Chapter 11
Intergenerational Justice in Protective and Resilience Investments with Uncertain Future Preferences and Resources

Louis Anthony (Tony) Cox Jr. and Emeline D. Cox

Abstract How much should each generation invest in building resilient infrastructure to protect against possible future natural disasters? If such disasters are infrequent, members of each generation may be tempted to defer investments in resilience and protective infrastructure (e.g., in building or improving dams and levees; retrofitting office and residential buildings; creating more robust transportation, power, and communications networks; etc.) in favor of consumption or growth. Succumbing to this temptation imposes risks on future generations of needlessly large losses or disproportionate need to invest in resilience. Yet, even the most dutiful and altruistic present generation has limited obligations to invest to protect future ones, especially if present investments in resilience reduce growth and future prosperity, or if the preferences, priorities, resources, and capabilities of future generations are highly uncertain. This paper discusses several different frameworks for clarifying how much each generation should invest in protection. Optimal economic growth models provide a well-developed technical framework for maximizing average or minimal expected social utility over time, but require consistency and cooperation over time that may not be psychologically or politically realistic. If investment decisions are viewed as a form of dynamic "dictator game" in which earlier generations choose how to allocate benefits between themselves and later generations, then insights from behavioral economics, risk psychology, and moral psychology suggest cues related to deservingness and trustworthiness that powerfully affect what is perceived as fair and right in such settings. A Rawlsian concept of justice (what investment decision rules would people choose from behind a veil of ignorance, in which no one knew what generation he or she would be born into?) solves the problems of over-discounting long-delayed and uncertain consequences that have frustrated some previous efforts to apply cost-benefit analysis to ethically charged issues involving intergenerational justice. We suggest several principles for applying insights from these different frameworks to investments in building resilient communities and mitigating natural disaster risks across generations.

L. Anthony (Tony) Cox Jr. (✉) • E.D. Cox
Cox Associates, 503 Franklin Street, Denver, CO, USA
e-mail: tcoxdener@aol.com

P. Gardoni et al. (eds.), *Risk Analysis of Natural Hazards*, Risk, Governance and Society 19, DOI 10.1007/978-3-319-22126-7_11

11.1 Introduction: How Much Care Does Each Generation Owe to Future Ones?

Each of us lives with decisions that our former selves have made—choices about what to read and study, whom to marry, where to live, what to work at, how hard to work, how much to consume, and what to do with any savings. Our satisfaction with the choices made by our prior selves may be mixed, but it is not unusual to find that our own current choices are only modestly responsive to the imagined preferences and evaluations of our future selves (Kahneman 2011). And so we may let slide New Year's resolutions that our future selves will predictably wish we had abided by; spend more on current consumption and less on savings or prudent investments or retirement plans than our censorious future selves will, predictably, deem optimal; and succumb to tempting but risky prospects in the heat of the present moment knowing that, with high probability, we will soon regret them as imprudent choices that we will (predictably) wish we had made differently.

Researchers investigating individual preferences and choices over time and under uncertainty have long noted, and created analytic models of, the dynamic inconsistency of individual plans and intentions formed under realistic (hyperbolically discounted) preferences for future rewards (*ibid*). More usefully, they have developed a large array of devices for improving the time-consistency of present decisions to help us make present choices that are less regrettable to our future selves. These devices range from freezing credit cards in block of ice to changing defaults and nudges (e.g., to opt-in or opt-out of an employer-sponsored savings or retirement plans such as Save More Tomorrow) to encouraging "premortem" thinking about how well-intended plans and projects might come to be seen in retrospect as predictable failures. Such imposition of dynamic consistency and rationality constraints and effortful ("System 2") deliberation to restrain and modify the choices urged by immediate intuitive and impulsive ("System 1") responses to the opportunities and temptations before us is at the forefront of current efforts to improve real-world decision-making for individuals by better integrating and coordinating the recommendations from Systems 1 and 2 (Kahneman 2011). It also corresponds to venerable traditions in personal ethics that extol the virtues of temperance, patience, prudence, thrift, steadfastness, and the like.

Social decision-making, too, raises questions of how best to trade-off the needs and preferences of current vs. future agents (Oliver et al. 2014). In many cases, these agents include members of present and future generations, rather than only our own selves at earlier and later dates. Crucial practical risk management policies and social investment decisions depend on how we trade-off their interests. For example,

- How large a tsunami, hurricane, earthquake or other natural disaster should nuclear power plant designers, levee designers, sea wall builders, and other construction and infrastructure engineers design for? Following experiences with extreme events such as the Fukushima tsunami or Hurricane Sandy, it is common for news stories, editorials, and politicians to urge that more should have been done and should be done now to provide for facilities that would better withstand the stresses of wind and water under such extreme conditions. Hindsight bias may make such criticisms and recommendations inevitable. But should engineers henceforth design for the most extreme events that are expected to occur on average once every 50 years, or 100 years, or 1000 years, or for some other level of extremity? Designing to protect against more extreme (rarer) events typically costs more, but the future risk reduction benefits purchased by these present investments in safety may be very uncertain, depending in part on future climate and weather and perhaps technological advances.

- How much should we invest in developing and stockpiling antibiotics and other countermeasures in case of an anthrax (or other) outbreak? The preventive measures cost money now. Whether they will be used before they expire, and how much good they will do if they are used, are typically quite uncertain. If the decision problem is stationary (i.e., looks the same starting from any point on time), then it may be reduced to a choice between (A) Creating and maintaining some level of stockpiled countermeasures, costing a certain amount per year; or (B) Foregoing the costs and benefits of doing so. In either case, some people alive now and others not yet born may share in the costs, risks, and benefits resulting from whatever decision is made.
- *How much should we spend now in seeking to postpone or mitigate future losses from climate change?* Is it better to commit a substantial fraction of present GDP to climate change mitigation efforts, or to invest in current economic growth and technological and institutional progress, so that future generations will be better equipped to deal with any problems that materialize?
- *How much should we spend per year on nuclear fusion reactor research?* If there is substantial uncertainty about whether and when fusion-generated power production might become practical (e.g., at least producing more power than it consumes), then how much should the current generation invest in fusion R&D? Under what conditions (if any) should it stop investing in favor of more promising alternative uses of current resources? If future generations were allowed to vote on or have input to current decisions, how (if at all) would that change the answers?

These examples, and countless others in the realms of public finance and investments, consumption of non-renewable resources, costly exploration and discovery initiatives, storage of radioactive wastes, and even military investments in long-term security and readiness, illustrate the challenges and importance of learning to address preferences and trade-offs in decisions with consequences that affect multiple generations. All of them can be posed generically as decisions about *how much care (whether measured in prevention, investment, or other expenditures) current decision-makers should take in order to benefit future recipients of the benefits of these decisions*. Making such decisions wisely requires constructive frameworks for determining both *what* we should decide to do now, and *how* we should decide what to do now, given that some of those who bear the consequences may not be alive yet to participate in the decisions.

Optimizing trade-offs between present costs of taking care (or investing in risk reduction) and uncertain future risk-reduction benefits requires an approach to decision-making different from traditional subjective expected utility (SEU)-maximizing decision analysis. The fact that there are multiple stakeholders, not all of whom are alive at present, moves the decision problem from the domain of a single decision-maker to the less well-developed domain of multiple agents, for whom justice and cooperation over time in the absence of opportunities for explicit collaboration and agreement may be important. Even defining clearly what "optimize" means in such settings is difficult. At the individual level, what seems the best choice in retrospect may be very different from what seemed best in prospect, when decisions had to be made and when many possible futures were being considered. In retrospect, evaluations of alternatives can be powerfully affected by hindsight bias, regret, blame-seeking, availability bias, and other potent psychological influences on how we judge previous decisions (Kahneman 2011). In prospect, our evaluations and judgments about what course of action seems best are influenced by different biases, including over-optimism, narrow framing, and over-confidence (*ibid*). Defining a single concept of "optimal" choice that satisfies the biases of both forward-looking and backward-looking evaluations may be impossible. At the group level, what most people prefer to do now may differ predictably from what most of them (or future generations) will later wish had been done now—and from what they may praise or blame current policy makers for doing or failing to do. Such differences can arise because of predictable shifts in beliefs and information, or from changes in tastes and preferences, as well as because of predictable differences in prospective and retrospective preferences and evaluations of alternatives.

The following sections present several frameworks for thinking constructively about how much we should spend now to protect against possible future natural disasters by investing in more resilient infrastructures and communities, precautionary measures, sustainable production and consumption practices and processes, and other instances of costly present care taken to create potential future risk-reducing benefits that will affect multiple generations. To simplify and clarify key issues, we first introduce simple, idealized models of multi-generational conflict and cooperation in the gradual consumption of a desirable resource, generically called "capital stock" or, more colloquially, "pie." This can be variously interpreted to illuminate issues such as savings vs. investment in growth of capital stock in economic growth models; optimal consumption of a scarce resource that many generations would like to share (e.g., oil, or susceptibility of microbial pathogens to current antibiotics, or military reputation with allies and with enemies, etc.); investment in a stock of "safety" in infrastructure (e.g., fraction of bridges replaced or renewed no more than a target number of years ago, or height of a levee or sea wall to protect against flooding during extreme weather events); or investments in precaution against future disasters (e.g., stockpiling of antibiotics or other countermeasures). We consider how one might answer the ethical questions of how big a piece of the remaining pie each successive generation should take, and what duties each generation has to add to the stock of pie or capital, if doing so is possible but costly. This multi-generational sharing of a pie, and generalizations in which the size of the pie can change due to the production and consumption decisions of each generation, and perhaps to random events, provide fruitful metaphors for clarifying both economic and ethical aspects of intergenerational justice that arise in many practical applications, including decisions about how much each generation should spend to help reduce risks from natural disasters for future generations. We compare several principles and constructive analytic and ethical frameworks for deciding what each generation should do, addressing both what decisions should be made and how they should be made. Finally, the main generalizable insights and results from applying these frameworks to problems of intergenerational cooperation and justice are summarized, with attention to how they might be applied to improve practical decisions in which intergenerational consequences and justice are key concerns. In particular, we discuss implications for how to make ethically defensible investments in protecting against natural hazards, as well as implications for proposed principles of sustainable consumption and production of safety benefits.

11.2 Simple Models of Intergenerational Justice: Sharing a Pie Over Time

An idealized model of intergenerational cooperation and decision-making often used in economics, philosophy, and game theory discussions of intergenerational justice is a pie-division problem (e.g., Bahr and Requate 2014; Schmidtz 2011). At the start of the problem, the first generation is endowed with a pie of fixed size, conventionally normalized to be 1, or 100 %. Members of that generation must decide how much of the pie to consume themselves, and how much to leave as a bequest to the next generation. Each successive generation in turn must then decide how much of the endowment of pie that it received from previous generations it will consume, and how much it will pass on to later generations. In the simplest case, the pie does not grow or shrink, except by the consumption decisions of each generation. Of course, generations may overlap, but the simplest models treat them as discrete; this does not substantially affect the analysis.

In variations of the problem, the pie not consumed grows at some rate (analogous to the savings and investment rate in economic growth models); or additions to the stock of pie may occur at random times (analogous to discoveries of new deposits of a non-renewable and valuable resource such as oil); or the pie may be perishable or may randomly shrink or disappear (analogous to depletion or exhaustion of a resource, or to random destruction of property by disasters). In each variation, the key question of how much to consume now remains. For modeling the issues that arise in investing to protect against natural disasters, it is useful to allow two different types of investment for unconsumed pie: (a) Invest in growth (i.e., save some pie, which may then grow at a certain rate); or (b) Invest in safety, which reduces the amount of pie lost when disasters occur.

11.3 Analytic Frameworks for Managing Cooperation Across Generations

This section briefly introduces several different analytic frameworks from economics, game theory, sociology, and philosophy for addressing the question of how much each generation should consume, save, and invest, where investment could be either in growth (as in traditional economic models of optimal consumption and savings/investment over time) or in safety.

11.3.1 Economic Growth Models

For the past century, classical and neoclassical economic models of optimal capital accumulation and growth, such as the Ramsey (1928), Solow (1956), Swan (1956), and Phelps (1961) models, have been used to answer questions about how best to manage a society's consumption, savings, and investment over time. These models describe the changing stock of capital (or "pie") over time by differential equations in continuous time or by difference equations in discrete time, such as

$$K(t+1) = (1+g)(1-c)[1-d(t)]K(t) \tag{11.1}$$

where $K(t)$ is the stock of capital at the start of period t; g is the fractional growth rate for unconsumed capital (typically between about $g = 0.02$ and $g = 0.10$ per year); c is the consumption fraction [and therefore $(1-c)$ is the savings fraction, often denoted by s]; and $d(t)$ is the random fraction of capital lost in period t due to disasters. In classical economics, the capital good, K, is often interpreted as something like "corn" that can be either consumed, thereby producing immediate utility from consumption; or saved and invested ("planted"), in which case it will grow at rate g and expand future opportunities for consumption and savings/investment. In more elaborate models, the values of c and g may also vary over time. Traditional deterministic economic models of optimal consumption, savings, and growth do not model losses from disasters (in effect, assuming that all $d(t) = 0$), but do model labor supply and population growth, in addition to capital growth. Utility in period t is usually assumed to be based on the amount consumed then, $cK(t)$. A central planner who wishes to maximize the net present value of utility (or, in some variations, the steady-state sustainable level of utility) can solve a dynamic optimization problem, constrained by the differential or difference equations for capital accumulation (and for growth in the population and labor force, if they are modeled), to obtain recommended levels of consumption, savings, and investment for each period (Ramsey 1928). Extensions to stochastic growth models, in which the growth rate of unconsumed capital, g, varies randomly over time, have been developed and analyzed using techniques such as stochastic dynamic programming to illuminate the trade-offs between present and future consumption implied by different growth rates (or distributions of growth rates, if they are random) and polices (Olson 2005).

Useful qualitative insights flow from such economic models, which characterize the set of possibilities for intergenerational choices and their probable consequences over time. Models of capital consumption and accumulation specify the physics, or rules of the game, within which intergenerational sharing of resources, consumption, and production take place. They also reveal qualitative properties of optimal policies. For example, analysis of growth models reveals conditions under which all optimal growth trajectories (i.e., sequences of states, consisting of consumption and savings/investment levels in each period, that jointly maximize discounted utility or similar objective functions) approach each other and stay close to each other along most of their lengths, regardless of the initial state. Such optimal growth path "turnpikes" exist for many deterministic and stochastic growth models. In simple models, the optimal policies have simple and intuitively appealing economic interpretations, with each period's consumption and investment being optimized by adjusting consumption levels to equate the marginal utility from further current consumption to the discounted expected marginal value of further investment (i.e., the product of the marginal productivity of investment and the marginal utility from consuming the incremental output next period) (*ibid*).

Perhaps most interestingly, stochastic growth models reveal that there may be a critical threshold level of the initial capital stock above which it is guaranteed (with probability 1) that the optimal policy will never exhaust the capital stock, but below which there is a risk (or, if the growth rate g is variable enough, a certainty) that even optimal management will eventually end with $K = 0$ (collapse). Such results have implications for the management of fisheries and other renewable resources. If extinction is a realistic possibility, then keeping the stock above the threshold needed to avoid extinction might well be an overriding priority, trumping all other considerations about what each generation owes to future ones. Including the possibility of disasters in growth models [$d(t) > 0$] can modify optimal policies, e.g., by providing reason to restrict current consumption to provide an adequate margin of safety. For a non-renewable resource (i.e., growth rate $g = 0$), the optimal policy for sharing the initial stock among generations may shift toward more consumption earlier on if disaster might destroy some or all of the remaining stock. Whether such generalizations hold depends on details of the model considered, such as whether utility of consumption exhibits increasing or decreasing marginal utility (or some of each, perhaps being S-shaped, e.g., if consuming very little oil per generation has zero or little marginal utility compared to consuming more, but consuming a lot also generates less utility per barrel than consuming less). Likewise, incorporating disaster mitigation opportunities into a detailed model requires specifying the cost curve or technology possibilities for spending $K(t)$ to reduce (shift leftward the distribution of) $d(t)$, and how long the effects of such expenditures last. For example, investing in higher sea walls or levees consumes more of the current capital stock that might otherwise have been spent on consumption or invested in economic growth, but reduces the probable losses from floods during the useful life of the walls or levees. Understanding the relation between present costs and future benefits, as modeled by the leftward shift in the loss terms $d(t)$ in future periods purchased by a current investment in safety, provides the essential technical information about possible costs and benefits of disaster risk reduction needed to decide what to do in a multi-period optimization model.

If the model includes population sizes and labor forces, and if a value function for reducing lives lost in the event of a disaster is included in the objective function (thus inviting the usual vexed questions about how to value statistical lives saved), then economic optimization models can deliver recommendations for consumption and investments (in growth and safety) in each period that take into account this evaluation of lives saved. The effects on optimized current consumption of allowing for potential disasters depend on model details; in various specific models, they range from consuming more now ("Get it while it lasts") to consuming less now in order to protect and expand future opportunities ("Safety first," "Make hay while the sun shines," i.e., invest while it is still productive to do so). If the objective function is smooth and exhibits diminishing marginal returns, then optimizing multi-period consumption, investment, and savings (e.g., via stochastic dynamic programming) typically requires equating the marginal returns from incremental expenditures of $K(t)$ on present consumption, on savings and capital growth, and on disaster mitigation, assuming that future decisions will also be made optimally.

For our purposes of comparing different frameworks for making intergenerational consumption, growth, and safety (i.e., disaster mitigation) investment decisions, the point of including the random disaster term $d(t)$ is not to study optimal growth policies in specific models, but only to point out that standard economic methods for studying optimal consumption and investment over time in models of capital and growth can easily be modified to investigate how the possibility of capital-destroying disasters—and of safety investments that mitigate them—changes optimal policies. Overall, the possibility of investing in precautions that stochastically reduce the damage done by disasters (sacrificing some present consumption or investment in growth by instead spending some of $K(t)$ to shift the distribution of $d(t)$ leftward, toward smaller values) provides an alternative to savings and investment as a way to increase capital stock and production and consumption possibilities over time.

In this framework, concerns for intergenerational justice are addressed implicitly, by making decisions in each period, or for each generation, to maximize the objective function (e.g., expected discounted utility) for all. A fundamental limitation of all such models is that no single objective function may correctly represent the preferences of different generations, or even of social planners living at different times. Not only might future tastes and preferences for consumption vs. safety trade-offs and future societal attitudes toward accepting or mitigating disaster risks differ from present ones, but also future generations whose wellbeing is discounted in present objective functions might wish that a different objective function had been optimized. (Worse, if current choices about wealth vs. safety affect the existence or sizes of future generations, then the hypothetical preferences of potential future individuals might be considered to matter in deciding what should be done now. But the hypothetical preferences of as-yet non-existent individuals provides at best a speculative basis for making present choices.) Allowing different generations to have, and to act on, different objective functions from the current generation's requires shifting our analytic perspective from multi-period economic optimization to game-theory to better understand how the choices of different generations interact over time.

11.3.2 Behavioral Game Theory Framework for Strategic Interactions Among Generations

Game-theoretic frameworks for deciding how much of its inherited stock of goods each generation should allocate to consumption and how much it should bequeath to its successors via investments in disaster mitigation measures and economic growth differ radically from optimal economic growth models (D'Albis and Ambech 2010; Balbus and Nowak 2008). They drop the fiction of a single dispassionate social planner who is willing and able to make decisions for each generation to maximize some overall objective function. The difference can be illustrated as follows. Suppose that the first generation, perhaps motivated by compassion or ethical considerations, intends to consume only a small share of the initial endowment of a non-renewable resource ("pie"), leaving the rest for posterity. If they somehow learned that the second generation plans to consume the entire remainder, passing nothing on, then generation 1 might be inclined to revise its initial generous intent, consuming more itself. But, if it turns out that generation 2's intent to consume everything is based on discovering that generation 3 plans to consume whatever is bequeathed to it, passing nothing on, then generation 1 might feel that generation 2 is not so undeserving after all. In short, what each generation concludes that it should do might well depend on what it expects subsequent generations to do. If a generation trusts that a plan that it initiates for enriching the future will be followed faithfully by at least the next few generations, then it may bequeath more than if it lacks such trust.

How such trust arises is better illuminated by behavioral game theory, experimental economics, psychology, and descriptions of what is sometimes called "social capital" (trustworthiness of others with whom one participates in dealings) than by the logical prescriptions of formal game theory models, in which trust plays no role. Real people are often far more altruistic and cooperative than models of purely rational behaviors and interactions predict or explain. For example, in the much-studied dictator game, one player is given an amount of money (or other desirable good) to divide between himself and a second player. The recipient has all the power in this game, and might be expected to keep everything for himself. But this is not what is observed in many experimental dictator games and variations (e.g., Bahr and Requate 2014): most dictators choose to bequeath substantial shares (e.g., 20 % or more) of the initial endowment to the other, powerless players, depending on what social norms are evoked by the contextual cues of the experiment, such as earning, sharing, pure giving, etc. (List 2007). To what extent generous sharing is observed in practice depends on many contextual factors, such as whether participants view their interactions in a market frame or in a gift frame; on whether taking as well as giving is included in the feasible set of actions; on whether the player selected as the dictator believes that the choice reflects his own skill or luck in a fair contest; on how often the situation is repeated. But purely selfish behavior is seldom observed (List 2007). The multi-generation division of an initial endowment can be viewed as an expansion and generalization of the dictator game in which each generation is in the position of the dictator in deciding how much to share with powerless future generations.

Although behavioral game theory and experiments provides insights into realistic behaviors that formal non-cooperative game theory (e.g., based on the understanding that all rational players will use sub-game perfect equilibrium (SPE) strategies if they exist and are unique) cannot, neither type of analysis is concerned primarily with clarifying what choices are most morally correct. Behavioral game theory and behavioral economics recognize that people (and some other primates) have intuitive and emotional responses to perceived fairness, equity, injustice, reciprocity, and altruism that are important drivers of decisions about when and how much to share in a variety of settings, including the dictator game and its variants. These personal moral intuitions and impulses are tremendously important in helping real people to cooperate more successfully than purely rational (SPE-implementing) agents can in many situations (e.g., one-shot and iterated Prisoner's Dilemma, the stag game, the centipede game, the trust game, and other staples of modern game theory) (Haidt 2012). Yet, they do not provide a coherent normative account of applied social or moral decision-making that can be used to obtain reliable moral guidance on important policy questions, such as how best to share the burdens and benefits of investments in disaster mitigation and in economic growth across generations. To the contrary, our moral intuitions are easily dumbfounded when situations appeal to several competing moral intuitions (Haidt 2012), as in alternative framings of tax breaks for dependents as providing disproportionate benefits to the wealthy, if they are used; or as imposing disproportionate costs on the wealthy, if they are not (Kahneman 2011). Likewise, the prescriptions from formal mathematical non-cooperative game theory models of interacting rational players (or generations, in our setting) are not intended to convey any moral authority: SPE solutions and refinements only guarantee that no player can get more of what it wants by changing its strategy, not that what any player wants is morally worthwhile.

11.3.3 Axiomatic Solution Concepts from Cooperative Game Theory

A different branch of game theory concerns itself directly with normative principles that participants in a cooperative enterprise might adopt to decide how to share the benefits from cooperation. This is axiomatic game theory, which defines and studies the logical relations among solution concepts in cooperative games, such as those involving bargaining or fair division of one or more goods. For example, the Shapley value, which assigns to each player the expected incremental value that the player creates by joining a coalition of other players (when the order in which the players join is random), is the unique way to allocate gains from cooperation that satisfies certain axioms, such as that players be treated symmetrically, without favoritism (each player's allocation depends only on what it contributes, and not on who the player is) and that the allocation procedure should be Pareto-efficient, allocating 100 % of the potential gains from cooperation. The Nash Bargaining solution, which maximizes the product of the utility gains of the players compared to the no-agreement outcome, is likewise the unique solution concept satisfying these two conditions and an additional two (that the outcome should not depend on the arbitrary choice of scales for expressing player utilities, and that expanding the opportunity set of utilities jointly available to the players by introducing new options should either leave the original outcome unchanged, or change it to one of the newly added options). Each of these and many other proposed solution concepts for sharing in the cooperative production and allocation of desired outcomes, can be justified as the unique concept implied by a set of more-or-less reasonable-seeming normative axioms, many of which directly represent principles of fairness such as symmetry (i.e., no favoritism), or maximizing the minimum payoff among players, as well as Pareto efficiency.

However, axiomatic cooperative game theory is vulnerable to its own version of moral dumbfounding: different proposed normative principles can conflict logically. This leads to impossibility theorems showing that no possible decision procedure can satisfy all the appealing principles (normative axioms) that one might want to require. For example, the Shapley value solution and the Nash Bargaining solution can prescribe different outcomes for the same situation, so that no outcome satisfies the normative principles proposed for both. In such cases, one might try to decide which principles should be sacrificed so that other (mutually consistent) ones can be preserved, or, equivalently, choose among alternative solution concepts. But there is no meta-ethical framework within axiomatic game theory to prescribe or justify which normative principles to keep and which to abandon when there are logical conflicts among them. Axiomatic theories may also under-determine outcomes in practice. For example, the Nash bargaining solution requires knowing what the disagreement outcome is, but it is not always clear how it should be determined, e.g., as the present *status quo*, which may incorporate the results of many historical injustices, or as an idealized initial position that treats all participants symmetrically.

Such questions about how one should, even in principle, implement the prescriptions of normative axiomatic solution concepts, open a gap between their mathematical implications and the information needed to act on them. Theories of justice developed by Rawls and his successors can help to close this gap by specifying a particular initial position from which (idealized, hypothetical) deliberation and selection of principles proceeds. They also provide a meta-ethical framework for reasoning about how societies should choose among rival normative principles (e.g., among different axiomatic cooperative game-theoretic solution concepts) to guide their subsequent applied choices and the rights, duties, and principles of fair distribution or redistribution over time that they should impose on themselves and on each other. These theories are explained next.

11.3.4 Intergenerational Justice

Philosophical discussions of intergenerational justice since Rawls (2001) have considered an idealized form of social contracting across generations. Participants from all generations are imagined to jointly agree on policies for sharing resources and investments over time, such as social savings rates, from a hypothetical "original position" behind a "veil of ignorance," in which no participant knows which generation (or other position within society) he or she will be born into (e.g., Meyer 2009; Manzini et al. 2010). Principles that would be agreed to from behind this veil of ignorance are defined as principles of justice, to which real policy-makers should adhere if they wish to make just decisions. This concept can be applied to inform each generation's consumption and investment decisions, including decisions about how much to consume and how much to invest in disaster mitigation and community resilience, or directly in economic growth, at different times. In the simple case of sharing a non-renewable resource ("pie") over time for a finite number of generations (which may be random, if the duration of the human race is uncertain), the allocation of the resource across generations recommended by such a Rawlsian criterion typically coincides with the allocation that would be achieved by social utility maximization in growth economics, although the two criteria of social utility maximization and Rawlsian justice may recommend different allocations of the resource across generations if the resource can be produced, thereby augmenting its supply, by generations that choose to do so (Llavadora et al. 2010).

Identifying just policies as those that would result from social contracting if all stakeholders started from a symmetric original position solves the problem of having earlier generations exploit their asymmetrically powerful position to the detriment of later generations, e.g., by consuming all of the initial stock of pie immediately. Because participants in the multi-generation social contract arrived at from the original position do not know which generation they will occupy, they are motivated to treat all generations fairly. This framework for inter-temporal justice also provides an alternative to discounting, thus avoiding the ethical problem posed by conventional discounting in cost-benefit analysis of under-weighting the costs and benefits borne by far-future generations compared to those borne by present or near-future generations (van Liederkerke 2004; Parfit 1984). From the original position, the interests of different generations are valued equally, and so any discounting would reflect only real asymmetries, such as in production opportunities (e.g., earlier investments in growth pay off over more years), and not a bias against later generations.

Other proposed features of intergenerational justice have been erected on these foundations by speculating about what people in the original position would agree to. Rawls himself argued that the primary duty of each generation is to bequeath just institutions to the next; once this has been fulfilled (via an "accumulation phase"), a frequently proposed secondary duty is that each generation should pass on to the next an endowment at least equivalent (in size or productivity) to the one it received, so that no generation's consumption reduces the opportunities or wellbeing of those that follow (e.g., Hamilton 1995). However, such proposed principles of *sustainability* in consumption and production invite scrutiny and skepticism (e.g., Wolf 2007), especially if they are asserted without careful qualification and reference to specific underlying economic growth models. For example, in the case of an initial endowment of pie that can only be consumed, and not produced, requiring that no generation's consumption should diminish the stock bequeathed to future generations would imply that none of the pie would ever be consumed—a Pareto-inefficient outcome that would not necessarily appeal to anyone, even behind the veil of ignorance. Similarly, for a model of multi-generational sharing of a renewable resource, Krautkraemer and Batina (1999) show that imposing a sustainability constraint of non-decreasing utility over time creates Pareto-inefficient stockpiling of the resource: everyone would prefer a usage pattern that allowed later generations to have lower utilities than earlier ones. Such conflicts between various proposed criteria for sustainability and what all members of all generations would prefer arise in many other models of intergenerational sharing (e.g., Hoberg and Baumgärtner 2011), although not in all if there is no uncertainty and if property rights, taxes, and transfer payments among generations are dexterously deployed to allow earlier generations to, in effect, purchase resource usage rights from later ones (Howarth and Norgaard 1990). But the frequent conflicts between sustainability principles and economic efficiency (e.g., what is unanimously preferred by all members of all generations) are perhaps unsurprising from a Rawlsian perspective on distributive justice, insofar as social contracting from an original position in which no participant knows what generation he or she will occupy removes any reason to favor the utility or opportunities of later generations over those of earlier ones.

Despite their considerable intellectual appeal, theories of intergenerational justice based on implicit social contracting behind a veil of ignorance are not free of philosophical and logical difficulties. For example, in such theories, it is not always clear how potential future people whose very existence may be affected by present production and consumption decisions should be treated (Meyer 2009). Llavadora et al. (2010) present models in which the possible extinction of humanity is considered as a key uncertainty about the future. They prove that it is optimal, in economic growth models with a policy objective of maximizing the minimum welfare across generations, weighted by the sizes of future populations (so that potential individuals, rather than entire discrete generations, are treated as the participants in the social contract) to ignore this uncertainty about continued survival in deciding how best to allocate consumption and investment over time. This result holds when the economy is sufficiently productive. In effect, the appropriate discount rate due to uncertainty about continued existence is then zero.

Likewise, if future people will have habits, expectations, and preferences that are shaped in part by current production and consumption decisions, then it may not be clear what preferences should be assumed in modeling what they would agree to in the original position. For example, if early generations derive high utility from consumption, and if later generations regard consumption as an unattractive materialistic addiction that thankfully became obsolete when excessive consumption triggered a collapse, then a Rawlsian social contract that allowed members of different generations to bring these preferences with them behind the veil of ignorance might lead to the conclusion that high early consumption is justified. (We acknowledge that it is not clear whether a preference for increasing or decreasing consumption over time can be appropriately debated behind the veil of ignorance, as it may require more preference information than ignorance allows.) But if more gradual consumption would lead to all generations putting a high value on it, then these different assumed preferences might imply that gradual consumption was the just pattern. If each alternative choice about patterns of production, consumption, and investment over time induces the generation-specific preferences needed to justify it (by making it the alternative that would be selected from the original position), then the original position loses its prescriptive power.

11.3.5 Sustainability, Protective Principles, and Fundamental Trade-Offs

Sustainability principles, usually requiring that resources or opportunities or utilities be non-decreasing over time (Hamilton 1995; Wolf 2007), are intended to make sure that current policies do not unjustly sacrifice the interests of powerless future generations in favor of the interests of present generations who currently have the power to make choices. Other approaches have the same goal of protecting the interests of future generations in current decision-making; these range from variations on Rawls's idealized social contracting (subordinating the difference principle, which implies accumulation would never begin if the first generation is also the least well-of, to the desirability of an "accumulation phase") (van Liederkerke 2004) to idealized Nash bargaining solutions in which each generation is assumed to be given the power to veto plans that it considers unacceptable (Manzini et al. 2010). In addition, some optimal growth models imply a sustainability condition, in which, in steady state, each generation passes on to the next the capital that it received from its predecessor, and this stock is maintained at a level that maximizes the utility from consumption per capita for members of each generation (Phelps 1961).

However, when realistic uncertainties about future preferences and choice sets and about the consequences of current decisions are taken into account, the capacity of current decision-makers to protect the interests of future generations effectively may be very limited. For example, Krysiak and Collado (2009) present models in which there is a trade-off between taking actions to protect future generations against risks and taking actions that all generations prefer, while Hoberg and Baumgärtner (2011) demonstrate similar trade-offs between sustainability and efficiency when irreversible policy decisions are made by earlier generations, trying to protect the interests of later ones, but later generations have better information about the (perhaps unforeseen and unintended) consequences of earlier policies.

More generally, sustainable production and consumption, economic efficiency (i.e., Pareto optimality), Rawlsian justice, and free democratic (non-dictatorial) choice procedures have all been proposed as desirable principles for guiding and constraining how societies should make decisions, both within and across generations. But careful analysis indicates that any two of these principles, when appropriately formalized for specific models relating choices to economic growth, can conflict. For example, free democratic choice procedures may lead to outcomes that no one favors, e.g., if different people have different beliefs about the probable consequences of alternative choices, and these probabilistic beliefs are used to help select policies (Nehring 2005). Similarly, conflicts between sustainability criteria and Pareto-efficiency (Wolf 2007; Hoberg and Baumgärtner 2011) and trade-offs between Pareto-efficiency and various measures of intergenerational equity or justice in resource allocations arise for many intergenerational decision processes (e.g., Krysiak and Collado 2009). No matter how well intended, efforts to protect the interests of future generation using the information available today risks creating outcomes that, in retrospect, no one favors; this is especially likely if today's choices have uncertain, irreversible consequences and if future preferences are uncertain (Hoberg and Baumgärtner 2011). Thus, fundamental tradeoffs must be made among these proposed desirable characteristics of collective choice procedures for managing investments in growth vs. protective measures and production and consumption of limited resources (including infrastructures that reduce the adverse impacts of natural disasters) over time and generations. Equivalently, impossibility theorems expressing the logical incompatibility of different sets of principles under stated conditions limit the possibilities for a satisfactory approach to intergenerational cooperation, including management of natural disaster risks.

11.3.6 Investing in Building Resilient Communities and Societies: An Emerging Framework

If the normative frameworks and principles for intergenerational justice we have considered so far—growth economics, behavioral game theory, axiomatic solution concepts for cooperative games, philosophical theories of intergenerational justice, and proposed principles of sustainability and protection of the interests of future generation—all lead to contradictions or unresolved difficulties, then what type of analysis might be used instead to provide practical guidance on how much to spend on consumption, investments in disaster mitigation, and investments in economic growth in each period? One emerging approach avoids such mathematical and theoretical arguments, instead emphasizing building the capacity of communities and societies to respond quickly and competently to new information and circumstances as they arise. This is the framework of *resilience* for communities and societies; it is still under development by many investigators (Bretherton and Ride 2011; Lucini 2014; Tierney 2013). Key concepts are that resilient communities should prepare to manage disasters effectively by accumulating the physical and social capitals needed to adapt and respond effectively and to manage mitigation and recovery efforts when needed. Physical capitals include transportation, telecommunications, power, and emergency medical infrastructures. Social capitals include training and preparation, ability of communities to organize effectively and act competently and autonomously when needed, self-reliant communities, and high trust and individual trustworthiness in cooperating to respond to disasters. Proponents of resilience often argue that communities and societies benefit in multiple ways from developing the capacity, responsibility, and self-reliance needed to improvise and take appropriate actions to deal with possibly unforeseen events. From this perspective, the obligation of each generation to the next may be viewed as bequeathing at least a minimal level of resilience, including the needed physical and social capitals.

Resilience frameworks are still a work in progress. They are sometimes conflated with proposed principles of sustainability, participatory decision-making, environmental justice, environmentalism, putative rights to safety and prosperity, and intergenerational equity for managing interlinked social-economic-ecological systems, often with little explicit discussion of the limitations, trade-offs, and contradictions among the principles espoused. Even without such overlays, however, it is useful to keep in mind the basic idea that investment in building community resilience may be a valuable alternative or complement to investments in economic growth or direct protective measures (e.g., better levees) when the possibility of occasional disasters is present.

11.4 Discussions: Principles for Applying the Frameworks to Improve Decisions and Policies

How might an engineer, planner, or policy maker apply insights from the preceding frameworks to improve practical disaster protection and mitigation decisions, such as how high to build a costly levee or sea wall, or how large (and rare) a tsunami, earthquake, flood, or hurricane to plan for in the design of nuclear power plants or other facilities, or how much to invest in a proposed community resilience or civil defense program? The following suggested principles seek to distil from the frameworks considered above implications to help guide practical decision-making when current choices have long-lasting consequences that may affect risks and benefits to future generations.

1. **Use wide framing**: Consider a wide range of alternatives to optimize benefits produced for resources spent, taking into account opportunity costs. Exploit different ways to reduce risk. The optimal economic growth framework implies that each method for maximizing a social objective function—whether by investing in economic growth, in reducing potential disaster-related losses, in less costly and more rapid and resilient recovery following disasters, or in other improvements that reduce risk or increase wellbeing—should be funded optimally in each period in competition and combination with the other approaches. In simple settings with diminishing marginal returns, for example, this typically requires funding each alternative up to the point where a different one starts to yield larger marginal returns in improving the objective function. Thus, a planner wondering whether to invest in a taller sea wall or barrier against flooding should ask not only "Is the extra cost of a taller barrier justified by the extra benefit from reduced risk?" but also "Could a larger reduction in risk (or, more generally increase in objective function) be achieved by not making it taller, and instead applying the resulting cost savings to other opportunities, such as relocating people or improving local early warning and transportation systems?" More generally, optimal provision of safety and other good requires considering opportunity costs and optimizing economic trade-offs, while avoiding narrow framing (Kahneman 2011) that considers only one approach at a time (e.g., investment in levees, but not change in zoning or land use). The optimization problems to be solved can be viewed as allocating each period's limited resources to a portfolio of alternative ways to increase the objective function, with one of those ways being to bequeath more to the next generation, which may have different opportunities.
2. **Follow golden-rule consumption and investment principles**: Do not over-invest (or under-invest) in protecting or benefitting future generations compared to the present one. Biases such as the affect heuristic, which judges alternatives and shapes perceptions of their attributes and desirability based on whether they elicit positive or negative emotional responses ("gut reactions") (Kahneman 2011) can encourage simplistic thinking that equates current consumption with

selfishness and greed (bad affect) and equates current investment to protect or benefit future generations with benevolence, virtuous self-restraint, and generosity (good affect). Optimal growth models, including ones with ethical and justice constraints, tell a more nuanced story. Under-consumption and over-accumulation of capital stocks to pass on to the future violate the golden-rule maxim of doing in each generation what one would want other generations to do to maximize sustainable utility (Phelps 1961). From this perspective, increasing saving and investment on behalf of the future is not necessarily always better. Instead, saving and investing at the golden-rule rate, and not more, maximizes the wellbeing of present and future generations. Thus, optimal economic growth theory weans us from a multi-generation zero-sum perspective, in which increased current consumption necessarily comes at a cost to future generations. Instead, it encourages a cooperative perspective in which members of different generations collaborate in maximizing the sustainable level of wellbeing.

3. **Use simple rules to help optimize current decisions**: Exploit qualitative properties of optimal policies to simplify practical decisions. The economic growth perspective can be implemented in detail if trustworthy mathematical or computational models are available representing the causal relation between choices and the probabilities of their consequences (immediate and delayed). Techniques such as stochastic dynamic programming can then be used to decide what to do in each period to maximize a social objective function. Mathematical and computational techniques and resulting solutions can become quite sophisticated and complex, but, in many settings, the optimal solutions have qualitative properties that can inform and improve practical decision-making with simple rules that take into account future effects, even when detailed models and numerical optimization results are not available. For example, both optimal growth and Rawslian justice might require first boosting economic productivity as quickly as possible to a level where desirable institutions can be sustained and passed on from one generation to the next. Once there, optimal growth policies often have simple characteristics, such as saving and investing just enough so that the marginal productivity of additional capital stock offsets (equals) its effective depreciation rate due to aging, population growth (which dilutes the capital-per-worker), and other causes, including occasional disasters or catastrophes (Phelps 1961). Risk management to jointly optimize consumption and investments in growth, disaster prevention and mitigation to maximize average utility of consumption per capita per period might require keeping capital stocks of renewable resources at or above certain threshold levels to avoid risk of collapse, which could reduce or eliminate their availability to subsequent generations (Olson 2005). Such simple characterizations of optimal growth and risk management policies can help to focus practical policy-making analysis and deliberation on a few key questions, such as whether the current savings and investment rate is clearly above or below the socially optimal rate [e.g., the golden rule rate, in a Solow growth model (Phelps 1961)]; or whether stocks of renewable resources are currently above or below safety-stock thresholds. The answers then suggest directions for remedial actions, such as increasing or

reducing investment, respectively. Pragmatic constraints may limit how much adjustment can be made how quickly. In short, knowledge of the qualitative properties of optimal policies, such as the existence of thresholds or of optimal rates of capital accumulation or investment, can produce simple decision rules (e.g., take action to increase investment or stock of a renewable resource if we are below the optimal level, or to decrease it if we are above the optimal level, where the optimal level is estimated from data on depreciation rates or renewal rates, respectively). Such simple rules can often help to answer the practical policy question of what to do next, even without explicit formulation, estimation, and solution of sophisticated multi-period optimization models.

4. **Do not discount the utility from future benefits**: Count lives saved in different generations equally and count utility received in different generations equally. In particular, do not discount future lives saved or future utility from benefits received simply because they are in the future. This follows from Rawlsian justice models that treat the interests of future participants in an extended multi-generational social contract symmetrically with present ones. It implies that benefits such as improvements in quality-of-life per person per year due to increased resilience and reduced anxiety, or greater consumption utility per capita-year, should not be discounted. In making cost-benefit comparisons, lives saved or life-years improved that accrue over the lifetime of a facility should all be counted equally according to such models of justice over time. The practical effect of this recommendation is to increase the present evaluation of benefits that flow from current decisions into the future, such as the benefits from risk reductions obtained via current investments in levees or in other protective or resilient infrastructure. Although multi-period optimization methods such as stochastic dynamic programming can still be used to decide what to do in detail, if credible models are available to support the required calculations, concern for intergenerational justice will modify the usual objective function of expected discounted social utility to give equal weights to life-saving or other intrinsically valued benefits received at different times.
5. **Consider the value of waiting**: Do not commit prematurely to expensive present actions with long-lasting or irreversible consequences. Trust future generations to help decide what is best. This principle requires current decision-makers to consider the potential value of seeking and using better information to improve decisions before committing resources or foreclosing other options. For example, it may be worthwhile for Federal regulators to let individual states experiment with new measures first, and to learn from the consequences, before deciding on a Federal policy that all states must follow. This cautious principle, of seeking to learn more before betting large-scale investment decisions with lasting consequences on what currently seems to be the best choice, follows from studies of fundamental trade-offs in making protective investments under uncertainty, such as the trade-off between investing in proposed measures to protect future generations against possible future harms vs. investing in other ways that, in retrospect, all members of all generations might prefer (Krysiak and Collado 2009; Hoberg and Baumgärtner

2011). Acknowledging realistic uncertainties about future costs, benefits, risk attitudes, preferences, technology alternatives, and opportunity costs highlights the potential value of seeking more information before committing to decisions with long-lasting or irreversible consequences, such as about the height of a levee, enactment of enduring regulation of carbon dioxide emissions, diminishment of economic growth rates in order to invest in protective measures, or consumption of non-renewable resources.

Our first principle above, wide framing, encourages planners confronted with a proposed costly measure to reduce risks to future generations to ask not only "*Is it worthwhile?*" in the sense that the proposed measure's benefits exceed its costs, but also "*Is there a cheaper way to achieve the same benefits?*" The latter question is typically a matter for engineers and economists. The answer is often that no one knows yet. If further research can reduce current uncertainties, then value-of-information (VOI) calculations from decision analysis can address the question of whether the benefits of that research, in terms of improving decisions and their probable outcomes, exceed the costs of doing it, including potential costs from delay. (Indeed, such VOI considerations are automatically included in stochastic dynamic programming whenever acquiring more information is a possible choice.) Thus, the planner should also ask a third question: "*Is it worthwhile to pay for better information before deciding whether to approve the proposed risk-reducing measure?*" When decisions have long-lasting consequences that affect future generations, the value of information acquired to help make the best decision—the decision that will be preferred in retrospect when future information becomes available—may be especially great.

Likewise, there may be a value to keeping options open, recognizing that the best choice based on current information may not still be seen as the best one when evaluated using future information. There can be a "real option" value to keeping options open until better information is available on which to act, even if delay is costly. Again, stochastic dynamic programming considers such real option values as well as VOI, and optimizes information collection and the timing of decisions, including ones with irreversible or long-lasting. However, the guiding principle of stochastic dynamic programming is the Bellman optimality principle: that each generation's (or period's) decisions are made optimally, *assuming* that all future generations' (or periods') decisions will likewise be made optimally (Olson 2005). Practical application of this principle across generations requires decision-makers in different generations to collaborate in implementing it consistently, with each generation acting accordingly, but having to trust other generations to do the same. Behavioral game theory suggests that such cooperation is far more likely than would be expected based on purely rational (System 2) responses, in part because of moral psychology and pro-social impulses (System 1 responses) (Haidt 2012) that make us eager to reciprocate the generosity of earlier generations by being, in our turn, equally generous to our successors. However, System 1 responses are notoriously non-quantitative, and are not designed to identify and optimize quantitative trade-offs (Kahneman 2011). Thus, deliberate investments in social capital, a culture of trustworthiness and effective cooperation, and building resilient communities, may help generations to collaborate more effectively over time in implementing long-term plans that benefit all of them.

11.5 Conclusions

Making decisions well over time is challenging for societies as well as for individuals. Deciding how best to allocate current resources to competing ends, of which protection of future generations against natural disasters is one, requires considering what future generations are likely to do to maintain and expand (or, perhaps, to draw down without replenishing) the stocks of infrastructure, social and economic capitals, and other protective investments bequeathed to them. Our moral intuitions often deliver altruistic and benign impulses toward others, including strangers separated from us in time or by geography. But they usually do not render finely calculated decisions about how to optimize trade-offs between our benefits and theirs. Nor do they identify the most efficient use of protective and other investments (e.g., in growth of economic prosperity, or in building community resilience) to accomplish desired trade-offs, or to carry out the prescriptions of ethical and justice theories to maximize the average or minimum wellbeing of members of present and future generations. Methods of multi-period optimization that have long been used in optimal economic growth models can accomplish these quantitative trade-off and optimization tasks. They can be adjusted to incorporate principles of intergenerational justice, such as assuring that the lives and utilities of future people are not discounted relative to those of people now living. In simple models, including the multi-generation pie-sharing example that we started with and in golden-rule optimal growth models (Phelps 1961), multi-period optimization leads to simple consumption and investment rules that also satisfy equity and sustainability conditions. In this way, System 2 methods can be used to help identify multi-generation investment plans to achieve ethical goals that System 1 might approve of. The results can often be expressed as simple decision rules that are useful for informing practical policy-making, such as taking actions to adjust levels of investments or of renewable resources toward desired target levels based on estimated marginal rates of return or renewal rates, respectively, perhaps with adjustments for the value of information and of keeping options open.

However, even when clear and simple rules can be identified for maximizing a social objective function, such as the average or minimum utility per capita in present and future generations, it takes cooperation across generations to implement them. In turn, this may require just and effective institutions, high social capital, and community resilience, as prerequisites for effective multi-generational cooperation in managing losses due to natural disasters or other causes. These insights suggest that successful efforts to improve intergenerational justice and efficiency must be rooted in a deep understanding of human social nature and cooperation over time, and of the possibilities for designing and maintaining effective cultures and institutions for promoting and sustaining such cooperation. They also require clear understanding of the goals that we seek to achieve in collaboration with other generations, and of the trade-offs that we want to make when those goals conflict.

The frameworks and principles we have discussed make a start at clarifying possible goals and trade-offs among them, and the implications of technical principles (especially, stochastic dynamic programming distributed over multiple generations) for achieving them. How to develop institutions and cultures that promote effective cooperation over time and across generations, as well as within them, without necessarily assuming that future generations will share our preferences and values, remains a worthy problem for both theoretical and applied investigation.

References

Bahr G, Requate T (2014) Reciprocity and giving in a consecutive three-person dictator game with social interaction. Ger Econ Rev 15(3):374–392

Balbus ÃL, Nowak AS (2008) Existence of perfect equilibria in a class of multigenerational stochastic games of capital accumulation. Automatica 44(6):1471–1479

Bretherton D, Ride A (2011) Community resilience in natural disasters. Palgrave Macmillan, New York

D'Albis H, Ambech S (2010) Fair intergenerational sharing of a natural resource. Math Soc Sci 59 (2):170–183

Haidt J (2012) The righteous mind: why good people are divided by politics and religion. Pantheon, New York

Hamilton K (1995) Sustainable development, the Hartwick rule and optimal growth. Env Resour Econ 5(4):393–411

Hoberg N, Baumgärtner S (2011) Irreversibility, ignorance, and the intergenerational equity-efficiency trade-off. University of Lüneburg Working Paper Series in Economics no. 198. www.leuphana.de/institute/ivwl/publikationen/working-papers.html

Howarth BR, Norgaard BR (1990) Intergenerational resource rights, efficiency, and social optimality. Land Econ 66(1):1–11

Kahneman D (2011) Thinking fast and slow. Farrar, Straus, and Giroux, New York

Krautkraemer AJ, Batina GR (1999) On sustainability and intergenerational transfers with a renewable resource. Land Econ 75(2):167–184

Krysiak FC, Collado IG (2009) Sustainability and its relation to efficiency under uncertainty. Economic Theory 41(2):297–315

List JA (2007) On the interpretation of giving in dictator games. J Polit Econ 115(3):482–493

Llavadora H, Roemer JE, Silvestre J (2010) Intergenerational justice when future worlds are uncertain. J Math Econ 46(5):728–761

Lucini B (2014) Disaster resilience from a sociological perspective: exploring three Italian earthquakes as models for disaster resilience planning. Springer, Switzerland

Manzini P, Mariotti M, Veneziani R (2010) Intergenerational justice in the Hobbesian state of nature. Economics Department Working Paper Series. Paper 108. http://scholarworks.umass.edu/econ_workingpaper/108

Meyer L (2009) Intergenerational justice. In: Zalta EN (ed) The Stanford encyclopedia of philosophy (Winter 2014 Edition). http://plato.stanford.edu/archives/win2014/entries/justice-intergenerational/

Nehring K (2005) The (im)possibility of a paretian rational. Mimeo. Available at http://www.sss.ias.edu/publications/papers/econpaper68.pdf

Oliver P, Hauser OP, Rand DG, Peysakhovich A, Nowak MA (2014) Cooperating with the future. Nature 511(7508):220–223. doi:10.1038/nature13530

Olson LJ (2005) Theory of stochastic optimal economic growth. http://faculty.smu.edu/sroy/olson-roy-handbook.pdf

Parfit D (1984) Reasons and persons. Oxford University Press, Oxford

Phelps E (1961) The golden rule of accumulation: a fable for growthmen. Am Econ Rev 51 (4):638–643. http://www.jstor.org/stable/1812790

Ramsey FP (1928) A mathematical theory of savings. Econ J 38(152):543–559

Rawls J (2001) Justice as fairness. Harvard University Press, Cambridge

Schmidtz D (2011) Nonideal theory: what it is and what it needs to be. Ethics 121(4):772–796. doi:10.1086/660816

Solow RM (1956) A contribution to the theory of economic growth. Q J Econ 70(1):65–94. doi:10.2307/1884513. JSTOR 1884513

Swan TW (1956) Economic growth and capital accumulation. Econ Rec (Wiley) 32(2):334–361. doi:10.1111/j.1475-4932.1956.tb00434.x

Tierney K (2013) “Only connect!” social capital, resilience, and recovery. Risk Hazards Crisis Public Policy 4(1):1–5

Van Liederkerke L (2004) Discounting the future: John Rawls and Derek Parfits’ critique of the discount rate. Ethical Perspect 11(1):72–83. http://www.sehn.org/tccpdf/tccdiscounting%20future.pdf

Wolf C (2007) Intergenerational justice. In: Frey RG, Wellman CH (eds) A companion to applied ethics, chap 21. Wiley, New York. http://www.public.iastate.edu/~jwcwolf/Papers/Wolf_Intergenerational_Justice.pdf