The Journal of Risk and Uncertainty, 31:3; 243–262, 2005

Relative Risk Aversion: What Do We Know?

DONALD J. MEYER donald.meyer@wmich.edu
Department of Economics, Western Michigan University, Kalamazoo, MI 49008

JACK MEYER* jmeyer@msu.edu
Department of Economics, Michigan State University, East Lansing, MI 48824

Abstract

The relative risk aversion measure that represents the risk preferences of a decision maker depends on the outcome variable that is used as the argument of the utility function, and on the way that outcome variable is defined or measured. In addition, the relationship between any two such relative risk aversion measures is determined by the relationship between the corresponding outcome variables. These well-known facts are used to adjust several reported estimates of relative risk aversion so that those estimates can be directly compared with one another. After adjustment, the significant variation in the reported relative risk aversion measures for representative decision makers is substantially reduced.

Keywords: risk aversion, relative risk aversion

JEL Classification: D81

When economists model choices made by expected utility maximizing decision makers, a variety of variables are used to represent the possible outcome of the decision being represented. Wealth or consumption are most often used, although a number of other variables, including profit, income, return, or rate of return, are sometimes selected. In addition to using different variables as the argument of the utility function, the various studies also frequently define or measure the chosen variable in different ways. Measures of wealth, for instance, often, but not always, exclude the value of human capital. Similarly, the return from investment in a portfolio of assets can be calculated either before or after taxes are paid on investment income.

These variations in the way the outcome variable of a risky choice is defined or measured significantly alter the relative risk aversion measure determined for the decision maker. The size can be changed by an order of magnitude, the slope can be made much steeper or flatter, and even the sign of the slope can be reversed. Although the measure of relative risk aversion is invariant to the unit in which the outcome variable is measured, this elasticity measure is sensitive to what is included or excluded when defining or measuring a variable. Recall that relative risk aversion for a particular outcome variable is the percentage change in marginal utility from that variable divided by the percentage change in the variable itself. Including or excluding components when measuring or defining a variable has a different impact on the denominator and numerator of this ratio. The analysis shows that whenever

*To whom correspondence should be addressed.

two outcome variables are anything other than proportional to one another, their relative risk aversion measures must differ, and the way they differ can be predicted if one knows how the variables are related to one another.

Consider a very simple example where Y represents the random *return* on an investment, and X represents the random *rate of return* on the same investment. By definition it is the case that $Y = 1 + X$. The rate of return X represents the gain or loss from investing a dollar, while Y includes the return of the original dollar with this gain or loss. This relationship between return and rate of return implies that the relative risk aversion measure for return is $\frac{(1+X)}{X}$ times the relative risk aversion measure for rate of return. Thus, when analyzing mutual fund annual performance, for instance, at the mean rate of return of 9%, relative risk aversion for return is approximately twelve times larger than that for rate of return. Furthermore, when the former has a slope of zero, the latter must have a positive slope.

Failure to fully recognize that the way the outcome variable is defined or measured impacts the measure of relative risk aversion has contributed significantly to the confusion that prevails concerning what is known about relative risk aversion measures, and risk preferences in general. There are instances where the results from one study are viewed as contradicting those from another, when, after adjusting for differences in outcome variables, the results support one another. The opposite occurs even more frequently. Evidence concerning the slope and level of relative risk aversion for one outcome variable is often cited as support for that same assumption for an entirely different outcome variable, even when the two relative risk aversion measures cannot be the same and must differ in a predictable way. Even after forty years of published research concerning measures of relative risk aversion, this confusion is viewed by some as so severe that the existing body of evidence is dismissed in favor of introspection.[1]

This variation in the choice of outcome variable and the way it is measured has also lead to the scattered and disconnected nature of the information that is available. The risk preferences of agricultural producers, or consumers, or stock market investors, are each examined separately even though, *a priori*, there is no reason to expect representative agents from these groups to differ in their attitudes toward risk. The estimated relative risk aversion measures across groups are perceived to differ, in part, because of a failure to distinguish between variation in outcome variables across groups, and variation in their risk attitudes.

The purpose of this study is to help remedy this situation by carefully reviewing a small portion of the information that others have provided concerning relative risk aversion measures for expected utility maximizing decision makers. The focus is on a careful examination of a few studies rather than a more cursory examination of many studies. In part this is because a longer paper on this topic is underway, but it also because this work is meant to provide a framework so that anyone can properly interpret any relative risk aversion information that is made available. In each of the studies examined, the reported measure of relative risk aversion is taken as it is given, and no attempt is made to evaluate the methodology employed in arriving at that estimate. Special attention, however, is paid to how the

[1] Feldstein and Ranguelova (2001), for instance, dismiss the existing evidence and suggest a thought experiment to allow readers to choose the appropriate assumption to make concerning the magnitude of relative risk aversion.

outcome variable is defined or measured, and the effect that this has on the reported measure of relative risk aversion.

The paper focuses on the case where outcome variables, Y and X, are related by a differentiable function $Y = Y(X)$. For this case, the relative risk aversion measures for X and Y, $R_u(X)$ and $R_v(Y)$, are related by $R_u(X) = R_v(Y)\frac{(X \cdot Y')}{Y} - \frac{Y'' \cdot X}{Y'}$. Using the definitions and measures of the outcome variable provided in published research, an attempt is made to determine the function $Y(X)$ relating pairs of outcome variables. This analysis is sometimes carried out within a single study where the "same" outcome variable is defined or measured in more than one way. At other times this involves comparison across two different studies where the outcome variables employed are not the same in the two studies. In the course of this discussion, it proves useful to choose a reference outcome variable to which all others can be related, and the Arrow-Pratt measure of wealth is identified and chosen to be this reference variable.

The basic conclusions drawn from the evidence reviewed here include the following. Theoretical discussion, comparative static results, and the evidence provided by Friend and Blume (1975), and others, indicates that relative risk aversion for wealth, using the narrow definition provided by Arrow (1965, 1971) and Pratt (1964), is likely larger than but near one and increasing slightly or constant. The evidence concerning relative risk aversion for consumption, where consumption is broadly measured such as aggregate consumption in national income account data or permanent consumption, indicates that at the mean, relative risk aversion for consumption is on the order of five times larger than that for Arrow-Pratt wealth, and likely to be rather steeply negatively sloped. This evidence comes from Constantinides (1990) and Barsky et al. (1997) and others. Moreover, the relationship between Arrow-Pratt wealth and these comprehensive measures of consumption implies that these two findings are not contradictory.

The paper has two parts. The first part provides a framework for comparing and contrasting information concerning relative risk aversion from the various theoretical and empirical studies with their differing outcome variables. After this first step, the second part of the analysis begins the task of gathering, and translating to this common scale, the information concerning risk aversion acquired during the past forty years. The work chosen for review are among the more prominent studies. A few papers are reviewed in considerable detail, while the others are discussed only briefly. Information obtained from theoretical analysis as well as empirical evidence is examined.

The main focus of this paper is on how to perform this translation of information concerning relative risk aversion to a common measurement scale. An important and perhaps unexpected secondary contribution, however, is the finding that the evidence reviewed here concerning relative risk aversion for various decision makers, and for a variety of outcome variables, is more unified and consistent than casual inspection would indicate. Both the slopes and magnitudes of relative risk aversion measures are less different across decision makers, and across a variety of decisions, than the unadjusted measures seem to imply. Table 2 near the end of the paper summarizes these findings.

The paper is organized as follows. First, in the next section, the framework of discussion and the notation used in the paper are given. Following this discussion are two separate sections reviewing information concerning relative risk aversion for wealth, and for consumption, respectively. Table 2 summarizes the findings of the analysis. Finally, the paper concludes with a brief summary of what is known concerning relative risk aversion, both level and slope, for representative agents making risky economic decisions.

1. Notation and framework of analysis

For forty years now, the measures of risk aversion defined by Pratt (1964) and Arrow (1965, 1971) have been used to describe the risk attitudes of expected utility maximizing decision makers. These measures apply most directly to a utility function that has a single argument, and this is the only case considered here. The expected utility maximizing decision maker under consideration is assumed to make choices, and these choices then combine with the realizations of random parameters to yield a single dimension random outcome, such as wealth or consumption, over which utility is defined. For the purposes of this work, the decision model need not be specified more precisely than this.

Initial discussion uses X and Y to represent two different random outcome variables, or perhaps the same variable measured or defined in two different ways. For a particular decision maker, $u(X)$ and $v(Y)$ denote the utility function whose expected value is maximized when the outcome variable is X or Y, respectively. Even though these two utility functions are associated with, and represent the risk preferences of a single decision maker, they are likely to differ due to the different outcome variables being considered. How $u(X)$ and $v(Y)$ are related depends on how the outcome variables themselves are related.

Given utility function $u(X)$ or $v(Y)$, the definitions of Pratt and Arrow specify two primary risk aversion measures employed in economic analysis, absolute and relative risk aversion. Let $A_u(X) = \frac{-u''(X)}{u'(X)}$ and $R_u(X) = \frac{-u''(X) \cdot X}{u'(X)}$ denote the absolute and relative risk aversion measure for utility function $u(X)$ respectively, with a similarly defined $A_v(Y)$ and $R_v(Y)$ denoting those for $v(Y)$. Later on Menezes and Hanson (1970) and Zeckhauser and Keeler (1970) define a measure of partial risk aversion, which is a more general version of relative risk aversion that permits the argument of utility be split into two components, one of which is random and the other is not.

Each of these three risk aversion measures has advantages when doing comparative static analysis in particular settings. Often a simple condition on one of the three measures, rather than a more cumbersome condition on either of the other two, is sufficient to demonstrate the desired result. For instance, when wealth increases, constant absolute risk aversion is sufficient to imply that the allocation to the risky asset remains fixed. Of course, assuming constant absolute risk aversion is the same as requiring that relative risk aversion increases proportionally to X. Similarly, constant relative risk aversion implies that the proportion of wealth allocated to the risky asset remains unchanged as wealth changes. Other such simplifications can result from using the partial risk aversion measure as demonstrated by Menezes and Hanson, and Zeckhauser and Keeler when introducing the definition, and later by Cheng et al. (1987).

The focus here is on representing and measuring the risk preferences of the decision maker, not on comparative static analysis of a particular decision. Any one of these three risk aversion measures is sufficient for this, and since simple arithmetic allows one to go from one to the others, only one of the three is needed. Relative risk aversion is selected as the focus here because this measure is an elasticity and is unaffected by changes in the unit in which the outcome variable is measured. Relative risk aversion is the elasticity of marginal utility, the percentage change in marginal utility divided by the percentage change in that outcome variable. Since relative risk aversion is not sensitive to the unit in which the outcome variable is measured, reported results from the various studies are more directly comparable than would be the case if either of the other measures had been used. Even so, as the analysis below demonstrates, the relative risk aversion measure is sensitive to what is included or excluded when defining or measuring a variable. Relative risk aversion measures have well known difficulties with zero and negative values for X and Y, and therefore only outcome variables such that $X > 0$ and $Y > 0$ are considered.

Although Y and X are different outcome variables, these outcome measures often can be related to one another. If so, the nature of the relationship between Y and X determines how the risk aversion measures for $u(X)$ must relate to those for $v(Y)$. When the relationship between Y and X is not stochastic it is denoted $Y = Y(X)$, and is assumed to be differentiable at least twice. This relationship could result from the definitions of X and Y themselves, as is the case when Y is return on an investment of one unit and X is rate of return on that same investment implying that $Y = 1 + X$. There are other sources for the relationship between Y and X. There are cases where Y and X are linked within an economic model, with the relationship resulting from optimizing behavior by the decision maker. For example, in a multi-period consumption model with random return on saved wealth, it is often possible to determine the optimal level of consumption, Y, for any value for wealth X.

An even more frequently observed source of a relationship between outcome variables in the work reviewed here is the use of different definitions or measures for a particular outcome. Empirical measures of the "same" variable often differ because of data limitations or differences in opinion of what should be included. Similarly, theoretical definitions depend on which other variables, such as tax variables, are included in the analysis and the definition can differ across studies. In either of these situations, it is usual for X and Y to be discussed as if they represent the same outcome variable, while careful examination indicates that they are not precisely the same and sometimes differ in significant ways.

For example, Y and X can each be referred to as the return from investing in a portfolio of risky and riskless assets, but Y is the outcome from a more complicated model that includes the taxation of investment income, while X represents return without consideration of taxes. The inclusion of a tax variable in the one definition of return and not the other alters the relative risk aversion measure that is obtained. Similarly, when measuring wealth, the value of human capital is often not observed, and housing is sometimes viewed as illiquid and consumption related, and thus some but not all wealth measures exclude these components. This omission affects, in a predictable way, the relative risk aversion measure that is estimated.

One of the main tasks in reviewing the literature providing information concerning relative risk aversion is to determine exactly how the outcome variable used in a particular study is defined and measured. Unfortunately, many published works do not give enough information to make this determination. The papers reviewed here were selected in part because they provide the needed information. In addition to knowing precisely what outcome variable is used, it is also important to be able to relate this variable to other outcome variables used elsewhere. Without this, the various reported findings cannot be appropriately compared and contrasted.

The only type of relationship between outcome variables considered in any detail in this analysis is when Y and X are deterministically related; that is, when $Y = Y(X)$, and each realization of the random outcome Y is determined by the corresponding realization of X. When Y and X are related this way, then $u(X) = v(Y(X))$ is the corresponding relationship between utility functions.[2]

When $u(X) = v(Y(X))$, the relationship between the risk aversion measures for $u(X)$ and $v(Y)$ is a simple calculation. $u(X) = v(Y(X))$ implies that

$$u'(X) = v'(Y) \cdot Y' \tag{1}$$

and

$$u''(X) = v''(Y) \cdot (Y')^2 + v'(Y) \cdot Y'' \tag{2}$$

This implies that

$$R_u(X) = \frac{-u''(X) \cdot X}{u'(X)} = \frac{[-v''(Y)(Y')^2 - v'(Y) \cdot Y'']X}{v'(Y) \cdot Y'} \tag{3}$$

$$R_u(X) = \frac{(-v''(Y)Y) \cdot X \cdot Y'}{v'(Y) \cdot Y} - \frac{Y'' \cdot X}{Y'} \tag{4}$$

$$R_u(X) = R_v(Y)\frac{(X \cdot Y')}{Y} - \frac{Y'' \cdot X}{Y'} \tag{5}$$

Equation (5) gives the relationship between the relative risk aversion measures for $u(X)$ and $v(Y)$. The last term on the right hand side of (5) is the equivalent of the relative risk aversion measure for the relationship $Y = Y(X)$. It is zero when the relationship between Y and X is linear so that $Y'' = 0$. Even so, unless Y is proportional to X, $R_u(X)$ and $R_v(Y)$ still differ. This case is discussed in detail next.

One simple form for $Y = Y(X)$ is frequently observed, and hence warrants special notation and attention. This is the case where Y and X are linearly related, that is, $Y = a + bX$. The relative risk aversion measures for $u(X)$ and $v(Y)$ then satisfy

$$R_u(X) = R_v(Y)\left[\frac{bX}{a + bX}\right] = R_v(Y)\left[\frac{bX}{Y}\right] \tag{6}$$

Notice that even though the function relating Y and X is linear and thus the term in (5) measuring the curvature of $Y(X)$ is zero, the two relative risk aversion measures are not the same or even a constant multiple of one another when the intercept in the linear relationship is not zero. Because this linear case plays an important role in the discussion that follows,

[2]More complex analysis results when the relationship between Y and X is not completely deterministic and these cases are only briefly discussed.

a formal statement of four of the implications of $Y = a + bX$ for the relationship between the two relative risk aversion measures is useful.

Theorem 1. *If $Y = a + b \cdot X$ with $a > 0$ and $b > 0$, then*
(a) $R_u(X) \leq R_v(Y)$
(b) $R'_u(X) > 0$ *when* $R'_v(Y) = 0$ *and* $R'_v(Y) < 0$ *when* $R'_u(X) = 0$
(c) $\frac{R'_u(X) \cdot X}{R_u(X)} = [\frac{bX}{Y}][\frac{R'_v(Y) \cdot Y}{R_v(Y)}] + [\frac{a}{Y}][1]$
(d) $\frac{R'_u(X) \cdot X}{R_u(X)} \geq \frac{R'_v(Y) \cdot Y}{R_v(Y)}$ *if and only if $v(Y)$ displays decreasing absolute risk aversion*

Proof: When a, b and X are each positive, the $\frac{bX}{a+bX}$ term in Eq. (6) is less than one and increasing in X and approaches one from below as X increases. This positive slope for $\frac{bX}{a+bX}$ implies that $R_v(Y)$ is negatively sloped when the slope of R_u(X) is zero and $R_u(X)$ is positively sloped when $R_v(Y)$ has zero slope. Thus, (a) and (b) are demonstrated. To demonstrate (c) and (d) differentiate (6) to obtain:

$$R'_u(X) = R'_v(Y)\frac{b^2 X}{Y} + R_v(Y)\frac{(bY - b^2 X)}{Y^2}$$

Dividing by $R_u(X)$ on the left side and on the right side by its equivalent $R_v(Y)[\frac{bX}{Y}]$ gives

$$\frac{R'_u(X)}{R_u(X)} = b\left[\frac{R'_v(Y)}{R_v(Y)}\right] + \frac{a}{XY}$$

Converting this to elasticity form gives

$$\frac{R'_u(X) \cdot X}{R_u(X)} = \left[\frac{bX}{Y}\right]\left[\frac{R'_v(Y) \cdot Y}{R_v(Y)}\right] + \left[\frac{a}{Y}\right][1]$$

To show (d) observe that decreasing absolute risk aversion for $v(Y)$ is equivalent to $R_v(Y)/Y$ decreasing in Y. Differentiating $R_v(Y)/Y$ shows that this is equivalent to $R'_v(Y) \leq R_v(Y)/Y$.
□

The reason for no clear statement in Theorem 1 concerning the relative sizes of the slopes of the two relative risk aversion measures is that one can simultaneously change the unit of measure for X and the magnitude of b so as to preserve the relationship $Y = a + bX$. This implies that the size of b is arbitrary and makes statements concerning the relative slopes of the two risk aversion measures difficult to formulate. As a consequence, it is more convenient to formulate statements concerning the relative rates of change of $R_u(X)$ or $R_v(Y)$ by making statements about their relative elasticity values. This is what properties (c) and (d) in Theorem 1 do. Property (c) indicates that the elasticity of $R_u(X)$ is a convex combination of the elasticity of $R_v(Y)$ and one. Thus, the elasticity of $R_u(X)$ is always between that for $R_v(Y)$ and one. Property (d) indicates that decreasing absolute risk aversion for v(Y) is necessary and sufficient for the elasticity of $R_v(Y)$ to be less than one, and therefore the elasticity of $R_u(X)$ is larger than that for $R_v(Y)$. These elasticity statements hold when the two relative risk aversion measures are both increasing, are both decreasing, and when $R_u(X)$ is increasing and $R_v(Y)$ is decreasing.

A linear relationship between Y and X can occur as an equilibrium condition in an economic model, and also from measurement or definitional disparities. For instance, in multi-period consumption models, certain utility functions imply that optimal consumption is a linear function of beginning of period wealth. Moreover, the intercept, a, in these models is positive, reflecting the fact that optimal consumption decisions ensure that consumption is less variable than wealth.[3] Thus, in those models, the relative risk aversion for utility from consumption is necessarily larger than relative risk aversion for utility from wealth. Also, the two distinct relative risk aversion measures can be such that the one is decreasing in consumption while the other is increasing in wealth. Meyer and Meyer (2004) use such a linear relationship between consumption and wealth to argue that the equity premium puzzle disappears when one uses a utility function for consumption that displays decreasing relative risk aversion, and furthermore show that such a utility function is consistent with relative risk aversion for wealth being constant and of a much smaller magnitude.

The linear relationship, $Y = a + bX$, can also result from measurement differences. Specifically this occurs when components are included with Y and are not included with X, and these excluded components are either independent of the level of X, the intercept portion, or are proportional to X. When the independent components excluded from X but included with Y are positive, so that Y is a broader and more inclusive measure than is X, then the conditions in Theorem 1 are satisfied and $R_v(Y)$ is larger than $R_u(X)$, and the elasticity of $R_v(Y)$ is smaller than that for $R_u(X)$ under decreasing absolute risk aversion (DARA).

A second type of relationship between Y and X that is sometimes observed is not completely deterministic. There can be random components that are included with Y but not X so that $Y = X + Z$, where Z is random. When Z and X are independently distributed the relationship between $u(X)$ and $v(Y)$ is given by $u(X) = E_Z v(Y) = E_Z v(X + Z)$. If Z has a zero mean, Z is referred to as background risk. The implications of background risk for the risk aversion properties of $u(X)$ and $v(Y)$ have been discussed in the literature. In the analysis here, the additional component, Z, is typically one with a positive support rather than a zero mean. When Z is random with positive support, the findings reported in Theorem 1 are unchanged. Adding such a Z increases the magnitude and decreases the elasticity of relative risk aversion, assuming DARA.[4]

To compare the relative risk aversion measures obtained in studies with two different or differently measured outcome variables, only the relationship between those two variables, Y and X, is needed. To summarize and examine that same information from a larger number of studies with many outcome variables, however, it is convenient to pick a particular reference outcome variable, X, and to relate all the Y_i from the various studies to this reference variable. Doing this allows the relative risk aversion information from any particular study to be adjusted to a common measurement scale.

[3]Gollier (2002) refers to this as time diversification.

[4]Pratt (1964) discusses the risk aversion properties of the convex combinations of utility functions that result when an additive independent variable is included.

The question then is which X to choose when defining this common scale? In a formal sense, it does not matter, but in a practical sense, familiarity is an important consideration. For this reason, wealth is chosen as the reference outcome variable. Furthermore, the definition of wealth that is used includes with wealth the value of all assets, and only those assets, whose quantities can be freely adjusted. This type of wealth is referred to as Arrow-Pratt (A-P) wealth.

A-P wealth is selected as the reference outcome variable for several reasons. The main reason, however, is that it is the variable used by Arrow and by Pratt in the majority of their analysis, and hence goes back to the earliest discussion of measures of relative risk aversion. This particular definition of wealth also continues to be that employed in theoretical analysis of portfolio behavior. With the selection of a reference variable, it is now appropriate to alter the general notation to reflect this. From here forward, W rather than X is used as the argument of utility $u(\cdot)$, and $R_u(W)$ denotes the relative risk aversion measure for $u(W)$. $R_u(W)$ is referred to as the Arrow-Pratt (A-P) measure of relative risk aversion. In the discussion that follows, attempts are made to adjust all estimates of relative risk aversion for any outcome variable, so that they provide direct information concerning $R_u(W)$. With this framework in place, the task of reviewing the available information concerning measures of relative risk aversion can begin.

2. Relative risk aversion for wealth

As was indicated earlier, Arrow and Pratt provide definitions of measures of absolute and relative risk aversion. These definitions are formulated for an outcome variable or argument of utility referred to as wealth, and this wealth is assumed to be freely allocable among the available assets. This section reviews a portion of the information concerning relative risk aversion for wealth that has accumulated since those definitions were introduced. The review begins with information relying primarily on theoretical analysis.

One way to determine the properties of a risk aversion measure, is to ask which properties lead to comparative static predictions that are consistent with observation, and which do not. This approach has been used extensively in the expected utility literature to determine plausible properties of the utility function, including those concerned with the magnitude and slope of relative risk aversion. Early on, Arrow argued that $R_u(W)$ displays increasing relative risk aversion (IRRA), and is approximately one. These two properties of $R_u(W)$ have become a reference point in the ensuing discussion.

IRRA for $R_u(W)$ is supported by the observation that in the portfolio model with one risky and one riskless asset, IRRA implies that increases in wealth lead to a decreased proportion of wealth invested in the risky asset, and that the opposite occurs under decreasing relative risk aversion (DRRA) for $R_u(W)$. Arrow cites evidence concerning the wealth elasticity of cash balance or money demand, interpreted as demand for the riskless asset, in support of IRRA.

Arrow's argument that $R_u(W)$ should be near one is based on a limiting argument.[5] He notes that for utility to be bounded from below as wealth approaches zero, $R_u(W)$ must

[5] It is important to recognize that Arrow's argument applies to any outcome variable and its utility function, and thus applies to all relative risk aversion measures.

approach a number less than one, and to be bounded from above as wealth becomes infinite, $R_u(W)$ must approach a number greater than one. Thus, $R_u(W)$, if constant should equal one, and when not constant must increase at least over some portion of the domain $[0, \infty]$. This argument does not, however, rule out values for $R_u(W)$ less than or greater than one except at the extremes. Empirical evidence concerning the magnitude and slope of $R_u(W)$ is mixed. Nonetheless, Arrow's conclusions concerning $R_u(W)$ have significantly influenced assumptions made concerning measures of relative risk aversion for wealth, and often inappropriately, for other outcome variables as well.

Cass and Stiglitz (1972) employ a more general portfolio model with many risky assets, and determine which aspects of Arrow's comparative static findings continue to hold. They are able to show that while IRRA is not sufficient to predict the effect of wealth changes on the exact composition of the risky portion of the portfolio, IRRA does lead to a similar conclusion concerning the relative share of the portfolio allocated to the riskless asset. Cass and Stiglitz also assume that all wealth is freely allocable across the various assets.

Another important comparative static finding associated with relative risk aversion for A-P wealth is presented by Hirshleifer and Riley (1992) (H-R). They show that in a state preference model satisfying the axioms of expected utility, relative risk aversion less than one implies that the price elasticity of demand for any state contingent claim is greater than one, and that contingent claims are gross substitutes rather than gross complements. Furthermore, the reverse is true. Relative risk aversion greater than one implies that state contingent claims are gross complements, and that own price elasticity is less than one. This comparative static result is interpreted as lending some support for the assumption that $R_u(W)$ is less than one. H-R assume that the initial endowment of contingent claims can be freely traded, an assumption consistent with that made for A-P wealth.

As a final statement pertaining to theoretical findings concerning relative risk aversion for wealth, it is noted that there exists a rather large body of work by a number of researchers, including (Hadar and Seo, 1993), who show that IRRA and $R_u(W) \leq 1$ lead to "sensible" comparative static findings. Hadar and Seo demonstrate, for instance, that a first degree stochastic dominant (FSD) improvement in the return to the risky asset, always leads to a increase in the allocation to that asset when $R_u(W)$ is less than one, a seemingly sensible result, while the opposite can occur when $R_u(W)$ is greater than one. Other similarly "sensible" comparative static findings can be listed as support for the assumption of $R_u(W) \leq 1$ and IRRA for $u(W)$, while little or no support of this type exists for the opposite assumptions.

Direct empirical evidence concerning relative risk aversion for wealth is provided in an early and often cited empirical investigation by Friend and Blume (1975) (F-B), and this is a study chosen for detailed review. Friend and Blume provide a very detailed description of their methodology and the variables employed. In their analysis, data concerning the composition of the asset portfolios of 2100 households are used to determine how allocation to risky and riskless assets varies with household wealth. The proportion invested in risky assets, and how this proportion changes with wealth, is interpreted as providing cross section evidence concerning the magnitude, and the increasing or decreasing nature, of relative risk aversion for utility for wealth for a representative decision maker.

When carrying out their empirical analysis, F-B present several different estimates of relative risk aversion measures, each estimate associated with a different measure of wealth.

Three of these estimates and their associated measures of wealth are discussed here. Two different decisions by F-B lead to the three wealth outcome variables, each of which differs from A-P wealth. First, they assume that households recognize the taxation of income from investing, and therefore F-B estimate relative risk aversion for expected utility from wealth after taxes are paid on investment income. Second, F-B sometimes choose to include, and other times exclude, the values of housing and human capital when measuring wealth. Determining the effect of these decisions on the measure of relative risk aversion is an important consideration when interpreting their findings.

Since F-B's explicit modeling of the taxation of investment income affects each of their estimates in the same way, this issue is discussed first. When an asset earns a rate of return and that rate of return is taxed, after tax income is likely the variable of main concern to the investor. F-B explicitly recognize this point and alter the simple portfolio model to include this feature. Phrasing this using the terminology and notation from section 1, F-B examine the risk taking properties of utility for outcome variable W_t, $v(W_t)$, where W_t is an after tax measure of wealth. How the risk aversion properties for this utility function relate to those for $u(W)$ is determined once the relationship between W_t and W is identified.

When all investment income is taxed at a constant rate t, after tax wealth to the decision maker is given by

$$W_t = W_0 + W_0(1-t)(\alpha \cdot r + (1-\alpha)\rho) \tag{7}$$

In this expression, W_0 is initial wealth, $0 \leq t < 1$ is the tax rate, r and ρ are the rate of return to the risky and riskless assets, respectively, and α is the proportion of wealth allocated to the risky asset.[6] Arrow-Pratt wealth W is this same expression with $t = 0$. This implies that

$$W_t = t \cdot W_0 + (1-t)W \tag{8}$$

showing that that W and W_t are linearly related. The intercept in this relationship is positive and reflects the fact that it is the income from wealth, not wealth itself, which is taxed. Inserting the intercept and slope values into (6), indicates that the relative risk aversion measures for these two outcome variables satisfy

$$R_u(W) = \frac{W(1-t)}{W_t} R_v(W_t) = \frac{(W_t - tW_0)}{W_t} R_v(W_t). \tag{9}$$

To convert the three estimates of $R_v(W_t)$ reported by F-B to the reference A-P relative risk aversion scale, $R_u(W)$, information concerning the variable adjustment factor $\frac{(W_t - tW_0)}{W_t}$ is needed. It is clear that this adjustment factor is always less than one. Two lines of reasoning argue that this factor is approximately $(1-t)$. First, at $W_t = W_0$, the adjustment factor is exactly $(1-t)$. Second, the adjustment factor is equal to $(1 + r_p - t - r_p \cdot t)/(1 + r_p - r_p \cdot t)$, when $r_p = (\alpha \cdot r + (1-\alpha)\rho)$ denotes the portfolio rate of return. A Taylor series expansion of this shows that for small values for r_p, the adjustment factor is approximately $(1-t)$.

[6]For this altered portfolio model that includes taxes, the relationship between IRRA and DRRA and the proportion invested in the risky asset is no longer that presented by Arrow. Now, it is possible that DRRA is associated with a smaller share of wealth being allocated to the risky asset as wealth increases.

When W_t rather than W is used as the outcome variable, the slope of relative risk aversion is also affected, but for portfolio rate of return values on the order of 10%, this effect can be shown to be quite small and is ignored here. As a consequence, to adjust the F-B estimates to $R_u(W)$, the procedure used here is to simply multiply by $(1-t)$, concluding that to a first approximation, F-B's use of after tax wealth biases their estimates of relative risk aversion upward by $1/(1-t)$.[7]

Average tax rates during the time period of the F-B study are approximately 50% for the high wealth categories, and marginal tax rates are even higher. Thus, the reported relative risk aversion levels for those wealth categories are reduced significantly when converted to an estimate of $R_u(W)$. The inferred reductions for the three wealth measures and for all wealth levels are reported in Table 1 given later in this section. Before presenting this table, however, the three wealth measures themselves are described in more detail.

The first wealth measure used by F-B includes the value of several categories of financial assets and is reported in their Table 1. This measure, denoted here as W_1, consists mainly of assets whose value is quite easily reallocated to other assets. If W_1 completely measures the conceptual W of Arrow and Pratt, then the only adjustment necessary to convert $R_v(W_1)$ to $R_u(W)$ is the tax adjustment just described. It is possible, however, that W_1 fails to include the value of all freely allocable assets, or does include assets which are not freely allocable. F-B suggest that the value of housing, which they purposely exclude from W_1, may be an important asset to consider even though it is more difficult to reallocate than are many of the assets included with W_1. Thus, F-B include housing as a risky asset in their second (and in the third) measure of wealth, denoted here as W_2, and reported in their Table 2.

Since all wealth categories have some housing equity[8] on average, F-B's treatment of housing equity as a risky investment implies that the average proportion of wealth held in risky assets is increased for all investors. This necessarily leads to lower estimates of relative risk aversion. In this portfolio allocation framework, omitting any risky asset leads to an overestimate of relative risk aversion while omission of a riskless asset leads to an underestimate. In addition to reducing relative risk aversion, including the value of the risky housing asset also changes the slope of the relative risk aversion measure because in the data analyzed, housing is a much larger share of wealth for low wealth categories than for those with higher wealth. Hence the reduction in relative risk aversion due to the inclusion of housing equity is most pronounced at low wealth levels. At the highest wealth levels, housing is such a small portion of wealth that the effect of including it is nearly zero. Our Table 1 also gives the F-B estimates of relative risk aversion for utility from W_2 and scales these estimates by $(1-t)$ to adjust for taxes.

In a third measure of wealth, W_3, along with the equity value of housing, F-B also include the value of human capital. They assume, however, that human capital is not a usual risky or riskless asset since its value cannot easily be reallocated to other assets. They use a

[7]One way to think of this is to observe that by imposing taxes on income from investing, the taxing authority shares the risk of investing with the investor. As a consequence the investor is willing to assume more risk, and evidence concerning the magnitude of the risk chosen no longer implies the same level of relative risk aversion that is implied in an investment setting without taxes. Sandmo (1971) makes this same point.

[8]Friend and Blume report and use both the equity and the gross value of housing as measures of housing wealth, but argue that the equity value is the better measure.

Capital Asset Pricing Model approach to take into account the portion of the riskiness of the return on human capital that cannot be diversified away. Thus, the riskiness of human capital impacts the allocation of other wealth among risky and riskless assets, even though human capital itself cannot be reallocated.

One way to represent this simply is to assume that human capital's return is like that of a fixed portfolio of the risky and riskless assets, that is, human capital earns a risky rate of return $r_h = (\beta \cdot r + (1 - \beta)\rho)$, where β is the correlation with return on risky assets, and r_h is the return to human capital. When human capital is included in the portfolio allocation model in this way, its effect on estimated risk aversion depends critically on the relative magnitudes of desired level for α and the given level for β in the fixed human capital portion of the portfolio. When β and α are equal, the inclusion of human capital has no effect on the choice of α because the riskiness of human capital is just what would have been chosen had the decision maker been able to freely reallocate that wealth among the risky and riskless assets. On the other hand, when the chosen α is less than β, the chosen α is lower than it would have been without human capital in order to compensate for the higher than desired risk level associated with human capital. The reverse occurs when the chosen α is larger than β. In this case the riskiness associated with human capital is less than the desired level, so to compensate, the investor chooses a portfolio of financial assets that is riskier than consideration of only the measure of relative risk aversion would imply. The consequence of this is that including human capital with financial assets, reduces the estimate of relative risk aversion when $\beta > \alpha$, and increases the estimate when $\beta < \alpha$.

When $r_h = (\beta \cdot r + (1 - \beta)\rho)$, both the return on human capital and that on the portfolio of financial assets are linear functions of r. This implies that W_3, the measure of wealth including human capital, is a linear transformation of W_2, the measure that excludes it. Straightforward but tedious algebra detail how the slope and intercept of this relationship depends on α and β, and on the relative sizes of financial wealth W_2, and human capital H. The scale factor $\frac{W_2}{(1-\frac{\beta}{\alpha})H+W_2)}$ transforms $R_v(W_3)$ in $R_v(W_2)$. As must be the case, no adjustment is necessary when α and β are equal or if H is zero. The magnitude of the adjustment is greater the larger H is relative to W_2. The scale factor is greater or less than one depending on whether $\beta > \alpha$ or $\beta < \alpha$. F-B report the data needed to calculate this scale factor, and these values along with the tax adjustment, are used in Table 1 to transform the F-B estimate of $R_v(W_3)$ into an estimate of $R_u(W)$, the A-P reference relative risk aversion scale.

In summary, Table 1 reports three of F-B's estimates of relative risk aversion for wealth for a representative decision maker. These estimates are all based on the same data, and differ in part because each uses a different measure of wealth. The unadjusted estimates are quite different from one another, and it is not clear which is the best estimate to use. The analysis here has taken these estimates as they are reported, and using the information provided concerning the various wealth measures, has made adjustments to each of these estimates to provide three alternative estimates of $R_u(W)$, the A-P relative aversion measure. While the adjusted estimates clearly are still not the same as one another, they are more similar than before adjustments were made.

There are several reasons for the remaining variation in the estimates of $R_u(W)$. First, logically it cannot be the case that all three wealth measures include exactly the complete set of assets that can be freely reallocated since the three measures include different assets.

Table 1. Friend and blume estimates of relative risk aversion.

	W_1		W_2		W_3	
Wealth level	$R_v(W_1)$	$R_u(W)$	$R_v(W_2)$	$R_u(W)$	$R_v(W_3)$	$R_u(W)$
1–10	7.02	6.39	2.98	2.76	2.69	3.06
10–100	3.32	3.01	2.79	2.52	2.83	3.74
100–200	2.67	2.28	2.76	2.36	3.27	1.96
200–500	2.62	2.14	2.56	2.15	3.27	2.82
500–1000	2.95	2.13	3.11	2.13	3.78	2.35
1000+	3.08	2.00	3.01	1.98	4.01	2.36

Second, F-B define five wealth categories at the outset and maintain these categories even as the wealth measure changes. Since investor wealth increases with the inclusion of additional assets, a large number of investors change wealth category as the wealth measure is changed. For example, there are 523 households with wealth in the 1–10 thousand range when the measure is W_1, but only 114 remain in this category for W_3, where the values of housing and human capital are included. Because the average of the relative risk aversion levels for these households is used to determine the representative relative risk aversion level for that wealth group, this procedure necessarily averages across a different set of households whenever the wealth measure is altered. All of this prevents the adjustment to a common scale from being exact.

Which of the various estimates of relative risk aversion for the group of investors surveyed is the best in a statistical, procedural and measurement sense is unclear. All estimates of $R_u(W)$ decline more with wealth than do the unadjusted estimates reported by F-B, and even though the adjusted estimated magnitudes are still greater than one, the deviation from one is considerably smaller than suggested by a casual reading of Friend and Blume's analysis.

Blake (1996) conducts a similar study using British data, but with a much more limited number of risky assets. Not including many assets that are included in the F-B study necessarily increases the estimated levels of relative risk aversion. In addition to asset differences, Blake's portfolio model uses rate of return rather than wealth as the outcome variable. Blake does not model or adjust for taxes. The first step in interpreting Blake's findings is to determine the relationship between wealth W and rate of return r. This is given in (7′) which is obtained from Eq. (7) assuming $t = 0$.

$$W = W_0 + W_0(\alpha \cdot r + (1 - \alpha)\rho) \tag{7′}$$

Rewriting this becomes:

$$W = W_0 + W_0(1 - \alpha)\rho + \alpha W_0 r \tag{8}$$

indicating that wealth and rate of return are linearly related. Using (6), this linear relationship implies that the respective relative risk aversion measures are related by $R_v(r) = R_u(W)\frac{\alpha r W_0}{W}$. The adjustment factor $\frac{\alpha r W_0}{W}$ is not a constant, so we evaluate it at the mean value for W, which is $W_0(1 + \bar{r})$. Using the data provided by Blake, his reported relative risk aversion values for $R_v(r)$ can be adjusted to provide information concerning $R_u(W)$.

Adjusting $R_v(r)$, one obtains an estimate for $R_u(W)$ which displays decreasing relative risk aversion, with magnitudes ranging from 16.8 at the lowest levels of wealth to .59 at the highest. The level for the median in the sample, taken as the average for the middle two wealth categories is reported in the summary Table 2 in the concluding section of the paper. Considerable more discussion and analysis is necessary to adequately portray the findings in the Blake paper, and is left to further research.

3. Relative risk aversion for consumption

In this section, a portion of the literature presenting information concerning relative risk aversion for utility from consumption is reviewed. As was the case for wealth, different measures of consumption have been used and this must be taken into account. In addition to adjusting for measurement differences, however, it is also necessary to link the various consumption measures to A-P wealth. Most often, the studies that estimate relative risk aversion for consumption do not themselves provide such a link. Thus, other sources for this information are used, and a brief discussion of a way to relate consumption and A-P wealth is included.

Theory provides less guidance concerning the slope or magnitude of relative risk aversion for consumption than it does for wealth. Arrow's bounded utility argument, which applies to all relative risk aversion measures, indicates a nearness to one in the limit, but other than this, no particular assumption concerning relative risk aversion for consumption, $R_v(C)$, is supported by comparative static analysis. Rothschild and Stiglitz (1971) do point out the importance of both the slope and magnitude of $R_v(C)$ when they show that an increase in the riskiness of the return on saved wealth leads to more consumption now and less later whenever $R_v(C)$ is nondecreasing and less than one, and that the opposite occurs when $R_v(C)$ is nonincreasing and greater than one. It is also the case that the most common assumption made in multi-period consumption models where utility is additive over time is that $R_v(C)$ is constant. This may be due, however, to the convenience of the power function when analyzing such models.

Several papers present a limited amount of empirical evidence concerning relative risk aversion for consumption although in most cases measuring risk aversion is not the primary focus of the research. Some of this work is discussed later in the section. The research that is reviewed in most detail is that by Barsky et al. (1997) (BJKS) who analyze the survey responses of 11,707 individuals to a sequence of hypothetical questions concerning gambles over lifetime income. These responses, and several assumptions concerning risk preferences, allow individuals to be placed in one of four categories defined by the magnitude of their relative risk aversion. Although the survey questions themselves are expressed as gambles over lifetime income, BJKS interpret the relative risk aversion information that is obtained as for permanent consumption. Thus, BJKS determine a range for the magnitude of $R_v(C)$ for each individual where C is interpreted as permanent consumption.

After determining the approximate magnitude of $R_v(C)$ for each individual on the basis of that person's survey responses, BJKS ask how relative risk aversion level varies with a wide range of demographic variables. Unfortunately for the purposes here, permanent consumption or lifetime income is not one of those demographic variables. Instead, two related variables, current wealth and current income, are discussed. BJKS group individuals into five equal size categories defined by either current wealth or current income, and find that average relative risk tolerance, the inverse of relative risk aversion, neither increases nor decreases significantly as the wealth or income level associated with the category changes. Average risk tolerances for the various levels of current wealth (income) range from .2318 to .2601 (.2310 to .2556). Risk tolerance decreases at first, but then increases slightly, as either current wealth or current income increases.

To convert this relatively constant risk tolerance finding into information concerning $R_v(C)$, requires at least two additional assumptions. First, it is assumed that grouping individuals by permanent consumption rather than by current income or current wealth would not affect the findings significantly. The fact that both current income and current wealth lead to similar findings lends some support for this assumption. Second, since only average risk tolerance estimates are reported, these averages must be converted to average relative risk aversion levels. This would be very simple if the inverse of the average were the average of the inverse, but Jensen's inequality tells us that it is just a lower bound. Based on the relationships reported by BJKS in their first table where they report both average risk tolerance and average relative risk aversion grouped by question response, it appears that average relative risk aversion can be as large as double the inverse of average risk tolerance. Thus, our interpretation of the BJKS finding is that $R_v(C)$ is nearly constant with a magnitude of somewhere between four and eight. BJKS do not model the relationship between permanent consumption and A-P wealth so other sources are used for that information.

Ogaki and Zhang (2001) do not provide an estimate for the magnitude of the relative risk aversion for consumption, but do provide information concerning the slope of this measure. Using utility function $u(C_F) = \frac{[(C_F - x)^{1-\alpha} - 1]}{(1-\alpha)}$, whose relative risk aversion measure is $\frac{\alpha C_F}{(C_F - x)}$, they develop a procedure for estimating the value for x. It is the case that when $x = 0$, relative risk aversion is constant, and otherwise is decreasing or increasing depending on whether x is positive or negative, respectively. In their analysis, C_F represents food consumption for households in Pakistan and India. Ogaki and Zhang obtain large positive and significant estimates of the value for x, and therefore conclude that relative risk aversion for food consumption for these individuals is decreasing. To get a sense for the rate of decrease that they find, Ogaki and Zhang's point estimate for x is approximately equal to 2/3 the average level of C_F in their sample. This implies that $R_v(C_F)$ has an elasticity equal to -2 at the average level of food consumption. Ogaki and Zhang do not relate food consumption C_F to A-P wealth in their analysis.

The macroeconomics literature dealing with asset pricing, the equity premium puzzle and the habit formation utility function has also determined approximate values for relative risk aversion for consumption. It does this by requiring risk aversion levels to be consistent with observed asset prices and interest rates. Some such models also directly relate consumption to A-P wealth within an optimization model. A prominent paper in this literature is one by Constantinides (1990), who assumes that the consumer's marginal utility from current consumption depends in part on past consumption levels. A model that leads to this habit level of consumption is formulated.

Utility for consumption takes the form $v(C) = \frac{(C-x)^{1-\alpha}}{(1-\alpha)}$ at each point in time, where x is the habit level of consumption whose value depends on past consumption levels. The multi-period utility from consumption is assumed to be separable and additive over time, allowing utility to be a function of a single argument. The relative risk aversion measure for this utility function $v(C)$ is $\frac{\alpha C}{(C-x)}$, and is identical in form to that of Ogaki and Zhang. Since the proposed values for x are positive, and large relative to C, this form for utility for consumption has a relative risk aversion measure which decreases at a fast rate.

Using annual data for aggregate consumption, and the return on risky and riskless assets for 1889–1978, Constantinides examines various possible parameter values for the process leading to the habit level of consumption, $x(t)$. He reports six pairs of values that are not rejected by the observed data. In his model, the optimal and the habit levels of consumption are random and vary with time, and thus the information provided concerns their mean values. Included in that information is the mean value for $R_u(W)$. The corresponding value for R_v(C), while not reported, can be calculated from the form chosen for $v(C)$ and the parameter values chosen in the calibration. Constantinides does assume that wealth can be freely allocated across assets, and thus the W in his analysis is interpreted as A-P wealth. No other wealth or sources of income are included in his model. In the data used, consumption is aggregate consumption and is treated here as the same as permanent consumption for the representative consumer.

The six parameterizations of the habit formation model that are consistent with the observed consumption and rate of return data imply mean values for $R_v(C)$ ranging from 10.5 to 15.7. The corresponding mean values for $R_u(W)$ range from 7.03 to 2.78. It is the case that $R_v(C)$ is between 1.5 and 6 times larger than $R_u(W)$ depending on which of the six particular forms of the model is considered. These findings are reported in the summary Table 2 identified by the value that Constantinides chooses for the parameter denoted a in the calibration.

Constantinides also reports that the habit level of consumption, x, on average is approximately 80% of consumption at the mode and that this is true for all six of the parameterizations. Using the functional form for $v(C)$, this can be used to calculate the elasticity of $R_v(C)$. This is reported in Table 2 for each of the six parameterizations discussed by Constantinides. These elasticity values range from -3.8 to $-$ 6.1 at the median, implying an even faster rate of decrease than that determined by Ogaki and Zhang.

Consumption and A-P wealth can be related to one another in at least two different ways. One is to specify a model where the optimal level of consumption depends on wealth. The Constantinides' model that was just reviewed, as well as those of Kimball and Mankiw (1989), and Meyer and Meyer (2004) each do this. Although the utility functions and other aspects differ across these models, in each case, the optimal level of consumption is demonstrated to be a linear function of wealth, $C = a+b\cdot W$, with a positive intercept. Partly for this reason, and also for simplicity, it is assumed here that the link between consumption and wealth takes this linear form. This implies that the adjustment to use when converting information concerning $R_v(C)$ to the reference scale involves multiplying $R_v(C)$ by $\frac{bW}{a+bW}$.

To determine this conversion factor, estimates of the relative sizes of the constant, a, and the term bW are needed. In the Constantinides' analysis, the habit level of consumption is approximately 80% at the mode, and $b \cdot W$ represents the remaining 20%. This implies that the conversion factor is .2 at that point. National income account data also indicates that this is a reasonable number. Disposable personal income, for instance, is approximately 80%

from compensation to labor and 20% from rate of return on wealth in the form of return to capital and equity. (Dornbusch and Fisher (1987), Kuznets (1946)). Evidence concerning the marginal propensity to consume from wealth also indicates that the $b \cdot W$ is likely to be quite small relative to a.

Assuming that $C = a + b \cdot W$, and that at the mean, $b \cdot W$ is 20% of consumption implies that $R_u(W) = R_v(C)/5$ at the mean. In addition, using part c) of Theorem 1, at the mean, the elasticity of $R_u(W)$ is a convex combination of the elasticity of $R_v(C)$ and one, with 80% of the weight on one. This implies that the elasticity of $R_u(W)$ is (.2) times that for $R_v(C)$ plus (.8). This is used in Table 2 to convert all estimates of elasticity for relative risk aversion for consumption into an estimate of elasticity of relative risk aversion for wealth. This conversion changes the very negative estimates of Constantinides and Ogaki and Zhang into an elasticity of $R_u(W)$ that is approximately zero. Thus, very severe DRRA for utility for consumption and CRRA for utility for A-P wealth are consistent findings in that analysis.

Applying this same adjustment to the BJKS finding converts values for $R_v(C)$ ranging from 4 to 8 to values for $R_u(W)$ ranging from .8 to 1.6. The BJKS finding that R_v(C) is nearly constant, however, indicates that $R_u(W)$ has a positive elasticity equal to +.8. Thus, their data is quite consistent with the Arrow conjecture that $R_u(W)$ is increasing and has a magnitude near one.

In summary, the link between consumption and A-P wealth suggests that $R_v(C)$ is significantly larger than $R_u(W)$. Moreover, elasticity of $R_u(W)$ is also larger than that for $R_v(C)$, and so much so that the evidence of a strongly declining $R_v(C)$ can be associated with a relatively flat or even increasing $R_u(W)$. This information and that provided in Section 2 is summarized in Table 2.

Table 2. Summary of the reviewed relative risk aversion estimates

	Size of R (Middle)		Elasticity of R (Middle)	
Author/Variable Name	Reported	Adjusted	Reported	Adjusted
Arrow/A-P Wealth	1	1	0 or +	0 or +
F-B/After Tax Wealth W_1	2.65	2.21	−0.02	−0.06
F-B/After Tax Wealth W_2	2.66	2.26	−0.08	−0.09
F-B/After Tax Wealth W_3	3.27	2.39	0	+0.36
Blake/Rate of Return	.23	4.72	−2.46	−0.83
BJKS/Permanent Consumption	4–8	.8–1.6	0	+0.8
Constantinides/Consumption(a = .1)	15.7	7.03	−6.1	−.42
Constantinides /Consumption(a = .2)	12.2	4.09	−4.6	−.12
Constantinides /Consumption(a = .3)	11.6	3.36	−4.3	−.06
Constantinides /Consumption(a = .4)	11.0	3.03	−4.0	0
Constantinides/Consumption(a = .5)	10.5	2.84	−3.8	+.04
Constantinides /Consumption(a = .6)	11.6	2.78	−4.3	−.06
Ogaki and Zhang/Food Consumption.	–	–	−2	+0.4

Table 2 summarizes the majority of the results reported in the paper. Since in many instances the relative risk aversion measure is not a constant, the magnitude reported is a middle level magnitude. For the F-B and Blake studies, there are six wealth categories so an average of the categories three and four is used to compute the magnitude, and the rate of change between these two categories is used to compute the elasticity. Constantinides reports values at the mode so these are used directly. For Arrow and BJKS, the reported magnitude is approximately constant. The wealth elasticity values are obtained as (.2) times the consumption elasticity plus (.8) as described earlier.

4. Summary and conclusions

This review attempts to consolidate and standardize a small portion of the information that is available concerning the risk attitudes of representative decision makers. Information concerning relative risk aversion for different outcome variables and various types of decision makers is used. When appropriately adjusted for differences in outcome variables, the reported relative risk aversion measures are more similar across decision makers and decisions than the unadjusted information indicates.

The data reviewed here suggests that for A-P wealth, a narrowly defined wealth measure, $R_u(W)$ is near, but larger than one and constant or increasing slightly. Evidence for this comes from estimates for relative risk aversion for wealth using various wealth measures, and also from relative risk aversion estimates for consumption.[9] This same evidence implies that for consumption, when broadly defined, R_v(C) is much larger, at least five times as large at the mean. Furthermore, relative risk aversion for consumption decreases rapidly so that values observed away from the mean can vary significantly. For consumption levels below the mean, R_v(C) could be 20, 30 or 50 as some have suggested.

Acknowledgments

The authors thank a referee, the editor and the editorial board for helpful comments. The comments of participants at the 2004 meeting of European Group of Risk and Insurance Economists, Marseille, France, and the Department Seminar at the National University of Ireland, Maynooth are also acknowledged.

References

Arrow, Kenneth J. (1965). *Aspects of the Theory of Risk Bearing*. Helsinki: Yrjö Jahnssonin Säätiö.
Arrow, Kenneth J. (1971). *Essays in the Theory of Risk Bearing*. New York: Markham Publishing.
Barsky, Robert B., F. Thomas Juster, Miles. S. Kimball, and Matthew D. Shapiro. (1997). "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *The Quarterly Journal of Economics* 112 (2), 537–579,

[9]Early drafts of this paper included a section discussing relative risk aversion for profit as estimated for agricultural producers. Because that literature is large, and this paper is already long, that material has been moved to a separate paper.

Blake, David. (1996). "Efficiency, Risk Aversion and Portfolio Insurance: An Analysis of Financial Asset Portfolios Held by Investors in the United Kingdom," *Economic Journal* 106(438), 1175–1192.

Cass, D. and J. E. Stiglitz. (1972). "Risk Aversion and Wealth Effects on Portfolios with Many Assets," *The Review of Economic Studies* 39(3), 331–354.

Cheng, Hsueh-Cheng, Michael J. P. Magill, and Wayne J. Shafer. (1987). "Some Results on Comparative Statics Under Uncertainty," *International Economic Review* 28 (2), 493–507.

Constantinides, George. M.. (1990). "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy* 98 (3), 519–543.

Dornbush, Rudiger and Stanley Fischer. (1987). *Macroeconomics*, 4th edition, New York: McGraw–Hill.

Feldstein, Martin and Elena Ranguelova. (2001). "Individual Risk in an Investment-Based Social Security System," *American Economic Review* 91 (4), 1116–1125.

Friend, Irwin and Marshall E. Blume. (1975). "The Demand for Risky Assets," *American Economic Review* 65 (5), 900–922.

Gollier, Christian. (2002). "Time Diversification, Liquidity Constraints, and Decreasing Aversion to Risk on Wealth," *Journal of Monetary Economics* 49 (7), 1439–1459.

Hadar, Josef and Tae Kun Seo. (1993). "Sensible Risk Aversion," Working Paper, Department of Economics, Southern Methodist University, Dallas, Texas.

Halek, Martin and Joseph G. Eisenhauer. (2001). "Demography of Risk Aversion," *The Journal of Risk and Insurance* 68 (1), 1–24.

Hirshleifer, Jack and John G. Riley. (1992). *The Analytics of Uncertainty and Information*. Cambridge: Cambridge University Press.

Kimball, Miles. S. and N. Gregory Mankiw. (1989). "Precautionary Saving and the Timing of Taxes," *Journal of Political Economy* 97 (4), 863–879.

Kocherlakota, N. R. (1996). "The Equity Premium: Its Still a Puzzle," *Journal of Economic Literature* 34, 42–71.

Kuznets, Simon E. (1946). *National Income: A Summary of Findings*. New York: National Bureau of Economic Research .

Mehra, Rajnish and Edward C. Prescott. (1985). "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15 (2), 145–61.

Menezes, C.F. and D.L. Hanson. (1970). "On the Theory of Risk Aversion," *International Economic Review* 11 (3), 481–487.

Meyer Donald. J. and Jack Meyer. (2004). "Risk Preferences in Multi-Period Consumption Models, The Equity Premium Puzzle, and Habit Formation Utility," Working Paper, Department of Economics, Michigan State University. (Forthcoming in the *Journal of Monetary Economics*).

Ogaki, Masao and Qiang Zhang. (2001). "Decreasing Relative Risk Aversion and Tests of Risk Sharing," *Econometrica* 69 (2), 515–526.

Pratt, John W. (1964). "Risk Aversion in the Small and in the Large," *Econometrica* 32 (12), 122–136.

Rothschild, Michael and Joseph E. Stiglitz. (1971). "Increasing Risk II: Its Consequences," *Journal of Economic Theory* 3 (1), 66–84.

Saha, Atanu and C. Richard Shumway. (1994). "Joint Estimation of Risk Preference Structure and Technology Using Expo–Power Utility," *American Journal of Agricultural Economics* 76 (2), 173–184.

Sandmo, Agnar. (1971). "On the Theory of the Competitive Firm Under Price Uncertainty," *American Economic Review* 61 (1), 65–73.

Zeckhauser, Richard and Emmett Keeler. (1970). "Another Type of Risk Aversion," *Econometrica* 38 (5), 661–665.