Uzbekistan | 2.23 | 1.48 | 6.84 |
| Vanuatu | 1.80 | 1.28 | 2.92 |
| Venezuela | 1.67 | 1.22 | 0.28 |
| Vietnam | 1.95 | 1.34 | 9.06 |
| Yemen | 1.69 | 1.16 | −0.53 |
| Zambia | 1.69 | 1.24 | 3.30 |
| Zimbabwe | 1.65 | 1.17 | . . . |
| World | 2.13 | 1.33 | 4.61 |

## REFERENCES

Arrow, Kenneth J. 1999. Discounting, morality, and gaming. In *Discounting and intergenerational equity*, ed. Paul Portney and John Weyant, 13–21. Washington, DC: Resources for the Future.

Arrow, Kenneth J., William R. Cline, Karl-Göran Maler, Mohan Munasinghe, R. Squitieri, and Joseph E. Stiglitz. 1996. *Intertemporal equity, discounting, and economic efficiency*. Cambridge: Cambridge University Press.

Arrow, Kenneth J., Maureen Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G. Newell, William D. Nordhaus, et al. 2012. How should benefits and costs be discounted in an intergenerational context? The views of an expert panel. Resources for the Future Discussion paper RFF DP 12-53. Resources for the Future, Washington, DC.

Arrow, Kenneth J., Maureen L. Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G.. Newell, William D. Nordhaus, et al. 2014. Should governments use a declining discount rate in project analysis? *Review of Environmental Economics and Policy* 8 (2): 145–63.

Carvalho, Carlo, Andrea Ferrero, and Fernanda Nechio. 2016. Demographics and real interest rates: Inspecting the mechanism. *European Economic Review* 88:208–26.

Cline, William R., et al. 1992. *The economics of global warming*. Institute for International Economics, Washington, DC.

Cropper, Maureen L., Mark C. Freeman, Ben Groom, and William A. Pizer. 2014. Declining discount rates. *American Economic Review* 104 (5): 538–43.

Dasgupta, Partha. 2007. *Human well-being and the natural environment*. New York: Oxford University Press.

———. 2008. Discounting climate change. *Journal of Risk and Uncertainty* 37:141–69.

Dasgupta, Partha, Karl-Goran Maler, and Scott Barrett. 1999. Intergenerational equity, social discount rates and global warming. In *Discounting and intergenerational equity*, ed. Paul Portney and John Weyant, 51–78. Washington, DC: Resources for the Future.

Drupp, Moritz, Mark Freeman, Ben Groom, and Frikk Nesje. 2018. Discounting disentangled. *American Economic Journal: Economic Policy* 10 (4): 109–34.

Eckstein, Otto. 1961. A survey of the theory of public expenditure criteria. In *Public finances: Needs, sources, and utilization*, ed. Universities-National Bureau, National Bureau of Economic Research, 439–504. Princeton, NJ: Princeton University Press.

Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins. 2017. A model of secular stagnation: Theory and quantitative evaluation. NBER Working paper w23093. National Bureau of Economic Research, Cambridge, MA.

Emmerling, Johannes, Ben Groom, and Tanja Wettingfeld. 2017. Discounting and the representative median agent. *Economics Letters* 161:78–81.

Fisher, Irving. 1930. *The theory of interest as determined by impatience to spend income and opportunity to invest it.* New York: Macmillan.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. Time discounting and time preference: A critical review. *Journal of Economic Literature* 40 (2): 351–401.

Freeman, Mark C., and Ben Groom. 2015. Positively gamma discounting: Combining the opinions of experts on the social discount rate. *Economic Journal* 125:1015–24.

Gagnon, Etienne, Benjamin K. Johannsen, and David Lopez-Salido. 2016. Understanding the new normal: The role of demographics. Discussion paper, Washington Board of Governors of the Federal Reserve System.

Gollier, Christian. 2013. *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton, NJ: Princeton University Press.

———. 2015. Discounting, inequality and economic convergence. *Journal of Environmental Economics and Management* 69:53–61.

Gollier, Christian, and Richard Zeckhauser. 2005. Aggregation of heterogeneous time preferences. *Journal of Political Economy* 113 (4): 878–96.

Goulder, Lawrence H., and Roberton C. Williams III. 2012. The choice of discount rate for climate change policy evaluation. Technical report, Resources for the Future, Washington, DC.

Groom, Ben, and David Maddison. 2018. New estimates of the elasticity of marginal utility for the UK. *Environmental and Resource Economics* 72 (4): 1–28.

Heal, Geoffrey M., and Antony Millner. 2014a. Agreeing to disagree on climate policy. *Proceedings of the National Academy of Sciences* 111 (10): 3695–98.

———. 2014b. Uncertainty and decision making in climate change economics. *Review of Environmental Economics and Policy* 8 (1): 120–37.

Jones, Charles I. 2016. Life and growth. *Journal of Political Economy* 124:539–78.

Jouini, Elyes, Jean-Michel Marin, and Clotilde Napp. 2010. Discounting and divergence of opinion. *Journal of Economic Theory* 145:830–59.

Kula, Erhun. 1985. An empirical investigation on the social time-preference rate for the United Kingdom. *Environment and Planning A* 17:199–212.

Lotka, A. J. 1922. The stability of normal age distributions. *Proceedings of the National Academy of Sciences* 9:339–45.

Millner, Antony. 2018. Non-dogmatic social discounting. Working paper University of California, Santa Barbara.

Millner, Antony, and Geoffrey Heal. 2018. Discounting by committee. *Journal of Public Economics* 167:91–104.

Nordhaus, William D. 2007. A review of the Stern review on the economics of climate change. *Journal of Economic Literature* 45 (3): 686–702.

———. 2008. *A question of balance: Weighing the options on global warming policies*. New Haven, CT: Yale University Press.

———. 2010. Economic aspects of global warming in a post-Copenhagen environment. *Proceedings of the National Academy of Sciences* 107 (26): 11721–26.

Nordhaus, William D., and Zili Yang. 1996. A regional dynamic general-equilibrium model of alternative climate-change strategies. *American Economic Review* 86 (4): 741–65.

Office of Management and Budget. 2003. Circular A-4: Regulatory analysis. OMB, Washington, DC.

Preston, Samuel H., ed. 1978. *The effects of infant and child mortality on fertility*. New York: Academic Press.

Preston, Samuel H., Christen Himes, and Mitchell Eggers. 1989. Demographic conditions responsible for population aging. *Demography* 26 (4): 691–704.

Rae, John M. A. 1834. *The sociological theory of capital*. New York: Macmillan.

Ramsey, F. P. 1928. A mathematical theory of saving. *Economic Journal* 38 (152): 543–59.

Roser, Max. 2019. Child and infant mortality. *Our world in data*. https://ourworldindata.org/child-mortality.

Samuelson, Paul A. 1958. An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of Political Economy* 66 (6): 467–82.

Schneider, Maik T., Christian P. Traeger, and Ralph Winkler. 2012. Trading off generations: Equity, discounting, and climate change. *European Economic Review* 56:1621–44.

Stern, Nicholas, et al. 2007. *The economics of climate change: The Stern report*. Cambridge: Cambridge University Press.

Weitzman, Martin L. 1998. Why the far-distant future should be discounted at its lowest possible rate. *Journal of Environmental Economics and Management* 36:201–8.

———. 2001. Gamma discounting. *American Economic Review* 91 (1): 260–71.

———. 2007. A review of the Stern review on the economics of climate change. *Journal of Economic Literature* 45 (3): 703–24.

———. 2009. On modeling and interpreting the economics of catastrophic climate change. *Review of Economics and Statistics* 91 (1): 1–19.

World Health Organization. 2016. WHO methods and data sources for life tables 1990–2015. Global Health Estimates Technical paper WHO/HIS/IER/GHE/2016.2 (May), WHO, Geneva.

Yaari, Menahem E. 1964. On the consumer's lifetime allocation process. *International Economic Review* 5 (3): 304–17.