

Saving and Growth with Habit Formation

Published in the American Economic Review, June 2000

Christopher D. Carroll
The Johns Hopkins University
ccarroll@jhu.edu

Jody Overland

David N. Weil
Brown University
david_weil@brown.edu

July 18, 2000

JEL D91 E21 O40

Keywords: habit formation, relative consumption, endogenous growth

We are grateful to Xavier Sala-i-Martin, Joseph Gruber, seminar participants at Yale, Wharton, University of Chicago, NYU, Dartmouth, Brown, and NBER, and especially to Harl Ryder for extensive comments and advice. The views expressed in this paper are those of the authors and do not necessarily represent the views of the Board of Governors or the staff of the Federal Reserve System.

The programs that generated all figures in this paper and in our companion paper are available at Carroll's homepage, <http://www.econ.jhu.edu/people/Carroll>.

Abstract

Saving and growth are strongly positively correlated across countries. Recent empirical evidence suggests that this correlation holds largely because high growth leads to high saving, not the other way around. This evidence is difficult to reconcile with standard growth models, since forward-looking consumers with standard utility should save less in a fast-growing economy because they know they will be richer in the future than they are today. We show that if utility depends partly on how consumption compares to a ‘habit stock’ determined by past consumption, an otherwise-standard growth model can imply that increases in growth can cause increased saving. (JEL D91 E21 O40)

1 Introduction

Economists have long known that saving rates and growth rates are positively correlated across countries. Hendrik S. Houthakker (1961, 1965) and Franco Modigliani (1970) presented initial empirical evidence long ago, and many subsequent papers have confirmed the correlation. The recent revival in empirical research on the determinants of economic growth has further reinforced these early findings.

This positive correlation has generally been interpreted as supporting standard growth models in which higher saving results in either temporarily higher growth (in a Solow-style model), or permanently higher growth (in a Rebelo-style endogenous growth model). However, a growing body of evidence suggests that this saving-to-growth causation is not the only factor, and possibly not even the primary factor, responsible for the positive correlation between saving and growth across countries. Instead, a large part of the causation appears to run in the other direction, from growth to saving. This is most evident in the case of the East Asian economies, which had high growth rates long *before* they had exceptionally high saving rates.

Causation running from growth to saving is problematic for standard growth models, in which consumption is determined by a representative agent with intertemporally separable preferences. For plausible parameter values, such models typically imply that higher growth should reduce the saving rate, not increase it. We show, however, that if a standard endogenous growth model is modified to allow for habit formation in consumption, the model can generate growth-to-saving causality that is qualitatively similar to that observed in the data.

The rest of the paper is structured as follows. In section I, we summarize the evidence in support of the empirical claims made above. In section II we present our model. Our key assumption is that people get utility from a comparison of their current level of consumption to the level that they are ‘accustomed to’ in a well-defined sense. Section III examines the dynamics of saving and growth in our habit-formation model and discusses the relationship between the model’s results and the empirical evidence. Section IV concludes.

2 The Empirical Relationship Between Saving and Growth

Ross E. Levine and David Renelt (1992) have shown that the investment rate is virtually the only variable that is robustly correlated with growth in cross-country data, a correlation that has generally been interpreted as indicating causality running from high investment to high growth. The correlation between *saving* and growth, in turn, has been interpreted as reflecting this same causal channel, with the additional linkage that high saving induces high investment for reasons that

are not entirely clear.¹ Thus, in this view, saving causes growth. By contrast, the original literature by Houthakker (1961, 1965) and Modigliani (1970) had precisely the opposite interpretation: Those authors argued that growth caused saving.²

Recent work has attempted to solve the identification problem in a variety of ways. Sebastian Edwards (1995) examines data from a panel of 36 countries over the period 1970-92. Using lagged population growth, openness, political instability, and other lagged variables as instruments, he concludes that the rate of output growth has a significant, positive effect on saving. Barry P. Bosworth's (1993) comprehensive summary of the available evidence on the determinants of saving, investment, and growth concludes that causality from growth to saving is much more robust than that from saving to growth. In a paper summarizing the conclusions from a recent three-year World Bank project on the determinants of saving and growth across the world, Norman V. Loayza et al. (2000) use a variety of instrumental variables techniques in a cross-section of countries to address the identification problem, and in every regression the instrumented growth rate is among the most robustly significant variables explaining the national saving rate. These results hold for OECD and LDC subsamples as well as for the full sample of countries.

Another way to address the identification problem is to look at microeconomic data, because cross-household differences in growth are not associated with the general equilibrium effects that bedevil interpretation of the growth-saving correlation in aggregate data. Christopher D. Carroll and David N. Weil (1994) present evidence from three separate household-level data sets showing that higher labor income growth is associated with a higher saving rate. Angus S. Deaton and Christina H. Paxson (1994) find some similarly supportive evidence for Taiwan. And Matthew D. Shapiro and Joel B. Slemrod (1995) find that consumers who expected faster income growth were *more* likely to save a temporary increase in income.

Perhaps the most compelling evidence, however, comes from the time pattern of the correlations between saving and growth *within* countries. Carroll and Weil (1994) find that, in the fast-growing, high-saving East Asian countries that account for much of the statistical significance of the cross-country growth-saving relationship, the pattern appears to have been one in which increases in growth *preceded* the rise in saving rates. For example, even after Korea was well into its period of rapid growth, a mainstream observer wrote an article asking "Why Do

¹The powerful empirical association between saving and investment was first emphasized by Martin S. Feldstein and Charles Y. Horioka (1980), but no consensus explanation has emerged

²Some recent growth literature has also questioned the wisdom of interpreting the investment-growth correlation as indicating causation running from the former to the latter. Robert E. Hall and Charles I. Jones (1999), for example, argue that most of the cross-sectional variation in output per capita is due to variation in the productivity with which factors are combined, rather than to differences in factor accumulation. Peter J. Klenow and Andres Rodriguez-Clare (1997) go further, arguing that differences in growth cross sectionally are similarly due to differences in the growth rate of productivity, rather than to transitional dynamics of factor accumulation. Under these interpretations, it is the endogeneity of investment rates, rather than growth rates, that is responsible for the correlation between the two.

Koreans Save ‘So Little?’” (Jeffrey G. Williamson (1979)). Similarly, the period of blistering income growth in Japan began in the late 1940s and early 1950s (prewar growth rates had been more moderate), yet Japan did not exhibit a particularly high saving rate until the 1960s and 1970s.³

The same relation also appears to hold when the situation is reversed: countries that experience a slowdown in economic growth generally experience subsequent declines in saving rates. This pattern is evident in the experience of the OECD countries over the past 25 years: In the wake of the productivity growth slowdown that dates from the early 1970s, national saving rates have declined throughout the OECD.

More formally, Carroll and Weil (1994) present Granger-causality tests for 38 countries for which they have good data, and show that increases in growth significantly precede increases in saving. Robert Dekle (1993) presents similar Granger-causality regressions for a group of fast-growing countries and finds that growth positively Granger-causes saving in every country in his sample. A more recent and more comprehensive Granger-causality exercise by Orazio P. Attanasio et. al. (2000) confirms the Carroll-Weil findings with a broader cross-section of countries and using somewhat different methodology, and Dani Rodrik (1999) presents similar findings using a more qualitative methodology.

Of course, the existence of evidence that growth has a positive effect on saving does not mean that the entire positive cross-country correlation between the two variables is due to the growth-to-saving channel. It is perfectly possible that differences in saving rates (due to preferences or policies) will affect growth, and at the same time that differences in growth (due to policies, say, or to the import of new technologies) will affect saving. There may be two structural relationships, and the relationship between the two variables in the data will depend on both. But while the standard Cass-Koopmans representative-agent growth model provides a firm theoretical foundation for why saving should affect growth, the positive effect of growth on saving is more problematic because in the standard permanent income model of consumption embedded in the model, higher expected income growth should lead to less saving, not more.⁴

Overlapping generations growth models provide a potential theoretical channel for growth-to-saving causality that is lacking in representative-agent models. Indeed, Modigliani (1970, 1986) has long argued that in fast-growing economies, young consumers who are in the saving phase of the life cycle will be much richer than old consumers in the dissaving phase, and so the average saving rate of a fast-growing OLG economy will be higher than that of a slow-growing OLG economy (the ‘ag-

³See Carroll and Weil (1994) for the data. Robert G. King and Levine (1994) also provide evidence that capital accumulation alone is neither necessary nor sufficient in the “take-off” to rapid growth.

⁴See Carroll and Weil (1994) for a numerical demonstration in the Cass-Koopmans growth model; below we derive analytical results which apply to both the Cass-Koopmans growth model and the Rebelo *AK* growth model.

gregation effect'). However, James Tobin (1967) showed long ago that Modigliani's argument relies on an assumption that individual consumers living in a fast-growing economy do not expect faster income growth than do individual consumers living in a slow-growing economy. That is, Modigliani assumed that a 50-year-old Taiwanese consumer would not have expected or experienced faster income growth over the past 25 years than a 50-year-old American. In Modigliani's framework, aggregate growth manifests itself in a rapid shift upward from generation to generation in the *level* of a lifetime income profile whose *slope* (i.e. the growth rate of an individual's income at any given age) remains constant. But Carroll and Lawrence H. Summers (1991) present evidence that strongly suggests that the rough empirical fact is that if aggregate productivity growth is one percent higher, then people of every age experience one percent faster income growth - the polar opposite of Modigliani's assumption. Under these circumstances, as Tobin (1967) showed and Carroll and Summers (1991) reconfirmed, the theoretical 'human wealth effect' (in which consumers anticipating fast income growth save less) greatly outweighs Modigliani's 'aggregation effect' and the theory's implication is that the aggregate correlation between saving and growth should be negative.

The 'aggregation effect' also fails on empirical grounds as an explanation of cross-country saving-growth correlations. In a series of recent papers, Deaton and Paxson (1994, 1997, 2000) and Paxson (1995) have shown for a broad set of countries that even if the countervailing 'human wealth effect' were zero, the 'aggregation effect' would not be able to explain the positive cross-country relationship between saving and growth, because the assumption that the young save and the old dissave is a poor approximation to actual empirical age-saving profiles. In fact, as shown in a recent volume edited by James M. Poterba (1994), age-saving profiles for most countries are surprisingly flat, so that no reallocation of wealth to high-saving age groups could produce the dramatic differences in saving rates observed across countries. For example, the differential between Japan's saving rate and that of the US cannot be explained simply by differences in the relative wealth of different age cohorts because *no* age cohort in the US saves as much as the *lowest-saving* cohort in Japan, so *no* reshuffling of wealth across cohorts in the US could raise US saving to Japanese levels.

Another theoretical channel that could explain the positive correlation between saving and growth relies on transition dynamics in the standard growth model, in combination with a sufficiently high intertemporal elasticity of substitution. Consider a country that starts off capital-poor, and therefore has a high marginal product of capital. If the intertemporal elasticity of substitution is high enough, the high interest rate will induce a high saving rate, and high saving combined with a high marginal product of capital will produce rapid growth. While theoretically possible, however, this story does not correspond to the empirical evidence. Carroll and Summers (1991) show that there is no empirical relationship between rates of return and growth rates in their OECD sample, and Dekle (1993) shows that real interest rates were never particularly high in Japan, but were higher in the low-

saving, low-growth 1920s than in the high-saving, high-growth 1960s. Furthermore, this story would imply that saving rates in the high-growth countries should have been higher *early* in the sample and declined over time - the exact opposite of the observed pattern.⁵

In sum, there is now a substantial and diverse body of research showing that higher growth robustly leads to higher saving, and that such a correlation is difficult to reconcile with standard growth models.

3 The Model

Although the growth literature of the past decade has explored many possible assumptions about the nature of the aggregate production function, much less attention has been paid to the utility function that the representative agent is assumed to maximize. Standard practice, following tradition in the consumption literature, has been to assume that utility is time separable, and usually of the Constant Relative Risk Aversion (CRRA) form.

Several recent empirical papers in the microeconomic consumption literature, however, have argued that habits may play an important role in determining consumption. Contributions include Huib van de Stadt et. al. (1985), Carroll and Weil (1994) and Deaton and Paxson (1994).⁶ A separate macroeconomic literature on asset pricing under habit formation has also been developing, with prominent contributions by Andrew B. Abel (1990), George M. Constantinides (1990), Abel (1999), Urban J. Jermann (1998) and John Y. Campbell and John H. Cochrane (1999). Finally, two very recent papers make the case that habit formation may be essential in understanding the high-frequency dynamics of aggregate consumption data in the US (Jeffrey C. Fuhrer (2000)) and several OECD countries (Fuhrer and Michael W. Klein (1998)). Of course, the idea underlying this literature – that through the process of habit formation, one’s own past consumption might influence the utility yielded by current consumption – is hardly new; see, for example, James S. Duesenberry (1949) or Alfred A. Marshall (1898) (see pp. 86-91 or 110-111).

The implications of habit formation for the aggregate relationship between saving and growth, however, have not previously been examined in a rigorous formal growth model, to the best of our knowledge. Following the standard procedure of starting simple, we explore in this paper the relationship between saving and growth in a nonstochastic, perfect foresight model. We also make the simplest pos-

⁵See Carroll and Weil (1994) for a fuller exposition of the inability of the standard model to explain the observed facts.

⁶Karen E. Dynan (forthcoming) has used household-level data to estimate a modified Euler equation implied by the model with habits and found no statistically significant evidence for habit formation. However, a recent literature has shown that Euler equation tests may not be reliable, even for the standard version of the model without habits (Ludvigson and Paxson (1997); Carroll (1997)). Furthermore, if there is systematic measurement error in consumption, Dynan’s test would be biased toward finding no habit effects even if habits were in fact important.

sible assumption about the aggregate production function: it is of the AK form shown by Sergio T. Rebelo (1991) to be the ultimate underlying structure of all endogenous growth models.⁷ Finally, we use a functional form for utility which nests the two polar cases where the agent cares only about the level of consumption (the habit stock is irrelevant), and where the agent cares only about how consumption compares to the habit stock (the level of consumption is irrelevant), allowing us to explicitly show how changing the degree of habit formation affects the behavior of the model.

3.1 The Individual’s Problem

Consider the problem of an individual who cares about consumption relative to a “habit stock” determined by past consumption, and who takes into account the effect of current consumption on the future habit stock.⁸ The instantaneous utility function that we use, originally introduced by Abel (1990), is

$$U(c, h) = \frac{(c/h^\gamma)^{1-\sigma}}{1-\sigma} \quad (1)$$

where h is the stock of habits, c is the instantaneous flow of consumption, σ is the coefficient of relative risk aversion, and γ indexes the importance of habits. If $\gamma = 0$ then only the absolute level of consumption is important (the standard CRRA model), while if $\gamma = 1$, then consumption relative to the habit stock is all that matters. For values of γ between zero and one, both the absolute and the relative levels are important. For example, if $\gamma = .5$, then a person with consumption of 2 and habit stock of 1 would have the same utility as a person with both consumption and habit stock equal to 4. Finally, we assume $0 \leq \gamma < 1$ and $\sigma > 1$.

The stock of habits evolves according to

$$\dot{h} = \rho(c - h). \quad (2)$$

Thus, the habit stock is a weighted average of past consumption, with the parameter ρ determining the relative weights of consumption at different times. We assume $0 \leq \rho$. The larger is ρ , the more important is consumption in the recent past. If

⁷This is a strong assumption, but it greatly simplifies the analysis in comparison with models with a neoclassical production function, e.g. Harl E. Ryder and Geoffrey M. Heal (1973). We believe that the case for habits would only be strengthened by moving to a neoclassical growth model, because the negative effect of growth on saving in the CRRA utility version of the model would be amplified by an enhanced human wealth effect.

⁸The problem can be thought of in two ways: either the representative household cares directly about its *own* past consumption, or atomistic households care about how their consumption compares to a lagged average ‘standard of living’ and a social planner takes account of the negative externality that each household’s consumption has on all the other households. A related paper (Carroll, Overland, and Weil (1997)), shows that behavior is qualitatively similar in a model with atomistic households in which the externality is not taken into account by decision-makers (sometimes called a model with ‘external habits’).

$\rho = .1$, for example, then the half-life with which habits would adjust toward a permanent change in c is approximately 7 years (because $e^{-.1t} = 0.5$ for $t \approx 6.93$). If $\rho = .3$, then the half-life is a bit over two years.⁹

One implication of the sluggish adjustment of habits is that the introduction of habit formation does not change the risk aversion properties of the utility function at any given instant of time, because the habit stock is effectively fixed at any point in time. Thus it remains appropriate to call the parameter σ the coefficient of relative risk aversion. However, because habits can and do move over finite intervals in response to consumption choices, it will no longer be true that the intertemporal elasticity of substitution over time is equal to the inverse of the coefficient of relative risk aversion. We return to this point below.

As noted above, the production function is

$$y = Ak. \quad (3)$$

We assume that capital depreciates at rate $\delta \geq 0$. The capital stock thus evolves according to¹⁰

$$\dot{k} = (A - \delta)k - c. \quad (4)$$

The individual maximizes a discounted, infinite stream of utility:

$$\int_0^\infty U(c, h)e^{-\theta t} dt, \quad (5)$$

and the current value Hamiltonian is

$$H = U(c, h) + \psi[(A - \delta)k - c] + \lambda\rho(c - h). \quad (6)$$

⁹Models in which habit formation is used to explain the equity premium rely on a high value of ρ , so that the habit stock remains close to the current level of consumption. For example in Constantinides (1990), the values of ρ considered range as high as .6. Similarly, in Abel (1990), the habit stock is equal to the previous year's consumption. By contrast, in the growth context examined here, we think that lower values of ρ are appropriate, so that transitional dynamics are stretched over a substantial period of time. Our baseline assumption for the numerical exercises below will be $\rho = .2$.

¹⁰We are treating the economy as closed to international borrowing and lending here, both because it is hard to make sense of endogenous growth models with international capital transactions and because the evidence in Feldstein and Horioka (1980), recently confirmed and updated in Attanasio et. al. (2000), suggests that most investment is ultimately financed internally. We do not know of any papers that have examined the implications of habit formation for the current account, but intuition suggests that habit formation models might perform better than standard models in explaining the reaction of the current account to productivity shocks. Reuven Glick and Kenneth S. Rogoff (1995) show that a permanent income model of aggregate consumption implies that positive productivity shocks should cause a deterioration in the current account because in equilibrium the economy's permanent income rises by more than the rise in current income as capital adjusts upward to take advantage of higher productivity. Glick and Rogoff find instead that consumption does not adjust as much as the model predicts. Habits might explain such slow adjustment of consumption. (We are grateful to Joseph Gruber for bringing this point to our attention).

Carroll, Overland, and Weil (1997) present the full solution to this problem with equations of motion relating consumption, the capital stock, and the habit stock. In the steady state, c , k , and h all grow at the same rate. Here we analyze the problem in terms of three ratios, c/h , \dot{c}/c , and k/h that are constant in steady state. Dynamics arise from departures of these “state-like variables” from their steady state values, as in Robert E. Lucas, Jr. (1988) and Casey B. Mulligan and Xavier Sala-i-Martin (1993).

The equations of motion are

$$\left[\frac{\dot{c}}{h}\right] = \frac{c}{h} \left(\frac{\dot{c}}{c} - \rho\left(\frac{c}{h} - 1\right)\right). \quad (7)$$

$$\begin{aligned} \left[\frac{\dot{c}}{c}\right] &= \alpha_0 + \alpha_1 \left[\frac{\dot{c}}{c}\right] \\ &+ \alpha_2 \left[\frac{c}{h}\right] + \alpha_3 \left[\frac{\dot{c} c}{c h}\right] \\ &+ \alpha_4 \left[\frac{\dot{c}}{c}\right]^2 + \alpha_5 \left[\frac{c}{h}\right]^2 \end{aligned} \quad (8)$$

$$\left[\frac{\dot{k}}{h}\right] = \frac{k}{h} \left(A - \delta - \rho\left(\frac{c}{h} - 1\right)\right) - \frac{c}{h}. \quad (9)$$

Equation (8) shows the change in the rate of consumption growth as a function of the level of consumption growth and the ratio of consumption to the habit stock (the coefficients $\alpha_0 \dots \alpha_5$ are functions of the taste and technology parameters; see the appendix (A.1) for the explicit version of the equation of motion for consumption as a direct function of taste and technology parameters). Note that this differs from the usual Euler condition that emerges from a Ramsey model in that the second time derivative of consumption is involved.¹¹ In intuitive terms, this result arises because the consumer’s utility is now affected by the growth rate of consumption (through the effect of that growth rate on c/h^γ) as well as the level of consumption, so the temporal evolution of the *growth rate* must satisfy an optimality condition, just as in the Ramsey model the temporal evolution of the *level* of consumption must satisfy an optimality condition. Intuitively, habit-forming consumers will desire to smooth consumption growth rates for essentially the same reasons that CRRA consumers desire to smooth levels of consumption.

Setting the three dynamic equations equal to zero determines the steady state of the model,¹²

¹¹In Carroll, Overland, and Weil (1997), we show that in the cases where $\gamma = 0$, so that habit stock has no effect on utility, or where $\rho = 0$, so that the habit stock is unchanging, equation (9) reduces to the first order condition from the standard Rebelo model.

¹²In our companion paper we show that there is a second, extraneous solution to these equations that is not related to optimal behavior.

$$\left(\frac{\dot{c}}{c}\right) = \frac{A - \delta - \theta}{\gamma(1 - \sigma) + \sigma}, \quad (10)$$

$$\left(\frac{c}{h}\right) = 1 + \frac{1}{\rho} \left(\frac{A - \delta - \theta}{\gamma(1 - \sigma) + \sigma} \right) \quad (11)$$

$$= 1 + \frac{1}{\rho} \left(\frac{\dot{c}}{c}\right), \quad (12)$$

$$\frac{k}{h} = \frac{1}{\rho} \left[\frac{\rho(\gamma(1 - \sigma) + \sigma) + (A - \delta - \theta)}{(A - \delta)[(1 - \sigma)\gamma + \sigma - 1] + \theta} \right]. \quad (13)$$

Equation (12) indicates that the rate at which the habit stock catches up with consumption, ρ , affects the steady state ratio of consumption to habit stock in an intuitive way: with a higher ρ and thus a faster catchup of the habit stock, the ratio of consumption to the habit stock gets closer to one.

Equation (10) shows the effect of the parameters on the steady state growth rate of consumption, which is also the steady state growth rate of capital, output, and the habit stock. Note that ρ does not affect the steady-state growth rate (although we show in our companion paper that the value of ρ does affect transitional dynamics). However, the other habit parameter, γ , which captures the extent to which consumers care about how consumption compare to habits, has an important effect on the steady-state growth rate. Higher values of γ will lead to a higher growth rate of consumption in the steady state (recall that earlier we assumed that $\sigma > 1$.)

One way to interpret this result is to think of habits as increasing the infinite-horizon value of the intertemporal elasticity of substitution. Indeed, the intertemporal elasticity of substitution in consumption is *defined* as the response of consumption growth to interest rates. Since the interest rate in this model is $A - \delta$, equation (10) implies that the infinite-horizon intertemporal elasticity of substitution in this model is $1/(\gamma(1 - \sigma) + \sigma)$, which for $\sigma > 1$ and $0 < \gamma < 1$ is strictly greater than the inverse of the coefficient of relative risk aversion $1/\sigma$. However, if we were to calculate the intertemporal elasticity of substitution with respect to *temporary* changes in the interest rate, we would discover that as the interval of the temporary change in the interest rate approaches zero, the intertemporal elasticity of substitution approaches $1/\sigma$.

The reason for the discrepancy between the short-horizon and the long-horizon elasticities is that over a sufficiently short interval the habit stock is effectively fixed, while over a sufficiently long interval the habit stock is effectively perfectly flexible. Intuitively, the gain or loss in utility associated with a given increase or decrease in consumption over a long horizon will be diminished by the associated movement in the habit stock; this reduction in the effective curvature of the utility function constitutes an increase in the effective intertemporal elasticity.¹³

¹³One implication of this result is that coefficient estimates obtained from regressions of con-

Another way to interpret the consumer's problem can be seen if we substitute the steady-state relationship between c , h , and growth into the expression c/h^γ (the object which is raised to the power $1 - \sigma$ to generate utility). From (12) we know that, designating the steady-state growth rate of the economy as g , in the steady-state

$$h = c/(1 + g/\rho).$$

This implies that

$$\begin{aligned} ch^{-\gamma} &= c[c/(1 + g/\rho)]^{-\gamma} \\ &= c^{1-\gamma}(1 + g/\rho)^\gamma, \end{aligned} \tag{14}$$

so we can think of the consumer as maximizing the utility from a geometrically weighted average of the level of consumption and (a linear function of) the growth rate of consumption. If the weight on habits is $\gamma = 0$ this expression just collapses to c and the consumer is maximizing utility from the level of consumption; if $\gamma = 1$ the consumer is maximizing only the utility which derives from the growth of consumption and the level is unimportant.

4 Implications of the Model for Saving and Growth

In this section, we take up the question of how allowing for habit formation changes the response of the economy to exogenous changes to productivity or capital. We show that allowing for habit formation can substantially change both the quantitative and qualitative response of saving to such events, and then we discuss the relationship of the model's results to the empirical evidence.

We begin by examining the steady-state relationship between saving and growth rates, then turn to transitional dynamics.

4.1 Steady States

The mathematical appendix shows that the derivative of the gross saving rate with respect to the growth rate of output will be positive only if

$$\sigma < 1 + \frac{\theta}{\delta(1 - \gamma)}. \tag{15}$$

sumption growth on, say, quarterly interest rates would not yield estimates of a 'pure' structural parameter in an economy composed of habit-forming consumers, because the estimated coefficient would in principle depend on the duration of interest rates as well as the level. Because the estimated coefficient should fall somewhere between the instantaneous elasticity (which is the inverse of σ) and the infinite-horizon elasticity, such empirical estimates of σ should understate both the instantaneous value of σ and the long-horizon intertemporal elasticity of substitution. Thus, as Constantinides (1990) and Campbell and Cochrane (1999) have shown, a model with habit formation can explain the equity premium puzzle by assuming very high instantaneous risk aversion while simultaneously avoiding some of the unattractive implications of a model with a correspondingly low intertemporal elasticity of substitution.

In the baseline model where habits do not matter ($\gamma = 0$), if $\theta = \delta$, for example, the relation between saving and growth is positive only if the instantaneous coefficient of relative risk aversion, σ , is less than two.¹⁴ Most evidence, however, suggests an instantaneous coefficient of relative risk aversion considerably greater than two.¹⁵ Note that habit formation (a choice of $0 < \gamma < 1$) increases the range of parameter values for which increases in the growth rate of output due to increases in the productivity parameter A are associated with a higher saving rate. For example, if $\theta = \delta$ and $\gamma = .75$, then $\frac{ds}{dg} > 0$ so long as $\sigma < 5$.

There are two ways to interpret the fact that habits make the relationship between saving and growth more positive. The first is that this is a consequence of the corresponding increase in the infinite-horizon intertemporal elasticity of substitution: habits make consumers more willing to postpone consumption in response to an increase in interest rates, and thus make the saving response to A stronger. The second interpretation derives from the earlier observation that introducing habits is like putting growth in the utility function. Increasing the value of A makes it possible for consumers to achieve higher steady-state growth rates. Since habit-forming consumers care directly about the growth rate of consumption, they will take advantage of a higher A partly to boost the steady-state growth rate (by increasing the saving rate). Regardless of the interpretation, it is clear that raising the level of habit formation can qualitatively change the relation between growth and saving in the steady state.

4.2 Dynamics

4.2.1 Policy Functions

In order to examine transition dynamics in our model, we derive policy functions tracing out the relationship between the state variable k/h and the optimal values of the control variable c/h . Similarly, we can trace the relationship between k/h and any transformation of the control variable along the optimal path. This amounts to graphing the optimal policy functions relating the state variable to each of the policy variables in question.

Figures 1 and 2 depict policy functions for the main variables of interest for several different values of γ , the parameter that determines importance of habits in utility, and σ , the coefficient of relative risk aversion. For each value of γ , the value of σ is chosen to keep the steady-state growth rate the same.¹⁶ The dots represent

¹⁴It turns out that this result is not unique to the endogenous growth model. In the Appendix we show that equation (15), with γ set to zero, must also hold in the Cass-Koopmans-Ramsey model if that model is to generate a positive steady-state relationship between saving and growth.

¹⁵Note that the choice of $\theta = \delta$ almost certainly *understates* the problem for the standard model, because in typical parameterizations θ is usually assumed to be considerably smaller than δ . For example, if $\theta = .03$ and $\delta = .09$ (relatively conventional choices), then the coefficient of relative risk aversion must be less than $4/3$ in order for the relationship between saving and growth to be positive.

¹⁶As can be seen in equations (10) and (13), the steady-state values of all of the “state like”

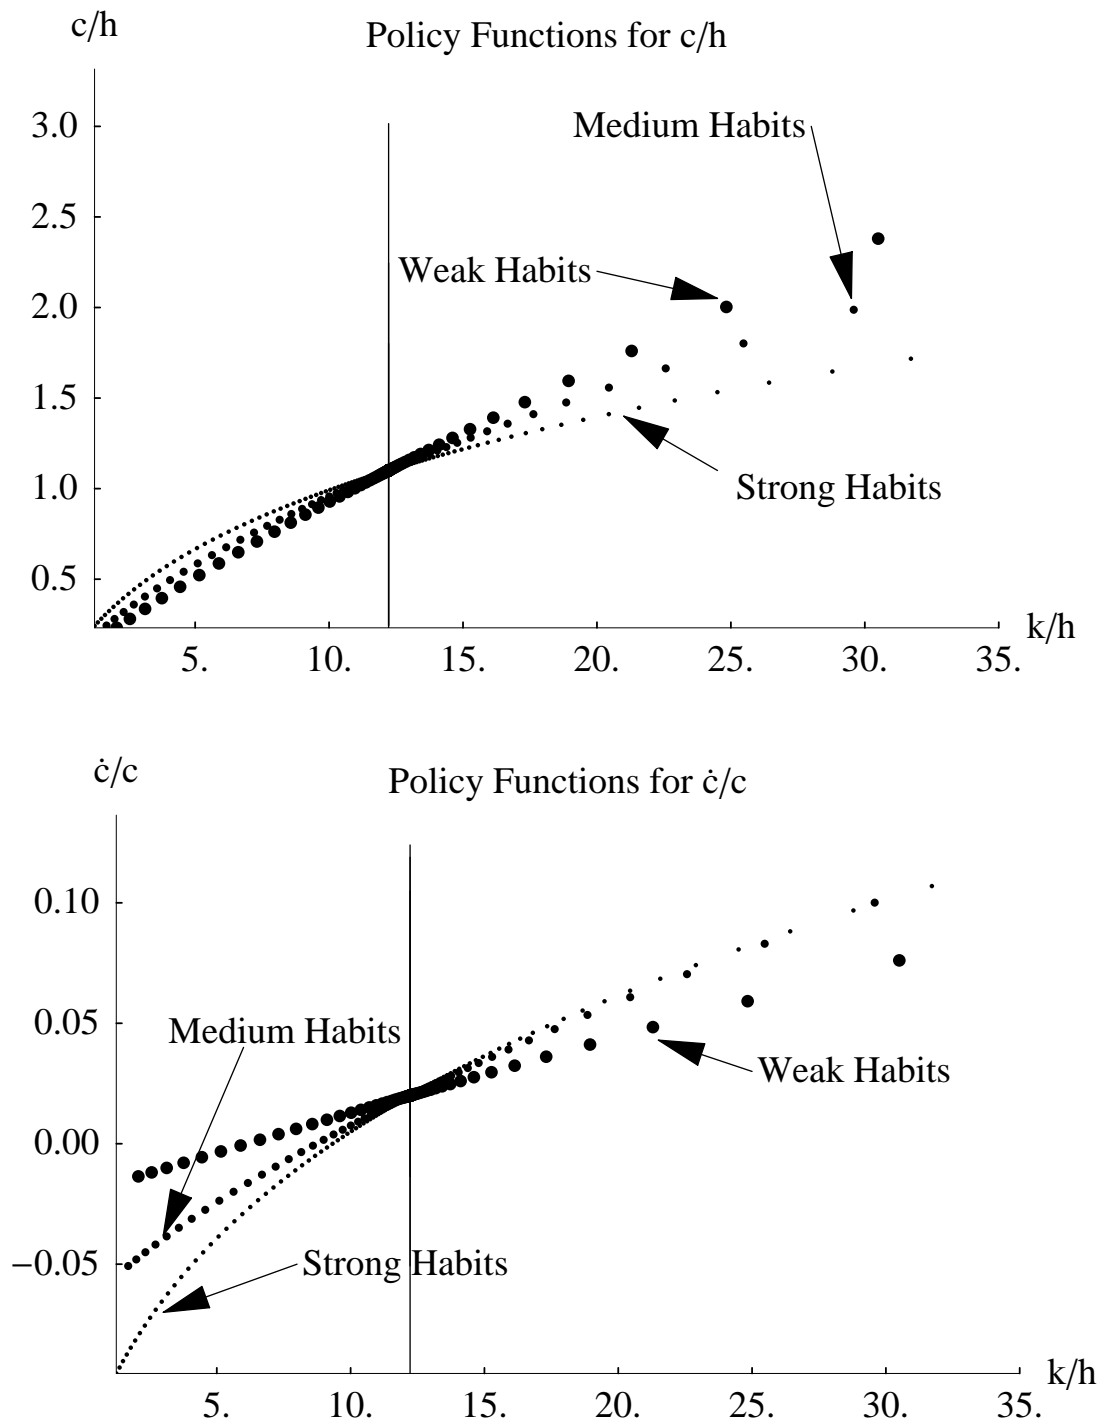


Figure 1: Policy Functions for c/h and \dot{c}/c

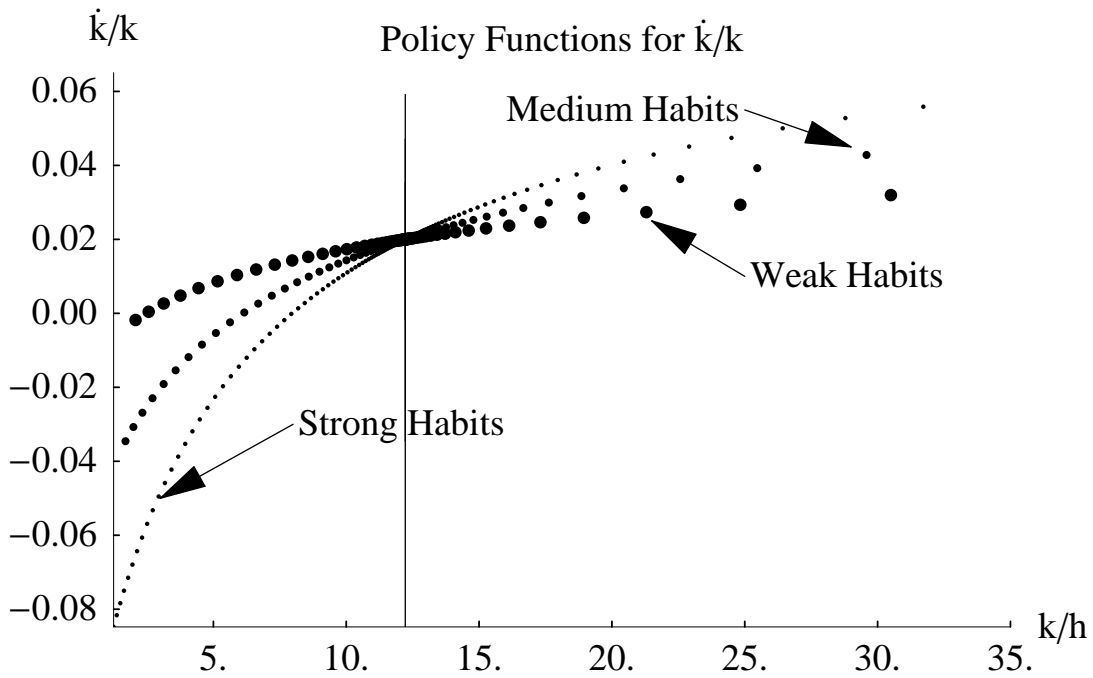
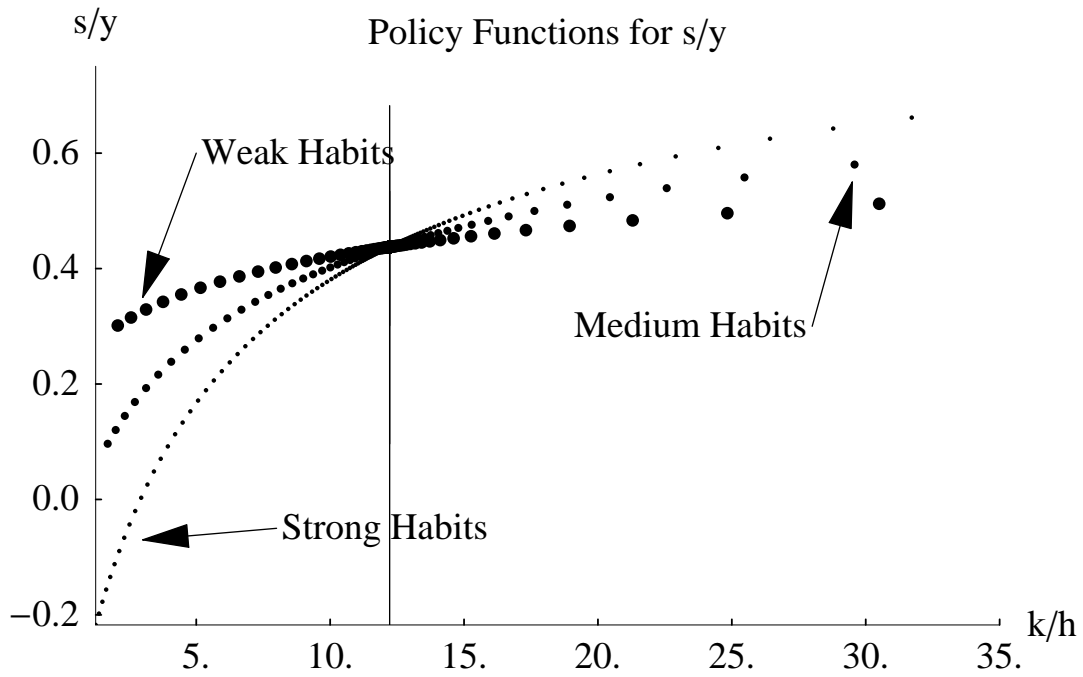


Figure 2: Policy Functions for s/y and \dot{k}/k

equally spaced points in time as the system evolves toward the shared steady state (where the three policy functions intersect). The largest dots correspond to policy functions for the case where habits are weakest, $\gamma = .25$, the medium-sized dots correspond to a medium degree of habit formation, $\gamma = .5$, and the smallest dots correspond to $\gamma = .75$.

The first set of figures shows that an economy that starts out ‘rich,’ in the sense of having a capital-to-habit-stock ratio above the steady state, will initially have both a higher-than-steady-state ratio of consumption to the habit stock and a higher-than-steady-state level of consumption growth. But the second set of figures shows that such a ‘rich’ economy will also have a higher-than-steady-state saving rate, implying a *lower*-than-steady-state ratio of consumption to capital (and therefore income). Thus, compared to the steady-state ratios, consumption is high relative to habits but low relative to income.

The intuition for this pattern is simple: habits tend to pull consumption toward the level of the habit stock and away from the steady-state ratio of consumption to income. If habits are low relative to capital (k/h is high), then consumption will be low relative to capital (the saving rate will be above the steady-state level). Another way to put this is that some of the economy’s good fortune is taken advantage of via a high level of consumption relative to habits, but growth-loving consumers use some of their good fortune to achieve an extended period of above-steady-state growth by saving more than the steady-state amount.

Comparing policy functions for different values of γ , several points stand out. First, for a given level of k/h which is above its steady state level, an economy in which the influence of habits on consumption is stronger will have a lower level of consumption (higher saving rate), and higher growth rates of consumption and output. This is because when habits matter more, the pull on consumption toward the habit stock is stronger. It is also clear from the spacing of the dots representing points in time that an economy with a lower degree of habit formation moves more rapidly toward the steady state for any given initial value of k/h . This is unsurprising because we know that in an economy with no habit formation at all there is no pull of c toward h ; thus the gap between c and h is larger and so h will adapt to c faster.

Thus an economy with a high γ (i.e. habits are particularly important to utility) and which starts out with a high k/h ratio will experience a prolonged period of having c/k below its steady state level, and therefore saving and growth above their steady-state levels. Conversely, in response to an initially low level of k/h , an economy with intense habit effects will preserve a higher initial level of consumption relative to habit stock, and will pay for its high level of consumption with an extended period of below-steady-state growth.

variables depend on γ and σ only through the term $\gamma(1 - \sigma) + \sigma$. We examine pairs of γ and σ that hold the value of this term at 3, which is the value consistent with a coefficient of relative risk aversion of 3 if γ were equal to zero.

4.2.2 The Dynamic Response to an Unanticipated Drop in Capital

To examine how habit formation affects co-movements of saving and growth, we consider the following experiment. The economy is in steady state, with A chosen such that the growth rate of output is 2 percent per year. In year 0, 10 percent of the capital stock is destroyed. Figure 3 shows the evolution of output growth and the saving rate following the shock. As shown above, so long as A remains unchanged, the steady state saving rate and growth rates will also be unaffected. However, the degree of habit formation importantly affects the transitional dynamics.

We consider pairs of values of γ and σ which hold the long run saving rate and growth rates constant. These are the same values used in Figures 1 and 2: $(\gamma = .25, \sigma = \frac{11}{3})$, $(\gamma = .5, \sigma = 5)$, and $(\gamma = .75, \sigma = 9)$.

For a high value of γ , the immediate effect of the destruction of the capital stock is to greatly reduce the saving rate. The reduction in saving is the result of the desire to maintain consumption relative to the existing habit stock. The low level of saving, in turn, reduces the growth rate of output. Both saving and growth return only slowly to their steady state levels. Not only is the effect on saving larger when γ is large, but it is also more persistent. When $\gamma = .75$, saving returns halfway to its steady state level after 11.6 years. For $\gamma = .5$ the half-life of is 6.8 years, and for $\gamma = .25$ it is 5.2 years. In the limit, when γ is zero, there is no effect of the drop in capital on either saving or growth. Of course, output will decline by 10 percent when the capital is destroyed. Thus, viewed over a time period that encompasses the drop, growth will fall and saving will be constant if $\gamma = 0$, while both growth and saving will fall if $\gamma > 0$. Note also that the persistent changes in growth observed when $\gamma > 0$ will be reflected in permanent differences in the level of output once countries have reached their steady state growth rates.

4.2.3 The Dynamic Response to a Change in A

Now consider the effect of an unanticipated but permanent change in A such that the steady-state growth rate of the economy moves from 1 percent to 2 percent per year. We choose parameters such that this change in productivity has a negative long-run effect on the saving rate. Figure 4 shows the paths of output growth and the saving rate following the change for economies with different levels of γ .

The figure shows that allowing for habit formation can make the short run effect of a change in A differ from the long run effect. Even though saving falls in the long run, it remains above its long run level during the transition to the new steady state, and even rises temporarily above its initial level under two of the three assumptions about the strength of habits. Growth also rises above its long-run level during the transition. The more powerful are habits, the larger and longer-lived are these transitional effects.

These experiments complement the discussion in section (4.1) which showed that allowing for habit formation substantially expands the set of parameter values for which increases in A will result in *long-run* increases in both growth and saving.

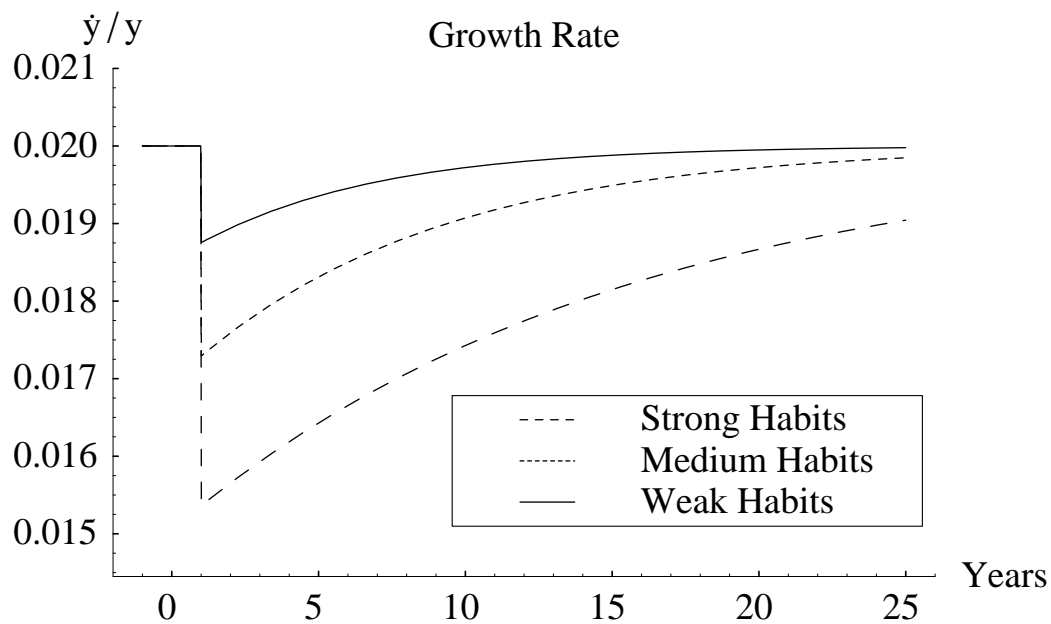
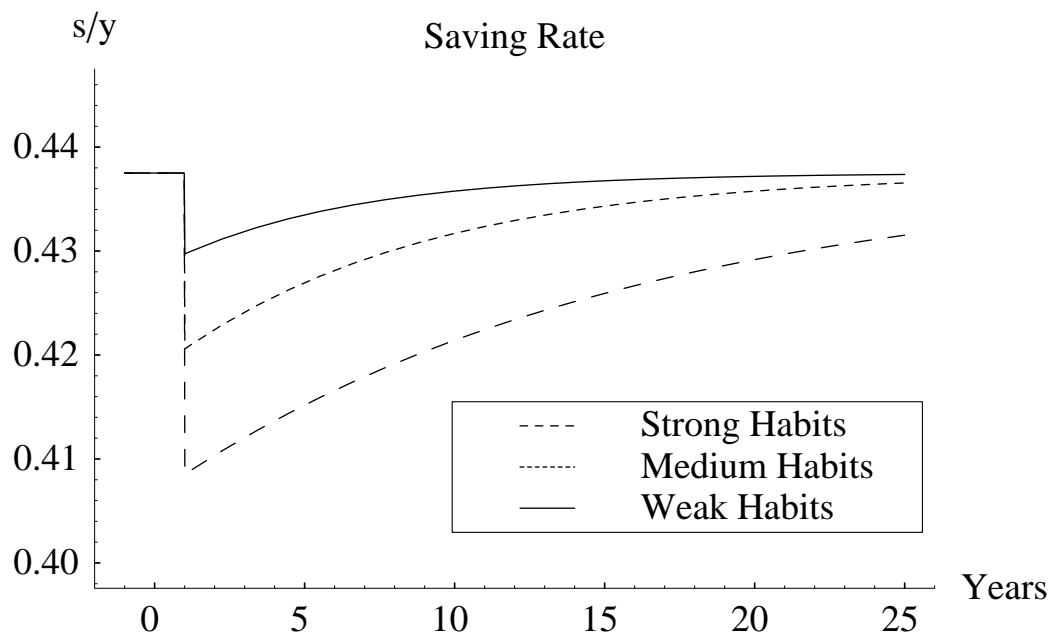


Figure 3: Dynamics Following Destruction of 10 Percent of k

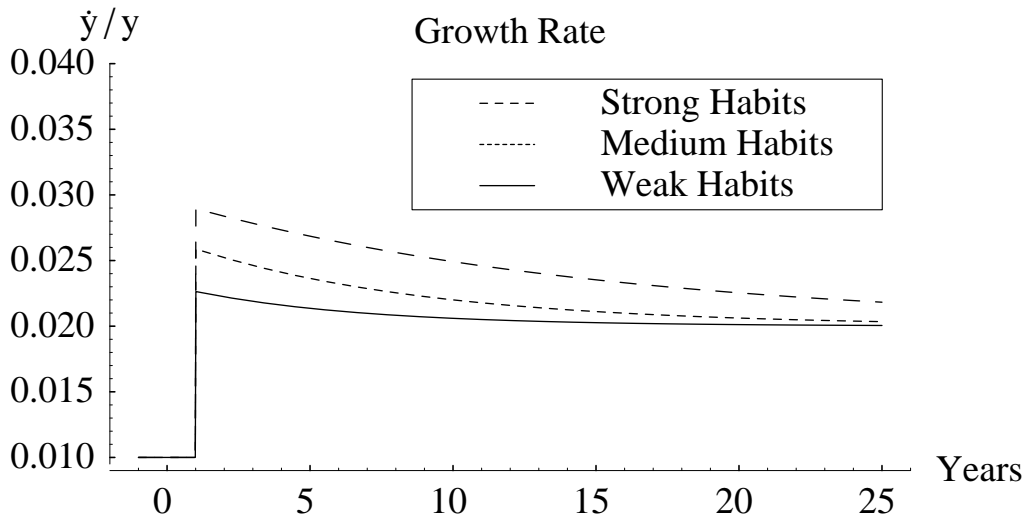
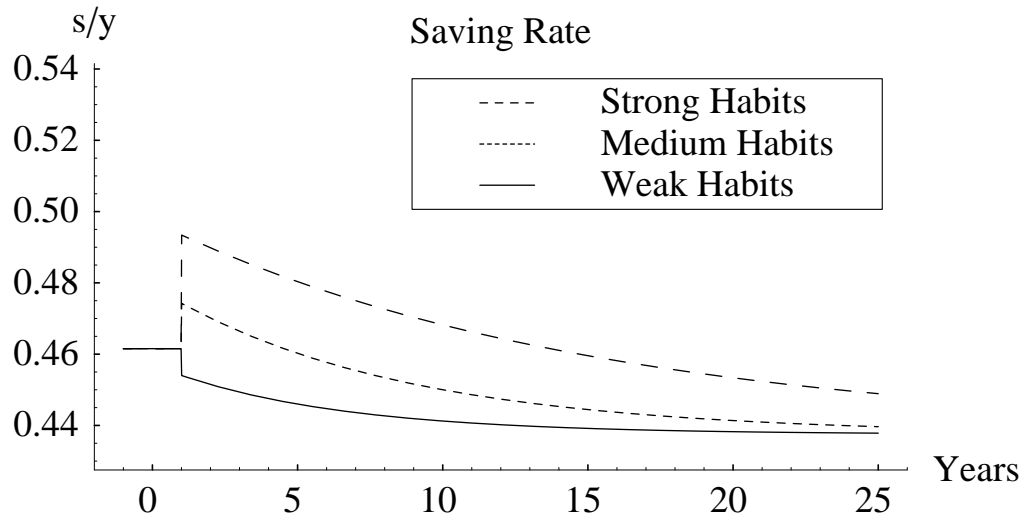


Figure 4: Dynamics Following an Unexpected Increase in A

Here we have shown that an increase in A will result in *short-run* increases in both growth and saving for an even larger set of parameter values.

4.3 Does the model match reality?

We have shown that the range of circumstances under which our model predicts a positive correlation between saving and growth is much greater than the range of circumstances under which the standard model predicts a positive correlation. These findings are directly consistent with growth-to-saving causality evidence like that presented by Edwards (1995), Bosworth (1993), and the microeconomic evidence in Carroll and Weil (1994) and Deaton and Paxson (1994).

Some of the most compelling evidence on causality running from growth to saving, however, came from timing: increases in growth precede increases in saving. We suggested that such a mechanism may be responsible for the high and rising saving rates in the fast-growing East Asian countries over the past thirty years. How does this result on timing fit with the model presented here?

The first step in answering this question is to determine how to interpret the experience of the East Asian countries in the context of a model of this kind. The most plausible interpretation is probably that much of the growth there resulted from the rapid import of technology from abroad. The evidence in William R. Easterly et al. (1993) suggests that the best way to model the growth experiences in the East Asian countries is as a series of positive shocks. Thus we might interpret the East Asian experience as a sequence of exogenous increases in the ‘broad capital’ embodied in k in our model. As the simulations above indicate, a series of positive shocks to k should result in a rising national saving rate. One prediction of our model is that saving rates in the East Asian countries should decline once those economies stop their technological convergence with more advanced economies. The decline in Japan’s national saving rate over the past two decades is therefore consistent with our model, but this is the only real test of the proposition thus far, because the other East Asian countries have not yet come close to technological convergence.

This discussion of the relationship between the theory and the East Asian experience is obviously rather loose, and a more careful analysis would certainly be a valuable topic for future research.¹⁷

5 Conclusion

Is habit formation an appealing explanation for the evidence that saving and growth are positively correlated across countries? In part, the answer comes down to the question of how countries differ: in preferences, technologies, initial conditions, or

¹⁷Carroll (2001) has recently shown that in a small open economy in which consumers receive idiosyncratic stochastic shocks to their income, habits can produce true Granger causality from growth to saving, but much more work remains to be done.

history of shocks. If one is willing to believe that differences in preferences, e.g. time preference rates, are the primary source of variation in saving rates across countries, then the positive cross-country correlation between average saving and average growth can be explained by a straightforward endogenous growth model without habit formation. However, one is then forced to address the question of why national saving rates change so dramatically over time; did Koreans simply become much more patient in the 1970s and 1980s? Furthermore, the limited available direct evidence does not support the proposition that differences in preferences explain cross-country differences in national saving rates.¹⁸ If countries differ instead in the level of the productivity parameter, A , in an endogenous growth model (or if they differ in their rates of productivity growth in a Ramsey model), then whether growth and saving will be positively correlated depends on the coefficient of relative risk aversion. For the standard model with values of this coefficient in the range usually considered, it should be the case that growth and saving are negatively correlated across countries. If one were committed to an A -based explanation for growth differentials across countries, the introduction of habit formation in utility could potentially help explain the observed data.

The argument that habit formation is important seems even stronger when dynamic evidence is considered. The empirical evidence cited above indicates that increases in growth tend to be followed by increases in saving. In this paper we have shown that habit formation can lead to a positive short-run response of saving to a favorable shock, even when there is no long-run effect of such a shock on saving. Finally, even if all countries have the same taste and technology parameters and thus the same steady-state growth rate, we show that allowing for habit formation in consumption leads to a positive correlation between saving and growth along transition paths to the steady state; in the endogenous growth model in the absence of habit formation, such transitions do not take place (that is, countries are always at their steady states), and so all countries would have the same saving and growth rates at all times.

¹⁸Carroll et al. (1994) show that immigrants to Canada from high-saving countries do not save more than immigrants from low-saving countries; Carroll et. al. (1999) present similar results for immigrants to the U.S.

A Appendix

A.1 The First Order Condition for Consumption

For expositional simplicity, equation (8) in the text presented a version of the equation of motion for consumption which hid the full complexity of that equation behind a set of coefficients α_0 to α_5 . The full expression for the equation of motion for consumption as a function of the model's taste and technology parameters is

$$\begin{aligned} \left[\frac{\dot{c}}{c}\right] &= \sigma \left(\frac{\dot{c}}{c}\right)^2 + \frac{\dot{c}}{c}(2\theta + \rho + \delta - A - 2\gamma\rho(1 - \sigma)) - \rho^2\gamma(\gamma(1 - \sigma) + 1) \left(\frac{c}{h}\right)^2 \\ &+ 2\gamma\rho(1 - \sigma)\frac{\dot{c}}{c}\frac{c}{h} + \left(\frac{\rho\gamma}{\sigma}\right)\frac{c}{h}(\rho\gamma(1 - \sigma)(2\sigma - 1) + \theta + \rho - \sigma(2\theta + \delta - A)) \\ &+ \frac{1}{\sigma}((\rho + \theta)(\theta + \delta - A) + \rho\gamma(1 - \sigma)(\rho(\gamma(1 - \sigma) + 1) - (2\theta + 2\rho + \delta - A))) \end{aligned} \quad (\text{A.1})$$

A.2 The Relation Between Saving and Growth in the Habit-Formation Model

This appendix derives the condition under which the steady-state relationship between saving and growth is positive, equation (15) in the paper. We consider changes in growth that are due to variation in the parameter A , which measures productivity. Define g as the steady state growth rate of output.

From equation (10)

$$\frac{dg}{dA} = \frac{1}{\gamma(1 - \sigma) + \sigma}.$$

The gross saving rate is

$$s = \frac{y - c}{y} \quad (\text{A.2})$$

$$= \frac{AK - c}{AK}. \quad (\text{A.3})$$

In steady-state, income, capital, and consumption all grow at the same rate g . With an AK production function with depreciation rate δ , gross saving must be enough to make the capital stock grow at rate g after depreciation:

$$s = \frac{(g + \delta)K}{AK} \quad (\text{A.4})$$

$$= \frac{(g + \delta)}{A}. \quad (\text{A.5})$$

Differentiating this expression with respect to g yields

$$\frac{ds}{dg} = \frac{A - (g + \delta) \frac{dA}{dg}}{A^2}.$$

The sign of ds/dg depends only on the numerator of this expression. Inverting equation (A.2) and substituting it into the numerator, this condition becomes

$$A - (g + \delta)(\gamma(1 - \sigma) + \sigma) > 0.$$

Finally, using equation (10) to substitute for g , we can re-write the condition as

$$A - \left(\frac{A - \delta - \theta}{\gamma(1 - \sigma) + \sigma} + \delta \right) (\gamma(1 - \sigma) + \sigma),$$

which reduces to (15) in the text.

A.3 The Relation Between Saving and Growth in the Cass-Ramsey-Koopmans Model

Robert J. Barro and Xavier Sala-i-Martin ((1995), equation 2.31) derive the gross saving rate in the steady state of a Ramsey model as (using our notation, assuming zero population growth, and calling the rate of labor-augmenting technical progress g)

$$s = \frac{\alpha(g + \delta)}{\delta + \theta + \sigma g}.$$

Differentiating this expression with respect to g yields

$$\frac{ds}{dg} = \frac{\alpha[(\delta + \theta + \sigma g) - (g + \delta)\sigma]}{(\delta + \theta + \sigma g)^2}.$$

The sign of this expression depends on the sign of the term in square brackets, and is positive if

$$\frac{\delta + \theta}{\delta} > \sigma,$$

which corresponds to equation (15) in the text if $\gamma = 0$.

References

- ABEL, A. B. (1990): “Asset Prices under Habit Formation and Catching Up with the Joneses,” *American Economic Review*, 80(2), 38–42.
- (1999): “Risk Premia and Term Premia in General Equilibrium,” *Journal of Monetary Economics*, 43(1), 3–33.
- ATTANASIO, O., L. PICCI, AND A. SCORCU (2000): “Saving, Growth, and Investment: A Macroeconomic Analysis Using a Panel of Countries,” *Review of Economics and Statistics*, 82(1).
- BARRO, R. J., AND X. S. I MARTIN (1995): *Economic Growth*. McGraw-Hill, New York.
- BOSWORTH, B. (1993): *Saving and Investment in a Global Economy*. The Brookings Institution, Washington, DC.
- CAMPBELL, JOHN Y., AND J. H. COCHRANE (1999): “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107(2), 205–251.
- CARROLL, C. D. (1997): “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second Order Approximation),” *NBER Working Paper Number 6298*, <http://www.econ.jhu.edu/people/ccarroll/death.pdf> .
- (2001): “‘Risky Habits’ and the Marginal Propensity to Consume Out of Permanent Shocks,” *Forthcoming, International Economic Journal*.
- CARROLL, C. D., J. R. OVERLAND, AND D. N. WEIL (1997): “Comparison Utility in a Growth Model,” *Journal of Economic Growth*, 2(4), 339–367, <http://www.econ.jhu.edu/people/ccarroll/compare.pdf> .
- CARROLL, C. D., C. RHEE, AND B. RHEE (1994): “Are There Cultural Effects on Saving? Some Cross-Sectional Evidence,” *The Quarterly Journal of Economics*, CIX(3), 685–700.
- (1999): “Does Cultural Origin Affect Saving Behavior? Evidence from Immigrants,” *Economic Development and Cultural Change*, 48(1), 33–50.

- CARROLL, C. D., AND L. H. SUMMERS (1991): "Consumption Growth Parallels Income Growth: Some New Evidence," in *National Saving and Economic Performance*, ed. by B. Douglas Bernheim, and John B. Shoven. Chicago University Press, Chicago.
- CARROLL, C. D., AND D. N. WEIL (1994): "Saving and Growth: A Reinterpretation," *Carnegie-Rochester Conference Series on Public Policy*, 40, 133–192.
- CONSTANTINIDES, G. M. (1990): "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98, 519–43.
- DEATON, ANGUS S., AND CHRISTINA H. PAXSON (1994): "Saving, Growth, and Aging in Taiwan," in *Studies in the Economics of Aging*, ed. by D. A. Wise. University of Chicago Press, Chicago.
- (1997): "The Effects of Economic and Population Growth on National Saving and Inequality," *Demography*, 34(1), 97–114.
- (2000): "Saving and Growth: Another Look at the Cohort Evidence," *Review of Economics and Statistics*, 82(1).
- DEKLE, R. (1993): "Raising Saving Rates: Lessons from the Japanese Experience," *Background Paper Prepared for The East Asia Miracle Project*.
- DUESENBERY, J. S. (1949): *Income, Saving, and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, MA.
- DYNAN, K. E. (forthcoming): "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*.
- EASTERLY, WILLIAM R., KREMER, MICHAEL R., PRITCHETT, LANT H., AND SUMMERS, LAWRENCE H. (1993): "Good Policy or Good Luck? Country Growth Performance and Temporary Shocks," *Journal of Monetary Economics*, 32(3), 459–83.
- EDWARDS, S. (1995): "Why are Saving Rates so Different Across Countries?: An International Comparative Analysis," NBER Working Paper 5097.
- FELDSTEIN, M. S., AND C. Y. HORIOKA (1980): "Domestic Saving and International Capital Flows," *Economic Journal*, 90, 314–29.
- FUHRER, J. C. (2000): "An Optimizing Model for Monetary Policy: Can Habit Formation Help?," *American Economic Review*, 90(3).
- FUHRER, J. C., AND M. W. KLEIN (1998): "Risky Habits: On Risk Sharing, Habit Formation, and the Interpretation of International Consumption Correlations," *NBER Working Paper No. 6735*.

- GLICK, R., AND K. S. ROGOFF (1995): "Global Versus Country-Specific Productivity Shocks and the Current Account," *Journal of Monetary Economics*, 35(1), 159–192.
- HALL, R. E., AND C. I. JONES (1999): "Why Do Some Countries Produce So Much More Output per Worker than Others?," *Quarterly Journal of Economics*, CXIV, 83–116.
- HOUTHAKKER, H. S. (1961): "An International Comparison of Personal Saving," *Bulletin of the International Statistical Institute*, 38, 55–60.
- (1965): "On Some Determinants of Saving in Developed and Underdeveloped Countries," in *Problems in Economic Development*, ed. by E. Robinson. MacMillan, London.
- JERMANN, U. J. (1998): "Asset Pricing in Production Economies," *Journal of Monetary Economics*, 42(2), 257–75.
- KING, ROBERT G., AND LEVINE, ROSS E. (1994): "Capital Fundamentalism," *Carnegie-Rochester Conference Series on Public Policy*, 40.
- KLENOW, P. J., AND A. RODRÍGUEZ-CLARE (1997): "The Neoclassical Revival in Growth Economics: Has It Gone Too Far?," *NBER Macroeconomics Annual, 1997*, pp. 73–103.
- LEVINE, ROSS E., AND DAVID RENELT (1992): "A Sensitivity Analysis of Cross-Country Growth Regressions," *American Economic Review*, 82, 942–963.
- LOAYZA, N., K. SCHMIDT-HEBBEL, AND L. SERVÉN (2000): "What Drives Saving Across the World?," *Review of Economics and Statistics*, 82(1).
- LUCAS, R. E. (1988): "On The Mechanics of Economic Development," *Journal of Monetary Economics*, 22(1), 3–42.
- LUDVIGSON, S., AND C. PAXSON (1997): "Approximation Bias in Euler Equation Estimation," *Manuscript, Federal Reserve Bank of New York*.
- MARSHALL, A. (1898): *Principles of Economics*. The MacMillan Company, New York, eighth edn.
- MODIGLIANI, F. (1970): "The Life Cycle Hypothesis of Saving and Inter-Country Differences in the Saving Ratio," in *Induction, Growth, and Trade: Essays in Honor of Sir Roy Harrod*, ed. by W. A. Eltis. Clarendon Press, London.
- (1986): "Life Cycle, Individual Thrift, and the Wealth of Nations," *American Economic Review*, 3(76), 297–313.

- MULLIGAN, CASEY B., AND XAVIER SALA-I-MARTIN (1993): "Transitional Dynamics in Two-Sector Models of Endogenous Growth," *Quarterly Journal of Economics*, CVIII(3), 739–773.
- PAXSON, C. H. (1995): "Saving and Growth: Evidence from Micro Data," *European Economic Review*, 40, 255–88.
- POTERBA, J. M. (ed.) (1994): *International Comparisons of Household Saving*. National Bureau of Economic Research, Chicago and London: University of Chicago Press.
- REBELO, S. T. (1991): "Long-Run Policy Analysis and Long-Run Growth," *Journal of Political Economy*, 99(3), 500–521.
- RODRIK, D. (1999): "Saving Transitions," *Manuscript, Harvard University*.
- RYDER, HARL, AND GEOFFREY HEAL (1973): "Optimal Growth with Intertemporally Dependent Preferences," *Review of Economic Studies*, 40, 1–31.
- SHAPIRO, M. D., AND J. B. SLEMROD (1995): "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding," *American Economic Review*, 85(1), 274–283.
- TOBIN, J. (1967): "Life Cycle Saving and Balanced Growth," in *Ten Economic Studies in the Tradition of Irving Fisher*, ed. by W. Fellner. Wiley, New York.
- VAN DE STADT, H., A. KAPTEYN, AND S. VAN DE GEER (1985): "The Relativity of Utility: Evidence from Panel Data," *Review of Economics and Statistics*, 67(2), 179–187.
- WILLIAMSON, J. G. (1979): "Why Do Koreans Save 'So Little'?", *Journal of Development Economics*, 6, 343–362.