Do positional goods inhibit saving?
Evidence from a life–cycle experiment

Nick Feltovich*
Department of Economics, Monash University
Clayton VIC 3800, Australia
+61 3 9902 0790
nicholas.feltovich@monash.edu

Ourega-Zoé Ejebu
University of Aberdeen Business School
Edward Wright Building
Aberdeen AB24 3QY, United Kingdom
zoejebu@yahoo.co.uk

January 24, 2014

Abstract

We investigate the effect of *positional goods* (goods for which one's consumption relative to others' matters) on saving, based on results from a life–cycle consumption/saving experiment. In a *Group* treatment, we allow inter–personal comparisons by assigning subjects to groups and displaying rankings based partly on consumption. A baseline *Individual* treatment is similar, but without the additional information. We find more *under–saving* (saving less than the optimal amount), and lower money earnings for subjects, in the Group treatment. Both effects are economically relevant, with magnitudes of roughly 6–7% of expected income and 7–8% of average earnings respectively. Additional analysis shows that the result is driven by those subjects who are *not* ranked in the top three in their group ("keeping up with the Joneses"), and males in particular.

Journal of Economic Literature classifications: D91, E21, D14, D03, G02.
Keywords: consumption, precautionary saving, inter–personal comparisons, positional goods, experiment.

*Corresponding author. Some of this research took place while Feltovich was at University of Aberdeen. Financial support from University of Aberdeen's College of Arts and Social Sciences is gratefully acknowledged. We thank Philip Grossman, Anmol Ratan, Joe Swierzbinski, participants at several conferences and seminars, the guest editors of this special issue, and three anonymous referees for helpful suggestions and comments that have greatly improved the paper.

1 Introduction

One of the main contributing factors of the recent global financial crisis and subsequent recession was the sub–prime mortgage crisis: a rise in mortgage defaults in many countries, notably the United States. Many explanations have been proposed for this rise in defaults (see, e.g., Demyanyk and Van Hemert (2011)), but one of the most important was the failure of many households to build up sufficient savings to withstand an adverse event. The US savings rate in 2005, at 1.5%, was well below those in Continental European countries such as France (11.4%), Germany (10.5%) and Italy (9.9%), though comparable to Canada at 2.2% and the UK at –1.2% (OECD (2011)). Low savings rates in English–speaking countries are not new (Elmer and Seelig (1998)), but their effects were fairly well hidden until growth in incomes and housing prices slowed markedly around this time.

It is worthwhile to understand why savings rates are so low in the US and other English–speaking countries. A recent book by Garon (2011) attributes the low US savings rate to a difference in culture relative to other countries: "In Germany, where households sock away much more than Americans, saving is sexy...[i]n Japan, which until recently boasted the world's highest household saving rates, saving was rarely sexy. But it has been stylish." (p. 1) The implication is clear: in the US, and by extension other English–speaking countries, saving is not sexy or stylish; consumption is. This point was made much earlier by Duesenberry (1949):

> "For any particular family the frequency of contact with superior goods will increase primarily as the consumption expenditures of others increase. When that occurs, impulses to increase expenditure will increase in frequency, and strength and resistance to them will be inadequate. The result will be an increase in expenditure at the expense of saving" (p. 27).

This phenomenon, which Duesenberry named the "demonstration effect", and which we often call "keeping up with the Joneses", is actually a combination of two distinct factors. First, consumption is a *positional good* (Hirsch (1976), Frank (1985a)): people care not only about how much they consume in an absolute sense, but also about where their consumption stands in comparison to others'. By contrast, saving – which is typically unobservable – is non–positional (Frank (1985b), Kosicki (1988)).

However, positional consumption alone is not sufficient to give rise to the demonstration effect. As Harbaugh (1996) has pointed out, rational (forward–looking) consumers with such preferences could actually do the opposite of what Duesenberry suggests, as the fear of "falling behind the Joneses" after a future negative income shock would reinforce an individual's precautionary motive for saving in the present. The demonstration effect additionally requires a degree of *myopia*: individuals, striving to emulate others' consumption in the present, ignore or under–weight the negative consequences their choices imply for the future. Since the main such consequence is lower saving leading to lower future consumption, such striving is actually counter–productive: consuming more now in an attempt to keep up with the Joneses makes one even more likely to fall behind them later!

There is growing evidence that individuals' preferences reflect relative as well as absolute economic variables.[1] Angelucci and DeGiorgi (2009) find that a programme of cash welfare payments to villagers in Mexico leads to increased consumption by those households that were *not* eligible for the payments. Kuhn et al. (2011) report that neighbours of lottery winners in the Netherlands are more likely to buy cars in the six months after the lottery date, and find a similar, though weaker, effect on purchases of exterior home renovations; however, they find no corresponding increases for nearly all other goods. Card et al. (2012) report that giving workers in the US access to

[1] There is less evidence that saving in particular is non–positional, though Kosicki (1987, 1988) shows that a model of positional consumption and non–positional savings can explain the relationship between saving and permanent income.

information about colleagues' salaries leads those with below–median salaries to have lower job and pay satisfaction. Taken together, these studies suggest that many people value relative as well as absolute standing, though clearly a variable needs to be observable in order for them to make inter–personal comparisons with respect to it.

A disadvantage of studies using field data, however, is that it can be difficult to identify the cause(s) of the observed pattern of behaviour; for example, lottery winners may share winnings with friends and family, some of whom might be neighbours. Angelucci and DeGiorgi (2009) attribute their results to increased gifts and loans as well as informal credit and insurance markets, and Kuhn et al. (2011) spend substantial time ruling out similar explanations for their results. Laboratory experiments may help address this difficulty, as they give the experimenter a much greater degree of control over decision variables (e.g., gifts to neighbours can be ruled out), and better data collection (consumption is known, rather than noisily imputed or elicited). As a result, there has been substantial experimental research into positional goods.[2] This research has produced the following stylised facts:

1. Positional concerns matter.
2. Positional concerns vary across goods; in general, goods that are more observable tend to be more positional.
3. People show more concern for those with positions above them than below them.
4. People show more concern for those who are "near" them in some sense (e.g., neighbours or co–workers).

The aim of this paper is to investigate the possibility that positional consumption can lead to low levels of saving. We do so by designing and conducting a laboratory experiment using a simple version of the life–cycle savings model (Modigliani and Brumberg (1954)), which is one of the work–horse models in theoretical and empirical macroeconomics (e.g., Ando and Modigliani (1963), Samuelson (1975), Summers (1981), Skinner (1988), Carroll (1997)). Subjects play the role of individuals, who know that they will live for 60 rounds ("years" in the experiment). In each round they receive a random income, then choose how much of their available funds to consume, with the remainder being saved. Consumption yields utility according to a concave per–round utility function; savings yield no utility directly, but serve as a buffer stock against low incomes in future years.

Optimal behaviour in this environment involves building up substantial savings over roughly the first half of "life", then dis–saving during the second half. However, previous research has shown that this decision problem is complex and difficult to solve, both theoretically in the sense that optimal behaviour must be approximated by numeric computations rather than derived analytically, and behaviourally in the sense that human subjects make frequent mistakes when placed into this situation. Experimental testing of this model dates back to the late 1980s, with Johnson, Kotlikoff and Samuelson's (1987) unincentivised questionnaire and Hey and Dardanoni's (1988) incentivised experiment. These two studies found systematic *over–saving* relative to optimal levels. Since then, though, the predominant result has been *under–saving* (Ballinger et al. (2003), Carbone (2005), Brown et al. (2009), Kitamura et al. (2012)), though Anderhub et al. (2000) observe approximately optimal levels of saving, and Carbone and Hey (2004) find evidence of both over– and under–saving.[3] Many of these studies have explained their results

[2]See, for example, Solnick and Hemenway (1998), Alpizar et al. (2005), Ferrer-i Carbonell (2005), Carlsson et al. (2007), Andersson (2008), and Grossman and Komai (2013) for particular studies, and see Knell (1999) for a discussion of the broader literature behind these stylised facts. Readers wishing to gather their own stylised facts might refer to Damianov (2011), who describes a classroom experiment, along with results from his own classes.

[3]Two other studies with related decision–making settings are relevant. Fatás et al. (2007) find over–saving in early rounds of an experiment with constant per–period income and uncertainty over the number of rounds (i.e., when the individual would "die"), while Ponzano and Ricciuti (2012) find under–saving in an experiment with production, exchange and investment as well as consumption and saving.

as arising from bounded rationality: inexperience (Ballinger et al. (2003), Brown et al. (2009), use of heuristics (Carbone (2005), Müller (2001)), truncated planning horizons (Ballinger et al. (2003), Carbone and Hey (2004), Carbone (2005)), and dynamically inconsistent preferences (Brown et al. (2009)).

In order to assess the effects of positional goods, our experiment has two treatments. In a baseline *Individual* treatment, subjects make their decisions in isolation; the only feedback they receive is a recapitulation of their own choices and the resulting outcomes. In a *Group* treatment, subjects are randomly assigned to groups with between five and seven members, and each subject is additionally shown the lifetime utilities of the three group members with the highest lifetime utilities (up to the current year), along with the subject's own rank. Other than this extra information, the Group and Individual treatments are identical.

Our procedures represent a fairly conservative attempt to frame consumption as positional. Subjects are paid based only on their induced lifetime utility function; there are no monetary or other physical incentives based on rank. However, there is previous evidence that people value high status per se – even in settings like ours where it is abstract in nature and yields no monetary rewards – to the point where they will actually sacrifice money for status (Huberman et al. (2004), Charness et al. (2014)). (Thus, rank is *salient* in the way the term is used by psychologists, but not in the way it is used by experimental economists.) Additionally, the previous studies mentioned above suggest that subjects' life–cycle decisions reflect substantial myopia. Under these circumstances, we would expect Duesenberry's demonstration effect to arise, with higher consumption and lower saving in the Group treatment than in the Individual treatment.

This is exactly what we find. Both non–parametric statistical tests on aggregate data, and standard regressions based on individual decisions, find that under–saving (saving less than the optimal amount) is more prevalent in the Group treatment than in the Individual treatment. Consistent with myopia, this under–saving is actually counter–productive in the long run: earnings (i.e., the actual amounts of cash subjects receive) are weakly significantly lower in the Group treatment. Both of these treatment effects are economically relevant, with the effect on under–saving on the order of 6–7% of expected income and the effect on earnings equal to roughly 7–8% of average earnings. (On the other hand, our evidence concerning over–saving is more equivocal; as expected, it is lower in the Group treatment than in the Individual treatment, but the difference is not significant.) Additional regressions indicate an interaction between our allowing inter–personal comparisons of consumption and subjects' positions in the rank–ordering. Specifically, we find that the effect of the Group treatment on under–saving is driven by lower–ranked subjects (i.e., those outside the top three in their group), as if they were trying to "keep up with the Joneses". The effect is particularly strong (under–saving of roughly 18% of expected income), and highly significant, for males, and less pronounced (9% of income) but still sometimes significant for females.

2 The experiment

The experiment implements, with some minor changes, the decision–making environment used by Ballinger et al. (2003), though our setting is also similar to most of the other life–cycle saving experiments discussed in the introduction. We present the theory first, followed by the experimental design, procedures and hypotheses.

2.1 The decision problem

There are 60 rounds (called *years* in our experiment). An individual begins with an endowment of 20,000 units of currency (called *francs*). In round t, the individual receives income y_t, which is drawn i.i.d. from a known distribution: equally likely to be 0 or 80,000 francs. After the income draw, the individual chooses a non–negative level of consumption c_t for the round, with the remainder saved for the next round. There is no borrowing, so it is not possible to consume more in a round than the individual's available funds (current income plus any savings from previous rounds). Individuals receive utility from their own consumption, but not (directly) from savings. The lifetime utility function is additively separable, with the corresponding per–round utility function given by

$$u(c_t) = k + \theta \frac{[\alpha(\epsilon + c_t)]^{1-\sigma}}{1-\sigma}.$$

This is a minor variation of a constant relative risk aversion (CRRA) utility function with coefficient σ. The parameter ϵ represents a fixed, exogenous source of income, and is restricted to be positive in order for u_t to be defined even when c_t is zero. The parameters $\alpha > 0$, $\theta > 0$ and k represent an affine transformation on the utility function, and thus have no theoretical impact on decision making; we use these simply to scale the utility function for the experiment, where we induce this utility function on subjects by making their cash earnings proportional to their lifetime utility. In the experiment, we set $\sigma = 3$, $\epsilon = 30000$, $\alpha = 1/10000$, $k = 34$ and $\theta = 1200$; the resulting utility function is shown in Figure 1.[4]

Figure 1: Utility function

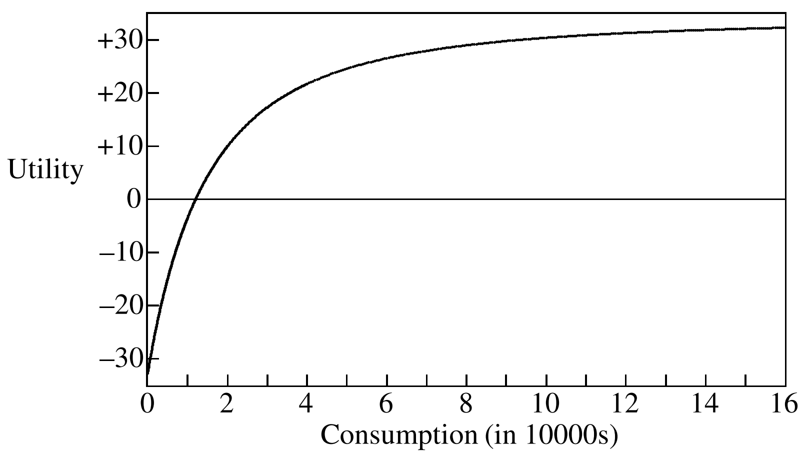

The individual's dynamic choice problem is

$$\max_{(c_t)} \left[u(c_1) + E_1 \sum_{t=2}^{60} u(c_t) \right] \text{ subject to } \begin{cases} c_t \in [0, w_t + y_t]; \\ w_t + y_t - c_t = w_{t+1}; \\ w_1 = 20000; \end{cases} \tag{1}$$

where w_t is wealth at the beginning of round t (i.e., the return from the previous round's savings), and $Prob(y_t = 0) = Prob(y_t = 80000) = \frac{1}{2}$. Note that there is a zero discount rate and a zero interest rate on savings; positive values for both of these would be more general (though zero real interest rates are approximately what we see in many developed countries currently), but would add additional complexity to an already complex decision problem.

The solution to this problem – the optimal choice of consumption (or equivalently savings) in each round – depends on the realised values of y_1, y_2, ..., y_{60}. For the experiment, we drew two different income streams in advance (see Table 1), and used each in about half of the experimental sessions. Compared to drawing new income

[4]As Ballinger et al. (2003) point out, either a no–borrowing constraint or a utility function with a positive third derivative would imply a precautionary motive for saving. Our setup, like theirs, has both.

Table 1: Characteristics of income streams used in the experiment

	Stream 1	Stream 2
Sequence of incomes, rounds 1–60 (1 = 80000)	1100111010	1011100010
	1111101101	0110011110
	1111101001	0011010001
	0010100101	0110011010
	0000101000	1100001111
	0100111100	0010011001
Number of 80000 incomes, rounds 1–10	6	5
rounds 11–20	8	6
rounds 21–30	7	4
rounds 31–40	4	5
rounds 41–50	2	6
rounds 51–60	5	4
all rounds	32	30
Avg. cash balance, optimal policy (000s)	257.46	164.03
Earnings from optimal policy (£)	22.17	21.24

streams for each session (or individual), our design facilitates comparisons across sessions, since all subjects facing Stream 1 have exactly the same opportunities, and likewise for Stream 2. For the same reason, it makes information about other subjects' performance more relevant, since any differences across subjects are the result of different decisions rather than different opportunities.[5] Stream 1 represents a more difficult problem for experimental subjects than Stream 2, as the accumulation of good rounds in the first half may provide a strong temptation to consume excessively, and the preponderance of bad rounds in the second half is likely to punish the kinds of myopic strategies that have often been seen in this literature (see the introduction).[6]

As usual for these models, there is no explicit analytical solution to our dynamic choice problem. Instead, we use numerical methods to find approximate solutions. Optimal behaviour is computed round–by–round, starting from the final round (where individuals should consume everything), and working backwards, using the optimal choices in round t for each possible level of available funds (and associated continuation lifetime utilities) to find the optimal choices in round t–1, until reaching the first round.[7] The results for the specific income streams we use are characterised in Figure 2 (with some additional information in Table 1). The figure shows the optimal balances

[5]Using the same income stream in multiple sessions raises the possibility that subjects taking part in later sessions might get information about the streams from friends who took part in earlier sessions. However, this seems not to have happened to any substantial extent, as regressions of earnings – similar to the ones reported in the right–most columns of Table 4, but with the session number included on the right–hand side – found the session number to have no significant effect. (Results available from the corresponding author upon request.)

[6]On the other hand, Stream 1 may be more important from a policy standpoint, since its pattern of incomes resembles a spell of unforeseen unemployment later in life.

[7]Our solution is approximate because we round levels of available funds, and savings amounts, to multiples of 100 francs, rather than letting them take on values of any multiple of 0.01 francs as in the experiment. This was done to allow computation of the (approximately) optimal policy without excessive demands on computing resources (calculating the exact policy would have involved manipulating a matrix with over 200,000,000 rows, as opposed to the 21,000 rows we needed for the approximate solution).

Figure 2: Optimal end–of–round cash balances (approx.)

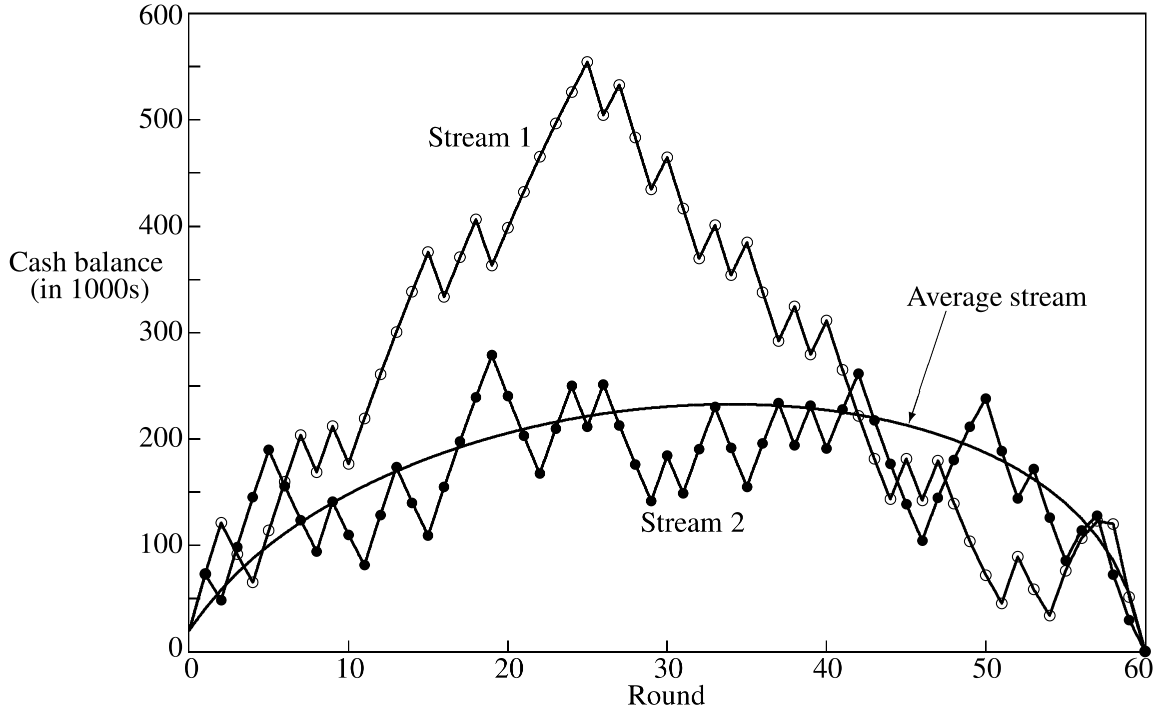

at the end of each round for each income stream. Also shown, for comparison, are average optimal balances based on one million randomly chosen income streams; this serves as an indication of optimal behaviour when facing a "typical" income stream. Under both streams we use, individuals should build up a buffer stock of savings over the first half of life – at its peak, almost 14 times expected per–round income for Stream 1 and about 7 times for Stream 2 – then dis–save over the second half of life.

Our setup is clearly a simplification of the life–cycle saving problem faced by real individuals. Besides the zero rate of interest on savings and the lack of access to loans, mentioned already, the income uncertainty is of a very simple form (i.i.d. and binary), there is no bequest motive, and no uncertainty about when "death" will occur. Many of these features could be changed to make our experiment more like real life, but at the likely cost of making the decision problem even more challenging for subjects, and leading to more decision errors, though in some cases, additional realism may actually benefit the subjects (e.g., changing the distribution from which income is drawn to a more realistic one might make the environment more forgiving of these errors).

2.2 Experimental treatments

In the experiment, human subjects faced this life–cycle decision problem with real monetary incentives; specifically, they were paid proportionally to their lifetime utility according to Equation (1) above. We use a 2x2 factorial design; besides the income stream faced by subjects (Stream 1 or Stream 2), our experiment varies the amount of feedback they receive. In the I[ndividual] treatment, subjects only receive information about their own performance. In the G[roup] treatment, subjects receive not only the I–treatment information, but also some information about other subjects' performance. Specifically, each G–treatment session is divided into groups of 5–7 subjects, and each subject in a group is assigned an ID number.[8] At the end of each round, the computer program ranks the subjects in each group according to their lifetime utilities up to that point (i.e., $\sum_{t=1}^{s} u_t(c_t)$ in round s), and the ID numbers and lifetime utilities of the three subjects in the group with the highest lifetime utilities are displayed on the screens of all of that group's members until the end of the following round, when they are updated. Subjects are also shown their own ID number and rank. This information is *salient* in the way psychologists use the term – it receives more attention than alternative pieces of information (though subjects are not told that they should value their position relative to others, how they should act on this information, or even that they should find it useful) – but it is not salient in the way experimental economists use the term, since a subject's rank has no direct impact on her monetary earnings from the experiment, nor any other kind of physical reward.

The G treatment is designed specifically to allow the kinds of inter–personal comparisons that play a role in decision making outside the lab (see the introduction for details). Observation of other subjects is limited to those in the same group, reflecting the likelihood that one's reference group tends to comprise people "near" the individual, such as relatives, friends, neighbours or co–workers. Rather than observing everyone in the group, subjects only see those group members with the highest lifetime utilities, capturing the tendency to devote more mental resources to envy of the better off, rather than relief at not being among the worse off. Having the ranking based on utility up to the current period, which depends on past and present consumption spending but not directly on savings, reflects the positional nature of consumption relative to savings, as well as the fact that many positional goods are durable, so that both past and present consumption spending matter. Basing the ranking on utility instead of explicitly on consumption reflects the noisy way that individuals' consumption is observed by others, and also makes the connection between behaviour and rank less transparent to subjects, lessening demand effects.[9]

2.3 Experimental procedures

The experimental sessions took place at the Scottish Experimental Economics Laboratory (SEEL) at the University of Aberdeen. Subjects were primarily undergraduate students from University of Aberdeen, with just over half (53%) females, and were recruited using the ORSEE web–based recruiting system (Greiner (2004)). No–one took part more than once. The experiment was run on networked personal computers, and was programmed using the z–Tree experiment software (Fischbacher (2007)). Subjects were asked not to communicate with other subjects except via the computer program. In the I treatment, no information was given to subjects about others' behaviour or results. In the G treatment, information about others was based on ID numbers that were randomly assigned at the beginning of a session, so that the behaviour or results of specific individuals could not be traced to their identities.

At the beginning of a session, subjects were seated in a single room and given written instructions.[10] The instructions were also read aloud to the subjects, in an attempt to make the rules of the experiment common knowledge. After the instructions were read, subjects' computer screens displayed a few demographic questions, which they

[8] Our instructions to subjects stated that group sizes could be between 5 and 9. This larger range was stated in order to insure against the possibility that a large number of no–shows would have left us with either 8 or 9 subjects, which would have made it impossible to form groups of 5–7. As it happened, groups of 5–7 were always possible, since all sessions had at least 14 subjects. We note that subjects were always informed of the true size of their group, from the first round on.

[9] A recent working paper by Carbone and Duffy (2013) examines the effect of giving subjects information about the consumption decisions themselves. Carbone and Duffy are primarily interested in deviations from optimal behaviour, rather than treatment effects, but their results seem to replicate ours in that the extra information raises consumption. A third treatment suggests that consumption does not increase further when subjects are paid directly based on relative consumption.

[10] The text of the instructions and some sample screen–shots are available as supplementary online material through the publisher's website. Other experimental materials, such as the z–Tree programs and the raw data, are available from the corresponding author upon request.

were asked to answer. Once all subjects had done so, the first round began.

Each round began with subjects being reminded of their savings from the previous round. After clicking a button, each was shown the realisation of income for the current round and prompted to choose that round's consumption, which could be any multiple of 0.01 francs between zero and the subject's available funds (savings plus current income), inclusive. Subjects were given 45 seconds (60 in the first three rounds) to enter their consumption, and were told in the instructions that a zero amount would be entered if time ran out.[11] During this time, part of the computer screen showed a "francs–to–pounds conversion calculator", where subjects could enter any allowable consumption choice and view the corresponding utility amount. Information about the utility function was also shown in the instructions, as both a table and a chart. Once all subjects had entered their consumption choices, they received feedback comprising consumption, savings, and utility for the round, and total and average utility up to the current round; this same information from each round was available in a history table at the top of subjects' screens. In the G treatment, a section of each subject's computer screen showed the top–3 table described in Section 2.2, along with the subject's own ID number and rank; in the I treatment, this portion of the screen remained blank.

At the end of the sixtieth round, subjects were paid, privately and individually. Each subject received, in British pounds, their average utility over all rounds (i.e., 1/60 times their lifetime utility); there was no show–up fee.[12] Parameters of the utility function were chosen so that a completely myopic strategy – consuming all available funds in every round – would earn nearly zero, while optimal choices would earn just over £20 (see Table 1). In fact, total earnings averaged £17.06 (with a median of £18.36), for a session that typically lasted about 90 minutes.[13]

2.4 Hypotheses

Our null hypothesis is that the additional information in the G treatment does not systematically affect behaviour.

Hypothesis 1 *There are no systematic differences between the I and G treatments.*

This hypothesis is entailed by rational behaviour, but full rationality is by no means required for it to be true. Indeed, given the complexity of the decision problem, it is unlikely that any subject will behave optimally in all rounds. However, if deviations from optimal behaviour are due to mistakes in calculating the optimal policy, there would be no reason to expect consumption or saving to differ systematically from the optimal policy, or across treatments.

Our alternative hypothesis is based on the possibility that subjects treat consumption as a positional good in the myopic manner described by Duesenberry (1949), so that they will tend to consume more and save less, earning lower payoffs in the long run.

Hypothesis 2 *Saving and payoffs are lower in the G treatment than in the I treatment.*

[11] Such "forced zeros" account for about 0.6% of observations, and as discussed in Note 19, the analysis is unaffected by whether these observations are included in the sample.

[12] We note that our experiment does not control for subjects' risk attitudes. The theoretical predictions presented in this paper are for risk–neutral subjects, on whom we've induced the expected–utility function shown in Section 2.1; utility maximisers with different risk attitudes will behave differently. However, individuals who are *more* risk averse will optimally consume less and save more, so in the probable case that most subjects are risk averse, our experiment will under–state the overall degree to which they under–save (relative to optimal behaviour). Of course, deviations from risk neutrality would not, on their own, imply any difference between our Individual and Group treatments.

[13] These averages include two subjects – one in each treatment – who incurred losses. As usual in lab experiments, we could not collect these subjects' losses from them, so their true earnings were zero. See Table 4 and the surrounding text for some evidence that our results are robust to whether these two outlier subjects are included in the analysis. We note that one unforeseen benefit of our restriction of observation in the G treatment to the three highest lifetime utilities is that a subject making a series of extremely sub–optimal choices is unlikely to be observed after the first few rounds.

3 Results

Table 2 shows the cells of our experimental design, and Table 3 displays some aggregate results. We consider five

Table 2: Treatment information

Feedback	Income stream	Sessions	Groups	Subjects
Individual	1	2	—	35
	2	2	—	29
Group	1	3	9	54
	2	3	8	52

measures of behaviour. One of these, "average earnings", is simply the actual cash amount earned by subjects at the end of an experimental session, in British pounds. The remaining four statistics compare observed saving behaviour with optimal behaviour given the subject's current available funds (savings plus the current round's income) and the round.[14] "Under–saving" is defined as the extent of sub–optimal saving, and is calculated as the amount actual saving is below the optimal saving choice, cut off at zero if actual saving is above or equal to optimal saving. Similarly, "over–saving", the extent of super–optimal saving, is the amount actual saving is above the optimal saving choice – or zero if less than or equal. "Saving" is measured as the difference between the actual amount saved and the optimal choice, and can be positive, negative or zero. (So, saving is equal to over–saving minus under–saving.[15]) "Absolute deviation" is measured as the absolute value of the difference between actual and optimal saving (i.e., the sum of under– and over–saving); this provides a measure of overall decision errors without regard to their direction. All of these measures except for earnings are in lab–francs, normalised to multiples of expected income by dividing by 40000, pooled over income streams (the regressions later in this section will control for the income stream) and reported both over all rounds and over the first half of the experiment (rounds 1–30), the latter being where subjects would be expected to build up savings. The table also displays p–values from non–parametric robust rank–order tests of differences between the two treatments, using the session–level data.[16]

There are clearly systematic differences between treatments: the Group treatment has more under–saving and lower earnings than the Individual treatment. These two effects are not only statistically significant at the 10% level or better, but also economically relevant; for example, average subject earnings are about 8% higher in the I treatment than in the G treatment. Of the other variables, only saving over all 60 rounds shows a (weakly) significant treatment effect, with a lower average in the Group treatment as expected.

[14] Like previous researchers, we distinguish between the ex ante optimal policy – based on optimal choices at every point from the beginning – and optimal behaviour given the subject's current available funds and the round. In particular, if a subject had previously deviated from the ex ante optimal policy, it will generally not be optimal to return to that policy. Our interest is in whether subjects are making optimal decisions given their circumstances, not whether they eventually return to a policy that had been optimal but is now suboptimal due to a prior deviation. Hence, we define "saving", "under–saving", "over–saving", and "absolute deviation" relative to the optimal saving amount given the subject's current available funds and the round, not the ex ante optimal policy. Also, "optimal" means conditional on the information the individual actually has (the round and her current available funds), and specifically *not* the realisation of future income draws. Indeed, a clairvoyant individual (or a lucky one) could achieve a higher lifetime utility than optimal behaviour would predict.

[15] The reason for treating over– and under–saving separately, rather than looking at saving exclusively, is that without this distinction, the welfare implications of changes to saving are ambiguous: a decrease in saving is welfare–reducing if it is an increase in under–saving, but welfare–enhancing if it is a decrease in over–saving.

[16] See Siegel and Castellan (1988) for descriptions of the non–parametric statistical tests used in this paper.

Table 3: Treatment aggregates

	Under–saving		Over–saving		Saving		Absolute deviation		Earnings
Rounds:	All	1–30	All	1–30	All	1–30	All	1–30	All
I treatment	0.264	0.371	0.135	0.076	–0.129	–0.295	0.398	0.447	17.88
G treatment	0.308	0.430	0.112	0.075	–0.195	–0.355	0.420	0.505	16.57
p–value*	0.012	0.087	0.36	0.50	0.076	0.35	0.28	0.20	0.087

**p–values refer to significance of differences between I and G treatments (see text for details).*

Figure 3 displays round–by–round average cash balances for all combinations of treatment and income stream (recall that Stream 1 had many high–income rounds in the first half and many low–income rounds in the second half, while Stream 2 was more "typical", with high and low incomes more evenly spread), along with the corresponding optimal cash balances. Overall, subjects save too little: over the middle third of the session (rounds 21–40), average

Figure 3: Average end–of–round cash balances (in 1000s), Individual and Group treatments

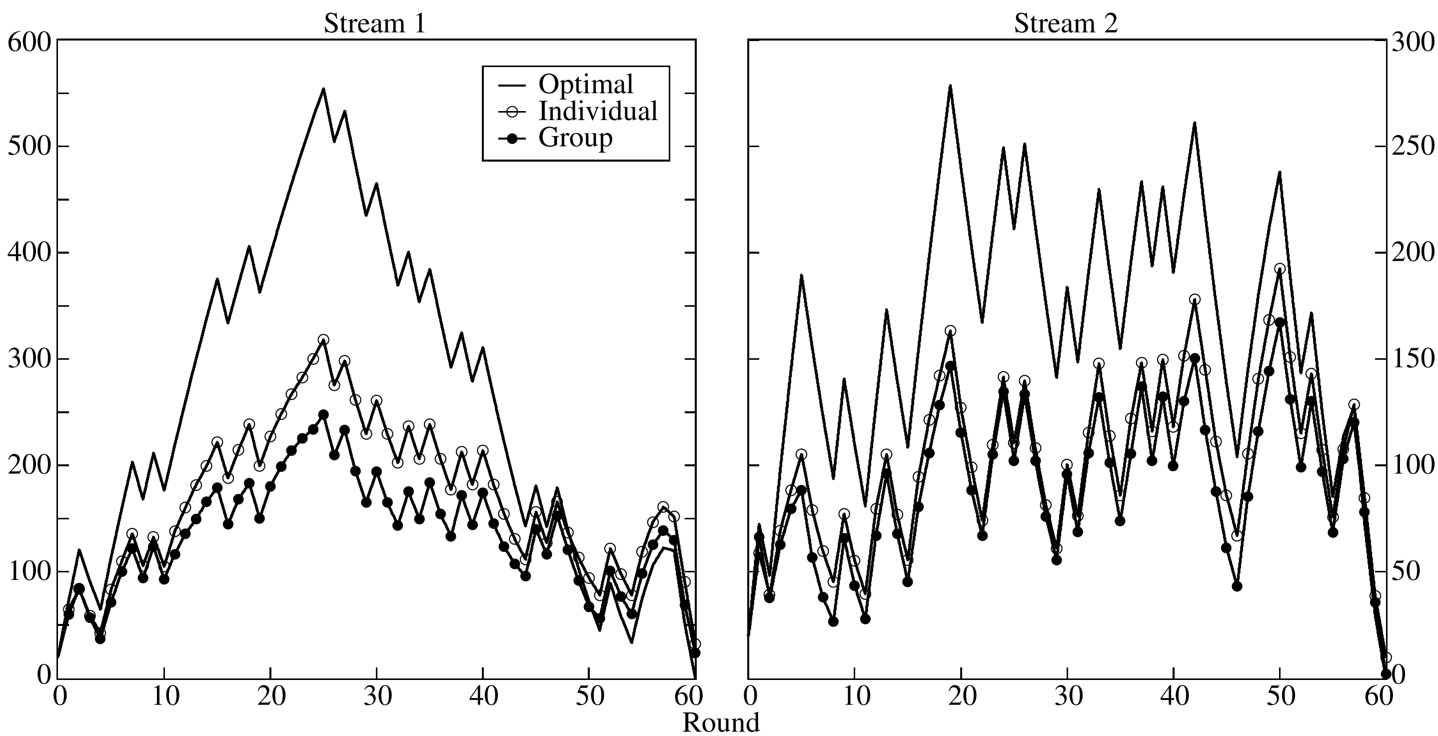

savings are only about half of what the optimal policy implies. Saving is visibly higher in the I treatment than the G treatment (averaging 58.1% versus 44.5% of the optimal amount over this time period in Stream 1, 56.1% versus 51.0% in Stream 2). These qualitative differences (more saving in I than G, less in both than under the optimal policy) are also present – though smaller qualitatively – in the first and last thirds of the session.

Figure 4 displays, for each individual subject, the number of rounds – out of the total of 60 – he/she under–saved (as the vertical distance from the horizontal axis) and the number he/she over–saved (the horizontal distance from the vertical axis). There is substantial between–subject heterogeneity: individual subjects range from over–

Figure 4: Individual subjects' number of rounds with over– and under–saving

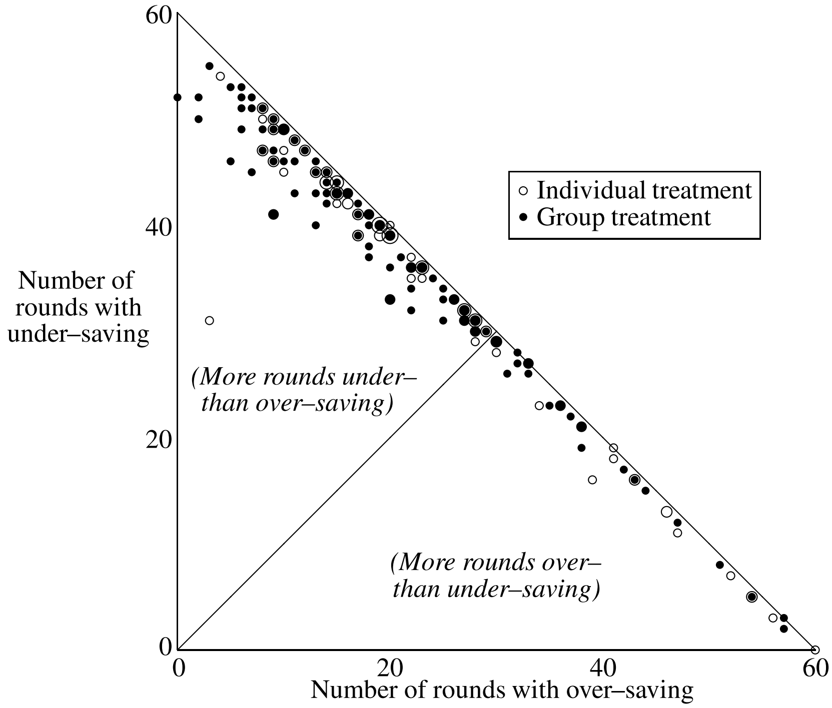

saving in every single round to under–saving in nearly all rounds, though under–saving clearly predominates in both treatments. Figure 5 shows that this heterogeneity is payoff–relevant, with a scatter–plot of each individual subject's under–saving in the first half of the experimental session (where subjects should be building up a buffer stock) and money earnings for the complete session (i.e., lifetime utility). The relationship between these is clear: more under–saving tends to be associated with lower earnings. This means that the differences in under–saving and earnings between G and I treatments seen in Table 3 are not coincidental; they are implications of the same treatment effect.

We next move to regressions. We estimate five pairs of models, with each pair having as the dependent variable one of the measures used in Table 3; within a pair of models, the only difference within a pair is whether or not two outlier subjects (the two with negative money earnings for the session, mentioned earlier in Note 13) are removed. For the equations with under–saving, over–saving, saving and absolute deviation as the dependent variable, the main right–hand–side variables are an indicator for the Group treatment (versus the Individual treatment), an indicator for Stream 2 (vs. Stream 1), the round number, and its square. We also include all two– and three–way interactions of these variables, to control for differing treatment effects according to the income stream and over time.[17] We use Stata (version 12) to estimate Tobit models for under–saving, saving and absolute–deviation, while we use a linear model for over–saving.[18] These models were estimated based on rounds 1–59 of the data–set (we leave out round 60, where rational subjects should consume all wealth, irrespective of its initial level), and incorporate individual–subject random effects. The model with earnings was similar, except that it used 60–round average utility (i.e., actual money earnings) as the left–hand–side variable, so there was only one data point for each subject. Thus, the round

[17] Leaving out the interaction terms, here and in Table 5 below, does not qualitatively change our results.

[18] We originally set out to use a Tobit for over–saving as well (with endpoints zero and the optimal consumption amount), but Stata was unable to maximise its likelihood function. We therefore report results from a linear model which, even given its obvious drawbacks, is arguably better than no model at all. The endpoints for the under–saving model are zero and the optimal savings, those for the saving model are minus the optimal savings and the available funds minus the optimal savings, and the absolute–deviation model has a left endpoint of zero and no right endpoint.

Figure 5: Individual subjects' average under–saving over rounds 1–30 and total money earnings from experiment

(Note: One point at (0, –25.07) not shown)

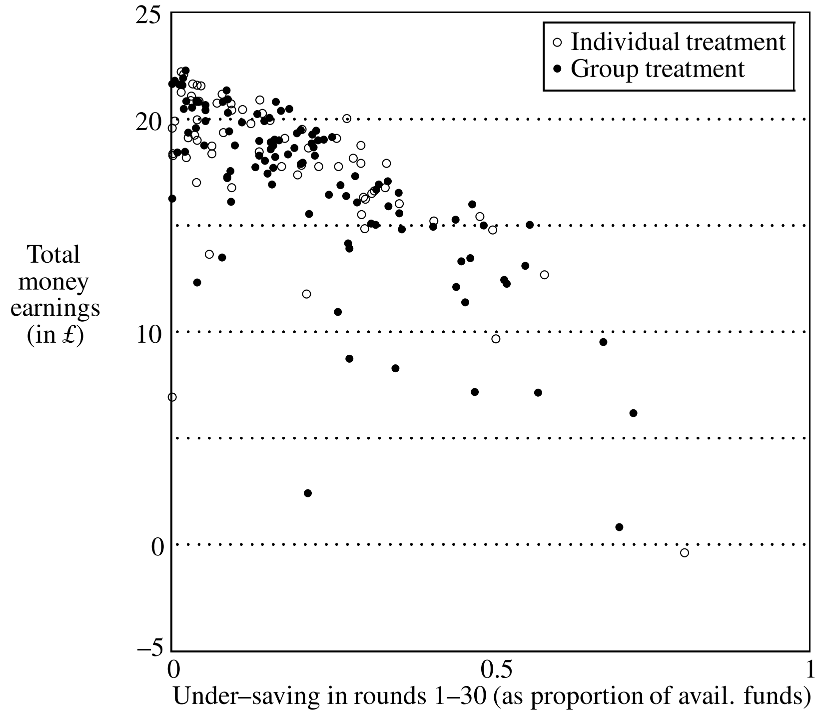

number and its interaction terms were omitted from the right–hand side, and a linear model was estimated.

Table 4 reports estimated average marginal effects (over observed variable values) and standard errors for each variable.[19] These results provide further support for the conclusions we drew from Table 3. When all subjects are

Table 4: Regression results – marginal effects, with standard errors in parentheses

Dependent variable:	Under–saving		Over–saving		Saving		Absolute deviation		Earnings	
Group treatment	0.058*	0.071**	–0.017	–0.044	–0.027	–0.040**	0.023	0.008	–1.317	–1.188**
	(0.031)	(0.029)	(0.042)	(0.028)	(0.022)	(0.017)	(0.032)	(0.023)	(0.812)	(0.593)
Income stream 2	–0.038	–0.036	–0.052	–0.023	–0.006	0.002	–0.085***	–0.064***	–0.336	–0.353
	(0.031)	(0.029)	(0.040)	(0.027)	(0.021)	(0.016)	(0.031)	(0.022)	(0.787)	(0.575)
Round	–0.006***	–0.006***	0.003***	0.002***	0.004***	0.004***	–0.003***	–0.003***		
	(0.0003)	(0.0003)	(0.0002)	(0.0002)	(0.002)	(0.0001)	(0.0002)	(0.0002)		
Outliers removed?	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Constant term?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Interaction terms?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	10030	9912	10030	9912	10030	9912	10030	9912	170	168
$\|ln(L)\|$	7781.46	7476.80	2626.47	107.18	8659.83	7472.86	7300.13	6044.68	516.93	457.09

* (**,***): *Coefficient significantly different from zero at the 10% (5%, 1%) level.*

[19] As mentioned in Section 2.3, the computer program imposed a zero consumption level if a subject made no decision before time expired. This happened in about 0.7% of observations in the G treatment and 0.4% of observations in the I treatment. These observations are included in the regression results reported in Table 4 and those in Table 5 below, but additional regressions (available from the corresponding author upon request) show that the results are unaffected by dropping these "imposed zeros".

kept in the sample, the marginal effect of the Group treatment on under–saving is positive and weakly significant ($p \approx 0.060$) as hypothesised; its point estimate suggests an average under–saving of about 6% of expected income (about 2300 lab–francs per round) compared to the Individual treatment. Not surprisingly, given the relationship between under–saving and earnings seen in Figure 5, the Group variable also has the hypothesised negative effect on average earnings, but just misses being significant at the 10% level ($p \approx 0.105$). Its point estimate suggests that a subject could expect to earn £1.32 less in the Group treatment than in the individual treatment, which is about 7.7% of the average earnings of £17.06. The marginal effects of the Group treatment on over–saving, saving and absolute deviation are all insignificant ($p > 0.20$). The treatment effects are stronger when the two outlier subjects are removed. The marginal effect of the Group variable on under–saving becomes a bit larger: 7% of expected per–round income (and more significant, $p \approx 0.015$). The effect on earnings becomes significant with outliers removed ($p \approx 0.045$), though smaller in magnitude: the effect size of £1.19 is about 6.8% of average earnings of £17.42.

Table 4 also shows that the effect of the income stream on these five measures is usually insignificant, with the exception of absolute deviation, which is lower under Stream 2, supporting our earlier conjecture that this "more representative" income stream may be easier for subjects to deal with. Also, the marginal effects of the round number are significant in each of the four pairs of regressions where it appears, and their signs suggest that over time, under–saving decreases and is only partially replaced by over–saving.

In Table 5, we look more closely at the relationship between observation and under–saving. A pair of Tobit regressions continues to use our under–saving measure as its dependent variable; for robustness, we also estimate a pair of probits with an indicator for positive under–saving on the left–hand side. Right–hand–side variables are as before, but with two new variables added (as well as interactions with the original variables and each other). First, we include an indicator for the subject's sex (female = 1).[20] Second, we add a new indicator variable, "Top–3 rank". In the Group treatment, the Top–3 rank variable has a value of one if the subject's average utility up to the previous round was in the top three in his/her group (and thus observable by the other group members in the current round). In the Individual treatment, subjects were not assigned to groups, so there was no perfectly analogous way to define this variable. We chose to set it to one if the subject's average utility was in the top $3m$ in the session, where m is the number of groups that would have been formed had it been a G session. This ensures that the overall proportion of ones for this variable in a session is the same in G and I sessions. In our experiment, G sessions were split into two groups if they had 14 or fewer subjects, or three if they had 15 or more subjects. So, in the one I session with 14 subjects, the "Top–3 rank" dummy is set to 1 if the subject was in the top 6, while in the three I sessions with 15 or more subjects, it is set to 1 if the subject was in the top 9.

To ease the exposition, we will use phrases like "high ranked" for those subjects in the top 3 in their group (or in the I treatment, the top $3m$ in their session), and "low ranked" for those subjects outside that group. As should be clear by now, this is not meant to imply that high–ranked subjects have made better or more rational choices; the opposite is often true, especially in early rounds. However, these high–ranked subjects are the exemplars for others, according to our conjectures about positional consumption.

The top section of results parallels those in Table 4: average marginal effects, standard errors and significance levels.[21] The first row shows that even after controlling for subjects' relative positions, the Group treatment indicator

[20] The relationship between demographic variables and saving was previously examined by Carbone (2005), who found no systematic associations in her life–cycle experiment.

[21] Additional regressions, not reported here but available from the corresponding author, show that including either additional demographic variables (age, economics classes taken, number of years living in the UK) or the *size* of the group – which affects the likelihood of being ranked in the top 3 and hence might affect under–saving – turns out not to change any of the qualitative results reported here, nor to have any significant effect.

Table 5: Regression results – marginal effects, with standard errors in parentheses

Dependent variable:	Under–saving (frac. of expected income)		Under–saving (indicator)	
Average marginal effects				
Group treatment	0.063**	0.073**	0.032	0.038
	(0.031)	(0.029)	(0.036)	(0.035)
Income stream 2	–0.040	–0.036	0.038	0.032
	(0.032)	(0.030)	(0.034)	(0.034)
Round	−0.006***	−0.006***	−0.007***	−0.007***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Top–3 rank	0.015	0.014	0.031**	0.030*
	(0.012)	(0.011)	(0.016)	(0.016)
Female	0.024	0.014	–0.007	–0.013
	(0.032)	(0.030)	(0.034)	(0.034)
Marginal effects at particular variable values				
Group treatment | not top–3 rank	0.120***	0.131***	0.092**	0.098**
	(0.032)	(0.030)	(0.042)	(0.041)
Group treatment | top–3 rank	0.011	0.021	–0.014	–0.009
	(0.034)	(0.032)	(0.037)	(0.036)
p–value of difference	$<$ 0.001	$<$ 0.001	0.001	0.001
Group treatment | male	0.086	0.090**	0.042	0.051
	(0.045)	(0.042)	(0.053)	(0.052)
Group treatment | female	0.059	0.077*	0.061	0.061
	(0.044)	(0.041)	(0.049)	(0.048)
p–value of difference	0.660	0.828	0.796	0.883
Group | male, not top–3	0.177***	0.181***	0.156**	0.167***
	(0.045)	(0.043)	(0.062)	(0.061)
Group | male, top–3	–0.007	–0.003	–0.077	–0.069
	(0.050)	(0.048)	(0.053)	(0.052)
Group | female, not top–3	0.086	0.105**	0.072	0.073
	(0.046)	(0.043)	(0.057)	(0.056)
Group | female, top–3	0.031	0.050	0.050	0.051
	(0.049)	(0.046)	(0.050)	(0.050)
p–value, males (top–3 vs. not)	$<$ 0.001	$<$ 0.001	$<$ 0.001	$<$ 0.001
p–value, females (top–3 vs. not)	0.103	0.085	0.636	0.629
p–value, male diff. vs. female diff.	0.006	0.005	0.001	0.001
Outliers removed?	No	Yes	No	Yes
Constant term?	Yes	Yes	Yes	Yes
Additional demo. variables?	Yes	Yes	Yes	Yes
Interaction effects?	Yes	Yes	Yes	Yes
N	9860	9744	9860	9744
$\|ln(L)\|$	7539.69	7242.56	5448.53	5388.57

* (**,***): *Coefficient significantly different from zero at the 10% (5%, 1%) level.*

continues to have a significant positive effect on the level of under–saving, though interestingly, its effect on whether or not subjects under–save is insignificant (though the point estimate has the expected positive sign). With the exception of the round number, the other variables have no significant effect on either measure of under–saving.

The remaining sections of the table show our primary interest, the marginal effect of the Group variable conditional on specific values of other variables. The second section highlights the interaction between the Group variable and our Top–3–rank variable. The marginal effect of the Group indicator is positive and significant for low–ranked subjects, and insignificant and much smaller for high–ranked subjects, in all four regressions. This suggests that the effect of the ranking information in the G treatment is driven primarily by the behaviour of low–ranked subjects, who under–save more when they are given information about others' consumption ("keeping up with the Joneses"), while high–ranked subjects behave largely the same whether they have this information or not. (Section 4 includes a discussion of other potential explanations for this result.) The difference–in–difference test between the effect of the G treatment on high–ranked versus low–ranked subjects is significant ($p < 0.01$) in all four regressions, confirming that high– and low–ranked subjects are affected differently by this extra information. By contrast, the third section of the table shows that the marginal effect of the Group variable is largely the same for both sexes: differences are statistically insignificant and economically negligible. That is, we do not find sex differences in how subjects *overall* are affected by the extra information in the G treatment.

The fourth section of the table shows the marginal effect of the G treatment conditional on the *combination* of sex and top–3 status. We continue to see that the effect of the Group treatment is limited to low–ranked subjects, but now we also can see that it is particularly pronounced in low–ranked males. The size of the effect on males outside the top–3 is roughly 18% of expected income, significantly larger than that for males in the top–3 ($p < 0.001$ in all four regressions). The corresponding difference for females is much smaller, and significant in only one of the four regressions. The difference–in–difference–in–difference between males and females is significant (at the 1% level or better) in all four regressions. Thus, while the second section of the table shows that low–ranked subjects respond to the added information in the Group treatment by trying to "keep up with the Joneses", this last section shows that males are significantly more prone than females to succumb to this temptation.

4 Discussion

We examine the effect of positional goods – goods for which individuals value relative as well as absolute holdings – on saving, using a lab experiment based on a 60–round life–cycle decision problem (Modigliani and Brumberg (1954)). In each round, subjects receive a randomly drawn income, then choose consumption and saving. This is a complex problem to solve, and previous experimental research (see the introduction) has found ample evidence of sub–optimal, and in particular myopic, behaviour in settings like this. In our baseline *Individual* treatment, subjects perform the task in isolation, with no feedback other than their own choices and the resulting outcomes. Our *Group* treatment makes inter–personal comparisons possible by additionally showing subjects the lifetime utilities of the three group members with the highest lifetime utilities (up to the current round), along with their own rank within the group. This extra information has no direct effect on monetary payments, but subjects may value the status associated with a high rank. Rank is somewhat manipulable: at the end of the experimental session, lifetime utility depends only on how well the subject has "solved" the consumption/savings problem, but during the session, total utility up to the current round can be raised by increasing consumption. Thus, myopic subjects who value status may succumb to the "demonstration effect" described by Duesenberry (1949) and others: increasing consumption will improve one's relative position now, but at the cost of savings, meaning that future consumption and hence future relative position will be harmed as a result.

Our experimental results are consistent with this pattern of behaviour, according to both non–parametric tests and parametric regressions. Overall, under–saving (saving less than what would be optimal given the round and the subject's available funds) is significantly *higher* in the Group treatment than in the Individual treatment. This under–saving is myopic: actual cash earnings (which are proportional to average induced utility over the entire 60 rounds) are weakly significantly *lower* in the Group treatment. Both effects are economically relevant, with magnitudes in the neighbourhood of 6–8% (of expected income and average earnings, respectively). Additional regressions show that our results are driven by the effect of the extra information on lower–ranked subjects (i.e., those outside the top three in their group) under–saving in an attempt to “keep up with the Joneses”, rather than, for example, a uniform increase in under–saving over all subjects, or attempts by the higher–ranked subjects to prevent others catching up with them. This keeping–up–with–the–Joneses effect is substantially stronger in males than in females, but it is sometimes significant for females also.

As with any experiment designed to test an implication of a theory, it is important to be alert to alternative explanations. One conceivable explanation for our results is “information overload” (Earl (1990)): the additional information given to subjects in the G treatment confuses or distracts them, so that the observed increase in under–saving in that treatment is simply an increase in decision errors. However, Tables 3 and 4 show that it is only errors in one direction (under–saving) that increase in the G treatment. There is no corresponding increase in errors in the other direction (over–saving); indeed, over–saving is *lower* in the G treatment than the I treatment, though the difference is not significant. The directional nature of this treatment effect makes it unlikely that it is due primarily to information overload or related factors.

Another plausible explanation is imitation: a subject in the G treatment sees someone else performing “better” (in the sense of having a higher rank), and adopts that subject’s strategy. If the highest–ranked subjects are the ones who have consumed the most up to the current round, imitation would be indistinguishable from the demonstration effect. To address the possibility that imitation drives our results, we start by examining the extent to which high rank is linked to high consumption. Table 6 presents the results of regressions, estimated using the G–treatment data from a single round (round 1, 10, 20, 30, 40 or 50), with average utility up to that round as the dependent variable and with end–of–round savings as the main explanatory variable (for a given income stream, total income from the first round to the current round is the same for all subjects, so end–of–round savings are a perfect proxy for total consumption up to that round). Additional right–hand–side variables are a dummy for Stream 2 and its product with end–of–round savings. As one would expect, at the end of the first round the highly–ranked subjects

Table 6: Regression results – average utility, G treatment, one outlier removed (marginal effects and standard errors)

Round:	1	10	20	30	40	50
End–of–round saving	-27.326^{***}	-0.678^{*}	0.085	0.257	0.141	0.521^{***}
	(1.641)	(0.393)	(0.150)	(0.177)	(0.159)	(0.167)
Income stream 2	-1.983^{***}	-5.867^{***}	-5.555^{***}	-5.790^{***}	-3.697^{***}	-2.575^{***}
	(2.107)	(1.127)	(0.393)	(0.781)	(0.857)	(0.834)
R–squared	0.75	0.30	0.41	0.42	0.19	0.12

* (**,***): *Coefficient significantly different from zero at the 10% (5%, 1%) level.*

are overwhelmingly the ones who have saved the least, or equivalently consumed the most, and even at the end of the tenth round this qualitative relationship continues to hold. Over time, however, subjects’ ranks increasingly reflect their ability to choose optimal consumption rather than high consumption. By the twentieth round, subjects with higher lifetime utilities actually tend to be those with lower consumption, and this continues in later rounds. Thus, imitation implies that a subject will *increase* consumption in early rounds and *decrease* it in later rounds.

In Table 7, we see that subjects do not behave this way. For this table, we replicate the models estimated for Table 5, but for only the second halves of subjects’ “lives”, when they would have had at least ten rounds of exposure to the environment where high rank reflects lower, not higher, total consumption. These results are qualitatively

Table 7: Regression results – identical to Table 5 but for rounds 31–59 only

Dependent variable:	Under–saving (as fraction of expected income)		Under–saving (indicator)	
Marginal effects at particular variable values				
Group treatment | not top–3 rank	0.069^{**}	0.074^{**}	0.031	0.036
	(0.033)	(0.030)	(0.055)	(0.055)
Group treatment | top–3 rank	0.008	0.012	–0.018	–0.017
	(0.030)	(0.027)	(0.050)	(0.050)
p–value of difference	0.073	0.044	0.401	0.372
Group | male, not top–3	0.153^{***}	0.147^{***}	0.102	0.113
	(0.050)	(0.046)	(0.085)	(0.086)
Group | male, top–3	–0.001	–0.003	–0.071	–0.069
	(0.043)	(0.038)	(0.069)	(0.069)
Group | female, not top–3	0.015	0.033	0.021	0.018
	(0.043)	(0.038)	(0.071)	(0.071)
Group | female, top–3	0.019	0.029	0.026	0.026
	(0.045)	(0.041)	(0.072)	(0.072)
p–value, males (top–3 vs. not)	0.002	0.001	0.049	0.039
p–value, females (top–3 vs. not)	0.929	0.929	0.952	0.921
p–value, male diff. vs. female diff.	0.021	0.019	0.135	0.110
Outliers removed?	No	Yes	No	Yes
Constant term?	Yes	Yes	Yes	Yes
Additional demo. variables?	Yes	Yes	Yes	Yes
Interaction effects?	Yes	Yes	Yes	Yes
N	4930	4872	4930	4872
$\|ln(L)\|$	3225.75	2936.16	2878.83	2846.10

* (**,***): *Coefficient significantly different from zero at the 10% (5%, 1%) level.*

similar to those using all rounds' data: we continue to see the "keeping up with the Joneses" effect on low–ranked subjects, particularly males. We do *not* see any evidence that subjects in the G treatment are over–saving, or even reducing their under–saving, even though this is how they would imitate higher–ranked subjects. We thus conclude that imitation, like information overload, is not driving our results. However, we acknowledge the possibility that still other explanations for our results may exist.

We would like to finish with a few comments about our results and what can be learned from them. First, we note that the behavioural effects we have observed were the result of a rather conservative implementation of positional concerns. The difference between our Group and Individual treatments was nothing more than a table on the computer screen showing other subjects' average utilities, and a few lines at the bottom of the instructions in fairly non–leading language (see Appendix A). Subjects with high relative status received no monetary payment for their rank, nor did they receive any tangible prize (as they did in Ball and Eckel (1998), Ball et al. (2001), and Kumru and Vesterlund (2010) for example). This suggests that it is fairly easy to induce positional concerns.

Second, we note a connection between the demonstration effect we observe and the "winner's curse" (over–bidding) observed in common–value auctions (Kagel and Levin (2002)). Several explanations have been suggested, including cognitive limits, risk aversion, and insensitivity to small payoff differences, but one intriguing possibility (Goeree et al. (2002)) is that subjects receive a "utility of winning". A subject with such a preference has an incentive to bid higher than the expected–payoff–maximising level, increasing the chance of winning but at the expense of profit. This closely resembles the sacrifice of money earnings to improve (temporarily) one's rank in our Group treatment, and suggests that displaying additional information about the high bidder in an auction may exacerbate the winner's curse.

References

Alpizar, F., Carlsson, F., Johansson-Stenman, O., 2005. How much do we care about absolute versus relative income and consumption? Journal of Economic Behavior & Organization 56, 405–421.

Anderhub, V., Güth, W., Müller, W., Strobel, M., 2000. An experimental analysis of intertemporal allocation behavior. Experimental Economics 3, 137–152.

Andersson, F.W., 2008. Is concern for relative consumption a function of relative consumption? Journal of Socio–Economics 37, 353–364.

Ando, A., Modigliani, F., 1963. The 'life cycle' hypothesis of saving: aggregate implications and tests. American Economic Review 53, 55–84.

Angelucci, M., DeGiorgi, G., 2009. Indirect effects of an aid program: how do cash transfers affect ineligibles' consumption? American Economic Review 99, 486–508.

Ball, S., Eckel, C.C., 1998. Stars upon thars: status and discrimination in ultimatum games. Working paper, Virginia Tech.

Ball, S., Eckel, C.C., Grossman, P.J., Zame, W., 2001. Status in markets. Quarterly Journal of Economics 116, 161–181.

Ballinger, T.P., Palumbo, M.G., Wilcox, N.T., 2003. Precautionary saving and social learning across generations: an experiment. Economic Journal 113, 920–947.

Brown, A.L., Chua, Z.E., Camerer, C.F., 2009. Learning and visceral temptation in dynamic saving experiments. Quarterly Journal of Economics 124, 197–231.

Carbone, E., 2005. Demographics and behaviour. Experimental Economics 8, 217–232.

Carbone, E., Duffy, J., 2013. Lifecycle consumption plans, social learning and external habits: experimental evidence. Working paper, University of Pittsburgh.

Carbone, E., Hey, J.D., 2004. The effect of unemployment on consumption: an experimental analysis, Economic Journal 114, 660–683.

Card, D., Mas, A., Moretti, E., Saez, E., 2012. Inequality at work: the effect of peer salaries on job satisfaction. American Economic Review 102, 2981–3003.

Carlsson, F., Johansson-Stenman, O., Martinsson, P., 2007. Do you enjoy having more than others? Survey evidence of positional goods, Economica 74, 586–598.

Carroll, C.D., 1997. Buffer–stock saving and the life cycle/permanent income hypothesis. Quarterly Journal of Economics 112, 1–55.

Charness, G., Masclet, D., Villeval, M.C., 2014. The dark side of competition for status. Management Science 60, 38–55.

Damianov, D.S., 2011. A classroom experiment on status goods and consumer choice. Journal of Economics and Finance Education 10, 1–13.

Demyanyk, Y., Van Hemert, O., 2011. Understanding the subprime mortgage crisis. Review of Financial Studies 24, 1848–1880.

Duesenberry, J.S., 1949. Income, Saving and the Theory of Consumer Behavior. Cambridge, MA: Harvard University Press.

Earl, P.E., 1990. Economics and psychology: a survey. Economic Journal 100, 718–755.

Elmer, P.J., Seelig, S.A., 1998. The rising long–term trend of single–family mortgage foreclosure rates. FDIC working paper 98–2.

Fatás, E., Lacomba, J.A., Lagos, F. Moro-Egido, A.I., 2007. Experimental tests on consumption, savings and pensions. Working paper.

Ferrer-i Carbonell, A., 2005. Income and well–being: an empirical analysis of the comparison income effect. Journal of Public Economics 89, 997–1019.

Fischbacher, U., 2007. z–Tree: Zurich toolbox for ready–made economic experiments. Experimental Economics 10, 171–178.

Frank, R.H., 1985a. Choosing the Right Pond: Human Behavior and the Quest for Status. New York: Oxford University Press.

Frank, R.H., 1985b. The demand for unobservable and other nonpositional goods. American Economic Review 75, 101–116.

Garon, S., 2011. Beyond Our Means: Why America Spends While the World Saves. Princeton University Press, Princeton, NJ.

Goeree, J.K., Holt, C.A., Palfrey, T.R., 2002. Quantal response equilibrium and overbidding in private–value auctions. Journal of Economic Theory 104, 247–272.

Greiner, B., 2004. An online recruitment system for economic experiments. In: Kremer, K., Macho, V. (Eds.). Forschung und wissenschaftliches Rechnen, GWDG Bericht 63. Göttingen: Gesellschaft für Wissenschaftliche Datenverarbeitung, 79–93.

Grossman, P., Komai, M., 2013. Within and across class envy: anti–social behaviour in hierarchical groups. Working paper, Monash University.

Harbaugh, H.R., 1996. Falling behind the Joneses: relative consumption and the growth–savings paradox. Economics Letters 53, 297–304.

Hey, J.D., Dardanoni, V., 1988. Optimal consumption under uncertainty: an experimental investigation. Economic Journal 98, 105–116.

Hirsch, F., 1976. Social Limits to Growth. Cambridge, MA: Harvard University Press.

Huberman, B.A., Loch, C.H., Önçüler, A., 2004. Status as a valued resource. Social Psychology Quarterly 67, 103–114.

Johnson, S., Kotlikoff, L.J., Samuelson, W., 1987. Can people compute? An experimental test of the life cycle consumption model. NBER working paper #2183.

Kagel, J.H., Levin, D., 2002. Common Value Auctions and the Winner's Curse. Princeton, NJ: Princeton University Press.

Kitamura, T., Yonezawa, Y., Nakasato, M. 2012. Saving behavior under the influence of income risk: an experimental study. Economics Bulletin 30, 1–8.

Knell, M., 1999. Social comparisons, inequality, and growth. Journal of Institutional and Theoretical Economics 155, 664–695.

Kosicki, G., 1987. A test of the relative income hypothesis. Southern Economic Journal 54, 422–434.

Kosicki, G., 1988. A note about savings as a 'nonpositional' good. Eastern Economic Journal 14, 271–276.

Kuhn, P., Kooreman, P., Soetevent, A. Kapteyn, A., 2011. The effects of lottery prizes on winners and their neighbors: evidence from the Dutch postcode lottery. American Economic Review 101, 2226–2247.

Kumru, C.S., Vesterlund, L., 2010. The effect of status on charitable giving. Journal of Public Economic Theory 12, 709–735.

Modigliani, F., Brumberg, R., 1954. Utility analysis and the consumption function: an interpretation of the cross–sectional data. In: Kurihara, K. (Ed.). Post–Keynesian Economics. New Brunswick, NJ: Rutgers University Press, 388–436.

Müller, W., 2001. Strategies, heuristics, and the relevance of risk–aversion in a dynamic decision problem. Journal of Economic Psychology 22, 493–522.

Organisation for Economic Cooperation and Development, 2011. OECD Factbook 2011–2012: Economic, Environmental and Social Statistics.

Ponzano, F., Ricciuti, R., 2012. An experimental AK model of growth. CESifo working paper #3744.

Samuelson, P.A., 1975. Optimum social security in a life–cycle growth model. International Economic Review 16, 539–544.

Siegel, S., Castellan, N.J., Jr., 1988. Nonparametric Statistics for the Behavioral Sciences. New York: McGraw–Hill.

Skinner, J., 1988. Risky income, life cycle consumption, and precautionary savings. Journal of Monetary Economics 22, 237–255.

Solnick, S.J., Hemenway, D., 1998. Is more always better? A survey on positional concerns. Journal of Economic Behavior & Organization 37, 373–383.

Summers, L.H., 1981. Capital taxation and accumulation in a life cycle growth model. American Economic Review 71, 533–544.

A Instructions from experiment [extra wording in G treatment]

Instructions

You are about to participate in a decision making experiment. Please read these instructions carefully, as the money you earn may depend on how well you understand them. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with the other participants – or anyone else – during the experiment.

The main part of the experiment consists of a sequence of decision problems. It will last for a total of 60 periods, which will be called “years”. Decisions will involve a fictional currency called “francs”. Francs can be converted to real money, as described below.

The decision problems involve consumption and saving over a lifetime of 60 years. You begin the first year with 20000 francs of savings. In each year, there are two equally likely possibilities: all of you are employed or all of you are unemployed. If you are employed, your income for the year is 80000 francs; if you are unemployed, your income is zero. After learning your income, you decide how much of your funds to spend on consumption, with the remainder being saved. The amount you spend on consumption is converted into real money – the amount of real money earned from a particular amount of consumption is shown overleaf, as a table and as a graph. The amount you save is not converted directly into money, but it is available for consumption in future years.

Sequence of events: The sequence of events in a year is as follows.
(1) Your computer screen displays your savings from the previous year.
(2) Your computer screen displays your income for the year – either 0 or 80000. You choose the amount you wish to consume. The amount you choose can be any multiple of 0.01 francs, between zero and your available funds (your savings plus your income).
(3) The year ends. The computer screen displays: your consumption and savings for the year, the money you earned for the year, and your total and average money earnings up to the current year.
After this, you go on to the next year.

Payments: At the end of the experimental session, you will be paid the average of your money earnings in all years. Payments are made privately and in cash.

Additional information:
(1) In the final year, there is no benefit to saving (since there are no future years after this). So, in the final year you should consume all of your francs.
(2) When it’s time to choose your consumption for a year, the bottom-right of your computer screen will display a francs-to-pounds conversion calculator. You can enter any allowable consumption amount into the calculator, and by clicking the “compute” button, the calculator will show the amount of money you would earn for the year from that consumption amount. The calculator is for your information only – the amounts you enter into it have no effect on your earnings, and they do not restrict the amount you eventually choose for your consumption.
(3) The top-right corner of your computer screen will display a timer that shows the time left to make a decision. If the timer reaches zero before you choose a consumption amount, a zero amount is entered for your consumption for the year.
(4) The top part of your computer screen will display the history of your results from all years you have completed, including your average money earnings up to the current year.
[In G treatment only: (5) The participants in this session have been divided into groups, with each group made up of 5-9 participants. Starting from the end of the first year, the bottom-left of your computer screen will display the ID numbers of the three participants in your group with the highest average money earnings, along with their earnings. This ranking will include you, if you have one of the three highest earnings in your group.]

Relationship between consumption amount and money earnings in a round:

Note that the consumption values listed here are only examples. You are not restricted to choosing these amounts; your consumption amount in a year can be any multiple of 0.01 francs, between 0 and the funds you have available in that year.

Consumption (francs)	Money earned (£)
0	−32.67
1000	−28.43
2000	−24.59
3000	−21.10
4000	−17.90
5000	−14.98
6000	−12.30
7000	−9.83
8000	−7.55
9000	−5.45
10000	−3.50
15000	+4.37
20000	+10.00
25000	+14.17
30000	+17.33
40000	+21.76
50000	+24.63
100000	+30.45
150000	+32.32
200000	+32.87
250000	+33.23
300000	+33.45

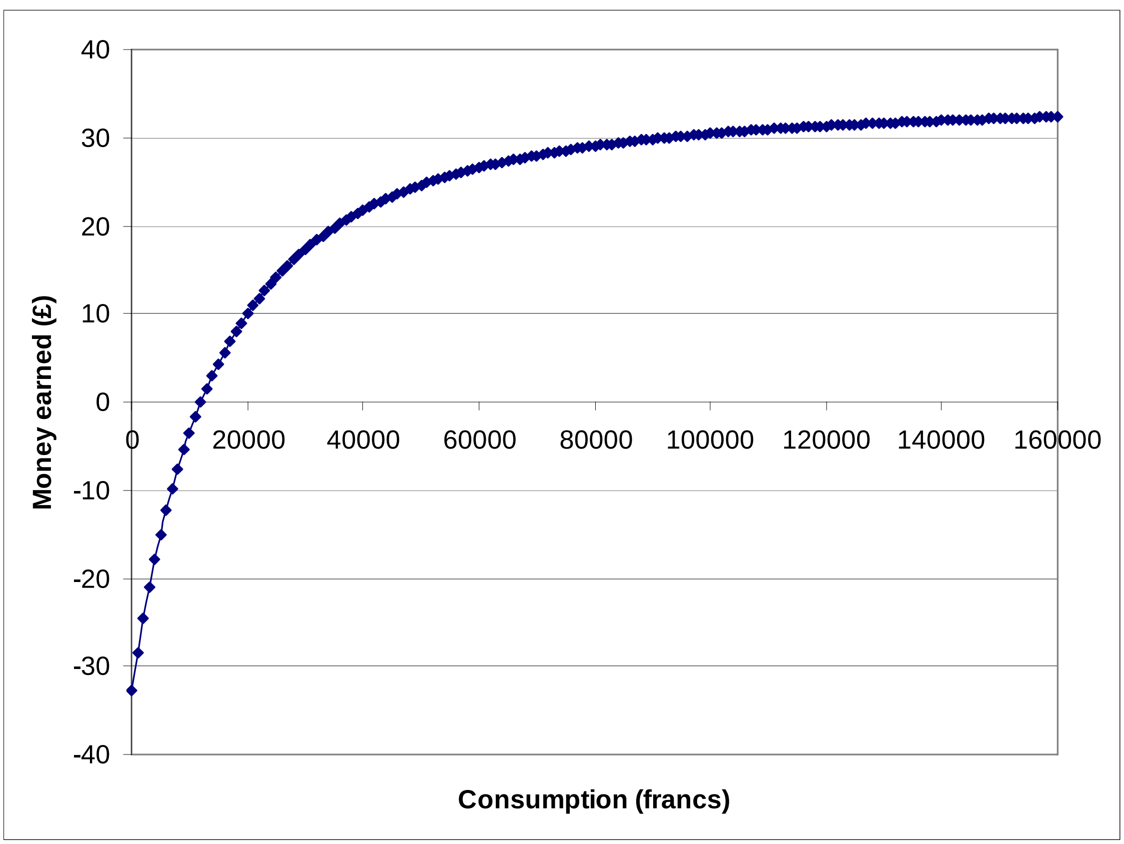

B Screenshots from experiment

Decision screen, I treatment:

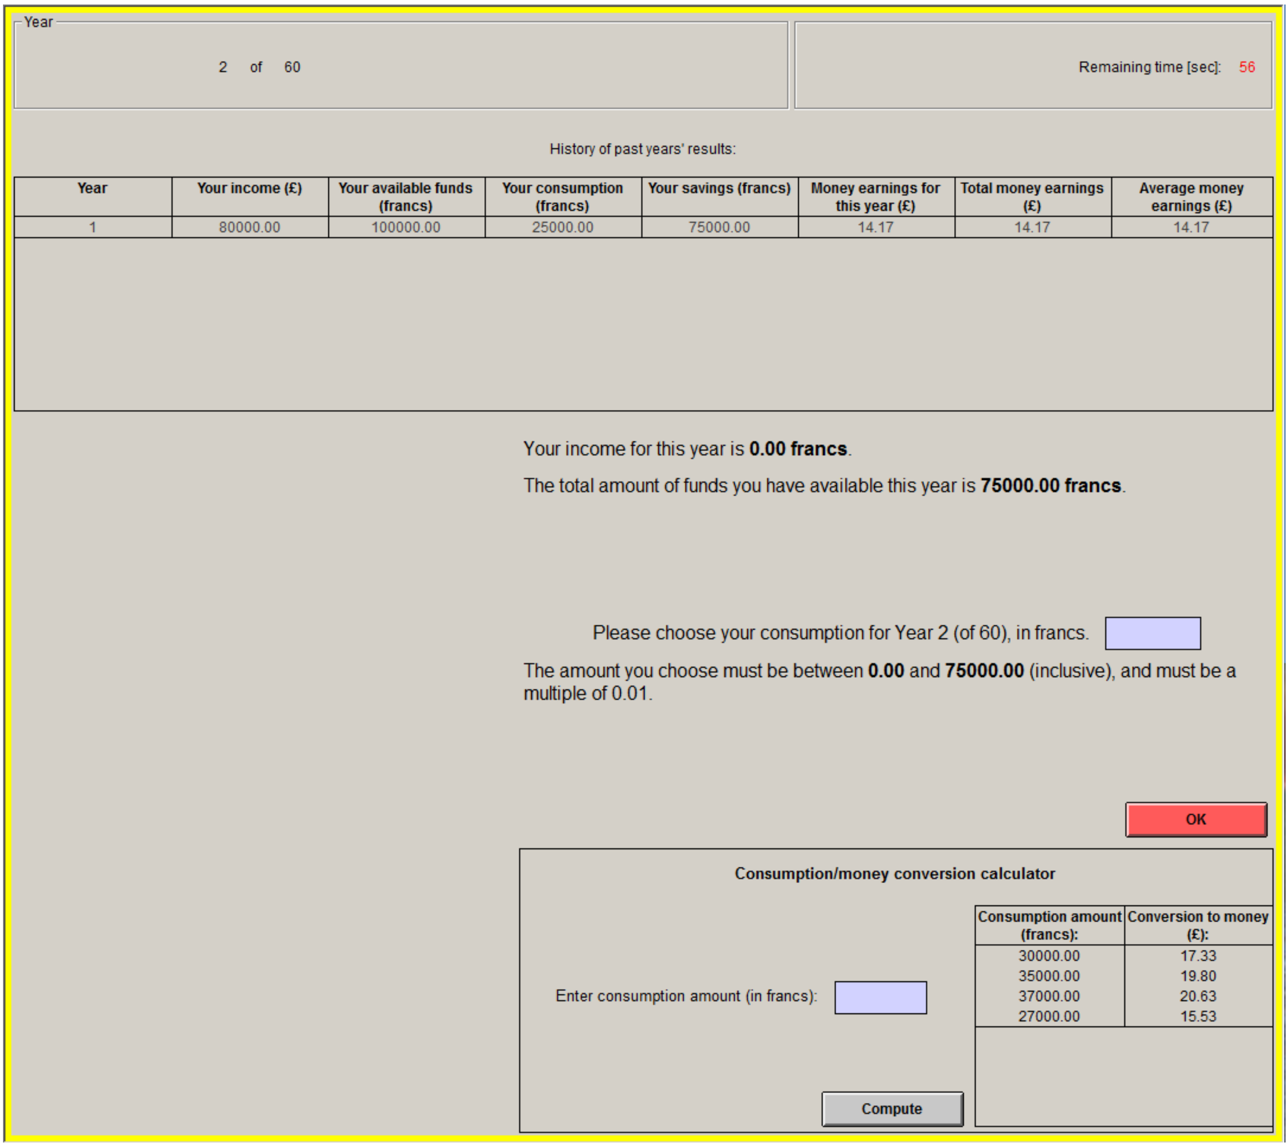

Feedback screen, I treatment:

Round

2 of 60

Remaining time [sec]: 14

History of past years' results:

Year	Your income (francs)	Your available funds (francs)	Your consumption (francs)	Your savings (francs)	Money earnings for this year (£)	Total money earnings (£)	Average money earnings (£)
1	80000.00	100000.00	25000.00	75000.00	14.17	14.17	14.17
2	0.00	75000.00	30000.00	45000.00	17.33	31.50	15.75

THIS YEAR'S RESULTS:

You had **75000.00 francs** available this year.

You chose to spend **30000.00 francs** on consumption, which converts into **£17.33**.

You have saved **45000.00 francs** for next year.

Your total earnings, including this year, are **£31.50**.

Your average earnings, including this year, are **£15.75**.

OK

Decision screen, G treatment:

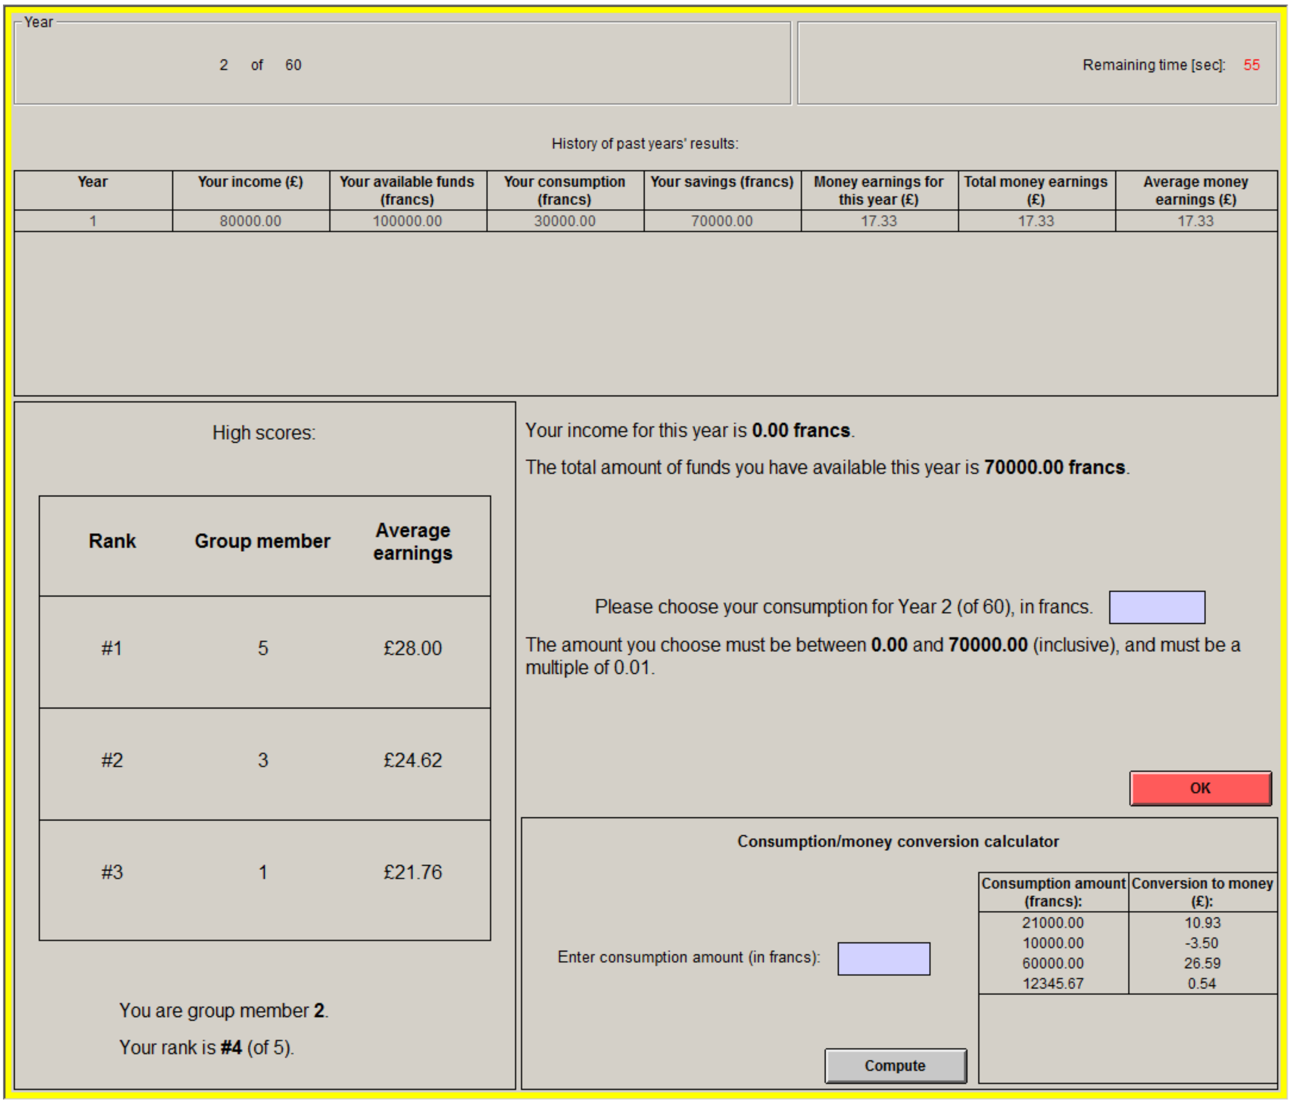
Year

2 of 60

Remaining time [sec]: 55

History of past years' results:

Year	Your income (£)	Your available funds (francs)	Your consumption (francs)	Your savings (francs)	Money earnings for this year (£)	Total money earnings (£)	Average money earnings (£)
1	80000.00	100000.00	30000.00	70000.00	17.33	17.33	17.33

High scores:

Rank	Group member	Average earnings
#1	5	£28.00
#2	3	£24.62
#3	1	£21.76

You are group member **2**.

Your rank is **#4** (of 5).

Your income for this year is **0.00 francs**.

The total amount of funds you have available this year is **70000.00 francs**.

Please choose your consumption for Year 2 (of 60), in francs.

The amount you choose must be between **0.00** and **70000.00** (inclusive), and must be a multiple of 0.01.

OK

Consumption/money conversion calculator

Enter consumption amount (in francs):

Compute

Consumption amount (francs):	Conversion to money (£):
21000.00	10.93
10000.00	-3.50
60000.00	26.59
12345.67	0.54

Feedback screen, G treatment:

Round

2 of 60

Remaining time [sec]: 13

History of past years' results:

Year	Your income (francs)	Your available funds (francs)	Your consumption (francs)	Your savings (francs)	Money earnings for this year (£)	Total money earnings (£)	Average money earnings (£)
1	80000.00	100000.00	30000.00	70000.00	17.33	17.33	17.33
2	0.00	70000.00	30000.00	40000.00	17.33	34.66	17.33

THIS YEAR'S RESULTS:

High scores:

Rank	Group member	Average earnings
#1	3	£23.19
#2	5	£20.71
#3	1	£17.97

You are group member **2**.

Your rank is **#4** (of 5).

You had **70000.00 francs** available this year.

You chose to spend **30000.00 francs** on consumption, which converts into **£17.33**.

You have saved **40000.00 francs** for next year.

Your total earnings, including this year, are **£34.66**.

Your average earnings, including this year, are **£17.33**.

OK