The basic model and some results from Martin Lettau and Harald Uhlig (1999): "Rules of Thumb versus Dynamic Programming", *American Economic Review* 89(1): 148–174.

Lecture notes for the course
"Bounded Rationality and Macroeconomics"

February 3, 2006

1 From the introduction

Lettau and Uhlig (1999) motivate their analysis by the observation that standard macroeconomic theory has difficulties to explain the empirically observed high sensitivity of consumption to transitory income. Standard modern macro theory describes the consumption-and-saving decisions of economic agents as the solution to a dynamic optimisation problem of maximising a concave utility function (usually assumed to be additively separable over time) subject to a sequence of budget restrictions. One (well-known) result is that agents are predicted to smooth consumption over time as well as across different states of the world.

Some researchers (e.g., Campbell and Mankiw, 1990) have proposed "rules-of-thumb" consumers as an explanation.[1] Various crucial questions arise, however, when assuming rule-of-thumb behaviour: What are the rules that researchers should endow their artificial agents with? If the agents are equipped with several rules, how do they decide which one to use in a given situation? In this context, it may become important to model the agents' evaluation of their rules, such that they can learn which rules are useful and which are not. Consequently and very importantly: why should the agents stick to rules of thumb which make them worse off compared to what they would be if they engaged in dynamic optimisation?[2]

[1] Such rules of thumb could, for instance, be: (*a*) "Save a constant amount of the current income, and consume all your savings as soon as they amount to at least *y* euros." (*b*) "Save all period income that exceeds *x* euros, and consume *p* percent of your savings when they exceed *y* euros." (*c*) "Always have at least *y* euros on your bank account."

[2] I dare to argue that these issues do *not* imply a disadvantage of rule-of-thumb models vis-à-vis models in which the agents are assumed to be "perfectly rational", often identical, fully informed, to possess a certain utility function, etc. Such assumptions are, in my eyes, no less special than the assumption, e.g., that agents evaluate their rules in a specific way (say, by using a linear function).

Lettau and Uhlig address exactly these questions. They show that when agents evaluate their rules' past performance via a specific algorithm, agents may stick to suboptimal rules of thumb *even if* the solution to the dynamic-programming problem is included in the set of rules they are equipped with.

The authors mention several motivations for studying rules of thumb in economic decision making: "partly the difficulty of explaining observed facts [by dynamic optimisation]—partly, because rules of thumb are an interesting paradigm in themselves, and partly, because they can be used as a computational tool" (p. 149).

We will see that the results that Lettau and Uhlig find depend strongly on the assumed set of rules which they equip the agents with. Additionally, the way in which the agents learn about the rules' past performance turns out to be crucial in generating their results. Therefore, it is important that Lettau and Uhlig can provide what they call a "'psychofoundation,' i.e., a foundation in psychology as well as in experimental evidence" (p. 149) for the mechanism they assume. As such a "psychofoundation" they mention (p. 150):

> First, the "Law of Effect" says that choices that have good outcomes in the past are more likely to be repeated in the future (E. Thorndike, 1898). Second, the "Power Law of Practice" states that the slope of the learning curve is decreasing with time (J. M. Blackburn, 1936). Our learning model is consistent with both laws. It is also consistent with "melioration," which means that people choose a strategy that on average gave the maximal payoff in the past rather than choosing what is optimal now: for experimental evidence, see Erev et al. (1997). … Gerald M. Edelman (1992) provides a useful reference on how a biologist and psychologist relates the functioning of the human brain to classifier system like structures, indicating that this paradigm may be promising indeed.

A little problematic in this context is, however, that "[t]here are surprisingly few experimental studies on dynamic economic decision problems" (p. 150).

2 From Section I—"The Sensitivity of Consumption to Transitory Income"

2.1 From Subsection A—"A Stylized Example"

The agent in Lettau and Uhlig's model faces an infinite-horizon dynamic decision problem. In each period, given some state of the world, the agent chooses some (feasible) action. The state in combination with the action then delivers instantaneous utility; afterwards, the state for next period is drawn.

The usual choice in economics for solving this type of optimisation problem is dynamic programming. Here, however, agents use *rules* instead: a rule maps states into actions. As one can easily imagine, the domain of a certain rule is usually not the entire set of possible states. Consider, e.g., rule (*a*) from footnote 1: "Save a constant amount of the current income, and consume all your savings as soon as they amount to at least y euros." This prescription can be split up into two different rules, each of which applies only to a certain domain of the entire state space: "Save a constant fraction of the current income" applies to all possible states, whereas "Consume all your savings" is, in this example, restricted to be applicable only in states that go along with having at least y euros on the bank account.

Lettau and Uhlig assume that the agent is endowed with a fixed set of rules; i.e., there is *no* invention of new rules and no recombination of existing ones via a genetic algorithm or the like[3]—all the agent does is to choose between the rules that she is endowed with. She does so based on past experience: with each rule, a real number, called "strength", is associated; this "strength" is a weighted average of past "rewards" from using the respective rule. The "rewards", in turn, are the sum of instantaneous utility and the discounted strength of the rule used in the subsequent period (how this works will be illustrated in detail by an example below). The latter idea is supposed to capture the dynamic nature of the agent's decision making: She is aware of the recurring nature of the choices she is making; however, unlike "economic man" she is unable to calculate the best course of action in a forward-looking fashion and instead resorts to rules of thumb.

Lettau and Uhlig's main point is the following: They can show that even when the optimal solution to the dynamic programming problem ("policy function") is among the rules which an agent has been equipped with, it can be the case that the agent uses a different rule! "This can happen if the suboptimal rule is only applicable in 'good' states, in which it is easy to generate high rewards: the strength of that rule will be correspondingly biased upward, possibly exceeding the strength of a universally applicable dynamic programming solution rule" (p. 151). Lettau and Uhlig call this the "good state bias".

This train of thought is illustrated in the paper by a by an example, designed to explain the excess sensitivity of consumption to transitory income. In this example the agent is equipped with two rules only:

1. The first rule implements the dynamic-programming solution.
2. The second rule is "spend everything". This rule is assumed to be applicable only when income is "high".

[3] This is unlike what is modelled in the paper by Arthur et al. (1996) that we discussed in the seminar session on January 27th, 2006.

The agent lives indefinitely long. She consumes $c_t \geq 0$ in period t; her discount rate is $0 < \beta < 1$. Her income y_t is given by an exogenous random sequence and can take on only two different values: $y_t \in \{\underline{y}, \overline{y}\}$. The sequence is assumed to be a Markov chain with transition matrix

$$P = \begin{pmatrix} p_{\overline{y}\overline{y}} & p_{\overline{y}\underline{y}} \\ p_{\underline{y}\overline{y}} & p_{\underline{y}\underline{y}} \end{pmatrix},$$

where $p_{yy'}$ is the probability of drawing y' in the upcoming period, given that y was drawn in the current period.

At the beginning of period t, the agent possesses wealth w_t, where initial wealth $w_0 = 0$ by assumption. Consequently, next period's wealth is given by $w_{t+1} = w_t + y_t - c_t$. A borrowing constraint ensures that $0 \leq c_t \leq w_t + y_t$. Written as a dynamic-programming problem, this means

$$v(w, y) = \max_{c \in [0, w+y]} \left(u(c) + \beta \sum_{y'} p_{yy'} v(w + y - c, y') \right).$$

The period utility function u fulfils the usual assumptions. Therefore, this problem is solved by a unique decision function ("policy function"), which we denote by $c^*(w, y)$. Note that $c^*(w, y) = 0$.

Consider now an agent who has two decision rules at her disposal, r_1 and r_2. r_1 coincides with the policy function resulting from the dynamic-programming problem: $r_1(w, y) = c^*(w, y)$. As already mentioned, it is applicable in the good state as well as in the bad state. r_2, in contrast, is assumed to be only applicable in the good state and amounts to consuming everything: $r_2(w, \overline{y}) = w + \overline{y}$. Hence, the choice that the agent faces is whether to use r_1 or r_2 in the good state.

Suppose that the agent has so far always used the "spend-it-all" rule r_2 when her income was high. Now she reconsiders this behaviour and does so by calculating the strengths z_1 and z_2 from the rules' past payoffs. Suppose that income was high at time t and low in the subsequent period $t + 1$. Thus, in t rule r_2 was used, and in $t + 1$ rule r_1 was used. Hence, the payoff attributed to using rule r_2 at t is $u(c_t) + \beta z_{1,t-1}$. If the income had been high in $t + 1$, triggering rule r_2, the payoff attributed to using r_2 in t would have been, consequently, $u(c_t) + \beta z_{2,t-1}$.

More generally, the strengths evolve as follows (compare Lettau and Uhlig's formula (13) on p. 160):

$$z_{i,t} = (1 - \gamma) z_{i,t-1} + \gamma (u_t + \beta z_{j,t-1}),$$

where i is the index of the rule used in t and j is the index of the rule used in $t + 1$. $\gamma \in [0, 1]$ is a weighting factor. For any rule k that was *not* used in t, the strength is not updated, i.e. $z_{k,t} = z_{k,t-1}$. Note the following remark on the timing: “Choosing an action takes place at the end of date t, whereas the updating step for the strength of the active classifier in period t takes place at the beginning of date $t + 1$” (p. 159).

The following elaboration on LETTAU AND UHLIG’s example of an agent who always uses rule 1 in the low-income state and rule 2 in the high-income state illustrates the mechanism outlined in the previous paragraph:

Example (compare Section IV of Lettau and Uhlig; here, $\gamma = 1$)

$t = 0$: $z_{1,0} = 0$
$z_{2,0} = 0$

$t = 1$: $y_1 = 10$
⇒ rule used: r_2 (by assumption)
⇒ instantaneous utility: $u(y_1) = u(10)$

$t = 2$: $y_2 = 5$
⇒ rule used: r_1
⇒ instantaneous utility: $u(y_2) = u(5)$
⇒ strength 1 not updated: $z_{1,1} = z_{1,0}$ (because r_1 was not used in $t = 1$)
⇒ strength 2 updated: $z_{2,1} = u(y_1) + \beta z_{1,0}$ (because r_2 was used in $t = 1$ and r_1 was used in $t = 2$)

$t = 3$: $y_3 = 5$
⇒ rule used: r_1
⇒ instantaneous utility: $u(y_3) = u(5)$
⇒ strength 1 updated: $z_{1,2} = u(y_2) + \beta z_{1,1}$ (because r_1 was used in $t = 2$ and also in $t = 3$)
⇒ strength 2 not updated: $z_{2,2} = z_{2,1}$ (because r_2 was not used in $t = 2$)

$t = 4$: $y_4 = 10$
⇒ rule used: r_2 (by assumption)
⇒ instantaneous utility: $u(y_4) = u(10)$
⇒ strength 1 updated: $z_{1,3} = u(y_3) + \beta z_{2,2}$ (because r_1 was used in $t = 3$ and r_2 was used in $t = 4$)
⇒ strength 2 not updated: $z_{2,3} = z_{2,2}$ (because r_2 was not used in $t = 3$)

And so on …

In this example, the probabilities of switching between the two rules are given by the transition probabilities of the two states. Therefore, if many periods are used for calculating the strengths, the agent solves approximately the following system of linear equations in z_1 and z_2:

$$z_1 = u(\underline{y}) + \beta(p_{\underline{y}\underline{y}} z_1 + (1 - p_{\underline{y}\underline{y}}) z_2); \tag{1}$$

$$z_2 = u(\overline{y}) + \beta(p_{\overline{y}\underline{y}} z_1 + (1 - p_{\overline{y}\underline{y}}) z_2). \tag{2}$$

If z_2 is larger than z_1, then the assumed behaviour (using rule 2 in the good and rule 1 in the bad state) can indeed be the outcome of the evaluation scheme because under this condition the agent has no reason to change her behaviour. Subtracting equation (1) from equation (2), we see that indeed

$$z_2 - z_1 = \frac{u(\overline{y}) - u(\underline{y})}{1 - \beta(p_{\underline{y}\underline{y}} - p_{\overline{y}\underline{y}})} > 0,$$

because u is increasing and $0 < \beta, p_{\underline{y}\overline{y}}, p_{\underline{y}\underline{y}} < 1$. "The intuition behind this result is this: rule r_2 may 'win' against rule r_1 since it only applies in 'good times' and thus 'feels better' on average than rule r_1. This 'good state bias' gives rule r_2 an intrinsic advantage when competing against the optimal rule r_1, which is applicable at all times" (p. 153).

Of course, in this example, spending everything in the good state implies that one has to spend also everything in the bad state—which is too extreme to be a description of people's actual behaviour. However, recall the example rules mentioned at the beginning of this subsection ("Save a constant fraction of the current income", applying to all possible states; and "Consume all your savings", applicable only in states that go along with having at least y euros on the bank account): they are quite similar in spirit. Hence, one can expect that LETTAU AND UHLIG's findings carry over to cases of less restrictive rules. In any case, this two-states–two-rules example is merely supposed to demonstrate the very core of the analysis. LETTAU AND UHLIG also simulate a more complicated model in which five states are incorporated, and the "spend-it-all" rule is weekend a little. Let's turn to this extension now.

2.2 From Subsection B—"Calibrated Calculations"

In a next step, LETTAU AND UHLIG extend their model and calibrate the extended version in order to match empirical data on the sensitivity of consumption to changes in income. Regarding their choice of the calibrated values, LETTAU AND UHLIG refer to HUBBARD ET AL. (1994, 1995). The latter authors estimate empirically the following earnings equation:

$$\log y_{it} = Z_{it}\beta + u_{it} + v_{it}.$$

The dependent variable, $\log y_{it}$, is the log of earnings. Z_{it} is a vector incorporating the usual explanatory variables (such as age, age squared, work experience, years of schooling etc.), while v_{it} is measurement error and u_{it} is an autocorrelated part.

According to LETTAU AND UHLIG (p. 154), "[a]s for u_{it}, Hubbard et al. (1994, 1995) fitted an AR(1) and found its its [sic] autocorrelation to equal $\rho = 0.955$ and its innovation variance to equal 0.033 for household heads without high-school education".

LETTAU AND UHLIG simulate $n = 5$ different states. These correspond to five different income levels y, chosen to be equally spaced in logs and normalised so that the mean income is unity. The Markov transition probabilities $p_{yy'}$ are chosen to equal $\rho - (1 - \rho)/n$ for $y = y'$ and $(1 - \rho)/n$ for all other y', which delivers an autocorrelation of ρ. The calibration of the further parameters is $\beta = 0.96$ and $R = 1.02$, with R being the real return on savings. The utility function was chosen to be of the constant-relative-risk-aversion (CRRA) type: $u(c) = c^{1-\gamma}/(1-\gamma)$, where the coefficient of relative risk aversion was set to $\gamma = 3$.

Like in the highly stylised example of the previous subsection, there is a representative agent equipped with two rules:

1. Rule 1 is again the implementation of the dynamic-programming solution. This rule is, again, assumed to be *always applicable*, i.e. in all $n = 5$ states.
2. Rule 2 is a convex combination of rule 1 and "spending everything", where λ is the weight on "spending everything". Hence, $\lambda = 0$ implements the dynamic-programming solution, and $\lambda = 1$ means that the entire current wealth is consumed. This rule is applicable only in a subset of states: whenever current wealth *and* current income exceed some respective cutoff levels.

The actions that the agent can implement are spending different percentages of her current total wealth. These percentages are restricted to lie on a grid of 40 possibilities in the interval [0, 1]; the grid is "somewhat denser close to the extreme values of 0 and 1" (p. 154). The total wealth is restricted to lie on 80 grid points—which are logarithmically evenly spaced, starting at 0.0176 and ending at 119.

LETTAU AND UHLIG's aim is to simulate excess sensitivity of consumption to transitory income; they try to get close to the empirical data by varying λ and the two cutoff levels. So how can one detect "excess sensitivity" of consumption to current income in empirical data? The "excess" is defined in comparison to the predictions of the theoretical model using dynamic optimisation. ZELDES (1989) regresses $\Delta \log c_{t+1} = \log c_{t+1} - \log c_t$ on $\log y_t$ and finds the coefficient on $\log y_t$ to be -0.07 for low-wealth individuals. That is, the higher y_t, the larger is the *reduction* in consumption in the subsequent period: people obviously spend a relative lot of an income rise immediately and then have to reduce consumption in the next period. Alternatively, one can regress $\Delta \log c_{t+1} = \log c_{t+1} - \log c_t$ on $\mathrm{E}_t[\Delta \log y_{t+1}]$.[4] The coefficient on $\mathrm{E}_t[\Delta \log y_{t+1}]$ is found to be around 0.4 by LUSARDI (1996).

[4] A prerequisite for this to yield reasonable results is, of course, that people form expectations in the same way (at least on average) as the econometrician calculates them!

TABLE 1—VARIATIONS IN λ (ZERO VARIANCE IN TRANSITORY INCOME COMPONENT)

Regression of Δ log c_{t+1} on constant and log y_t					
λ =	0.0	0.25	0.5	0.75	1.0
wealth ≥ 0	−0.004	−0.027	−0.034	−0.039	−0.046
wealth ≥ 0.02	−0.004	−0.023	−0.027	−0.028	−0.030
wealth ≥ 1	−0.004	−0.022	−0.024	−0.026	−0.028
wealth ≥ 3	−0.004	−0.017	−0.018	−0.020	−0.022
wealth ≥ 10	−0.004	−0.012	−0.013	−0.014	−0.016

Regression of Δ log c_{t+1} on constant and $E_t[\Delta \log y_{t+1}]$					
λ =	0.0	0.25	0.5	0.75	1.0
wealth ≥ 0	0.095	0.590	0.761	0.864	1.031
wealth ≥ 0.02	0.095	0.518	0.602	0.614	0.667
wealth ≥ 1	0.095	0.488	0.543	0.575	0.619
wealth ≥ 3	0.095	0.373	0.411	0.447	0.489
wealth ≥ 10	0.095	0.272	0.296	0.322	0.361

Average fraction of time in percent that agent applies the rule					
λ =	0.0	0.25	0.5	0.75	1.0
wealth ≥ 0	40.0	40.0	40.0	40.0	40.0
wealth ≥ 0.02	39.7	39.3	39.2	39.1	19.8
wealth ≥ 1	39.3	38.4	26.2	16.2	13.4
wealth ≥ 3	37.6	22.4	10.8	7.3	5.7
wealth ≥ 10	30.8	5.2	2.9	2.1	1.8

Notes: The suboptimal rule 2 is applicable whenever wealth exceeds the level indicated and whenever income does not fall below the median income $y = 0.835$. The transitory income component ν_{it} has zero variarnce.

Table 1 from Lettau and Uhlig (1999, p. 156).

Recalling that $\lambda = 0$ corresponds to "perfect rationality"—i.e., implementation of the solution to the dynamic-programming problem—in LETTAU AND UHLIG's simulations, one can see from Table 1 the "excess sensitivity" phenomenon:

1. A "perfectly rational" agent creates a coefficient of only −0.004 (instead of the empirically found −0.07) in a regression of $\Delta \log c_{t+1}$ on $\log y_t$. That is, also a "standard" agent reacts to an income rise by increasing current consumption and then consuming a little less in the next period than in the current period, but the difference between the two consumption levels is much lower than found in empirical data.
2. The coefficient in the regression of $\Delta \log c_{t+1}$ on $\mathrm{E}_t[\Delta \log y_{t+1}]$ would be only 0.095 for a "perfectly rational" agent (instead of the empirically found 0.4).

In their five-states–two-rules framework, LETTAU AND UHLIG try to replicate the empirical findings by varying λ and by varying the cutoff level of wealth, above which rule 2 can be applied. They argue, however, that they rather aim at a value of −0.02 (not −0.07) in the regression of $\Delta \log c_{t+1}$ on $\log y_t$ (see p. 155 for the reasons that they provide).

It turns out—see Table 1—that even with moderate deviations from "perfect rationality", the authors get regression coefficients that are quite close to the values they aim at. For example, with $\lambda = 0.25$ and the wealth cutoff level (above which rule 2 can be applied) being equal to 1, the coefficients generated by Lettau and Uhlig's simulations are −0.022 and 0.488, respectively. With $\lambda = 0.25$ and the wealth cutoff level being equal to 3, the coefficients are −0.017 and 0.373, respectively.

References

Arthur, W. Brian, John H. Holland, Blake LeBaron, Richard G. Palmer and Paul Tayler (1996): "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market". Santa Fe Institute Working Paper no. 96-12-093. Available online: *http://www.santafe.edu/s/publications/Working-Papers/96-12-093.ps*. Also published in: W. Brian Arthur; Steven N. Durlauf; David Lane (eds.) (1997): *The Economy as an Evolving Complex System II*. Addison-Wesley: Reading.

Campbell, John Y. and N. Gregory Mankiw (1990): "Permanent Income, Current Income, and Consumption", *Journal of Business and Economic Statistics*, vol. 8, no. 3, pp. 265–279.

Hubbard, R. Glenn, Jonathan Skinner and Stephen P. Zeldes (1994): "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving", *Carnegie–Rochester Conference Series on Public Policy*, vol. 40, no. 4, pp. 59–125.

Hubbard, R. Glenn, Jonathan Skinner and Stephen P. Zeldes (1995): "Precautionary Saving and Social Insurance", *Journal of Political Economy*, vol. 103, no. 2, pp. 360–399.

Lettau, Martin and Harald Uhlig (1999): "Rules of Thumb versus Dynamic Programming", *American Economic Review*, vol. 89, no. 1, pp. 148–174.

Lusardi, Annamaria (1996): "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets", *Journal of Business and Economics Statistics*, vol. 14, no. 1, pp. 81–90.

Zeldes, Stephen P. (1989): "Consumption and Liquidity Constraints: An Empirical Investigation", *Journal of Political Economy*, vol. 97, no. 2, pp. 305–346.