



ELSEVIER

Contents lists available at ScienceDirect

## Journal of Economic Behavior &amp; Organization

journal homepage: [www.elsevier.com/locate/jebo](http://www.elsevier.com/locate/jebo)

# Lifecycle consumption plans, social learning and external habits: Experimental evidence<sup>☆</sup>

Enrica Carbone<sup>a</sup>, John Duffy<sup>b,\*</sup><sup>a</sup> *Seconda Università di Napoli, Italy*<sup>b</sup> *University of California, Irvine, United States*

## ARTICLE INFO

*Article history:*

Received 23 September 2013

Received in revised form 10 July 2014

Accepted 13 July 2014

Available online 12 August 2014

*JEL classification:*

C91

C92

D11

D91

E21

*Keywords:*

Consumption

Intertemporal optimization

Social learning

Lifecycle models

External habit formation

Experimental economics

## ABSTRACT

We report results from a laboratory experiment exploring the extent to which individuals can solve a deterministic, intertemporal lifecycle consumption optimization problem and the effect of revealing social information on past average consumption amounts; as all individuals have identical induced preferences and lifetime incomes, such social information could be useful in solving for the optimal consumption path. Instead, we find that the provision of social information on past average levels of consumption results in a greater deviation of consumption from both the unconditional and the conditionally optimal paths. We find some improvement in consumption planning relative to the conditional optimum when social concerns (external habits) are explicitly incorporated into subject's period utility functions as in external habit formation preference specifications. Our results on the effects of social information on consumption behavior may help to explain the phenomenon of over-consumption and under-saving that has been observed in many developed countries in recent decades as social information on the behavior of others has become more readily available.

© 2014 Elsevier B.V. All rights reserved.

## 1. Introduction

In dynamic, intertemporal models of lifecycle consumption, it is standard to assume that agents can solve for the optimal consumption/savings plan over their expected lifetimes given all available information on income and prices. Indeed, this assumption underlies all modern micro-founded models of household behavior, beginning with Modigliani and Brumberg's (1954) lifecycle theory of consumption and Friedman's (1957) permanent income hypothesis. The possibility that individuals might condition their choices on the decisions made by other, similarly situated individuals is typically excluded (e.g., via the representative agent assumption) though there are specifications of preferences where *habit* levels of consumption as determined by the choices of other individuals do enter into agents' utility functions and therefore affect individuals' consumption choices. However, there is little doubt that individuals often look to their peers when deciding how much to consume or to save and indeed, the notion that individuals make such social comparisons was an important part of earlier theories of consumption behavior by Veblen (1899) and Duesenberry (1949). The idea that the utility derived from

<sup>☆</sup> We thank the co-editor, Ragan Petrie and two anonymous referees for helpful comments and suggestions that greatly improved the paper.

\* Corresponding author. Tel.: +1 412 605 7006.

E-mail addresses: [enrica.carbone@gmail.com](mailto:enrica.carbone@gmail.com) (E. Carbone), [duffy@uci.edu](mailto:duffy@uci.edu) (J. Duffy).

consumption depends on ones' own consumption relative to that of others has been formalized in models of external habit formation or "keeping up with the Joneses" preferences, e.g., by [Abel \(1990\)](#). Nevertheless, there is little in the way of micro-level evidence as to whether and how such social information on the consumption choices of others actually affects an individual's own consumption and savings choices. In this paper we take a first step toward understanding how social information on the consumption choices made by contemporary peers affects an individual's own lifecycle consumption decisions.

Specifically, we examine the impact of social information on consumption and savings decisions over the lifecycle by designing and analyzing data from a controlled laboratory experiment. The control of the laboratory provides us with several benefits that are not available to researchers working with non-experimental field data. In particular, since we can control endowments, interest rates and the utility derived from consumption, we are able to make sharp predictions as to the optimal path of consumption that individuals should follow. Further, the control of the laboratory allows us to manipulate the information that individuals have available to them when making decisions so that we can assess whether information about the decisions of others, i.e., "the Joneses," really matters. Finally, and perhaps most importantly, as we endow all agents with the *same* lifetime lengths and income process and we eliminate all uncertainty, we ensure that social comparisons with others are potentially useful, i.e., that individuals could find it relevant to consider the decisions made by their peers, "the Joneses". By contrast, it is not possible ensure that the Joneses are so similarly situated in studies involving field data. We note further that dynamic optimization problems are difficult to solve (the solution to the one we study in this paper is solved numerically) and so it seems entirely plausible that agents having difficulty solving such problems might reasonably look to the decisions of other, similarly situated individuals in formulating their consumption and savings plans.

Our main finding is that, instead of helping, social information on the consumption and savings plans of peers can have a detrimental effect on an individual's consumption and savings choices and thus his or her welfare, relative to the optimally chosen path. Specifically, when social information on average consumption choices is provided, subjects' consumption and savings plans depart *further* from the optimal path relative to an environment without social information on the consumption choices of others. Intuitively, observing that one's own consumption is below average may entice the observer to consume more than is individually optimal in a race to keep up with the Joneses. This finding is potentially important for understanding the dramatic decline in national savings rates that has taken place in many developed countries (France, Italy, Japan, Spain, the U.K. and the U.S.) since the 1970s as documented, e.g., by [Dobrescu et al. \(2012\)](#), as the availability of information on the consumption decisions of peers has likely increased over this same time period.

In addition to exploring the impact of information on peer behavior for lifecycle consumption and saving choices, we also explore how a leading theory of social preferences with regard to consumption choices – the theory of "external habit formation" – fares in the laboratory. According to this theory, an individual's utility from consumption depends on their consumption choices relative to some external reference level of consumption, known as the "habit level" of consumption. Here we take this reference level to be the same, economy-wide prior period average level of consumption that we use in our social information treatment. Indeed, this choice for the reference point is a standard specification in the external habit formation literature. Specifically, we examine consumption decisions when the period utility function is modified so that individuals derive utility from consumption relative to the prior average level of consumption by all individuals in the economy. We are interested in whether modifying the period utility function so that agents' utility from consumption depends explicitly on the average consumption of other similarly situated individuals, yields any improvement in consumption planning relative to the case where this same reference level of consumption has no direct utility consequences.

## 2. Related literature

There already exists an experimental literature examining whether and how individuals can optimally solve dynamic, intertemporal lifecycle consumption and savings problems. See, for example, [Hey \(1988\)](#), [Hey and Dardanoni \(1988\)](#), [Johnson et al. \(2001\)](#), [Ballinger et al. \(2003, 2011\)](#), [Carbone and Hey \(2004\)](#), [Carbone \(2006\)](#), [Fehr and Zych \(1998, 2008\)](#), [Brown et al. \(2009\)](#), [Feltovich and Ejebu \(2013\)](#) and [Meissner \(2014\)](#). A general finding of this literature is that, over the lifecycle, individuals initially consume too much (or save too little) relative to the optimal path so that toward the end of their lifecycle, savings are too low and consumption is below the optimal path. While our baseline treatment involving an individual, intertemporal lifecycle consumption problem has much in common with the design used in these prior studies, one important difference is that the dynamic optimization problem that we study is *non*-stochastic, making it perhaps the simplest environment yet studied in this literature. Specifically, the entire income sequence over the lifecycle is perfectly known at date 0 as is the constant interest rate on savings as well as the length of the planning horizon. We deliberately chose such a simple environment because we wanted to minimize the role of uncertainty (and uncontrollable attitudes toward risk) and focus attention instead on the question of whether and how optimal consumption plans can be achieved, and the role played by information on the consumption decisions of other similarly situated individuals.

The role of social learning in the formation of lifecycle consumption plans has been previously addressed in experimental studies by [Ballinger et al. \(2003\)](#) and [Brown et al. \(2009\)](#). Those authors report that observation by subjects of the *prior lifecycle consumption plans* of other subjects (who faced the exact same planning problem and horizon) enables the observer subjects to form lifecycle consumption plans that are closer to the optimal intertemporal path relative to the case of no observation. [Brown et al. \(2009\)](#) refer to this kind of social learning as "intergenerational imitation." However, this is just *one* kind of social learning that may be operative. Another form of social learning is that individuals look to the decisions

of their contemporaries or peers in deciding how much to consume and save in every period, a kind of social learning we might term “contemporaneous peer imitation.” This is the type of social learning that we pursue in this paper. [Feltovich and Ejebu \(2013\)](#) have also recently examined peer-to-peer social learning in a lifecycle consumption/savings task. However, they only show subjects information on the highest *payoffs earned* among a cohort of subjects and do not reveal information on actual consumption/savings decisions. By contrast, we study whether and how information on contemporaries’ consumption decisions, specifically information on the economy-wide average level of consumption, affects individual consumption decisions. Specifically we ask whether average information on the consumption decisions of others is used and if so, whether it improves individual decision-making in the direction of the optimal consumption path. Intuitively, since dynamic intertemporal optimization can be computationally difficult, it may be that provision of information on average consumption behavior is a useful aid in solving the intertemporal planning problem.

Our focus on average consumption has another motivation as we also use this same statistic in a third, “external habit” formation treatment. The motivation for this external habit formation treatment comes from the observation that many researchers use external habit formation specifications for the period utility function to capture other concerns regarding social concerns, i.e., keeping up with the Joneses as first formalized by [Abel \(1990\)](#); see [Schmitt-Grohé and Uribe \(2008\)](#) for a survey. The claim is that such specifications for utility aid in explaining a variety of macroeconomic phenomena such as the equity premium puzzle ([Campbell and Cochrane, 1999](#)) or the hump-shaped response of consumption to various expansionary shocks ([Christiano et al., 2005](#)) that cannot be explained with standard, self-regarding utility function specifications. While there is some experimental evidence exploring *internal* habit formation specifications (e.g., [Fehr and Zych, 1998, 2008](#); [Brown et al., 2009](#)) where the habit level of consumption is *internally unique* to each *individual*, there is little in the way of micro-level evidence for the external habit preference class, where the habit level of consumption is not unique to the individual but instead depends on the decisions of his or her peers. Here we do not address where such preferences come from but rather whether they come more naturally to laboratory subjects than preference functions that do not explicitly incorporate concerns for relative consumption.

### 3. Model and experimental design

The model we implement in the laboratory is a standard, workhorse finite-horizon, deterministic model of lifecycle consumption and savings decisions. Each agent’s goal is to:

$$\max_{\{c_t\}} \sum_{t=1}^T u(c_t) \quad (1)$$

subject to

$$c_t + s_t = y_t + (1 + r)s_{t-1}, \text{ and } s_0 = s_T = 0. \quad (2)$$

Here,  $c_t$ ,  $s_t$  and  $y_t$  denote period  $t$  consumption, savings and income, respectively, and  $r > 0$  is the exogenous fixed and known real rate of interest; thus we are considering a partial equilibrium environment.

In one of our experimental treatments, subjects are also shown the ex-post average level of consumption, denoted by:

$$C_{t-1} = \frac{1}{N} \sum_{i=1}^N c_{t-1}.$$

In particular, they are shown this economy-wide average level of consumption for period  $t - 1$  after period  $t - 1$  has concluded but prior to making their consumption choice for period  $t$ . For period 1, no information on average past consumption was reported.

In this finite horizon lifecycle problem, there is no discounting, no uncertainty regarding the income process or the interest rate, and individuals can be presumed to have perfect foresight (at least as a benchmark). Borrowing is not allowed and there is no government/taxation. We deliberately chose to consider this simple, deterministic intertemporal optimal consumption framework not because it is realistic, but because it is both easy to explain to subjects and relatively easier to solve than more complicated, stochastic versions of the same model. If subjects have trouble with this simple, deterministic framework then one can expect that such difficulties will also carry over to more complicated (e.g., stochastic) intertemporal environments. Further, we wish to focus on the role played by observation of the decisions of others and induced external habits, and thus we wish to minimize the role played by other factors, e.g., uncertainty regarding the income process, infinite horizons, liquidity constraints, etc.

#### 3.1. Parameterization

In all cases we chose to work with a population of  $N = 10$  subjects who interact together for a finite lifetime of  $T = 25$  periods. One can think of model periods as approximating one or two year periods in real time. We set the real interest rate on savings equal a constant,  $r = .05$  in all of our treatments; the historical average annual real return on U.S. treasury bonds is 3.5 percent and on U.S. equities it is 6.5 percent so a 5 percent real rate of return on savings seemed to be a good compromise.

The individual income process was also assumed to be constant, with  $y_t = \bar{y} = 10$  for all  $t = 1, 2, \dots, 25$  periods. A constant and certain income level is unrealistic, of course, but we wanted to reduce any uncertainty arising from the determination of permanent income so as to better focus attention on the role played by social information in the formulation of consumption and savings plans.

The period utility function that we induced is specialized to the constant absolute risk aversion (CARA) exponential class:

$$u(x) = \kappa - \frac{1}{R} e^{-Rx}$$

where  $R$  denotes the coefficient of absolute risk aversion. As there is no uncertainty in treatments 1 and 2 of our experiment, the function  $u(x)$  is essentially a means of converting from consumption units into monetary payoffs.<sup>1</sup> We wanted to a concave functional form for  $u$  to insure that the optimal consumption path was unique. We set  $\kappa = 15$  and  $R = .10$ . These choices were made so as to ensure that payoffs were always strictly positive and that differences in monetary payoffs from different consumption amounts were sufficiently salient to subjects.

The parameterization of the environment was public knowledge to all  $N$  agents – that is all agents knew that every other participant had the same income process and faced the same decision problem involving the same induced utility function. Given this information it is (in principle) possible to calculate the optimal path. More importantly, as all individuals were similarly situated (in terms of income, induced preferences), it is reasonable to conjecture that individuals might find information on the *average* level of consumption useful in thinking about their own consumption and saving plans.<sup>2</sup>

### 3.2. Experimental design and hypotheses

Our experiment involves two main treatments<sup>3</sup>:

**Treatment 1** Control: no information on  $C_{t-1}$  is provided.

**Treatment 2** Social Information: information on  $C_{t-1}$  is provided prior to each individual's time  $t$  choice of  $c_t$ .

Under both treatment conditions, the optimal lifecycle consumption path is determined as follows. First, repeated substitutions using the budget constraint (2) yields the lifetime budget constraint:

$$\sum_{t=1}^{T=25} \frac{c_t}{(1+r)^{t-1}} \leq \sum_{t=1}^{T=25} \frac{y}{(1+r)^{t-1}}. \quad (3)$$

Second, the first order necessary conditions for a maximum to the problem ((1) and (2)) are given by the Euler equation,

$$u'(c_t) = (1+r)u'(c_{t+1}), \quad (4)$$

which relates consumption at every date  $t$  to consumption at date  $t+1$  and must be satisfied along the optimal equilibrium path. With  $T=25$ , there are 24 such Euler equations. The optimal consumption path is found by simultaneously solving these 24 Euler equations together with the lifetime budget constraint (3) to obtain 25 different optimal consumption amounts for the 25 periods of an agent's life. Fig. 1 depicts the optimal consumption path for the parameterization of the economy that we study along with the constant income level of 10 per period. The optimal consumption path is increasing since there is a positive real rate of return and no discounting. Notice that the optimal consumption path requires subjects to consume less than their period income of 10 in each of the first 10 periods and to consume more than their period income of 10 in each of the remaining 15 periods.

We hypothesize that since subjects face no uncertainty with respect to their income or the rate of return on savings that they will choose a path for consumption that approximates this optimal path as it maximizes their earnings in the experiment. An auxiliary hypotheses is that subjects get closer to the optimal consumption path with experience (we describe below how subjects acquire experience with lifecycle consumption planning in our experiment).

The only difference between treatment 1 and treatment 2 is that in the latter treatment subjects are informed of the average consumption amount by all  $N$  subjects in the prior period,  $t-1$ , i.e.,  $C_{t-1}$ , before making consumption choices in period  $t$ . However, since  $C_{t-1}$  does not directly enter into subjects' utility (payoff) functions, the optimal consumption path remains the same as for treatment 1 as shown in Fig. 1. Nevertheless, we hypothesize that providing social information on the group average consumption amount may enable subjects to get closer to the optimal consumption path. Intuitively, one may think of this environment as approximating a computational system involving distributed parallel processing (i.e.,

<sup>1</sup> We assume that subjects are risk neutral with regard to monetary payoffs; if they are not, this would add another layer of complexity to the function mapping from consumption to monetary payoffs.

<sup>2</sup> Another possible specification would be for each agent  $i$  to observe the average lagged consumption of the other  $N-1$  agents excluding agent  $i$ . However, using this specification, each agent would potentially be observing different past average values for consumption, which makes the interpretation of such average information more difficult.

<sup>3</sup> Later, in Section 5, we also introduce a third treatment where the external habit level of consumption,  $C_{t-1}$  directly enters into the period utility function.

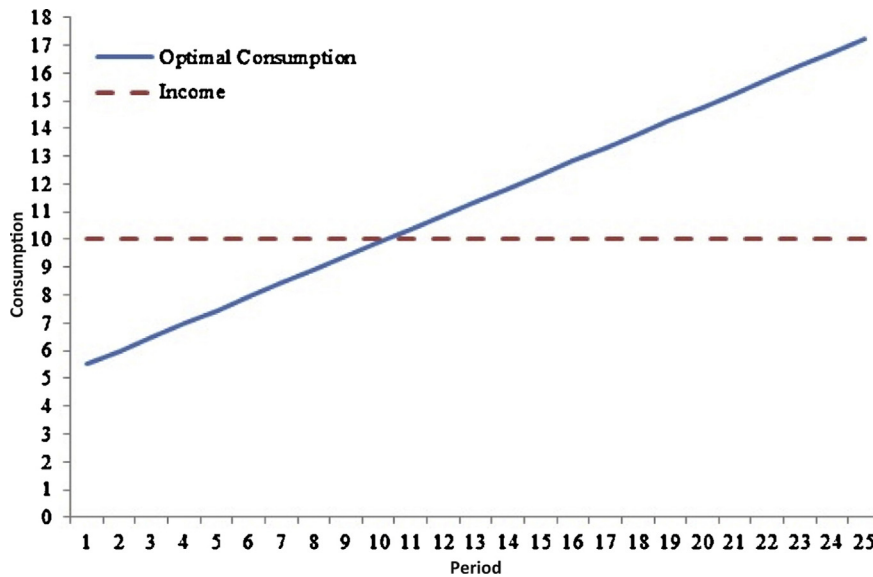


Fig. 1. Optimal lifecycle consumption path for treatments 1 and 2.

the processors are the 10 individual human subjects); since all 10 subjects are homogeneous in terms of their lifetime endowments and payoff objective (and this fact is public knowledge) it seems reasonable that they might assess their own period  $t$  consumption choice relative to the average consumption choice made for that same period and in doing so it is conceivable that they may collectively arrive at the optimal decision. Of course, the freely provided information on the past average consumption choice can also be ignored, in which case performance in treatment 2 should be no worse than in treatment 1.

Note that in our experimental design, information on the past average level of consumption is endogenous to each group of 10 subjects (each session). An alternative approach would be to provide subjects with information on the period-by-period past average level of consumption by *other* subjects in a *different* session of the same treatment. In our view, the latter approach is a less credible means of conveying social information on the consumption choices of *peers*. Also, the habit formation literature (referenced in Section 2) models the habit level of consumption as an endogenous process and we wanted to be consistent with the approach used in that literature.

Each experimental session involved 10 subjects with no prior experience in any treatment of our experiment. That is, we adopt a “between-subjects” design. Subjects were recruited from the undergraduate population of the University of Pittsburgh. At the start of a session subjects were given written instructions which were also read aloud in an effort to make those instructions public knowledge.<sup>4</sup> Subjects had to correctly answer a number of comprehension questions prior to the start of the experiment to ensure that they had a good understanding of the written instructions. Neutral language was used throughout. Subjects were told that they would participate in two “sequences” each of length 25 periods. They were instructed that for each sequence they would receive an endowment of 10 “tokens” ( $y = 10$ ) at the start of each of the 25 periods of that sequence. They were further instructed that in each period of a sequence they could “convert” any number of their available tokens for that period into money earnings (this conversion act is meant to constitute “consumption,” though we did not refer to it as such). Subjects were told that tokens not converted into money earnings, (i.e., tokens that were saved) would earn “interest” at the rate of 5 percent per period. They were also informed that any tokens held beyond the final, 25th period of a sequence would have no redemption value. In making token conversion decisions, subjects were instructed to consider how their token conversions converted into money earnings. The mapping from the number of tokens converted in period  $t$ ,  $c_t$ , to a subject’s “point” earnings for period  $t$  was given by our parameterization of the chosen utility function,  $u(c_t) = 15 - 10e^{-10c_t}$ . Points were converted into money (US dollars) at the fixed and known rate of 1 point = \$0.06 and this final mapping from the number of tokens converted into money amounts was the one that was actually shown to subjects, i.e., subjects were told that their money earnings were given by  $.06u(c) = .9 - .6e^{-1c}$ . While we reported to subjects this monetary payoff formula, we also provided them with a payoff table indicating how various amounts of token conversions they could make would translate into money earnings in each period. In addition to payoff tables, we also provided subjects with a figure illustrating the payoff function. In addition to payoff tables for token conversions (consumption), we also provided subjects with a table explaining how various amounts of tokens saved at time  $t$ ,  $s_t$ , would result in additional token amounts  $r \times s_t$  in the next period. Prior to making a token conversion decision, subjects were informed of the total tokens they

<sup>4</sup> Copies of these instructions including tables, figures and record sheets are given in Appendix B of the Supplementary Material.

**Table 1**

Averages of mean squared deviations of consumption from the optimal path, all 25 periods and 5-period non-overlapping subsamples for both sequences of each session of treatments 1–2.

Treatment-session	All 25 periods	Periods 1–5	Periods 6–10	Periods 11–15	Periods 16–20	Periods 21–25
1-1	26.74	13.01	18.58	17.57	23.81	60.75
1-2	16.07	11.17	13.19	8.32	19.06	28.60
1-3	20.24	13.42	13.89	28.48	16.88	28.52
1-All	21.02	12.53	15.22	18.12	19.92	39.29
2-1	45.33	29.78	42.94	30.37	42.30	81.25
2-2	27.60	10.67	7.334	6.18	10.66	103.16
2-3	57.40	16.27	23.77	39.66	72.64	134.68
2-All	43.44	18.90	24.68	25.40	41.87	106.37

had available which was always at least  $y = 10$  tokens. For periods  $1 < t \leq 25$ , subjects could have additional tokens available to them in the amount  $(1+r)s_{t-1}$  depending on whether they had saved any of their token balance in the previous period, i.e., if  $s_{t-1} > 0$ . Thus in periods  $1 < t \leq 25$ , subjects had  $y + (1+r)s_{t-1}$  tokens available to convert into money earnings in that period  $t$ . In the initial period 1 of each sequence they only had their endowment of  $y = 10$  tokens. Savings were restricted to be non-negative; there was no possibility of borrowing or lending tokens to other participants.

Subjects's monetary earnings accumulated over all 25 periods of each sequence (lifetime). Thus, subject  $i$ 's cumulative monetary payoff for a given sequence was  $\sum_{t=1}^{25} .06u(c_t^i)$ . Two sequences of 25 periods were played so as to determine whether experience mattered or not. Subjects were told that at the end of the session, one of the two 25-period sequences would be randomly selected and they would be paid their total monetary earnings from the one chosen sequence together with a \$5 show-up payment in cash and in private. Average total earnings (including a \$5 show-up payment) were \$21.87 for treatment 1 (standard deviation, 0.30) and \$21.48 (standard deviation 1.41) for treatment 2. Each session was completed in approximately 1.25 hours.

#### 4. Findings

We have data from three sessions of each of our two main treatments. As each session involved 10 subjects, we have data from  $3 \times 2 \times 10 = 60$  subjects. We begin by noting that, as a basic test of rationality, most of our subjects understood that they should consume all of their wealth in the 25th and final period of each lifecycle (or sequence). Indeed, all but 2 of the 30 subjects did this in treatment 1 and all but 3 of 30 subjects did this in treatment 2. Our main results concern the consumption decisions of *all* subjects in each of our two main treatments.<sup>5</sup> We report these main results as a number of different findings. Our first (and main) finding concerns the effect that social information has on consumption choices relative to the optimal path:

**Finding 1.** Social information on average past average consumption amounts causes individual consumption choices to deviate further away from the unconditional optimal path relative to the control treatment where such social information is absent.

Support for Finding 1 comes from Fig. 2a and b and Table 1. Fig. 2a shows the path of the average consumption choices in the second and final 25-period lifetime of each of the three sessions of Treatment 1, which is labeled 'Avg. Cons'.<sup>6</sup> Fig. 2b shows the same for Treatment 2. Also shown in each figure is the time 0 unconditional optimal path for consumption (taken from Fig. 1), labeled 'Optimal' and the average *conditionally* optimal path for the second and final 25-period lifetime in each session of each treatment (more on the conditionally optimal path below). Table 1 confirms the impression conveyed from a comparison of Fig. 2a and b that the MSD of consumption from the unconditional optimal path is lower in treatment 1 than in treatment 2.

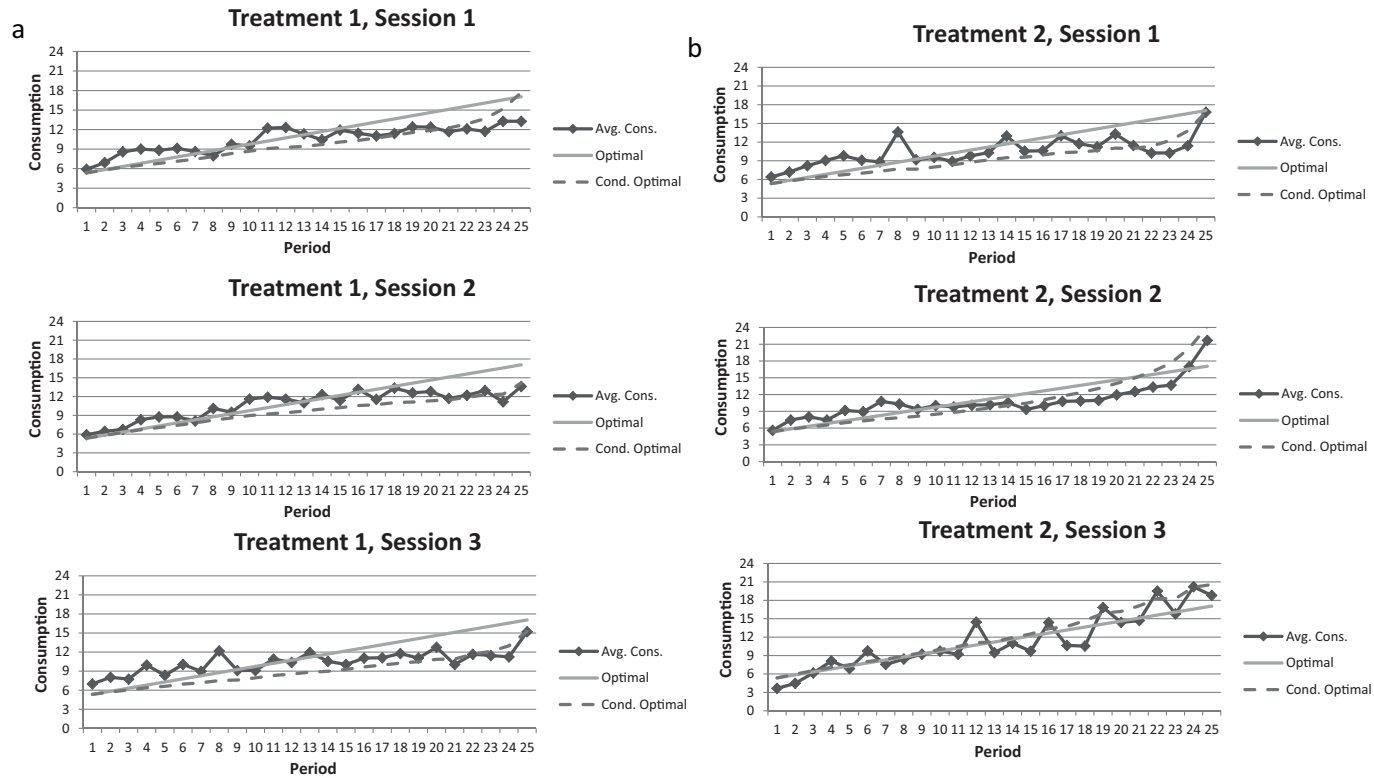
Table 1 reports averages of the mean squared deviations between individual subjects' consumption paths (for both 25-period sequences) and the unconditional optimal consumption path as shown in Fig. 1 for both treatments over all 25 periods and over 5 period non-overlapping intervals for the three sessions of each treatment

To determine whether these differences are statistically significant, we conducted pairwise Wilcoxon–Mann–Whitney (WMW) tests of the null hypothesis of no significant differences in MSDs between our two treatments over all rounds as well as for 5-round non-overlapping intervals. These results are reported in Table 2.

The results indicate that despite the larger MSD from the unconditional optimal path for treatment 2 relative to treatment 1, we cannot reject the null of no difference in MSD from the optimal consumption path between treatments 1 and 2. Of course, a difficulty with considering the MSD from the unconditional optimal consumption path is that mistakes made in early periods (e.g., under-saving) become compounded over time. A more reasonable measure of subjects' performance is

<sup>5</sup> We have verified that the results reported below are unaffected if we remove those few subjects who do not pass the rationality test of consuming all of their wealth in the final period.

<sup>6</sup> We show only the second lifetime to make the graphs clearer. The data analysis, however, makes use of both 25-period lifetimes.



**Fig. 2.** (a) Average consumption (second sequence) in the three sessions of treatment 1 relative to optimal and conditionally optimal consumption. (b) Average consumption (second sequence) in the three sessions of treatment 2 relative to optimal and conditionally optimal consumption.

**Table 2**

Wilcoxon–Mann–Whitney tests of differences in MSD from the optimal consumption path across various rounds of the two treatments.

Rounds	Treatment 1 vs. 2
1–5	$z = -1.446, Pr >  z  = 0.1483$
6–10	$z = -1.275, Pr >  z  = 0.2025$
11–15	$z = -1.352, Pr >  z  = 0.1765$
16–20	$z = -1.413, Pr >  z  = 0.1576$
21–25	$z = -0.887, Pr >  z  = 0.3750$
All 1–25	$z = -1.429, Pr >  z  = 0.1531$

**Table 3**

Averages of mean squared deviations of consumption from the conditionally optimal path, all 25 periods and 5-period non-overlapping subsamples for both sequences of each session of treatments 1–2.

Treatment-session	All 25 periods	Periods 1–5	Periods 6–10	Periods 11–15	Periods 16–20	Periods 21–25
1-1	18.85	15.51	27.90	19.19	10.91	20.71
1-2	11.47	13.58	19.84	9.71	8.37	5.86
1-3	15.45	15.09	19.09	29.31	8.99	4.77
1-All	15.26	14.73	22.28	19.40	9.43	10.45
2-1	34.07	31.32	48.53	28.14	28.92	33.44
2-2	27.79	12.71	12.46	12.92	15.88	85.00
2-3	367.59	18.92	37.96	81.49	167.45	1532.14
2-All	143.15	20.98	32.98	40.85	70.75	550.19

to consider the MSD of consumption from the *conditionally* optimal path. The conditionally optimal consumption path for each subject  $i$  is constructed as follows: Subject  $i$  enters each period  $t$  with some cash on hand (COH) which consists of his or her endowment income,  $y = 10$  (the same for all  $i$ ), plus subject  $i$ 's gross return on savings from the prior period,  $(1.05)s_{t-1}^i$ . We thus treat each individual's COH for period  $t$  as though it were the initial wealth level that the subject brought to solving a reduced,  $T - t + 1$  period consumption planning problem, and we calculate the optimal consumption and savings plan for subject  $i$  conditional on subject  $i$ 's  $COH_t^i$ , as of the start of period  $t$ . In the final period  $T$ , it is optimal for all subjects to consume *all* of their  $COH_T$ . In essence, we use only the current period optimal consumption amount,  $c_t^{i*}$ , given the current period  $COH_t^i$  as our measure of the *conditionally optimal consumption* amount for subject  $i$  and using this value we calculate the MSD as  $(c_t^i - c_t^{i*})^2$  for each subject over all 25 rounds of each lifetime. Considering the MSD of consumption choices from the *conditionally* optimal path we have the following finding:

**Finding 2.** Consumption is significantly further from the conditionally optimal path in treatment 2 than in treatment 1.

Support for Finding 2 comes again from Fig. 2a and b and from Table 3 which, differently from Table 1 reports averages (for both 25-period sequences) of the MSD from the *conditionally* optimal consumption path, the derivation of which was described above. First we calculated each subjects' conditionally optimal consumption path at each time period and we then calculated that subjects' MSD from the conditionally optimal path. The MSDs reported in Table 3 are the mean values of those individual MSDs.

As Table 3 reveals, considering the mean squared deviation of consumption from the *conditionally* optimal path serves to reinforce and strengthen the earlier finding of Table 1 as summarized in Finding 1. In the case of conditionally optimal behavior we observe that treatment 2 has a much greater MSD from the conditionally optimal path as compared with treatment 1. Indeed, WMW tests as reported in Table 4 confirm the impression given by Table 3 as well as Fig. 2a and b that the null hypothesis of no significant differences in MSDs between treatments can now be rejected in pairwise comparisons between treatments 1 and 2 ( $p < .01$ ) in favor of the alternative that the MSD from the conditionally optimal path is *greater* in treatment 2 than in treatment 1.

Our interpretation of the findings of Tables 1–4 is straightforward: social information on the consumption decisions of others as summarized by the past average consumption level,  $C$ , lead to greater departures of consumption from the unconditional or conditionally optimal path, due to a desire by subjects to keep up with the social norm (the Joneses). As this social norm information is *not* present in treatment 1, it does not impact on consumption decisions in that treatment.

**Table 4**

Wilcoxon–Mann–Whitney tests of differences in MSD from the conditionally optimal consumption path across various rounds of the two treatments.

Rounds	Treatment 1 vs. 2
1–5	$z = -4.145, Pr >  z  = 0.0000$
6–10	$z = -3.749, Pr >  z  = 0.0002$
11–15	$z = -4.252, Pr >  z  = 0.0000$
16–20	$z = -3.103, Pr >  z  = 0.0019$
21–25	$z = -3.058, Pr >  z  = 0.0022$
All 1–25	$z = -4.152, Pr >  z  = 0.0000$



**Table 5**  
Random effects regression estimates of consumption decisions.

Dependent variable: consumption			
Variable	Treatment 1 only coefficient (st. error)	Treatment 2 only coefficient (st. error)	Treatments 1–2 combined coefficient (st. error)
Constant	5.60*** (1.477)	2.67** (1.026)	6.79*** (0.943)
seq2	−0.220 (0.485)	0.170 (0.412)	0.00 (0.436)
tr2			−4.20 <sup>†</sup> (1.756)
Period	0.20*** (0.028)	0.21** (0.073)	0.22*** (0.051)
Period × seq2	0.02 (0.033)	0.00 (0.037)	0.01 (0.034)
Period × tr2			−0.02 (0.081)
Prioravgcons	0.080 (0.152)	0.42 <sup>†</sup> (0.197)	
Prioravgcons × tr2			0.42 <sup>†</sup> (0.198)
Wealth	0.07*** (0.007)	0.03 <sup>†</sup> (0.014)	0.04** (0.015)
No. observations	1440	1440	2880
No. subjects	30	30	60
R <sup>2</sup>	0.208	0.169	0.174

Notes: Robust standard errors clustered by session, appear in parentheses.

\*  $p < 0.05$ .

\*\*  $p < 0.01$ .

\*\*\*  $p < 0.001$ .

#### 4.1. Social information, learning and wealth effects on consumption decisions

In an effort to better understand the impact of social information and other potential factors on individual consumption decisions, we turn to a regression analysis of individual consumption decisions that is reported in Tables 5–7. All results reported in these tables are from generalized least squares (GLS) random effects regressions with clustering of robust standard errors on group/session identity. In addition to a constant term, the explanatory variables in these regressions are as follows: “seq2” is a dummy variable indicator for whether the consumption decision was made in the second 25-period sequence of the session; “period” refers to the period number within each 25-period lifetime horizon; period × seq2 captures the interaction between the period number and the second sequence dummy variable, seq2; prioravgcons is the one-period-lagged average consumption level ( $C_{t-1}$ ) that was visible at time  $t$  to subjects in treatment 2 only; finally, “wealth” refers to each subject’s available cash on hand at the start of period  $t$  which is given by:  $10 + (1.05)s_{t-1}$ . We also report on regressions combining data from both treatments 1 and 2. In these regressions, shown in the right-most column of these tables, “tr2” represents a dummy variable for treatment 2 (treatment 1 is the baseline), “period × tr2” and “prioravgcons × tr2” capture the interaction between the period number and prioravgcons, respectively, with the treatment dummy variable, tr2.

Table 5 reports on a regression analysis where the dependent variable is the individual consumption amount (tokens converted) in each period by all subjects. The regression results using the consumption data of treatment 1 only (first column) reveal that the coefficient on the constant term for this treatment is 5.6 which is a little higher than 5.5 for the unconditional optimal path as shown in Fig. 1, and that there appears to be no effect from experience as the coefficient on the dummy variable, “seq2” is not significantly different from zero. We further observe that the coefficient on the period variable is

**Table 6**  
Random effects regression estimates of the difference between actual and optimal consumption decisions.

Dependent variable: deviation of consumption from the unconditional optimal path			
Variable	Treatment 1 only coefficient (st. error)	Treatment 2 only coefficient (st. error)	Treatments 1–2 combined coefficient (st. error)
Constant	0.74 (1.477)	−2.20 <sup>†</sup> (1.026)	1.86 <sup>†</sup> (0.759)
seq2	−0.22 (0.485)	0.17 (0.412)	0.13 (0.114)
tr2			−4.21 <sup>†</sup> (1.793)
Period	−0.29*** (0.028)	−0.28*** (0.073)	−0.26*** (0.036)
Period × seq2	0.02 (0.033)	0.00 (0.037)	
Period × tr2			−0.02 (0.081)
Prioravgcons	0.08 (0.152)	0.42 <sup>†</sup> (0.197)	
Prioravgcons × tr2			0.42 <sup>†</sup> (0.204)
Wealth	0.07*** (0.007)	0.03 <sup>†</sup> (0.014)	0.04** (0.015)
No. observations	1440	1440	2880
No. subjects	30	30	60
R <sup>2</sup>	0.2373	0.0627	0.1124

Notes: Robust standard errors clustered by session, appear in parentheses

\*  $p < 0.05$ .

\*\*  $p < 0.01$ .

\*\*\*  $p < 0.001$ .

**Table 7**

Random effects regression estimates of the difference between actual and conditionally optimal consumption decisions.

Dependent variable: deviation of consumption from the conditionally optimal path			
Variable	Treatment 1 only coefficient (st. error)	Treatment 2 only coefficient (st. error)	Treatments 1–2 combined coefficient (st. error)
Constant	2.52 (1.615)	11.14** (3.808)	6.89*** (1.525)
seq2	−0.32 (0.234)	−1.27 (0.853)	0.63 (0.398)
tr2			2.13 (1.602)
Period	−0.20*** (0.040)	−0.13 (0.123)	−0.13*** (0.014)
Period × seq2	0.03** (0.012)	0.17 (0.125)	−0.02 (0.081)
Period × tr2			−0.02 (0.081)
Prioravgcons	0.24 (0.211)	−0.32 (0.376)	
Prioravgcons × tr2			−0.28 (0.294)
Wealth	−0.05*** (0.012)	−0.22*** (0.052)	−0.18** (0.071)
No. observations	1440	1440	2880
No. subjects	30	30	60
R <sup>2</sup>	0.1254	0.4424	0.3757

Notes: Robust standard errors clustered by session, appear in parentheses. \* $p < 0.05$ .\*\*  $p < 0.01$ .\*\*\*  $p < 0.001$ .

positive and significantly different from zero indicating a slope coefficient of approximately 0.23, which is considerably lower (flatter) than the slope of the optimal consumption path shown in Fig. 1, which has a slope of 0.49. Note also that the variable period × seq is insignificantly different from zero. Thus, relative to the unconditional optimal path, subjects begin both sequences of treatment 1 consuming more than is optimal in the early periods of their lives and consequently they have less wealth so that they are able to consume less over time, in the later periods of their lives as indicated by the flatter than optimal slope coefficient on actual consumption over time (see also Fig. 2a). In addition, we observe that the coefficient on prioravgcons, the prior average level of consumption  $C_{t-1}$ , is *not* significantly different from zero, which is not surprising as subjects in treatment 1 did not have access to this social information when making their consumption decisions. On the other hand, the coefficient on wealth, while small, is positive and significant, indicating that subjects' available cash on hand was a determinant in their consumption choices.

Regression results using the data of treatment 2 indicate that subjects start out consuming much less than the optimal level – the intercept coefficient is 2.67 and significantly different from zero and the coefficient on the period variable is again statistically significant and lower than optimal at 0.21. The main difference however, between treatments 1 and 2 is that the consumption of subjects in treatment 2 is now largely determined by the prior average level of consumption as indicated by the positive and significant coefficient of 0.42 on this variable in the regression reported in Table 5. This high weight placed on the prior, economy-wide average level of consumption confirms the important role of social information in consumption decision-making and helps us to better understand the larger deviation of consumption from the unconditional or conditionally optimal paths in treatment 2 relative to treatment 1 as reported in Findings 1–2. As in treatment 1, in treatment 2 experience seems not to matter as the coefficients on both the dummy variable, “seq2” and the period interaction variable, period × seq2, are not significantly different from zero. Wealth in treatment 2 continues to have a small, significant effect on consumption decisions, but it has less weight as compared with treatment 1.

The final regression combining consumption data from both treatments 1 and 2 serves mainly to confirm the individual treatment regression analyses. We note that all of our regression results remain robust if we eliminate from the specification those variables that are not statistically significant.<sup>7</sup>

In Table 6 we repeat the same regression exercise of Table 5 but here we use as the dependent variable the *deviation* of consumption from the unconditional (time 0) optimal path (as shown in Fig. 1) for each period. Looking at the deviation, as opposed to the mean squared or absolute deviation as we did in Tables 1 and 2, tells us more about the impact of each explanatory variable on departures from the optimal path, e.g., the *sign*, positive or negative, of such departures. Note first that for treatment 1, the initial deviation from the unconditional optimal path is not significantly different from zero while for treatment 2 it is significantly negative as indicated by the coefficient estimate on the constant term. The coefficient on the period variable is significantly negative for both treatments, indicating that the deviation from the unconditional optimal path was decreasing over the 25 period lifecycle which is evidence of some learning behavior. However, there is again no difference in the intercept or slope coefficients in the second 25 period sequence as compared with the first as indicated by the insignificant coefficients on the seq2 and period × seq2 variables.<sup>8</sup> Perhaps most importantly, social information on the prior average consumption of group members again results in a statistically significant and positive deviation of consumption away from the unconditional optimal path only in treatment 2; in treatment 1, where this information is not

<sup>7</sup> Those regression results are available upon request.<sup>8</sup> The insignificance of lifecycle experience in solving consumption/savings plans has also been observed in other experimental studies by Brown et al. (2009) and Meissner (2014).

provided, it has no effect in explaining deviations of consumption from the optimal path. This is consistent with the finding reported in Table 5 that consumption in treatment 2 is positively and significantly affected by increases in lagged average consumption. Finally, we again observe that higher wealth (cash on hand) serves to increase the deviation of consumption from the unconditional optimal path. We summarize the main finding as follows:

**Finding 3.** Subjects' consumption choices depend positively on their current wealth and steadily increase over the 25 periods of a lifecycle. Consumption is significantly affected by the prior average consumption amount in treatment 2 but not in treatment 1 (where that information was not provided). Specifically, in treatment 2, the higher is past average consumption, the greater is an individual's current own consumption and this dependency accounts for much of the deviation of consumption from the unconditional optimal path.

Finally, in Table 7 we report on a regression analysis of the deviation of individual consumption choices from the *conditionally* optimal path (as described previously) using the same explanatory variables as in the regressions reported in Tables 5 and 6. In this regression analysis we observe that the initial deviation of consumption from the conditionally optimal path is again significantly positive only for treatment 2, as indicated by the coefficient on the constant term. We further observe that the deviation of consumption from the conditionally optimal path decreases as the period number increases suggesting again some evidence for learning.

The main differences in our regression analysis of deviations from the conditionally optimal path as reported in Table 7 are (1) the coefficient on *prioravgcons* is now insignificant for both treatments 1 and 2, and (2) the coefficient on *wealth* is now significantly negative for both treatments. To understand these findings, recall that the *conditionally* optimal level of consumption each period depends only on the individual's own current period cash on hand, i.e., on the variable referred to as *wealth* in the regression analysis. As this *wealth* level changes from period to period, the conditionally optimal amount of consumption changes period by period as well and our regression analysis is studying the deviation of consumption choices from this conditionally optimal level. Hence, there is less of a role to be played by longer-term past average consumption considerations in understanding deviations from the conditionally optimal path and a greater role to be played by variations in the *wealth* variable in explaining such deviations. The finding that the coefficient on the *wealth* term is negative in Table 7 while it was positive in Table 6 can be explained as follows. High levels of *wealth* (cash on hand) may result from large departures from the *unconditional* optimal path (as we found in Table 6); for instance, subjects who saved without consuming much for many periods only to periodically consume all of their savings. On the other hand, high levels of *wealth* in any single period are evidence of some predilection to have saved in the immediate prior period. Conditioning on such *wealth* levels as is done in the regression analysis reported in Table 7, the savers (those with high *wealth*) are more likely to behave according to the conditionally optimal prediction than are those with less *wealth*, e.g., the non-savers, and this may account for the negative coefficient on the *wealth* variable in Table 7.

## 5. External habit formation

In this section we report on a third treatment where we replace the period utility function,  $u(c_t)$ , used in treatments 1 and 2 with an external habit formation specification, namely,

$$u(c_t, C_{t-1}) = u(c_t - \alpha C_{t-1})$$

where as before,  $C_{t-1}$  denotes the economy-wide average level of consumption and  $\alpha > 0$  is a parameter measuring the intensity of the external habit for current period consumption. As noted in Section 2, this external "habit formation" specification is a standard means of injecting social concerns about relative consumption in the literature on consumption behavior (see, e.g., Schmitt-Grohé and Uribe, 2008 for a survey) and so it seems reasonable to explore the extent to which laboratory subjects can solve the lifecycle, intertemporal consumption planning problem when they have such an *explicitly induced* concern for the economy-wide average level of consumption,  $C_{t-1}$ .<sup>9</sup>

In this version of the model, the optimal consumption path is again found by maximizing (1) subject to (2) but where  $u(c_t - \alpha C_{t-1})$  is used in place of  $u(c_t)$  and  $C_0$  is set equal to 0. For this version of the model we kept the parameterization of the environment exactly the same as in treatments 1–2. In particular, we continued to set  $r = .05$  and  $y = 10$  for all  $T = 25$  periods. We again used the CARA specification,  $u(x) = \kappa - (1/R)e^{-Rx}$ , where  $\kappa = 15$ ,  $R = .10$  but where  $x$  is now equal to  $c_t - \alpha C_{t-1}$ . We chose to set  $\alpha = .5$ , the midpoint of the feasible range ( $0 \leq \alpha \leq 1$ ). The solution is found in the same manner as before, by combining 24 Euler equations of the form:

$$u'(c_t - \alpha C_{t-1}) = E_t[(1+r)u'(c_{t+1} - \alpha C_t)],$$

with the lifetime budget constraint to get a sequence of optimal consumption amounts in each of the 25 periods of an individual's lifetime. Notice that in this case, subjects must form expectations in period  $t$  of the value of  $C_t$ , which is not

<sup>9</sup> We chose to work with past average specification for the external habit level of consumption rather than (say) the contemporaneous average,  $C_t$  (which is also studied in the literature) to enable comparisons with our treatment 2 and so that subjects would be operationally able to condition their behavior on the habit level of consumption. Using instead the contemporaneous average,  $C_t$ , the external habit level for consumption is *endogenously* and simultaneously determined by subjects' time  $t$  consumption decisions, and thus subjects would not be able to condition on this statistic.




Fig. 3. Optimal lifecycle consumption paths for the model parameterization when  $\alpha = .5$  compared with the case where  $\alpha = 0$ .

known until period  $t + 1$ . The optimal consumption path in this third treatment is constructed similarly to the case where  $\alpha = 0$ ; in constructing this path we assume that agents take the sequence of external habit-level consumption amounts  $\{C_t\}_{t=0}^{24}$  as if it were exogenously determined when solving for their optimal consumption decisions.<sup>10</sup> Fig. 3 provides a comparison of the unconditional optimal consumption path in the case where  $\alpha = .5$  with the original formulation used in treatments 1 and 2 where  $\alpha = 0$ .

Notice that when  $\alpha = .5$ , the optimal consumption path has a lower intercept and a steeper slope than the optimal consumption path when  $\alpha = 0$ . Intuitively, the explicit inclusion of the  $C_{t-1}$  term in the period utility function means that savings must now initially be greater (consumption must initially be lower) in the external habit formation treatment so that the individual may realize a positive utility from consumption later on in life when  $C_{t-1}$  is also larger. We hypothesized that the inclusion of the external habit formation term in the utility function would indeed lead subjects to save more in the earlier periods of their lives relative to the two treatments where  $\alpha = 0$ , and thus, that external habit formation might work to correct a tendency toward over-consumption as in our earlier treatments.

Aside from the change in  $\alpha$  from 0 to 0.5, the model economy of treatment 3 is identical to that of treatments 1 and 2. The experimental procedures were also similar to those used in treatments 1–2. As in treatment 2, subjects in treatment 3 were shown the past average level of consumption,  $C_{t-1}$ , on their computer screens prior to making their consumption decision in each period  $t = 1, 2, \dots, 25$ , with  $C_0 = 0$ . The main difference between treatments 2 and 3 is that subjects in treatment 3 had an explicit incentive to condition their consumption/savings decision on the prior average level of consumption as this level now directly affected their payoff each period. While we reported to subjects the monetary payoff formula for treatment 3, as in treatments 1–2 we also provided subjects with a payoff table indicating how various amounts of token conversions they could make would translate into money earnings in each period. However, by contrast with treatments 1–2, the payoff table used in treatment 3 was two-dimensional showing monetary earnings for a large number of possible  $\{c_t, C_{t-1}\}$  pairs. In addition to payoff tables, we also provided subjects with a 3-dimensional figure showing the payoff surface as a function of  $c_t$  and  $C_{t-1}$ .<sup>11</sup> In all other respects, the procedures were the same as for treatments 1–2. We conducted three sessions of treatment 3, each involving 10 subjects with no prior experience in our experiment who were drawn again from the undergraduate population of the University of Pittsburgh. Each session again consisted of decision-making in two 25-period lifetimes and each subject's earnings in one of these two lifetimes was randomly chosen for payment at the end of each session. Average total earnings for treatment 3 (including a \$5 show-up payment) were \$18.23 (standard deviation, 1.27) for an approximately 1.25 hour session.

### 5.1. Treatment 3 findings

Table 8 reports averages of the mean squared deviations between individual consumption paths and the unconditional or conditionally optimal consumption path (for both 25 period sequences) over various periods of the lifecycle. Similar to

<sup>10</sup> With  $N = 10$  subjects this seemed a reasonable assumption.

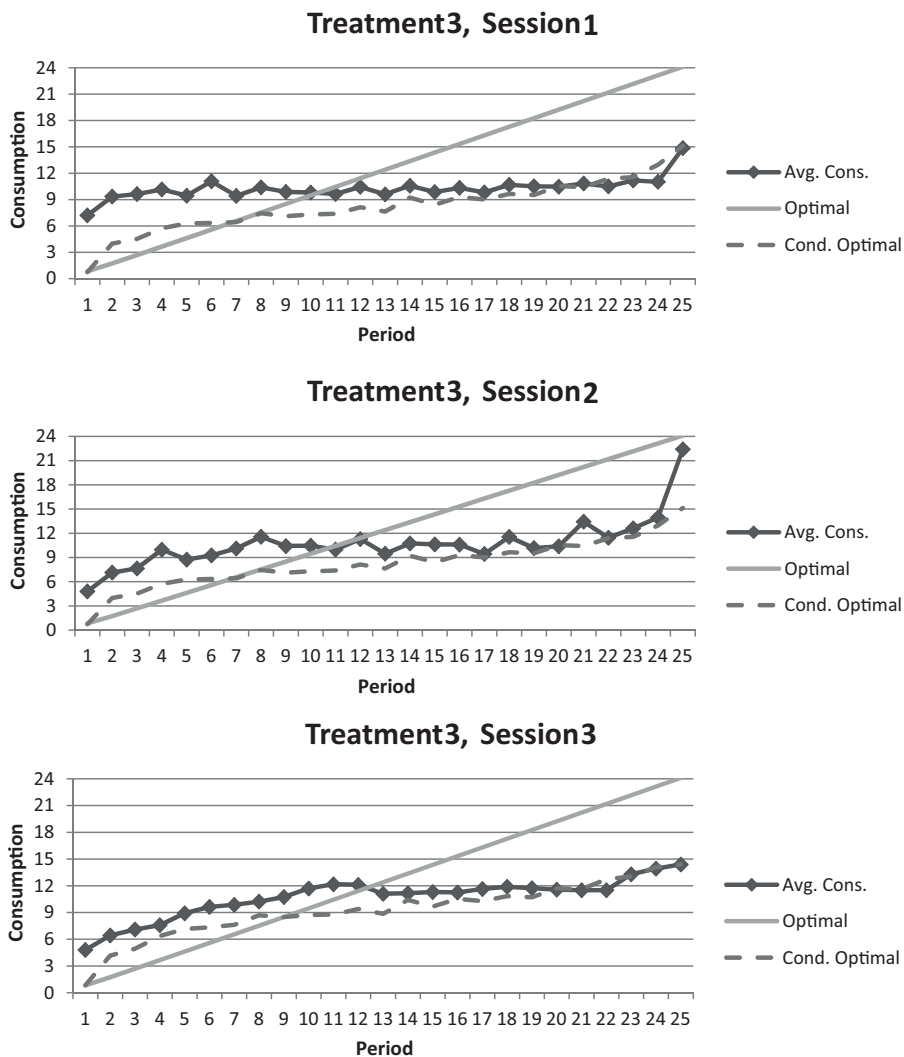
<sup>11</sup> See the written instructions for the figures presented to subjects.

**Table 8**

Averages of mean squared deviations of consumption from the unconditional and conditionally optimal path, all 25 periods and 5-period non-overlapping subsamples for both sequences of each session of treatment 3.

Treatment-session	All 25 periods	Periods 1–5	Periods 6–10	Periods 11–15	Periods 16–20	Periods 21–25
<i>Deviation of consumption from unconditional optimum</i>						
3-1	48.33	42.10	12.68	8.56	53.32	124.99
3-2	45.78	33.05	18.16	19.09	53.21	105.40
3-3	53.18	31.05	13.78	9.40	41.70	169.95
3-All	49.10	35.40	14.87	12.35	49.41	133.45
<i>Deviation of consumption from conditional optimum</i>						
3-1	11.41	24.94	13.59	8.90	2.94	6.67
3-2	18.03	19.40	18.45	14.53	11.66	26.12
3-3	21.16	20.19	16.31	21.05	42.02	6.25
3-All	16.87	21.51	16.12	14.83	18.87	13.01

Fig. 2a and b, Fig. 4 shows the path of average consumption choices in the second and final 25-period lifetime of the three sessions of Treatment 3, labeled 'Avg. Cons.' along with the unconditional optimal path for consumption (taken from Fig. 3) and the mean conditionally optimal consumption path for treatment 3. Note that for treatment 3, the conditionally optimal consumption path for each agent  $i$  as of period  $t$  depends on both the available cash on hand,  $COH_t^i$ , and the value of lagged



**Fig. 4.** Average consumption (second sequence) in the three sessions of treatment 3 relative to optimal and conditionally optimal consumption.

**Table 9**

Wilcoxon–Mann–Whitney tests of differences in MSD from the unconditional and conditionally optimal consumption path across various rounds of treatments 3 vs. treatments 1–2.

Rounds	Treatment 1 vs. 3	Treatment 2 vs. 3
<i>Comparison of deviation of consumption from unconditional optimum</i>		
1–5	$z = -12.977, \Pr >  z  = 0.0000$	$z = -10.999, \Pr >  z  = 0.0000$
6–10	$z = -17.807, \Pr >  z  = 0.0000$	$z = -15.591, \Pr >  z  = 0.0000$
11–15	$z = -18.539, \Pr >  z  = 0.0000$	$z = -16.560, \Pr >  z  = 0.0000$
16–20	$z = -12.551, \Pr >  z  = 0.0000$	$z = -10.932, \Pr >  z  = 0.0000$
21–25	$z = -13.716, \Pr >  z  = 0.0000$	$z = -12.527, \Pr >  z  = 0.0000$
All 1–25	$z = -16.600, \Pr >  z  = 0.0000$	$z = -14.623, \Pr >  z  = 0.0000$
<i>Comparison of deviation of consumption from conditional optimum</i>		
1–5	$z = 0.244, \Pr >  z  = 0.8069$	$z = 4.368, \Pr >  z  = 0.0000$
6–10	$z = -1.466, \Pr >  z  = 0.1427$	$z = 2.375, \Pr >  z  = 0.0176$
11–15	$z = -1.174, \Pr >  z  = 0.2405$	$z = 3.080, \Pr >  z  = 0.0021$
16–20	$z = -1.391, \Pr >  z  = 0.1643$	$z = 1.788, \Pr >  z  = 0.0738$
21–25	$z = 0.187, \Pr >  z  = 0.8520$	$z = 3.231, \Pr >  z  = 0.0012$
All 1–25	$z = -0.825, \Pr >  z  = 0.4091$	$z = 3.357, \Pr >  z  = 0.0008$

average consumption  $C_{t-1}$ , which are known at the start of period  $t$ .<sup>12</sup> Thus unlike in treatment 2, the conditionally optimal path in treatment 3 explicitly conditions on the value of  $C_{t-1}$ .

By comparison with Tables 1–3 and Fig. 2a and b, we note several differences. First, the MSD from the unconditional optimal path is much greater in Treatment 3 as compared with Treatments 1–2. Further, these differences are statistically significant according to Wilcoxon–Mann–Whitney tests as reported in Table 9. An explanation for the greater deviation of consumption from the unconditional optimal path in treatment 3 relative to the other two treatments is that the optimization problem in treatment 3 is considerably more difficult than in the other two treatments, as subjects have to consider their payoffs as a function of two variables,  $c_t$  and  $C_{t-1}$ , rather than just their own consumption choices.

Regarding deviations from the conditionally optimal path, we find that treatments 1 and 3 have roughly similar MSDs from the conditionally optimal path while treatment 2 has a much greater MSD from the conditionally optimal path. Indeed, WMW tests as reported in Table 9 confirm that the null hypothesis of no significant differences in MSDs between treatments cannot be rejected in pairwise comparisons between treatments 1 and 3 ( $p = .41$ ) but can be rejected in pairwise comparisons between treatments 1 and 2 and between treatments 2 and 3 ( $p < .01$ ) in favor of the alternative that the MSD from the conditionally optimal path is greater in treatment 2 than in the other two treatments. We summarize these findings as follows

**Finding 4.** Consumption is furthest from the unconditional optimal consumption path in the external habit formation treatment 3. Regarding deviations from the conditionally optimal path, treatments 1 and 3 are not significantly different from one another while treatment 2 exhibits a significantly larger deviation from the conditionally optimal path than the other two treatments.

Recall that the optimal consumption path in treatment 3 requires that subjects save much more in the first part of their 25-period lifetime than in treatments 1 or 2 – see again Fig. 3. That is, the optimal consumption path in the external habit formation treatment has a lower intercept and a steeper slope than in treatments 1 and 2 where the external habit level of consumption does not directly enter into the period utility function. Consequently, if subjects initially approach the intertemporal, lifecycle planning problem similarly across all three treatments, then they are destined to do worse relative to the unconditional optimal path in treatment 3 as compared with the other two treatments. This is our explanation for the first part of Finding 4 that deviations from the unconditional optimal path are significantly larger in treatment 3 than in the other two treatments. On the other hand, as Finding 4 further indicates, the social information on average consumption provided in treatment 2 causes consumption to depart further from the conditionally optimal path due to a desire by subjects to keep up with the social norm (the Joneses), as reflected in the past average level of consumption. As the information on prior period average consumption is *not* present in treatment 1, it does not impact on consumption choices in that treatment. In treatment 3, the conditionally optimal consumption level depends on both the current period wealth (COH) *and* on the prior period past average consumption level,  $C_{t-1}$  (as this is now explicitly included in the period utility function) and consequently subjects are incentivized to take proper account of this social norm level of consumption in making their consumption decisions as that information directly affects their payoffs. As a consequence of this explicit need to condition consumption choices on the level of prior average consumption, deviations from the conditionally optimal path in treatment 3 are smaller than in treatment 2 and no worse than in treatment 1. A regression analysis of consumption behavior in treatment 3 similar to that reported in Tables 5–7 is available in Appendix A of the Supplementary Material for the interested reader.

<sup>12</sup> Of course in the next period,  $t+1$ , subject  $i$  may have an amount of cash on hand,  $COH_{t+1}^i$ , that is not optimal from the perspective of period  $t$  and thus we re-calculate subject  $i$ 's optimal path *each* period over their ever-shortening planning horizon.

## 6. Conclusions

The role of social information and habits on intertemporal consumption and savings plans is thought to be important but is poorly understood. In this paper we have taken a first step toward understanding the role of *peer* influences on consumption behavior by designing and reporting on a laboratory experiment involving a simple lifecycle consumption planning problem. Our model has a fixed interest rate and no income uncertainty making it a relatively simple framework in which to evaluate intertemporal decision making. Furthermore, all agents are identical in terms of period income and preferences, thus enhancing the attractiveness as well as the relevance of peer influence. We find in our baseline treatment 1 that subjects have difficulty solving a 25-period lifecycle consumption optimization problem despite the absence of any uncertainty and given two opportunities to go about it. In particular, they tend to save too little relative to the optimal path (i.e., they over-consume), a phenomenon that has been found in other studies as well. Importantly, the provision of social information about the average consumption decisions of other, similarly situated agents in our treatment 2 leads to further deviations from the optimal path (Finding 1) and still further deviations from the conditionally optimal path – a more forgiving metric that conditions on subjects' current available cash on hand (Finding 2). Our interpretation of these findings is that subjects over-react to the social information on average consumption that is provided in treatment 2 resulting in larger deviations from the conditionally optimal path relative to treatment 1 where such social information is absent. In other words, social information can be detrimental to intertemporal decision-making when it serves to reinforce over-consumption and reduces savings relative to the optimal path. As suggested in the introduction, the latter finding might help to account for declining savings rates in developed countries since the 1970s as information on the consumption habits of others has become more readily available. Finally, in a third, “habit formation” treatment, we explicitly incorporate concerns for consumption relative to the prior past average level of consumption into subject's utility functions as in the external habit formation literature. We find that subjects have difficulty formulating optimal consumption plans in this environment as well, though their decisions are not any further from the conditionally optimal path than in the baseline treatment 1.

Our findings should be interpreted with some caution since the environment we study is one in which there is little or no uncertainty, as all individuals face the same lifetime income streams and rates of return on savings and individuals only differ from one another at any date in time according to their accumulated wealth (cash on hand). In more realistic settings, with greater uncertainty and/or heterogeneity in incomes and rates of return on savings, it could be that the role played by social influences is much reduced as individuals recognize that such comparisons may not have much relevance to their own intertemporal consumption and savings plans. Alternatively, individuals may also become more selective about the group of peers with whom they form social comparisons. We leave the evaluation of peer effects in such richer environments to future research.

## Appendix. Supplementary Material

Appendices A and B associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jebo.2014.07.010>.

## References

- Abel, A.B., 1990. Asset prices under habit formation and catching up with the Joneses. *Am. Econ. Rev.* 80 (2), 38–42.
- Ballinger, T.P., Palumbo, M.G., Wilcox, N.T., 2003. Precautionary saving and social learning across generations: an experiment. *Econ. J.* 113, 920–947.
- Ballinger, T.P., Hudson, E., Karkoviata, L., Wilcox, N.T., 2011. Saving behavior and cognitive abilities. *Exp. Econ.* 14, 349–374.
- Brown, A.L., Chua, Z.E., Camerer, C.F., 2009. Learning and visceral temptation in dynamic saving experiments. *Q. J. Econ.* 124, 197–231.
- Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. *J. Polit. Econ.* 107, 205–251.
- Carbone, E., Hey, J.D., 2004. The effect of unemployment on consumption: an experimental analysis. *Econ. J.* 114, 660–683.
- Carbone, E., 2006. Understanding intertemporal choices. *Appl. Econ.* 38, 889–898.
- Christiano, L., Eichenbaum, M., Evans, C., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. *J. Polit. Econ.* 113, 1–45.
- Dobrescu, L.I., Kotlikoff, L.J., Motta, A., 2012. Why aren't developed countries saving? *Eur. Econ. Rev.* 56, 1261–1275.
- Duesenberry, J.S., 1949. *Income, Saving, and the Theory of Consumer Behavior*. Harvard University Press, Cambridge, MA.
- Fehr, E., Zych, P.K., 1998. Do addicts behave rationally? *Scand. J. Econ.* 100, 643–662.
- Fehr, E., Zych, P.K., 2008. Intertemporal choice under habit formation. In: Plott, C.R., Smith, V.L. (Eds.), *Handbook of Experimental Economics Results*, vol. 1. Amsterdam, North-Holland, pp. 923–928.
- Feltovich, N., Ejebu, O.-Z., 2013. Do Positional Goods Inhibit Saving? Evidence From a Life-cycle Experiment, Working Paper. Monash University.
- Friedman, M., 1957. The permanent income hypothesis. In: Friedman, M. (Ed.), *A Theory of the Consumption Function*. Princeton University Press (for NBER), Princeton, NJ, pp. 20–37.
- Hey, J.D., Dardanoni, V., 1988. Optimal consumption under uncertainty: an experimental investigation. *Econ. J.* 98, 105–116 (supplement).
- Hey, J.D., 1988. A pilot experimental investigation into optimal consumption under uncertainty. In: Maital, S. (Ed.), *Applied Behavioural Economics*. New York University Press, New York.
- Johnson, S., Kotlikoff, L.J., Samuelson, W., 2001. Can people compute? An experimental test of the life-cycle consumption model. In: Kotlikoff, L.J. (Ed.), *Essays on Saving, Bequests, Altruism, and Life-cycle Planning*. MIT Press, Cambridge, MA, pp. 335–385.
- Meissner, T., 2014. Intertemporal Consumption and Debt Aversion: An Experimental Study, Working Paper. Technical University of Berlin.
- Modigliani, F., Brumberg, R.H., 1954. Utility analysis and the consumption function: an interpretation of cross-section data. In: Kurihara, K.K. (Ed.), *Post-Keynesian Economics*. Rutgers University Press, New Brunswick, NJ, pp. 388–436.
- Schmitt-Grohé, S., Uribe, M., 2008. Habit persistence. In: Durlauf, S.N., Blume, L.E. (Eds.), *The New Palgrave Dictionary of Economics*, 2nd ed. Palgrave Macmillan, Basingstoke, UK.
- Veblen, Thorstein, 1899 [1973]. *The Theory of the Leisure Class*. Boston, MA, Houghton Mifflin.