Liquidity Constraints and Precautionary Saving

Christopher D. Carroll†
ccarroll@jhu.edu

Miles S. Kimball‡
mkimball@umich.edu

August 30, 2001

Abstract

Economists working with numerical solutions to the optimal consumption/saving problem under uncertainty have long known that there are quantitatively important interactions between liquidity constraints and precautionary saving behavior. This paper provides the analytical basis for those interactions. First, we explain why the introduction of a liquidity constraint increases the precautionary saving motive around levels of wealth where the constraint becomes binding. Second, we provide a rigorous basis for the oft-noted similarity between the effects of introducing uncertainty and introducing constraints, by showing that in both cases the effects spring from the concavity in the consumption function which either uncertainty or constraints can induce. We further show that consumption function concavity, once created, propagates back to consumption functions in prior periods. Finally, our most surprising result is that the introduction of additional constraints beyond the first one, or the introduction of additional risks beyond a first risk, can actually reduce the precautionary saving motive, because the new constraint or risk can 'hide' the effects of the preexisting constraints or risks.

Keywords: liquidity constraints, consumption function, uncertainty, stochastic income, precautionary saving

JEL Classification Codes: C6, D91, E21

† Department of Economics, The Johns Hopkins University, Baltimore MD 21218-2685.
‡ Department of Economics, The University of Michigan, Ann Arbor, MI 48109-1220.

We are grateful to Mark Huggett for suggesting the proof of lemma 1, which is a substantial improvement over our original proof, to Luigi Pistaferri for an insightful discussion of the paper, to conference participants at the "International Savings and Pensions Conference" in Venice in 1999, and to Misuzu Otsuka for meticulous proofreading and mathematical checking.

1 Introduction

In the past decade, numerical solutions to the optimal consumption/saving problem have become the standard theoretical tool for modelling consumption behavior. Numerical solutions have become popular because analytical solutions are not available for realistic descriptions of utility and uncertainty, nor for the plausible case where consumers face both liquidity constraints and uncertainty.

A drawback to numerical solutions, however, is that often it is difficult to determine why results come out the way they do. A leading example of this problem crops up in the relationship between precautionary saving behavior and liquidity constraints. At least since Zeldes (1984), economists working with numerical solutions have known that liquidity constraints can induce precautionary saving even by consumers with quadratic utility functions that provide no inherent precautionary saving motive. Simulations have also sometimes found that liquidity constraints boost the effect of risk on saving even when the utility function already induces a precautionary saving motive.[1] On the other hand, simulation results have sometimes seemed to suggest that liquidity constraints and precautionary saving are substitutes rather than complements. For example, Samwick (1995) has shown that unconstrained consumers with a precautionary saving motive in a retirement saving model behave in ways qualitatively and quantitatively similar to the behavior of liquidity constrained consumers facing no uncertainty.

This paper provides the theoretical tools needed to make sense of the interactions between liquidity constraints and precautionary saving. These tools provide a rigorous theoretical foundation that can be used to clarify the reasons for the numerical literature's apparently contrasting findings.

For example, one of the paper's simpler points is a proof that when a liquidity constraint is added to the standard consumption problem, the resulting value function exhibits increased prudence around the level of wealth where the constraint becomes binding. (Kimball (1990) defines prudence of the value function and shows that it is the key theoretical requirement to produce precautionary saving.) The essential logic for why a liquidity constraint can induce precautionary saving is relatively straightforward. Constrained agents have less flexibility in responding to shocks because the effects of the shocks cannot be spread out over time; thus risk has a bigger negative effect on expected utility (or value) for constrained agents than for unconstrained agents. The precautionary saving motive is heightened by the desire (in the face of risk) to make such constraints less likely to bind.

At a deeper level, we also show that the effect of a constraint on prudence is an example of a more general theoretical result: Prudence is induced by concavity of the consumption function. Since a constraint causes consumption concavity around the point where the constraint binds, adding a constraint necessarily boosts prudence

[1]For a detailed but nontechnical discussion of simulation results on the relation between liquidity constraints and precautionary saving, see Carroll (2001).

around that point. We show that this concavity-boosts-prudence result holds not just for quadratic utility functions but for any utility function in the Hyperbolic Absolute Risk Aversion (HARA) class (which includes Constant Relative Risk Aversion, Constant Absolute Risk Aversion, and most other commonly used forms).

These results tie in closely with findings in our previous paper, Carroll and Kimball (1996), which shows that within the HARA class, the introduction of uncertainty causes the consumption function to become strictly concave (in the absence of constraints) for all but a few carefully chosen combinations of utility function and uncertainty. Indeed, taken together, the results of the two papers can be seen as establishing rigorously the sense in which precautionary saving and liquidity constraints are very close substitutes.[2] In this paper, in fact, we provide an example of a specific kind of uncertainty that (under CRRA utility, in the limit) induces a consumption function that is pointwise identical to the consumption function that would be induced by the addition of a liquidity constraint.

We further show that, once consumption concavity is induced (either by a constraint or by uncertainty), it propagates back to periods before the period in which the concavity is first created.[3] But in the quadratic utility case the propagation is rather subtle: the prior-period consumption rules are concave (and prudence is higher) at any level of wealth from which it is possible that the constraint will bind, but also possible that it may not bind. Precautionary saving takes place in such circumstances because a bit more saving can reduce the probability that the constraint will bind.

The fact that precautionary saving arises from the *possibility* that constraints might bind may help to explain why such a high percentage of households cite precautionary motives as the most important reason for saving (Kennickell and Lusardi (1999)) even though the fraction of households who report actually having been constrained in the past is relatively low (Jappelli (1990)).

Our final theoretical contribution is to show that the introduction of further liquidity constraints beyond the first one may actually *reduce* precautionary saving by 'hiding' the effects of the preexisting constraint(s); identical logic implies that uncertainty can hide the effects of a constraint, because the consumer may need to save so much for precautionary reasons that the constraint becomes irrelevant. For example, a typical perfect foresight model of retirement consumption for a consumer with Social Security income implies that the legal constraint on borrowing against Social Security benefits will cause the consumer to run assets down to zero, then set consumption equal to income for the remainder of life. Now consider adding the possibility of large medical expenses near the end of life (e.g. nursing home fees). Under reasonable assumptions the consumer may save enough against this risk to render the constraint irrelevant.

[2]See Fernandez-Corugedo (2000) for a related demonstration that 'soft' liquidity constraints bear an even closer resemblance to precautionary behavior.

[3]Our previous paper showed that the concavity induced by uncertainty propagated backwards, but the proofs in that paper cannot be applied to concavity created by a liquidity constraint.

The rest of the paper is structured as follows. To fix notation and ideas, the next section presents a very brief review of the logic of precautionary saving in the standard case (without liquidity constraints). The third section sets out our general theoretical framework. The fourth section shows that concavity of the consumption function heightens prudence. The fifth section shows how concavity, whether induced by constraints or uncertainty, propagates to previous periods. Section 6 shows how the introduction of a constraint creates a precautionary saving motive for consumers with quadratic utility, and how that precautionary motive propagates backwards; it also shows that the introduction of additional liquidity constraints beyond the first constraint does *not* necessarily further increase (and can even reduce) the precautionary motive at any given level of wealth. The next section examines the effects of introducing a constraint when utility is of the CRRA form, and contains our example in which a constraint and uncertainty have identical effects on the consumption function. It uses this example to make the point that introduction of uncertainty can hide the effects of constraints or preexisting uncertainty. The final section concludes.

2 A Brief Review

We begin with a very brief review of the logic of precautionary saving in the two-period case; with minor modifications this two-period model is directly applicable to the multiperiod case when the second period utility function is interpreted as the value function arising from optimal behavior from time $t+1$ on.

Consider a consumer with initial wealth w_t who anticipates uncertain future income y_{t+1}. This consumer solves the unconstrained optimization problem[4]

$$\max_{\{c_t\}} \ u(c_t) + E_t\left[V_{t+1}(w_t - c_t + \tilde{y}_{t+1})\right], \tag{1}$$

or, equivalently,

$$\max_{\{s_t\}} \ u(w_t - s_t) + E_t\left[V_{t+1}(s_t + \tilde{y}_{t+1})\right]. \tag{2}$$

The familiar first-order condition for this problem is to set $u'(c_t) = E_t[V'_{t+1}(w_t - c_t + \tilde{y}_{t+1})]$ or, equivalently, $u'(w_t - s_t) = E_t[V'_{t+1}(s_t + \tilde{y}_{t+1})]$.

Figure 1 shows a standard example of this problem in which both u and V_{t+1} are Constant Relative Risk Aversion (CRRA) utility functions. The consumer is assumed to start period t with amount of wealth w_t. The horizontal axis represents the choice of how much the consumer saves in period t, and the upward-sloping curve labelled

[4]Here and henceforth we use a $\sim$ to designate those variables inside an expectations operator whose value is uncertain as of the date at which the expectation is taken. Hence, since y_{t+1}'s value is uncertain as of time t, it is written as $\tilde{y}_{t+1}$ on the RHS of (1).

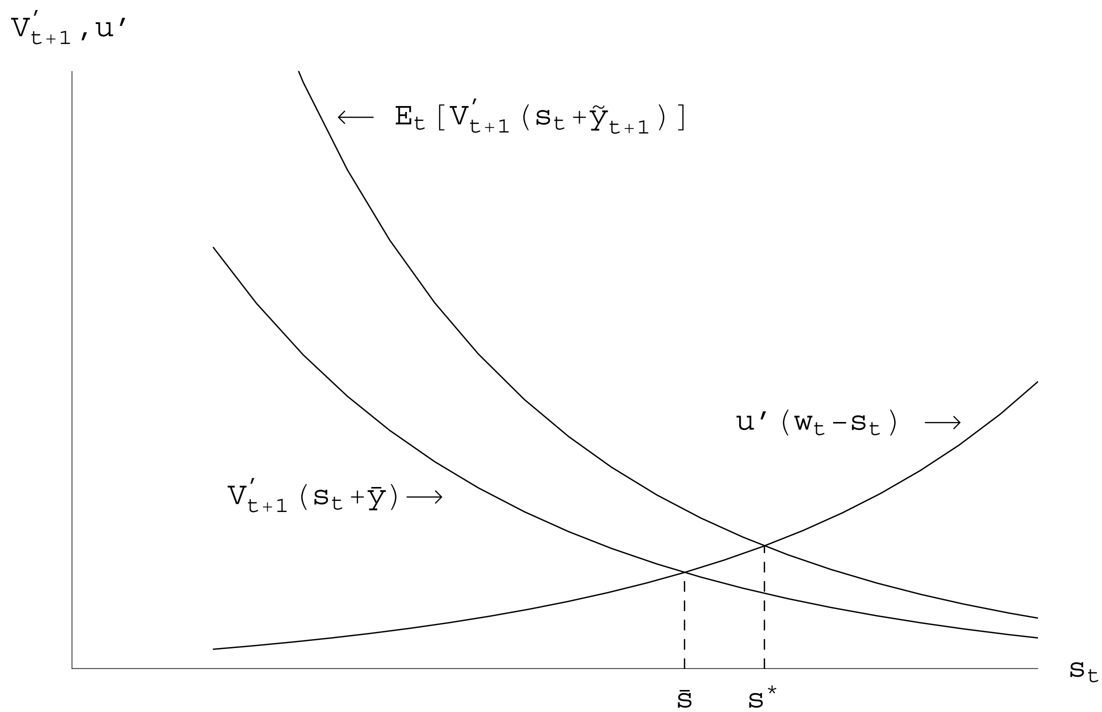

Figure 1: Determining Consumption in the Two Period Case Given Initial Wealth w_t

$u'(w_t - s_t)$ reflects the period-t marginal utility of the consumption $(w_t - s_t)$ associated with that choice of saving. The downward-sloping curve labelled $V'_{t+1}(s_t + \overline{y})$ reflects the marginal value the consumer would experience in period $t+1$ as a function of saving s_t in the previous period if she were perfectly certain to receive income $\overline{y} = E_t[\tilde{y}_{t+1}]$ in period $t + 1$. This curve is downward-sloping as a function of s_t because the more the consumer saves in period t, the more is available for consumption in period $t + 1$ and thus the lower is the marginal utility of spending in $t+1$. In this perfect-certainty case, the utility-maximizing level of consumption is found at the point of intersection between the $u'(w_t - s_t)$ and the $V'_{t+1}(s_t + \overline{y})$ curves, i.e. the level of saving that equalizes the current and future marginal utility of consumption. In the CRRA case where the period-utility functions $u(c)$ and $V_{t+1}(w_{t+1})$ are identical, the optimal solution is to consume exactly half of total lifetime resources in the first period; the point labelled $\bar{s}$ reflects this level of saving.

In the case where period $t + 1$ income is uncertain, first-period marginal utility must be equated to the expectation of the second-period marginal value function. That expectation will be a convex combination of the marginal values associated with each possible outcome, where the weights on each outcome are given by the probability of that outcome. For illustration, suppose there is a 0.5 probability that the consumer will receive income $\overline{y} + \eta$ and a 0.5 probability that she will receive income $\overline{y} - \eta$. Since the

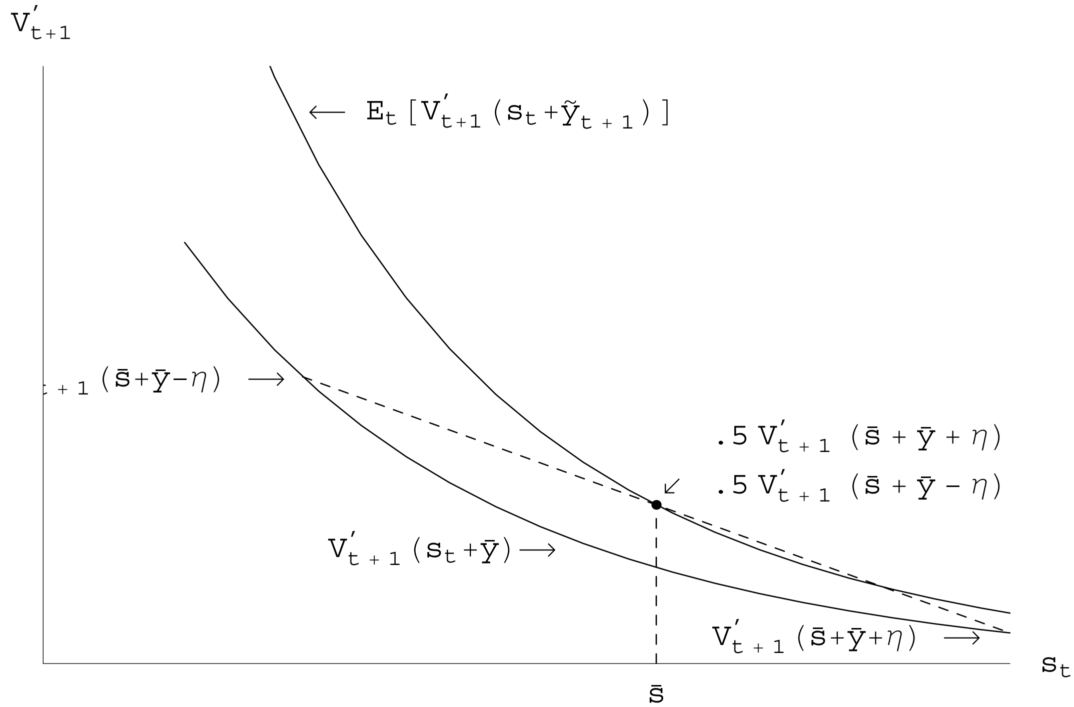

Figure 2: Construction of $E_t[V'_{t+1}]$

probability of each outcome is 1/2, the consumer's expected marginal value function for each s_t will be traced out by the midpoint of the line segment connecting $V'_{t+1}(s_t+\bar{y}+\eta)$ and $V'_{t+1}(s_t+\bar{y}-\eta)$. Figure 2 illustrates the construction of the $E_t[V'_{t+1}(s_t+\tilde{y}_{t+1})]$ curve; for example, if the consumer chooses to save $s_t = \overline{s}$, then her expected marginal value in the second period is given by $.5V'_{t+1}(\bar{s}+\overline{y}+\eta)+.5V'_{t+1}(\bar{s}+\overline{y}-\eta)$, as shown in the figure.

The expected marginal value function traced out by this convex combination of the good and bad outcomes is reproduced and labelled $E_t[V'_{t+1}(s_t+\tilde{y}_{t+1})]$ in figure 1. The optimal level of saving s^* under uncertainty is simply the level of s_t at the intersection of $u'(w_t-s_t)$ and $E_t[V'_{t+1}(s_t+\tilde{y}_{t+1})]$, where the first order condition is satisfied. The magnitude of precautionary saving is the amount by which saving rises from the riskless case ($\bar{s}$) to the risky case (s^*).

Figure 2 illustrates the simple point that the magnitude of precautionary saving is related to the degree of convexity of the marginal value function. Jensen's inequality guarantees that if V'_{t+1} is strictly convex, then $E_t[V'_{t+1}(s_t+\tilde{y}_{t+1})] > V'_{t+1}(s_t+E_t[\tilde{y}_{t+1}])$ and consequently the intersection with $u'(w_t-s_t)$ will occur at a higher value of first-period saving. Clearly, if V'_{t+1} were linear (as is true in the case of quadratic utility in the absence of liquidity constraints), mean-zero risks in period $t+1$ would not affect the expectation of the marginal value function, because the curve generated by the 'convex combination' would lie atop the original marginal value function. Thus, the convexity in the marginal value function creates a precautionary saving motive.

Formally, Kimball (1990) shows that the prudence of the value function (defined as $-V'''(w)/V''(w)$) measures the convexity of the marginal value function at w and therefore the intensity of the precautionary saving motive at that point. To be precise, given two different value functions $V(w)$ and $\hat{V}(w)$, if the absolute prudence of $\hat{V}(w)$ is greater than for $V(w)$—that is, if $-\hat{V}'''(w)/\hat{V}''(w) > -V'''(w)/V''(w)$– then the addition of a risk causes a greater rightward shift of expected $\hat{V}(w)$ than of expected $V(w)$. As figure 2 suggests, a greater rightward shift tends to produce a greater increase in precautionary saving.

To analyze the multiperiod case, we need to be able to characterize the degree of convexity of the marginal value function or the prudence of the value function itself.[5]

3 The Setup

Before stating and proving our main theorems, we need to lay out the basic setup of the consumption/saving problem with many periods. Consider a consumer who faces some future risks but is not subject to any current or future liquidity constraints. Assume that the consumer is maximizing the time-additive present discounted value of utility from consumption $u(c)$. Denoting the (possibly stochastic) gross interest rate and time preference factors as $\tilde{R}_t \in (0, \infty)$ and $\tilde{\beta}_t \in (0, \infty)$, respectively, and labelling consumption c_t, stochastic labor income y_t, and gross wealth (inclusive of period-t labor

[5] In order to use the prudence of the value function to gauge the effect of a risk in labor income at time $t+1$, we implicitly assume that this risk is independent of all the other risks realized in periods beyond $t+1$ that are already built into the shape of V_{t+1}. In other words, the effect of labor income on the value function must work entirely through its effect on wealth at time $t+1$. There are two possible approaches when the realization of y_{t+1} is correlated with future risks, incomes, or rates of return. First, each period could be decomposed into two transitions, one where the information is revealed about the distribution of future incomes, rates of return, etc. and a second where the labor income at time $t+1$ is revealed. The other approach, which, when possible, is more powerful, is to capitalize all the future effects of a shock into wealth at time $t+1$. This approach is possible when the news revealed is mathematically equivalent to a particular effect on the quantity of an asset in the model.

income) w_t, the consumer's problem can be written as:[6,7]

$$V_t(w_t) = \max_{\{c_t\}} u(c_t) + E_t\left[\sum_{s=t+1}^{T}\left(\prod_{j=t+1}^{s}\tilde{\beta}_j\right)u(\tilde{c}_s)\right] \quad (3)$$
$$s.t. \quad w_{t+1} = R_{t+1}(w_t - c_t) + y_{t+1}.$$

As usual, the recursive nature of the problem makes this equivalent to the Bellman equation:

$$V_t(w_t) = \max_{\{c_t\}} u(c_t) + E_t[\tilde{\beta}_{t+1}V_{t+1}(\tilde{R}_{t+1}(w_t - c_t) + \tilde{y}_{t+1})]. \quad (4)$$

Defining

$$\Omega_t(s_t) = E_t[\tilde{\beta}_{t+1}V_{t+1}(\tilde{R}_{t+1}s_t + \tilde{y}_{t+1})] \quad (5)$$

where $s_t = w_t - c_t$ is the portion of period t resources saved, this becomes[8]

$$V_t(w_t) = \max_{\{c_t\}} u(c_t) + \Omega_t(w_t - c_t). \quad (6)$$

It is also useful to define $\breve{c}_t(\mu_t)$, $\breve{s}_t(\mu_t)$, and $\breve{w}_t(\mu_t)$ as:

$$\breve{c}_t(\mu_t) = u'^{-1}(\mu_t), \quad (7)$$
$$\breve{s}_t(\mu_t) = \Omega_t'^{-1}(\mu_t), \quad (8)$$
$$\breve{w}_t(\mu_t) = V_t'^{-1}(\mu_t). \quad (9)$$

In words, $\breve{c}_t(\mu_t)$ ('c-breve') indicates the level of consumption which yields marginal utility μ_t (note the mnemonic convenience of defining marginal utility as the Greek

[6] We allow for a stochastic discount factor because some problems which contain a stochastic scaling variable (such as permanent income) can be analyzed more easily by dividing the problem through by the scale variable; this division induces a term that effectively plays the role of a stochastic discount factor.

[7] The analysis here is similar in some respects to the analysis in Carroll and Kimball (1996); see that paper for more detailed discussion of the methods used below.

[8] For notational simplicity we express the value function $V_t(w_t)$ and the expected discounted value function $\Omega_t(s_t)$ as functions simply of wealth and savings, but implicitly these functions reflect the entire information set as of time t; if, for example, the income process is not i.i.d., then information on lagged income or income shocks could be important in determining current optimal consumption. In the remainder of the paper the dependence of functions on the entire information set as of time t will be unobtrusively indicated, as here, by the presence of the t subscript. For example, we will call the policy rule in period t which indicates the optimal value of consumption $c_t(w_t)$. In contrast, because we assume that the utility function is the same from period to period, the utility function has no t subscript.

letter spelled mu), $\breve{s}_t(\mu_t)$ indicates the level of end-of-period savings[9] in period t that yields a discounted expected marginal value of μ_t, and $\breve{w}_t(\mu_t)$ indicates the level of beginning-of-period wealth that would yield marginal value of μ_t assuming optimal (though potentially constrained) disposition of that wealth between consumption and saving.[10] In the absence of a liquidity constraint in period t, these definitions imply that for an optimizing consumer whose optimal choice of consumption in period t yields marginal utility μ_t,

$$c_t = \breve{c}_t(\mu_t), \tag{10}$$
$$s_t = \breve{s}_t(\mu_t), \tag{11}$$
$$w_t = \breve{w}_t(\mu_t). \tag{12}$$

In the presence of a liquidity constraint that requires $s_t \geq 0$, equation (11) becomes:

$$s_t = \max[0, \breve{s}_t(\mu_t)]. \tag{13}$$

Note that the budget constraint $w_t = c_t + s_t$ allows us to write:

$$\breve{w}_t(\mu_t) = \breve{c}_t(\mu_t) + \max[0, \breve{s}_t(\mu_t)]. \tag{14}$$

4 Prudence and the Concavity of the Consumption Function

Our ultimate goal is to understand the relationship between liquidity constraints and precautionary saving behavior. As noted above, the magnitude of precautionary saving depends on the absolute prudence of the value function. We begin this section by showing that the absolute prudence of the value function will be greater whenever the consumption function is concave (as opposed to linear); later we will tie constraints to concavity (and therefore to prudence) by showing that the imposition of liquidity constraints concavifies the consumption function.

Our analysis of the concavity of the consumption function is couched in general terms, and therefore applies whether the source of consumption concavity is liquidity constraints or something else. This generality is useful, because there is a compelling candidate for the 'something else': uncertainty. Carroll and Kimball (1996) show that the introduction of uncertainty into an optimization problem without preexisting uncertainty or constraints causes the consumption function to become strictly concave

[9] We use the word 'savings' to indicate the level of wealth remaining in a period after that period's consumption has occurred; 'savings' is therefore a stock variable, and is distinct from 'saving' which is the difference between income and consumption.

[10] We chose the slightly unusual breve accent (˘) because of its rough resemblance to the shape of marginal utility μ, which is the argument for the breve-accented functions.

for most combinations of utility function and uncertainty. Our treatment here will therefore alternate between discussion of the effects of imposing liquidity constraints and the effects of introducing uncertainty. Our treatment thus provides the analytical foundation for the qualitative similarity between the effects of liquidity constraints and of uncertainty that has been known from simulation results since Zeldes (1984).

4.1 When Is the Consumption Function Linear?

Our method in this section will be to compare prudence in a *baseline* case where the consumption function $c_t(w_t)$ is linear to prudence in a *modified* situation in which the consumption function $\hat{c}_t(w_t)$ is concave.

Carroll and Kimball (1996) prove that for utility functions in the HARA class, in the absence of liquidity constraints the consumption function will be linear ($c_t''(w_t) = 0$) only in three cases: when utility is of the Constant Relative Risk Aversion (CRRA) form $u(c) = c^{1-\gamma}/(1-\gamma)$ and the only future risk is multiplicative (i.e. rate-of-return risk);[11] when utility is of the Constant Absolute Risk Aversion (CARA) form $u(c) = -(1/a)e^{-ac}$ and the only future risk is additive (i.e. labor income risk); and when the utility function is quadratic, $u(c) = -(\alpha/2)(c-\kappa)^2$.[12] Thus, the natural baseline cases to consider are the three HARA cases where the consumption function is linear.

4.2 How Does Concavity of the Consumption Function Heighten Prudence?

4.2.1 The CRRA Case

Our first baseline $c_t(w_t)$ will be the linear consumption function that arises under CRRA utility in the absence of labor income risk or constraints.[13] Below (in section 6) we show that imposing a constraint makes the consumption function in the constraint-modified situation $\hat{c}_t(w_t)$ concave. Similarly, Carroll and Kimball (1996) show that the addition of labor income risk renders the risk-modified consumption rule concave. In either case it is possible to show that as wealth approaches infinity the consumption rule in the modified situation approaches the consumption rule in the baseline situation. When the experiment is the imposition of a liquidity constraint,

[11] The consumption function when utility is CRRA with a shifted origin, $u(c) = (c-\kappa)^{1-\gamma}/(1-\gamma)$, is *not* linear when there is multiplicative risk, as can be seen from the first order condition for the penultimate period of life when $\beta = 1$: $(c_{T-1} - \kappa)^{-\gamma} = E_{T-1}[(\tilde{R}_T(w_{T-1} - c_{T-1}) - \kappa)^{-\gamma}]$ implying that $c_{T-1} - \kappa = \{E_{T-1}[(\tilde{R}_T(w_{T-1} - c_{T-1}) - \kappa)^{-\gamma}]\}^{-1/\gamma}$ which has no linear solution for c_{T-1} unless $\kappa = 0$.

[12] See section 4.2.3 for a demonstration that the consumption rule is linear under quadratic utility in the presence of both labor income and rate-of-return risk.

[13] The analysis below goes through even if there is rate-of-return risk in the problem, so long as the rate-of-return risk is not modified when the labor income risk is added.

the reason $\hat{c}_t(w_t)$ approaches $c_t(w_t)$ is that as wealth approaches infinity the constraint becomes irrelevant because the probability that it will ever bind becomes zero. When the treatment is the addition of labor income risk, $\hat{c}_t(w_t)$ approaches $c_t(w_t)$ because as wealth approaches infinity the portion of future consumption that the consumer plans on financing out of the uncertain labor income stream becomes vanishingly small.[14] Formally, we can capture both the liquidity constraint and the precautionary saving cases with the assertion that

$$\lim_{w_t \to \infty} \hat{c}(w_t) - c(w_t) = 0.$$

Theorem 1 *Consider an agent who has a utility function with $u'(c) > 0$, $u''(c) < 0$, $u'''(c) > 0$ and nonincreasing absolute prudence $-u'''(c)/u''(c)$ in two different situations. If optimal consumption in the baseline situation is described by a neoclassical consumption function $c_t(w_t)$ that is linear, while optimal behavior in the modified situation (indicated by a hat) is described by a concave neoclassical consumption function $\hat{c}_t(w_t)$ and if $\lim_{w_t \to +\infty} \hat{c}_t(w_t) - c_t(w_t) = 0$, then at any given level of wealth w_t the value function in the modified situation exhibits greater absolute prudence than the value function in the baseline situation. Prudence in the modified situation is strictly greater at w_t than in the baseline situation if and only if the consumption function is strictly concave at some wealth level at or above w_t.*

Proof. By the envelope theorem, the marginal value of wealth is always equal to the marginal utility of consumption as long as it is possible to spend *current* wealth for *current* consumption. That is,

$$V_t'(w_t) = u'(c_t(w_t)) \tag{15}$$
$$\hat{V}_t'(w_t) = u'(\hat{c}_t(w_t)). \tag{16}$$

Differentiating each of these equations with respect to w_t,[15]

$$V_t''(w_t) = u''(c_t(w_t))c_t'(w_t) \tag{17}$$
$$\hat{V}_t''(w_t) = u''(\hat{c}_t(w_t))\hat{c}_t'(w_t). \tag{18}$$

Taking another derivative can run afoul of possible discontinuity in $\hat{c}_t'(w_t)$, but to establish intuition it is useful to consider first the case where $\hat{c}_t''(w_t)$ exists; we will

[14] Since in the CRRA case the proportionate effect of risk on consumption depends on the *square* of the standard deviation of the risk relative to wealth, as this ratio gets small as wealth approaches infinity, the absolute size of the effect of the risk in reducing consumption approaches zero.

[15] Since $\hat{c}(w_t)$ is concave, it has left-hand and right-hand derivatives at every point, though the left-hand and right-hand derivatives may not be equal. Equation (18) should be interpreted accordingly as applying to left-hand and right-hand derivatives separately. (Reading (18) in this way implies that $\hat{c}_t'(w_t^-) > \hat{c}_t'(w_t^+)$; therefore $V''(w_t^-) < V''(w_t^+)$).

then adapt the proof for the case where $\hat{c}''_t(w_t)$ does not exist. For the baseline linear consumption function,

$$V'''_t(w_t) = u'''(c_t(w_t))[c'_t(w_t)]^2 + u''(c_t(w_t))[c''_t(w_t)] \tag{19}$$
$$= u'''(c_t(w_t))[c'_t(w_t)]^2, \tag{20}$$

where the second line follows because with a linear consumption function $c''_t(w_t) = 0$. Thus,

$$\text{Absolute Prudence} = -\frac{V'''_t(w_t)}{V''_t(w_t)} = \left(\frac{-u'''(c_t(w_t))}{u''(c_t(w_t))}\right) c'_t(w_t).$$

In the modified situation with a concave consumption function, where $\hat{c}''_t(w_t)$ exists,

$$\hat{V}'''_t(w_t) = u'''(\hat{c}_t(w_t))[\hat{c}'_t(w_t)]^2 + u''(\hat{c}_t(w_t))[\hat{c}''_t(w_t)] \tag{21}$$
$$-\frac{\hat{V}'''_t(w_t)}{\hat{V}''_t(w_t)} = -\left(\frac{u'''(\hat{c}_t(w_t))[\hat{c}'_t(w_t)]^2 + u''(\hat{c}_t(w_t))[\hat{c}''_t(w_t)]}{u''(\hat{c}_t(w_t))\hat{c}'_t(w_t)}\right) \tag{22}$$
$$-\frac{\hat{V}'''_t(w_t)}{\hat{V}''_t(w_t)} = \left(\frac{-u'''(\hat{c}_t(w_t))}{u''(\hat{c}_t(w_t))}\right) \hat{c}'_t(w_t) - \frac{\hat{c}''_t(w_t)}{\hat{c}'_t(w_t)}. \tag{23}$$

As can be seen from Figure 3,[16] the assumption that the two consumption functions converge asymptotically, $\lim_{w_t \to +\infty} \hat{c}_t(w_t) - c_t(w_t) = 0$, together with the linearity of $c_t(w_t)$ and concavity of $\hat{c}_t(w_t)$, guarantees that the marginal propensity to consume is higher and the level of consumption lower in the modified situation: Thus $\hat{c}'_t(w_t) \geq c'_t(w_t)$ and $\hat{c}_t(w_t) \leq c_t(w_t)$. The inequalities are strict if there is any strictness to the concavity of $\hat{c}_t(\cdot)$ at any level of wealth above w_t.

In conjunction with the assumption of nonincreasing absolute prudence of the utility function, $\hat{c}_t(w_t) \leq c_t(w_t)$ implies that

$$\frac{-u'''(\hat{c}_t(w_t))}{u''(\hat{c}_t(w_t))} \geq \frac{-u'''(c_t(w_t))}{u''(c_t(w_t))}. \tag{24}$$

[16]This figure was generated using simulation programs written for Carroll (2001); these programs are available on Carroll's web page. The parameterization is as follows. The coefficient of relative risk aversion is $\rho = 2$, the time preference factor is $\beta = 0.95$, the gross interest factor is $R = 1.04$, the growth factor for permanent income is $G = 1.01$. The stochastic process for transitory income for $\hat{c}(w)$ involves a small probabilitly (0.005) that income will be zero; if it is not zero, then the transitory shock is lognormally distribuuted with standard deviation of 0.2. Both rules reflect the limit as the number of remaining periods of life approaches infinity.

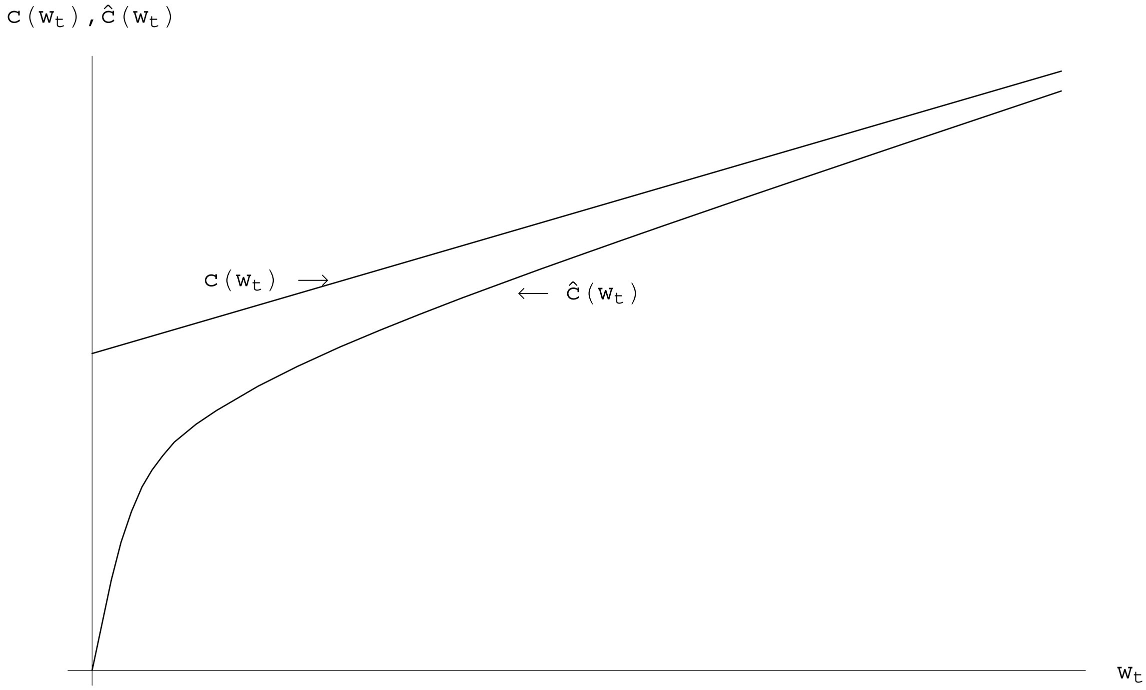

Figure 3: Consumption Functions in the Baseline and Modified Cases

Therefore, where $\hat{c}_t''(w_t)$ exists,

$$-\frac{\hat{V}_t'''(w_t)}{\hat{V}_t''(w_t)} = \left(\frac{-u'''(\hat{c}_t(w_t))}{u''(\hat{c}_t(w_t))}\right)\hat{c}_t'(w_t) - \underbrace{\overbrace{\hat{c}_t''(w_t)}^{\leq 0} / \overbrace{\hat{c}_t'(w_t)}^{>0}}_{\leq 0} \tag{25}$$

$$\geq \left(\frac{-u'''(c_t(w_t))}{u''(c_t(w_t))}\right)c_t'(w_t) \tag{26}$$

$$= -\frac{V_t'''(w_t)}{V_t''(w_t)}. \tag{27}$$

That is, concavity of $\hat{c}_t(w_t)$ along with $\lim_{w_t \to \infty} c_t(w_t) - \hat{c}_t(w_t) = 0$ implies that the absolute prudence of $\hat{V}(w_t)$ is greater than the absolute prudence of $V(w_t)$.

Even when the absolute prudence of the utility function is constant, (26) is strict whenever either (1) $\hat{c}_t(\cdot)$ is strictly concave at some level of wealth above w_t (because, with weak concavity everywhere, strict concavity anywhere above w_t implies that $\hat{c}_t'(w_t) > c_t'(w_t)$); or (2) $\hat{c}_t(\cdot)$ is strictly concave exactly *at* w_t (because strict concavity at w_t implies that $-\frac{\hat{c}_t''(w_t)}{\hat{c}_t'(w_t)} > 0$). Conversely, if $\hat{c}_t(\cdot)$ is linear at w_t and all higher levels of wealth, (26) clearly holds with equality. We can summarize by saying that the inequality (26) which expresses the result of the theorem is *strict* if and only if $\hat{c}_t(\cdot)$ is strictly concave at or above w_t.

What if $\hat{c}_t''(w_t)$ and $\hat{V}_t'''(w_t)$ do not exist? In that case, greater prudence of $\hat{V}$ than V is defined as $\hat{V}'$ being an increasing, convex function of V', or equivalently, $\frac{\hat{V}_t''(w_t)}{V_t''(w_t)}$ being a decreasing function of w_t.[17] By (17) and (18),

$$\frac{\hat{V}_t''(w_t)}{V_t''(w_t)} \equiv \frac{u''(\hat{c}_t(w_t))}{u''(c_t(w_t))}\frac{\hat{c}_t'(w_t)}{c_t'(w_t)}. \tag{28}$$

The second factor, $\frac{\hat{c}_t'(w_t)}{c_t'(w_t)}$, is clearly decreasing (it declines monotonically toward 1). As for the first factor, note that nonexistence of $\hat{V}_t'''(w_t)$ and/or $\hat{c}_t''(w_t)$ do not spring from nonexistence of either $u'''(c)$ or $\hat{c}_t'(w_t)$,[18] so to discover whether $\frac{\hat{V}_t''(w_t)}{V_t''(w_t)}$ is decreasing we can simply differentiate:

$$\frac{d}{dw_t}\left(\frac{u''(\hat{c}_t(w_t))}{u''(c_t(w_t))}\right) = \frac{u'''(\hat{c}_t(w_t))\hat{c}_t'(w_t)u''(c_t(w_t)) - u''(\hat{c}_t(w_t))u'''(c_t(w_t))c_t'(w_t)}{[u''(c_t(w_t))]^2}. \tag{29}$$

Since the denominator is always positive, this will be negative if the numerator is negative, i.e. if

$$u'''(\hat{c}_t(w_t))u''(c_t(w_t))\hat{c}_t'(w_t) \leq u''(\hat{c}_t(w_t))u'''(c_t(w_t))c_t'(w_t) \tag{30}$$

$$\left(\frac{u'''(\hat{c}_t(w_t))}{u''(\hat{c}_t(w_t))}\right)\hat{c}_t'(w_t) \leq \left(\frac{u'''(c_t(w_t))}{u''(c_t(w_t))}\right)c_t'(w_t) \tag{31}$$

$$\underbrace{\left(\frac{-u'''(\hat{c}_t(w_t))}{u''(\hat{c}_t(w_t))}\right)}_{\text{Absolute prudence at } \hat{c}_t(w_t)}\hat{c}_t'(w_t) \geq \underbrace{\left(\frac{-u'''(c_t(w_t))}{u''(c_t(w_t))}\right)}_{\text{Absolute prudence at } c_t(w_t)}c_t'(w_t). \tag{32}$$

Recall that $\hat{c}(w_t) \leq c(w_t)$ (see figure 3), so the assumption of nonincreasing absolute prudence tells us that the absolute prudence term on the LHS of (32) is greater than that on the RHS. But by the assumption of concavity of $\hat{c}_t(w_t)$ we also know that $\hat{c}'(w_t) \geq c'(w_t)$. Hence both terms on the LHS are greater than or equal to the corresponding terms on the RHS.

Thus, combining all of the factors involved in comparing the prudence of $\hat{V}_t(w_t)$ to the prudence of $V_t(w_t)$, we have shown that the value function in the modified situation will exhibit strictly greater prudence at any given w_t than the value function in the baseline situation if and only if $\hat{c}_t(w_t)$ is strictly concave at w_t or at some level of wealth above w_t.

---

[17] To see this, use the implicit function theorem as in Pratt (1964), remembering that $\hat{V}_t'(w_t)$ exists, at least in the sense of right-hand and left-hand derivatives. Note that the theorem at hand is about guaranteeing that concavifying the *argument* of V' on the inside from w_t to $c_t^{-1}(\hat{c}_t(w_t))$ will have the same effect as convexifying V' on the *outside* by some increasing convex function.

[18] It is possible that $\hat{c}_t'(w_t)$ may be discontinuous at specific values of w_t, but in this case the argument below goes through when taking either the right or the left derivatives.

4.2.2 The Exponential Case

The assumption that $\lim_{w_t \to \infty} \hat{c}_t(w_t) - c_t(w_t) = 0$ holds true if consumers have CRRA utility and if the difference between the baseline and the modified situations is the addition of either labor income risk or a liquidity constraint. However, if the consumer's utility function is of the CARA form, a labor income risk simply shifts the entire consumption function down by an equal amount at all levels of w_t, and so the level of consumption in the modified case does not approach the level in the baseline case as wealth approaches infinity. We therefore need a modified version of the theorem to apply in this case.

Corollary 1 *Consider an agent who has a utility function with $u'(c) > 0$, $u''(c) < 0$, $u'''(c) \geq 0$ and nonincreasing absolute prudence $-u'''(c)/u''(c)$ in two different situations. If the baseline situation has a neoclassical consumption function $c_t(w_t)$ that is linear, while the modified situation (indicated by a hat) has a concave neoclassical consumption function $\hat{c}_t(w_t)$ and $\lim_{w_t \to +\infty} \hat{c}'_t(w_t) - c'_t(w_t) = 0$ with $\lim_{w_t \to +\infty} \hat{c}_t(w_t) - c_t(w_t) \leq 0$, then the value function in the modified situation has greater absolute prudence at w_t than does the value function for baseline situation. The inequality of prudence is strict if the modified consumption function is strictly concave at or above w_t.*

The proof of the corollary follows the proof of the main theorem, except where $\lim_{w_t \to +\infty} \hat{c}_t(w_t) - c_t(w_t) = 0$ and concavity of $\hat{c}_t(w_t)$ were used to demonstrate that $\hat{c}'_t(w_t) \geq c'_t(w_t)$ and that $\hat{c}_t(w_t) \leq c_t(w_t)$; here we assume the second, and the first follows from concavity of $\hat{c}_t(w_t)$, linearity of $c_t(w_t)$, the assumption that $\lim_{w_t \to \infty} \hat{c}'(w_t) - c'(w_t) = 0$, and the fact that $\lim \hat{c}_t(w_t) - c_t(w_t) \leq 0$.

4.2.3 The Quadratic Case

The quadratic case requires a somewhat different approach. First, the limit $w_t \to \infty$ is not as meaningful, since it goes beyond the bliss point. Second, since $u'''(\cdot) = 0$, strict inequality between the prudence of $\hat{V}$ and the prudence of V will hold only *at* those points where $\hat{c}_t(\cdot)$ is strictly concave.

To gain intuition for the quadratic problem, consider the Euler equation in the second-to-last period of a lifetime that ends at T, under the assumption that there is no chance that wealth in period T will be greater than the bliss-point level of consumption:[19]

$$u'(c_{T-1}) = E_{T-1}\left[\tilde{\beta}_T\tilde{R}_T u'(\tilde{R}_T(w_{T-1}-c_{T-1})+\tilde{y}_T)\right] \quad (33)$$

$$\alpha(\kappa - c_{T-1}) = E_{T-1}\left\{\tilde{\beta}_T\tilde{R}_T\alpha\left(\kappa-\left[\tilde{R}_T(w_{T-1}-c_{T-1})+\tilde{y}_T\right]\right)\right\} \quad (34)$$

$$c_{T-1} = \left(\frac{E_{T-1}[\tilde{\beta}_T\tilde{R}_T^2 w_{T-1}]+E_{T-1}[\tilde{\beta}_T\tilde{R}_T\tilde{y}_T]+\kappa(1-E_{T-1}[\tilde{\beta}_T\tilde{R}_T])}{1+E_{T-1}[\tilde{\beta}_T\tilde{R}_T^2]}\right) \quad (35)$$

As this equation indicates, in the quadratic case in the absence of liquidity constraints, the solution exhibits certainty equivalence with respect to risks to labor income y_T. An interesting subtlety is that even though the solution is linear in wealth, it does *not* exhibit certainty equivalence with respect to rate-of-return risk, since the level of consumption is related to the expectation of the *square* of the gross return, in a way that implies that an increase in rate-of-return risk increases the marginal propensity to consume. Finally, note that interactions between rate-of-return risk and income risk can cause the consumption function to shift up or down by a potentially large amount.

Recall now from equation (28) that greater prudence of $\hat{V}(w_t)$ occurs if

$$\frac{\hat{V}_t''(w_t)}{V_t''(w_t)} \equiv \frac{u''(\hat{c}_t(w_t))}{u''(c_t(w_t))}\frac{\hat{c}_t'(w_t)}{c_t'(w_t)} \quad (36)$$

$$= \frac{\hat{c}_t'(w_t)}{c_t'(w_t)} \quad (37)$$

is a decreasing function of w_t (the second line follows because for quadratic utility $u''(c)$ is a constant).

Thus, prudence of the value function can be increased in the quadratic case only by something that causes the marginal propensity to consume to decrease as wealth rises. We will show below that in the quadratic case $\hat{c}_t'(w_t)$ experiences a discrete decline at points where a future liquidity constraint potentially becomes binding. Note, however, that an increase in rate-of-return risk, while it increases the level of the MPC compared to the baseline case, does *not* induce a *declining* MPC in wealth: The MPC is higher everywhere, but constant. Thus, rate-of-return risk does not induce an increase in prudence in the quadratic case because $u'''(c) = 0$ for quadratic utility functions (cf. equation (32)).

Corollary 2 *Consider an agent who has a quadratic utility function in two different situations. If the baseline situation has a neoclassical consumption function $c_t(w_t)$ that is linear over some range $w_t < \bar{\omega}$, while the modified situation has consumption function $\hat{c}_t(w_t)$ that exhibits a declining marginal propensity to consume for $w_t < \bar{\omega}$,*

[19] If there is a chance that w_T could exceed the bliss point, then the kink point in the period-T consumption rule can impart concavity to the period-$T-1$ consumption rule.

then prudence of $\hat{V}_t(w_t)$ will be greater than prudence of $V_t(w_t)$ at points where $\hat{c}_t'(w_t)$ declines.

The proof is simply to note that equation (37) is a declining function of w_t only at points where $\hat{c}'(w_t)$ declines with w_t.

5 The Recursive Propagation of Consumption Concavity

The preceding sections make clear the significance of a concave consumption function for the prudence of the value function. Now, we provide conditions guaranteeing that if the consumption function is concave in period $t + 1$, it will be concave in period t and earlier, whatever the source of that concavity may be.

Carroll and Kimball (1996) show that in the absence of liquidity constraints, uncertainty will cause the consumption function to become concave, and that this concavity is propagated to earlier periods. The crucial element in the proof is to show that the value function satisfies the differential inequality

$$V'''(w)V'(w)/[(V''(w))^2] \geq k \tag{38}$$

which holds if the utility function is in the HARA class, which that paper views as a utility function satisfying

$$u'''(c)u'(c)/[(u''(c))^2] = k. \tag{39}$$

The HARA utility functions with positive, nonincreasing absolute prudence satisfy this equation with $k \geq 1$, quadratic utility satisfies it with $k = 0$, while the imprudent HARA utility functions satisfy it with $k < 0$.

For reasons that will become evident, it will be more convenient in this paper to work with an alternative to (39) as our definition of the HARA class; here we view the HARA class as those utility functions with nonnegative, nonincreasing absolute prudence that (after normalization) satisfy either (1) $u'(c) = \kappa - c$, with the domain of c limited to $c < \kappa$ (the quadratic case); (2) $u'(c) = (c - \kappa)^{-\gamma}$ with $\gamma \geq 0$ and the domain of c limited to $c > \kappa$ (the main case); or (3) $u'(c) = e^{-ac}$ with $a > 0$ (the exponential case).

Our goal in this section is to generalize the Carroll and Kimball (1996) results on the propagation of consumption concavity to encompass the case where consumption concavity may arise from the possibility of future liquidity constraints, rather than from the presence of uncertainty. Since (as we show below) constraints can cause V'' to be discontinuous and V''' to fail to exist entirely, the proof strategy of Carroll and Kimball (1996) involving condition (38) will not work. Instead, the central issue in our new proof will involve whether the value function exhibits what we will call "property CC". (The mnemonic is that "CC" stands for "concave consumption".)

Definition 1 *A function $F(x)$ has property CC in relation to a utility function $u(c)$ with $u' > 0$, $u'' < 0$ iff $F'(x) = u'(\psi(x))$ for some monotonically increasing concave function ψ.*

Thus, to say that property CC holds for a value function $V_t(w_t)$ is to say that there exists a concave $\psi(w_t)$ such that

$$V_t'(w_t) \;=\; u'(\psi(w_t)).$$

But the envelope theorem tells us that

$$V_t'(w_t) = u'(c_t(w_t)), \tag{40}$$

so property CC holding for $V_t(w_t)$ is equivalent to having a concave consumption function $\psi(w_t) = c_t(w_t)$.[20]

It is easy to show by taking derivatives that if $V(w)$ satisfies property CC, then when $V'''(w)$ exists this condition reduces to the differential inequality (38), with $k = 0$ in the quadratic case, $k = 1 + (1/\gamma)$ in the main case and $k = 1$ in the exponential case.

5.1 Horizontal Aggregation

First we establish that property CC of the value function is preserved through the process we call "horizontal aggregation," in which the utility from optimal current consumption and the expected utility from optimal saving are aggregated to yield the value function for current wealth.[21]

Lemma 1 *If $\Omega_t(s_t)$ has property CC in relation to u, then $V_t(w_t)$ has property CC in relation to u, whether or not a liquidity constraint holds at the end of period t.*

We begin by showing concavity in the case where there is no liquidity constraint; we will then show that incorporating a current-period constraint does not disturb concavity.

Designate the amount of consumption that would occur in the absence of a current constraint c_t^*. In the unconstrained case, the first order condition for the problem implies that

$$u'(c_t^*) \;=\; \Omega_t'(s_t) \tag{41}$$

$$\;=\; u'(\psi_t(s_t)) \tag{42}$$

[20]Remember that the envelope theorem depends only on being able to spend *current* wealth on *current* consumption, so it holds whether or not there is a liquidity constraint.

[21]We call the intertemporal summing of utility 'horizontal aggregation' because it is easy to visualize as the sum of a series of (expected) marginal values laid out horizontally through time. See Carroll and Kimball (1996) for a more detailed justification of this terminology.

for some increasing concave ψ_t. Taking u'^{-1} of both sides yields

$$\begin{aligned} c_t^* &= \psi_t(w_t - c_t^*) && (43) \\ \psi_t^{-1}(c_t^*) &= w_t - c_t^* && (44) \\ w_t &= \psi_t^{-1}(c_t^*) + c_t^*. && (45) \end{aligned}$$

Since the inverse of an increasing concave function is an increasing convex function,[22] ψ_t^{-1} is an increasing convex function. Since the sum of an increasing linear function c_t^* and an increasing convex function $\psi_t^{-1}(c_t^*)$ is an increasing convex function, $w_t(c_t^*)$ is an increasing convex function. Finally, since the inverse of an increasing convex function is an increasing concave function, $c_t^*(w_t)$ is an increasing concave function.

Thus, in the absence of a period-t liquidity constraint, property CC of $\Omega_t(s_t)$ implies property CC of $V_t(w_t)$.

Note now that when there is a liquidity constraint that requires actual consumption c_t to be less than total resources w_t, actual consumption will be given by the lesser of the unconstrained amount of consumption and the total amount of resources,

$$c_t(w_t) = \min[c_t^*(w_t), w_t].$$

But the min operator applied to two concave functions preserves concavity. Hence even when there is a binding constraint at the end of period t, the consumption rule is concave, implying that $V_t(w_t)$ satisfies property CC.

5.2 Vertical Aggregation

Our next result states that property CC is preserved when expectations are taken.[23]

Lemma 2 *If $V_{t+1}(w_{t+1})$ has property CC and R_{t+1} is always nonnegative, then the function $\Omega_t(s_t)$ defined by equation (5) has property CC.*

Unfortunately, separate proofs are needed for the three classes of HARA utility functions specified above.

We begin by simplifying the problem by assuming that $\beta_{t+1} = R_{t+1} = 1$. This is for expositional clarity only; the steps below all go through with stochastic β_{t+1} and R_{t+1}.

[22]To see this, flip an increasing concave function through the 45 deg line.

[23]We refer to the taking of expectations as 'vertical aggregation' because it is easy to visualize as the vertical stacking and summation of all possible outcomes at a point in time, weighted by their probabilities. Again, see Carroll and Kimball (1996) for a more detailed justification of this terminology.

5.2.1 The Quadratic Case

In the quadratic case, property CC of $V_{t+1}(w_{t+1})$ means that $V'_{t+1}(w_{t+1}) = \kappa - c_{t+1}(w_{t+1})$ for an increasing concave function $c_{t+1}(w_{t+1})$, implying that $V'_{t+1}(w_{t+1})$ is a decreasing convex function. Since $\Omega'_t(s_t)$ is a positive linear combination of decreasing convex functions, it must itself be a decreasing convex function. Property CC of $\Omega'_t(s_t)$ holds if

$$\begin{aligned}\Omega'_t(s_t) &= u'(\phi_t(s_t)) \\ &= \kappa - \phi_t(s_t)\end{aligned}$$

for some concave ϕ_t. But we can simply define $\phi_t(s_t) = \kappa - \Omega'_t(s_t)$ which is clearly an increasing concave function because it is a constant minus a decreasing convex function. Hence $\Omega'_t(s_t)$ satisfies property CC with respect to a quadratic utility function.

5.2.2 The Main Case

In the main case, property CC of $V_{t+1}(w_{t+1})$ means that

$$V'_{t+1}(w_{t+1}) = u'(\psi_{t+1}(w_{t+1})) \quad (46)$$
$$= \psi_{t+1}(w_{t+1})^{-\gamma} \quad (47)$$

for some concave ψ_{t+1}, which under the simplification $R_{t+1} = \beta_{t+1} = 1$ implies that

$$\Omega'_t(s_t) = E_t\left[V'_{t+1}(s_t + \tilde{y}_{t+1})\right] \quad (48)$$
$$= E_t\left[\psi(s_t + \tilde{y}_{t+1})^{-\gamma}\right]. \quad (49)$$

Concavity of ψ implies that

$$\psi(s_t + y_{t+1}) \geq p\psi(s_1 + y_{t+1}) + (1-p)\psi(s_2 + y_{t+1}) \quad (50)$$

for all y_{t+1} if $s_t = ps_1 + (1-p)s_2$ with $p \in [0, 1]$. Since this holds for all y_{t+1}, we know that

$$\left\{E_t\left[\psi(s_t + \tilde{y}_{t+1})^{-\gamma}\right]\right\}^{-1/\gamma} \geq \left\{E_t\left[\{p\psi(s_1 + \tilde{y}_{t+1}) + (1-p)\psi(s_2 + \tilde{y}_{t+1})\}^{-\gamma}\right]\right\}^{-1/\gamma}. \quad (51)$$

Now we need to use Minkowski's inequality, which says that

$$\left\{E_t\left[(\tilde{a}_{t+1} + \tilde{b}_{t+1})^{-\gamma}\right]\right\}^{-1/\gamma} \geq \left\{E_t[\tilde{a}_{t+1}^{-\gamma}]\right\}^{-1/\gamma} + \left\{E_t[\tilde{b}_{t+1}^{-\gamma}]\right\}^{-1/\gamma} \quad (52)$$

for $\gamma > 0$ if a_{t+1} and b_{t+1} are positive.[24]

[24]For a proof, see Hardy, Littlewood, and Polya (1967), page 146, Theorem 198, equation (6.13.2).

Minkowski's inequality implies that

$$\left\{ E_t \left[\{\overbrace{p\psi(s_1+\tilde{y}_{t+1})}^{=\tilde{a}_{t+1}} + \overbrace{(1-p)\psi(s_2+\tilde{y}_{t+1})}^{=\tilde{b}_{t+1}}\}^{-\gamma} \right] \right\}^{-1/\gamma}$$

$$\geq \left\{E_t[\{p\psi(s_1+\tilde{y}_{t+1})\}^{-\gamma}]\right\}^{-1/\gamma} + \left\{E_t[\{(1-p)\psi(s_2+\tilde{y}_{t+1}\}^{-\gamma}]\right\}^{-1/\gamma}$$

$$= p\left\{E_t[\psi(s_1+\tilde{y}_{t+1})^{-\gamma}]\right\}^{-1/\gamma} + (1-p)\left\{E_t[\psi(s_2+\tilde{y}_{t+1})^{-\gamma}]\right\}^{-1/\gamma}$$

$$= p\Omega_t^{'}(s_1)^{-1/\gamma} + (1-p)\Omega_t^{'}(s_2)^{-1/\gamma}. \quad (53)$$

Combining (49) with (51) and (53),

$$\{\Omega_t^{'}(s_t)\}^{-1/\gamma} \quad \geq \quad p\{\Omega_t^{'}(s_1)\}^{-1/\gamma} + (1-p)\{\Omega_t^{'}(s_2)\}^{-1/\gamma}. \quad (54)$$

Thus $\{\Omega_t^{'}(s_t)\}^{-1/\gamma}$ is concave, and $\Omega_t(s_t)$ exhibits property CC.

5.2.3 The Exponential Case

In the exponential case, to show property CC for $\Omega_t(s_t)$ we need

$$\Omega_t^{'}(s_t) \quad = \quad e^{-a\phi_t(s_t)} \quad (55)$$

$$-(1/a)\log\Omega_t^{'}(s_t) \quad = \quad \phi_t(s_t) \quad (56)$$

for concave ϕ_t. Clearly, ϕ_t will be concave if $-\log\Omega_t^{'}(s_t)$ is concave. Again assuming $\beta_{t+1} = R_{t+1} = 1$ for expositional clarity,

$$\log\Omega_t^{'}(s_t) \quad = \quad \log E_t[\exp(-\psi(s_t+\tilde{y}_{t+1}))]. \quad (57)$$

Property CC of V_{t+1} implies concavity of ψ. Thus,

$$-\psi(s_t+y_{t+1}) \quad \leq \quad -p\psi(s_1+y_{t+1}) - (1-p)\psi(s_2+y_{t+1}) \quad (58)$$

$$\log E_t[e^{-\psi(s_t+\tilde{y}_{t+1})}] \quad \leq \quad \log E_t[e^{-p\psi(s_1+\tilde{y}_{t+1})-(1-p)\psi(s_2+\tilde{y}_{t+1})}], \quad (59)$$

where $s_t = ps_1 + (1-p)s_2$.

The arithmetic-geometric mean inequality implies that for positive a and b, if $\bar{a} = E_t[\tilde{a}]$ and $\bar{b} = E_t[\tilde{b}]$, then

$$E_t\left[(\tilde{a}/\bar{a})^p(\tilde{b}/\bar{b})^{1-p}\right] \leq E_t\left[p(\tilde{a}/\bar{a}) + (1-p)(\tilde{b}/\bar{b})\right] = 1, \quad (60)$$

implying that

$$E_t[\tilde{a}^p\tilde{b}^{1-p}] \leq \bar{a}^p\bar{b}^{1-p}. \quad (61)$$

Substituting in $a = e^{-\psi(s_1+y_{t+1})}$ and $b = e^{-\psi(s_2+y_{t+1})}$, this means that

$$E_t[e^{-p\psi(s_1+\tilde{y}_{t+1})-(1-p)\psi(s_2+\tilde{y}_{t+1})}] \leq \left\{E_t[e^{-p\psi(s_1+\tilde{y}_{t+1})}]\right\}^p \left\{E_t[e^{-(1-p)\psi(s_2+\tilde{y}_{t+1})}]\right\}^{1-p} \quad (62)$$

$$\log E_t[e^{-p\psi(s_1+\tilde{y}_{t+1})-(1-p)\psi(s_2+\tilde{y}_{t+1})}] \leq p \log E_t[e^{-p\psi(s_1+\tilde{y}_{t+1})}] + (1-p) \log E_t[e^{-(1-p)\psi(s_2+\tilde{y}_{t+1})}]. \quad (63)$$

Inequalities (59) and (63) together prove convexity of $\log \Omega_t'(s_t)$ and concavity of $-\log \Omega_t'(s_t)$, so that $\Omega_t(s_t)$ satisfies property CC.

5.3 Recursion

Repeated application of Lemma 1 and Lemma 2 implies that if the value function in period t exhibits property CC, then the value functions in all previous periods will also exhibit property CC.

So far we have shown that weak concavity of the consumption function in period $t+1$ will be propagated into previous periods. We now examine how strict concavity is propagated.

5.4 Definition of Strict and Borderline Concavity at a Point

Definition 2 *A function $F(x)$ has property strict CC over the interval between x_1 and $x_2 > x_1$ in relation to a HARA utility function $u(c)$ with nonnegative, nonincreasing prudence iff*

$$F'(x) = u'(\psi(x))$$

for some increasing function $\psi(x)$ which satisfies strict concavity over the interval from x_1 to x_2, defined by

$$\psi(x) > \frac{x_2 - x}{x_2 - x_1}\psi(x_1) + \frac{x - x_1}{x_2 - x_1}\psi(x_2) \quad (64)$$

for all $x \in (x_1, x_2)$.

Definition 3 *A function $F(x)$ has property borderline CC over the interval from x_1 and x_2 if equation (64) holds with equality.*

Definition 4 *A function $F(x)$ has property CC (strict or borderline, respectively) at a point x if there exists a δ such that if $x \in (x_1, x_2)$ and $|x_2 - x_1| < \delta$ then the function exhibits property CC (strict or borderline, respectively) over the interval from x_1 to x_2.*

Intuitively, these definitions are the formal apparatus necessary to handle value functions that have a kink point at which the slope of the marginal value function jumps from one value to another, as will occur (for example) in the transition between levels of wealth where a constraint is not binding and where it is binding.

Note that if a function has property CC globally, then it will have either strict or borderline CC at every point.

In order to understand our approach here it will be useful to step back for a moment to preview the next few steps in the paper. The next section will show that, starting with a setup in which there are no liquidity constraints, the introduction of a first liquidity constraint that binds at the end of period $t+1$ imparts strict concavity to the consumption function at the period-$t+1$ level of wealth $w_{t+1} = \omega^{\#}$ where the constraint begins to bind. What we are constructing in the present discussion is the apparatus to determine how that strict concavity at $V_t(\omega^{\#})$ is propagated back to $\Omega_t(s_t)$, $V_t(w_t)$, and so forth *in the absence of other constraints*. Later in the paper (in section 6.5) we examine what happens when additional constraints are added to the problem (for example, a constraint that might bind at the end of period t).

5.5 Horizontal Aggregation of Strict and Borderline CC

Recall that we defined 'horizontal aggregation' as the propagation of properties from the end-of-period value function $\Omega_t(s_t)$ to the maximized beginning-of-period value function $V_t(w_t)$. We begin with an $\Omega_t(s_t)$ function that has property CC globally. Given this, we can prove the following lemma.

Lemma 3 *If $\Omega_t(s_t)$ exhibits property strict CC at level of saving s_t then $V_t(w_t)$ exhibits property strict CC at the (unique) level of wealth w_t such that optimal consumption at that level of wealth yields $s_t = w_t - c_t(w_t)$.*

If $\Omega_t(s_t)$ exhibits strict CC at a specific point s_t, then for any $s_1 < s_t < s_2$ which are arbitrarily close to s_t we can write

$$\Omega_t'(s_t) \quad = \quad u'(\psi(s_t)) \tag{65}$$

for some monotonically increasing function $\psi(s_t)$ for which

$$\psi(ps_1 + (1-p)s_2) \quad > \quad p\psi(s_1) + (1-p)\psi(s_2) \tag{66}$$

holds for $0 < p < 1$. Now take ψ^{-1} of both sides, yielding

$$ps_1 + (1-p)s_2 \quad < \quad \psi^{-1}(p\psi(s_1) + (1-p)\psi(s_2)). \tag{67}$$

Now use the first order condition from the maximization problem to find the levels of consumption corresponding to s_1 and s_2:[25]

$$
\begin{aligned}
u'(c) &= \Omega_t'(s) && (68)\\
&= u'(\psi(s)) && (69)\\
c &= \psi(s) && (70)\\
\psi^{-1}(c) &= s. && (71)
\end{aligned}
$$

Substituting (70) and (71) into (67) yields

$$
p\,\overbrace{\psi^{-1}(c_1)}^{s_1}+(1-p)\,\overbrace{\psi^{-1}(c_2)}^{s_2} \;<\; \psi^{-1}(pc_1+(1-p)c_2) \tag{72}
$$

which means that ψ^{-1} satisfies the definition of a convex increasing function in a neighborhood from c_1 to c_2 around c_t. But recall our derivation earlier (equation (45)) that

$$
\begin{aligned}
w_t &= \psi^{-1}(c_t)+c_t && (73)\\
\omega_t(c_t) &\equiv \psi^{-1}(c_t)+c_t. && (74)
\end{aligned}
$$

Since $\omega_t(c_t)$ is the sum of the increasing convex function and an increasing linear function, it is itself an increasing convex function, so by the definition of an increasing convex function we have

$$
\begin{aligned}
p\omega_t(c_1)+(1-p)\omega_t(c_2) &> \omega_t(pc_1+(1-p)c_2) && (75)\\
\omega_t^{-1}(pw_1+(1-p)w_2) &< pc_1+(1-p)c_2 && (76)\\
c_t(pw_1+(1-p)w_2) &< pc_t(w_1)+(1-p)c_t(w_2) && (77)
\end{aligned}
$$

where (76) follows from (75) because the inverse of an increasing convex function is an increasing concave function and (77) follows because the definition of ω_t^{-1} implies that it yields the level of consumption that satisfies the first order condition of the maximization problem for the given level of wealth. Thus, $c_t(w_t)$ satisfies the definition of a strictly concave function in the neighborhood of c_t.

Lemma 3 was stated for points s_t at which $\Omega_t(s_t)$ exhibits property strict CC. What about points at which $\Omega_t(s_t)$ exhibits property borderline CC? It turns out that the

[25]This first order condition holds with equality if there are no constraints that apply in the current period. It does *not* hold with equality at every point if there is a constraint in force at the end of the current period, because in that case there will be a level of wealth $\omega^{\#}$ at which the constraint becomes binding and below which all levels of wealth lead to zero savings; hence when there is a constraint at the end of period-t there is not a one-to-one mapping from s_t to a unique corresponding c_t and w_t. As noted above, we defer to later sections discussion of what happens when a such an additional constraint is imposed.

exact same steps can be employed, substituting equality signs for inequalities, to show that if $\Omega_t(s_t)$ exhibits borderline CC at s_t then $V_t(w_t)$ exhibits borderline CC at the level of wealth w_t that optimally leads to savings s_t.

The structure of the argument here is identical to the structure used for demonstrating horizontal aggregation of nonstrict CC in section 5.1; the only complications are the necessity to be careful about the definition of concavity and convexity over intervals in a neighborhood, and the restriction to cases where there is no liquidity constraint at the end of period-t.

5.6 Vertical Aggregation of Strict and Borderline CC

We are interested in whether strict CC applies to $\Omega_t(s_t)$ at point s_t.

Lemma 4 *If from a given value of period-t savings s_t it is possible that a value of period-$t+1$ wealth w_{t+1} could arise at which the period-$t+1$ value function exhibits property strict CC, then $\Omega_t(s_t)$ will exhibit property strict CC at s_t.*

In the quadratic case, $u'(c_{t+1}(w_{t+1})) = V'_{t+1}(w_{t+1})$ is linear in w_{t+1} except around points where the consumption function exhibits strict concavity; around such points strict concavity of $c_{t+1}(w_{t+1})$ implies strict convexity of $u'(c_{t+1}(w_{t+1}))$. Thus $\Omega'_t(s_t)$ is a positive linear combination of functions some of which are linear and some of which are strictly convex at s_t, and since the sum of functions some of which are strictly convex and some of which are linear is strictly convex, $\Omega'_t(s_t)$ exhibits strict CC at s_t.

In the main case, suppose the stochastic income process consists of n possible values of y_{t+1} indexed by i, and write the simplified $\Omega'_t(s_t)$ as

$$\Omega'_t(s_t) \;=\; \sum_{i=1}^{n} p_i \psi(s_t + y_i)^{-\gamma}. \tag{78}$$

Now note that concavity of ψ implies that

$$\sum_{i=1}^{n} p_i \psi(s_t + y_i)^{-\gamma} \;<\; \sum_{i=1}^{n} p_i \left\{(1/2)\left((\psi(s_t + \delta + y_i) + \psi(s_t - \delta + y_i)\right)\right\}^{-\gamma} \tag{79}$$

if $\psi(s_t + y_i)$ is strictly concave for any one of the possible realizations of y. This means that (51) will be a strict inequality if V_{t+1} exhibits property strict CC at any level of w_{t+1} reachable with a positive probability from s_t. But if (51) holds with strict inequality, then the remaining chain of inequalities from (51) through (54) yields strict concavity of $\{\Omega'_t(s_t)\}^{-1/\gamma}$ implying that Ω_t exhibits property strict CC at point s_t.

In the exponential case, (59) will be a strict inequality for s_t if $\psi(w_{t+1})$ is strict at any w_{t+1} reachable from s_t, and again the remaining inequalities lead to the conclusion that $\Omega_t(s_t)$ has property strict CC at s_t.

The foregoing arguments again hold in the case of a general stochastic distribution for y_{t+1}, β_{t+1} and R_{t+1}, with the additional implication that a general distribution for y_{t+1} and stochasticity of R_{t+1} makes more values of w_{t+1} reachable from a given s_t and thus expands the number of values of s_t where $\Omega_t(s_t)$ exhibits strict concavity.

5.7 Recursion for Strict and Borderline Concavity

Recursive application of horizontal and vertical aggregation of strict CC implies that the value function $V_s(w_s)$ will exhibit property strict CC at any value of w_s such that there is any possibility in any future period ($t > s$) of a level of wealth w_t occurring at which $V_t(w_t)$ exhibits property strict CC.

6 Liquidity Constraints and Prudence in the Quadratic Utility Case

Our results thus far have demonstrated that concavity of consumption functions and prudence of value functions, once created, are propagated back through time to previous periods' consumption and marginal value functions. We now turn to the question of how liquidity constraints can *create* strict convexity of the marginal value function.

6.1 Introducing the First Constraint

We begin with the question of how introducing a single liquidity constraint induces precautionary saving when utility is quadratic. The purpose of this section is to show that the introduction of a liquidity constraint that applies between the current period and the next period convexifies the marginal value function for the current period and all prior periods, thus providing the key theoretical requirement for a precautionary saving motive. The essence of the proof will be to show that a liquidity constraint introduces a 'kink' in the marginal value function at the point where the constraint binds, and that the kink will be propagated back to the marginal value functions in all prior periods.[26]

Before proceeding to the proofs, we need a definition.

Definition 5 *Define* $(\omega_t^{\#}, \mu_t^{\#})$ *as the "activation point" of the liquidity constraint under consideration in period* t*, i.e.* $\mu_t^{\#} = \Omega_t'(0)$*, and define two potential values* $\mu_1 = V_t'(\omega_1) < \mu^{\#} < \mu_2 = V_t'(\omega_2)$*.*

[26] For a simple analysis of how liquidity constraints cause a kink in the decision rule and thereby induce precautionary saving in a three period model, see Besley (1995).

That is, μ_1 is the marginal value of wealth at some level of wealth ω_1 above $\omega^{\#}$ where the liquidity constraint is not binding, and μ_2 is the marginal value of wealth at some level of wealth below $\omega^{\#}$, where the constraint *is* binding.

We are now in a position to state the following result:

Lemma 5 *If the period utility function is quadratic and there are no liquidity constraints that could bind after period $t+1$, the imposition of a liquidity constraint in period $t+1$ induces strict convexity of the period $t+1$ marginal value function between any two points ω_1 and ω_2 which lie on opposite sides of the "activation point" $\omega^{\#}$ of the liquidity constraint.*

Proof. Carroll and Kimball (1996) analyze the unconstrained problem for utility functions in the HARA class, i.e. those functions which satisfy the condition $u'''u'/u''^2 = k \geq 0$. Integrating this in the form $\frac{u'''}{u''} = k\frac{u''}{u'}$ yields the equation $u'' = -Au'^k$. But because $\breve{c}'(\mu) = 1/u''(c)$ (differentiate $u'(\breve{c}(\mu))$ with respect to μ),

$$\frac{\breve{c}'(\mu_2)}{\breve{c}'(\mu_1)} = \frac{1/u''(c_2)}{1/u''(c_1)} \tag{80}$$

$$= \frac{u''(c_1)}{u''(c_2)} \tag{81}$$

$$= \frac{u'(c_1)^k A}{u'(c_2)^k A} \tag{82}$$

$$= \left(\frac{\mu_2}{\mu_1}\right)^{-k} \tag{83}$$

or $\breve{c}'(\mu_2) = (\frac{\mu_2}{\mu_1})^{-k}(\breve{c}'(\mu_1))$ (this corresponds to equation (10) in Carroll and Kimball (1996)). The quadratic utility case corresponds to $k = 0$, which implies that

$$\breve{c}'(\mu_2) = \breve{c}'(\mu_1). \tag{84}$$

In the quadratic utility case where there are no future constraints, similar results can be derived for $\breve{s}$ (corresponding to Carroll and Kimball (1996) equation (11)), implying that $\breve{s}'(\mu_2) = \breve{s}'(\mu_1)$. Finally, these can be combined as:

$$\breve{c}'(\mu_2) + \breve{s}'(\mu_2) = \breve{c}'(\mu_1) + \breve{s}'(\mu_1). \tag{85}$$

Now recall that when there is a constraint that binds at the end of the current period, $\breve{w}(\mu) = \breve{c}(\mu) + \max[\breve{s}(\mu), 0]$. The first derivative is $\breve{c}'(\mu) + \chi(\mu)\breve{s}'(\mu)$ where both $\breve{c}'$ and $\breve{s}'$ are negative and χ is a zero-one indicator function for whether $\mu > \mu^{\#}$. But by assumption $\breve{s}(\mu_2) < 0$. As a result we know that $\breve{w}'(\mu_2) = \breve{c}'(\mu_2) > \breve{c}'(\mu_2) + \breve{s}'(\mu_2)$ implying

$$\breve{w}'(\mu_2) = \breve{c}'(\mu_2) > \breve{c}'(\mu_2) + \breve{s}'(\mu_2). \tag{86}$$

Combining this with (85) yields

$$\breve{w}'(\mu_2) > \breve{w}'(\mu_1). \tag{87}$$

Now recall that $\breve{w}(\mu) = V_t^{'-1}(\mu)$ so that $\breve{w}'(\mu) = 1/V''(\omega)$. Therefore, $V''(\omega_1) > V''(\omega_2)$. That is, the (negative) slope of the marginal value function is strictly shallower at levels of wealth above $\omega^{\#}$ than at levels of wealth below $\omega^{\#}$, which is precisely the condition required for strict convexity of the marginal value function in the neighborhood of $\omega^{\#}$. The discrete change in the slope of $V'_{t+1}(w_{t+1})$ at $w_{t+1} = \omega^{\#}$ is the formal definition of the 'kink' in the marginal value function.

Less formally, the logic here is essentially as follows. At any level of wealth below the point $\omega^{\#}$ at which the constraint begins to bind, all incremental wealth is devoted to extra current consumption. The decline in marginal value with extra wealth is exactly as steep as the decline in marginal utility with extra consumption. This is captured by the fact that, below the constraint cutoff, $\breve{c}'(\mu) = \breve{w}'(\mu)$. However, when wealth is above $\omega^{\#}$, an increment to wealth can be spread between the present and the future, and the decline in total marginal value is therefore strictly less than when all of the extra wealth had to be consumed in the present.

6.2 The Effect of the Kink(s)

Now consider the effect of the introduction of liquidity constraints on the response of the expected marginal value function Ω_t to risk. Our primary interest in this paper is in the effects on precautionary saving behavior of the *introduction* of risk to a situation without risk. But analyzing the more general case of increases in risk helps to clarify the theoretical issues, as well as being of interest in its own right.

We need to begin by defining the support of a mean preserving spread in next period's wealth w_{t+1}. To motivate our definition, consider the following example.

Recall the definition of the expected marginal value function,

$$\Omega'_t(s_t) \quad = \quad E_t[\tilde{\beta}_{t+1}\tilde{R}_{t+1}V'_{t+1}(\tilde{w}_{t+1})]. \tag{88}$$

For expositional simplicity, suppose again that the interest rate and time preference factors are nonstochastic and that both are equal to one, $R = \beta = 1$. Suppose further that income shocks can only take the form of a two-point mean-zero risk. That is, if the size of the income shock is ν, then

$$\hat{\Omega}'_t(s_t, \nu) \quad = \quad E_t[V'_{t+1}(s_t + \tilde{y}_{t+1})] \tag{89}$$

$$\quad = \quad 0.5\left\{V'_{t+1}(s_t - \nu) + V'_{t+1}(s_t + \nu)\right\}. \tag{90}$$

We wish to consider the effects on the marginal utility of saving of an increase in the degree of uncertainty about y_{t+1}, which corresponds to an increase in the size of

ν. Suppose that the marginal value function for period $t+1$ takes the piecewise linear form depicted in figure 4, corresponding to the form that we have just shown arises in a quadratic utility problem with a single liquidity constraint that applies in period $t+1$, with an activation point $\omega^{\#}$.

Consider the value of ending period t with a specific amount of savings $s_t = A$ under the initial assumption that income in period $t+1$ is nonstochastically equal to zero, $\nu = 0$. In this case,

$$\begin{aligned} \hat{\Omega}_t^{'}(A,0) &= 0.5\left\{V_{t+1}^{'}(A+0) + V_{t+1}^{'}(A-0)\right\} \quad (91)\\ &= V_{t+1}^{'}(A). \quad (92)\end{aligned}$$

Now consider the effect of increasing the size of the income risk to $\nu = \epsilon$ where $\epsilon > 0$ but $A + \epsilon < \omega^{\#}$. In this case, as the figure illustrates, the addition of the risk has no effect on the expected marginal value of saving A:

$$\begin{aligned} \hat{\Omega}_t^{'}(A,\epsilon) &= 0.5\left\{V_{t+1}^{'}(A-\epsilon) + V_{t+1}^{'}(A+\epsilon)\right\} \quad (93)\\ &= V_{t+1}^{'}(A) \quad (94)\end{aligned}$$

because the expected marginal value function $V_{t+1}^{'}(w_{t+1})$ is linear over the entire range spanned by the possible values of w_{t+1} that arise from saving $s_t = A$.

Now consider the effect of a larger risk $\eta > \epsilon$ where $A+\eta > \omega^{\#}$. It is clear from the figure that increasing the size of the risk from ϵ to η increases the expected marginal value of saving amount A, i.e. $\hat{\Omega}_t^{'}(A,\eta)$ is strictly greater than $\hat{\Omega}_t^{'}(A,\epsilon)$.

Define the level of initial period-t wealth that would have resulted in the optimal amount of period-t savings being A in the absence of income risk as $w^0 = \breve{w}_t(\hat{\Omega}_t^{'}(A,0))$. Now note that since the increase in risk from $\nu = 0$ to $\nu = \epsilon$ has no effect on the marginal utility of savings ($\hat{\Omega}_t^{'}(A,0) = \hat{\Omega}_t^{'}(A,\epsilon)$), the first order conditions of the problem continue to be satisfied after this increase in risk at $c_t = w^0 - A$ and $s_t = A$, so period-t consumption does not change in response to the increase in risk from 0 to ϵ. However, when the risk is increased further to $\nu = \eta$, the marginal utility associated with saving amount A now becomes strictly higher than it was before. This means that the problem's first order conditions for initial wealth w^0 are no longer satisfied at the original level of period-t consumption. In order to satisfy the FOC's, it is necessary to find a new level of consumption that generates a marginal utility partway toward the new higher marginal utility of saving - which is to say, the level of consumption that now satisfies the FOC's will have to be lower. Thus, the increase in risk from ϵ to η induces a (precautionary) decline in the level of consumption in period t.

Note that the critical issue is whether the additional risk 'interacts' with the kink at the activation point. Consider figure 5. In this case the original situation involves an equal chance of w_{t+1} ending up at point A or at point B. Now consider again the

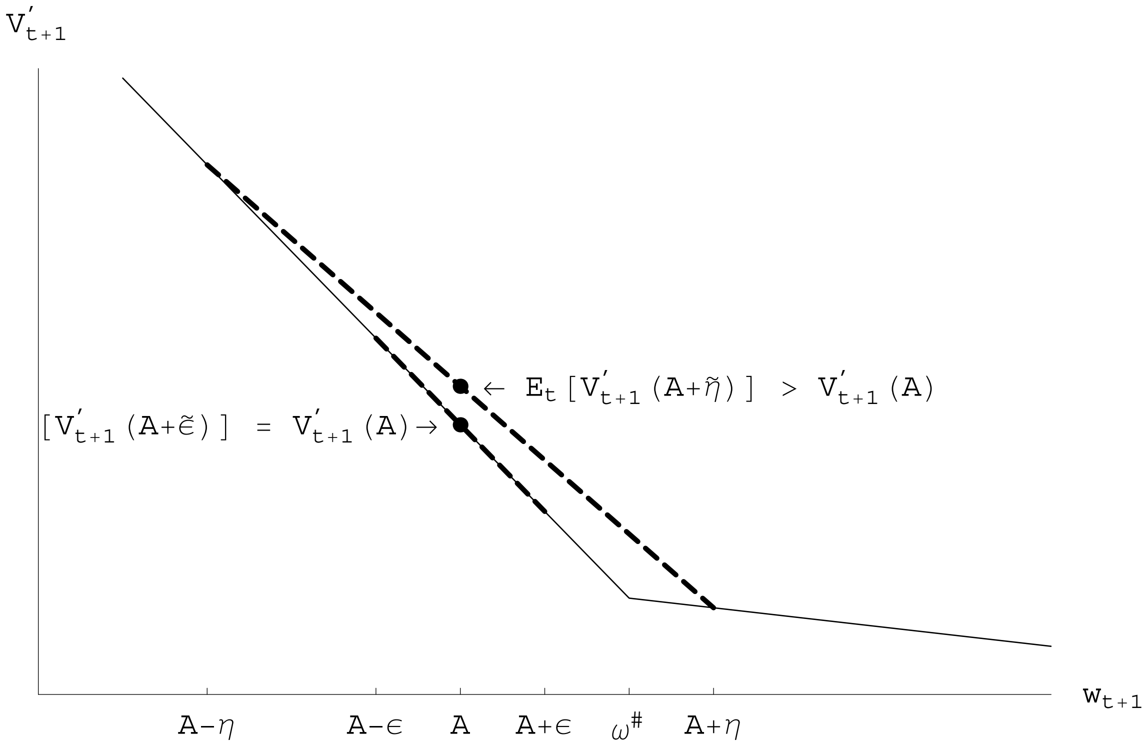

Figure 4: Effect of Adding Mean-Zero Risks of Size ϵ and $\eta > \epsilon$

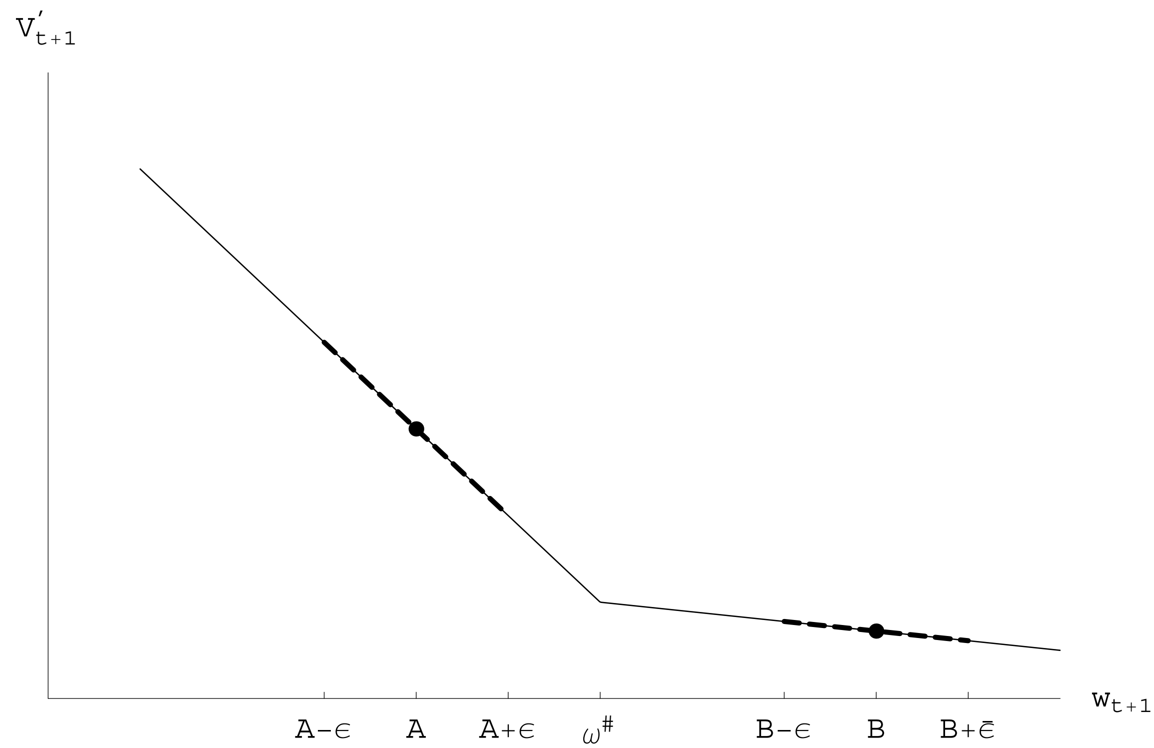

Figure 5: Example of When Adding Risk $\tilde{\epsilon}$ Does Not Induce Prudence

effect of adding a small mean zero risk $\pm\epsilon$ to each of the original outcomes A and B. Although the two points A and B are on opposite sides of the kink, the *change* in the risk does not interact with the kink. The expected marginal value conditional on either A or B does not change, and so the overall expected marginal value does not change.

We now define the formal concept of the support of a mean-preserving spread, which allows us to say when a mean-preserving spread (including the special case of a newly introduced mean-zero risk) interacts with a kink.

Definition 6 *In an interval* $[\underline{w}, \overline{w}]$ *such that* $F_1(\underline{w}) = F_2(\underline{w}) = 0$ *and* $F_1(\overline{w}) = F_2(\overline{w}) = 1$*, let the distribution* F_2 *be a mean preserving spread of* F_1*; that is, if we define* $G_1(w) = \int_{\underline{w}}^{w} F_1(\omega)d\omega$ *and* $G_2(w) = \int_{\underline{w}}^{w} F_2(\omega)d\omega$*, then* $G_2(w) \geq G_1(w)$ *and* $G_2(\overline{w}) = G_1(\overline{w})$*.*

Definition 7 *The open support of the mean preserving spread is the set* $\{w|G_2(w) > G_1(w)\}$*. The support is the closure of the open support.*

In figure 4, the support of the mean preserving spread going from $\nu = 0$ to $\nu = \epsilon$ is the region from $A - \epsilon$ to $A + \epsilon$, and the support of the mean preserving spread going from $\nu = 0$ to $\nu = \eta$ is the region from $A - \eta$ to $A + \eta$. The support of the mean-preserving spread caused by going from $\nu = \epsilon$ to $\nu = \eta$ is the union of the region from $A - \eta$

to $A - \epsilon$ and the region from $A + \epsilon$ to $A + \eta$. In figure 5, the support of the mean preserving spread is the union of the region from $A - \epsilon$ to $A + \epsilon$ and the region from $B - \epsilon$ to $B + \epsilon$.

We are now in position to state the critical lemma.[27]

Lemma 6 *For a given level of saving s_t, let Ψ be the open support of a mean preserving spread in w_{t+1}, and let $\mathcal{W}_{t+1}$ be the set of points at which $V'_{t+1}(w_{t+1})$ exhibits strict CC. Then the expected marginal value of saving $\Omega'(s_t)$ is strictly increased by the mean preserving spread iff $\Psi \cap \mathcal{W} \neq \emptyset$.*

Proof. The lemma is proven using integration by parts. Dropping the $t+1$ subscripts for clarity, the change in the expectation of next period's value function as a result of the mean preserving spread is:

$$\begin{aligned}
&\int_{\underline{w}}^{\overline{w}} V'(w)dF_2(w) - \int_{\underline{w}}^{\overline{w}} V'(w)dF_1(w) \qquad (95)\\
&= V'(\overline{w}) \underbrace{[F_2(\overline{w}) - F_1(\overline{w})]}_{=0} - V'(\underline{w}) \underbrace{[F_2(\underline{w}) - F_1(\underline{w})]}_{=0} \\
&\quad - \int_{\underline{w}}^{\overline{w}} [F_2(w) - F_1(w)]V''(w)dw \\
&= -V''(\overline{w}) \underbrace{[G_2(\overline{w}) - G_1(\overline{w})]}_{=0} \\
&\quad +V''(\underline{w}) \underbrace{[G_2(\underline{w}) - G_1(\underline{w})]}_{=0} \\
&\quad + \int_{\underline{w}}^{\overline{w}} [G_2(w) - G_1(w)]dV''(w) \\
&= \int_{\underline{w}}^{\overline{w}} [G_2(w) - G_1(w)]dV''(w).
\end{aligned}$$

This integral expresses the proposition of the lemma, because the integral will be positive only if there is some set of points at which $G_2(w) > G_1(w)$ and $dV''(w) > 0$. These are the points where the mean preserving spread interacts with the convexity of the marginal value function. Note that the integrals here are well defined even if V'' is discontinuous.[28]

With this lemma in hand, the actual theorem is trivial, simply by focusing on the introduction of a mean-zero risk as a special case of mean preserving spreads.

[27] In the proof below, choose $\overline{w}$ strictly above the suprema of both Ψ and $\mathcal{W}$ and $\underline{w}$ strictly below the infima of both Ψ and $\mathcal{W}$.

[28] Figure 5 gives the essential intuition for how the concept of the support of a mean preserving spread can differ from the convex hull of the support of a risk, in the following sense. For a mean preserving spread, since $G_2(\omega) \geq G_1(\omega)$, when a point ω_0 is *not* in the support of a mean preserving spread,

Theorem 2 *The introduction of a mean-zero risk in period $t+1$ will induce a (precautionary) increase in saving s_t at a given level of wealth w_t even if utility is quadratic, so long as the initially optimal level of saving s_t before introduction of the risk is such that the introduction of the risk leads to a probability $0 < p < 1$ that the liquidity constraint will bind in period $t + 1$.*

Proof: By the previous lemma, the introduction of the risk increases the expected marginal utility of saving, which induces the consumer to save more and consume less.

As we showed in Lemma 5, strict convexity of the marginal value function is generated by transition from a liquidity constraint being binding to non-binding as wealth rises. The introduction of a mean-zero risk raises the expected marginal value of saving in period t if the outcome of that risk affects the probability that the period $t + 1$ liquidity constraint will bind.

To restate in a slightly different way, the lemma shows that, from any level of wealth such that after the risk is introduced the liquidity constraint will bind for bad outcomes of the risk but will not bind for better outcomes, the introduction of the risk increases the marginal utility of saving $\Omega^{'}(s_t)$ by interacting with the convexity of V'_{t+1}. Furthermore, since we are considering the case where utility is quadratic and where there are no liquidity constraints beyond period $t + 1$, the convexity of V'_{t+1} all comes from the activation point, so the introduction of a risk that does not span the activation point does not affect the marginal utility of saving. For the more complex case of a mean-preserving spread as opposed to the introduction of a risk, there is no substitute for the concept of the support of a mean-preserving spread.

The theorem just proven is somewhat backwards: we set out to show that the addition of liquidity constraints induces precautionary saving, but theorem 2 starts out with liquidity constraints and adds a risk. Of course, the bottom line is that when both constraints and risks are present, there will be a precautionary saving motive, but when only risks and no constraints are present quadratic utility implies that there is no precautionary motive, so theorem 2 leads trivially to the theorem that we initially set out to prove.

Theorem 3 *The imposition of a liquidity constraint that binds in period $t + 1$ induces precautionary saving for consumers in period t at all levels of wealth w_t such that at the optimal level of saving s_t the probability that the constraint will bind depends on the outcome of the risk.*

$G_2(\omega_0) = G_1(\omega_0)$, and G_2 is tangent to G_1 at ω_0. The tangency implies either that $F_1(\omega_0) = F_2(\omega_0)$ or that the interval $[F_1(\omega_0^-), F_1(\omega_0^+)]$ is a subset of the interval $[F_1(\omega_0^-), F_2(\omega_0^+)]$. This means that both the distribution F_1 and the distribution F_2 can be sliced into two parts at ω_0 such that the two upper parts have the same mass, the two lower parts have the same mass, the upper part of F_2 is a mean preserving spread of the upper part of F_1 and the lower part of F_2 is a mean preserving spread of the lower part of F_1. Thus, the implication of the theorem is that if you have a kink that is between two separate mean preserving spreads, those mean preserving spreads will not interact with that kink to create an increase in the marginal utility of saving.

The next question we want to address is the extent to which the precautionary saving motive induced in period t propagates back to prior periods, in the absence of further liquidity constraints.

6.3 Main Theorem

Theorem 4 *Introducing a liquidity constraint that applies in period* $t+1$ *to the quadratic utility optimization problem induces a strictly convex marginal value function in period* $s \leq t$ *at any level of wealth* w_s *such that, when the constraint is introduced, there is a probability* $0 < p < 1$ *that the constraint will bind in period* $t + 1$.

Proof. First, note that, by Lemma 5, introducing the liquidity contraint in period $t + 1$ imparts a kink to the marginal value function $V_{t+1}^{'}$ at the point $\omega^{\#}$ where the constraint begins to bind. Define the set $\mathcal{S}_t$ as the set of points s_t such that if the period-t consumer saves s_t there is some probability $0 < p < 1$ that the constraint will bind in period $t + 1$. Then by Lemma 4, $\Omega_t^{'}$ is strictly convex at all points in $\mathcal{S}_t$. Now for each point $s_t \in \mathcal{S}_t$ find the corresponding level of initial wealth such that the level of saving s_t is optimal, i.e. $w_t - c_t(w_t) = s_t$, and call the full set of such points $\mathcal{W}_t$. Then by Lemma 3, $V_t^{'}$ is convex at all points in $\mathcal{W}_t$. Continued iteration using these lemmas demonstrates that for any $s \leq t$, $V_s'(w_s)$ is strictly convex at any value of w_s such that there is both a positive probability that the period $t+1$ constraint will bind, and a positive probability that it will not bind.

Note that if the risks have discrete distributions, with quadratic utility this recursion yields a piecewise linear marginal value function for all $s \leq t$, where the kinks are all associated with the points at which the future liquidity constraint begins to bind. Continuous risk distributions tend to smooth out the kinks.

6.4 Introducing Many Liquidity Constraints and Background Risks All at Once

The theorems above indicate that adding many liquidity constraints and background risks simultaneously will make the consumption function concave. If the consumption function was linear to begin with, by making the consumption function concave, the addition of many liquidity constraints and background risks unambiguously raises the prudence of the value function. In the simplest case of a quadratic utility function, the addition of many liquidity constraints and background risks makes prudence positive-an unambiguous increase when prudence was zero to begin with. This increase in prudence implies that all of the liquidity constraints and background risks as a group make consumption and saving react more strongly to the primary risk. Since interactions, as a matter of logic, go both ways, this also implies that adding liquidity constraints and background risks has a bigger negative effect on consumption if there is a primary risk in place than in the absence of the primary risk.

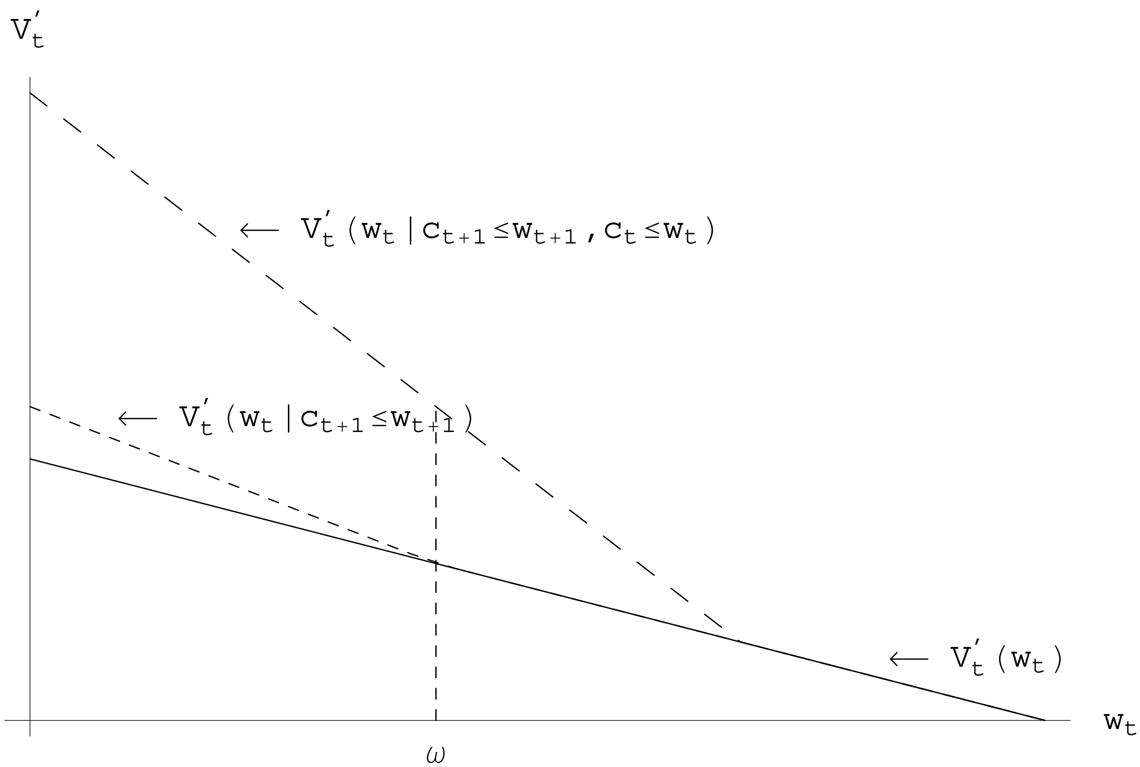

Figure 6: How a Liquidity Constraint Today Can 'Hide' A Future Constraint

6.5 Does Adding Successive Constraints Further Increase Prudence?

A natural next question is whether adding successive further constraints after the first one necessarily intensifies the precautionary saving motive.[29] The surprising answer, in general, is no. The reason is that if a constraint exists in period $t+n$ which convexifies the marginal value function at point $w_t = \omega$, introducing a constraint that applies in period t can 'hide' the kink in the initial period-t value function caused by the period $t+n$ constraint. This point is illustrated by figure 6.

The curve labelled $V_t^{'}(w_t)$ reflects the marginal value function if there are no future liquidity constraints. $V_t^{'}(w_t|c_{t+1} \leq w_{t+1})$ is the marginal value function when there is a constraint in period $t+1$. As we argued above, the constraint in period $t+1$ introduces a kink into the marginal value function in period $t+1$ and all earlier periods. In the figure, the point labelled ω designates the level of w_t at which the constraint in period $t+1$ kinks the period-t marginal value function. $V_t^{'}(w_t|c_{t+1} \leq w_{t+1}, c_t \leq w_t)$ is the

[29]This is, in effect, a question about a triple cross-derivative between the primary risk and two (sets of) liquidity constraints and/or background risks, which helps explain why the question and its answer are more subtle than one might initially guess.

marginal value function when there are constraints in both period t and period $t+1$.

The point of the figure is that when the period-t constraint is not in force but the period-$t+1$ constraint is in force (that is, for $V_t^{'}(w_t|c_{t+1} \leq w_{t+1})$), the fact that the period-t marginal value function is kinked at point ω implies that the marginal value function exhibits prudence at the kink point ω, as argued above. However, when both constraints are in force, the marginal value function (now $V_t^{'}(w_t|c_{t+1} \leq w_{t+1}, c_t \leq w_t)$) is linear in the neighborhood of ω. Thus the imposition of the constraint in period t has the effect of 'hiding' the constraint at $t+1$, and so adding the new constraint *reduces* the prudence of the value function with respect to risks around wealth level ω.

The intuition is as follows. In the absence of the period-t constraint, for levels of wealth $w_t < \omega$, the period-t consumer would borrow enough from period $t+1$ that the period $t+1$ consumer would become constrained (that is why the period-t marginal value function is kinked at ω). Imposing the period-t liquidity constraint prevents period-t consumers with $w_t < \omega$ from borrowing so much and causes such consumers to enter period $t+1$ with enough wealth that the liquidity constraint between $t+1$ and $t+2$ is no longer relevant.

Note a crucial feature of the liquidity constraint that 'hides' the subsequent constraint: the 'hiding' happens for points at which the marginal value function increases as a consequence of the introduction of the new constraint. Since an increase in the marginal value function corresponds to an increase in the value of saving, the new constraint unambiguously increases *total* saving, even though it reduces precautionary saving.

6.6 The Bliss Point and Consumption Concavity in the Quadratic Case

To this point in our analysis of the quadratic utility case, we have implicitly been assuming we are examining behavior at levels of wealth low enough that there is no possibility future wealth will ever be large enough for consumption to equal the bliss point beyond which extra consumption yields negative utility. However, we are now in a position to understand that this is an implausible assumption if there are many periods of life remaining.

The crucial insight here comes from considering the consumption function in the last period of life T. The consumption rule will be

$$c_T(w_T) \;\; = \;\; \min[w_T, \kappa] \tag{96}$$

where κ is the bliss point. But this is obviously an example of a strictly concave consumption function, with concavity at the point $w_T = \kappa$. Thus, $c_{T-1}(w_{T-1})$ will be strictly concave at any level of wealth such that there is a possibility (however remote) that w_T will exceed κ. Similarly for $c_{T-2}(w_{T-2})$, and so forth. So even in the baseline case for quadratic utility, if there is future uncertainty (in either labor income or the rate of return) the consumption function $c_{T-n}(\omega)$ will be strictly concave over an ever-expanding range of values of ω as the number of periods remaining in life n increases.

The upshot is that even the most extreme compromise economists have been willing to make for the sake of tractability (quadratic utility with no liquidity constraints) does not yield the desired payoff of a linear consumption function if there is any substantial amount of uncertainty and there are many periods of life remaining, except for levels of wealth so far below the bliss point that even the most wildly favorable realizations of uncertainty could not result in sufficient wealth ever to permit blisspoint consumption. We hope that this will help extinguish any remaining embers of enthusiasm for the use of quadratic utility functions as a tool for practical economic modelling.

7 Liquidity Constraints and Prudence for CRRA Utility

We now turn to the question of whether adding a first liquidity constraint to a previously unconstrained optimization problem with risky future income globally increases the prudence of the value function for problems where the initial value function already exhibits positive prudence. Once again, the answer is not necessarily. The reason, once again, is that a liquidity constraint can 'hide' certain points on the marginal value function that are exposed if the constraint is not present.

We consider a problem in which a consumer with CRRA utility faces a future income risk but no future liquidity constraints. Note first that the Inada condition of the utility function will necessarily induce an Inada condition in the value function $V_{t+1}(w)$ at some point $\underline{w}$, i.e. $\exists \ \underline{w}$ such that $\lim_{w \downarrow \underline{w}} V'_{t+1}(w) = \infty$.[30] Suppose for simplicity that the time preference and interest factors are equal to one (the result generalizes to the case of stochastic interest and time preference rates considered above).

Consider first consumer A for whom income in period $t+1$ is nonstochastically equal to $\overline{y}$, and who has amount of wealth w_t in period t. Suppose that this consumer faces a liquidity constraint that prevents borrowing against future income. Consumer A's maximization problem is:

$$\max_{\{c_t^A\}} \quad u(c_t^A) + V_{t+1}(s_t^A + \overline{y}) \qquad (97)$$
$$\text{s.t. } s_t^A = w_t - c_t^A \geq 0.$$

Now consider consumer B, a non-liquidity-constrained consumer with the same $u(c_t)$, V_{t+1}, and initial wealth w_t but whose income has a small probability p of going

[30] The argument in this section actually applies to any utility function with an Inada condition; we use CRRA as our example because of its familiarity.

to $\underline{w}$ next period. If this event does not occur, then income will be the same as in the constrained case, $\overline{y}$. This consumer solves

$$\max_{\{c_t^B\}} \quad u(c_t^B) + pV_{t+1}(s_t^B + \underline{w}) + (1-p)V_{t+1}(s_t^B + \overline{y}) \tag{98}$$
$$\text{s.t. } s_t^B = w_t - c_t^B.$$

We wish to show that, as the probability p of the bad shock approaches zero, the behavior of the unconstrained consumer facing the risk becomes arbitrarily close to the behavior of the constrained consumer.

First, consider the case in which the constrained consumer's initial wealth w_t is large enough so that this consumer would, if unconstrained, have saved a positive amount. That is, consider the case where the liquidity constraint does not bind. In this case the first order condition for the first consumer is $u'(w_t - s_t^A) = V_{t+1}'(s_t^A + \overline{y})$. The FOC for consumer B, on the other hand, is $u'(w_t - s_t^B) = pV_{t+1}'(s_t^B + \underline{w}) + (1-p)V_{t+1}'(s_t^B + \overline{y})$. But clearly $\lim_{p\downarrow 0}[pV_{t+1}'(s_t^B + \underline{w}) + (1-p)V_{t+1}'(s_t^B + \overline{y})] = V_{t+1}'(s_t^B + \overline{y})$.

Since saving is determined uniquely by the FOC's, this implies that $\lim_{p\downarrow 0} s_t^B = s_t^A$, i.e. as p approaches zero the saving of the consumer facing the risk becomes arbitrarily close to the saving of the constrained consumer.

Now, consider the case where consumer A would be constrained. By the definition of 'constrained,' this consumer spends her full available resources w_t, and the marginal utility of spending in period t exceeds the marginal utility of having more income in the next period,

$$u'(w_t) > V_{t+1}'(\overline{y}). \tag{99}$$

Note first that if consumer B were to save exactly 0 and then experienced the bad income shock in period $t+1$, consumer B's expected utility would be $-\infty$. Hence saving an amount less than or equal to 0 is ruled out.

What we need to show now is that if consumer B were to choose to save any amount *greater* than 0, say $\delta > 0$, then as p approaches zero there will always come some point at which consumer B could improve her utility by saving less.

Begin by noting that if consumer B saves fixed amount δ (rather than the 0 that A saves), consumer B's marginal utility in period t will be $u'(w_t - \delta)$. But for any fixed δ, $\lim_{p\downarrow 0}[pV_{t+1}'(\delta + \underline{w}) + (1-p)V_{t+1}'(\delta + \overline{y})] = V_{t+1}'(\overline{y} + \delta)$. But we know from equation (99) and from concavity of the utility function that $u'(w_t - \delta) > u'(w_t) > V_{t+1}'(\overline{y})$. Hence we know that as $p \downarrow 0$ there must come a point at which the consumer can improve her total utility by shifting some resources from the future to the present, i.e. by saving less. Since this argument holds for any $\delta > 0$, this argument demonstrates that as p goes to zero there is no positive level of saving which would make the consumer better off. Hence saving goes to zero.

Thus, we have shown that whether the two consumers start with a wealth position at which the constrained consumer would like to save, or start with a wealth position at which the constrained consumer would not like to save, as p goes to zero the level of saving of the unconstrained consumer facing the risk becomes arbitrarily close to the level of saving of the constrained consumer. Hence, in the limit this kind of risk is indistinguishable from a liquidity constraint.

We showed in section 6 that introducing a liquidity constraint introduces a 'kink' in the value function at the point $\omega^{\#}$ where the constraint begins to bind. The arguments in that section are easy to extend to the CRRA case considered here. Recalling that the prudence of the value function is defined as $-\frac{V'''}{V''}$ it is clear that the discrete jump in the value of V'' at the kink point implies infinite prudence exactly at the kink.

Now consider the implications of these arguments. In the limit as $p \downarrow 0$ a future risk with the character described above becomes indistinguishable from a liquidity constraint in the implied consumption function, and therefore in the implied marginal value function $V'(\omega) = u'(c_t(\omega))$. Hence introducing a liquidity constraint in period t when there is a preexisting risk of this kind is essentially indistinguishable from introducing a second liquidity constraint when there is already a preexisting constraint. There is no reason that a point which was a kink before imposing the new liquidity constraint will necessarily remain a kink point after imposing the new constraint. Since the prudence of the value function at the kink point was infinite before the constraint was introduced and may be finite after the constraint is introduced, the introduction of the constraint could reduce the prudence of the value function at the level of wealth corresponding to the kink. This period's constraint can 'hide' the effects of future risk by making the consumer save so much that those future risks are less consequential (from the standpoint of their effects on precautionary saving) than before the liquidity constraint was introduced.

8 Conclusion

The central message of this paper is that the effects of precautionary saving and liquidity constraints are very similar to each other, because both spring from the concavity of the consumption function. The paper provides an explanation the apparently contadictory results that have emerged from simulation studies, which have sometimes seemed to indicate that constraints intensify precautionary saving motives, and sometimes have found constraints and precautionary behavior to be substitutes.

Our results may have important applications even beyond the traditional consumption/saving problem in which the results were derived. The precautionary-saving effect of liquidity constraints may apply in many circumstances where a decision-maker faces the possibility of future liquidity constraints which raise the marginal value of an extra dollar of cash. Thus, firms which are not currently liquidity constrained may engage in precautionary saving if they believe there is some risk that constraints may bind in the future. Governments that worry about whether they will always be able to borrow on international markets may engage in precautionary saving even in periods when they are unconstrained. The logic could even apply to central banks charged with the responsibility of maintaining stable exchange rate regimes; the possibility of a run on the currency might induce ‘precautionary’ holdings of international reserves that are larger than a risk-neutral central bank would hold. Of course, these are all ideas that have appeared, at least informally and sometimes formally, in the relevant literatures. But this paper provides a general logic which can be applied to clarify precisely when and why one should expect such effects to emerge.

References

Besley, Timothy (1995): "Savings, Credit, and Insurance," in *Handbook of Development Economics, Volume III*, ed. by Jere Behrman, and T.N. Srinivasan, chap. 36, pp. 2123–2207. Elsevier.

Carroll, Christopher D. (2001): "A Theory of the Consumption Function, With and Without Liquidity Constraints (Expanded Version)," *NBER Working Paper Number W8387*.

Carroll, Christopher D., and Miles S. Kimball (1996): "On the Concavity of the Consumption Function," *Econometrica*, 64(4), 981–992.

Fernandez-Corugedo, Emilio (2000): "Soft Liquidity Constraints and Precautionary Savings," *Manuscript, Bank of England*.

Hardy, Godfrey Harold, John E. Littlewood, and George Polya (1967): *Inequalities*. Cambridge University Press, second edn.

Jappelli, Tullio (1990): "Who is Credit Constrained in the U.S. Economy?," *Quarterly Journal of Economics*, 105(1), 219–34.

Kennickell, Arthur B., and Annamaria Lusardi (1999): "Assessing the Importance of the Precautionary Saving Motive: Evidence from the 1995 SCF," *Manuscript, Board of Governors of the Federal Reserve System*.

Kimball, M. S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica*, 58, 53–73.

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

Samwick, Andrew A. (1995): "The Limited Offset Between Pension Wealth and Other Private Wealth: Implications of Buffer-Stock Saving," *Manuscript, Department of Economics, Dartmouth College*.

Zeldes, Stephen P. (1984): "Optimal Consumption with Stochastic Income," Ph.D. thesis, Massachusetts Institute of Technology.