ENVIRONMENT

How Should We Value the Future?

Determining the social cost of carbon is as much a matter of ethics as it is of science and economics.

BY DAVID KEMP AND PETER VAN DOREN

Immediately following his second inauguration, President Donald Trump signed Executive Order 14154, which in part directs federal agencies to stop considering the "social cost of carbon" (SCC) when assessing the effects of regulations. The social cost of carbon dioxide and other greenhouse gases shown to contribute to climate change is an estimate of the monetized damages caused by emitting an additional ton of the carbon-equivalent pollutant. It serves as a benchmark against which the costs of regulations are evaluated. A higher estimated SCC implies that more aggressive climate policies are warranted; a lower SCC suggests a more moderate approach.

EO 14154 charges that the SCC "is marked by logical deficiencies, a poor basis in empirical science, politicization, and the absence of a foundation in legislation." A *New York Times* article summarized the response to the order from scientists, environmentalists, and economists, many of whom argued that, by ignoring the SCC, the Trump administration is refusing to acknowledge the scientific reality of climate change (Friedman, 2025). For example, Michael Greenstone, a former chief economist on the Obama-era Council of Economic Advisers, said that the new order means that "'feelings, not facts' would guide federal policy."

The truth, however, is that the SCC has always rested as much on feelings as on facts. The estimated damages of greenhouse gases are highly sensitive to normative, non-scientific judgments. Even if researchers agree on the basic facts of climate change and its consequences, estimates of the SCC and the level of government action it implies vary wildly depending on their choice of a few subjective assumptions.

David Kemp is a research fellow for climate and energy policy at C3 Solutions. Peter Van Doren is a Cato Institute senior fellow and editor of *Regulation*.

This is why an observer of federal climate policy may feel whiplash from the dramatic swings in the SCC between successive administrations. Under President Barack Obama, the first official SCC estimate was $42 per ton of carbon dioxide in 2020. During Trump's first term, that value was revised down to $7 per ton or less. The Biden administration raised it to $190 per ton.

Much of this volatility stems from differences in discount rates. Discounting is a core economic concept used to value money at different periods in time. The higher the rate at which future outcomes are discounted, the less we value those outcomes today. In the climate change context, where greenhouse gases can persist in the atmosphere for hundreds of years and today's emissions can affect people far into the future, discounting takes on an important ethical role. It dictates how we weigh intergenerational tradeoffs: How much present-day well-being should we forgo to benefit future generations, and vice versa?

Contrary to the view that estimating the SCC is a purely objective exercise, the choice of discount rate is inherently unscientific. Climate science can estimate the physical effects of emissions, and economics can attempt to convert those effects into monetary damages, but only ethics can tell us how much the damages to future generations should matter to us today.

These ethical questions have been debated since before there was any broad consensus on anthropogenic climate change. Leading intellectuals, including philosopher John Rawls and Nobel laureate economist Kenneth Arrow, contributed to these debates. Yet, decades later, we remain no closer to a definitive answer about how to value the well-being of future generations relative to our own.

In practice, researchers often avoid these ethical questions by falsely presenting discount rate choices as the result of

GUIDOVROLA / GETTY IMAGES (CRYSTAL BALL), AERIALPERSPECTIVE WORKS / GETTY IMAGES (POWER PLANT)

scientific inquiry. Others, such as the Trump administration, have dismissed SCC estimation altogether. Neither approach is appropriate. The normative dimensions of climate policy should be transparent so that policymakers and the public can better understand the consequences of their decisions.

THE CHALLENGE OF INTERGENERATIONAL CLIMATE DAMAGE

By affecting the Earth's climate, greenhouse gas emissions create damages whose marginal costs are not borne by emitters. Government intervention—either a carbon tax or regulation—may improve welfare if benefits exceed costs. But to select an appropriate carbon tax or evaluate the net benefits of other policy options, policymakers need to know the marginal damages of the emissions. This is the SCC, a measure of the incremental damages created by emitting one additional ton of carbon dioxide or an equivalent pollutant into the atmosphere.

If the benefits and costs of emitting the pollutant were confined to a single time period or generation, estimating the marginal damages would be difficult enough. But greenhouse gases can remain in the atmosphere for centuries. Any estimate of the SCC must ask how present generations value damages that will occur long after they have died.

The SCC is typically calculated using Integrated Assessment

Models (IAMs), so-called because they integrate socioeconomic projections, climate models estimating how emissions cause temperature change, and economic damage functions relating the temperature increases to environmental harm. The final step in IAMs is discounting: converting the future marginal damages of a ton of pollution to a present value.

Discounting reflects the economic concept that a dollar today is worth more than a dollar in the future. Imagine you are given a choice: You can either receive $100 today or in one year. Most people would take the money offered now over the money in the future. How much higher would the future amount need to be for the chooser to be ambivalent between it and the present $100? If, for example, you have no preference between receiving $100 today or $105 in a year, then your discount rate is 5 percent (*Future Value* = *Present Value* × (100% + *Discount Rate*).

The discount rate can be used to estimate the present value of money received later. Given our implied discount rate of 5 percent, the present value of $100 received in a year's time is $95.24 (*Present Value* = *Future Value* × $\frac{1}{100\% + \textit{Discount Rate}}$).

One of the central challenges of climate economics is determining the appropriate discount rate to apply when evaluating damages that will occur in the far future. This is a purely ethical choice, reflecting how much we value the well-being of future generations relative to our own. On one hand, it is difficult to defend steeply discounting future harms solely because they occur later in time. On the other hand, it is not obvious how much of our current well-being we should be willing to sacrifice for the sake of future people—especially when, as long as economic growth continues at historical rates, those future people are likely to be significantly wealthier than we are.

The results of IAMs are fragile to the assumed discount rates. The rate chosen not only encapsulates challenging ethical judgments, but it also dictates the results of any analysis. The most prominent example is Stern (2006), commonly called the *Stern Review*, written for the government of the United Kingdom. It recommended aggressive global action to avoid the damages of climate change. Its estimates of future damages were above those of other economists at the time and previously. Many of them responded with critiques. A prominent one, Nordhaus (2007), used its own IAM, the Dynamic Integrated Climate-Economy (DICE) model, and found smaller future damages that implied a more moderate approach to climate policy. Nordhaus noted that the difference between his findings and Stern's was not driven by different perspectives on the scientific, economic, or modeling dynamics of climate change, but resulted from Stern's reliance on a much lower discount rate of 1.4 percent compared to Nordhaus's baseline discount rate of 5.5 percent.

Similarly, the large changes in the official SCC estimates used by the US federal government are largely caused by discount rate changes. For example, in 2021 the Biden administration updated the Obama-era methodology to estimate an SCC of $51 in 2020 at a 3 percent discount rate. A mere two years later, the administration announced that a new methodology, based largely on an IAM developed by researchers at the environmental think tank Resources for the Future (RFF) and the University of California, Berkeley, named Greenhouse Gas Impact Value Estimator (GIVE), determined a new SCC of $190.

The primary difference between these two calculations? The new estimate relied on a discount rate of 2 percent. At the old 3 percent rate, the GIVE model finds an SCC that is larger, but after discounting is much more in line with the previous estimates.

A simple example illustrates the importance of discounting and why IAMs come to such divergent conclusions given relatively small differences in discount rates. Imagine that scientists know climate change will cause damages of $100 trillion in 100 years. Fortunately, they have also found a solution that will require a one-time investment today of $10 trillion. What should we do? At a rate of 2 percent, the future damages are worth about $14 trillion in present value ($100 *trillion* × $(\frac{1}{100\% + 2\%})^{100\,\textit{years}}$). Thus, the most cost-effective strategy is to invest in the solution today at $10 trillion to avoid damages of $14 trillion. At a rate of 3 percent, however, the damages are $5 trillion in present value. The avoided damages are not worth the required investment.

HOW SHOULD WE DISCOUNT THE FUTURE?

Given the significance of the discount rate and IAM's sensitivity to small changes, how the rate is determined is vitally important. Academic debates over how society should derive a social discount rate began more than 60 years ago about the evaluation of the costs and benefits of public investment in projects with long-term benefits, such as water infrastructure. The participants understood that discounting concerned both technical and ethical implications.

As knowledge and concern over climate change grew, discounting and related ethical decisions were quickly identified as core questions. Early climate policy discussions explored how discount rates should be derived. The Second Assessment Report of the Intergovernmental Panel on Climate Change (IPCC), for example, included a chapter on discounting lead-authored by Arrow (Arrow et al., 1996).

Economists have frequently acknowledged that discounting is ultimately an ethical rather than scientific or economic question. Arrow (1999) took "the problem for discounting for projects with payoffs in the far future (climate change, nuclear waste disposal) to be largely ethical." Similarly, Greenstone and Carleton (2022) concede that "because climate change involves intergeneration tradeoffs, it raises difficult economic, scientific, philosophical, and legal questions regarding equity across long periods of time. Thus, there is no scientific consensus about the correct approach to discounting for the SCC." They cite Gollier and Hammit (2014), who noted that "there is no consensus about what would be the socially desirable level of the discount rate. We must recognize that economics cannot provide a full answer to this question, which involves deep ethical issues."

Despite these acknowledgments, the practical necessity of selecting a discount rate has led to two broad approaches, which Arrow et al. (1996) labeled "descriptivist" and "prescriptivist." "Descriptivists" derive rates from observed market behavior, while the "prescriptivist" approach sets discount rates based on normative (i.e., ethical) reasoning.

Descriptive discounting / The appeal of descriptive discount rates is their purported objectivity. Two main arguments support this approach. First, market interest rates represent how society actually values present versus future consumption. Second, descriptivists treat the discount rate as the opportunity cost of the resources spent on climate policy. Investment in mitigation could be spent elsewhere. If the returns of alternative investments exceed the benefits of mitigation, then pursuing mitigation could make future generations worse off.

Consider again our earlier scenario: Scientists project that climate change will cost $100 trillion in 100 years, and we could spend $10 trillion today to avoid that damage. Should we spend the money now, or invest it instead? The answer depends on the expected rate of return: At a 3 percent annual return, $10 trillion would grow to over $190 trillion in 100 years. That would be more than enough to offset the damages with money to spare. At a 2 percent return, however, $10 trillion would grow to only about $70 trillion in 100 years, making mitigation a better investment. In this sense, the discount rate is used to identify an arbitrage opportunity. If an alternative investment yields more than the avoided damages, then mitigation is inefficient.

A challenge in descriptive discounting is determining which market interest rate is most appropriate. Many economists have recently settled on the use of inflation-adjusted yields of US Treasury bonds, often considered a "risk-free" rate that isolates society's impatience for present benefits from other factors.

This approach has been adopted by the federal government. In 2003 the Office of Management and Budget (OMB) directed federal agencies to use a 3 percent discount rate for evaluating regulations affecting private consumption. The rate was based on the average real yields for 10-year Treasury bonds from 1973 to 2003. Since then, real returns have declined, averaging closer to 2 percent. This prompted the OMB to revise discount rates downward and justified the Biden administration's use of a 2 percent discount rate to estimate a much higher SCC.

Descriptive discounting appears empirical, but different methods applied to the same observations can produce drastically different rates. In recent work, for example, Nordhaus uses a preferred discount rate of 4.5 percent, substantially higher than the 2 percent rate used by the Biden administration. Both rates are based on the same observed US Treasury yields. Divergent views on how discount rates should incorporate other factors, such as uncertainty and risk, can lead to wildly different rates. The Nordhaus–Biden difference also suggests the possibility that normative views of discounting determine methodological choices, undermining the alleged empiricism of descriptive discounting.

Even if market interest rates do represent the fundamental time preferences of individuals, it's not clear that they are answering the right question. Interest rates represent *intrapersonal* choices about consumption and saving over individuals' own lifetimes, not how people weigh *intergenerational* tradeoffs. And, importantly, they represent only the preferences of present people, not the preferences of future generations.

The opportunity cost argument also falters upon close examination. As the IPCC's Fifth Assessment Report (2014) argues, "We do not observe safe assets with maturities similar to those of climate impacts, so the arbitrage argument cannot be applied." There are no real markets for investing today that can transfer resources securely across centuries, and it may not be possible to create such mechanisms.

Prescriptive discounting / The prescriptive approach begins from the premise that ethical reasoning, not observed behavior, should dictate how we value the future. It is typically grounded in a simple equation based on the work of British mathematician and economist Frank Ramsey, who in a 1928 paper asked how much income a nation should save for the future. Although the "Ramsey formula" itself did not appear in his paper, it was later derived from his model and has since become the cornerstone of modern economic thinking on how to discount future well-being.

The Ramsey formula expresses the discount rate (r) as:

$$r = \rho + \eta \times g$$

where:

- ρ is the pure rate of time preference (PRTP), or how much more we value consumption today than in the future;
- η is a measure of inequality aversion (or sometimes risk aversion);
- and g is the rate of economic growth.

In simple terms, this formula suggests we discount the future for two reasons: The first, the PRTP, is a measure of impatience, or how much we prefer immediate consumption to future consumption. A higher ρ means we place greater value on the present relative to the future. The second reason is the expectation that future people will be wealthier ($\eta \times g$), reflecting the idea that a dollar is worth less to a richer person than to a poorer one because of diminishing returns. The parameter η is our societal aversion to inequality and its value dictates how highly we weight the well-being of poorer people relative to richer people. If we care about equity (i.e., a high η) and expect future people to be wealthier (i.e., g is positive), then we should discount their well-being more.

The ethical considerations of the formula are expressed by the values chosen for the PRTP and inequality aversion. A National Academies of Sciences, Engineering, and Medicine report (NASEM, 2017) notes that "while g is determined by the performance of the economy and is observable (ex post), ρ and η are never observable, but require an ethical judgment."

As a result, the choice of parameter values typically represents the ethical viewpoints of the researcher. Prescriptive approaches have traditionally favored rates lower than those implied by descriptive approaches based on the belief that observed rates place too little value on the future.

There are, however, legitimate reasons why the appropriate social discount rate could be higher or lower than the market rate. On one hand, market rates do not reflect the preferences of future people who stand to benefit from long-term investments. Because they value those benefits more than current generations, their exclusion means observed rates may be too high. On the other hand, today's rates mostly reflect the preferences of wealthier individuals who are more able to save and invest. They do not capture the perspective of poorer people, who have less ability to save and place greater importance on immediate benefits. Their absence suggests market rates are too low. This may be especially relevant in a global context, where discount rates in developing countries are much higher than in developed countries. Thus, when evaluating global costs and benefits of climate policies from an ethical standpoint, the social discount rate should reflect both the interests of future generations and those of the present-day poor in low-income countries. Because market rates reflect neither, it is unclear whether the rate should be adjusted upward or downward.

The most famous application of prescriptive discounting is in the *Stern Review*. It chose a near-zero pure rate of time preference (0.1 percent) and low inequality aversion, which, combined with an average projected economic growth of 1.3 percent, resulted in an average discount rate of 1.4 percent. Much of the critique of the *Review* focused on those parameter choices.

Nordhaus (2007), for instance, argued:

> The *Review* takes the lofty vantage point of the world social planner, perhaps stoking the dying embers of the British Empire, in determining the way the world should combat the dangers of global warming. The world, according to Government House utilitarianism, should use the combination of [pure rate of time preference] and consumption elasticity [i.e., inequality aversion] that the *Review*'s authors find persuasive from their ethical vantage point.

This is the primary drawback of prescriptive discounting: It risks technocratic overreach. Researchers impose their own ethical judgments into their analyses even though those values may not reflect public preferences.

A drawback of prescriptive discounting is researchers impose their own judgments into their analyses even though those values may not reflect public preferences.

Ethics of the Ramsey formula / Although much ink has been spilled exploring the ethical dimensions of different discount rates and values of the Ramsey parameters, there is no consensus on the right values to use. Taken to their extremes, both the pure rate of time preference and inequality aversion suggest potentially troubling conclusions. Ultimately, the Ramsey formula is an oversimplification of an incredibly complex ethical dilemma. Exploring different values of the parameters and their implications, however, helps illuminate some of the tradeoffs.

Pure rate of time preference When the ethical implications of the PRTP are considered, especially at large values, it is difficult to defend. For example, assuming that there will be no future growth (i.e., $g = 0$ in the Ramsey equation), a PRTP of 1.5 percent implies that a person's well-being in 50 years is worth about half our own. The postscript to the *Stern Review* observes, "In other words, a grandparent would tell a grandchild that simply because the latter's consumption flow came later (e.g., 50 years) in time than his or her own consumption flow, it would be correct to assign a value of less than half to it in thinking about the consequences of actions today." The implications grow more severe the further out we look: A person's well-being in 300 years would be valued at just 1 percent of a present person.

However, a PRTP of zero also leads to some troubling conclusions. It demands that we value other people's well-being as equal to our own. Under certain assumptions, including a positive rate of return on investment and an infinite time horizon, Koopmans et al. (1964) found that this can imply that society should defer all consumption for future generations. These findings have been challenged based on their assumption of an infinite time horizon, but their underlying intuition remains the same: A zero PRTP can lead to excessive sacrifice for the distant future.

Many prominent scholars, including Rawls, Arthur Pigou, and Robert Solow, have argued that zero is the only morally defensible PRTP. Ramsey himself famously wrote that a non-zero PRTP is "a practice which is ethically indefensible and arises merely from the weakness of the imagination." Gollier and Hammit (2014) even liken a non-zero PRTP and time-based discounting to racism or sexism.

Others have disagreed. In a paper on the ethics of the *Stern Review*, Beckerman and Hepburn (2007) note that people routinely weigh the interests of others unequally. For example, parents may value the well-being of their children more than strangers, and nations may value the well-being of citizens more than foreigners. Likewise, it is not necessarily unethical for current generations to discount the well-being of people who are not yet born. They contend:

> The argument made by the *Review* rests on an (implicit) analogy: if a parent should not discount the welfare of their children and grandchildren, nor should the global decision maker discount the welfare of future generations. Yet it is not necessarily unethical to place decreasing weight on the well-being of our increasingly distant descendants, corresponding to a positive utility discount rate.

Convinced by Koopmans' arguments, Arrow (1999) argued for a positive PRTP. In his critique of the *Stern Review*, Nordhaus (2007) disagrees with the use of a near-zero PRTP. To illustrate the effect of different discount rates, he utilizes a PRTP of 1.5 percent to better match his preferred observed market interest rates.

In practice, the appropriate PRTP is likely somewhere between the extremes. A PRTP of 0.5 percent, for example, would suggest that we value the well-being of people in 50 years close to 80 percent of our own well-being; the well-being of people in 300 years would be valued at 20 percent. Whether this represents an acceptable compromise between undue sacrifice and indifference toward the future is a normative judgment.

Inequality aversion According to the Ramsey formula, the discount rate also depends on expected future economic growth and the relative wealth of future generations. Stated simply, as people become wealthier, an additional dollar contributes less to their well-being. In the context of climate change, suppose a ton of carbon causes $100 in undiscounted future damages. That would imply that we should be willing to pay $100 today to avoid those damages. But if future generations are expected to be wealthier than we are, then the avoided $100 in damages is worth less to them than the $100 we spend today. This suggests that we should be willing to pay less than the full amount of undiscounted damages. How much less depends on the assumed level of inequality aversion.

This conception relies on the assumption that future people will be richer. This is a plausible assumption given historical growth trends, but it is not guaranteed. The Ramsey framework accommodates this uncertainty. A low growth rate results in a lower relative discount rate. In the worst case, where economic growth is negative and future people are worse off, the discount rate could even be negative, implying that their well-being should be weighted more heavily than our own.

Still, most economists operate under the assumption of positive growth. If that's the case, the question of climate mitigation becomes one of redistribution. As Schelling (1995) noted, investing resources in climate mitigation

> now out of our own incomes will benefit people in the future who are expected to be better off than we are—an unaccustomed direction for redistributing income! Both within countries and among countries, we expect civilized governments to redistribute toward the poorer countries and toward the poorer elements of their own populations.

Schelling suggests that we should value damages according to the relative wealth of future generations.

But relative wealth is not the only consideration. Causing harm to others, even if they are wealthier, creates moral questions. If Mexico were to build factories along the US border that polluted Texas, would it be acceptable to argue that no compensation is owed simply because Americans are richer on average? Revesz and Shahabian (2011) additionally argue that discounting based on growth is problematic for two reasons: First, aggregate trends in global wealth may obscure important distributional effects. Even if future global average wealth is higher, those suffering climate damages in the future may still be poorer than the people paying for mitigation today. If mitigation costs are mostly paid by developed countries like the United States and damages primarily fall on low-income countries like Bangladesh, then climate policies may still transfer wealth from the rich to the poor. Second, they argue that discounting based on future growth assumes that the marginal damages will matter less to wealthier people. However, if environmental quality is valued more highly the richer people get, then discounting for economic growth is misguided.

In practice, however, most estimates of the SCC do discount based on growth. That is, higher growth implies future damages matter less to the present. The degree to which future damages are discounted depends on the inequality aversion parameter in the Ramsey formula. Similar to the PRTP, extreme values of inequality aversion can produce questionable implications. The *Stern Review* adopted an inequality aversion of 1, which drew criticism. Dasgupta (2007), for instance, agreed with the *Review*'s choice of a near-zero PRTP but argued that the inequality aversion assumption imposes unrealistic burdens on the present. He calculated that this parameter combination implies that we should save 97.5 percent of aggregate output for the future.

Though other commentators have noted that Dasgupta's critique assumes no technological progress, his broader concern remains: An inequality aversion of 1 implies that a proportional change in consumption carries the same ethical weight regardless of a person's wealth. To illustrate the implications, we can build on an example used by Dasgupta. He imagines two individuals, one consuming $360 per year, less than $1 per day, and another consuming $36,000 per year, slightly below the median after-tax income in the United States and higher than in most European countries. With an inequality aversion of 1, a $3.60 (1 percent) loss to the poorer person is treated as ethically equivalent to a $360 (1 percent) loss to the richer person. In practical terms, this means one of the world's poorest people going without food for several days is weighed the same as an average person in a wealthy country being unable to buy an Xbox.

Higher levels of inequality aversion imply more weighting toward the poor. At an inequality aversion of 2, Dasgupta finds that 1 percent loss to the poor person is equivalent to a 50 percent loss to the richer one. An inequality aversion of 3 is equivalent to a 93 percent loss. Most people, he argues, would find these implied tradeoffs unreasonable.

However, these extreme implications stem from assuming a 100-to-1 wealth gap. If we instead assume the richer person is 10 times wealthier (roughly the increase in global per capita GDP since 1870), then the implied tradeoff is less extreme. At inequality aversions of 2 and 3, the 1 percent loss to the poor person equals a 9 percent or a 43 percent loss to the rich person, respectively. At a 10-to-1 consumption ratio, the tradeoffs may be considered more reasonable.

In other words, the ethical justification for any given level of inequality depends on both a judgment about the tradeoff between rich and poor people and the expected relative wealth of future generations. Based on the implied savings rates of different levels of inequality aversion, Dasgupta argued that values in the range of 2 to 3, and perhaps higher, are warranted. A broader literature suggests values between 0.5 and 4 are within a reasonable range. Ultimately, inequality aversion requires an ethical decision about how to value the well-being of people with different relative levels of wealth across both space and time.

HOW ETHICAL CHOICES CHANGE THE SCC

Given these open ethical questions, how do different parameter values affect the estimated social cost of carbon? This question can be explored using the GIVE model, one IAM underlying the Biden administration's updated SCC estimates in 2023. Developed using an open-source framework, GIVE facilitates replication and testing alternative assumptions.

GIVE was created in response to shortcomings of an earlier generation of IAMs, including the DICE model developed by Nordhaus and the Policy Analysis of Greenhouse Effect (PAGE) model used in the *Stern Review*. NASEM (2017) identified areas for improvement in these models. Researchers at RFF and the University of California, Berkeley, developed GIVE to incorporate these insights, as described in Rennert et al. (2022).

The model's primary innovation is its explicit treatment of uncertainty in key assumptions. For example, while DICE relies on single forecasts of expected future economic growth, GIVE models economic growth as a distribution of possible trajectories based on expert judgment from 10 leading economists. Other inputs—including population growth, emissions, climate sensitivity, and economic damage functions—are similarly treated as probabilistic. Rather than relying on point estimates, GIVE runs 10,000 simulations using different combinations of inputs sampled from these distributions, a method known as Monte Carlo simulation. This allows the researchers to analyze how uncertainty in the core assumptions translates into uncertainty in the estimated SCC.

This method still requires that the researchers make crucial assumptions about the acceptable range of the key variables and the relationships between them. For this analysis, however, we take all non-discounting assumptions in GIVE as given so that we can focus on the ethical judgments embedded in discounting.

To estimate the SCC, the GIVE model is run twice. The first run uses global population and economic growth projections to forecast a baseline path of greenhouse gas emissions. These emissions are then translated into physical effects, such as temperature change and sea level rise, which are then converted into annual monetized damages. The second run adds a one-ton pulse of carbon emissions in 2020 and calculates the additional damages in each year through 2300. (To maintain consistency with the estimates reported by Rennert et al., we retain 2020 as the starting year of the simulations.)

Figure 1 displays the median annual projected gross domestic product per capita and marginal damages across the thousands of trial runs. It is important to emphasize that estimated damages are determined by the forecast economic growth: The same physical effects result in higher damages in wealthier futures simply because the world is richer. For example, if sea levels rise by the same amount in two different futures, the cost of protecting Miami will be higher in the scenario with higher per capita GDP because the assets at risk are more valuable and adaptation measures (such as construction and labor) are more expensive.

The final step of the model is discounting. The marginal damages in each year are converted to their 2020 value, and the SCC is the total of these discounted future damages. GIVE differs from earlier models by applying a variable discount rate. While models like DICE apply a constant discount rate to a fixed growth path, GIVE links its discount rate to each simulation's growth trajectory using the Ramsey formula. This means marginal damages in a given year are discounted at a

Figure 1

Median per Capita Consumption and Marginal Damages, 2020–2300

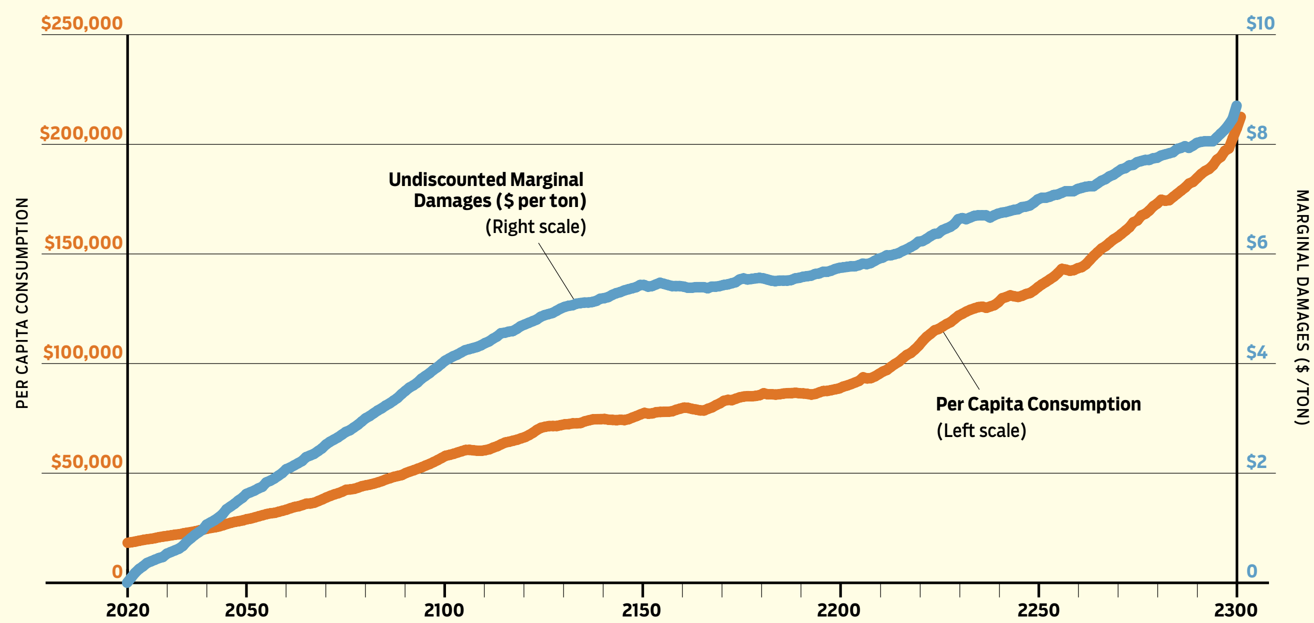

Source: Authors calculations using GIVE model.

rate determined by the simulated growth rate in that year and the assumed values for PRTP and inequality aversion.

This, however, is where ethical judgments are incorporated. Rennert et al. (2022) use a descriptive approach. They choose a "near-term" discount rate based on the observed interest rates of 10-year US Treasury bonds (their preferred near-term rate is 2 percent, but they also report the results when using 1.5, 2.5, and 3 percent). The values of PRTP and inequality aversion are then calibrated according to this near-term rate and projections of long-term growth.

This methodology has a scientific pretense, but the values of PRTP and inequality aversion still depend on the assumption that observed near-term rates reflect society's ethical preferences toward intergenerational tradeoffs. In general, while GIVE incorporates methods to account for uncertainty and risk (such as "stochastic" discounting and declining discount rates), these do not eliminate the need for normative judgments. They can illustrate how uncertainty shapes the consequences of ethical choices, but they cannot replace the ethical decisions themselves.

The ethical implications of different parameter values are embedded in the estimated SCC even if the descriptive approach obscures them. As Greenstone (2021) observes in a comment on the Rennert et al. methodology:

> They ignore the available evidence of the values of these parameters and instead use observed interest rates to govern the choice of the parameters. This creates an inconsistency between their approach and the large body of literature that has, for example, estimated values for η that range from 1 to 4 but are generally centered around 2.... So this aspect of their approach has practical appeal, but it is not built on an especially solid foundation of evidence.

This raises a central question: How does the SCC change when PRTP and inequality aversion are varied to reflect alternative ethical views?

To answer this, we replicate GIVE using the baseline assumptions from the Rennert et al. 2 percent and 3 percent near-term discount rate scenarios. The 2 percent case uses a PRTP of 0.2 percent and an inequality aversion of 1.24, resulting in an SCC of $183 per ton. The 3 percent case uses a PRTP of 0.8 percent and an inequality aversion of 1.57, producing an SCC of $78 per ton.

We then vary the PRTP and inequality aversion across a range of values cited in the literature. Lower values (i.e., PRTP of 0.1 percent and inequality aversion of 1) reflect the assumptions of the *Stern Review*, whereas the higher end values are closer to alternative assumptions proposed by Nordhaus, Dasgupta, and others.

Table 1 reports the results of GIVE using different combinations of PRTP and inequality aversion. As expected, the estimated SCC is highly sensitive to the assumed discounting parameters. Across plausible parameter combinations, the SCC ranges from $39 to $279 per ton, a roughly seven-fold difference.

What do these parameter values imply about the intergen-

ENVIRONMENT

erational tradeoffs of climate policy? One way to assess this is by comparing the SCC to the total *undiscounted* marginal damages caused by a ton of carbon. But this comparison is misleading for two reasons. First, because the damages are a function of economic growth, in trial runs with very high-growth, future damages become extremely large in absolute dollar terms simply because societies are wealthier. While GIVE's variable discount rate accounts for this by applying higher discounting to high-growth futures (in line with the Ramsey formula), the undiscounted damages do not reflect this adjustment and therefore place disproportionate weight on extreme growth scenarios. Second, welfare economics suggests we should care more about the proportional changes in well-being, not the absolute dollar losses. The median undiscounted marginal damage in the year 2300 from a ton of carbon emitted in 2020 is about $9, yet the median projected global GDP per capita in that year is roughly $200,000, over 11 times the 2020 level of about $18,000 per capita. Adjusting for the difference in overall well-being, the equivalent damage is much smaller.

To provide a more meaningful benchmark, we scale the annual damages for projected economic growth. In other words, the marginal damage (*MD*) in each year (*t*) is adjusted according to the proportional change in per capita GDP (*C*) between that year and 2020 ($\text{Growth-adjusted } MD_t = MD_t \times \frac{C_{2020}}{C_t}$). After this adjustment, the median damages in 2300 are equivalent to 66¢ of 2020-level income. Figure 2 illustrates the comparison between the undiscounted damages and the growth-adjusted damages over time.

Table 1

SCC Estimated by GIVE Using Alternative Assumptions for Discounting Parameters

2020 $ per ton of carbon

	Inequality aversion		
PRTP	η =1	2	3
ρ = 0.1 percent	279	114	88
0.5 percent	171	77	59
1 percent	104	51	39

Source: Authors' calculations using the GIVE model and assumptions described by Rennert et al. (2022).

Summing these growth-adjusted damages across all years yields a total of $319 per ton of carbon emissions, representing the expected long-run economic harm in constant welfare terms of emitting a ton of carbon today. The central question of climate policy is: How much of this future damage should we charge people today? The answer depends on how we weigh the future against the present and could be as high as the full amount of damages, $319 per ton. IAMs make a choice about this tradeoff through their selection of the discount rate.

For GIVE, this choice can be illustrated by comparing the estimated SCC to the growth-adjusted benchmark. Figure 3 shows the median growth-adjusted damages alongside the same stream of damages discounted using the ethical param-

Figure 2

Median Undiscounted and Growth-Adjusted Annual Marginal Damages, 2020–2300

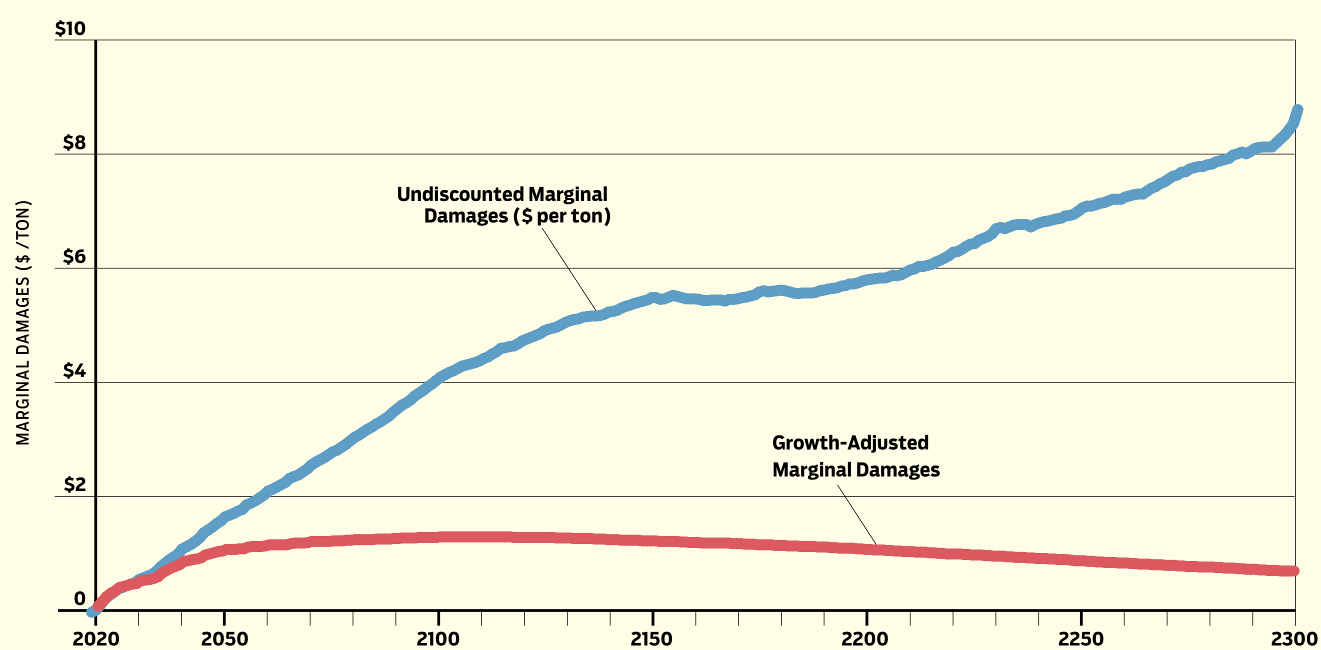

Source: Authors calculations using GIVE model.

Figure 3

Median Growth-Adjusted Damages and Rennert et al. Discounted Marginal Damages, 2020–2300

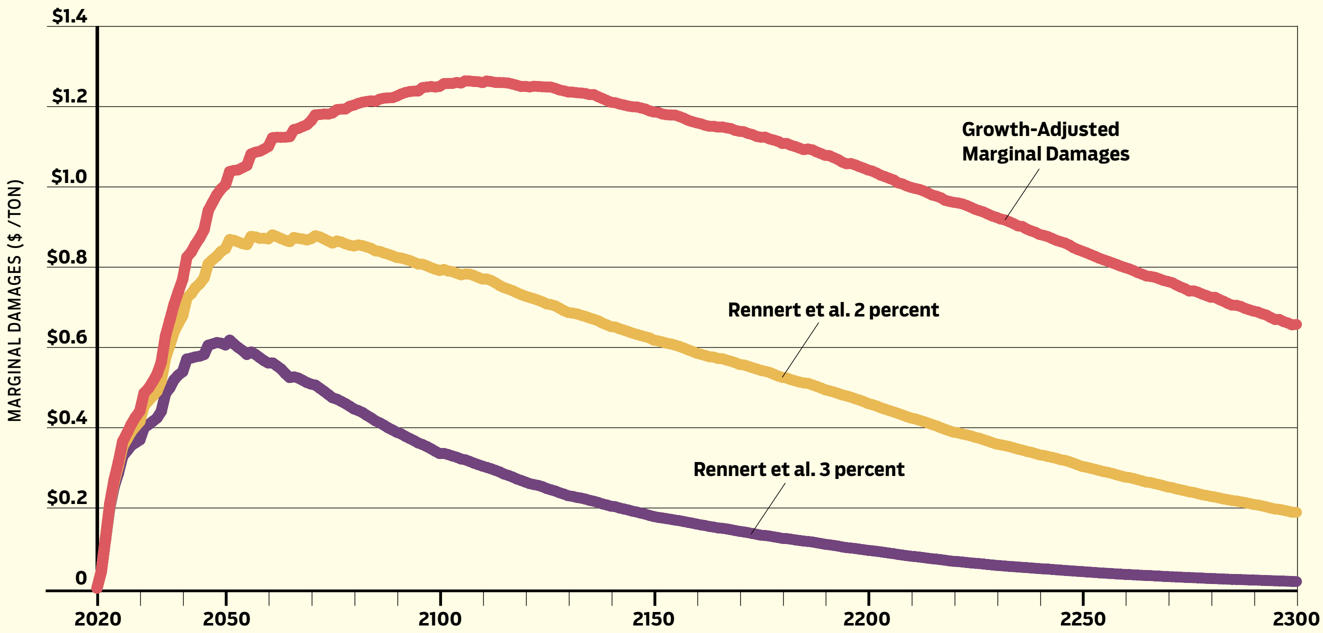

Source: Authors calculations using GIVE model.

eters assumed by the Rennert et al. 2 and 3 percent scenarios. The discounted lines show how much of the future harm we are willing to internalize today. Across all 280 years, the 2 percent case's SCC of $183 represents 57 percent of the growth-adjusted total. In the 3 percent case, the $78 SCC represents 24 percent.

IMPLICATIONS FOR CLIMATE POLICY

What does this comparison suggest for deriving discount rates? Obscuring the ethical dimensions of discounting does not remove them. The fact that Rennert et al. calibrate their 2 and 3 percent near-term discount rates from observations does not make the underlying ethical judgments any more or less appropriate.

In practice, however, empirical discount rates are used as a cudgel against any disagreement on discounting. For example, in response to the updated Biden administration SCC, Heritage Foundation researchers explored the sensitivity of the GIVE model to higher discount rates (Frei et al., 2025). They justified their choices in descriptivist terms, citing the opportunity cost of climate investment at higher, observed market rates. While we disagree with their reasoning, it is as legitimate as the reasoning underpinning the GIVE model and the Biden administration SCC increase. RFF researchers Prest and Rennert (2025), however, dismissed the much lower Frei et al. results from higher discount rates out-of-hand. Prest and Rennert allege that the "inappropriately high discount rates" (e.g., 5 percent) are "out of line with the economic literature and market interest rates."

This line of argument misses the deeper point. The market rates and economic "consensus" Prest and Rennert cite are only tangentially related to the ethical questions at hand. The central truth of discounting remains: From an economic perspective, any given discount rate is as legitimate as any other. As Arrow et al. (1996) wrote some three decades ago, "Although economists can make no special claim to professional expertise in questions of ethics, they have developed rigorous methods for analyzing the implications of ethical judgments." Models like GIVE can clarify how different ethical choices affect the SCC, but they and other empirical methods cannot resolve which ethical values are right.

Conversations about the ethical implications of climate policies should happen. To its credit, the *Stern Review* made its ethical assumptions explicit and created public debate. The fact that, nearly two decades later, we have yet to reach consensus is not a reason to bury the ethical issues. Instead, an analysis that relies on much lower or higher discount rates should confront these issues: Why is a discount rate of 2 percent and an SCC of $183 per ton more ethically acceptable than a discount rate of 3 percent and an SCC of $78?

Given the stakes, individuals may reasonably come to different conclusions about the ethically appropriate discount rate. The larger difficulty is in aggregating these ethical preferences. Descriptivists are right to criticize prescriptive approaches as undemocratic. The answer is not to place responsibility for determining our ethical obligations in the hands of experts. But despite its appeal of relying on observed behavior, descriptivist discount rates are just as misguided because these observations do not address the underlying ethics.

So how should society decide the right way to value the future? Beckerman and Hepburn (2007) mention other possible avenues for determining discount rates: "There are a range of intermediate approaches including the use of stated preference surveys, behavioural experiments, and methods to reveal the social preferences inherent in our social institutions." Yet, none of these methods has emerged as a replacement for the descriptive or prescriptive approaches. While some work has explored how differing normative perspectives might be aggregated, any such approach to aggregating ultimately depends on knowing which ethical preferences to include and how to weight them.

In the meantime, social preferences are likely expressed indirectly. To the extent that the 2024 presidential election reflected the public's stance on the tradeoffs of climate policy, voters implicitly chose a high discount rate and minimal policy over the alternative. A presidential election is, of course, a blunt way to resolve ethical questions, but it is arguably no worse at aligning climate policy with public values than selective use of market rates or relying on technocrats' ethics.

CONCLUSION

The ethical questions raised by discounting are not a flaw in climate policy; they are the fundamental issue of the climate change debate. The consequences of today's emissions will be felt most by people who have not yet been born, while the benefits of fossil fuel use are enjoyed by people who are alive today. Determining how much present generations should sacrifice to reduce harm to future generations is not a scientific calculation; it is a moral one.

Why, after decades of debate, is there no consensus about how to discount the future? Because there is no universally accepted answer to the ethical questions climate change poses. Reasonable people can disagree about how to value future generations or weigh the interests of the rich and poor. While integrated assessment models can help clarify how different ethical assumptions affect tradeoffs between future and present well-being, they cannot choose the appropriate tradeoff.

We should not obscure these ethical dimensions by ignoring the problem or pretending that science alone can tell us the right answer. If our actions impose harm on others, whether they are alive now or in the distant future, we have an obligation to reckon with that harm. But the ethical reasoning should be made explicit. Hiding it behind a veneer of scientific objectivity does not mean the ethical implications disappear. Any argument for a particular course of action, whether advocating for aggressive mitigation or a more restrained approach, should openly discuss the ethical consequences of the choices it endorses. R

READINGS

- Arrow, Kenneth J., 1999, "Discounting, Morality, and Gaming," in *Discounting and Intergenerational Equity*, Peter R. Portney and John P. Weyant, eds., Routledge, pp. 13–21.
- Arrow, Kenneth J., et al., 1996, "Intertemporal Equity, Discounting, and Economic Efficiency," in *Climate Change 1995: Economic and Social Dimensions of Climate Change, Contribution of Working Group III to the Second Assessment Report of the Intergovernmental Panel on Climate Change*, James P. Bruce et al., eds., Cambridge University Press, pp. 125–144.
- Barrage, Lint, and William Nordhaus, 2024, "Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model," *Proceedings of the National Academy of Sciences* 121(13): e2312030121.
- Beckerman, Wilfred, and Cameron Hepburn, 2007, "Ethics of the Discount Rate in the Stern Review on the Economics of Climate Change," *World Economics* 8(1): 187–210.
- Carleton, Tamma, and Michael Greenstone, 2022, "A Guide to Updating the US Government's Social Cost of Carbon," *Review of Environmental Economics and Policy* 16(2): 196–218.
- Cherrier, Beatrice, and Pedro Garcia Duarte, 2024, "How the 'Ramsey Formula' Came to Define Time Discounting in Economics (1950–2000)," SSRN, September 1.
- Dasgupta, Partha, 2007, "Commentary: The Stern Review's Economics of Climate Change," *National Institute Economic Review* 199(1): 4–7.
- Frei, Alexander, et al., 2025, "Calculating the 'Social Cost of Carbon' with the GIVE Model: An EPA Model Not Ready for Prime Time," Special Report no. 308, Heritage Foundation, January 24.
- Friedman, Lisa, 2025, "What's the Cost to Society of Pollution? Trump Says Zero," *New York Times*, May 10.
- Gollier, Christian, and James K. Hammitt, 2014, "The Long-Run Discount Rate Controversy," *Annual Review of Resource Economics* 6(1): 273–295.
- Greenstone, Michael, 2021, "Comment on Kevin Rennert et al., 'The Social Cost of Carbon: Advances in Long-Term Probabilistic Projections of Population, GDP, Emissions, and Discount Rates,'" *Brookings Papers on Economic Activity*, Fall 2021: 276–294.
- Koopmans, Tjalling C., 1960, "Stationary Ordinal Utility and Impatience," *Econometrica* 28(2): 287–309.
- Koopmans, Tjalling C., et al., 1964, "Stationary Utility and Time Perspective," *Econometrica* 32(1–2): 82–100.
- NASEM (National Academies of Sciences, Engineering, and Medicine), 2017, *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide*, National Academies Press.
- Nordhaus, William D., 2007, "A Review of the Stern Review on the Economics of Climate Change," *Journal of Economic Literature* 45(3): 686–702.
- Prest, Brian C., and Kevin Rennert, 2025, "When Put to the Test, Higher Social Cost of Carbon Stands Firm," Resources for the Future, February 18.
- Rennert, Kevin, et al., 2022, "Comprehensive Evidence Implies a Higher Social Cost of CO_2," *Nature* 610(7933): 687–692.
- Revesz, Richard L., and Matthew R. Shahabian, 2011, "Climate Change and Future Generations," *Southern California Law Review* 84(4): 1097–1163.
- Schelling, Thomas, 1995, "Intergeneration Discounting," *Energy Policy* 23(4–5): 395–401.
- Stern, Nicholas, 2006, *Stern Review on the Economics of Climate Change*, Cambridge University Press.
- Weisbach, David A., and Cass R. Sunstein, 2008, "Climate Change and Discounting the Future: A Guide for the Perplexed," Reg-Markets Center Working Paper no. 08-19, August 2008.

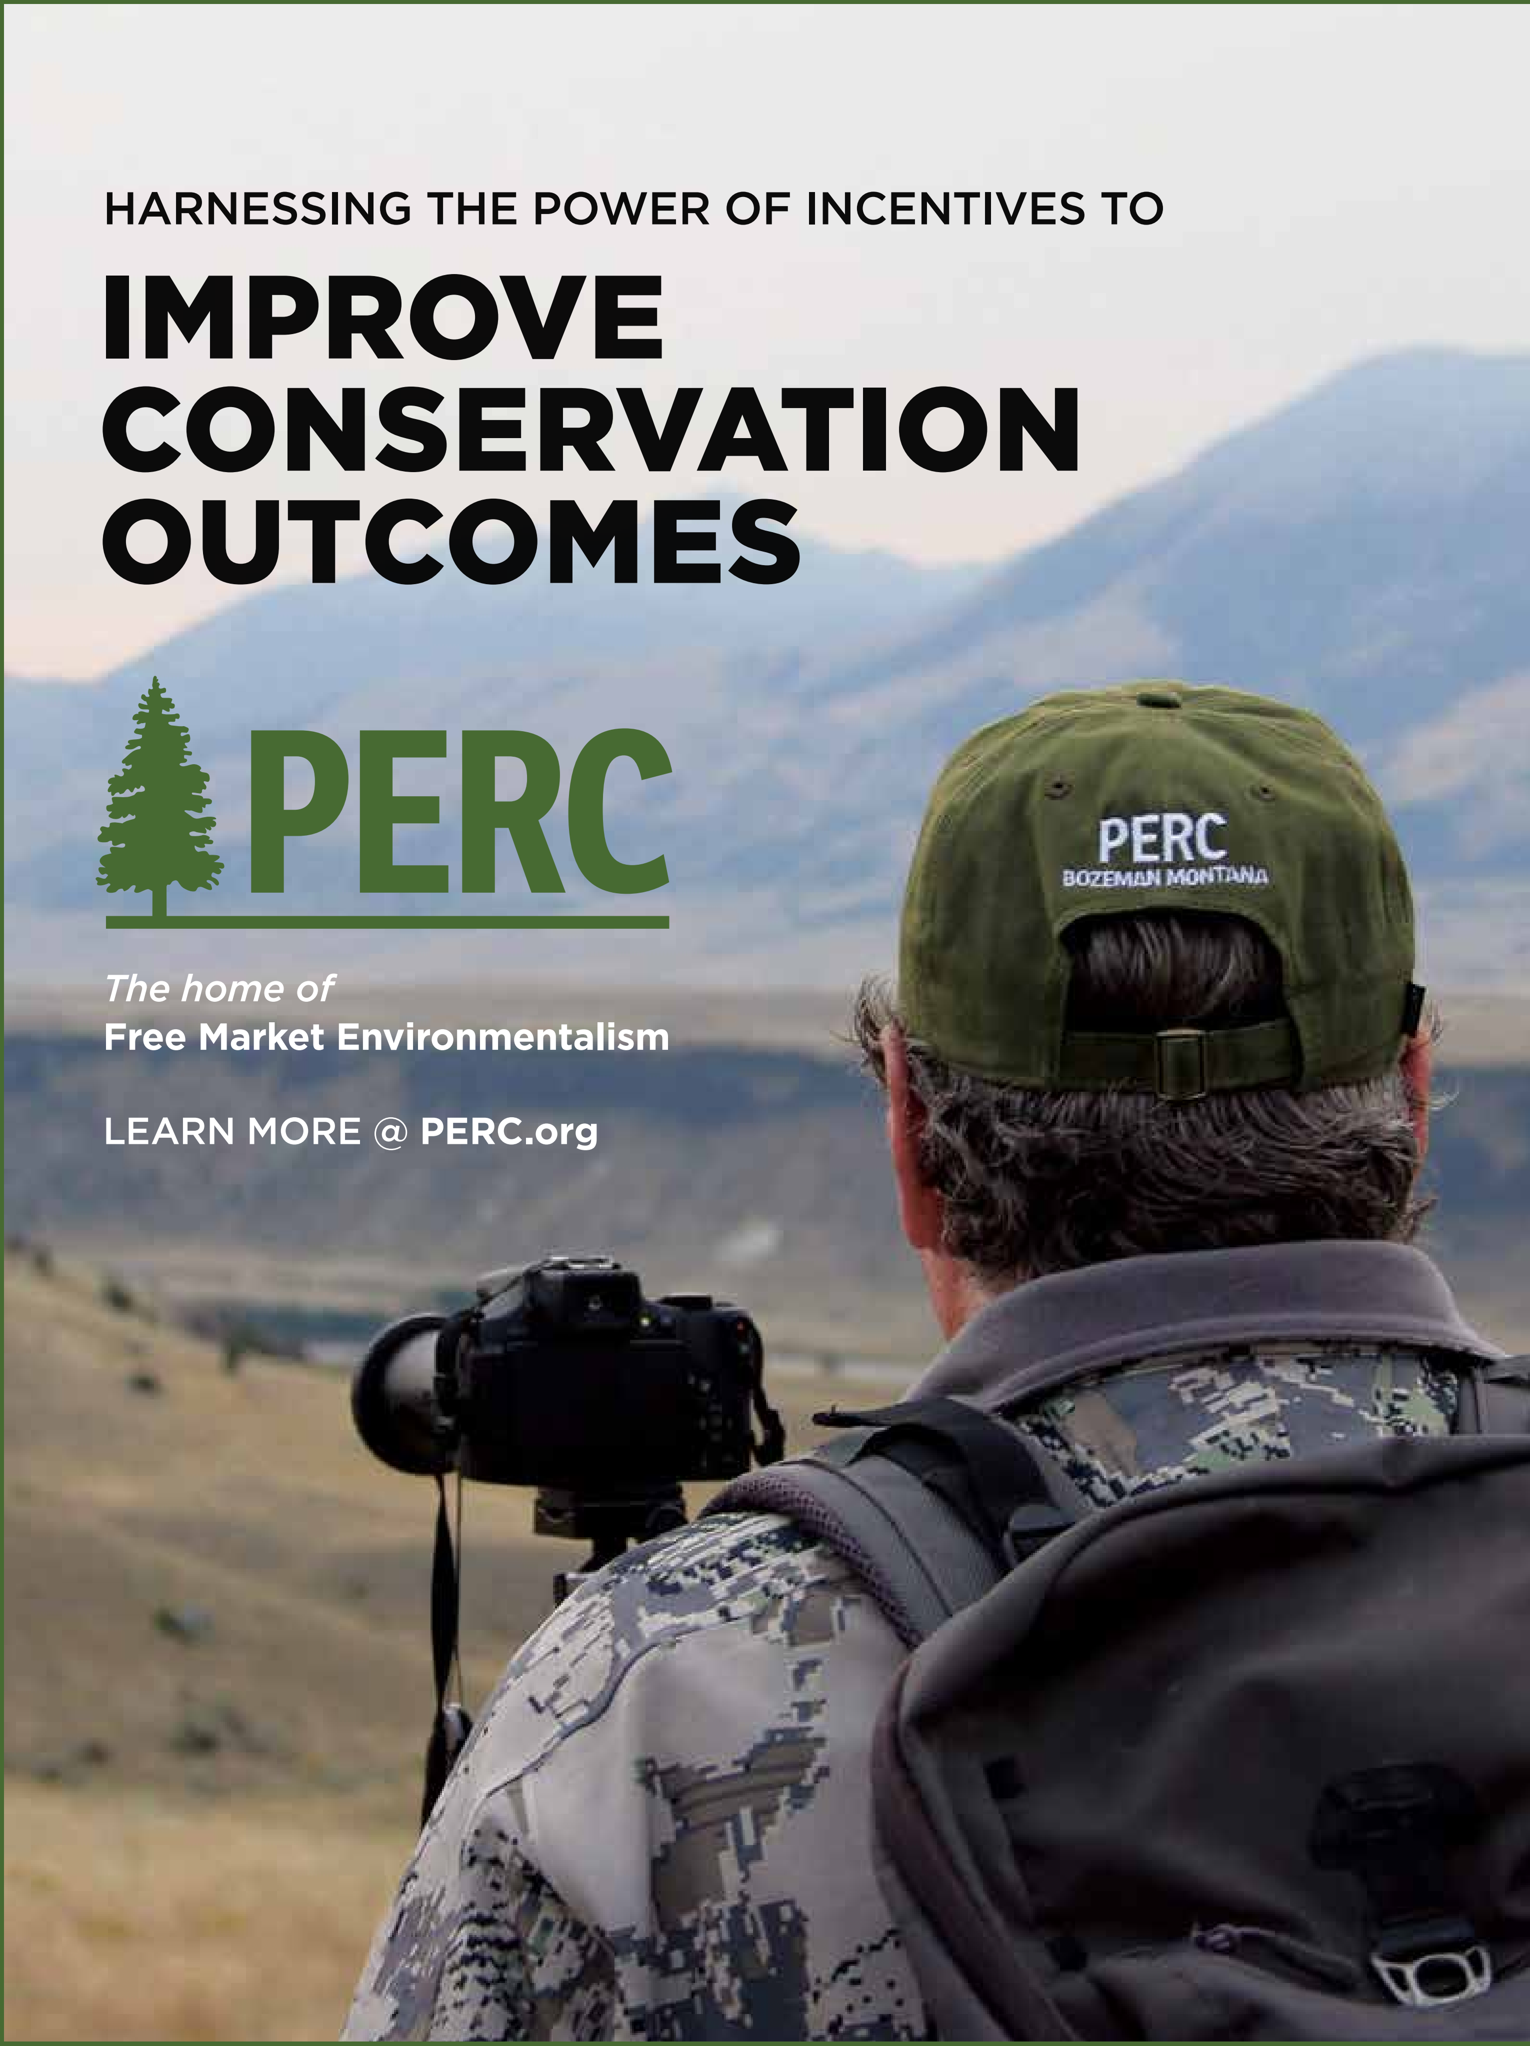
HARNESSING THE POWER OF INCENTIVES TO
IMPROVE CONSERVATION OUTCOMES
PERC
The home of Free Market Environmentalism
LEARN MORE @ PERC.org
PERC
BOZEMAN MONTANA