UNIVERSITY OF MANNHEIM

empirical and quantitative methods graduate school of economic & social sciences

CDSE

Center for Doctoral Studies in Economics
University of Mannheim

Discussion Paper No. 71

Repeated Moral Hazard with History-Dependent Preferences

Sebastian Koehne*

August 2009

*Sebastian Koehne, Department of Economics, University of Mannheim, 68131 Mannheim, Germany, skoehne@mail.uni-mannheim.de

Financial support from the Deutsche Forschungsgemeinschaft through Graduiertenkolleg Risk and Liquidity in Financial, Commodity, and Factor Markets is gratefully acknowledged.

Graduate School of Economic and Social Sciences GESS · University of Mannheim · gess.uni-mannheim.de
CDSE · University of Mannheim · cdse.uni-mannheim.de

Director: Prof. Dr. Ernst-Ludwig von Thadden · Department of Economics · University of Mannheim · D-68131 Mannheim,
Phone: +49(0621)1812034 · Fax: +49(0621)1812042

Repeated Moral Hazard with History-Dependent Preferences

Sebastian Koehne*

University of Mannheim

This Version: August 4, 2009

(First Version: January 29, 2008)

Abstract

This paper introduces history-dependent preferences into the principal-agent framework. In this setup, the Inverse Euler equation breaks down. The paper characterizes optimal contracts and shows that the wedge between the principal's rate of return to saving and the agent's shadow rate of return is *not* positive in general. However, the wedge is positive given a rather weak assumption on the agent's marginal rate of intertemporal substitution. This points out an intimate link between the sign of the wedge and regressive/progressive taxation of wealth. Finally, the paper explores differences between the two most common models of history-dependent preferences.

Keywords: repeated moral hazard, principal-agent problem, history-dependent preferences, habit formation, intertemporal wedge, optimal taxation

JEL Classification: D82, E21, H21, J31

*Address: Department of Economics, University of Mannheim, 68131 Mannheim, Germany. Email: skoehne@mail.uni-mannheim.de. An earlier version of this paper circulated under the title "Repeated Moral Hazard under Habit Formation". I would like to thank Ernst-Ludwig von Thadden and Nicola Pavoni for many helpful discussions. I received further helpful comments from Felix Bierbrauer, Patrick Bolton, Mathias Dewatripont, Martin Hellwig, Heiko Karle, Dezsoe Szalay, conference participants at the SFB/TR 15 Summer School 2007 Bronnbach, EBIM Workshop 2007 Bielefeld, ENTER Jamboree 2009 London and seminar participants at the University of Mannheim, University College London and Max Planck Institute Bonn.

1 Introduction

The Inverse Euler equation, initially discovered by Diamond and Mirrlees (1978) and Rogerson (1985), recently generalized by Golosov, Kocherlakota, and Tsyvinski (2003), is arguably one of the most important findings in the dynamic contracting literature. The equation is based on the central observation that the agent's utilities can be intertemporally adjusted without changing the incentives to exert effort. As a consequence, the principal's marginal costs of providing utility today must be equal to the conditional expected marginal costs of providing utility tomorrow. For history-independent preferences, the marginal costs of providing utility are given by the agent's inverse marginal utility of consumption. The Inverse Euler equation is therefore an immediate consequence of the previous observation.

However, the Inverse Euler equation requires a strong assumption regarding the agent's preferences as the above argument shows. The present paper introduces a more general preference class into the principal-agent framework. Following similar ideas in the asset pricing literature, the agent's preferences are allowed to be history-dependent. More precisely, the agent's second-period consumption utility may be a function of both the first and second period level of consumption. Similarly, the agent's second-period effort disutility may depend on both the first and second period level of effort. In this setup, the simple link between the principal's marginal costs of providing utility and the agent's inverse marginal utility of consumption breaks down and thus the Inverse Euler equation does not longer apply. The goal of this paper is to characterize optimal contracts in this environment and study the robustness of the lessons drawn from the Inverse Euler equation. This task is important not only from a theoretical point of view, but also from a more applied perspective, considering that history-independent preferences might be too narrow to describe and predict economic behavior (Helson 1964, Frederick and Loewenstein 1999).

I derive two key properties of optimal contracts for history-dependent preferences. First, by a simple extension of Rogerson's (1985) argument, I show that the marginal costs of increasing the agent's utility are still constant over time. This yields an intertemporal condition that is related to the Inverse Euler equation, but includes effects of current wages on future preferences.

In a second step, I explore how saving distorts the incentive problem. In models with history-independent preferences, the Inverse Euler equation implies that optimal contracts feature a positive *intertemporal wedge*: The principal's rate of return to saving exceeds the shadow rate of return associated with the agent's consumption scheme. Therefore, the agent would choose to save if he had access to the same savings technology as the principal. In the present setup, by contrast, the intertemporal wedge can be positive or negative. This implies that the agent's access to the savings technology must be distorted too. However, it is possible that the crucial problem is that of borrowing, rather than saving.

To understand this result, note how saving changes the agent's future incentives to exert effort. First of all, saving generates a future payoff. Due to concavity, the marginal utility of this payoff is high whenever future consumption is low. By this means, saving increases the relative attractiveness of the states in which the agent is poor. For history-dependent preferences, this wealth effect of saving is complemented by a second effect: Saving reduces the agent's consumption habit. Without imposing any structure on the latter, the effect of saving on the attractiveness of states with low consumption is ambiguous. As a consequence, the intertemporal distortion can go both ways.

However, I show that negative intertemporal wedges are not generic. The intertemporal wedge is positive if the agent's marginal rate of intertemporal substitution is increasing in future consumption. Put differently, this condition states that the agent's willingness to give up future consumption is increasing in the level of future consumption. Hence it is not history-independent utility, but more generally the variation of the marginal rate of intertemporal substitution that yields positive wedges.

In addition to these general insights on intertemporal optimality, I discuss the two most common specifications of history-dependent preferences in more detail. Both specifications assume that the agent's second-period consumption utility is decreasing in first-period consumption. This assumption is often referred to as habit formation. For the additive formulation of habit formation (e.g. Constantinides 1990), the marginal rate of intertemporal substitution increases in future consumption. Therefore, optimal contracts feature a positive intertemporal wedge, analogous to the time-separable model. An important difference compared to the time-separable model is the increased intertemporal slope of optimal wage schemes. For the multiplicative specification of habit formation (e.g. Abel 1990), the intertemporal wedge is not positive in general. If the agent's degree of risk aversion is sufficiently small compared to the importance of habits, then the intertemporal wedge becomes negative in that model.[1]

This paper is not the first to argue that an individual's preferences may be sharply influenced by the personal history. Earlier theoretical works on history-dependent preferences include Ryder and Heal (1973), Abel (1990), Constantinides (1990), and Campbell and Cochrane (1999), among others. Frederick and Loewenstein (1999) review the substantial body of empirical research supporting this approach. For instance, workers' self-reported well-being is often strongly related to recent changes in pay, but not so much to absolute levels of pay (Clark 1999, Grund and Sliwka 2007). More generally, the psychological literature documents that the repetition of any stimulus tends to reduce the perception and response to it (Helson 1964). Interpreting consumption as a stimulus, this suggests that past consumption has an important (negative) influence on current well-being.

Models with history-dependent preferences have contributed to the solution of empirical puzzles related to consumption behavior (excess sensitivity puzzle, excess smoothness puzzle), asset pricing (equity premium puzzle) and the relationship between savings and growth; see Messinis (1999) for a review. The present paper studies the impact of such preferences on principal-agent problems. To the best of my knowledge, this question has not been explored so far.

In a simultaneous paper, Grochulski and Kocherlakota (2008) study a multi-period hidden information problem with preferences that are not time-separable. They consider a dynamically evolving idiosyncratic skill process instead of a moral hazard environment. The basic intertemporal properties of optimal allocations are similar in the two settings, however. The focus of Grochulski and Kocherlakota's work greatly differs from the present one. Their aim is to construct a tax system that decentralizes optimal allocations. They show that taxes on current wealth need to depend on future labor incomes, which contrasts with the results obtained for standard preferences (Kocherlakota 2005). By considering a two-period model, I abstract from this finding. The central contribution of my paper is the precise characterization of the intertemporal wedge. The results provide a close link between positivity of the intertemporal wedge and regressivity of wealth taxes akin to Kocherlakota (2005). In such tax systems, given a skill realization in the first-period, the tax rate on wealth brought into the second period is a

[1] I also investigated how wages depend on the habit parameter under multiplicative habits. However, in contrast to the additive model, the comparative statics were not analytically tractable.

negatively-sloped affine transformation of the agent's marginal rate of intertemporal substitution. My findings imply that if the marginal rate of intertemporal substitution increases with second-period consumption, then the intertemporal wedge will be positive. Hence, regressive tax rates on wealth cannot be associated with negative intertemporal wedges. A similar argument shows that progressive tax rates on wealth cannot be associated with positive intertemporal wedges.

The present paper can also be seen as a complement to the literature on effort persistence, which extends the Rogerson (1985) model by allowing for a *technology* that is not time-separable; see Mukoyama and Sahin (2005), Kwon (2006), Jarque (2008), and Hopenhayn and Jarque (2009). In that setup, the Inverse Euler equation is not affected. Hence, effort persistence and history-dependent preferences have fundamentally different effects on the shape of optimal contracts. The reason is that effort persistence alters the information structure of the model, but leaves the preference structure unchanged.

The paper proceeds as follows: Section 2 describes the setup of the model. Section 3 provides a general characterization of intertemporal optimality. Sections 4.1 and 4.2 explore two specific forms of history-dependent preferences. Section 5 concludes.

2 Model

I study a two-period moral hazard problem in which the agent's preferences are history-dependent: Second-period consumption utility may depend on consumption in both periods. Similarly, second-period effort disutility may depend on effort in both periods. This setup generalizes the time-separable model introduced by Rogerson (1985).

2.1 Preferences

The relationship between principal (P) and agent (A) lasts for two periods. The principal maximizes expected profits. She has access to a linear savings technology at the rate R_2, thus her discount factor equals $1/R_2$.

The agent has von-Neumann-Morgenstern preferences and maximizes the expected value of

$$u_1(c_1) - v_1(e_1) + \beta(u_2(c_2, c_1) - v_2(e_2, e_1)),$$

where c_t denotes consumption, e_t represents effort, and $\beta \in (0,1]$ is the discount factor. The functions u_1, u_2, v_1, v_2 are twice continuously differentiable. Effort disutility v_t is strictly increasing and weakly convex in e_t. Consumption utility u_t is strictly increasing and strictly concave in c_t. In addition, I impose two restrictions on the relationship between u_2 and c_1. First, A's lifetime consumption utility is strictly increasing in c_1: $u_1' + \beta \partial_{c_1} u_2 > 0$.[2] Second, to rule out boundary solutions, I suppose $\lim_{c_1 \to 0} (u'(c_1) + \beta \partial_{c_1} u_2(c_2, c_1)) = \infty$ for all $c_2 > 0$. The agent cannot save or borrow, hence his consumption equals his wage in each period.

I call A's preferences **history-independent** (or time-separable) if u_2 does not depend on c_1 and v_2 does not depend on e_1. Otherwise, I call A's preferences **history-dependent**. In the latter case, I will refer to c_1 and e_1 as A's second-period **habit level** of consumption and effort, respectively.

2.2 Technology

In each period, A exerts a hidden work effort and thereby generates a publicly observable stochastic output. Output realizations for period 1 are denoted $x_1(i)$, $i = 1, \ldots, N$, with associated probabilities $\pi_i(e_1) > 0$. Second-period output realizations are denoted $x_2(j)$, $j = 1, \ldots, N$, with associated probabilities $\pi_j(e_2) > 0$.[3]

2.3 Contracts

A **contract** is a specification (w, e) of wages $w = (w_1(i), w_2(i,j))_{i,j}$ and effort levels $e = (e_1, e_2(i))_i$. Here, e_1 is the recommended effort for period 1, $e_2(i)$ the recommended effort for period 2 given that $x_1(i)$ has realized in period 1. Similarly, $w_1(i)$ is the wage paid in period 1, $w_2(i,j)$ is the wage paid in period 2, given that outputs $(x_1(i), x_2(j))$ have realized.

P offers a contract at the beginning of period 1. She has to respect A's participation constraint

$$\sum_{i,j} \pi_i(e_1)\pi_j(e_2(i)) \Big(u_1(w_1(i)) - v_1(e_1) + \beta \left[u_2(w_2(i,j), w_1(i)) - v_2(e_2(i), e_1) \right] \Big) \geq \underline{U}. \tag{PC}$$

In addition, since effort is not observed by P, contracts must satisfy the incentive compatibility

[2] Of course, since u_2 is strictly increasing in c_2, A's lifetime utility is also strictly increasing in c_2.

[3] I could easily allow for a production technology that is not separable over time, but this would not add any interesting insights here.

constraint

$$e \in \underset{(e_1', e_2'(i))}{\operatorname{argmax}} \sum_{i,j} \pi_i(e_1')\pi_j(e_2'(i)) \Big(u_1(w_1(i)) - v_1(e_1') + \beta \left[u_2(w_2(i,j), w_1(i)) - v_2(e_2'(i), e_1') \right] \Big). \tag{IC}$$

A contract (w, e) is called **optimal** if it maximizes P's expected profits

$$\sum_{i,j} \pi_i(e_1)\pi_j(e_2(i)) \Big(x_1(i) - w_1(i) + \frac{1}{R_2} \left[x_2(j) - w_2(i,j) \right] \Big)$$

subject to the incentive compatibility constraint (IC) and the participation constraint (PC).

3 General results on intertemporal optimality

To simplify notation, I denote expectations conditional on first-period information $x_1(i)$ by $E_i[\,\cdot\,]$ in this section. For example, I write

$$\sum_j \frac{\pi_j(e_2(i))}{\partial_{c_2} u_2(w_2(i,j), w_1(i))} = E_i\left[\frac{1}{\partial_{c_2} u_2}\right],$$

and so on.

3.1 Modification of the Inverse Euler equation

The principal can divide wages over time in the following way: The total reward for good performance in period 1 need not be given immediately, but may also be included in second-period wages (which makes the second-period wages dependent on the first-period output). If at least one of the parties is risk-averse, as in the model considered here, then it is clearly beneficial to make use of this option.

This reasoning yields the following intertemporal condition.

Proposition 1. *Let (w, e) be an optimal contract. Then for all $i = 1, \ldots, N$*

$$\frac{1}{\beta R_2} E_i\left[\frac{1}{\partial_{c_2} u_2}\right] = \frac{1}{u_1'(w_1(i))} \left(1 - \frac{1}{R_2} E_i\left[\frac{\partial_{c_1} u_2}{\partial_{c_2} u_2}\right]\right). \tag{1}$$

Proof. The proof is a straightforward extension of the argument in Rogerson (1985). Let $h_1(\,\cdot\,)$

be the inverse of $u_1(\,\cdot\,)$ and $h_2(\,\cdot\,, c_1)$ be the inverse of $u_2(\,\cdot\,, c_1)$, c_1 fixed. For $\epsilon \in \mathbb{R}$, construct a wage scheme w^ϵ from w as follows. Fix an arbitrary first-period output $x_1(i)$. For this output, increase date 1 utility by ϵ and reduce date 2 utility by ϵ/β. Formally, set

$$w_1^\epsilon(k) := \begin{cases} w_1(k) & \text{for } k \neq i, \\ h_1(u_1(w_1(i)) + \epsilon) & \text{for } k = i, \end{cases} \tag{2}$$

$$w_2^\epsilon(k,j) := \begin{cases} w_2(k,j) & \text{for } k \neq i, \\ h_2(u_2(w_2(i,j), w_1(i)) - \epsilon/\beta, w_1^\epsilon(i)) & \text{for } k = i. \end{cases} \tag{3}$$

Notice that the incentive and participation constraints under w^ϵ and $w = w^0$ coincide. Hence, a necessary condition for optimality of (w, e) is that expected wage payments are minimal at $\epsilon = 0$. That is, $\epsilon = 0$ must minimize

$$h_1(u_1(w_1(i)) + \epsilon) + \frac{1}{R_2} E_i\left[h_2(u_2(w_2(i,j), w_1(i)) - \epsilon/\beta, w_1^\epsilon(i))\right], \tag{4}$$

which implies the first-order condition

$$0 = \frac{1}{u_1'(w_1(i))} + \frac{1}{R_2} E_i\left[-\frac{1}{\beta\,\partial_{c_2} u_2(w_2(i,j), w_1(i))} - \frac{\partial_{c_1} u_2(w_2(i,j), w_1(i))}{u_1'(w_1(i))\,\partial_{c_2} u_2(w_2(i,j), w_1(i))}\right]. \tag{5}$$

Equivalently,

$$\frac{1}{\beta R_2} E_i\left[\frac{1}{\partial_{c_2} u_2(w_2(i,j), w_1(i))}\right] = \frac{1}{u_1'(w_1(i))}\left(1 - \frac{1}{R_2} E_i\left[\frac{\partial_{c_1} u_2(w_2(i,j), w_1(i))}{\partial_{c_2} u_2(w_2(i,j), w_1(i))}\right]\right). \tag{6}$$

□

Proposition 1 states that P's marginal cost of increasing A's utility conditional on output $x_1(i)$ is constant over time: The cost of increasing u_2 by ϵ/β is approximately equal to

$$\frac{\epsilon}{\beta R_2} E_i\Big[\frac{1}{\partial_{c_2} u_2}\Big]. \tag{7}$$

The cost of increasing u_1 by ϵ while keeping u_2 constant is approximately equal to

$$\frac{\epsilon}{u_1'} - \frac{\epsilon}{R_2 u_1'} E_i\Big[\frac{\partial_{c_1} u_2}{\partial_{c_2} u_2}\Big], \tag{8}$$

where the first term denotes the increase of period 1 consumption required to raise u_1 by ϵ, and the second term represents the increase of period 2 consumption needed to keep u_2 at the former level. Equation (1) shows that at any optimal contract these two costs are the same.

If preferences are history-independent, with $u_1 = u_2 = u$, then equation (1) collapses to

$$\frac{1}{u'(w_1(i))} = \frac{1}{\beta R_2} E_i \left[\frac{1}{u'(w_2(i,j))} \right]. \tag{9}$$

This well-known result from Rogerson (1985) is often called the Inverse Euler equation. The relation between inverse marginal utilities described in equation (1) is slightly different, thus we cannot attach the same label here. Yet, equations (1) and (9) have exactly the same interpretation: For any optimal contract, the principal's marginal costs of providing utility conditional on a given output in period 1 must be equal across periods.

3.2 The intertemporal wedge

In dynamic models with hidden effort choices, there is typically an intertemporal distortion due to the influence of savings on the incentive problem. The intertemporal distortion in models with history-independent preferences is well understood. Diamond and Mirrlees (1978), Rogerson (1985) and, in a more general framework, Golosov, Kocherlakota, and Tsyvinski (2003) show that there is a positive intertemporal wedge in such models: At optimal contracts, the principal's rate of return to saving exceeds the shadow rate of return associated with the agent's consumption scheme.[4]

This subsection characterizes the intertemporal wedge for preferences that are history-dependent. Recall that $E_i[\,\cdot\,]$ denotes expectations conditional on first-period information $x_1(i)$, i.e., with respect to probability weights $\pi_j(e_2(i))$. Similarly, $\mathrm{cov}_i(\,\cdot\,,\,\cdot\,)$ represents covariances with respect to $\pi_j(e_2(i))$.

Consider the difference between P's rate of return to saving and A's shadow rate of return,

$$R_2 - \frac{u_1'(w_1(i)) + \beta E_i\left[\partial_{c_1} u_2(w_2(i,j), w_1(i))\right]}{\beta E_i\left[\partial_{c_2} u_2(w_2(i,j), w_1(i))\right]}. \tag{10}$$

It is convenient to divide this difference by P's rate of return, R_2, and define the **intertemporal**

[4]The shadow rate of return is defined as the rate at which the agent does not wish to save or borrow.

wedge (given output $x_1(i)$) as

$$IW_i := 1 - \frac{u_1'(w_1(i)) + \beta E_i\left[\partial_{c_1} u_2(w_2(i,j), w_1(i))\right]}{\beta R_2 E_i\left[\partial_{c_2} u_2(w_2(i,j), w_1(i))\right]}. \tag{11}$$

The intertemporal wedge is characterized as follows.

Proposition 2. *Consider an optimal contract. Then, conditional on first-period output $x_1(i)$, the intertemporal wedge equals*

$$IW_i = 1 - \frac{R_2 - \operatorname{cov}_i\left(\partial_{c_1} u_2,\, (\partial_{c_2} u_2)^{-1}\right)}{R_2 - R_2 \operatorname{cov}_i\left(\partial_{c_2} u_2,\, (\partial_{c_2} u_2)^{-1}\right)}. \tag{12}$$

Hence, the intertemporal wedge is positive if and only if

$$\operatorname{cov}_i\left(R_2 \partial_{c_2} u_2 - \partial_{c_1} u_2,\, \frac{1}{\partial_{c_2} u_2}\right) < 0. \tag{13}$$

Proof. Solve the first-order condition (1) for u_1' to get

$$u_1' = \beta \frac{R_2 - E_i[\partial_{c_1} u_2\, (\partial_{c_2} u_2)^{-1}]}{E_i[(\partial_{c_2} u_2)^{-1}]}. \tag{14}$$

Using

$$E_i[\partial_{c_1} u_2\, (\partial_{c_2} u_2)^{-1}] = \operatorname{cov}_i\left(\partial_{c_1} u_2,\, (\partial_{c_2} u_2)^{-1}\right) + E_i[\partial_{c_1} u_2] E_i[(\partial_{c_2} u_2)^{-1}], \tag{15}$$

we find

$$u_1' = \beta \frac{R_2 - \operatorname{cov}_i\left(\partial_{c_1} u_2,\, (\partial_{c_2} u_2)^{-1}\right)}{E_i[(\partial_{c_2} u_2)^{-1}]} - \beta E_i[\partial_{c_1} u_2]. \tag{16}$$

Substitute this into definition (11) to obtain

$$IW_i = 1 - \frac{R_2 - \operatorname{cov}_i\left(\partial_{c_1} u_2,\, (\partial_{c_2} u_2)^{-1}\right)}{R_2 E_i\left[\partial_{c_2} u_2\right] E_i[(\partial_{c_2} u_2)^{-1}]}. \tag{17}$$

Now the result follows from

$$E_i[\partial_{c_2} u_2] E_i[(\partial_{c_2} u_2)^{-1}] = 1 - \operatorname{cov}_i\left(\partial_{c_2} u_2,\, (\partial_{c_2} u_2)^{-1}\right). \tag{18}$$

□

Proposition 2 proves that the sign of the intertemporal wedge coincides with the sign of the negated covariance

$$-\operatorname{cov}_i\left(R_2\partial_{c_2}u_2 - \partial_{c_1}u_2,\ \frac{1}{\partial_{c_2}u_2}\right). \tag{19}$$

Intuitively, this expression captures the hypothetical impact of saving on incentives: The first variable in the covariance, $R_2\partial_{c_2}u_2 - \partial_{c_1}u_2$, is A's marginal benefit of saving. This benefit consists of the marginal payoff of saving, $R_2\partial_{c_2}u_2$, plus the marginal utility of reducing the habit level, $-\partial_{c_1}u_2$.[5] The second variable, $(\partial_{c_2}u_2)^{-1}$, is monotonic in A's second-period wage. If these two variables covary negatively, then saving tends to increase A's utility in the states with low wages relative to the states with high wages. In other words, saving "hedges" against incentives, or "tightens" the incentive compatibility constraint.

Simply put, Proposition 2 thus shows that the intertemporal wedge is positive if and only if saving "hedges" against the incentive scheme. This insight extends the intuition from Golosov, Kocherlakota, and Tsyvinski (2003) to the case of history-dependent preferences. Their explanation for the emergence of a positive intertemporal wedge, given history-independent preferences, is that saving has an adverse effect on incentives (Golosov, Kocherlakota, and Tsyvinski 2003, p.577). For this reason, the social cost of saving is higher than the agent's individual cost. It is optimal to equate the social marginal cost and benefit of saving, but not the agent's individual marginal cost and benefit. Therefore, at optimal contracts, the agent's individual marginal benefit of saving exceeds the agent's individual marginal cost.

In the present setup, the intertemporal wedge can be negative as the next proposition shows. However, a rather weak assumption guarantees that the wedge is positive. To formulate the assumption, define the **marginal rate of intertemporal substitution** as

$$MRS(c_1, c_2) := \frac{u_1'(c_1) + \beta\partial_{c_1}u_2(c_2, c_1)}{\beta\partial_{c_2}u_2(c_2, c_1)}. \tag{20}$$

Proposition 3. *Let (w, e) be optimal. Let the first-period output be $x_1(i)$ and suppose $w_2(i, j) \neq w_2(i, j')$ for some j, j'.*[6] *Then we have the following. (i) The intertemporal wedge is not positive in general. (ii) The intertemporal wedge is positive if A's marginal rate of intertemporal*

[5] The agent's marginal cost of saving, u_1', does not show up here, because it is independent of the realization of second-period output.

[6] The case with $w_2(i, j) = w_2(i, j')$ for all j, j' is not interesting. It corresponds to the implementation of minimal effort.

substitution is increasing in second-period consumption.

Proof. (i) See Example 2 on page 19.

(ii) The intertemporal wedge is positive if and only if

$$\beta R_2 E_i\left[\partial_{c_2} u_2\right] - \left(u_1' + \beta E_i\left[\partial_{c_1} u_2\right]\right) > 0. \tag{21}$$

This is the case if and only if

$$E_i\left[\partial_{c_2} u_2 \left(R_2 - \frac{u_1' + \beta\partial_{c_1} u_2}{\beta\partial_{c_2} u_2}\right)\right] > 0. \tag{22}$$

Proposition 1 implies that the random variable

$$R_2 - \frac{u_1' + \beta\partial_{c_1} u_2}{\beta\partial_{c_2} u_2} \tag{23}$$

is centered. Hence, equation (22) can be rewritten as

$$\mathrm{cov}_i\left(\partial_{c_2} u_2,\, R_2 - \frac{u_1' + \beta\partial_{c_1} u_2}{\beta\partial_{c_2} u_2}\right) > 0. \tag{24}$$

This is equivalent to

$$\mathrm{cov}_i\left(\partial_{c_2} u_2,\, \frac{u_1' + \beta\partial_{c_1} u_2}{\beta\partial_{c_2} u_2}\right) < 0. \tag{25}$$

By concavity of u_2, the partial derivative $\partial_{c_2} u_2$ is decreasing in second-period consumption. By assumption, the marginal rate of intertemporal substitution is increasing in second-period consumption. This implies that condition (25) is satisfied. □

To understand the first part of Proposition 3, suppose that consumption and habit are substitutes for A's second-period utility. That is, suppose $\partial_{c_1}\partial_{c_2} u_2(c_2, c_1) < 0$ for all c_1, c_2. In this case, the hypothetical "hedging value" of saving captured by (19) is ambiguous. On the one hand, the marginal payoff of saving, $R_2\partial_{c_2} u_2$, is due to concavity high whenever the second-period wage is low. On the other hand, the marginal utility of reducing the habit level, $-\partial_{c_1} u_2$, is low whenever the second-period wage is low. Hence, the "hedging value" of saving can be positive or negative, depending on which effect dominates. Therefore the intertemporal wedge can go both ways.

The second part of Proposition 3 shows that the above case is not generic. Under the assumption that A's willingness to give up future consumption for current consumption increases in the level of future consumption, the intertemporal wedge is positive. This identifies the key factor in Rogerson's (1985) result. It is not history-independent utility, but the variation of the marginal rate of intertemporal substitution that matters for the sign of the intertemporal wedge.

The next section studies two specific forms of history-dependent preferences. For additive habit formation, the marginal rate of intertemporal substitution increases in future consumption. For multiplicative habits, consumption and habit can be substitutes and the latter property does not generally hold. As a consequence, the intertemporal wedge may be negative in that setup.

4 Additive and multiplicative habit formation

In this section, I explore the two most common models of history-dependent consumption preferences: additive and multiplicative habit formation.[7] Among many others, Constantinides (1990), Lahiri and Puhakka (1998), and Campbell and Cochrane (1999) follow the first approach. Prominent examples of the multiplicative approach are Abel (1990) and Carroll, Overland, and Weil (1997, 2000).

Both specifications impose two additional conditions to the agent's preferences: Utility in the second period is decreasing in first-period consumption. Secondly, the marginal rate of intertemporal substitution decreases as habits become more important.[8] Notice, however, that the results derived in the previous section did not rely on any of these two properties.

[7]As in the previous section, the form of the effort disutility functions does not play a role for the questions addressed here.

[8]For multiplicative habits, there is a range of the consumption space in which this property is violated (Wendner 2003). I will only consider cases where the contracts do not fall into this range.

4.1 Additive habits

Let $u : \mathbb{R}_+ \to \mathbb{R}$ be twice continuously differentiable, with $u' > 0$, $u'' < 0$, $\lim_{c \to 0} u'(c) = \infty$. Set

$$\begin{aligned} u_1(c_1) &: = u(c_1), \\ u_2(c_2, c_1) &: = u(\tilde{c}_2), \quad \tilde{c}_2 := c_2 - \gamma c_1. \end{aligned} \tag{26}$$

Equivalently, we can write $\tilde{c}_2 = (1-\gamma)c_2 + \gamma(c_2 - c_1)$. Hence, **effective consumption** in period 2 is a weighted average of absolute consumption, c_2, and absolute consumption minus the habit level, $c_2 - c_1$. The parameter $\gamma \in [0, 1]$ controls the importance of habits: The higher the value of γ, the more the agent cares about how period 2 consumption relates to period 1 consumption.

The first important finding is that Rogerson's (1985) insight on intertemporal wedges extends to this setup.

Proposition 4. *Let A's preferences be given by* (26)*, let* (w, e) *be an optimal contract and suppose that for each* i *there exist* j, j' *such that* $w_2(i, j) \neq w_2(i, j')$*. Then the intertemporal wedge is positive.*

Proof. Note $MRS(c_1, c_2) = u'(c_1)/(\beta u'(c_2 - \gamma c_1)) - \gamma$. This shows that $MRS(c_1, c_2)$ is increasing in c_2. Hence, the intertemporal wedge is positive by Proposition 3. □

The remainder of this subsection discusses the comparative statics of optimal wage profiles in the habit parameter. For history-independent preferences, the intertemporal slope of wages is determined by the curvature of inverse marginal utility as Rogerson (1985) has shown. If inverse marginal utility is convex, then expected wages are nonincreasing over time. This convexity requirement is satisfied in many generic cases, for instance for CRRA utility with a coefficient of relative risk aversion larger than one.

If the agent forms habits, then the intertemporal slope of wages does not only depend on the curvature of inverse marginal utility, but also on the importance of habits. This tends to shift wages to the end of the relationship as we will see now.

For logarithmic utility, the functional relation between the slope of wages and the habit parameter γ is particularly simple.

Example 1. Let $u(c) = \log(c)$, $\beta R_2 = 1$. Then equation (1), the condition for intertemporal optimality of rewards, takes the form

$$E_i\left[\frac{1}{u'(\tilde{w}_2(i,j))}\right] = \frac{1}{u'(w_1(i))}(1+\beta\gamma), \tag{27}$$

which implies

$$\frac{E_i[w_2(i,j)]}{w_1(i)} = 1 + (1+\beta)\gamma. \tag{28}$$

Thus, expected wages are constant over time if the agent does not form habits ($\gamma = 0$). If he does form habits ($\gamma > 0$), then expected wages are increasing over time. Moreover, the ratio between expected wages paid in period 2 and period 1 is increasing in the habit parameter γ.

In the above example, expected wages are increasing over time for any positive habit parameter. More generally, no matter what utility function is chosen, expected wages will increase over time if the habit parameter is sufficiently large, as the following proposition shows.

Proposition 5. *Let $\gamma = 1$. If (w, e) is an optimal contract, then $w_2(i,j) > w_1(i)$ for all i, j.*

Proof. For $\gamma = 1$, second-period utility takes the form $u_2(w_2(i,j), w_1(i)) = u(w_2(i,j) - w_1(i))$. Therefore, the statement follows immediately from the assumption $\lim_{c\to 0} u'(c) = \infty$. $\square$

The ratio between expected wages at date 2 and date 1 is a monotonic function of the habit parameter γ in the log-utility example. This result can be generalized. However, for non-logarithmic utility, the intertemporal condition (1) cannot be solved for wage levels. The shape of wages can only be studied if one includes the incentive and participation constraint in the analysis. To keep things tractable, I study the simplest possible case: two effort levels, $\{l, h\}$, and two outputs, $\{x_L, x_H\}$, $x_L < x_H$. Probability distributions satisfy $\pi_H(l) < \pi_H(h)$, and effort costs are $v_1(l) = v_2(l) = 0$, $v_1(h) = v_2(h) = v > 0$. To simplify notation, write $w_i = w_1(i)$, $w_{ij} = w_2(i,j)$ and set $v' := v/(\pi_H(h) - \pi_H(l))$.

Suppose P wants to implement effort h in both periods (i.e., $x_H - x_L$ is sufficiently large).

Then optimal contracts are characterized as follows:

$$\frac{1+\gamma/R_2}{u'(w_H)} - \pi_H(h)\frac{1}{\beta R_2 u'(w_{HH}-\gamma w_H)} - (1-\pi_H(h))\frac{1}{\beta R_2 u'(w_{HL}-\gamma w_H)} = 0 \tag{29}$$

$$\frac{1+\gamma/R_2}{u'(w_L)} - \pi_H(h)\frac{1}{\beta R_2 u'(w_{LH}-\gamma w_L)} - (1-\pi_H(h))\frac{1}{\beta R_2 u'(w_{LL}-\gamma w_L)} = 0 \tag{30}$$

$$u(w_{HH}-\gamma w_H) - u(w_{HL}-\gamma w_H) - v' = 0 \tag{31}$$

$$u(w_{LH}-\gamma w_L) - u(w_{LL}-\gamma w_L) - v' = 0 \tag{32}$$

$$u(w_H) + u(w_{HL}-\gamma w_H) + (1+\beta)v'\pi_H(l) - v' - \underline{U} = 0 \tag{33}$$

$$u(w_L) + u(w_{LL}-\gamma w_L) + (1+\beta)v'\pi_H(l) - \underline{U} = 0. \tag{34}$$

Here, (29),(30) are the intertemporal optimality conditions, (31),(32) ensure incentive compatibility of second-period effort, (33),(34) ensure incentive compatibility of first-period effort, and together with (31),(32) also imply that the participation constraint is satisfied.

This system of six equations can be separated into two subsystems with three equations each—one system determining w_H, w_{HL}, w_{HH}, and one determining w_L, w_{LL}, w_{LH}. The comparative statics in the habit parameter γ are as follows.

Proposition 6. *Let (w, e) be the optimal contract in the two-effort two-output model described above. Then for each i, j, the ratio $w_2(i,j)/w_1(i)$ is increasing in γ.*

Proof. Let $i = H$. (The case $i = L$ is exactly analogous.) Applying the implicit function theorem to the system of equations (29),(31),(33) yields

$$\begin{aligned}
\frac{\partial w_H}{\partial \gamma} &= -\frac{\beta^2}{D} u'(w_H)u'(\tilde{w}_{HL})^3 u'(\tilde{w}_{HH})^3 \\
\frac{\partial w_{HH}}{\partial \gamma} &= w_H + \frac{\beta}{D} u'(w_H)u'(\tilde{w}_{HL})^3 u'(\tilde{w}_{HH})^2 [u'(w_H) - \gamma\beta u'(\tilde{w}_{HH})] \\
\frac{\partial w_{HL}}{\partial \gamma} &= w_H + \frac{\beta}{D} u'(w_H)u'(\tilde{w}_{HL})^2 u'(\tilde{w}_{HH})^3 [u'(w_H) - \gamma\beta u'(\tilde{w}_{HL})],
\end{aligned}$$

where

$$\begin{aligned}
D = &-(1-p_H(h))u'(w_H)^3 u'(\tilde{w}_{HH})^3 u''(\tilde{w}_{HL}) \\
&- u'(\tilde{w}_{HL})^3[(R_2+\gamma)\beta^2 u'(\tilde{w}_{HH})^3 u''(w_H) + p_H(h)u'(w_H)^3 u''(\tilde{w}_{HH})] \\
&> 0.
\end{aligned}$$

Hence, the expression $\frac{\partial w_{HH}}{\partial \gamma} w_H - w_{HH} \frac{\partial w_H}{\partial \gamma}$ is equal to

$$w_H^2 + \frac{\beta}{D} u'(w_H) u'(\tilde{w}_{HL})^3 u'(\tilde{w}_{HH})^2 \left[w_H u'(w_H) + \beta \tilde{w}_{HH} u'(\tilde{w}_{HH}) \right],$$

which is positive. This shows that the ratio w_{HH}/w_H is increasing in γ. Monotonicity of w_{HL}/w_H can be seen analogously. $\square$

The intuition underlying Proposition 6 is straightforward. If the habit parameter γ increases, A demands a higher compensation for the "comparison effects" generated by date 1 wages. This makes wages paid at date 1 more costly relative to wages paid at date 2, hence P substitutes.

4.2 Multiplicative habits

Let $u : \mathbb{R}_+ \to \mathbb{R}$ be twice continuously differentiable, with $u' > 0$, $u'' < 0$, $\lim_{c \to 0} u'(c) = \infty$. Let $\gamma \in [0, 1]$, $b > 0$. Set

$$\begin{aligned} u_1(c_1) :&= u(c_1), \\ u_2(c_2, c_1) :&= u(\hat{c}_2), \quad \hat{c}_2 := c_2/(b + c_1)^\gamma. \end{aligned} \tag{35}$$

Effective consumption in period 2 can be rewritten $\hat{c}_2 = c_2^{1-\gamma}(c_2/(b + c_1))^\gamma$. That is, effective consumption is a Cobb-Douglas aggregate of absolute consumption, c_2, and absolute consumption relative to the habit level, $c_2/(b+c_1)$. The parameter γ is the weight attached to the latter term. For an agent with a higher value of γ, the comparison between date 2 consumption and date 1 consumption thus gets more important in this sense.

The introduction of $b > 0$ in the habit term is a technical necessity for this specification. Note that for $b = 0$ one obtains

$$u'(c_1) + \beta \partial_{c_1} u_2(c_2, c_1) = u'(c_1) - \beta \gamma c_1^{-\gamma - 1} c_2 u'(\hat{c}_2). \tag{36}$$

Hence, in this case the assumption $u'(0) = \infty$ is no longer sufficient for the condition

$$\forall c_2 > 0 \quad \lim_{c_1 \to 0} \left(u'(c_1) + \beta \partial_{c_1} u_2(c_2, c_1) \right) = \infty, \tag{37}$$

which is needed to make sure that solutions are interior. For $b > 0$, this problem does not arise.

The next proposition is the main result of this subsection. It shows that multiplicative habits change the key insight of Rogerson's (1985) time-separable model. The intertemporal wedge is not generally positive.

Proposition 7. *Let A's preferences be given by* (35)*, let* (w, e) *be an optimal contract and suppose that for each* i *there exist* j, j' *such that* $w_2(i, j) \neq w_2(i, j')$*. Then we have the following.*
(i) The intertemporal wedge is not positive in general.
(ii) The intertemporal wedge is positive if A's coefficient of relative risk-aversion, $-\hat{c}_2 \frac{u''(\hat{c}_2)}{u'(\hat{c}_2)}$*, is bounded below by 1.*

Proof. (i) See Example 2 on page 19.
(ii) The assumption $-\hat{c}_2 \frac{u''(\hat{c}_2)}{u'(\hat{c}_2)} \geq 1$ is equivalent to

$$- u'(c_2(b + c_1)^{-\gamma}) - c_2(b + c_1)^{-\gamma} u''(c_2(b + c_1)^{-\gamma}) \geq 0. \tag{38}$$

Multipliying this by $\gamma(b + c_1)^{-1-\gamma}$, we have

$$- \gamma(b + c_1)^{-1-\gamma} u'(c_2(b + c_1)^{-\gamma}) - c_2\gamma(b + c_1)^{-1-2\gamma} u''(c_2(b + c_1)^{-\gamma}) \geq 0, \tag{39}$$

which is equivalent to $\partial_{c_1}\partial_{c_2}u_2 \geq 0$. In combination with $\partial_{c_2}\partial_{c_2}u_2 < 0$, this implies that $MRS(c_1, c_2)$ is increasing in c_2. Hence, the intertemporal wedge is positive by Proposition 3. $\square$

The reason why the intertemporal wedge is not generally positive has already been pointed out at the end of Section 3.2: While A's marginal benefit of saving covaries negatively with future wages in the time-separable model, this is no longer the case for the model with multiplicative habits. If consumption and habit are substitutes, which happens exactly when the coefficient of relative risk aversion is smaller than one, then the covariance may be positive. In this case, saving would strengthen the link between A's utility and future wages, opposite to the time-separable model.

To prove the first part of Proposition 7, I now study an example in which the coefficient of relative risk aversion is smaller than one. Note that for multiplicative habits optimal contracts can only be determined numerically (except for logarithmic utility).

Example 2. Consider the two-effort two-output problem described in the previous subsection. The optimal contract can be easily characterized by a system of equations analogous to the system (29)–(34) for additive habits. Assume $R_2 = \beta = 1$, $u(c) = \frac{1}{1-\rho}c^{1-\rho}$, $\rho = 0.4$, $b = 0.3$, $v = 0.1$, $\pi_H(l) = 0.25$, $\pi_H(h) = 0.5$, $\underline{U} = 2$.

Table 1 depicts optimal wages for different values of the habit parameter γ. Figures 1 and 2 display wages graphically. We see that first-period wages $w_1(i)$ decrease in the habit parameter γ, whereas second-period effective wages $\hat{w}_2(i,j)$ increase in γ.

Figure 3 shows the intertemporal wedge for each output realization of period 1, $x_1(i) \in \{x_L, x_H\}$. For $\gamma = 0$, preferences are history-independent and thus the intertemporal wedge is positive, of course. The wedge declines in γ, and eventually becomes negative. Hence, there is a cutoff value of γ, below of which the intertemporal wedge is positive for both first-period output realizations, and above of which the intertemporal wedge is negative for at least one output realization. Define γ^* as this cutoff value. In summary, we have the following.

Observation 1. The intertemporal wedge decreases in γ and can become negative (Figure 3).

The intuition is as follows. Since $\rho < 1$, the marginal habit effect of saving, $-\partial_{c_1} u_2 = \gamma(b + c_1)^{\gamma(\rho-1)-1} c_2^{1-\rho}$, is increasing in second-period consumption c_2. If γ is sufficiently large, this will outweigh the negative relation between the marginal payoff of saving, $R_2 \partial_{c_2} u_2 = (b + c_1)^{\gamma(\rho-1)} c_2^{-\rho}$, and second-period consumption. In that case, A's marginal benefit of saving, $R_2 \partial_{c_2} u_2 - \partial_{c_1} u_2$, will covary positively with second-period wages, hence the hypothetical "hedging value" of saving will be negative.

The size of the covariance between the marginal payoff of saving and second-period wages depends crucially on A's risk aversion. The above discussion points out that the habit effect of saving will dominate the wealth effect if the coefficient ρ of relative risk aversion is small. This suggests that for small values of ρ, a relatively small habit parameter γ will suffice to make the intertemporal wedge negative. Figure 4 shows the cutoff value γ^*, above of which the intertemporal wedge is negative for at least one realization of first-period output, for varying coefficients ρ of relative risk aversion. Indeed we see the following.

Observation 2. The cutoff value γ^* decreases as ρ gets smaller. Moreover, γ^* approaches zero as ρ goes to zero. (Figure 4)

γ	$w_1(H)$	$\hat{w}_2(H,H)$	$\hat{w}_2(H,L)$	$w_1(L)$	$\hat{w}_2(L,H)$	$\hat{w}_2(L,L)$
0	0.57615	0.75031	0.42881	0.42444	0.58109	0.29632
0.1	0.50713	0.83061	0.49353	0.36826	0.63594	0.34152
0.2	0.40427	0.96563	0.60437	0.29368	0.72352	0.41315
0.3	0.24524	1.22542	0.82304	0.19830	0.87988	0.53872
0.4	0.10551	1.54966	1.10302	0.10978	1.09680	0.71313
0.5	0.05256	1.72965	1.26090	0.06041	1.24748	0.83940
0.6	0.03091	1.82740	1.34723	0.03648	1.33626	0.91567
0.7	0.02001	1.88827	1.40119	0.02380	1.39268	0.96471
0.8	0.01377	1.92978	1.43806	0.01638	1.43138	0.99855
0.9	0.00988	1.95983	1.46479	0.01171	1.45945	1.02320
1	0.00730	1.98254	1.48501	0.00861	1.48068	1.04189

Table 1: Optimal (effective) wages for different values of the habit parameter γ

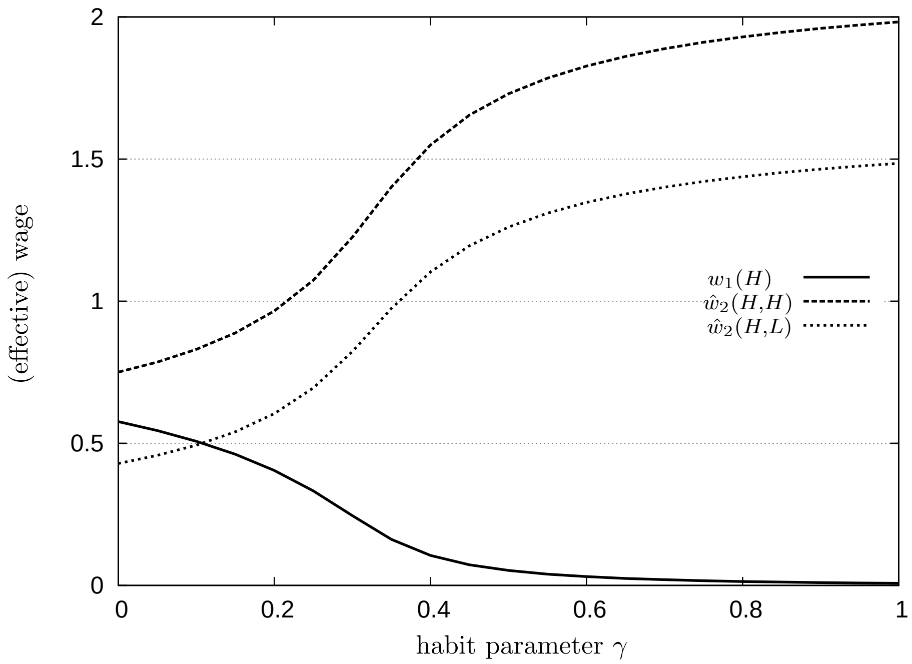

Figure 1: (Effective) wages, given high output in period 1

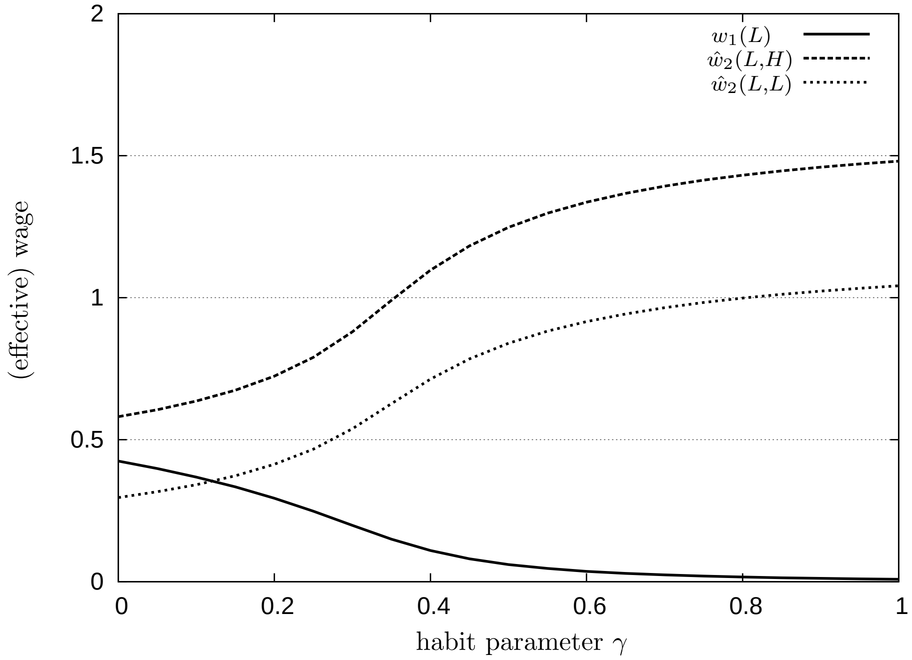

Figure 2: (Effective) wages, given low output in period 1

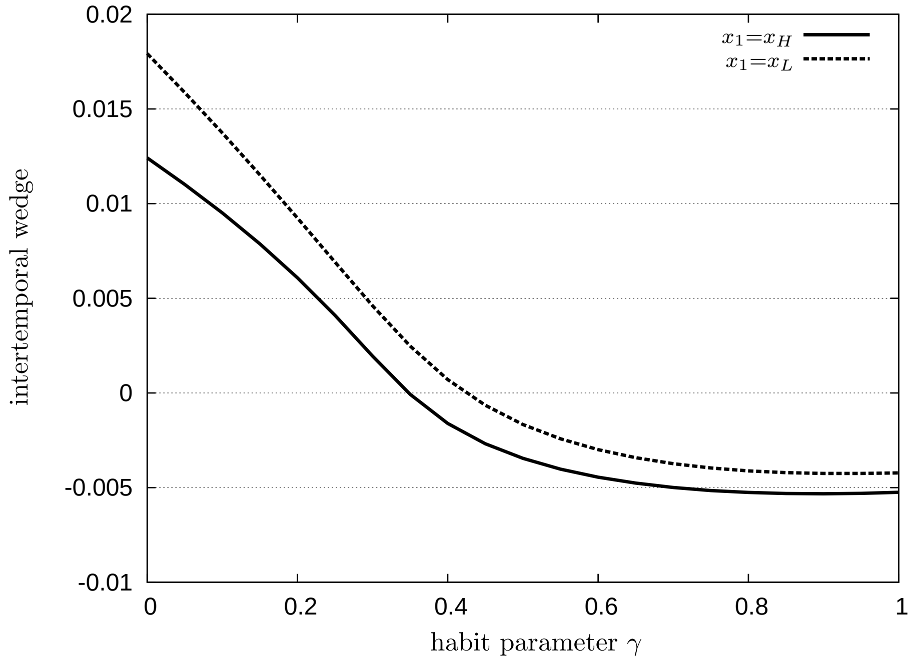

Figure 3: Intertemporal wedge for high and low output in period 1. The wedges are negative if the habit parameter is sufficiently large.

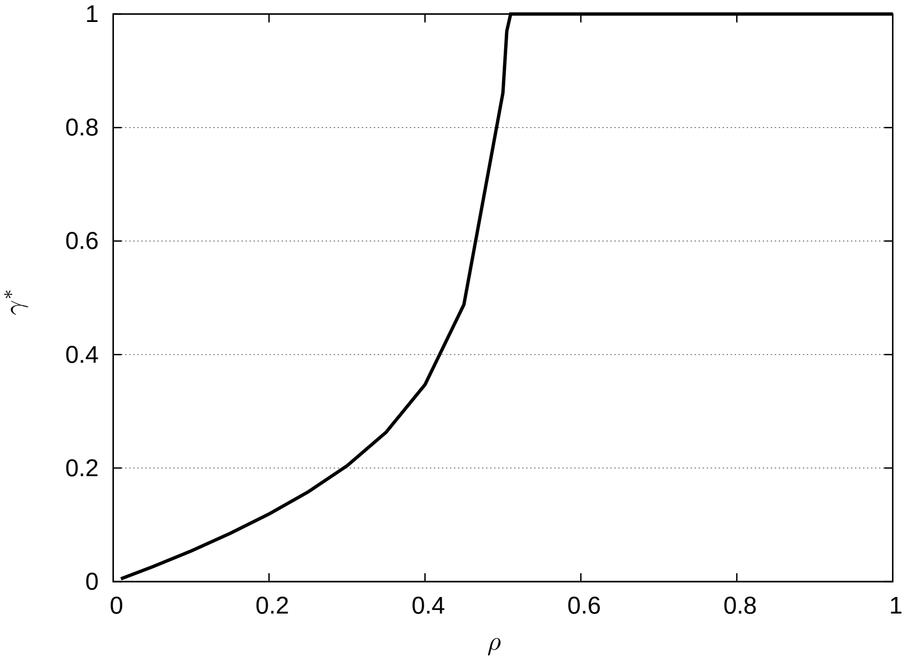

Figure 4: Maximum value of the habit parameter for which both intertemporal wedges are positive. This value is increasing in the coefficient of relative risk aversion, ρ.

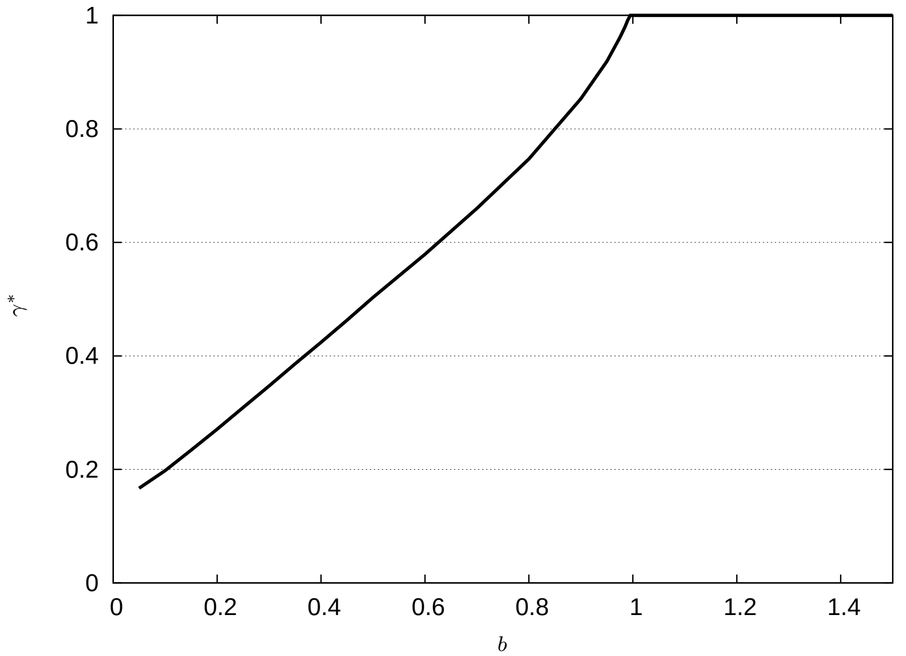

Figure 5: Maximum value of the habit parameter for which both intertemporal wedges are positive. This value is increasing in the intercept of the habit process, b.

In other words, Figure 4 shows that the intertemporal wedge is negative if the habit parameter γ is sufficiently large compared to the coefficient ρ of relative risk aversion. Finally, we note the following.

Observation 3. There exist coefficients $\rho < 1$ with $\gamma^* = 1$ (Figure 4).

Hence, the highest possible habit parameter ($\gamma = 1$) does not for all risk aversion coefficients $\rho < 1$ lead to negative intertemporal wedges. As Figure 5 shows, the set of parameters with negative intertemporal wedges can be increased by decreasing the size of the constant b. This seems due to the fact that b has a dampening effect on changes in the habit level: The larger the size of b, the smaller is the relative change of the habit level $b + c_1$ given a reduction of consumption c_1 by one unit.

5 Concluding remarks

This paper introduces history-dependent preferences into the principal-agent framework. The paper shows how the Inverse Euler equation must be modified to account for effects of current wages on future preferences. Moreover, the paper characterizes the intertemporal wedge. The key condition to obtain a positive wedge is that the marginal rate of intertemporal substitution increases in future consumption. This identifies the driving force of the results by Diamond and Mirrlees (1978), Rogerson (1985), and Golosov, Kocherlakota, and Tsyvinski (2003).

The results in this paper characterize the efficient way to induce a given effort profile. For this question, history-dependence of consumption preferences is the only relevant difference to the standard model. A complementary and also very interesting question is what effort profile the principal wants to implement. In that context, history-dependence of effort preferences will make an important difference. However, the latter question cannot be addressed in the general environment used in this paper. One has to think about concrete applications of moral hazard problems and tailor models to those applications. Moreover, one needs a better understanding of how previous effort choices may influence people's current preferences to work hard. This creates interesting directions for future research.

References

Abel, A. B. (1990): "Asset Prices under Habit Formation and Catching up with the Joneses," *American Economic Review*, 80(2), 38–42.

Campbell, J. 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., J. Overland, and D. N. Weil (1997): "Comparison Utility in a Growth Model," *Journal of Economic Growth*, 2(4), 339–67.

——— (2000): "Saving and Growth with Habit Formation," *American Economic Review*, 90(3), 341–355.

Clark, A. E. (1999): "Are wages habit-forming? evidence from micro data," *Journal of Economic Behavior & Organization*, 39(2), 179–200.

Constantinides, G. M. (1990): "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98(3), 519–43.

Diamond, P. A., and J. A. Mirrlees (1978): "A model of social insurance with variable retirement," *Journal of Public Economics*, 10(3), 295 – 336.

Frederick, S., and G. Loewenstein (1999): "Hedonic Adaptation," in *Well-being: The foundations of hedonic psychology*, ed. by D. Kahneman, E. Diener, and N. Schwarz, pp. 302–329. Russell Sage Foundation Press.

Golosov, M., N. Kocherlakota, and A. Tsyvinski (2003): "Optimal Indirect and Capital Taxation," *Review of Economic Studies*, 70(3), 569–587.

Grochulski, B., and N. Kocherlakota (2008): "Nonseparable Preferences and Optimal Social Security Systems," July 2008. Mimeo. http://www.econ.umn.edu/∼nkocher/grokoch08-07-01-PM.pdf.

Grund, C., and D. Sliwka (2007): "Reference-Dependent Preferences and the Impact of Wage Increases on Job Satisfaction: Theory and Evidence," *Journal of Institutional and Theoretical Economics (JITE)*, 163(2), 313–335.

HELSON, H. (1964): *Adaptation-Level Theory*. Harper & Row New York.

HOPENHAYN, H., AND A. JARQUE (2009): "Unobservable Persistent Productivity and Long Term Contracts," *Review of Economic Dynamics*, forthcoming.

JARQUE, A. (2008): "Repeated Moral Hazard with Effort Persistence," Working Paper 08-04, Federal Reserve Bank of Richmond.

KOCHERLAKOTA, N. R. (2005): "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation," *Econometrica*, 73(5), 1587–1621.

KWON, I. (2006): "Incentives, Wages, and Promotions: Theory and Evidence," *RAND Journal of Economics*, 37(1), 100–120.

LAHIRI, A., AND M. PUHAKKA (1998): "Habit Persistence in Overlapping Generations Economies under Pure Exchange," *Journal of Economic Theory*, 78(1), 176–186.

MESSINIS, G. (1999): "Habit Formation and the Theory of Addiction," *Journal of Economic Surveys*, 13(4), 417–42.

MUKOYAMA, T., AND A. SAHIN (2005): "Repeated moral hazard with persistence," *Economic Theory*, 25(4), 831–854.

ROGERSON, W. P. (1985): "Repeated Moral Hazard," *Econometrica*, 53(1), 69–76.

RYDER, JR., H. E., AND G. M. HEAL (1973): "Optimum Growth with Intertemporally Dependent Preferences," *Review of Economic Studies*, 40(1), 1–33.

WENDNER, R. (2003): "Do habits raise consumption growth?," *Research in Economics*, 57(2), 151–163.