

Working Paper Series
Congressional Budget Office
Washington, D.C.

Governmental Risk Taking Under Market Imperfections

Michael Falkenheim
Congressional Budget Office
michael.falkenheim@cbo.gov

Working Paper 2021-07

June 2021

To enhance the transparency of the work of the Congressional Budget Office and to encourage external review of that work, CBO's working paper series includes papers that provide technical descriptions of official CBO analyses as well as papers that represent independent research by CBO analysts. Papers in this series are available at <http://go.usa.gov/xUzd7>.

This work benefited from the guidance of Sebastien Gay. Jeffrey Kling reviewed the document. Helpful comments were provided by Christopher Adams, William Carrington, Wendy Edelberg (formerly of CBO), Mark Hadley, Wendy Kiska, John Kitchen, Sam Papenfuss, Jeffrey Perry, Chad Shirley, Robert Sunshine, and David Torregrosa. Useful input was also received from Coleman Bazelon of the Brattle Group, Hening Bohn of University of California at Santa Barbara, Deborah Lucas of the Massachusetts Institute of Technology, Donald Marron of the Urban Institute, Damien Moore of Moody's Analytics, and Loiuise Sheiner of the Brookings Institution. Vinay Maruri provided research assistance. Elizabeth Schwinn was the editor.

Abstract

An extensive literature debates whether market prices should be used to measure the benefits and costs of risk in government activities or whether the government should be treated as risk neutral. This paper explores the benefits and costs of governmental risk taking in formal models of market imperfections, in which the government serves as an intermediary between different stakeholders in its finances. Some stakeholders cannot participate in markets, either because they belong to future generations or because they have no funds to invest and face borrowing constraints. The cost of risk for those government stakeholders might not be equal to market price under *laissez-faire* but will be the same as market prices under Pareto-efficient policy, which creates inframarginal benefits. In an overlapping generations model, the market price of risk might understate or overstate the cost of risk that is shifted by the government to future generations, depending on whether uncertainty is driven by permanent or temporary shocks to technology. Permanent shocks to technology lead the market price of risk to understate the cost of risk to future generations, whereas temporary shocks cause it to overstate the cost of such risk.

Keywords: government policy, uncertainty, risk premiums

JEL Classification: E32, G10, G12, G18, H50, H83

Contents

Summary	1
The Government's Cost of Risk	3
The Cost of Risk in an Endowment Model With Market Imperfections	4
Notation and Assumptions	5
<i>Laissez-faire</i> Equilibrium.....	6
Defining the Market Price of Risk	7
Pareto-Efficient Policy and the Market Price of Risk.....	8
Government Investment, Risk Shifting, and Welfare.....	11
Factors That Determine the Direction of Risk Shifting Under Pareto-Efficient Policy	14
The Cost of Risk for Overlapping Generations With Uncertainty	15
Notation and Assumptions	15
The Price of Risk Under <i>Laissez-faire</i>	16
The Price of Risk Under Pareto-Efficient Policy.....	17
Figures.....	24
Tables	26
Appendix: Derivation of Selected Equations.....	28
Reference List	31

Summary

The government takes on risk in many ways. It issues loans and loan guarantees, which expose it to the risk of high default rates. It insures bank deposits and pension funds, which expose it to the risk of higher-than-expected costs when bank failures spike and pension investments underperform. It charges taxes, which produce revenues that may be systematically higher or lower than average, depending on the economy's performance. Government spends funds on safety net programs that cost more in bad economic times than in good economic times. Those factors create the risk that deficits will be higher when the economy performs poorly and lower when it performs well. That systematic risk, which cannot be eliminated through pooling or diversification, is often referred to as market risk. Market risk commands a premium under finance theory, and idiosyncratic risks that can be eliminated in a diversified pool do not command that premium.

There is long-standing debate about how the government should capture risk in its budgeting and in benefit-cost analysis. Some observers argue that government should incorporate the cost of risk by using market prices, reflecting the premium that financial market participants would charge to bear the same risk. If government stakeholders had the same views and preferences about risk as private investors, then they would view the market risk associated with uncertain cash flows as having the same cost. When the government extends credit or engages in other activities with risky cash flows, the associated market risk of those obligations is effectively passed along to households and businesses in either current or future generations. When the government faces higher-than-expected costs on credit programs, for example, it must eventually offset the effect of those costs by cutting spending or raising taxes. Alternatively, government could operate with a higher level of debt, but then it would face higher interest costs and the debt would crowd out private investment.

Other observers suggest that because of the government's special nature, it does not require such a risk adjustment. Initially, some of those observers proposed that the government can spread idiosyncratic risks more broadly than private investors could. More recently, others have argued that the government has a superior ability to connect with future generations or overcome borrowing constraints. Those arguments have in common the proposition that government can reduce the cost of risk by sharing it efficiently with citizens who cannot participate in current markets, either because they have not yet been born or because they have little financial wealth and face borrowing constraints.

This paper formally analyzes the proposition that market imperfections cause the price of risk in the market to differ from the cost of risk for some government stakeholders. It also analyzes the effect of governmental risk taking on both the market price of risk and the cost of that risk to the government's stakeholders. Using both an endowment economy model and an overlapping generations model, the paper shows that under a *laissez-faire* policy in which the government does not intervene in markets, the market price of risk can differ from the cost of risk to

government stakeholders who cannot participate in markets. The market price of risk reflects the preferences and endowments of those who can participate in the market, not necessarily all government stakeholders.

The endowment model illustrates how governmental risk taking can redistribute risk among government stakeholders, some of whom can participate in financial markets and some of whom cannot, affecting the market price of risk in the process. If governmental policy achieves a Pareto-efficient distribution of risk, then it places the exact amount of risk on government stakeholders that makes them indifferent to taking on additional risk at market prices. Accordingly, a Pareto-efficient distribution aligns the market price with the marginal cost of risk for all government stakeholders.

The model also clarifies some concepts related to governmental risk taking. First, it clearly defines the costs and benefits of governmental risk taking and makes a distinction between those benefits and costs and the effects of government programs. For example, the model makes a distinction between the benefit of the government's taking on the risk of student loan default and the benefit of promoting higher education. Second, the model distinguishes between inframarginal and marginal benefits and costs. In the presence of market imperfections, a Pareto-efficient policy creates inframarginal net benefits by taking risk off the hands of some stakeholders and placing it on others; those actions result in a situation in which additional governmental risk taking has no marginal net benefit.

The overlapping generation analysis compares the risk premium under *laissez-faire* and Pareto-efficient policies when the economy is subject to shocks that affect the generations unequally. Under Pareto-efficient policy, the risk of shocks is shared among generations to a degree that aligns the market price of risk to its cost for each generation. Under *laissez-faire*, the market price of risk depends only on the risks facing the generation that is working and saving for retirement, whose consumption and investment trade-offs determine the market price of risk. If that generation is exposed to temporary technological shocks that do not affect future generations as heavily, then the market price of risk will be higher under *laissez-faire* than under Pareto-efficient policy. If that generation is exposed to permanent shocks that are magnified over generations, then the market price of risk will be lower under *laissez-faire* than under Pareto-efficient policy.

The implications of those results ties to literature on Mehra and Prescott's (1985) equity premium puzzle, in particular to the proposition that equity risk premiums might derive from long-run risks described by Bansal and Yaron (2005) and other studies. Under their view of risk premiums, the findings in this analysis would imply that the cost of risk for future generations could be higher than indicated by the market price, and thus the market risk premium under Pareto-efficient policy could be higher than under *laissez-faire*. If, in contrast, risk premiums derive to a greater extent from temporary shocks, as implied by Campbell and Cochrane (1999),

then future generations will be less exposed than current generations, and market risk premiums will be lower under Pareto-efficient policy than under *laissez-faire*.

The Government's Cost of Risk

Debate over the government's cost of risk (or government discounting) has shifted somewhat since it began in the 1960s and 1970s. Initially, arguments hinged on whether the government had a superior capacity to diversify. Diamond (1967) and Hirshleifer (1964 and 1966) argued that government investments should account for the market cost of risk; they relied on the assumption that imperfections in markets do not have a significant effect on the price of risk. Jorgenson, Vickrey, Koopmans, and Samuelson (1964) and Arrow and Lind (1970), in contrast, argued for evaluating government investments as if they were risk free because the government has a large, diversified portfolio. Sandmo (1972) argued that the main difference between the two views related to the independence of government exposures, noting that Arrow and Lind's theorem required government risks not to be correlated with the overall economy.

As Sandmo argued, the assumption of independent risk seems implausible for government exposures. The government faces risks of higher-than-average defaults and claims in credit and insurance programs that depend on business, employment, and housing market outcomes that clearly relate to the overall economy. Government spending in safety net programs and entitlement programs as well as government revenues have systematic relationships with the economy that come from either their formulas or their dependence on aggregate income. Moreover, financial markets have grown increasingly able to pool and spread risk among their participants, such that the idiosyncratic risks identified by Arrow and Lind are probably as well nullified through diversification in private markets as they could be through public pools. Under those circumstances, only market risk is likely to command a premium in financial markets, and government's finances seem to be exposed to market risk.

Some observers have argued that government's ability to connect with future generations or its tools to overcome the borrowing constraints faced by private households and businesses reduce its cost of risk.¹ Some of those advantages derive from the finite horizons of mortal individuals and the limitations imposed by those horizons on intergenerational risk sharing.² Bohn (1999, 2009, and 2010) and Orszag (2000 [in response to Bazelon and Smetters], 1999) cite results for the overlapping generations model to argue that the government can improve Pareto efficiency through intergenerational risk sharing. Bohn (2010) argues that "certain aggregate risks have social cost below their market prices—namely, risks to which future generations are less exposed

¹ See Sastry and Sheiner (2015) for a useful summary of this literature.

² Contrary to how they are often described, the inefficiencies that arise in overlapping generations models are not related to incomplete markets (see Shell 1971).

than current market participants.” Gordon and Varian (1988) argue that the government has a superior ability to share risk between generations and that doing so can improve welfare for its citizens. Ball and Mankiw (2007) and Gollier (2008) show how such sharing can be accomplished through pension schemes—for example, by having Social Security invest in equity.

Although not all market participants are wealthy, they are on average drawn from higher parts of the wealth distribution. Such limited participation in financial markets also casts doubt on market prices as a measure of risk for all government stakeholders. Mankiw and Zeldes (1991) argue that the equity premium puzzle can be at least partially explained by such limited participation given the fact that only 25 percent of households invest in the stock market. (The equity premium puzzle arises from the observation that equity premiums are higher than what would be consistent with the observed volatility of aggregate of consumption under most estimates of the average level of risk aversion.) Mankiw and Zeldes argued that equity premiums should relate to the volatility of market participants’ consumption, not to those who do not own equity, and found volatility to be higher for relatively wealthy households. Under that point of view, the market price of risk might overstate the cost to government stakeholders, although that effect might be offset by poor people’s greater aversion to risk, as documented in Shaw (1996).

Proponents of the use of market prices in federal decisionmaking have acknowledged the possibility of market imperfections but have argued that such imperfections are not significant enough to affect their conclusions. Lucas and Phaup (2010), following Kaplow (2007), argue that efficiency as an issue is separate from the issue of how much importance to assign the welfare of each generation; that is the question of intergenerational equity. However, intergenerational risk sharing is itself an efficiency question that is separable from intergenerational equity (how to weight the welfare of different generations), as shown in Bohn (2009). Moreover, shifting risk from one group to another is not trivial. Shifting risk to those whose cost of bearing it is lower is one of the main functions of financial markets and can lead to Pareto-efficient improvements in welfare.

This paper sheds light on the literature while taking a slightly different concept of risk. Rather than analyzing government’s cost of risk, as has been done in previous analyses, this paper treats the government as a mediator between different stakeholders, one that redistributes risk among them. The paper analyzes under what conditions the price of risk for all stakeholders is equal to the market price.

The Cost of Risk in an Endowment Model With Market Imperfections

Much of the previous literature suggested that the government should be seen either as risk neutral or as facing the same cost of risk as private investors. This paper suggests a more nuanced view. The models analyzed in this paper show that in the presence of market

imperfections, the government can theoretically overcome those imperfections by transferring risk from people who have a higher cost of bearing it to those whose cost is lower. And yet there is a limit to the net benefits that can be achieved through such risk transfers because they narrow the gap between the marginal cost of risk for market participants and the cost for other government stakeholders. With a Pareto-efficient amount of risk bearing, the market price of risk can be aligned with its cost for all government stakeholders.

In finance theory, the “marginal investor” represents the preferences and endowments of all unconstrained market participants, not of a single person. Individual households and private entities can adjust their level of risk to the point at which its marginal cost equals the risk premium. Given those adjustments, all individuals are indifferent to their last unit of risk at the margin, just as all consumers of apples would find the last apple they consume to be exactly worth the price. Breeden and Litzenberger (1978) formalize that notion with a model of complete markets in which individuals trade Arrow–Debreu (1954) securities that pay one dollar in every possible future state. Their analysis shows that a common utility function can represent all agents under minimal assumptions about their preferences and beliefs.

Notation and Assumptions

The endowment model in this analysis maintains Breeden and Litzenberger’s (1978) framework while introducing the assumption that markets are incomplete. Specifically, the model assumes that some individuals cannot participate in financial markets but that they are stakeholders in the government’s finances. Individuals are denoted by the subscript k . The set of individuals $k \in M$ participates in markets that trade consumption at time 0 and among a finite set of states $s \in S$ and times $t \in T$.

Each consumer $k \in M$ that participates in the market is endowed with consumption in different states and times. It can trade units of consumption with other consumers that have access to markets. At time 0, the consumer does not know which state will prevail in the future; he or she will learn the answer after trading has taken place. All individuals have the same subjective probability for each state, π_{ts} , at the time when trading takes place. The endowment of consumer k is y_0^k at time 0. Its endowment in state s at time t is y_{ts}^k . The consumer can sell that endowment to other market participants at prices ϕ_{ts} and buy consumption in each state at those same prices. The total value of consumer k ’s endowment at those prices is therefore $\sum_t \sum_s \phi_{ts} y_{ts}^k$, and the total value of its consumption, $\sum_t \sum_s \phi_{ts} c_{ts}^k$, is limited to that amount, such that its budget constraint is $\sum_t \sum_s \phi_{ts} c_{ts}^k \leq \sum_t \sum_s \phi_{ts} y_{ts}^k$. Its utility function, $u^k(c_{ts}^k, t)$, is concave but might differ from that of other consumers.³ Under those conditions each consumer maximizes the weighted average of

³ Some of the notation used here is from Breeden, Litzenberger, and Jia (2015), and the assumptions were originally introduced in Breeden and Litzenberger (1978).

his or her utility among states in each time period, summed across time periods, subject to a budget constraint:

$$\max_{c_0^k, c_{ts}^k, \forall t, s} u^k(c_0^k) + \sum_t \sum_s \pi_{ts} u^k(c_{ts}^k, t) \quad (1)$$

$$\text{s. t. } \phi_0 c_0^k + \sum_t \sum_{s \in S_t} \phi_{ts} c_{ts}^k \leq \phi_0 y_0^k + \sum_t \sum_{s \in S_t} \phi_{ts} y_{ts}^k$$

Laissez-faire Equilibrium

Solving the maximization problem in equation (1) will result in state prices ϕ_{ts} that reflect the total endowments of all market participants in each state and time as well as their combined preferences—the summed y_{ts}^k and all the u^k functions for the set of individuals who participate in asset markets. Specifically, under the first-order conditions, the following “bang for buck” equation will hold:

$$\frac{1}{\phi_0} \frac{\delta u^k(c_0^k)}{\delta c} = \frac{\pi_{ts}}{\phi_{ts}} \frac{\delta u^k(c_{ts}^k, t)}{\delta c}, \forall k, t, s \quad (2)$$

The ratio of the marginal utility of consumption in each state multiplied by the probability of that state to the state price will equal the marginal utility of consumption in time 0 divided by the price of consumption in period 0, for all times and states.

The total consumption of all individuals will equal their total endowment for each state and time:

$$\sum_{k \in M} c_{ts}^k \leq \sum_{k \in M} y_{ts}^k, \forall t, s \quad (3)$$

Breeden and Litzenberger (1978) have shown that every consumer’s consumption in this model is ordered among states in the same way as aggregate consumption. That ordering occurs because all consumers face the same order for the ratios of probabilities for state prices $\frac{\pi_{ts}}{\phi_{ts}}$.

Under equation (2), the order of the marginal utility of consumption is inversely related to the order of that ratio; because the second derivative of that utility is negative (by assumption), the marginal utility is a monotonic function of consumption. Under those conditions, a common aggregate utility function u^M can represent the utility functions of all market participants.

Under the first theorem of welfare economics, market equilibrium leads to a Pareto-efficient allocation. Therefore, the state prices characterized by equations (2) and (3) are associated with a level of consumption for market participants in which none can be made better off without making another worse off.

Individuals who are not able to participate in financial markets (individuals $k \notin M$), are outside the process for determining state prices or the shape of u^M . In that situation, trades between the other government stakeholders and the market participants might be Pareto-improving. It might

be possible for market participants and nonparticipants to gain if trade were possible between them, but that is not possible because of incomplete markets.

Defining the Market Price of Risk

The market price of risk at time t can be defined in terms of the state prices ϕ_{ts} as the difference between two bundles of Arrow–Debreu securities: a risk-free bundle and a risky bundle. The risk-free bundle consists of one unit of every state-time security at time t and thus pays one unit of consumption with certainty. Its price is the sum of the state prices at time t :

$$p_t^{\text{Risk Free}} = \sum_s \phi_{ts} \quad (4)$$

The risky bundle of Arrow–Debreu securities pays the same amount of consumption on average—one unit—but in an amount that directly varies with aggregate consumption, which is the sum of the consumption of the two agents. Define aggregate consumption as c_{ts} and its average value of in time t as \bar{c}_t . The price of the risky security is the following:

$$p_t^{\text{Risky}} = \sum_s \phi_{ts} \frac{c_{ts}}{\bar{c}_t} \quad (5)$$

The cost of risk is the difference between those two prices:

$$p_t^{\text{Risk}} = p_t^{\text{Risk Free}} - p_t^{\text{Risky}} = \sum_s \phi_{ts} \frac{\bar{c}_t - c_{ts}}{\bar{c}_t} \quad (6)$$

Combining that expression with expression (2)⁴,

$$p_t^{\text{Risk}} = -\phi_0 \sum_s \pi_{ts} \left(\frac{\frac{\delta u^M(c_{ts,t}^M)}{\delta c} - \frac{\overline{\delta u^M(c_{ts,t}^M)}}{\delta c}}{\frac{\delta u^M(c_0^M)}{\delta c}} \right) \left(\frac{c_{ts} - \bar{c}_t}{\bar{c}_t} \right) \quad (7)$$

Which translates into the following:

$$p_t^{\text{Risk}} = -\phi_0 \text{Cov} \left(\frac{\frac{\delta u^M(c_{ts,t}^M)}{\delta c}}{\frac{\delta u^M(c_0^M)}{\delta c}}, \frac{c_{ts}}{\bar{c}_t} \right) \quad (8)$$

Under expression (8), the cost of risk is equal to a negative constant times the covariance of aggregate consumption and the market participant’s marginal utility of consumption. That covariance is directly related to the volatility of consumption and the consumer’s risk aversion; higher levels of risk aversion and volatility lead to a higher risk premium.

⁴ See Appendix for derivation.

Pareto-Efficient Policy and the Market Price of Risk

Policy could enable Pareto-improving transfers between market participants and other government stakeholders, affecting the market price of risk in the process. Pareto-efficient policy will thus result in a different allocation of risk than the *laissez-faire* equilibrium and lead to a different market price of risk.

Our model assumes that Pareto-efficient policy allocates consumption among market participants and other stakeholders in each state and time to maximize the total weighted utility of the individual consumers. Just as a single utility function $u^M(c_{ts}^M, t)$ can combine the functions of each market participant, as shown in the previous section, so too can a single utility function represent the utility of all other government stakeholders. Let $u^O(c_{ts}^O, t)$ represent the aggregate utility of other government stakeholders given their aggregate consumption c_{ts}^O in state s at time t . The other stakeholder cannot trade securities in the market; that stakeholder's price of risk represents a hypothetical price that individual would be willing to accept to take on an additional unit of risk. For the other stakeholder, the price of risk would be expressed by equation (8) with c_{ts}^O and c_0^O in place of c_{ts}^M and c_0^M .

Pareto-efficient policy optimizes the weighted average of the utility of two representative consumers with those utility functions: one agent representing consumers who participate in financial markets, denoted M , and one representing other government stakeholders who cannot participate in financial markets, denoted O .

Given that they have a separate utility function, other government stakeholders could be more or less risk-averse than market participants. Other government stakeholders may be either members of future generations or members of current generations who lack the ability to participate in financial markets. Those who cannot participate in financial markets might be relatively poor and thus closer to subsistence levels of consumption, making them more averse to losing a given share of consumption than the wealthier market participants. That higher level of risk aversion might make it Pareto-efficient to transfer risk from the other government stakeholder to the market participant. However, there might be an offsetting effect if that stakeholder's endowments vary less than market participants. Based on the findings of Mankiw and Zeldes (1991), households that participate in financial markets have total consumption that varies more than those that do not, suggesting that nonparticipants in the market have less risky endowments in the real world.

Pareto-efficient policy weights the utility of those two agents in the following objective function:

$$\max_{c_0^M, c_0^O, c_{ts}^M, c_{ts}^O, \forall t, s} w_M u^M(c_0^M) + w_O u^O(c_0^O) + \sum_t \sum_s \{w_M u^M(c_{ts}^M, t) + w_O u^O(c_{ts}^O, t)\} \quad (9)$$

$$\text{s. t. } c_0^M + c_0^O \leq y_0^M + y_0^O$$

$$c_{ts}^M + c_{ts}^O \leq y_{ts}^M + y_{ts}^O, \forall s, t$$

The solution to that problem yields a Pareto-efficient allocation that incorporates the preferences and endowments of both the market participant and the other government stakeholder who cannot participate in the market.

The second welfare theorem describes in very general terms what kind of policy might achieve such a Pareto-efficient allocation and what such a policy might do to the price of risk. Under the second welfare theorem, any Pareto-efficient allocation can be achieved as a market equilibrium after a set of transfers between individuals. The prices in that market equilibrium make all individuals unwilling to trade at the Pareto-efficient allocation. Such a policy would transfer consumption between individuals across states and times. After those transfers, under the market equilibrium, no individual would be willing to take on more or less risk at the market price. In other words, the price of risk would measure both market participants' and other government stakeholders' willingness to accept, on the margin, an additional unit of risk.

Thus, a Pareto-efficient policy is a set of state-contingent transfers that reallocate endowments between the two representative agents to achieve the solution to (9). That policy could consist of the government insuring groups of stakeholders against the idiosyncratic risks associated with specific states of the world or with the government taking on market risk by investing in risky securities.

One effect of such a policy is to diversify idiosyncratic risks for the market participant and the other stakeholder. Another is to shift systematic risk between them (see Figure 1). Under *laissez-faire*, the separate endowments of the two agents correlate with the total endowment of both together. But that correlation is not perfect. Certain states of the world are associated with relatively low consumption for the market participant but not for the other stakeholder, and others are associated with the opposite result.

Another effect of Pareto-efficient policy is to shift overall risk from the market participant to the other stakeholder (or the other way around), depending on which direction improves welfare. Those two effects of policy can be seen through the difference between the endowments of the two representative agents and their consumption under Pareto-efficient policy. The endowment of the market participant in each state can be decomposed into a systematic component and an idiosyncratic component as follows:

$$\frac{\frac{\delta u^M(\widetilde{c}_{ts,t}^M)}{\delta c}}{\frac{\delta u^M(c_0^M)}{\delta c}} = \alpha + \beta \frac{\frac{\delta u^O(\widetilde{c}_{ts,t}^O)}{\delta c}}{\frac{\delta u^O(c_0^O)}{\delta c}} \quad (10)$$

$$\text{And } y_{ts}^M = \widetilde{c}_{ts}^M + \varepsilon_{ts}; \quad y_{ts}^O = \widetilde{c}_{ts}^O - \varepsilon_{ts},$$

where ε_{st} is an idiosyncratic component of each agent's endowment that is unrelated to aggregate consumption. The systematic levels of consumption \widetilde{c}_{ts}^M and \widetilde{c}_{ts}^O are the components ordered with total consumption.

Under equation (10), the parameters α and β are the results of a projection of the marginal utility of the market participant $\frac{\delta u^M(c_{ts,t}^M)}{\delta c}$ on the marginal utility of the other government stakeholder $\frac{\delta u^O(c_{ts,t}^O)}{\delta c}$.

$$\text{Under the Pareto-efficient allocation, } \frac{\frac{\delta u^M(c_{ts,t}^M)}{\delta c}}{\frac{\delta u^M(c_0^M)}{\delta c}} = \frac{\frac{\delta u^O(c_{ts,t}^O)}{\delta c}}{\frac{\delta u^O(c_0^O)}{\delta c}},$$

and therefore $\alpha = 0$, $\beta = 1$, and $\varepsilon_{ts} = 0$ for all s and t . The terms α , β and ε represent three deviations of the *laissez-faire* levels of consumption from that allocation. The term ε represents idiosyncratic risk in the division of consumption between the market participant and other government stakeholder under *laissez-faire*. The term α measures an inefficient allocation of risk-free consumption across time, and the term β an inefficient allocation of risk between the two stakeholders.

A positive value of ε_{ts} relates to a state in which the market participant is relatively well endowed and the other stakeholder poorly endowed, in a way that is independent of systematic risk. A negative value of ε_{ts} corresponds to the opposite case. Both participants will gain from trading consumption in states in which they are relatively well endowed for consumption in states in which their endowment is relatively small. Accordingly, the market participant transfers an amount equal to ε_{ts} to the other stakeholder in all cases in which $\varepsilon_{ts} > 0$ and receives a transfer of $-\varepsilon_{ts}$ in all cases in which $\varepsilon_{ts} < 0$.

This component, ε_{st} , is termed “idiosyncratic,” although it represents systematic risk for the market participant under *laissez-faire*. It is idiosyncratic from the perspective of the total endowment of both the market participant and the other government stakeholder. It creates excess variation in the consumption of the market participant under *laissez-faire* than what is related to their risks in common with the other stakeholder. A real-world example of an idiosyncratic risk of that sort might be a recession that affects the consumption of the current generation but does not having lasting effects on output and consumption that reach future generations.

The parameter α measures an inefficiency in the allocation of consumption over time that has nothing to do with uncertainty. Under a positive α , the marginal utility of consumption of the market participant is relatively high in at time t compared with that participant's marginal utility of consumption in time 0, implying that the participant's average marginal utility of consumption

is low at time t . The other government stakeholder's consumption is relatively high at time t . Both agents can gain in the case of a positive α by exchanging a risk-free security at time t for consumption at time zero.

The parameter β captures the amount of systematic risk faced by the market participant relative to the risk faced by the other government stakeholder. If β is greater than one, then the market representative's endowment will have more systematic risk than the other stakeholder's, the market participant will be more risk-averse, or both. Under that situation, the market cost of risk will be higher in *laissez-faire* than under Pareto-efficient policy and will go down as a result of Pareto-efficient policy shifting risk away from the market participant. If β is less than one, Pareto-efficient policy will do the opposite. It will shift risk from the other stakeholder to the market participant, raising the cost of risk in the process.

Under equation (11), the price of risk for the market participant after diversification and intertemporal substitution has taken place is a multiple β of the price of risk for the other government stakeholder; that is:

$$p_{\text{Risk}}^{\text{M}} = \beta p_{\text{Risk}}^{\text{O}} \quad (11)$$

Therefore, if the market participant shifts one unit of risk to the other government stakeholder, the market participant's utility will go up by β times the amount that the utility of the other stakeholder goes down. The market participant could shift risk by selling the risky security to the other government stakeholder and receiving a risk-free security in return. Under the incomplete markets assumption of this model, that sale is not possible on financial markets but can be enabled by the government.

The cost of risk under *laissez-faire* will be different from the cost under Pareto-efficient policy because of those two policy effects. The diversification effect will lower the cost of risk under Pareto-efficient policy relative to the cost under *laissez-faire* whereas the risk-shifting effect could go either way. If Pareto-efficient policy shifts risk from the market participant to the other stakeholder, then it would cause the cost of risk to be unambiguously lower than it would be under *laissez-faire* because the diversification and risk-shifting effects would be in the same direction. If, in contrast, Pareto-efficient policy shifts risk from the other stakeholder to the market participant, then the total effect of policy on the market cost of risk could be positive or negative. It would be positive when the effect of risk shifting outweighs the diversification effect, and negative when the opposite was true.

Government Investment, Risk Shifting, and Welfare

The parameter β represents the ratio of the cost of risk to the market participant to the cost of risk to the other government stakeholder. If $\beta > 1$, the government can improve everyone's welfare by transferring risk from the market participant to the other government stakeholder through purchases of the risky asset that are financed by sales of the risk-free security. Under the

opposite case, in which $\beta < 1$, Pareto-efficient policy improves welfare by doing the opposite, and the following analysis will work in reverse. The transfer of risk from market participants to other government stakeholders reduces the amount that market participants are exposed to risk, on net, and increases the risk held by other government stakeholders. In *laissez-faire*, the cost of risk for market participants will be higher than the cost of risk for other government stakeholders. The benefit of governmental risk taking is equal to market participants' reduced cost of risk.

In *laissez-faire* (in which the government does not transfer any risk), the marginal cost of risk for market participants and the marginal benefit of governmental risk taking exceed the marginal cost of risk to the government stakeholders. As more risk is taken, the gap between the marginal cost and benefit of risk taking narrows until Pareto-efficient policy is reached, at which point the marginal cost and benefit are equal. At that point, the cost of risk to the government equals the market price. The total net benefit of governmental risk taking at the point of Pareto-efficient policy can be expressed as a triangle encompassed by the y-axis and the two lines representing the marginal cost of risk for market participants and the marginal cost of risk for other government stakeholders (see Figure 2).

The benefits and costs of governmental risk taking are distinct from programmatic benefits and costs, and the term "Pareto-efficient policy" in this paper refers to the Pareto-efficient level of risk transfer, not to Pareto-efficient policy with respect to program activities. Risk taking and program activities are logically separate from each other. The benefits and costs of governmental risk taking can be clearly isolated by considering them against a benchmark in which the program exists, but a private entity bears the risk of the activity and receives fixed compensation from the government for doing so. The benefits and costs of program activities are distinct from the costs and benefits of risk taking; they are clearest when comparing the outcomes of a program against a different benchmark in which the program does not exist, and the private market holds the risk.

For example, consider student loan programs. Programmatic benefits and costs are the increase in enrollment that is enabled by student loans: the cost of education versus the higher productivity and better outcomes of those who can attend college because they took out student loans to help pay for it. In contrast, governmental risk taking in student loans benefits private investors by insuring them against the risk of the activity but places the cost of risk on government stakeholders. Those are not benefits for the borrower but rather for private investors who would have taken the risk of lending to students if the government had not done so.

The government could take risk by buying a risky security financed by the sale of a risk-free security. Its counterparty in such a transaction is always the market participant. The settlement of transactions produces a net payment from the government to the market participant in bad times, and a net inflow in good times. To finance or distribute those settlement cash flows, the

government taxes or transfers money to both individuals. The combined result of the settlement cash flows and the taxes and transfers is to pass risk from the market participant to the other government stakeholder.

Assume that the government sells x risk-free securities and purchases $(1 + \pi)x$ risky securities. That risk taking results in a net cash flow from the government to the market participant that varies by state. In states whose endowments exceed the average endowment for all states, there will be a net inflow to the government; in the state with below-average endowments, there will be a net outflow. The government's sale and purchase in such a case has the effect of exposing the government finances to risk and, at the same time, reducing the amount of risk that the market participant faces through its own wealth.

The government passes the state-contingent cash flows, with a set of state-contingent transfers, to both the market participant and the other government stakeholder. Those transfers can be made by using many different formulas. Any formula that transfers some of those state-contingent cash flows to the other government stakeholder will have a similar general effect. The formula will lead to a net transfer of risk from the market participant to the other government stakeholder. All the risk in the government's purchase of risky securities comes from the market participant and none from the other government stakeholders; in contrast, some of the exposure to that risk through government finances falls on the other government stakeholders. So, on net, government investment in risky assets reduces risk for market participants and raises risk for other stakeholders.

Defining the share of transfers that fall upon other government stakeholders as ρ , the net transfer from market participants to other government stakeholders is the following function of x :

$$\tau(s, x) = \rho x(c_{st} - \bar{c}_t) \quad (12)$$

That transfer affects the market price of risk, as observed in the following version of (9), with c_{ts}^M set equal to the endowment of the market participant plus the transfer associated with risk shifting.

$$p_t^{\text{Risk}} = -\phi_0 \text{Cov} \left(\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c}, \frac{y_{ts}^M - \rho x(c_{ts} - \bar{c}_t)}{\bar{c}_t} \right), \frac{\delta u^M(c_0^M)}{\delta c} \quad (13)$$

The derivative of the risk premium with respect to x is as follows:⁵

⁵ See Appendix for derivation.

$$\frac{\delta p_t^{\text{Risk}}}{\delta x} = \frac{\phi_0 \rho}{\frac{\delta u^M(c_0^M)}{\delta c} \bar{c}_t} E \left[\frac{\delta^2 u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c^2} (c_{ts} - \bar{c}_t)^2 \right] \quad (14)$$

That derivative is always less than zero given that $\frac{\delta^2 u^M}{\delta c^2} < 0$ and $(c_{ts} - \bar{c}_t)^2 > 0$. Risk transfers from the market participant to the other government stakeholder will reduce the market risk premium.

Factors That Determine the Direction of Risk Shifting Under Pareto-Efficient Policy

Government stakeholders who cannot participate in the markets fall mainly into two categories: members of the current generation who do not have financial wealth, and future generations. Whether Pareto-efficient policy shifts risk to or away from those stakeholders will depend on their level of risk aversion and how much systematic risk is in their endowment relative to the endowments of market participants. Members of the current generation without financial assets are likely to be more risk-averse than market participants but are also likely to be less exposed to market risk. Members of the current generation who lack financial wealth will generally be less exposed to market shocks than more wealthy households.

Risk will shift from the more risk-averse agent to the less risk-averse agent, all else being equal. So, for example, if other government stakeholders are more risk-averse than market participants, Pareto-efficient policy will move more risk toward government stakeholders. For this purpose, an agent is more risk-averse than another agent if the utility function of the other agent can be expressed as a concave function of the first agent's utility function.

If an agent is endowed with a larger amount of risk, then more of that risk will be shifted from the first to the second agent (or less risk will be shifted away). Specifically, suppose that the market participant's endowment of systematic risk is a mean-preserving shift of a benchmark case. It follows that the other stakeholder's endowment will be a mean-preserving contraction of what systematic risk under the benchmark case. In that situation, the β of the market participant will be higher, and more risk will be shifted from the market participant to the other participant.

Depending on who they are, other government stakeholders may be more or less risk-averse than market participants. If government stakeholders who cannot participate in today's financial markets are primarily of future generations, then they may be relatively wealthy and further away from subsistence, in which case they will be less risk averse. In contrast, if other government stakeholders are alive but not participating in financial markets, then they will likely be relatively poor and close to subsistence, in which case they might be more risk averse. (It is also possible that aversion to risk is a cause, as well as a consequence, of low levels of wealth).

The Cost of Risk for Overlapping Generations With Uncertainty

Government policies not only share risk more effectively, as explored in the previous section, but also can reduce risk to aggregate consumption by smoothing it over time. The overlapping generations framework is ideal for capturing the cost of government risk in a context in which government policy facilitates such intertemporal smoothing. This paper builds on the model specified in Diamond (1965), adding the possibility of shocks to the level of productivity. Under this framework, government policy can counter the effect of negative shocks by implementing policies that change the level of savings and investment, which have the effect of smoothing lifetime consumption and sharing risk more efficiently between generations.

Under the overlapping generations model, the market price of risk depends on the young generation's expectations of its marginal utility of consumption in each state when it is old and investment returns are realized. The implications depend on the level of risk aversion of the representative agents in the model and on whether shocks to technology are temporary or permanent. Under central estimates of risk aversion, temporary shocks to the path of technology cause the price of risk to be higher under *laissez-faire* than it would be under Pareto-efficient policy and higher than the cost (or risk) for the following generation, to whom the government could transfer risk. Permanent shocks have the opposite effect. They cause the market price of risk to be lower than is consistent with Pareto-efficient government policy.

Notation and Assumptions

Each generation is represented by someone who lives through two periods. In the first period, that person works, saves, and invests; in the second period, he or she lives off any investments and any transfers that might be received. Successive generations overlap for one period.

Let c_t^y represent the consumption of the younger generation in period t , and c_t^o represent the consumption of the older generation in that same period. All consumers have constant relative risk aversion (CRRA) utilities and a discount factor β such that the total utility of the generation that is young in period t is as follows:

$$u(c_t^y, c_{t+1}^o) = \frac{1}{1-\theta} \left[c_t^{y1-\theta} + \beta c_{t+1}^{o1-\theta} \right] \quad (15)$$

The parameter θ represents risk aversion and the desire for smooth consumption between youth and old age for each generation. Each representative individual is endowed with one unit of labor when he or she is young and can invest in capital that fully depreciates over one period. The production function is Cobb Douglas:

$$f(k_t) = A_t k_t^\omega \quad (16)$$

with the term A_t representing the level of technology, which is subject to uncertainty.

Because income shares of each factor are constant under Cobb Douglas technology, in equilibrium total wages in *laissez-faire* are $(1 - \omega)A_t k_t^\omega$ and capital income is $\omega A_t k_t^\omega$. Wages go entirely to the younger generation and capital income to the older. Wages fund the younger generation's consumption and its investment in the next period's capital as well as its purchase of risk-free securities b_t :

$$(1 - \omega)A_t k_t^\omega = c_t^y + k_{t+1} + b_t \quad (17)$$

The younger generation at time t will divide its income between consumption and investment in the next period's capital. To capture the market price of risk, the analysis will assume that it also has access to risk-free securities returning r_t , but under *laissez-faire* the total supply of b_t will be set to zero. When the younger generation reaches old age in period $t + 1$, its consumption comes from its investment in risk-free securities and productive capital.

$$c_{t+1}^o = \omega A_{t+1} k_{t+1}^\omega + b_t r_t \quad (18)$$

The Price of Risk Under *Laissez-faire*

The younger generation will thus solve the following Lagrangian:

$$\begin{aligned} \max_{c_t^y, c_{t+1}^o, b_t, k_{t+1}} \int \frac{1}{1-\theta} \left\{ c_t^{y1-\theta} + \beta E_t [c_{t+1}^{o1-\theta}] \right\} + \lambda_{1t} [A_t k_t^\omega - c_t^y - b_t - k_{t+1}] \\ + \lambda_{2t} [E_t (A_{t+1} k_{t+1}^\omega) + b_t r_t - c_{t+1}^o] \end{aligned} \quad (19)$$

The first-order conditions for that problem result in the following relationships:

$$1 = \beta E_t \left[\left(\frac{c_{t+1}^o}{c_t^y} \right)^\theta R_{t+1} \right] \quad (20)$$

in which $R_{t+1} \equiv \omega A_{t+1} k_{t+1}^{\omega-1}$ and

$$r_t = \beta^{-1} E_t \left[\left(\frac{c_{t+1}^o}{c_t^y} \right)^\theta \right] \quad (21)$$

Those two expressions define the expected return on capital and the risk-free rate. The risk premium is the difference between $E_t[R_{t+1}]$, the expected return on capital, and the risk-free rate, r_t . After some rearranging:

$$E_t[R_{t+1}] - r_t = \frac{\text{Cov}(R_{t+1}, c_{t+1}^{o\theta})}{E_t(c_{t+1}^{o\theta})} \quad (22)$$

That equation says that the risk premium is equal to the ratio of the covariance of the return on capital to the marginal utility of consumption, scaled by the average marginal utility. The risk premium in *laissez-faire* will depend on the level of capital k that determines the average level of

return and any uncertainty about the technological factor A . Those factors determine how much those returns might vary.

The Price of Risk Under Pareto-Efficient Policy

Equation (22) shows that the risk premium will depend on how much the older generation's consumption and investment returns covary. Under *laissez-faire*, consumption of the older generation and investment returns are one and the same. Under Pareto-efficient policy, the older generation's consumption might come partly from government transfers and will not be entirely tied to investment returns. That does not necessarily mean that consumption will be less risky.

Assume that Pareto-efficient policy maximizes a weighted sum of the utility of different generations, with no assumption about those weights other than that they are positive. Let w_t be the weight on the utility of the generation born in period t . Pareto-efficient policy is based on the solution to the following problem:

$$\begin{aligned} \max_{c_t^y, c_t^o, k_{t+1} \forall t} \frac{1}{1-\theta} \left\{ w_0 c_t^{o1-\theta} + \sum_{t=1}^{\infty} w_t \left[c_t^{y1-\theta} + \beta E_t [c_{t+1}^o]^{1-\theta} \right]^{\frac{1-\chi}{1-\theta}} \right\} \\ \text{s. t. } c_t^y + c_t^o + k_{t+1} \leq A_t k_t^\omega \quad \forall t \end{aligned} \quad (23)$$

The constraint requires that consumption by the younger generation, consumption by the older generation, and investment in the next period's capital do not exceed current output.

The terms χ and θ are both risk-aversion parameters, but they measure different aspects of preferences using the utility function developed by Epstein and Zinn (1989), which separates preferences for intertemporal smoothing and risk aversion. In that case, they relate to diminishing returns to progressively higher consumption throughout a lifetime versus preferences for equal consumption throughout that lifetime. The term θ , from the utility function of each generation, measures the desire for smooth consumption across those generations' youth and old age. The term χ , in contrast, measures the aversion to unequal lifetime consumption by those generations. In the case that $\chi = \theta$, the utility function is additive across a combination of generations and time periods.

Pareto-efficient policy will distribute total consumption at time t between the younger and the older generation according to the following formula:

$$\frac{c_t^y}{c_t^o} = \left(\frac{w_t}{\beta w_{t-1}} \left[\frac{U_t}{U_{t-1}} \right]^{\frac{\theta-\chi}{1-\theta}} \right)^{\frac{1}{\theta}} \quad (24)$$

in which $U_t \equiv c_t^{y1-\theta} + \beta E_t [c_{t+1}^o]^{1-\theta}$ is the expected lifetime utility of the generation that is young in period t . In the situation that $\chi = \theta$, the right side of the expression simplifies to a

constant. In that case, the younger and older generations have a fixed share of consumption in the period in which they are both alive. In other instances, the share of consumption varies on the basis of the relative lifetime utilities of the two generations.

Total consumption and investment react to shocks in the technology component A_{t+1} , according to the Euler equation:⁶

$$\lambda_t = \beta E_t[\lambda_{t+1} \omega A_{t+1} k_{t+1}^{\omega-1}] \quad (25)$$

In that equation, the term λ_t represents the shadow price of the budget constraint at time t , and thus the marginal utility of consumption of either generation in that time period. It says that the marginal utility of consumption of a dollar at time t equals the expected marginal utility of that same dollar in the next period if it has been invested in capital and is yielding a return.

The risk premium under Pareto-efficient policy will depend on how much the marginal utility of consumption by the generation that is old at time $t+1$ varies with the return on capital. Consider that consumption by that generation can be decomposed as follows:

$$c_{t+1}^o = \frac{c_{t+1}^o}{c_{t+1}} \frac{c_{t+1}}{y_{t+1}} y_{t+1} \quad (26)$$

That is, the consumption of the older generation c_{t+1}^o is equal to its share of total consumption $\frac{c_{t+1}^o}{c_{t+1}}$ times consumption as a share of output (one minus the savings rate) $\frac{c_{t+1}}{y_{t+1}}$ times output y_{t+1} .

The older generation's consumption will be less volatile if the two consumption shares offset the volatility coming from output; that is, if consumption as a share of gross domestic product (GDP) rises after a negative shock to output, the older generation's share of consumption rises, or both. The older generation's consumption will be more volatile if the opposite occurs.

The price of risk will depend on those relationships as well. That price will go up under Pareto-efficient policy if it causes the consumption of the older generation to become more volatile, and it will go down if Pareto-efficient policy causes it to be less volatile. The price of risk will thus go down if consumption rises as a share of GDP because of a negative shock, if the older generation's share of consumption rises, or both. Those effects on risk premiums are evident in

⁶ All shocks in this model are in technology, which aligns the analysis with the view of fluctuations taken in the theory of real business cycles. Fluctuations might also come from shocks to the demand in combination with nominal rigidities, as explored in the neo-Keynesian theory. The findings of this analysis for temporary shocks are most likely also relevant for such demand shocks, although future research is needed to confirm that.

the following approximation of the risk premium, which uses the approximation formula for the covariance of products of random variables in Bohrnstedt and Goldberger (1969).⁷

$$\text{Price of Risk} = \frac{\text{Cov}\left(R_{t+1}, \left(\frac{c_{t+1}^o}{c_{t+1}} \frac{c_{t+1}}{y_{t+1}} y_{t+1}\right)^\theta\right)}{E_t(c_{t+1}^o)^\theta} \approx \frac{\text{Cov}(R_{t+1}, (y_{t+1})^\theta)}{E_t(y_{t+1}^\theta)} + \frac{\text{Cov}\left(R_{t+1}, \left(\frac{c_{t+1}^o}{c_{t+1}}\right)^\theta\right)}{E_t\left(\left[\frac{c_{t+1}^o}{c_{t+1}}\right]^\theta\right)} + \frac{\text{Cov}\left(R_{t+1}, \left(\frac{c_{t+1}}{y_{t+1}}\right)^\theta\right)}{E_t\left(\left[\frac{c_{t+1}}{y_{t+1}}\right]^\theta\right)} \quad (27)$$

Under *laissez-faire*, the older generation's consumption is a constant share of output ω , so the price of risk is equal to the first term in equation (27). Therefore, the difference between the price of risk under *laissez-faire* and Pareto-efficient policy relates to the second and third terms. Those terms relate to how the older generation's share and rate of consumption respond to Pareto-efficient policy.

A shock in A_{t+1} when the system is in steady state leads consumption and capital to follow a transition path to a new steady state. The nature of those paths is dependent on the intertemporal elasticity of substitution, the inverse of θ , the relationship between χ and θ , and whether shocks are permanent.

Additive Preferences. In the case in which $\chi = \theta$, the younger and older generations have a fixed share of consumption during the period when they are both alive. In that situation, the only effect on the risk premium comes from the change in the saving rate. The intertemporal elasticity of substitution determines the relative size of the wealth and substitution effects that compose the saving response.

Under central estimates of the intertemporal elasticity of substitution, temporary and permanent shocks have opposite implications for the saving rate and thus for risk premiums. Empirical estimates of θ , the inverse of the elasticity, are usually greater than one, although they vary widely (see Hall (1998), Attanasio and Weber (1993), Scholz, Seshadri and Khitatrakun (2006), and Engelhardt and Kumar (2009)). Those estimates suggest that the intertemporal elasticity of substitution is less than one. Under such preferences, a one-period downward shock to A_{t+1} would lower the saving rate under Pareto-efficient policy. Such policies would blunt the effect of the shock on the older generation's consumption. Under a set of government policies that

⁷ See Appendix for derivation.

generated such a Pareto-efficient response, investors would demand lower premiums in compensation for temporary shocks than they would demand under *laissez-faire* (see Table 1).

A permanent downward shock would have the opposite effect. It would raise the saving rate, magnifying the initial shock's effect on consumption and raising the level of risk facing investors. It would raise the saving rate because a shock to technology would lead to an even larger effect on output in later periods than in the initial period in which it occurs because capital would be shed in the transition to a new steady state. The fact that the effect builds over time, outside the time horizon of current investors, causes future generations to be more exposed to the shock than current generations and thus to demand an even higher risk premium for that shock.

The implications of this analysis depend heavily on whether risk premiums arise primarily from the threat of permanent shocks or temporary shocks. If risk premiums arise primarily from permanent shocks, then the cost of risk for future generations could be higher than the market price of that risk. Studies of the time-series properties of GDP, starting with Nelson and Plosser (1982), generally fail to reject the possibility that they would follow a unit root, which would imply that shocks to GDP are permanent. Under Cochrane's (1990) nonparametric analysis, about half of the fluctuations in GDP seem to be permanent, with temporary fluctuations wearing off over a generation.⁸ Moreover, the key question for interpreting this analysis is not whether all or even most shocks to the economy are permanent, but rather whether the risk premium can be accounted for with the subset of shocks that are permanent. Bansal and Yaron (2005) explain the equity risk premium puzzle of Mehra and Prescott (1985) with a small component of shocks that lead to persistent changes in the growth rates of GDP. Under that model, a drop in GDP growth not only leads to a permanent drop in the GDP trend, but also is a harbinger of additional years of low growth in GDP; each year produces a permanent drop in the level of GDP. Those authors assume that the economy is also affected by a temporary shock that accounts for most of the year-to-year variance in GDP, but that variance does not generate a significant risk premium.

The Bansal and Yaron (2005) study is part of a larger literature that attributes risk premiums to persistent shocks to the trend of growth in economic activity. Hansen and Sargent (2021) argue that tenuous beliefs and ambiguity about those shocks can make the premiums associated with long-run risks larger than they would be if the parameters governing permanent shocks were known with certainty. Alvarez and Jermann (2004, 2005) offer additional evidence that permanent fluctuations in consumption account for most of the observed risk premiums.

An opposing view is held by Campbell and Cochrane (1999), who argue that habit formation can explain why small temporary fluctuations can have outsized effects on risk premiums. In their

⁸ An update of Cochrane's calculations using three additional decades of data did not significantly change those results.

model, utility is a function of surplus consumption beyond some threshold, as in the endowment model in this study. The authors assume a time-varying threshold that gravitates toward recent consumption levels. As a result, consumers are strongly averse to risk of consumption falling below recent levels. Their approach was expanded in Boldrin, Christiano and Fisher (2001), which embedded habit formation in a real business cycle model. Besides confirming the ability of habit formation to explain the equity risk premium puzzle, their analysis showed that habit formation could explain many stylized facts about business cycles and their relationship to asset prices. Santos and Veronesi (2010) offered a more skeptical view of habit formation, arguing that its implications for the cross section of risk premiums were inconsistent with the estimated premiums on “value” and “growth” stocks.

Because a time period in the overlapping generations model is a generation, the term “temporary” has a different meaning in this paper than it generally has in macrofinance literature. A temporary shock could affect an entire generation and therefore might be considered a persistent shock in another context. Those assumptions were adapted from real business cycle models in which “temporary” shocks generally do not last more than a business cycle. Moreover, the literature on the equity risk premium puzzle usually views temporary shocks as those that last for a short period.

That difference in time frame complicates the interpretation of the model’s findings, in particular of the finding that permanent shocks lead market risk premiums to be lower under *laissez-faire* than they would be under Pareto-efficient policy. That finding follows from an amplification process. A permanent shock to total factor productivity leads to a period of transition to a new steady-state level of capital that amplifies the initial shock to technology’s consumption and output. In the short run, only the technology shock affects capital. In the long run, the change in steady-state capital affects output in the same direction, amplifying the initial shock. If growth in total factor productivity stalls for only one year, the trend in total factor productivity will shift downward, and the long-run effect on output will be realized within a few decades under a model with empirically estimated responses in investment, such as Jorgenson (1963), or adjustment costs, such as Tobin (1969). Thus, future generations will be exposed to risk about as equally as the current generation. In that situation, the premiums for market risk will not much overstate the cost of permanent shocks to future generations.

Alternatively, if total factor productivity growth falls in a persistent way, then the cumulative effect of that development might be amplified over generations. Robert Gordon (2016) suggests that productivity can stall persistently when a revolution in productivity peters out. The model of Bansal and Yaron (2005) is also consistent with amplification over generations. The authors find that the persistent component of growth rates has a serial correlation of 0.98 on an annual basis. A shock that decays by a factor of 0.98 annually would decay by about one-half over 35 years, and have a long-run effect on the level of GDP that would be almost twice as large as the effect on GDP during the lives of the current generation.

Nonadditive Preferences. If either $\chi > \theta$ or $\chi < \theta$, then preferences become nonadditive across generations and time periods. Each of those cases introduces important considerations for how the cost of risk in *laissez-faire* might compare with its cost under Pareto-efficient policy.

The parameter χ has a slightly different interpretation in this analysis than it does in the infinitely lived representative-agent framework in which it is typically used. In that representative-agent framework, the parameter χ represents the agent's preferences between risks to its lifetime consumption that are separate from its preference for smoothed consumption. In this model, the parameter χ is in the optimization problem governing Pareto-efficient policy but does not enter into the decisions of each generation as it evaluates its own choices. The parameter represents how much happier each generation might have been if it had been born with a different opportunity set than the one it actually has, and not how that generation chooses from within the opportunity set it does have. The importance of χ can be illustrated by the preferences related to the following pair of choices:

- **Falling consumption in old age.** Two generations have an equal amount of lifetime consumption, but both experience a 10 percent drop in consumption in their old age.
- **Inequality between generations.** Both generations have an equal amount of consumption in old age and youth, but one generation's consumption is 10 percent lower than the other's in both periods.

Under a situation in which $\chi > \theta$, a 10 percent drop in old-age consumption by two separate generations would have less effect on total weighted utility a 10 percent drop between generations, which produces inequality between them. That preference for equal outcomes among generations would match the estimated parameters of the “risks for the long run” paper by Bansal and Yaron (2005). In a situation in which $\chi < \theta$, inequality between generations would have a lower welfare cost. In that situation, it would be better to have unequal consumption between generations than to have both generations experience a drop in consumption. That situation would match in spirit the habit-formation preferences of Campbell and Cochrane (1999). Under habit formation, a drop in consumption from youth to old age for both generations would be more costly than having one generation be relatively poor from the beginning and never know what it was missing.

Following Weil (1990), this analysis explores the case in which the shock to technology (to A_{t+1}) is log-normally distributed. As he demonstrated, little can be said in general about the case of $\chi \neq \theta$ without such a distributional assumption. Under log normality, the relative levels of χ and θ affect the magnitude of the effects of Pareto-efficient policy but not their sign (see Table 2). That is, permanent shocks cause the price of risk to be higher under Pareto-efficient policy than under *laissez-faire* regardless of whether $\chi > \theta$ or $\chi < \theta$. Instead, the price of risk changes by an amount that varies depending on $\chi - \theta$. Temporary shocks cause the price of risk

to be lower under Pareto-efficient policy than under *laissez-faire*, but by an amount that varies with the relative levels of χ and θ .

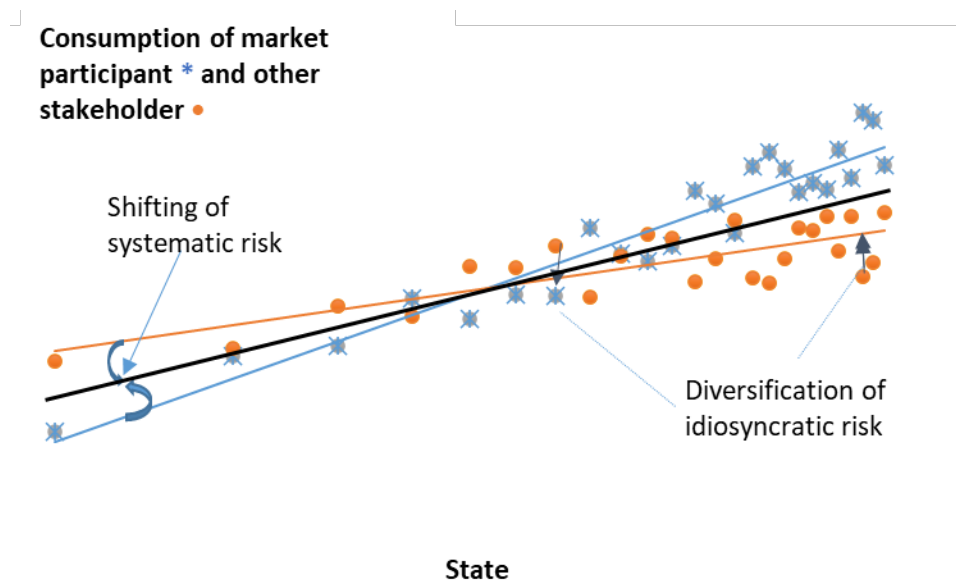
In the case that $\chi > \theta$, Pareto-efficient policy will tend to shift more of the burden of a permanent shock to the generation that is old at time $t + 1$. It does so because in the case that $\chi > \theta$, inequality between generations is more costly than a drop in consumption between youth and old age for a given generation. The older generation will be partly spared the effect of the shock on its lifetime consumption, whereas the younger generation will be fully exposed to that shock. Pareto-efficient policy will shift even more risk to the older generation at time $t + 1$ than would be the case in which $\chi = \theta$. Anticipating the effect of that policy at time t , when it is young, that generation will require an even larger risk premium to invest in risky assets.

Under the opposite case, where $\chi < \theta$, Pareto-efficient policy will be more concerned with smoothing consumption between youth and old age of each generation than reducing inequality between generations. They will shift less of the risk of a permanent shock to the generation that is old at time $t + 1$ than under the case in which $\chi = \theta$, to avoid as large a drop in consumption of the older generation. As a result, the older generation bears less of the risk of a permanent shock than under the case in which $\chi = \theta$, but still bears more than it would under *laissez-faire*. As a result, it will charge a risk premium that is between the *laissez-faire* case and the one in which $\chi = \theta$.

A similar logic applies to temporary shocks. Pareto-efficient policy shifts the risk of temporary shocks from the older to the younger generation at $t + 1$ if the intertemporal elasticity of substitution is less than one. That risk shifting lowers the risk premium by reducing the volatility of consumption faced by the generation that is old in time t . In the case of $\chi > \theta$, Pareto-efficient policy will shift relatively less risk to the generation that is young at time $t + 1$, and it will shift relatively more under the case in which $\chi < \theta$, than it would under the case in which $\chi = \theta$. Thus, risk premiums associated with temporary shocks will go down by less under Pareto-efficient policy if $\chi > \theta$ and by more if $\chi < \theta$.

Figures

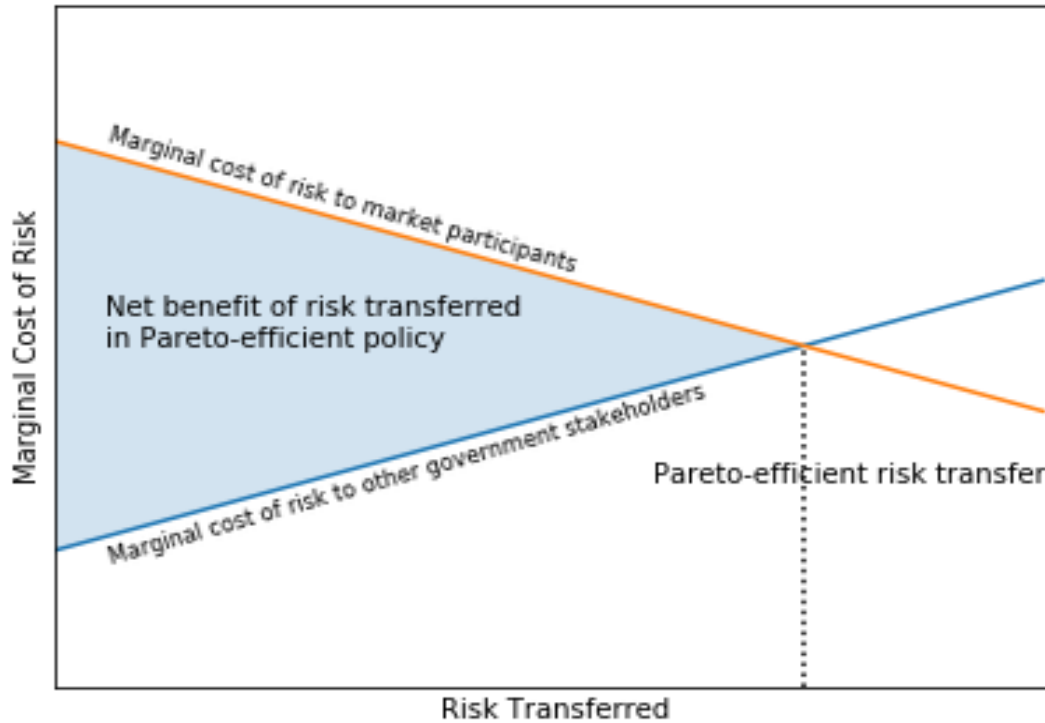
Figure 1.
Diversification and Risk Shifting Between Market Participant and Other Government Stakeholder



Data source: Congressional Budget Office.

The figure depicts a possible allocation of consumption of the market participant and other stakeholder under *laissez-faire*. Under that example, the market participant is more exposed to systemic risk than the other stakeholder, as that participant's consumption rises with total endowment at a higher slope. The two agents are exposed to equal and opposite idiosyncratic risks, represented by the gaps between the consumption in each state and the lines representing the systematic component of their consumption. Pareto-efficient policy would lead to diversification in which those equal and opposite gaps are traded off. Following diversification, risk shifting could be Pareto-improving and is represented by the shifting of the line slopes.

Figure 2.
The Benefits and Costs of Risk Transfer Associated With Governmental Risk Taking



Data source: Congressional Budget Office.

This figure depicts a case in which market participants are endowed with more risk than other government stakeholders. If those stakeholders were endowed with more risk than market participants, risk shifting would occur in the opposite direction, and the labels for market participants and other government stakeholders would be reversed.

Tables

Table 1.
The Effects of Permanent and Temporary Shocks, Assuming Central Estimates of the Intertemporal Elasticity of Substitution

	Effect on Savings Rate	Effect on Uncertainty of Consumption of Young Generation in Next Period	Implications for Gap Between the Cost of Risk to Future Generations and Market Price
Temporary Shocks to A_{t+1}	A negative shock lowers the savings rate.	Uncertainty of consumption is lower under Pareto-efficient policy than under <i>laissez-faire</i> .	Leads the cost of risk to future generations to be lower than the market price under <i>laissez-faire</i> .
Permanent Shocks to A_{t+1}	A negative shock raises the savings rate.	Uncertainty of consumption is higher under Pareto-efficient policy than under <i>laissez-faire</i> .	Leads the cost of risk to future generations to be higher than the market price under <i>laissez-faire</i> .

Data source: Congressional Budget Office.

The intertemporal elasticity of substitution is the inverse of the parameter θ from the representative consumer's utility function and is less than one under central estimates. If $\theta > 1$, the intertemporal elasticity of substitution is less than 1, and those implications hold. If $\theta < 1$, the intertemporal elasticity of substitution is greater than 1, and those implications are reversed, and if $\theta = 1$ (log utility), they are nullified because the saving rate is constant. Empirical estimates of θ are usually greater than 1, although they range widely. See Hall (1998), Attanasio and Weber (1993), Scholz, Seshadri and Khitatrakun (2006), Engelhardt and Kumar (2009).

Table 2.
The Effects of Different Risk-Aversion Preferences on Risk Premiums

	Implications for the Decrease in the Price of Risk Associated with Temporary Shocks	Implications for the Increase in the Price of Risk Associated with Permanent Shocks
Risk Aversion χ is larger than preference for intertemporal smoothing θ .	Pareto-efficient policy lowers risk premiums by less than under the case in which $\chi = \theta$.	Pareto-efficient policy raises risk premiums by more than under the case in which $\chi = \theta$.
Risk Aversion χ is smaller than preference for intertemporal smoothing θ .	Pareto-efficient policy lowers risk premiums by more than under the case in which $\chi = \theta$.	Pareto-efficient policy raises risk premiums by less than under the case in which $\chi = \theta$.

Data source: Congressional Budget Office.

Appendix: Derivation of Selected Equations

The following is a derivation of three equations: (7), (14) and (27):

Equation (7) is derived by combining two equations, equation (2)

$$\frac{1}{\phi_0} \frac{\delta u^k(c_0^k)}{\delta c} = \frac{\pi_{ts}}{\phi_{ts}} \frac{\delta u^k(c_{ts}^k, t)}{\delta c}$$

with equation (6)

$$p_t^{\text{Risk}} = p_t^{\text{Risk Free}} - p_t^{\text{Risky}} = \sum_s \phi_{ts} \frac{\bar{c}_t - c_{ts}}{\bar{c}_t}$$

as follows,

$$p_t^{\text{Risk}} = \phi_0 \sum_s \pi_{ts} \frac{\frac{\delta u^k(c_{ts}^k, t)}{\delta c}}{\frac{\delta u^k(c_0^k)}{\delta c}} \frac{\bar{c}_t - c_{ts}}{\bar{c}_t}$$

using equation (2) to substitute for ϕ_{ts} . That expression can be modified as follows:

$$p_t^{\text{Risk}} = \phi_0 \sum_s \pi_{ts} \frac{\left\{ \frac{\delta u^k(c_{ts}^k, t)}{\delta c} - E \left[\frac{\delta u^k(c_{ts}^k, t)}{\delta c} \right] \right\} + E \left[\frac{\delta u^k(c_{ts}^k, t)}{\delta c} \right]}{\frac{\delta u^k(c_0^k)}{\delta c}} \frac{\bar{c}_t - c_{ts}}{\bar{c}_t}$$

$$p_t^{\text{Risk}} = \phi_0 \sum_s \pi_{ts} \frac{\left\{ \frac{\delta u^k(c_{ts}^k, t)}{\delta c} - E \left[\frac{\delta u^k(c_{ts}^k, t)}{\delta c} \right] \right\}}{\frac{\delta u^k(c_0^k)}{\delta c}} \frac{\bar{c}_t - c_{ts}}{\bar{c}_t} + \phi_0 \frac{E \left[\frac{\delta u^k(c_{ts}^k, t)}{\delta c} \right]}{\frac{\delta u^k(c_0^k)}{\delta c}} \sum_s \pi_{ts} \frac{\bar{c}_t - c_{ts}}{\bar{c}_t}$$

Because $\sum_s \pi_{ts} c_{ts} = \bar{c}_t$, $\sum_s \pi_{ts} \frac{\bar{c}_t - c_{ts}}{\bar{c}_t} = 0$. The second term disappears, leaving equation (7):

$$p_t^{\text{Risk}} = -\phi_0 \sum_s \pi_{ts} \left(\frac{\frac{\delta u^M(c_{ts}^M, t)}{\delta c} - \overline{\frac{\delta u^M(c_{ts}^M, t)}{\delta c}}}{\frac{\delta u^M(c_0^M)}{\delta c}} \right) \left(\frac{c_{ts} - \bar{c}_t}{\bar{c}_t} \right)$$

Equation (14) is the first derivative of the cost of risk with respect to the level of risk transfer x . The cost of risk conditional on x is given by equation (13):

$$p_t^{\text{Risk}} = -\phi_0 \text{Cov} \left(\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c}, \frac{c_{ts}}{\bar{c}_t} \right)$$

Given that $\frac{\delta u^M(c_0^M)}{\delta c} = \phi_0$ and \bar{c}_t is constant with respect the state s , we can rewrite the equation as

$$\begin{aligned} & \frac{-\phi_0}{\frac{\delta u^M(c_0^M)}{\delta c} \bar{c}_t} \left\{ \text{Cov} \left(\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c}, c_{ts} \right) \right\} \\ &= \frac{-1}{\bar{c}_t} \left\{ \text{Cov} \left(\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c}, c_{ts} \right) \right\} \end{aligned}$$

Based on the formula for covariance $\text{Cov}(x, y) = E[xy] - E[x]E[y]$,

$$\begin{aligned} &= \frac{-1}{\bar{c}_t} \left\{ E \left[\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c} c_{ts} \right] - E \left[\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c} \right] \bar{c}_t \right\} \\ &= \frac{-1}{\bar{c}_t} E \left[\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c} (c_{ts} - \bar{c}_t) \right] \end{aligned}$$

Now, the derivative of the risk premiums with respect to x ,

$$\frac{\partial}{\partial x} \frac{-1}{\bar{c}_t} \left\{ \text{Cov} \left(\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c}, c_{ts} \right) \right\}$$

is equal to

$$\begin{aligned} &= \frac{\partial}{\partial x} \frac{-1}{\bar{c}_t} E \left[\frac{\delta u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c} (c_{ts} - \bar{c}_t) \right] \\ &= E \left[\frac{\delta^2 u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c^2} \right] (-\rho(c_{ts} - \bar{c}_t)(c_{ts} - \bar{c}_t)) \end{aligned}$$

by the chain rule. That equation simplifies to equation (14):

$$= \rho E \left[\frac{\delta^2 u^M(y_{ts}^M - \rho x(c_{ts} - \bar{c}_t), t)}{\delta c^2} (c_{ts} - \bar{c}_t)^2 \right]$$

Equation (27) applies the approximation formula in Bohrnstedt and Goldberger (1969) to the covariance of products of random variables.

$$\text{Cov}[xy, uv] \approx E[x]E[u]\text{Cov}[y, v] + E[x]E[v]\text{Cov}[y, u] + E[y]E[u]\text{Cov}[x, v] + E[y]E[v]\text{Cov}[x, u]$$

In this study, the formula is used to approximate the covariance of one variable with the product of three, as in $\text{Cov}[x, uwz]$. The approximation is accomplished by setting y to one and v to wz . If $y = 1$, then $\text{Cov}[y, v]$ and $\text{Cov}[y, u]$ both equal 0, and the expression above simplifies to

$$\text{Cov}[x, uv] \approx E[u]\text{Cov}[x, v] + E[v]\text{Cov}[x, u] .$$

$$\text{Similarly, } \text{Cov}[x, v] = \text{Cov}[x, wz] \approx E[w]\text{Cov}[x, z] + E[z]\text{Cov}[x, w],$$

and note that $E[v] = E[wz] = E[w]E[z] + \text{Cov}[w, z]$.

Substituting produces the following expression:

$$\begin{aligned} \text{Cov}[x, uwz] &\approx E[u]\{E[w]\text{Cov}[x, z] + E[z]\text{Cov}[x, w]\} + \{E[w]E[z] + \text{Cov}[w, z]\}\text{Cov}[x, u] \\ &\approx E[u]E[w]\text{Cov}[x, z] + E[u]E[z]\text{Cov}[x, w] + E[w]E[z]\text{Cov}[x, u], \end{aligned}$$

dropping the term $\text{Cov}[w, z] \text{Cov}[x, u]$ in favor of symmetry.

That approximation is applied to approximate $\text{Cov}\left(R_{t+1}, \left(\frac{c_{t+1}^o}{c_{t+1}} \frac{c_{t+1}}{y_{t+1}} y_{t+1}\right)^\theta\right)$ by setting x to R_{t+1} , and u, w and z to $\left(\frac{c_{t+1}^o}{c_{t+1}}\right)^\theta$, $\left(\frac{c_{t+1}}{y_{t+1}}\right)^\theta$, and $(y_{t+1})^\theta$, respectively.

$$\begin{aligned} &\text{Cov}\left(R_{t+1}, \left(\frac{c_{t+1}^o}{c_{t+1}} \frac{c_{t+1}}{y_{t+1}} y_{t+1}\right)^\theta\right) \\ &\approx E\left[\left(\frac{c_{t+1}^o}{c_{t+1}}\right)^\theta\right] E\left[\left(\frac{c_{t+1}}{y_{t+1}}\right)^\theta\right] \text{Cov}(R_{t+1}, (y_{t+1})^\theta) \\ &+ E\left[\left(\frac{c_{t+1}}{y_{t+1}}\right)^\theta\right] E[(y_{t+1})^\theta] \text{Cov}\left(R_{t+1}, \left(\frac{c_{t+1}^o}{c_{t+1}}\right)^\theta\right) \\ &+ E\left[\left(\frac{c_{t+1}^o}{c_{t+1}}\right)^\theta\right] E[(y_{t+1})^\theta] \text{Cov}\left(R_{t+1}, \left(\frac{c_{t+1}}{y_{t+1}}\right)^\theta\right) \end{aligned}$$

Dividing that equation by $E\left[\left(\frac{c_{t+1}^o}{c_{t+1}} \frac{c_{t+1}}{y_{t+1}} y_{t+1}\right)^\theta\right] \approx E\left[\left(\frac{c_{t+1}^o}{c_{t+1}}\right)^\theta\right] E\left[\left(\frac{c_{t+1}}{y_{t+1}}\right)^\theta\right] E[(y_{t+1})^\theta]$ results in formula (27).

Reference List

- Alvarez, Fernando, and Urban Jermann (2004). "Using Asset Prices to Measure the Cost of Business Cycles," *Journal of Political Economy*, vol. 122, no. 6, pp. 1223–1256.
- Alvarez, Fernando, and Urban Jermann (2005). "Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth," *Econometrica*, vol. 73, no. 6, pp. 1977–2016.
- Arrow, Kenneth J., and Gerald Debreu (1954). "Existence of an Equilibrium for a Competitive Economy," *Econometrica*, vol. 22, no. 3, pp. 265–290.
- Arrow, Kenneth J., and Robert Lind (1970). "Uncertainty and the Evaluation of Public Investment Decisions," *American Economic Review*, vol. 60, pp. 364–378.
- Attanasio, Orazio and Guglielmo Weber (1993). "Consumption growth, the Interest Rate and Aggregation," *Review of Economic Studies*, vol. 60, no. 3, pp. 631–649.
- Ball, Lawrence, and N. Gregory Mankiw (2007). "Intergenerational Risk Sharing in the Spirit of Arrow, Debreu, and Rawls, with Applications to Social Security Design," *Journal of Political Economy*, vol. 115, no. 4, pp. 523–547.
- Bansal, Ravi, and Amir Yaron (2005). "Risks for the Long Run: A Potential Resolution of the Asset Pricing Puzzles," *Journal of Finance*, vol. 59, no. 4, pp. 1481–1509.
- Bazonel, Coleman, and Kent Smetters (1999). "Discounting Inside the Washington D.C. Beltway," *Journal of Economic Perspectives*, vol. 13, no. 4, pp. 213–228.
- Bohn, Henning (1999). "Should the Social Security Trust Fund Hold Equities? An Intergenerational Welfare Analysis," *Review of Economic Dynamics*, vol. 2, no. 3, pp. 666–697.
- Bohn, Henning (2009). "Intergenerational Risk Sharing and Fiscal Policy," *Journal of Monetary Economics*, vol. 56, no. 6, pp. 805–816.
- Bohn, Henning (2010). "Comment on 'The Cost of Risk to the Government and Its Implications for Federal Budgeting,'" in Deborah Lucas, *Measuring and Managing Federal Financial Risk*. Cambridge, MA: National Bureau of Economic Research.
- Bohrnstedt, George W., and Arthur S. Goldberger (1969). "On the Exact Covariance of Products of Random Variables," *Journal of the American Statistical Association*, vol. 64, no. 328, pp. 1439–1442.
- Boldrin, M., L. J. Christiano, and J. D. Fisher (2001). "Habit Persistence, Asset Returns, and the Business Cycle," *American Economic Review*, vol. 91, no. 1, pp. 149–166.
- Breeden, Douglas T., and Robert H. Litzenberger (1978). "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, vol. 51, no. 4, pp. 621–51.
- Breeden, Douglas T., Robert H. Litzenberger and Tingyan Jia (2015). "Consumption-Based Asset Pricing, Part 1: Classic Theory and Tests, Measurement Issues, and Limited Participation," *Annual Review of Financial Economics*, vol. 7, pp. 35–83.
- Campbell, John Y., and John H. Cochrane (1999). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, vol. 107, no. 2, pp. 205–241.

- Diamond, Peter (1965). "National Debt in a Neoclassical Growth Model," *American Economic Review*, vol. 55, no. 5, pp. 1126–1150.
- Diamond, Peter (1967). "The Role of the Stock Market in a General Equilibrium Model With Technological Uncertainty," *American Economic Review*, vol. 57, no. 3, pp. 759–76.
- Engelhardt, Gary V., and Anil Kumar (2009). "The Elasticity of Intertemporal Substitution: New Evidence from 401(k) Participation," *Economic Letters*, vol. 103, no. 1, pp. 15–17.
- Epstein, Larry G., and Stanley E. Zin (1989). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, vol. 57, no. 4, pp. 937–969.
- Gollier, Christian. (2008) "Intergenerational Risk-Sharing and Risk-Taking of a Pension Fund," *Journal of Public Economics*, vol. 92 no. 5–6, pp. 1463–1485.
- Gordon, Robert (2016). *The Rise and Fall of American Growth*. Princeton, NJ: Princeton University Press.
- Gordon, Roger H., and Hal R. Varian (1988). "Intergenerational risk sharing," *Journal of Public Economics*, vol. 37, no. 2, pp. 185–202.
- Hall, Robert E. (1998). "Intertemporal Substitution in Consumption," *Journal of Political Economy*, vol. 96, no. 2, pp. 339–357.
- Hansen, Lars Peter, and Thomas Sargent (2021). "Macroeconomic Uncertainty Prices When Beliefs Are Tenuous," *Journal of Econometrics*, vol. 223, no. 1, Pages 222–250.
- Hirshleifer, Jack (1964). "Efficient Allocation of Capital in an Uncertain World," *American Economic Review*, vol. 54, no. 3, pp. 72–85.
- Hirshleifer, Jack (1966). "Investment decisions Under Uncertainty: Applications of the State Preference Approach," *Quarterly Journal of Economics*, vol. 80, no. 2, pp. 252–77.
- Jorgenson, D. W. (1963). "Capital Theory and Investment Behavior," *American Economic Review*, vol. 53, no. 2, pp. 247–259.
- Jorgenson, D. W., W. Vickrey, T. C. Koopmans, and P. A. Samuelson (1964). "Discussion," *American Economic Review*, vol. 54, no. 3, pp. 93–96.
- Kaplow, L. (2007). "Discounting Dollars, Discounting Lives: Intergenerational Distributive Justice and Efficiency," *University of Chicago Law Review*, vol. 74, no. 1, pp 79–118.
- Lucas, D., and M. Phaup (2008). "Reforming credit reform," *Public Budgeting and Finance*, vol. 28, no. 4, pp. 90–110.
- Lucas, Deborah, and Martin Phaup (2010). "The Cost of Risk to the Government and Its Implications for Federal Budgeting" in Deborah Lucas, *Measuring and Managing Federal Financial Risk*. Cambridge, MA: National Bureau of Economic Research, pp. 29–54.
- Mankiw, N.G., and S. P. Zeldes (1991) "The Consumption of Stockholders and Nonstockholders," *Journal of Financial Economics*, vol. 29, no. 1, pp. 97–112.
- Mehra, R., and E. C. Prescott (1985). "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, vol. 15, no. 2, pp. 145–61.
- Orszag, Peter (2000). "Correspondence," *Journal of Economic Perspectives*, vol. 14, no. 2, pp. 233–239.

- Sandmo, A. (1972). "Discount Rates for Public Investment Under Uncertainty," *International Economic Review*, vol. 13, no. 2, pp. 287–302.
- Santos, T., and P. Veronesi (2010). Habit Formation, the Cross Section of Stock Returns and the Cash-Flow Risk Puzzle," *Journal of Financial Economics*, vol. 98, no. 2, pp. 385–413.
- Sastry, P., and L. Sheiner (2015). "Credit Scoring and Scoring of Risk," Hutchins Center for Fiscal and Monetary Policy Background Paper. Washington, DC: Brookings Institution.
- Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. (2006) "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy*, vol. 114 no. 4, pp. 607–43.
- Shell, Karl (1971). "Notes on the Economics of Infinity," *Journal of Political Economy*, vol. 79, no. 5, pp. 1002–1011.
- Shaw, Kathryn (1996). "The Empirical Analysis of Risk Aversion and Income Growth," *Journal of Labor Economics*, vol. 14, no. 4, pp. 626–53.
- Tobin, J. (1969). "A General Equilibrium Approach to Monetary Theory," *Journal of Money, Credit, and Banking*, vol. 1, no. 1, pp. 15–29.
- Weil, Phillipe (1990). "Nonexpected Utility in Macroeconomics," *Quarterly Journal of Economics*, vol. 105, no. 1, pp. 29–42.