Do You Save More or Less in Response to Bad News? A New Identification of the Elasticity of Intertemporal Substitution

Schmidt, Lawrence and Toda, Alexis Akira

University of Chicago, University of California, San Diego

26 May 2015

Do You Save More or Less in Response to Bad News? A New Identification of the Elasticity of Intertemporal Substitution*

Lawrence D. W. Schmidt† Alexis Akira Toda‡

This Version: January 17, 2017

Abstract

We define the elasticity of intertemporal substitution (EIS) for general recursive preferences and identify a sharp comparative static from a general dynamic portfolio choice problem. In the homothetic case, if the EIS is smaller (larger) than 1, an investor will increase (decrease) current consumption in response to bad news about the future. Examples of bad news include if (i) she becomes more risk averse, (ii) investment opportunities shrink, (iii) investment returns become riskier, or (iv) she becomes more uncertain about the distribution of returns. Bad news effectively raises the price of future continuation utility, which produces the same qualitative changes in savings rates as lowering the interest rate.

Keywords: elasticity of intertemporal substitution, optimal portfolio problem, recursive preference

JEL codes: D91, E21, G11.

1 Introduction

A growing body of theoretical literature assumes that investors have recursive (non-additive) preferences, particularly the constant relative risk aversion, constant elasticity of intertemporal substitution (CRRA/CEIS) specification studied in Kreps and Porteus (1978), Epstein and Zin (1989), and Weil (1989). The so called Epstein-Zin utility function allows two different parameters to separately govern an agent's attitude over risky gambles and the willingness to smooth consumption over time—namely, the relative risk aversion (RRA) and the elasticity of intertemporal substitution (EIS). These two features are mechanically linked when the agent has additively separable constant relative risk aversion (CRRA) preferences.

There remains a considerable debate with respect to "reasonable" choices for the EIS. In finance, where the use of Epstein-Zin preferences is most common,

*We thank David Cashin, Henrik Kleven, Gianluca Violante, and seminar participants at University of Chicago, Hitotsubashi University International Corporate Strategy, University of Virginia Darden School of Business, and 2015 Econometric Society World Congress for comments and feedback.

†Department of Economics, University of Chicago. Email: ldwschmidt@uchicago.edu

‡Department of Economics, University of California San Diego. Email: atoda@ucsd.edu

the vast majority of papers assume that the EIS is weakly larger than one. Authors' rationale for such a choice is almost entirely pragmatic. When EIS > 1, the equity premium becomes larger, the risk free rate becomes low and stable, price-dividend ratios are pro-cyclical, and variance carries a negative price of risk. These fairly uncontroversial features of the data are much more difficult to reproduce when the EIS is less than one.[1] Assuming EIS > 1 also gets a number of comparative statics about the behavior of asset prices over the business cycle "right". However, to many observers, such an assumption is extremely puzzling, because the majority of empirical estimates of the EIS are smaller, often considerably so, than one.

This paper develops a number of comparative statics for an investor's optimal consumption-savings decision in a general portfolio problem with recursive preferences. We generalize the definition of the EIS to this setting and show that the relationship between the EIS and 1 has sharp implications for the investor's optimal response to bad news, which we formally define as any exogenous change in preferences and/or the investment opportunity set which lowers the indirect utility (continuation value) of reinvested wealth. Next, we provide a number of examples of bad news. We show that, when EIS > 1, the investor will *increase* her current consumption if (i) current or future risk aversion *increases*, (ii) current or future investment opportunities *shrink*, (iii) the investment environment becomes *riskier*, or (iv) the investor becomes *more uncertain* about the distribution of returns, while the opposite is the case when EIS < 1.

Thus, our main result is that changes in *future* investment opportunities yield qualitatively similar changes in current consumption as risk-free changes in contemporaneous returns. The intuition for this result is simple. The intertemporal decision involves the choice between contemporaneous consumption and the indirect utility of future wealth. All else constant, higher expected returns increase the relative price of current consumption, making it less costly for the agent to obtain the same level of continuation utility. Better future investment opportunities, from the agent's perspective, have the same effect qualitatively. As such, one can partially identify the EIS via these comparative statics.

An improvement in investment opportunities has two opposing effects on an investor's consumption-savings decision. The first is a substitution effect. More attractive (*e.g.*, higher return or safer) investment opportunities raise the relative price of current consumption, increasing savings. The second is a wealth effect. The improvement effectively makes the investor richer, which decreases her motivation to accumulate savings. The same intuition applies for decreases in risk aversion or reductions in uncertainty about the distribution of returns. Our result says that the substitution effect dominates the wealth effect if and only if the EIS is greater than one. Thus, these comparative statics for the consumption-savings behavior provide a new method for testing this critical assumption about preferences. Throughout, we place little structure on the problem, making it clear this comparative static is a robust feature of the preferences themselves. In particular, investors have general time and state dependent recursive preferences, and we only impose a few mild regularity conditions on asset returns.

We conclude by providing practical guidance on how to leverage these com-

[1] See Gârleanu and Panageas (2015) for a model that matches various features of financial data with EIS < 1.

parative statics in order to test whether the EIS is greater than, less than, or equal to unity. While our tests only partially identify the EIS, the generality of our theoretical framework provides several practical advantages; one could point-identify the EIS via these sources of variation with more structural assumptions. First, in contrast to standard methods of estimating the EIS, one need not take a precise stand on the definition of the "return on wealth" or the "risk free rate" when implementing the tests.[2] Despite our use of recursive preferences, one need not specify all of the inputs to the portfolio choice problem in order to make inferences. Moreover, given proper instruments and panel data, these tests are valid even if consumption is measured with error, some components (*e.g.*, human capital) of wealth are unobservable, and general equilibrium effects cause risk premia to change over time. In some cases, the tests only require repeated cross-sectional data. Finally, implementation is straightforward, since these tests involve simple differences-in-differences estimators.

2 The EIS debate

A fairly large empirical literature seeks to estimate the EIS, which is often defined as the elasticity of expected consumption growth to an exogenous increase in the real risk-free interest rate. Essentially all of these papers use the intertemporal Euler equation from a consumption-based model to derive a testable expression for the response of consumption growth to changes in the risk-free rate. Despite this common starting point, there remains a substantial debate about its magnitude. For example, Havranek et al. (2015)—who collect 2,735 estimates of the EIS from 169 published studies—find that the standard deviation of published estimates of the EIS from 33 articles in the top 5 economics journals is 1.4 (even after excluding several outliers).

When sufficient conditions hold for a representative agent to exist—particularly complete markets—the Euler equation can be estimated with aggregate data. Under further assumptions, the EIS can be identified via an instrumental variables regression of aggregate consumption growth on the risk-free rate, and its reciprocal is identifiable by regressing the risk-free rate on expected aggregate consumption growth. Using this approach, Hall (1988), a particularly influential paper, concludes that the EIS is at most 0.1. Other notable examples, such as Hansen and Singleton (1983) and Campbell and Mankiw (1989), also obtain very low estimates of the EIS. Campbell (2003), who also surveys this literature, presents a number of estimates of the EIS for a panel of countries, the vast majority of which are substantially less than one.

Many authors argue that the EIS may not be consistently estimable with aggregate data due to a variety of reasons such as market incompleteness and binding borrowing constraints, and instead use disaggregated data, which usually results in higher estimates. Attanasio and Weber (1993) estimate elasticities ranging from 0.3 to 0.8, while Beaudry and van Wincoop (1996) estimate an EIS around 1. Moreover, Mankiw and Zeldes (1991) and Vissing-Jørgensen (2002) estimate an EIS which is considerably higher for those who participate in financial markets—the only group for whom the Euler equation can be expected to

[2]See Mulligan (2002) for discussion and evidence on the role of different return concepts on estimates of the EIS.

hold with equality—though most estimates are still smaller than one.[3] Gruber (2013), who identifies the EIS via cross-sectional and time-series heterogeneity in tax rates, obtains estimates of around 2. However, not all estimates of EIS using micro data are large. Cashin and Ueyama (2016), who identify the EIS from the natural experiment of the 1997 Consumption Tax increase in Japan, obtain an estimate of 0.21;[4] Best et al. (2015) identify the EIS from the discrete jumps in the U.K. mortgage interest rate schedule and estimate it to be 0.05—0.25. Finally, even if the EIS is indeed large, there still remains the possibility of publication bias: Havránek (2015) argues that estimated values of EIS in the literature are biased upwards because researchers tend to selectively report statistically significant, large positive estimates.

To highlight several mechanisms which make the identification problem so challenging, we briefly discuss the linearized Euler equation, which plays a central role in most empirical estimates of the EIS.[5] Schmidt (2014) shows that the following restriction approximately holds in a setting with Epstein-Zin preferences, arbitrary jump-diffusion dynamics, and incomplete markets:

$$\mathrm{E}_t[\Delta \log c_{t+1}] \approx \psi\, \mathrm{E}_t[r_{t+1}] + (\psi - 1)\vartheta_t + (\psi - 1)\, \mathrm{E}_t[\nu^*_{t+1}]. \tag{2.1}$$

Here ψ is the EIS, γ is the coefficient of relative risk aversion, $\Delta \log c_{t+1}$ is consumption growth, and r_{t+1} is the continuously-compounded, real return on wealth. The first term derives a tight link between expected consumption growth and the expected return on wealth, which depends on the EIS. However, there are two additional terms, ϑ_t and ν^*_{t+1} which capture investors' preferences over the higher moments of the distribution of shocks to the aggregate state vector (*e.g.*, stochastic volatility) and the distribution of uninsurable, idiosyncratic shocks to wealth, respectively.[6] When markets are complete and shocks to investment opportunities are i.i.d., the last two terms in (2.1) are zero.

Since $\mathrm{E}_t[\Delta \log c_{t+1}]$ and $\mathrm{E}_t[r_{t+1}]$ are unobservable, it is conventional to estimate the EIS $\hat{\psi}$ via the following instrumental variables regression

$$\Delta \log c_{t+1} = \hat{\psi} r_{t+1} + u_{t+1},$$

[3] Guvenen (2006) provides a simple model with heterogeneous agents and limited participation in which the EIS estimated from aggregate data is considerably lower than the value which is relevant for pricing financial assets.

[4] Cashin and Ueyama (2016) also point out the importance of nonseparability across the nonstorable nondurable goods, storable nondurable goods, and durable goods. Their EIS estimate becomes 0.91 when the *intra*temporal substitution is restricted to be equal to the EIS (the separable case).

[5] Most papers ignore approximation errors associated with the log-lineariziation. Carroll (2001) and Ludvigson and Paxson (2001) argue that these errors can have a non-trivial effect on estimates of the EIS.

[6] In the Schmidt (2014) model, each agent's return on wealth is $r_{i,t+1} = r_{t+1} + \eta^i_{t+1}$, where η^i_{t+1} is an idiosyncratic shock which is ex-ante i.i.d. across agents and satisfies $\mathrm{E}\left[\exp(\eta^i_{t+1}) \,\middle|\, \mathcal{F}_{t+1}\right] = 1$, where $\mathcal{F}_{t+1}$ is the filtration containing aggregate information. The term $\nu^*_{t+1} := \frac{1}{1-\gamma} \log \mathrm{E}\left[\exp((1-\gamma)\eta^i_{t+1}) \,\middle|\, \mathcal{F}_{t+1}\right]$ is a certainty equivalent over higher moments of η^i_{t+1}. ϑ_t is a Jensen's inequality term which is similar to a certainty equivalent, defined as

$$\vartheta_t := \frac{1}{1-\gamma} \log \mathrm{E}_t \left[\exp\left((1-\gamma)(r_{t+1} - \mathrm{E}_t[r_{t+1}] + \nu^*_{t+1} - \mathrm{E}_t[\nu^*_{t+1}]) + \tfrac{1-\gamma}{\psi-1}\rho(wc_{t+1} - \mathrm{E}_t[wc_{t+1}])\right)\right],$$

where ρ is a linearization constant and wc_t is the log wealth-consumption ratio. If the conditional volatility of returns is driven by a single variable σ^2_t, ϑ_t is proportional to σ^2_t.

using variables which are observable at time t as instruments for r_{t+1}. Instruments are necessary since the innovation in $\Delta \log c_{t+1}$ is generally correlated with the innovation in r_{t+1}.

What can go wrong when trying to estimate the EIS using this approach? Suppose that we can indeed find an instrument z_t for expected returns which is uncorrelated with unexpected consumption growth. First, if shocks to investment opportunities are not i.i.d. and/or if markets are incomplete and the distribution of idiosyncratic shocks is state dependent, $\mathrm{E}_t[\nu^*_{t+1}]$ and ϑ_t in (2.1) are omitted variables in the regression. In this case,

$$\operatorname*{plim}_{T\to\infty} \hat{\psi} = \psi + (\psi - 1)\frac{\mathrm{Cov}[z_t, \vartheta_t + \mathrm{E}_t[\nu^*_{t+1}]]}{\mathrm{Cov}[z_t, r_{t+1}]}.$$

The weight of evidence in the empirical literature suggests that expected returns, uncertainty about investment opportunities, and idiosyncratic risk are all countercyclical—rising in recessions. Then, if $\gamma > 1$, we would expect ϑ_t and $\mathrm{E}_t[\nu^*_{t+1}]$ to be negatively correlated with $\mathrm{E}_t[r_{t+1}]$. If, for example, z_t was perfectly correlated with $\mathrm{E}_t[r_{t+1}]$, we would then expect $\hat{\psi}$ to be biased downwards if the true EIS is greater than 1 and vice versa for an EIS less than 1.

General equilibrium considerations suggest that, in practice, it is likely to be difficult to find instruments for $\mathrm{E}_t[r_{t+1}]$ which are uncorrelated with $\mathrm{E}_t[\nu^*_{t+1}]$ or ϑ_t in the time series. For example, Bansal and Yaron (2004) provide evidence that the presence of stochastic volatility, which is proportional to ϑ_t, can impart a substantial downward bias on estimates of the EIS obtained using the method in Hall (1988).[7] Schmidt (2014) argues that predictability in ν^*_{t+1} is quite substantial, suggesting that the third term can also affect the ability to estimate the EIS consistently. Slow-moving shocks to time-preferences, as studied in Albuquerque et al. (2012), are likely to cause similar challenges.

Second, we need to have the correct measures of consumption growth Δc_{t+1} and the return on wealth r_{t+1}, both of which are potentially controversial. For example, if our measure of Δc_{t+1} is too smooth or our measure of the expected return on wealth is too volatile, our estimate of ψ will be biased towards zero.[8] In a panel setting, a number of sample selection criteria are required and measurement errors are non-trivial. This potential bias is a problem because we tend to agree that the EIS is positive but disagree about its magnitude.

Finally, even if we have an instrument and the correct measures of consumption and wealth, the estimation is further complicated by the weak instrument problem (Yogo, 2004).

The lack of a consensus as to the magnitude of the EIS is troubling, since the relationship between the EIS and unity plays a central role in affecting dynamics of most theoretical models. With recursive preferences (which nest additive CRRA), this relationship affects not only the quantitative, but also the *qualitative* predictions of a model. For example, in the Bansal and Yaron (2004) long-run risk model, when EIS > 1, investors are willing to pay a premium to hedge against bad news about future economic growth rates. Valuation ratios

[7] However, Beeler and Campbell (2012) argue that this bias is relatively small using a different instrumental variables approach.

[8] The former could plausibly occur if our measure includes the consumption of households for whom the Euler equation does not hold, while the latter could occur if our measure of the return on wealth does not correctly incorporate the (unobservable) return on human capital.

are pro-cyclical, the equity premium is high, and the risk-free rate is low and stable. Setting EIS < 1 changes the basic intuition for the model and reverses each of these predictions. A hedging premium becomes a discount, valuation ratios become countercyclical, the equity premium becomes smaller, and the risk-free rate is high and volatile. Similar changes occur to asset prices and quantity dynamics in production-based models—see, for example, Kaltenbrunner and Lochstoer (2010) and Croce (2014). Further, Drechsler and Yaron (2011) also argue that, with Epstein-Zin preferences, an EIS > 1 is critical to match the level and cyclicality of the variance risk premium, a well-known feature of equity index option prices.

The remainder of this paper demonstrates that it is feasible to test the crucial distinction about the relationship between the EIS and unity using an indirect approach. We derive a number of comparative statics for how an investor's current consumption responds to news about future conditions. Consistent with the discussion above, many qualitative predictions depend crucially whether the EIS is greater than or less than one. In particular, when the EIS < 1, an investor decreases consumption in response to bad news. The opposite is the case when the EIS > 1.

3 EIS, continuation value, and consumption

3.1 Two period model without uncertainty

To build intuition, in this section we present a simple two period model with no uncertainty. Consider an economy with two goods and let $U(x_1, x_2)$ be a utility function. The elasticity of substitution is commonly defined by

$$\psi = -\frac{\mathrm{d}\log(x_1/x_2)}{\mathrm{d}\log(U_1/U_2)}, \tag{3.1}$$

where $U_l = \partial U/\partial x_l$ is the marginal utility of good l. The idea behind this definition is as follows. Suppose the price of good l is $p_l > 0$ and the agent faces a budget constraint. By the first-order condition, we have $U_l = \lambda p_l$, where $\lambda > 0$ is the Lagrange multiplier. Since $U_1/U_2 = p_1/p_2$, (3.1) implies

$$\psi = -\frac{\mathrm{d}\log(x_1/x_2)}{\mathrm{d}\log(p_1/p_2)}, \tag{3.2}$$

which is exactly the elasticity of the consumption ratio with respect to the relative price—how the agent is willing to substitute between the two goods in response to a price change. Note that while (3.1) is defined using only the preference, (3.2) requires a model, that is, how goods are traded.

Now call good 1 "consumption at time 1" and good 2 "consumption at time 2". In a two period model, consumption at time 2 is the same as the continuation utility (in consumption equivalent); hence instead of writing $U(x_1, x_2)$, let us write $f(c, v)$, and call c current consumption, v continuation value, and f aggregator. Then (3.1) becomes

$$\psi = -\frac{\mathrm{d}\log(c/v)}{\mathrm{d}\log(f_c/f_v)}, \tag{3.3}$$

and we call ψ the *elasticity of intertemporal substitution* (EIS) of the aggregator f.[9] Clearly the numerical value of EIS is invariant to a monotonic transformation of the utility function (aggregator). To see this, let $g(c, v) = F(f(c, v))$, where F is strictly increasing and differentiable. Then by the chain rule we have $g_c = F' f_c$ and $g_v = F' f_v$, so $g_c/g_v = f_c/f_v$.

The definition (3.3) is compatible with the usual definition of EIS, which is the elasticity of consumption growth with respect to the risk-free rate. To see this, let us write $c_1 = c$ and $c_2 = v$. Then the consumption growth is c_2/c_1. If p_t denotes the price of consumption at time $t = 1, 2$, then by the first-order condition and the definition of the risk-free rate R_f, we obtain $f_c/f_v = p_1/p_2 = R_f$. Therefore (3.3) becomes

$$\psi = -\frac{\mathrm{d}\log(c_1/c_2)}{\mathrm{d}\log R_f} = \frac{\mathrm{d}\log(c_2/c_1)}{\mathrm{d}\log R_f},$$

which is precisely the elasticity of intertemporal substitution as commonly defined (elasticity of consumption growth with respect to the risk-free rate).

Our main theoretical result is that if an agent faces a bad investment opportunity (which we give a precise definition in the next section), she will decrease current consumption if and only if EIS < 1. In the two period model with no uncertainty, the statement is as follows. Throughout the rest of the paper, in order to avoid corner or multiple solutions, we assume that the aggregator $f(c, v)$ is strictly increasing in each argument, strictly quasi-concave, and satisfies the Inada condition.

Proposition 3.1. *Consider an agent with endowment* (e_1, e_2)*. Suppose that the agent can save or borrow at gross risk-free rate* $R_f > 0$*. Let* $c = c_1$ *be the current consumption. Then*

$$\left(\frac{1}{c} + \frac{R_f}{R_f(e_1 - c) + e_2}\right)\frac{\partial c}{\partial R_f} = \frac{1}{R_f}\left(\frac{R_f(e_1 - c)}{R_f(e_1 - c) + e_2} - \psi\right). \tag{3.4}$$

In particular, if there is no future endowment ($e_2 = 0$)*, then*

$$\left(\frac{1}{c} + \frac{1}{e_1 - c}\right)\frac{\partial c}{\partial R_f} = \frac{1}{R_f}(1 - \psi),$$

so $\partial c/\partial R_f \gtreqless 0$ *according as* $\psi \lesseqgtr 1$*.*

As a concrete example, consider the Epstein-Zin constant elasticity of intertemporal substitution (CEIS) aggregator

$$f(c, v) = \left((1 - \beta)c^{1-1/\psi} + \beta v^{1-1/\psi}\right)^{\frac{1}{1-1/\psi}}, \tag{3.5}$$

where $0 < \beta < 1$ is the discount factor and $\psi > 0$ is the EIS.[10] The budget set is

$$v \le R_f(e_1 - c) + e_2 \iff R_f c + v \le R_f e_1 + e_2.$$

[9] Since $\log(f_c/f_v)$ is a function, not a variable, the notation (3.3) is not rigorous. We can make sense of (3.3) by writing as $\frac{1}{\psi} = -\frac{\mathrm{d}\log(f_c/f_v)}{\mathrm{d}x}$, where x is a variable, c is a constant, and v is determined such that $x = \log(c/v) \iff v = c\mathrm{e}^{-x}$. (It does not matter whether we treat c or v as a constant.)

[10] When $\psi = 1$, we can take the limit of (3.5) as $\psi \to 1$ to obtain

$$f(c, v) = \exp((1 - \beta)\log c + \beta \log v) = c^{1-\beta}v^{\beta}.$$

Using calculus, we can easily solve for the optimal consumption bundle, which is

$$(c, v) = \left(\frac{(1-\beta)^{\psi} R_f^{-\psi} (R_f e_1 + e_2)}{(1-\beta)^{\psi} R_f^{1-\psi} + \beta^{\psi}}, \frac{\beta^{\psi} (R_f e_1 + e_2)}{(1-\beta)^{\psi} R_f^{1-\psi} + \beta^{\psi}} \right). \tag{3.6}$$

Figure 1 plots the budget sets, indifference curves, and optimal consumption bundles for different values of the interest rate when there is no future endowment ($e_2 = 0$), $\beta = 1/2$, and the interest rates are 0%, 25%, and 50%. Different columns correspond with different choices of the EIS ψ—1/2, 1, and 2, respectively—which is constant for every interest rate.

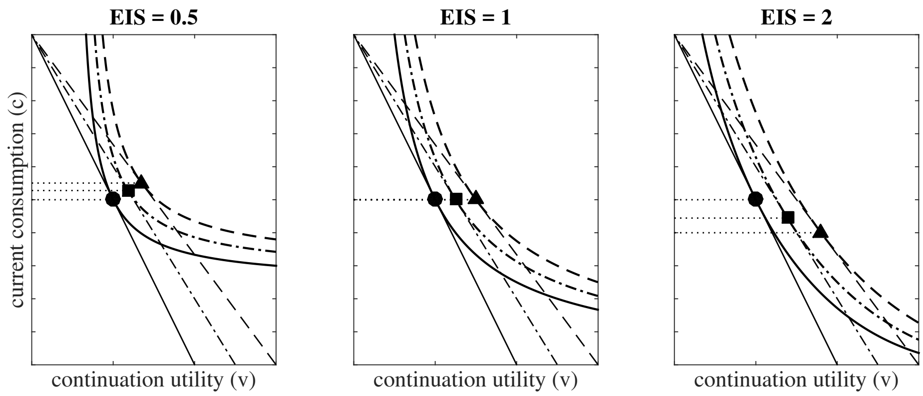

Figure 1: Intuition for comparative statics result–CEIS aggregator example

Note: This figure characterizes the optimal intertemporal consumption plan of an investor with Epstein-Zin preferences, where the aggregator $f(c, v)$ is defined by (3.5) with $\beta = 1/2$, for different values of the risk free rate R_f. Columns correspond with $\psi = 1/2, 1, 2$, respectively.

In the middle panel, the EIS exactly equals one, and the CEIS aggregator is Cobb-Douglas. An agent with Cobb-Douglas preferences spends a constant fraction of wealth on each good, and, as such, current consumption stays the same regardless of the value of the interest rate. When the EIS is less than unity (left panel), the current consumption and continuation utility are relative complements, so the agent consumes more of both goods. The opposite is the case in the right panel, in which the two goods are relative substitutes. Whether the two goods are complements or substitutes depends entirely on the EIS.

Note that after substituting the budget constraint into the aggregator, the optimal consumption problem reduces to

$$\max_c f(c, v(e_1 - c)),$$

where $v(w) = R_f w + e_2$ is the continuation value of the reinvested wealth w. When the risk-free rate increases, so does the continuation utility v. Our key observation is that the continuation utility may change from many other reasons than the change in the interest rate, and those changes affect current consumption in a similar way to (3.4), where the sign of the change in consumption depends crucially on whether EIS $\gtrless 1$. This observation enables us to identify whether EIS $\gtrless 1$ through the direction of the change in current consumption.

What can go wrong with using the usual definition of the EIS,

$$\hat{\psi} = \frac{\mathrm{d}\log(c_2/c_1)}{\mathrm{d}\log R_f}, \tag{3.7}$$

which is the elasticity of consumption growth with respect to the risk-free rate? Consider an agent with Epstein-Zin preference as in (3.5). Suppose that there is a borrowing constraint, and the agent can only pledge fraction $0 \le \alpha < 1$ of future endowment in order to borrow. Assuming that the borrowing constraint is binding, consumption is then $c_2 = (1-\alpha)e_2$ and $c_1 = e_1 + \frac{\alpha e_2}{R_f}$. Hence by (3.7), it follows that

$$\hat{\psi} = R_f \frac{\mathrm{d}}{\mathrm{d}R_f} \log\left(\frac{(1-\alpha)e_2}{e_1 + \frac{\alpha e_2}{R_f}}\right) = \frac{\alpha e_2}{R_f e_1 + \alpha e_2} < 1. \tag{3.8}$$

In order for the borrowing constraint to bind, the agent must be willing to consume less at $t = 2$ without the borrowing constraint. Hence by (3.6), the borrowing constraint binds if and only if

$$\frac{\beta^\psi (R_f e_1 + e_2)}{(1-\beta)^\psi R_f^{1-\psi} + \beta^\psi} \le (1-\alpha)e_2 \iff \frac{R_f e_1}{\alpha e_2} \le \frac{1-\alpha}{\alpha}\left(\frac{1-\beta}{\beta}\right)^\psi R_f^{1-\psi} - 1. \tag{3.9}$$

Hence by (3.8) and (3.9), when the borrowing constraint binds, we obtain

$$\frac{\alpha}{1-\alpha}\left(\frac{\beta}{1-\beta}\right)^\psi R_f^{\psi-1} \le \hat{\psi} < 1. \tag{3.10}$$

This argument shows that it is problematic to use the usual definition of EIS (3.7) to estimate it when borrowing constraints might be binding. According to (3.10), when the borrowing constraint is binding, $\hat{\psi}$ is always less than 1, independent of the true value ψ. Furthermore, if the borrowing constraint is tight ($\alpha \to 0$) or the agent is impatient ($\beta \to 0$), the lower bound for $\hat{\psi}$ approaches 0. Given this result, it is not surprising that the empirical estimates of EIS based on aggregate consumption are close to zero.

3.2 Dynamic model with uncertainty

Next we consider the general case. Time is finite and is denoted by $t = 0, 1, \dots, T$. All random variables are defined on a probability space $(\Omega, \mathcal{F}, P)$.

Consider a single agent (investor) that has recursive preferences defined over finite consumption plans from time t onwards $\{c_{t+s}\}_{s=0}^{T-t}$ (where $t = 0, 1, \dots, T$), constructed as follows. The terminal utility is $U_T = u_T(c_T)$, where $u_T : \mathbb{R}_+ \to \mathbb{R}_+$ is increasing (typically $u_T(c) = c$). Given the recursive utility at time $t+1$, denoted by U_{t+1}, the time t recursive utility is defined by

$$U_t = f_t(c_t, \mathcal{M}_t(U_{t+1})), \tag{3.11}$$

where $f_t : \mathbb{R}_+^2 \to \mathbb{R}_+$ is the aggregator, c_t is consumption, and $\mathcal{M}_t(U_{t+1})$ is the certainty equivalent of the distribution of time $t+1$ utility conditional on time t information (typically $\mathcal{M}_t(X) = \phi_t^{-1}(\mathrm{E}_t[\phi_t(X)])$ for some strictly increasing, concave function $\phi_t : \mathbb{R}_+ \to \mathbb{R}$). The terminal utility function u_T, the aggregator f_t, and the certainty equivalent $\mathcal{M}_t$ may be time dependent as well as state dependent.

Consider a general optimal consumption-portfolio problem. Given the initial wealth at time t, denoted by $w_t > 0$, the agent decides how much to consume (c_t), and how to allocate the remaining wealth $w_t - c_t > 0$ across assets. For now the details of the problem does not matter—for example, there may or may not be labor income, transaction costs, portfolio constraints, etc.—we only assume that a sequence of optimal consumption rules and value functions exist.

Let $V_t(w)$ be the value function at time t, given wealth w. Since the problem can be quite general, $V_t(w)$ may depend not only on time and current wealth but on other state variables (*e.g.*, expected stock market return and volatility in U.S., inflation in Brazil, current and past labor income, unemployment rate, etc.). By the principle of optimality and the definition of the recursive utility (3.11), we obtain the Bellman equation

$$V_t(w) = \max_c f_t(c, \mathcal{M}_t(V_{t+1}(w'))), \tag{3.12}$$

where w' is the initial wealth at time $t+1$ determined by the budget constraint. Given all state variables except for wealth, the value of $\mathcal{M}_t(V_{t+1}(w'))$ depends only on the reinvested wealth $w - c$. Let $v_t(w - c)$ be this value and call it the *continuation value function* for short, where again v_t may be time and state dependent. Then the Bellman equation (3.12) becomes

$$V_t(w) = \max_c f_t(c, v_t(w - c)). \tag{3.13}$$

In this setting, we can show a similar result to Proposition 3.1. Since the problem is now essentially static, let us suppress the time subscript and let

$$\psi = -\frac{\mathrm{d}\log(c/v)}{\mathrm{d}\log(f_c/f_v)}$$

be the EIS of aggregator f, as in (3.3). In the simple two period model, we perturbed the risk-free rate. In the general case, since we do not have any structure on the problem, we need to perturb the continuation value function v directly. Let h be any function and

$$c(\alpha; h) = \arg\max_c f(c, (v + \alpha h)(w - c))$$

be the optimal consumption rule when v is perturbed in the direction of h by a small amount $\alpha \in \mathbb{R}$.[11] Then we obtain the following result.

Theorem 3.2. *Let everything be as above. Then*

$$\left(\frac{1}{c} + \frac{v'}{v} - \psi\frac{v''}{v'}\right)\frac{\partial}{\partial\alpha}c(0; h) = \frac{h}{v} - \psi\frac{h'}{v'}, \tag{3.14}$$

where all functions are evaluated at $w - c$. In particular, if $h(w) = w$ (so we perturb v by adding a linear function), then

$$\left(\frac{1}{c} + \frac{v'}{v} - \psi\frac{v''}{v'}\right)\frac{\partial}{\partial\alpha}c(0; h) = \frac{1}{v'}(\varepsilon(w - c) - \psi), \tag{3.15}$$

[11] Mathematically, we are considering Gâteaux derivatives.

where $\varepsilon(w) = \mathrm{d}\log v(w)/\mathrm{d}\log w = wv'(w)/v(w)$ is the elasticity of continuation value with respect to reinvested wealth. If in addition v is concave ($v'' \le 0$), then $\partial c/\partial\alpha \gtrless 0$ according as $\psi \lesseqgtr \varepsilon$.

Proposition 3.1 is clearly a special case of Theorem 3.2. In the setting of Proposition 3.1, the continuation value function is $v(w) = R_f w + e_2$, which is affine in w. Set $h(w) = w$. Then $(v + \alpha h)(w) = (R_f + \alpha)w + e_2$, so changing α is the same as changing R_f. Since $v' = R_f > 0$, $v'' = 0$, and $\varepsilon(w) = wv'(w)/v(w) = \frac{R_f w}{R_f w + e_2}$, it follows from (3.15) that

$$\left(\frac{1}{c} + \frac{R_f}{R_f(w-c)+e_2}\right)\frac{\partial c}{\partial R_f} = \frac{1}{R_f}\left(\frac{R_f(w-c)}{R_f(w-c)+e_2} - \psi\right),$$

which is precisely (3.4) with $w = e_1$.

3.3 Characterization of "bad news"

By Theorem 3.2, assuming that the continuation value function v is concave, we can identify whether $\psi \lesseqgtr \varepsilon$ according as $\mathrm{d}c/\mathrm{d}\alpha \gtrless 0$ at $\alpha = 0$ when $v(w)$ is perturbed to $v(w) + \alpha w$. More generally, it would be convenient if we can determine under what conditions the continuation value $v(w)$ increases or decreases. To this end we introduce the following definition.

Definition 3.3. *Bad news* at time t is any exogenous change that makes v_t smaller.

The following proposition shows that increases in risk aversion, decreases in investment opportunities, and riskier or lower expected investment returns all constitute bad news. For concreteness let $R_{t+1}(\theta)$ be the gross return on a portfolio $\theta \in \Theta_t$, where Θ_t is the set of admissible portfolios at time t, and y_{t+1} be the labor income at time $t+1$. Suppose that the certainty equivalent $\mathcal{M}_t$ takes the form

$$\mathcal{M}_t(U) = \phi_t^{-1}(\mathrm{E}_t[\phi_t(U)]), \tag{3.16}$$

where $\phi_t : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing and concave.

Proposition 3.4. *The following events constitute bad news at time t:*

1. *the agent becomes more risk averse,* i.e., *ϕ_t changes to $\tilde{\phi}_t$, where $g = \tilde{\phi}_t \circ \phi_t^{-1}$ is increasing and concave (Pratt, 1964).*
2. *the investment opportunity shrinks,* i.e., *Θ_t changes to $\tilde{\Theta}_t \subset \Theta_t$.*
3. *the portfolio return and/or labor income become riskier: $\mathrm{E}_{t+1}[\tilde{R}_{t+1}(\theta)] \le R_{t+1}(\theta)$ and $\mathrm{E}_{t+1}[\tilde{y}_{t+1}] \le y_{t+1}$. (For this case we need to assume that $\phi_t \circ V_{t+1}$ is concave.)*

Proof.

Case 1: Agent becomes more risk averse. Let w be the reinvested wealth at time t and $w' = R_{t+1}(\theta)w + y_{t+1}$ be the initial wealth at time $t+1$ determined by the budget constraint. By the definition of the continuation value function, we have $v_t(w) = \max_\theta \mathcal{M}_t(V_{t+1}(w'))$. Therefore

$$\begin{aligned}
\tilde{\phi}_t(\tilde{v}_t(w)) &= \max_\theta \mathrm{E}_t[\tilde{\phi}_t(V_{t+1}(w'))] && (\because \text{Definition of } \tilde{v}_t) \\
&= \max_\theta \mathrm{E}_t[g(\phi_t(V_{t+1}(w')))] && (\because g = \tilde{\phi}_t \circ \phi_t^{-1}) \\
&\le \max_\theta g(\mathrm{E}_t(\phi_t(V_{t+1}(w')))) && (\because \text{Jensen's inequality}) \\
&= g\left(\max_\theta \mathrm{E}_t(\phi_t(V_{t+1}(w')))\right) && (\because g \text{ monotone}) \\
&= g(\phi_t(v_t(w))) && (\because \text{Definition of } v_t) \\
&= \tilde{\phi}_t(v_t(w)). && (\because g = \tilde{\phi}_t \circ \phi_t^{-1})
\end{aligned}$$

Applying $\tilde{\phi}_t^{-1}$ to both sides, we obtain $\tilde{v}_t \le v_t$.

Case 2: Investment opportunity shrinks. Suppose the portfolio constraint shrinks to $\tilde{\Theta}_t \subset \Theta_t$. Then by the definition of v_t, we have

$$\begin{aligned}
\tilde{v}_t(w) &= \max_{\theta \in \tilde{\Theta}_t} \phi_t^{-1}(\mathrm{E}_t[\phi_t(V_{t+1}(w'))]) \\
&\le \max_{\theta \in \Theta_t} \phi_t^{-1}(\mathrm{E}_t[\phi_t(V_{t+1}(w'))]) = v_t(w).
\end{aligned}$$

Case 3: Environment becomes riskier. By the definition of v_t, law of iterated expectations, and Jensen's inequality, we have

$$\begin{aligned}
\phi_t(\tilde{v}_t(w)) &= \max_\theta \mathrm{E}_t[\phi_t(V_{t+1}(\tilde{R}_{t+1}(\theta)w + \tilde{y}_{t+1}))] \\
&= \max_\theta \mathrm{E}_t[\mathrm{E}_{t+1}[\phi_t(V_{t+1}(\tilde{R}_{t+1}(\theta)w + \tilde{y}_{t+1}))]] \\
&\le \max_\theta \mathrm{E}_t[\phi_t(V_{t+1}(\mathrm{E}_{t+1}[\tilde{R}_{t+1}(\theta)w + \tilde{y}_{t+1}]))] \\
&\le \max_\theta \mathrm{E}_t[\phi_t(V_{t+1}(R_{t+1}(\theta)w + y_{t+1}))] = \phi_t(v_t(w)).
\end{aligned}$$

Applying ϕ_t^{-1} to both sides, we obtain $\tilde{v}_t \le v_t$. □

We can provide even more examples of "bad news". In Proposition 3.4, the agent uses a single objective or subjective probability measure to calculate the certainty equivalent $\mathcal{M}_t$. However, it is possible that she is uncertain about the probability measure itself. To capture such ambiguity, following Hayashi and Miao (2011) and Ju and Miao (2012), we assume that the certainty equivalent is of the form

$$\mathcal{M}_t(U_{t+1}) = \varphi_t^{-1}(\mathrm{E}_{\mu_t}[\varphi_t(\phi_t^{-1}(\mathrm{E}_{\pi_t}[\phi_t(U_{t+1})]))]), \tag{3.17}$$

where ϕ_t and φ_t capture risk aversion and ambiguity aversion, respectively. $\pi_t \in \mathcal{P}_t$ is the subjective probability measure over the state space, and μ_t is the subjective probability measure over the set of the underlying stochastic process $\mathcal{P}_t$. When $\varphi_t = \phi_t$, (3.17) reduces to (3.16), where the expectation is taken over $\mu_t \circ \pi_t$. If the agent is infinitely ambiguity averse, then (3.17) reduces to

$$\mathcal{M}_t(U_{t+1}) = \phi_t^{-1}\left(\min_{\pi_t \in \mathcal{P}_t} \mathrm{E}_{\pi_t}[\phi_t(U_{t+1})]\right), \tag{3.18}$$

the multi-priors model introduced by Gilboa and Schmeidler (1989) and generalized to the intertemporal setting (without the separation of EIS from risk aversion) by Epstein and Schneider (2003) and (with the three-way separation between EIS, risk aversion, and ambiguity aversion) by Hayashi (2005).

In this setting, the following proposition shows that increases in ambiguity aversion or model uncertainty are treated as bad news.

Proposition 3.5. *The following events constitute bad news at time t:*

4. *the agent becomes more ambiguity averse,* i.e., *φ_t changes to $\tilde{\varphi}_t$, where $g = \tilde{\varphi}_t \circ \varphi_t^{-1}$ is increasing and concave.*
5. *the agent is infinitely ambiguity averse ($\mathcal{M}_t$ takes the form* (3.18)*), and the set of subjective probability measures expands to $\tilde{\mathcal{P}}_t \supset \mathcal{P}_t$.*

Proof. Analogous to the proof of Proposition 3.4. □

4 Homothetic case

We have seen so far that whether consumption increases or decreases after learning a bad news crucially depends on whether the EIS is larger than or less than the elasticity of the continuation value with respect to wealth. However, this result cannot be directly used to determine the magnitude of the EIS ψ, because (i) since the functional form of v is unknown without imposing additional structure, so is its elasticity ε, and (ii) exogenous changes in the agent's environment do change v, but there is no reason why v should change by a linear function αw. By imposing additional structure—homogeneity in wealth—we can make Theorem 3.2 applicable.

4.1 Model

Preferences Let f_t be the aggregator and $\mathcal{M}_t$ be the certainty equivalent that defines the recursive utility. In this section we maintain the following assumptions.

Assumption 1. *Terminal utility is proportional to consumption: $u_T(c) = b_T c$ for some random variable $b_T > 0$. The aggregator $f_t : \mathbb{R}^2_+ \to \mathbb{R}_+$ is upper semi-continuous, weakly increasing in both arguments, strictly quasi-concave, and homogeneous of degree 1,* i.e., *$f_t(\lambda c, \lambda v) = \lambda f_t(c, v)$ for all $\lambda > 0$.*

As before, the aggregator f_t can be time and state dependent. Assumption 1 essentially says that preferences are homothetic. The reason we assume $u_T(c) = b_T c$ is for generality. If the agent cares only about consumption, then $u_T(c) = c$ (hence $b_T = 1$) would be a natural choice. Alternatively, the agent may have a bequest motive. If preferences over the last period's consumption and bequest are homothetic, then by maximizing over the consumption decision, the last period's value function will be linear in wealth. $u_T(c) = b_T c$ can be interpreted as such a terminal value function.

Assumption 2 (CRRA certainty equivalent)**.** *The certainty equivalent $\mathcal{M}_t$ exhibits constant relative risk aversion (CRRA),* i.e.,

$$\mathcal{M}_t(U) = \begin{cases} \mathrm{E}_t[U^{1-\gamma_t}]^{\frac{1}{1-\gamma_t}}, & (\gamma_t \neq 1) \\ \exp\left(\mathrm{E}_t[\log U]\right), & (\gamma_t = 1) \end{cases} \tag{4.1}$$

where $\gamma_t > 0$ is the coefficient of relative risk aversion.

Although it is not trivial that we should define the CRRA certainty equivalent for the case $\gamma_t = 1$ by using the exponential and logarithmic functions, Lemma A.1 in the Appendix shows that it is indeed a natural definition. Again the relative risk aversion (RRA) coefficient $\gamma_t > 0$ may be time and state dependent.

The recursive preferences satisfying Assumptions 1 and 2 nest Epstein-Zin CRRA/CEIS (constant relative risk aversion/constant elasticity of intertemporal substitution) preferences, which obtain by setting $u_T(c) = c$, $f_t(c, v) = ((1-\beta)c^{1-1/\psi} + \beta v^{1-1/\psi})^{\frac{1}{1-1/\psi}}$, and $\gamma_t = \gamma$, where $\psi > 0$ is the elasticity of intertemporal substitution and $0 < \beta < 1$ is the discount factor. Of course, we get the standard additive CRRA preference when $\psi = 1/\gamma$.

Clearly we can allow for ambiguity aversion by using the form (3.17) with $\phi_t(x) = \frac{x^{1-\gamma_t}}{1-\gamma_t}$ and $\varphi_t(x) = \frac{x^{1-\eta_t}}{1-\eta_t}$, where $\gamma_t, \eta_t > 0$ are the relative risk aversion and ambiguity aversion coefficients.

Investment opportunity There are finitely many assets indexed by $j \in J = \{1, \dots, J\}$. Asset returns are exogenous from the perspective of the investor, *i.e.*, the investor is a price taker. Let $R^j_{t+1} \geq 0$ be the *gross* return on asset j between time t and $t+1$, which is a random variable. Let θ^j be the fraction of wealth invested in asset j. $\theta^j > 0$ (< 0) means a long (short) position in asset j. Let $\theta = (\theta^1, \dots, \theta^J)$ be the portfolio, where $\sum_{j=1}^J \theta^j = 1$. At each point in time, the investor may be constrained in the asset position she can take. For example, she may face margin or short sales constraints. Let $\Theta_t \subset \mathbb{R}^J$ be the set of feasible portfolios at time t, which again can be time and state dependent.

Let $\theta \in \Theta_t$ be a portfolio. The gross return on this portfolio between time t and $t+1$ is denoted by

$$R_{t+1}(\theta) = \sum_{j=1}^J R^j_{t+1}\theta^j. \tag{4.2}$$

If short sales are allowed, it may be the case that $R_{t+1}(\theta) \leq 0$ in some states, leaving the agent with negative wealth. We rule out this possibility by letting the agent with negative wealth bankrupt and get utility $-\infty$, so she chooses only portfolios that satisfy $R_{t+1}(\theta) > 0$ almost surely. By redefining the portfolio constraint if necessary, we assume that $R_{t+1}(\theta) > 0$ almost surely for all $\theta \in \Theta_t$.

Assumption 3. *The portfolio constraint $\Theta_t \subset \mathbb{R}^J$ is nonempty, compact, and convex, and $R_{t+1}(\theta) > 0$ almost surely for all $\theta \in \Theta_t$.*

Assumption 3 implies that there is some limit to short sales (compactness) and that portfolios are infinitely divisible (convexity). These assumptions are quite natural in a developed capital market.

Optimal consumption-portfolio decision The investor is endowed with initial wealth (capital) $w_0 > 0$ in period 0 but nothing thereafter. Her budget constraint is therefore

$$w_{t+1} = R_{t+1}(\theta_t)(w_t - c_t), \tag{4.3}$$

where w_t is wealth at the beginning of time t, c_t is consumption, and $\theta_t \in \Theta_t$ is the portfolio. The objective of the agent is to maximize the recursive utility defined by (3.11) subject to the budget constraint (4.3) and the portfolio constraint $\theta_t \in \Theta_t$.

Let $V_t(w)$ be the value function of the agent with wealth w at time t. Substituting the budget constraint into the definition of recursive utility, we obtain the Bellman equation

$$V_t(w) = \max_{\substack{0 \le c \le w \\ \theta \in \Theta_t}} f_t\left(c, \mathrm{E}_t[V_{t+1}(R_{t+1}(\theta)(w-c))^{1-\gamma_t}]^{\frac{1}{1-\gamma_t}}\right). \tag{4.4}$$

The following lemma, which generalizes the classic portfolio problems of Samuelson (1969) and Hakansson (1970, 1971), shows that the optimal consumption-portfolio rule and the value function are linear.

Lemma 4.1. *Suppose Assumptions 1–3 hold. Define* $\{b_t\}_{t=0}^{T-1}$ *by*

$$b_t = \max_{0 \le \tilde{c} \le 1} f_t\left(\tilde{c}, (1-\tilde{c}) \max_{\theta \in \Theta_t} \mathrm{E}_t[(b_{t+1}R_{t+1}(\theta))^{1-\gamma_t}]^{\frac{1}{1-\gamma_t}}\right) \tag{4.5}$$

and let $\tilde{c}_t, \theta_t$ *be maximizers of* (4.5). *Then the value function is* $V_t(w) = b_t w$, *the optimal consumption rule is* $c = \tilde{c}_t w$, *and the optimal portfolio is* θ_t.

Proof. If $t = T$, there is no portfolio decision to make and the value function is $V_T(w) = \max_{0 \le c \le w} b_T c = b_T w$. Assume that the claim holds for time $t+1$ onwards. Then

$$\begin{aligned}
V_t(w) &= \max_{\substack{0 \le c \le w \\ \theta \in \Theta_t}} f_t\left(c, \mathrm{E}_t[V_{t+1}(R_{t+1}(\theta)(w-c))^{1-\gamma_t}]^{\frac{1}{1-\gamma_t}}\right) \\
&= \max_{\substack{0 \le c \le w \\ \theta \in \Theta_t}} f_t\left(c, \mathrm{E}_t[(b_{t+1}R_{t+1}(\theta)(w-c))^{1-\gamma_t}]^{\frac{1}{1-\gamma_t}}\right) \\
&= \max_{0 \le c \le w} f_t\left(c, (w-c) \max_{\theta \in \Theta_t} \mathrm{E}_t[(b_{t+1}R_{t+1}(\theta))^{1-\gamma_t}]^{\frac{1}{1-\gamma_t}}\right) \\
&= w \max_{0 \le \tilde{c} \le 1} f_t\left(\tilde{c}, (1-\tilde{c}) \max_{\theta \in \Theta_t} \mathrm{E}_t[(b_{t+1}R_{t+1}(\theta))^{1-\gamma_t}]^{\frac{1}{1-\gamma_t}}\right) = b_t w,
\end{aligned}$$

where the first line is the Bellman equation (4.4), the second line is by the induction hypothesis $V_{t+1}(w) = b_{t+1}w$, the third line is by the monotonicity and the upper semi-continuity of f, and the fourth line is by the homotheticity of f and (4.5). ☐

Remark. The acute reader may have noticed that in order to justify the above proof, we need in addition the continuity of $\mathrm{E}_t[(b_{t+1}R_{t+1}(\theta))^{1-\gamma_t}]$ with respect to θ so that the maximum with respect to $\theta \in \Theta_t$ is attained. We can prove the continuity using the Dominated Convergence Theorem, provided that the integrand is bounded above by some integrable function. This is the case, for example, if all random variables take finitely many values. In Lemma 4.1 we implicitly assume that some sufficient condition for continuity is satisfied.

Remark. Since by Assumption 1 the aggregator f is strictly quasi-concave, the optimal consumption rule $\tilde{c}_t$ is unique. Since the portfolio constraint Θ_t is convex and the objective function for the optimal portfolio problem

$$\frac{1}{1-\gamma_t} \mathrm{E}_t[(b_{t+1}R_{t+1}(\theta))^{1-\gamma_t}]$$

is concave, the optimal portfolios form a convex set. If, in addition, there are no redundant assets (so asset returns $R^1_{t+1}, \dots, R^J_{t+1}$ are linearly independent), then the optimal portfolio is unique.

4.2 Response of consumption to "bad news"

Armed with Lemma 4.1, we apply Theorem 3.2 to the homothetic case. By homogeneity, the value function is linear: $V_t(w) = b_t w$. By the proof of Lemma 4.1, the continuation value function is

$$v_t(w) = w \max_{\theta \in \Theta_t} \mathrm{E}_t[(b_{t+1} R_{t+1}(\theta))^{1-\gamma_t}]^{\frac{1}{1-\gamma_t}} =: \rho_t w, \tag{4.6}$$

which is also linear. Then the elasticity of continuation value is identically equal to 1, and any exogenous change in the agent's environment affects only the coefficient of the continuation value. Since $v_t(w) + \alpha w = (\rho_t + \alpha)w$, perturbing the continuation value by a linear function is equivalent to changing ρ. Therefore by Theorem 3.2 we immediately obtain the following results.

Corollary 4.2. *Suppose Assumptions 1–3 hold. Then*

$$\left(\frac{1}{c} + \frac{1}{w-c}\right)\frac{\partial c}{\partial \rho} = \frac{1}{\rho}(1-\psi). \tag{4.7}$$

In particular, $\partial c/\partial \rho \gtreqless 0$ according as $\psi \lesseqgtr 1$.

Proof. Since $v(w) = \rho w$, we have $v' = \rho$, $v'' = 0$, and $\varepsilon = wv'(w)/v(w) = 1$. Substituting these quantities into (3.15) and evaluating at $w - c$, we obtain (4.7). □

Corollary 4.3. *Let $s \le t$. Then $\partial c_s/\partial \rho_t \gtreqless 0$ according as $\psi \lesseqgtr 1$.*

Proof. Since the aggregator is monotonic, when ρ_t goes up or down, so does ρ_s for all $s \le t$. Hence the conclusion holds by Corollary 4.2. □

Corollary 4.3 shows that when the agent learns that the coefficient of the continuation value will become smaller at some future date, the agent will decrease (increase) consumption in every prior date if $\psi < 1$ ($\psi > 1$). Furthermore, since $v_t(w) = \rho_t w$, according to Definition 3.3 a bad news is any exogenous event that decreases ρ_t. By Propositions 3.4 and 3.5, increases in risk or ambiguity aversion, reduction in investment opportunities, and riskier or more uncertain environment all constitute bad news.[12]

So far, we have derived comparative statics of the saving rate of a single investor, who takes asset returns as given. The same result holds in a general equilibrium model studied by Toda (2014). In this model, there are a continuum of ex ante identical agents indexed by $i \in I = [0,1]$. There are J constant-returns-to-scale stochastic saving technologies indexed by $j \in J = \{1, \dots, J\}$. The gross investment return of agent i on technology j between time t and

[12]Working with additive CRRA preferences, many papers have found that increased risk increases or decreases savings depending on whether relative risk aversion is greater or less than 1 (Phelps, 1962; Levhari and Srinivasan, 1969; Merton, 1969; Sandmo, 1970; Rothschild and Stiglitz, 1971). Corollary 4.2 is considerably more general, and what matters is the EIS, not risk aversion.

$t+1$ is denoted by $A^j_{i,t+1}$. Thus if agent i invests capital K at the end of period t, she will collect capital $A^j_{i,t+1}K$ at the beginning of period $t+1$. The subscript i indicates that the investor may face idiosyncratic risk. Let $\mathcal{F}_{it}$ be agent i's information set and $\mathcal{F}_t = \bigcap_i \mathcal{F}_{it}$ be the σ-algebra generated by the aggregate variables. Assume that agents are symmetric, so A^j_{it} is i.i.d. across agents conditional on $\mathcal{F}_t$, and that the idiosyncratic risk is transitory (since A's are *rates* of return, the shocks are permanent in *levels*), so the distribution of $A^j_{i,t+1}$ conditional on $\mathcal{F}_{it}$ is the same as the one conditional on $\mathcal{F}_t$. One can also allow for arbitrarily many assets in zero net supply, whose dividends are $\mathcal{F}_t$-measurable.

Under these assumptions and using a similar argument to Toda (2014), we can show that there exists a unique equilibrium, the zero net supply assets are not traded in equilibrium, and that the optimal consumption-portfolio rule is the one in Lemma 4.1 with $R_{t+1}(\theta)$ replaced by

$$R_{i,t+1}(\theta) = \sum_{j=1}^{J} A^j_{i,t+1}\theta^j.$$

Since the structure of the general equilibrium model is identical to that of a single agent problem, Proposition 3.4 continues to hold. We note this result in the following corollary.

Corollary 4.4. *The following constitutes bad news:*

3' *the vector of investment returns $\mathbf{A}_{i,t+1} = (A^1_{i,t+1}, \dots, A^J_{i,t+1})$ becomes riskier to (second-order stochastically dominates) $\tilde{\mathbf{A}}_{i,t+1}$, so*

$$\mathrm{E}\left[\tilde{\mathbf{A}}_{i,t+1} \,\middle|\, \mathbf{A}_{i,t+1}\right] \leq \mathbf{A}_{i,t+1}.$$

5 Empirical identification strategies

Here, we will discuss how to develop formal econometric tests of the comparative statics developed in the previous section. Given our above concerns about the identification issues caused by general equilibrium effects, we will emphasize tests which rely on the availability of panel data on consumption/savings.

In the framework developed thus far, the subscript on the value function may be interpreted as indexing either age or calendar time. To make a clearer distinction between the two, we will index the value function by a, setting $A := T$, and add a second subscript t for calendar time. $V^i_{a,t}(w)$ will denote the value function of agent i, who is a years old with wealth w at calendar time t. We will do the same with risk aversion coefficients and portfolio constraints, which can depend on calendar time and age. The i notation allows for cross-sectional heterogeneity in preferences, constraints, and/or investment opportunities. Next, we place structure on the state space governing the common and individual-specific variation in determinants of the portfolio problem over time.

Assumption 4. *The following statements are true:*

1. *The aggregate state of the economy at time t is denoted by s_t, where $\{s_t\}$ is an exogenous, stationary, finite-dimensional, Markov process.*

2. *Conditional on s_t, all cross-sectional variation in preferences and investment opportunities is completely characterized by the Markovian random vector $\{\alpha_{it}\}$, which satisfies $\alpha_{i,t+1} \perp \alpha_{j,t+1} | s_{t+1}, s_t, \alpha_{it}, \alpha_{jt}$ for all $i \neq j$. The distribution of $\alpha_{i,t+1}$ given individual i's information set at time t only depends on α_{it} and s_t.*

Assumption 4.1 imposes stationarity restrictions on the variation in investment opportunities over calendar time. Changes in s_t generate common time series variation in realized returns and characteristics of the future portfolio choice problem across investors. Next, Assumption 4.2 allows for heterogeneity of a fairly general form. For example, investors may differ in their degrees of risk aversion, portfolio constraints, and return distributions. While, at first glance, heterogeneity in returns might seem like an unnatural assumption, differences in tax rates, portfolio management fees and/or transaction costs are all capable of generating return differentials in the data. These cross-sectional differences may be persistent over time (via the Markov structure of α_{it}), and changes in the aggregate states can interact with changes in the cross-sectional distribution of α in general ways.

When Assumptions 1–4 hold, it immediately follows that the value function is linear and takes the form: $V^i_{a,t}(w) = b_a(\alpha_{it}, s_t)w$. Analogously, the continuation value per dollar of future wealth satisfies $\rho^i_{a,t} = \rho_a(\alpha_{it}, s_t)$. If consumption, wealth, and $\rho^i_{a,t}$ were observable, then (4.7) implies a very straightforward way to estimate the EIS, provided that it is stable over time or across individuals. To see this, rewriting (4.7), we obtain

$$\psi_t = 1 - \frac{\partial \log \frac{c_t}{w_t - c_t}}{\partial \log \rho_t},$$

so EIS is 1 minus the elasticity of consumption-to-savings ratio with respect to ρ. Thus constant savings rates would imply a unit EIS.

Most estimates of the EIS from the literature predominantly rely on (calendar) time series variation in rates of return. These approaches compare savings behavior in different aggregate states, which are characterized by different values of s_t and s_{t+k}. If, for example, investment opportunities are more attractive in s_t relative to s_{t+k} and other aspects of the portfolio problem are held constant, then we can partially identify the EIS by testing whether similar types of agents save more or less at time $t+k$ or t.

Given our results above, one can use many different sources of identifying variation which involve different realizations of $\{\alpha_{it}\}$ across agents, rather than the aggregate state s_t. Suppose that, due to a bad realization of α_{it}, investor i learns at t that her future investment opportunities will look less attractive—*e.g.*, because her capital income will be taxed at a higher rate or her future portfolio constraints tighten—relative to a different investor j, who had a similar value of $\alpha_{j,t-1}$ (so $\alpha_{j,t-1} \approx \alpha_{i,t-1}$) but whose realization of α_{jt} left her expectations about the individual-specific component of future investment opportunities unchanged. Then, regardless of the realization of the aggregate state s_t, if both investors have an EIS < 1, our model would predict that her savings rate would increase relative to agent j—*i.e.*, $\log(c_{it}/w_{it}) - \log(c_{jt}/w_{jt}) \leq 0$, whereas $\log(c_{i,t-1}/w_{i,t-1}) - \log(c_{j,t-1}/w_{j,t-1}) = 0$. Further note that, if both agents held similar portfolios (as would be the case if $\alpha_{i,t-1} = \alpha_{j,t-1}$), then $\Delta w_{it} = \Delta w_{jt}$, which further implies a testable restriction which may be written in terms of consumption growth: $\log(c_{it}/c_{i,t-1}) - \log(c_{jt}/c_{j,t-1}) \leq 0$. We could also substitute out consumption using the budget constraint and test equivalent conditions on savings/reinvestment rates.

These simple comparative statics lend themselves naturally to differences-in-differences estimation procedures, and are testable with panel or repeated cross-sectional data on consumption and wealth. For concreteness, imagine that the econometrician can observe a binary variable $d_{it} := d_t(\alpha_{it})$ which identifies whether or not an individual receives bad news about future investment opportunities at time t. In our example above, $d_t(\alpha_{it}) = 1$ and $d_t(\alpha_{jt}) = d_{t-1}(\alpha_{j,t-1}) = d_{t-1}(\alpha_{i,t-1}) = 0$. If the "treated" and "control" groups both have similar levels of wealth at $t-1$, and variation in d_{it} is independent of all other individual characteristics (particularly preferences and the realized return on wealth from $t-1$ to t), then one can simply compare the consumption levels of the two groups.[13] If wealth is observable or we have panel or pseudo-panel data, then the requisite identifying assumptions are even weaker. Finally, measurement errors in either consumption or wealth are no cause for concern if they are orthogonal to the instrument d_{it}.

6 Conclusion

Essentially any dynamic model involves a tradeoff between consumption today and consumption tomorrow, which is characterized by the EIS. Both positive predictions and normative implications from dynamic models depend fundamentally on this tradeoff. For example, agents' preferences about macroeconomic stabilization and predicted responses to changes in monetary and/or fiscal policy changes are fundamentally linked to the EIS. These estimates also play a critical role in determining the costs of distortionary capital taxation. Yet, despite its central role in these calculations, there remains a substantial debate about its magnitude. Our new empirical identification methodology can enable researchers to provide a new perspective on this longstanding debate.

[13] If we have panel data on consumption, then we can compare consumption growth of treated individuals with similar controls, eliminating the requirement that initial wealth levels are the same. The same idea works with repeated cross-sections if the source of identifying variation is related to spatial variation—for instance, if one state raises its capital tax rate relative to another. One could compare average consumption levels in treated states with control states after the law change, relative to the levels of similar individuals in the period prior to the change.

A Proofs

Proof of Proposition 3.1. By the budget constraint we have $c_2 = R_f(e_1 - c_1) + e_2 > 0$. Therefore by the definition of EIS,

$$\begin{aligned}\psi &= \frac{\partial \log(c_2/c_1)}{\partial \log R_f} = R_f \frac{\partial}{\partial R_f}(\log(R_f(e_1 - c) + e_2) - \log c) \\ &= R_f \left(\frac{e_1 - c - R_f \frac{\partial c}{\partial R_f}}{R_f(e_1 - c) + e_2} - \frac{1}{c}\frac{\partial c}{\partial R_f} \right) \\ \iff R_f &\left(\frac{1}{c} + \frac{R_f}{R_f(e_1 - c) + e_2} \right) \frac{\partial c}{\partial R_f} = \frac{R_f(e_1 - c)}{R_f(e_1 - c) + e_2} - \psi.\end{aligned}$$

Dividing both sides by $R_f > 0$, we obtain (3.4). □

Proof of Theorem 3.2. Let $c = c(\alpha; h)$, $v = (v + \alpha h)(w - c)$, and $c_\alpha = \mathrm{d}c/\mathrm{d}\alpha$. Then by the chain rule, at $\alpha = 0$ we obtain

$$\frac{\mathrm{d}}{\mathrm{d}\alpha} \log(c/v) = \frac{c_\alpha}{c} - \frac{-v' c_\alpha + h}{v}.$$

Since $c = c(\alpha; h)$ solves

$$\max_c f(c, (v + \alpha h)(w - c)),$$

by the first-order condition we have

$$f_c - f_v(v' + \alpha h') = 0 \iff \log \frac{f_c}{f_v} = \log(v' + \alpha h'),$$

where v', h' are evaluated at $w - c$. Again by the chain rule,

$$\frac{\mathrm{d}}{\mathrm{d}\alpha} \log \frac{f_c}{f_v} = \frac{-v'' c_\alpha + h'}{v'}.$$

By the definition of EIS, we obtain

$$\begin{aligned}\psi &= -\frac{\frac{c_\alpha}{c} - \frac{-v' c_\alpha + h}{v}}{\frac{-v'' c_\alpha + h'}{v'}} \\ \iff \frac{\psi}{v'}(-v'' c_\alpha + h') &= -\frac{c_\alpha}{c} + \frac{1}{v}(-v' c_\alpha + h) \\ \iff \left(\frac{1}{c} + \frac{v'}{v} - \psi \frac{v''}{v'} \right) c_\alpha &= \frac{h}{v} - \psi \frac{h'}{v'},\end{aligned}$$

which is (3.14). If $h(w) = w$, then the right-hand side becomes

$$\frac{w - c}{v(w - c)} - \psi \frac{1}{v'(w - c)} = \frac{1}{v'(w - c)}(\varepsilon(w - c) - \psi),$$

where $\varepsilon(w) = wv'(w)/v(w)$ is the elasticity of v. Since the value function is increasing in wealth, we have $v' > 0$. Therefore if $v'' \le 0$, then $\frac{1}{c} + \frac{v'}{v} - \psi \frac{v''}{v'} > 0$ unambiguously, so $\mathrm{d}c/\mathrm{d}\alpha \gtreqless 0$ according as $\psi \lesseqgtr \varepsilon$. □

Lemma A.1. *Let X be an almost surely positive random variable and suppose that $\mathrm{E}[X^r]$ is finite for $0 < |r| < \epsilon$ and $\mathrm{E}[\log X]$ is finite. Then*

$$\lim_{r \to 0} \mathrm{E}[X^r]^{\frac{1}{r}} = \exp(\mathrm{E}[\log X]).$$

Proof. It suffices to prove when X is a discrete random variable (simple function) since the Lebesgue integral of a measurable function is defined by the limit of the integrals of approximating simple functions. Suppose that X takes values $x_1, \dots, x_N$ with probability $p_1, \dots, p_N$, and let

$$f(r) = \log \mathrm{E}[X^r] = \log \left(\sum_{n=1}^{N} p_n x_n^r \right).$$

Since $f(0) = \log \left(\sum_{n=1}^{N} p_n \right) = 0$, it follows that

$$\begin{aligned} \lim_{r \to 0} \frac{1}{r} \log \mathrm{E}[X^r] &= \lim_{r \to 0} \frac{f(r) - f(0)}{r} = f'(0) \\ &= \left. \frac{\sum_{n=1}^{N} p_n x_n^r \log x_n}{\sum_{n=1}^{N} p_n x_n^r} \right|_{r=0} = \sum_{n=1}^{N} p_n \log x_n = \mathrm{E}[\log X]. \end{aligned}$$

Therefore $\lim_{r \to 0} \mathrm{E}[X^r]^{\frac{1}{r}} = \exp(\mathrm{E}[\log X])$. □

Lemma A.2. *Let X be an almost surely positive random variable and define $\phi : \mathbb{R} \to [0, \infty]$ by*

$$\phi(r) = \begin{cases} \mathrm{E}[X^r]^{\frac{1}{r}}, & (r \neq 0) \\ \exp(\mathrm{E}[\log X]). & (r = 0) \end{cases}$$

Then ϕ is increasing in $r \in \mathbb{R}$.

Proof. Let $p, q > 1$ be numbers such that $1/p + 1/q = 1$. Let $\|f\|_p = \left(\int |f|^p \, \mathrm{d}\mu \right)^{\frac{1}{p}}$ denote the L^p norm of a function f. Let $f = X^r$, $g = 1$, and $s = pr$. By Hölder's inequality $\|fg\|_1 \leq \|f\|_p \|g\|_q$, we obtain

$$\mathrm{E}[X^r] \leq \mathrm{E}[X^{pr}]^{\frac{1}{p}} \iff \begin{cases} \mathrm{E}[X^r]^{\frac{1}{r}} \leq \mathrm{E}[X^s]^{\frac{1}{s}}, & (r > 0) \\ \mathrm{E}[X^r]^{\frac{1}{r}} \geq \mathrm{E}[X^s]^{\frac{1}{s}}. & (r < 0) \end{cases}$$

Noting that $s = pr \gtrless r$ according as $r \gtrless 0$ since $p > 1$, it follows that $\mathrm{E}[X^r]^{\frac{1}{r}}$ is increasing in r for $r \in (-\infty, 0)$ and $r \in (0, \infty)$. Since $\phi(r)$ is continuous at $r = 0$ by Lemma A.1, it is increasing on $\mathbb{R}$. □

References

Rui Albuquerque, Martin S. Eichenbaum, and Sergio Rebelo. Valuation risk and asset pricing. NBER Working Paper, 2012. URL `http://www.nber.org/papers/w18617`.

Orazio P. Attanasio and Guglielmo Weber. Consumption growth, the interest rate and aggregation. *Review of Economic Studies*, 60(3):631–649, July 1993. doi:10.2307/2298128.

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, August 2004. doi:10.1111/j.1540-6261.2004.00670.x.

Paul Beaudry and Eric van Wincoop. The intertemporal elasticity of substitution: An exploration using a US panel of state data. *Economica*, 63(251): 495–512, August 1996. doi:10.2307/2555019.

Jason Beeler and John Y. Campbell. The long-run risks model and aggregate asset prices: An empirical assessment. *Critical Finance Review*, 1(1):141–182, 2012. doi:10.1561/104.00000004.

Michael Carlos Best, James Cloyne, Ethan Ilzetzki, and Henrik Jacobsen Kleven. Interest rates, debt and intertemporal allocation: Evidence from notched mortgage contracts in the United Kingdom. Staff Working Paper 543, Bank of England, 2015. URL `http://www.bankofengland.co.uk/research/Documents/workingpapers/2015/swp543.pdf`.

John Y. Campbell. Consumption-based asset pricing. In George M. Constantinides, Milton Harris, and René M. Stultz, editors, *Handbook of the Economics of Finance*, volume 1B, chapter 13, pages 803–887. Elsevier, Amsterdam, 2003. doi:10.1016/S1574-0102(03)01022-7.

John Y. Campbell and N. Gregory Mankiw. Consumption, income, and interest rates: Reinterpreting the time series evidence. In Olivier Blanchard and Stanley Fischer, editors, *NBER Macroeconomics Annual*, volume 4, pages 185–246. MIT Press, Cambridge, MA, 1989.

Christopher Dixon Carroll. Death to the log-linearized consumption Euler equation! (and very poor health to the second-order approximation). *Advances in Macroeconomics*, 1(1), 2001. doi:10.2202/1534-6013.1003.

David Cashin and Takashi Ueyama. Measuring intertemporal substitution in consumption: Evidence from a VAT increase in Japan. *Review of Economics and Statistics*, 98(2):285–297, May 2016. doi:10.1162/REST_a_00531.

Mariano Massimiliano Croce. Long-run productivity risk: A new hope for production-based asset pricing? *Journal of Monetary Economics*, 66:13–31, September 2014. doi:10.1016/j.jmoneco.2014.04.001.

Itamar Drechsler and Amir Yaron. What's vol got to do with it. *Review of Financial Studies*, 24(1):1–45, 2011. doi:10.1093/rfs/hhq085.

Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, November 2003. doi:10.1016/S0022-0531(03)00097-8.

Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, July 1989. doi:10.2307/1913778.

Nicolae Gârleanu and Stavros Panageas. Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3):670–685, June 2015. doi:10.1086/680996.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989. doi:10.1016/0304-4068(89)90018-9.

Jonathan Gruber. A tax-based estimate of the elasticity of intertemporal substitution. *Quarterly Journal of Finance*, 3(1):1350001, 2013. doi:10.1142/S2010139213500018.

Fatih Guvenen. Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. *Journal of Monetary Economics*, 53(7):1451–1472, October 2006. doi:10.1016/j.jmoneco.2005.06.001.

Nils H. Hakansson. Optimal investment and consumption strategies under risk for a class of utility functions. *Econometrica*, 38(5):587–607, September 1970. doi:10.2307/1912196.

Nils H. Hakansson. Optimal entrepreneurial decisions in a completely stochastic environment. *Management Science*, 17(7):427–449, March 1971. doi:10.1287/mnsc.17.7.427.

Robert E. Hall. Intertemporal substitution in consumption. *Journal of Political Economy*, 96(2):339–357, April 1988.

Lars Peter Hansen and Kenneth J. Singleton. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91(2):249–265, April 1983.

Tomas Havranek, Roman Horvath, Zuzana Irsova, and Marek Rusnak. Cross-country heterogeneity in intertemporal substitution. *Journal of International Economics*, 96(1):100–118, May 2015. doi:10.1016/j.jinteco.2015.01.012.

Tomáš Havránek. Measuring intertemporal substitution: The importance of method choices and selective reporting. *Journal of the European Economic Association*, 13(6):1180–1204, December 2015. doi:10.1111/jeea.12133.

Takashi Hayashi. Intertemporal substitution, risk aversion, and ambiguity aversion. *Economic Theory*, 25(4):933–956, 2005. doi:10.1007/s00199-004-0508-2.

Takashi Hayashi and Jianjun Miao. Intertemporal substitution and recursive smooth ambiguity preferences. *Theoretical Economics*, 6(3):423–472, September 2011. doi:10.3982/TE843.

Nengjiu Ju and Jianjun Miao. Ambiguity, learning, and asset returns. *Econometrica*, 80(2):559–591, March 2012. doi:10.3982/ECTA7618.

Georg Kaltenbrunner and Lars A. Lochstoer. Long-run risk through consumption smoothing. *Review of Financial Studies*, 23(8):3190–3224, 2010. doi:10.1093/rfs/hhq033.

David M. Kreps and Evan L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200, January 1978. doi:10.2307/1913656.

David Levhari and T. N. Srinivasan. Optimal savings under uncertainty. *Review of Economic Studies*, 36(2):153–163, April 1969. doi:10.2307/2296834.

Sydney C. Ludvigson and Christina Paxson. Approximation bias in linearized Euler equations. *Review of Economics and Statistics*, 83(2):242–256, 2001. doi:10.1162/00346530151143789.

N. Gregory Mankiw and Stephen P. Zeldes. The consumption of stockholders and nonstockholders. *Journal of Financial Economics*, 29(1):97–112, March 1991. doi:10.1016/0304-405X(91)90015-C.

Robert C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics*, 51(3):247–257, August 1969. doi:10.2307/1926560.

Casey B Mulligan. Capital, interest, and aggregate intertemporal substitution. NBER Working Paper, 2002. URL `http://www.nber.org/papers/w9373.pdf`.

Edmund S. Phelps. The accumulation of risky capital: A sequential utility analysis. *Econometrica*, 30(4):729–743, October 1962. doi:10.2307/1909322.

John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32 (1-2):122–136, January-April 1964. doi:10.2307/1913738.

Michael Rothschild and Joseph E. Stiglitz. Increasing risk II: Its economic consequences. *Journal of Economic Theory*, 3(1):66–84, March 1971. doi:10.1016/0022-0531(71)90034-2.

Paul A. Samuelson. Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51(3):239–246, August 1969. doi:10.2307/1926559.

Agnar Sandmo. The effect of uncertainty on saving decisions. *Review of Economic Studies*, 37(3):353–360, July 1970. doi:10.2307/2296725.

Lawrence D. W. Schmidt. Climbing and falling off the ladder: Asset pricing implications of labor market event risk. 2014. URL `http://ssrn.com/abstract=2471342`.

Alexis Akira Toda. Incomplete market dynamics and cross-sectional distributions. *Journal of Economic Theory*, 154:310–348, November 2014. doi:10.1016/j.jet.2014.09.015.

Annette Vissing-Jørgensen. Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy*, 110(4):825–853, August 2002. doi:10.1086/340782.

Philippe Weil. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3):401–421, November 1989. doi:10.1016/0304-3932(89)90028-7.

Motohiro Yogo. Estimating the elasticity of intertemporal substitution when instruments are weak. *Review of Economics and Statistics*, 86(3):797–810, August 2004. doi:10.1162/0034653041811770.