DO CONSUMERS GAMBLE TO CONVEXIFY

Thomas F. Crossley
Hamish Low
Sarah Smith

Do Consumers Gamble to Convexify?

Thomas F. Crossley

Koç University, University of Cambridge and Institute for Fiscal Studies (IFS)

Hamish Low

University of Cambridge and IFS

Sarah Smith

University of Bristol and IFS

May 2013[*]

Abstract

The combination of credit constraints and indivisible consumption goods may induce some risk-averse individuals to gamble to have a chance of crossing a purchasing threshold. One implication of this is that income effects for individuals who choose to gamble are likely to be larger than for the general population. Using UK data on gambling wins, other windfalls and durable goods purchases, we show that winners display higher income effects than non-winners but only amongst those likely to be credit-onstrained. This is consistent with credit-constrained, risk-averse agents gambling to convexify their budget set.

JEL classifications: D12, E21, D81, L83

Keywords: Gambling, Lotteries, Consumption, Durables

"On Friday September 4th 1994, the freezer belonging to Gloria and Steve Kanoy of Weere's Cove suddenly and mysteriously broke down. Distraught, the couple set off the next day in search of a new one. Stopping for gas at Lake Raceway, 607 Main Avenue, they decided to buy a Lotto ticket..."

Virginia Lottery winner awareness campaign, quoted in Clotfelter and Cook (1990)

1. Introduction

Why do risk-averse individuals gamble? One explanation, first developed by Ng (1965), is that discreteness in spending (such as the freezer in the quote above) or in labour supply opportunities can induce local non-concavities in the value functions of risk-averse agents. This generates local risk-loving behaviour and makes it rational to gamble in order to have a chance of crossing the threshold required to finance a lumpy purchase. Bailey et al (1980) argued that access to credit markets made such gambling irrational, but Hartley and Farrell (2002) showed theoretically that rational gambling might still occur where borrowing and lending rates differ, where capital market imperfections exist, or if individuals' time preference rates differ from interest rates. The aim of this paper is to test using data from the UK whether this theory is true: do some consumers gamble to convexify?[1]

This convexification hypothesis may help to explain some of the demand for gambling and lotteries, particularly among the poor and likely credit-constrained: the desire for convexification provides a rationalization for that demand despite individuals being risk averse. Gambling which may be in part motivated by a desire for convexification includes not just state run lotteries and sports betting but also prize-linked savings products.[2] In many developing countries, there is also the interesting case of rotating savings and credit associations (ROSCAs), discussed by Besley et al (1993). These are a micro-finance initiative in which groups of individuals make regular contributions to a fund, the total amount of which is allocated to one member each cycle

---

[1] One natural question is why do households not use the stock market to increase their risk and to gamble because the expected return on the stock market is clearly high. The point is that the stock market does not offer a discrete jump in payoffs. We do not think it is worth pursuing the question of why low income households do not use the derivative markets to take risks.

[2] In the UK, these are premium bonds, a government bond which pays a return in the form of an entry to a prize draw.

via a lottery. Handa and Kirton (1999) provide evidence from Jamaica that people use their allocation from the ROSCA to buy durable goods.

On the other hand, gambling to convexify cannot be the only motive for gambling because people who are obviously not credit-constrained still gamble, and these individuals must have some other motive.[3] Nevertheless, low income households spend a much larger share of their income on gambling than high income households, as we document below.

To highlight the mechanisms at work, we first develop a simple model where consumers choose whether or not to play a lottery, and then after the outcome of the lottery is known, whether or not to buy an indivisible good. The only consumers who play the lottery are those who are close to the threshold of being able to buy the indivisible good. A lottery win then enables the purchase of the indivisible good.

As noted above, the strength of this incentive for gambling will be diminished if agents can borrow at reasonable rates, so that the path of non-durable consumption can be unaffected by the timing of indivisible purchases. The need to gamble to convexify is also diminished if there are many indivisible goods so that the indivisibility is less "lumpy", or if there are uninsurable income shocks which provide some convexification. All this means that the importance of the convexification hypothesis is an empirical question.

To look for evidence that consumers gamble to convexify we use data from the British Household Panel Survey. Our empirical strategy is effectively a "difference-in-difference" design (with individual fixed effects), contrasting estimated within-person income effects for lottery windfalls with those for other windfalls (specifically inheritances) among those that are credit-constrained with those that are not. We use the group who are not credit-constrained to control for more general differences in responses by windfall type – including the degree to which alternative windfalls are anticipated, or psychological feelings attached to different sources of windfall. We also use data on financial expectations to examine directly whether inheritances are more anticipated than lottery wins. There is no evidence in these data that this is the case.

---

[3] For example, Tufano (2008) and Kearney et al (2010) have emphasized the entertainment aspect of prize-linked savings products in explaining their potential attraction.

Our main result is that, among individuals who are credit-constrained, purchases of consumer durable goods are more responsive to a lottery win than to receipt of other windfall income: among the credit-constrained, the income effect of an inheritance is less than 15% of the income effect of a gambling win of the same size. By contrast, there is no difference in the income effects of different windfalls for the unconstrained.

This finding is exactly what we would expect if credit-constrained consumers gamble to convexify, as first suggested by Ng (1965). As a further test, we examine the effects of non-lottery windfalls on individuals who can be inferred to have played the lottery but not had large winnings. For the subset of these individuals who are credit-constrained, purchases of consumer durable goods are more responsive to non-lottery windfall income than purchases by non-players: those who play the lottery exhibit larger income effects than those that do not play. Thus, consistent with our hypothesis, it is not the source of the money (lottery versus other windfall) that matters, but rather that lotteryplayers are in different economic circumstances than non-players.

This result has a number of important implications. First, a number of authors, notably Imbens et al. (2001), have used lotteries to estimate income effects in labour supply and consumption demands. Income effects are central to policy evaluation, but standard estimation procedures suffer from a lack of plausibly exogenous variation in income. Imbens et al. exploit the fact that lottery winnings provide random variation in income among lottery players. Of course, as Imbens et al. recognise, those who play the lottery may differ from those who do not. The key point we make is that the "convexification" hypothesis provides a compelling economic rationale for expecting this to be the case. Moreover, they will differ not just in general characteristics, but in particular, in their propensity to change a discrete choice in response to an income shock. The desire to convexify generates a demand for lotteries among precisely those individuals who are close to the threshold of a discrete decision (for example, purchasing a durable or retiring) and who therefore will display large income effects. The resulting threat to the external validity of the estimates is analogous to randomization bias (Heckman and Smith, 1995) in randomized trials: those who *participate* in the random allocation of treatment are systematically different from those who do not. An empirical strategy of measuring income effects based on lottery winnings may thus overestimate

the average response to a more broadly distributed windfall. Interestingly, Imbens et al. report that lottery winnings appear to have larger effects on discrete margins, such as retirement.

Non-convexities due to the discreteness of choices pose a major technical challenge to researchers trying to model those choices structurally with dynamic programming models. One way to overcome this problem has been to assume that individuals facing such non-convexities play wealth lotteries (Hansen, 1985; Rogerson, 1988; Lentz and Traneas, 2004). A second implication of our results is that this is not simply a technical convenience, but it captures the way that individuals actually behave when faced with non-convexities.

An outline of the rest of the paper is as follows. In the next section we develop the theoretical framework that guides our analysis. In Section 3 we examine the implications of the model for the resulting income effects if lotteries are endogenously chosen. Section 4 describes our data and empirical framework. Section 5 presents our main results, and Section 6 concludes.

2. A Model of Gambling to Finance Indivisible Purchases

Our model is a one period model with two stages.[4] At the start of the period (in the first stage), agents have cash on hand x_1. They first make a decision about whether or not to buy at most one lottery ticket: $l \in \{0,1\}$, where the price of the lottery ticket is 1. They then discover whether or not they have won. The lottery ticket is actuarially fair:[5] an agent holding a ticket wins $1/q$ with probability q, so that net winnings are $(1-q)/q$ with probability q and -1 with probability $1-q$. Net winnings augment an agent's cash-on-hand. Thus, $x_2 = x_1$ if a ticket is not purchased, but if a ticket is purchased, disposable cash-on-hand will be $x_2 = x_1 + (1-q)/q$ with probability q and $x_2 = x_1 - 1$ with probability $1-q$.

---

[4] This means we can abstract from borrowing and saving. As discussed later, the ability to borrow and save is likely to reduce the need to gamble to convexify. We exploit this difference in our estimation procedure, but we abstract from this in our model to make the motive for gambling transparent.

[5] We could introduce a penalty for gambling and make the gamble actuarially unfair, but this would simply act to offset the motive to gamble caused by the non-convexity.

After lottery winnings are revealed, individuals decide, in the second stage, how to allocate their spending between a divisible consumption good and an indivisible consumption good. Agents can buy at most one unit of the indivisible good ($d \in 0,1$) at price p. In our empirical work, the indivisible goods will be consumer durables. Without borrowing or saving, consumption of the divisible good is just $x_2 - dp$. Individuals maximize utility, which depends on the consumption of divisible and indivisible goods: $v(x_2 - dp, d) = u(x_2 - dp) + \eta d$; η is a preference parameter. We assume that $u'(\cdot) > 0$, $u''(\cdot) < 0$ and $u(0) + \eta < u(p)$, where this last condition specifies that the individual will not buy the indivisible good if this implies 0 consumption of the divisible good. [6]

We solve this simple model by backward induction. Define $V_2^{d=1}(x_2) = u(x_2 - p) + \eta$ and $V_2^{d=0}(x_2) = u(x_2)$. The indivisible good is purchased if and only if $V_2^{d=1}(x_2) \geq V_2^{d=0}(x_2)$, ie. $u(x_2 - p) + \eta \geq u(x_2)$.

Result 1 (single-crossing): *There is a unique x_2^* such that the indivisible good is purchased if and only if $x_2 \geq x_2^*$. x_2^* is implicitly defined by $u(x_2^* - p) + \eta = u(x_2^*)$.*

Proof: Uniqueness follows from the fact that

$$\frac{\partial V_2^{d=0}(x_2)}{\partial x_2} = u'(x_2) < \frac{\partial V_2^{d=1}(x_2)}{\partial x_2} = u'(x_2 - p) \quad \forall x_2 \tag{1}$$

which in turns follows from the concavity of $u(\cdot)$.

This difference in the derivative of the conditional value functions implies that the unconditional value function is non-concave because the derivative changes discretely at the point where the two value functions cross. This is illustrated in Figure 1.

[Figure 1, Single Crossing, About Here]

---

[6] The additive separability assumed here is not necessary. It is however necessary to restrict the degree of substitutability between durable and non-durable consumption.

Turning to the first stage, in which the decision to gamble is taken, let $V_1^{l=1}(x_1)$ be the value of purchasing the lottery ticket and $V_1^{l=0}(x_1)$ the value of not gambling. A lottery ticket is purchased if and only if $E\left[V_1^{l=1}(x_1)\right] - V_1^{l=0}(x_1) \geq 0$. Note that:

$$V_1^{l=0}(x_1) = \max[u(x_1 - p) + \eta), u(x_1, 0)]$$
$$= \begin{cases} u(x_1 - p) + \eta & \text{if } x_1 \geq x_2^* \\ u(x_1) & \text{if } x_1 < x_2^* \end{cases} \quad (2)$$

and

$$E\left[V_1^{l=1}(x_1)\right] = q \max\left[u\left(x_1 - p + (1-q)/q\right) + \eta, \ u(x_1 + (1-q)/q)\right]$$
$$+ (1-q) \max\left[u(x_1 - p - 1) + \eta, \ u(x_1 - 1)\right]$$

Result 2: *Lottery tickets are not purchased outside the interval* $\left[x_2^* - \dfrac{1-q}{q}, x_2^* + 1\right]$.

Proof: See appendix.

The intuition behind this result is straightforward. If $x_1 > x_2^* + 1$ the agent purchases the indivisible good regardless of the outcome of the lottery. Thus only $V_2^{d=1}$ is relevant, and the concavity of $V_2^{d=1}$ (which is inherited from the concavity of $u(\cdot)$) ensures that the agent does not gamble. If $x_1 < x_2^* - (1-q)/q$ the agent does not purchase the indivisible good regardless of the outcome of the lottery. Thus only $V_2^{d=0}$ is relevant, and the concavity of $V_2^{d=0}$ (which is inherited from the concavity of $u(\cdot)$) ensures that the agent does not gamble. The bounds, $x_2^* - (1-q)/q$ and $x_2^* + 1$ are illustrated in Figure 2.

[Figure 2, Regions]

Corollary 1: A lottery winner always purchases the indivisible good.

Proof: Since lottery tickets are never bought if $x_1 < x_2^* - (1-q)/q$, a lottery winner (with net winnings $(1-q)/q$) always has $x_2 \geq x_2^*$.

Corollary 2: A lottery player that does not win does not purchase the indivisible good.

Proof: Since lottery tickets are never bought if $x_1 \geq x_2^* + 1$, any unsuccessful lottery player (with net winnings -1) always has $x_2 < x_2^*$.

Result 3: *There exists a compact region,* $x_1 \in \left[\underline{x_1}, \overline{x_1} \right]$, *which contains* x_2^* $\left(\underline{x_1} < x_2^* < \overline{x_1} \right)$, *in which the agent will purchase a lottery ticket.*

Proof: See appendix.

From *Result 2*, we know that $x_2^* - (1-q)/q \leq \underline{x_1} < \overline{x_1} \leq x_2^* + 1$. Within these bounds, the size of the region $x_1 \in \left[\underline{x_1}, \overline{x_1} \right]$ depends on parameter values (η, q and the curvature of of $u(\cdot)$).

Together, *Corollaries 1* and *2*, and *Result 3* imply that the state space (of cash on hand) can be divided into three regions. A region $x_1 \leq \underline{x_1}$ in which the agent does not buy a lottery ticket and does not buy the indivisible good; a region $\underline{x_1} < x_1 \leq \overline{x_1}$ in which the agent buys a lottery ticket and then buys a durable if and only if she wins the lottery; and a region $x_1 > \overline{x_1}$ in which the agent does not buy the lottery ticket but does buy the indivisible good. This is illustrated in Figure 2.

What this simple model illustrates is that lottery players are very likely to be close to the margin of a discrete decision. This implication that lottery players are gambling to convexify is tested in section 4 by looking at the income effects of gamblers and non-gamblers. In the next section we show the implications of our model for using lotteries to estimate income effects. Estimated income effects form the basis of subsequent empirical tests.

3. Implications for Estimating Income Effects

In this section, we show that the estimated income effect (i.e. the effect of income on the purchase of an indivisible good) associated with an endogenously chosen lottery will be a biased estimate of the population average income effect. To show this, we consider the follow thought experiment, which we refer to as a randomly assigned lottery: a random fraction (λ) of the population is compelled to buy the lottery ticket, and no other tickets are available. This thought experiment holds the number of tickets constant, but removes the element of choice from gambling. This leads to a measure of the income effect from lottery winnings when the lottery ticket purchase is random, and thus to a population average income effect.

We consider two cases, corresponding to two different data structures. In the first case, as in Imbens et al. (2001), income effects are estimated by comparing lottery winners and lottery losers (i.e. people who play the lottery, but lose). In the second case, the comparison is between winners and non-winners; the latter includes both losers and non-players. We show that in both cases the extra spending by winners is a biased estimate of the population average income effect (which would be the income effect arising from a truly exogenous windfall.)

3.1. Comparing Winners and Losers.

First, consider the comparison between lottery winners and lottery losers. In the model developed above, an agent always buys the indivisible good if they are a lottery winner (*Corollary 1*). Thus in this model, in which lottery playing is a choice, the probability that a lottery winner purchases the indivisible good is one: $Prob(d=1|winner,choice)=1$. We condition the probability on "choice" to indicate that playing the lottery was a decision taken by the individual. These probabilities are summarized in Table 1.

In the case of the randomly allocated lottery on the other hand, lottery winners who have net winnings of $(1-q)/q$ giving $x_2 = x_1 + (1-q)/q$, will purchase the indivisible good if $x_1 \geq x_2^* - (1-q)/q$. Thus

$$Prob(d=1|winner,rand)=1-F\left(x_2^*-(1-q)/q\right).$$

This implies winners from the chosen lottery are more likely to purchase the indivisible good than winners of the randomly allocated lottery:

$$\Pr ob(d =1\,|\, winner, choice) - \Pr ob(d =1\,|\, winner, rand) = F\left(x_2^* - \frac{1-q}{q} \right) \geq 0$$

(3)

The intuition behind this result is straightforward. In the case of the randomly allocated lottery, some winners will come from below the lower threshold and will not have enough cash on hand to buy the divisible good even if they win. This difference tends to zero as q becomes increasingly small: if there is a lottery prize that is very large but with a very small probability of winning, it is in the interest of everyone with cash-on-hand below x_2^* to gamble to convexify, and all winners will chose to buy the indivisible good, whether the lottery ticket was chosen or randomly allocated.

For those that chose to play, but lost, the probability that the non-winner purchases the indivisible good is zero: $Prob(d =1\,|\,lost, choice) = 0$. By comparison, among the losers in the randomly allocated lottery are some people with cash on hand above the upper threshold $\left(x_2^* +1\right)$ who will have enough cash on hand to purchase the divisible good even if they lose, i.e. $\Pr ob(d =1\,|\,lost, rand) = 1 - F\left(x_2^* +1\right)$. The probability of purchase is therefore lower among losers of the endogenous lottery than among losers of the randomly allocated lottery:

$$\Pr ob(d =1\,|\, loser, choice) - \Pr ob(d =1\,|\, loser, rand) = F\left(x_2^* +1\right) - 1 \leq 0 \quad (4)$$

Putting together the differences in purchase probabilities between winners and the difference in purchase probabilities between losers, it is clear that the income effect in the case of the endogenous lottery suffers from upward bias compared to income effect from the randomly allocated lottery. The latter is an unbiased estimate of the population average income effect. An expression for the bias is given in the final row of the second [d] column of Table 1. The size of the bias becomes smaller as the range in which tickets are bought becomes larger.

3.2. Comparing Winners and Non-winners

Non-winners comprise both non-players and losers. In the case of the endogenous lottery, those who choose not to play are those with $x_1 < \underline{x_1}$ or $x_1 > \overline{x_1}$ while losers are the fraction $1 - q$ of lottery players, who all have $\underline{x_1} < x_1 \leq \overline{x_1}$. Of these non-winners, only agents with cash on hand $x_1 > \overline{x_1}$ buy the indivisible good. Let λ denote the faction of lottery players: $\lambda = F\left(\overline{x_1}\right) - F\left(\underline{x_1}\right)$, where $F(\cdot)$ is the cumulative distribution of cash on hand (x_1) in the population. Thus

$$Prob(d = 1 \mid non-winner, choice) = \frac{1 - F\left(\overline{x_1}\right)}{1 - Prob(winner)} = \frac{1 - F\left(\overline{x_1}\right)}{1 - q\lambda} \quad (5)$$

The effect of winning the lottery (relative to non-winners) on the probability of indivisible good purchase is therefore the difference:

$$Prob\left(d = 1 \middle| winner, choice\right) - Prob\left(d = 1 \middle| non-winner, choice\right) = 1 - \frac{1 - F\left(\overline{x_1}\right)}{1 - q\lambda}$$

$$= \frac{F\left(\overline{x_1}\right) - q\lambda}{1 - q\lambda} \quad (6)$$

In the randomly assigned lottery, non-winners comprise those that were randomly allocated a ticket but did not win, and those that were not allocated a ticket. The former are fraction $(1 - q)\lambda$ of the population, have net winnings of -1 and purchase the durable if $x_1 \geq x_2^* + 1$. The latter are fraction $(1 - \lambda)$ of the population, have net winnings of 0 and purchase the durable if $x_1 \geq x_2^*$. The overall fraction of the population that are non-winners is, as before, $(1 - q)\lambda + (1 - \lambda) = 1 - q\lambda$. Thus the fraction of non-winners who purchase the durable is:

$$Prob(d = 1 \mid non-winner, rand) = \frac{Prob(d = 1, non-winner, rand)}{Prob\left(non-winner, rand\right)}$$

$$(7)$$

$$= \frac{(1 - q)\lambda\left(1 - F\left(x_2^* + 1\right)\right) + (1 - \lambda)\left(1 - F\left(x_2^*\right)\right)}{1 - q\lambda}$$

This can be interpreted more easily if we approximate $F(x_2^* + 1)$ by $F(x_2^*)$: this is a good approximation if the price of the lottery ticket, 1, is small compared to typical cash-on-hand. The probability then becomes:

$$Prob(d = 1 \mid non-winner, rand) = \frac{\left(1 - F\left(x_2^*\right)\right)(1 - q\lambda)}{1 - q\lambda} \tag{8}$$

The denominator is the fraction of the population who are not winners. The first part of the numerator, $1 - F\left(x_2^*\right)$, is the fraction of all individuals whose cash-on-hand means they would purchase the durable regardless of the lottery. Some of these individuals will be winners when the lottery tickets are randomly allocated and so this fraction is multiplied by the fraction of the population that are not winners.

By contrast, when the purchase of the lottery ticket was a choice, none of those individuals who would purchase regardless of the lottery choose to buy lottery tickets and so they are all non-winners, and the probability of purchasing the durable among non-winners is given by equation (7).

To aid interpretation of the difference between equation (10) and (7), approximate $F(\bar{x}_1)$ by $F(x_2^*)$ (recall from Results 2 and 3 that $x_2^* < \bar{x}_1 \leq x_2^* + 1$). This gives a difference in the probability of purchase among non-winners from the chosen and random lotteries of:

$$Prob(d = 1 \mid non-winner, choice) - Prob(d = 1 \mid non-winner, rand) = \frac{\lambda q \left(1 - F\left(x_2^*\right)\right)}{1 - q\lambda} \geq 0$$

(9)

The probability of purchasing the durable among the non-winners from the chosen lottery is higher. This arises from a subtle composition effect because the group of non-winners comprises two sets of individuals: those who did not have a ticket and those that had a losing ticket. Some of those who were non-winners by choice (ie chose not to have a ticket because they would have purchased the durable anyway) became random lottery winners, and this reduces the number of purchasers of the durable among those who were not winners.

When we compare winners and non-winners of the lottery, the probability of purchasing the durable is higher among both winners and losers when the lottery ticket is chosen. This implies that the effect of winning the (chosen) lottery on durable purchases may be greater or smaller than the effect of winning on purchases when the allocation is random. The net effect is given by:

$$\frac{F\left(x_2^* - \frac{1-q}{q}\right) - \lambda q\left(1 - \left(F\left(x_2^*\right) - F\left(x_2^* - \frac{1-q}{q}\right)\right)\right)}{1-\lambda q} \tag{10}$$

When the probability of winning, q, is small and the prize large, the first term on the numerator is close to zero and the negative composition effect might dominate. On the other hand, as q gets larger, the second term tends to λq and the net effect is positive: the effect of the winners of the chosen lottery being more likely to purchase the durable dominates.

These two offsetting differences can be highlighted by calculating numerically the size of the effects in our simple model, at particular parameter values. We assume log utility for consumption, and consider a high and low value for the utility of the durable (η). Figure 3 shows the difference in the probability of purchase between the chosen and random lotteries. For these parameters, estimates of the effect of a windfall on the purchase of the durable from lotteries will overestimate the effect of a random windfall except for very small values of q.

[Figure 3, Simulation, About Here]

3.3 Discussion

The aim of this model was to highlight the differing income effects arising from different sorts of windfall gain. In particular, the effect of a windfall on indivisible purchases is likely to be larger if the windfall arises from a lottery that the household has chosen to participate in. This arises because the indivisibility means households have an incentive to gamble to convexify. However, the strength of this incentive will be diminished if capital markets are well functioning, and so agents can borrow or save, because this allows the path of non-durable consumption to be unaffected by the timing of indivisible purchases (Bailey et al., 1980; Hartley and Farrell, 2002). The need to

gamble to convexify is also diminished if there are multiple indivisible goods so that the indivisibility is less "lumpy", or if there are uninsurable income shocks which provide some convexification. This means that the importance of the convexification hypothesis is an empirical question, and in the remainder of this paper we assemble empirical evidence on this question. While data on lottery players cleanly identifies the income effect among players, this data structure is incapable of shedding light on the convexification hypothesis that the income effects among lottery players are different from those of nonplayers. Thus we turn to a general population survey. We can however also use our general survey to approximate the data set focused only on lottery players used by Imbens et al (2001) (see the discussion of the small winnings test in Sections 4 and 5).

4. Empirical Framework and Data

4.1. Empirical framework

We adopt a reduced form empirical approach directly motivated by the discussion in the previous section where we showed that endogenous gambling results in income effects that are biased compared to those based on an exogenous windfall. Our main empirical strategy is to examine a difference-in-difference in income effects. In particular, we compare the effect of windfalls on purchases of indivisible goods between lottery winners[7] and those who receive another type of windfall – namely an inheritance, and we compare these differences between those who are likely to be credit-constrained and those who are not.

In the context of the model above, inheritances are intended to approximate the randomly allocated lottery. The assumption is *not* that inheritances are random across the population, but that they are exogenous with respect to the distance between cash on hand (x_1) and the critical value (x_2^*), conditional on controls and individual fixed effects. Note that the critical value will vary in the population and over time for a given individual

---

[7] The BHPS question actually asks about all gambling wins. In practice, 79% of all spending on gambling is on the UK National Lottery, according to the Expenditure and Food Survey. This is a general household survey that is unlikely to capture serious gamblers, but it is similar to the BHPS sample. "Lottery wins" is therefore a shorthand for all gambling wins.

according to tastes and needs. It is the latter variation that we exploit in our panel data analysis.

Our empirical work focuses on durable goods, which are inherently indivisible; our prediction is that durable purchases will respond differently to a lottery win than to an inheritance because the people who receive lottery winnings are on average closer to the threshold of a discrete purchase. This is the convexification hypothesis and it is a selection effect: different windfalls are received by different people.

If we find that durable purchases respond differently to a lottery win than to an inheritance, we need to rule out alternative interpretations. First, it could be that some other kind of selection is operative: lottery winners differ from inheritors, but not because of a need for convexification. Second, it may be that the same individuals respond to different kinds of windfall differently: it is the source of money, rather than the individual, that matters.

With respect to alternative selection stories, lottery winners may be systematically different to those who inherit with respect to other unobservable characteristics such as tastes for durables, risk aversion, or impatience. We deal with this by estimating a fixed effects regression model which isolates the within individual variation in windfalls and in durables purchases and so allows us to control for the confounding effect of unobservable characteristics on durable purchases.

Even conditional on fixed effects in the levels of durables demand, lottery winners may respond differently to income shocks for some reason other than the convexification hypothesis. We control for other differences in income effects between lottery players and inheritors by comparing responses to a lottery win and an inheritance across two types of households: those that are credit-constrained and those that are not. As noted in the previous section, we would not expect unconstrained households to use a lottery as a means of financing indivisible purchases when they have savings or are able to borrow because of the relatively high cost of gambling.[8] Focusing on those who are credit-constrained in this way, the group of people who are not credit-constrained allows us to control for differences in income effects between lottery winners and inheritors that are common to credit-constrained and unconstrained individuals. Thus, as noted above, our

---

[8] The expected return in the case of the National Lottery, for example is only 0.5.

empirical strategy amounts to examining a "difference-in-differences" in income effects: we compare estimated income effects between lottery wins and inheritances, for the credit-constrained and the unconstrained.[9]

The bottom line is that the convexification hypothesis is a selection mechanism that operates on variables (the need for durables, cash on hand) that vary through time for a given individual as their economic circumstances change, and further that operates only for the credit-constrained. By allowing for fixed effects in estimating income effects, and by double-differencing income effects (across the constrained and unconstrained, and across inheritors and lottery winners), we rule out any alternative selection mechanism which operates on time invariant unobservables, and any mechanism which is not limited to the credit-constrained. It is still possible (if improbable) that there is an alternative, time-varying selection mechanism that operates only on gamblers who are credit-constrained. We address this possibility by implementing a falsification test involving non-durable consumption. We return to the details of this test below.

With respect to the source of the money being the key difference, inheritances may be anticipated, as discussed by Hurst and Lusardi (2004). However, we present evidence below showing that household financial expectations and consumption do not adjust in anticipation of an inheritance, suggesting that at least the timing and amount of inheritance may not be anticipated.

Another way in which the source of the money may matter is if it affects what people feel that they can spend the money on. This idea was termed "emotional accounting" by Levav and McGraw (2009) and nicely summarised by Epley and Gneezy (2007) in the following way: "although all dollars are created equal, one may feel a pang of reluctance at spending grandma's inheritance on a new sports car, but little reluctance spending casino earnings doing the same." They showed in lab experiments that $200 hypothetically received by students from an ill uncle was less likely to be spent on a frivolous item than $200 from a rich uncle. In our case, compared to a "lucky" lottery win, an inheritance may come with negative associations causing consumers to avoid frivolous or hedonic purchases so as not to exacerbate any bad feelings.

---

[9] Note that this is a difference-in-difference in spending *responses*; thus it could equally be described as a triple-difference in spending *levels*.

We deal with this possibility in two ways. First, our difference-in-difference strategy controls for this possibility as long as "emotional accounting" (or other psychological explanations) operate similarly for those who are credit-constrained and those who are not.

Second, we pursue a second empirical test of the convexification hypothesis. The basis of the test is this: among the people who received an inheritance there are likely to be some who were gambling to convexify, but who lost the (endogenously-selected) gamble. We would expect these people to behave like the typical person winning the gamble rather than like the typical person receiving an inheritance. We exploit the fact that, while we do not observe people spending money on gambling, we do observe people who win small amounts (defined as less than £100). These amounts are not enough, typically, to finance consumer durables directly but they do allow us to identify people who have gambled. Thus we test whether the income effect of inheritances is larger for credit-constrained individuals who we know were gambling because we observe that they they had small winnings. If this is the case, then it makes it clear that it is who receives the windfall that matters, rather than the source of the windfall, and thus the explanation must be a selection story like the convexification hypothesis.

As a final empirical strategy which addresses both the possibility that the source of money matters and alternative selection stories which affect only those who are credit-constrained and gamble, we implement a falsification test. The convexification hypothesis suggests that among the credit-constrained, lottery players (and hence winners) are selected with respect to their marginal propensity to spend on *indivisible* goods. Our falsification test compares the income response of spending on divisible nondurable consumption (food at home and in restaurants) to lottery wins and inheritances. If we see a difference in income effects between credit-constrained lottery winners and credit-constrained inheritors, but only for the indivisible goods in our data and not for indivisible goods, it will provide further support for the convexification hypothesis.

To implement our first, difference-in-difference test, we estimate an empirical model along the following lines:

$$d_{it} = (\beta_1 + \beta_2 Q_i) Lot_{it} + (\gamma_1 + \gamma_2 Q_i) Inh_{it} + \alpha' X_{it} + u_{it} \qquad (11)$$

Where d_{it} is measure of durable purchases by agent household i at time t; $Q_i = 1$ if the agent is credit-constrained, and equals 0 otherwise; Lot_{it} and Inh_{it} are financial windfalls from lottery wins and inheritances, respectively; X_{it} is a vector of other variables that might affect purchase of durables, including age, composition of household (couple, number of kids), home-ownership status, permanent income (proxied by spending on food), employment status, financial expectations, year dummies, and u_{it} is a random error term. A possible limitation of regression adjustment is that, except in the special case of discrete independent variables and a fully-saturated model, it does not impose common support. That is, the linearity of the regression adjustment allows estimation of counterfactuals for treated units for whom there are no similar control units. To check this, we estimate propensity score models for the treatment group (credit-constrained, lottery winners) versus each control group (with our difference-in-difference design, there are multiple control groups). We then re-estimate our regression-adjusted difference-in-difference model imposing common support.

Previous empirical literature has shown that durables respond to unexpected windfalls (see Keeler and Abdel-Ghany, 1985), so we would expect that $\beta_1 = \gamma_1 \geq 0$. The theoretical considerations developed in the previous section suggest that, among credit-constrained households, selection into playing the lottery will lead to differential responses to a lottery win compared to other windfalls, in other words that $(\beta_1 + \beta_2) \neq (\gamma_1 + \gamma_2)$.

Our falsification test estimates a parallel model, with measures of nondurable (divisible) spending as the dependent variable.

To estimate our second, "small winnings" test, we estimate the following empirical model:

$$d_{it} = (\beta_1 + \beta_2 Q_i) Lot_{it} + (\gamma_1 + \gamma_2 Q_i) Inh_{it} + (\delta_1 + \delta_2 Q_i) SmLot_{it} \times Inh_{it} + \alpha' X_{it} + u_{it}$$
(12)

where $SmLot_{it} = 1$ if someone receives a lottery win of less than £100, and equals 0 otherwise. Our hypothesis is that, among credit-constrained households, those who receive a medium-sized inheritance *and* also a small lottery win will not behave like

those who only received a medium-sized inheritance, ie. $(\delta_1 + \delta_2) \neq 0$, but rather will have the larger income responses of those who receive a medium-sized lottery win, i.e. $(\beta_1 + \beta_2) = ((\gamma_1 + \gamma_2) + (\delta_1 + \delta_2))$.

4.2 Data

Our main analysis uses data taken from the British Household Panel Survey (BHPS) from 1997 – 2006 since this contains information on both durable purchases and financial windfalls. Beginning in 1991, this survey has annually interviewed members of a representative sample of around 5,500 households. On-going representativeness of the non-immigrant population is maintained by using a "following rule" – i.e. by following original sample members (adult and children members of households interviewed in the first wave) if they move out of the household or if their original household breaks up.[10] We select single and two-adult households where the head is aged 20 – 70. Our analysis sample contains information on 6,148 households (29,886 observations).

4.3 Descriptive Statistics

Consumer Durables

We focus on durables that they are largely unchanged over the period and that are genuinely "lumpy" to purchase new. This means we exclude, for example, VCRs which were becoming increasingly obsolete towards the end of the period, and microwaves and CD players where the typical expenditure is fairly low. We include televisions, fridge/ freezers, washing machines, tumbledriers, dishwashers and home computers. On average, 36% households had purchased at least one of these six durables over the previous year; 12% purchased two or more. This is a set of basic durables that most households seek to replace on a regular basis.

In principle, households could potentially smooth their spending on new durables. One possibility is renting, although this may be easier for some durables (televisions, for example) than for others (fridge-freezers). Also, most rental companies have a minimum rental period of 12 or 18 months and require a credit check, so the option of renting may not be open to everyone. Similarly, hire purchase companies also require a credit check

---

[10] The survey incorporated booster samples from Scotland and Wales in 1999 and Northern Ireland in 2001, but we restrict our sample to original sample members.

and may charge high interest rates if the repayments are made over a long period. Compared to these alternatives, buying a lottery ticket may not be an unattractive option.[11]

Credit Constraints

We define credit-constrained households as those with no (income from) savings or investments. This is a broad definition by which around half of all households are constrained and will include some households who are not credit-constrained in that they can borrow, even if they have no savings.[12] The benefit of our approach is that it yields reasonable sample sizes in each of group. In sensitivity analysis, we have used narrower definitions that exclude anyone who owns their home and anyone with household income in the top two-third of the distribution and found similar results.

Lottery wins and inheritances

Since 1997, the BHPS has asked individuals whether they have received any of the following financial windfalls in the previous 12 months: a gambling win, an inheritance, a life insurance payment, a pension lump sum, a personal accident claim or a redundancy payment. Our comparison focuses on gambling wins (referred to here as lottery wins since this is likely to be the case for most) and inheritances since the other windfalls may largely be anticipated (such as pension lump-sums), as we show below, and/or may be associated with events that directly affect the purchase of durables (such as redundancy payments).[13]

In the sample as a whole, 21 per cent reported Lottery winnings, while 5 per cent reported an inheritance. However, the average amounts received in the two cases are very different. The mean (median) Lottery winningswas £290 (£40) compared to £29,949 (£5,000) in the case of inheritances. This is not surprising given the structure of National

---

[11] There are rental outlets that specifically target those with poor credit histories which do not require a formal credit check, only five references. The advertised APR is 30%, but additional insurance which consumers are "strongly advised" to take out typically increases the effective rate of interest to more than 100% (Collard and Kempson, 2005).

[12] Young and Waldron (2008) show that 16% of the UK population is credit-constrained, according to self-reported constraints in the amount that they could borrow, including both *perceived* constraints that discouraged them from applying for credit, and *actual* constraints where the household was prevented from borrowing either by the unavailability of credit or its high price. This is similar to Jappelli (1990) for the US who found that c. 20% of US households are credit-constrained based on survey evidence that they have been refused credit, or put off applying for fear of refusal.

[13] We exclude any inheritances that are linked to widow(er)hood, i.e. deaths within the household that may have an immediate effect on durable purchase.

Lottery payouts.[14] However, this raises issues for our analysis; in particular, how to ensure that we pick up the response to a lottery win compared to inheritance and not responses to different sized windfalls. Landsberger (1966) and Keeler and Abdel-Ghany (1985), for example, show that the size of the windfall affects what people do with it, with smaller windfalls being more likely to be spent.

Our approach is to focus on "medium-sized" windfalls of between £100 and £5,000. Anyone who receives a windfall of more than £5,000 in any wave is dropped from the analysis and in our initial analysis we ignore small (< £100) lottery wins and inheritances. Focusing on medium wins seems appropriate given our interest in consumer durables: larger wins may be associated with more widespread lifestyle changes such as moving house, while smaller wins may not be enough to finance the purchase of the white goods we focus on. Furthermore, restricting windfalls to this narrower range makes the average lottery win more comparable in size to the average inheritance. Within the range £100 - £5,000, lottery wins are still smaller on average than inheritances, as shown in Table 2, but the difference is much smaller. In sensitivity analysis (details available on request), we found similar results with narrower ranges of £100 - £1000 and £1,001 - £5,000.

[Table 2, Descriptive Statistics, about here]

Table 2 summarises separately average windfall payments for those who are credit-constrained and those who are not. Many of those who are not credit-constrained receive windfalls from lottery wins, and indeed a higher proportion than among those who are credit-constrained. This is not inconsistent with people gambling to convexify, but is a reminder that this is only one of several possible motives for gambling. The BHPS does not contain information on who has gambled and lost. To provide direct evidence on who gambles and how gambling varies with total expenditure, we use data from the 2007 UK Expenditure and Food Survey. Figure 4 shows that budget shares on gambling decline markedly with total expenditure, consistent with the need to gamble to convexify being concentrated among low income groups. Figure 4 also shows that the fraction of households with positive gambling expenditure is around 40% across a wide

---

[14] The odds of winning £10 are 1:57, compared with odds of 1:1,031 to win around £100, 1:55,490 to win around £1,000, 1:2,330,636 to win around £100,000 and 1:13,983,817 to hit the jackpot.

range of incomes, again consistent with the idea of there being more than one motive for gambling. In fact, the existence of lottery winners who are not credit-constrained is necessary for the difference-in-difference strategy described above.

[Figure 4, about here]

Returning to Table 2, within the range we focus on (£100 - £5,000) there is no statistically significant difference in average windfall size between those who are potentially credit-constrained and those who are not. Also, there is no statistically significant difference in household income between those who receive a medium-sized lottery win and those who receive a medium-sized inheritance. This is reassuring for our difference-in-difference specification.

By contrast, the first row of Table 2 shows that there are clear differences across groups in the raw numbers of how many durables are being purchased. While the average number of durables being purchased is about 0.5 for the unconstrained, this number rises to 0.7 for the credit-constrained who have had a windfall due to a lottery win (but not for the credit-constrained inheritors). The aim of the detailed empirical analysis below is to understand how much of this difference in the raw numbers is due to the economic circumstances of those who have chosen to gamble.

5. Empirical Results

5.1. Main Results

Our main results, addressing the question "Do durable purchases respond differently to lottery wins than to inheritances?", are shown in Table 3. Column (1) presents the results of an OLS regression of durable purchases on lottery winnings, inheritances, and additional controls. We model the number of durables purchased during the previous twelve months as the dependent variable, but results are very similar using a binary indicator for whether or not the household purchased any durables. We include both lottery wins and inheritances as amount won (in £'00s) to deal with the fact that, even within our narrower range of "medium-sized" wins, the typical lottery win is quite a bit smaller than the typical inheritance.

The results indicate a significant response of durables purchases to a lottery win, but not to an inheritance. This is consistent with the convexification hypothesis, but, as

described in the previous section this finding is also consistent with a number of alternative interpretations. In Column (2) we interact the windfall variables with a dummy variable indicating whether the agent is credit-constrained. This corresponds to equation (13) above and implements our main "difference-in-different" test. The results show that the propensity to purchase durables out of (endogenously-selected) lottery wins is significantly greater than the propensity to purchase durables out of an (exogenously-determined) inheritance for households that are credit-constrained. By contrast, there is no significant difference among unconstrained households. In Column (3) we show that these results are robust to the inclusion of individual fixed-effects.

[Table 3 about here]

The convexification hypothesis is a selection mechanism that operates on variables (the need for durables, cash on hand) that vary through time, and that operates only for the credit-constrained. The results in Column (3), which allow for fixed effects in estimating income effects, and which double-difference income effects (across the constrained and unconstrained, and across inheritors and lottery winners), rule out any alternative selection mechanism which operates on time invariant unobservables, and any selection mechanism which is not limited to the credit-constrained.

Column (4) repeats these estimates but imposing common support in the probability, given characteristics, of being a credit-constrained lottery winner (a propensity score). The estimates are very similar to those in column (3).

[Table 4 about here]

Our results provide statistically significant evidence that consumers behave as our theory predicts: consumers gamble to convexify. Table 4 draws out the magnitudes of these estimates. A £500 gambling win results in a 0.09 increase in the number of durables purchased within the year. This is an increase of 19% of the purchase rate of those who are neither winners nor inheritors. The corresponding numbers for a £500 inheritance are a 0.01 increase in the number, which is a 2% increase in the purchase rate. As discussed above there may be differences between lotteries and inheritances which are common to credit-constrained and unconstrained consumers (for example, issues of emotional accounting or anticipation). The bottom two rows of table 4 show that the difference in magnitudes between gambling wins and inheritances are little affected by subtracting the

effect that is common to the constrained and unconstrained. The overall message of this table is that, among the credit-constrained, the income effect of an inheritance is less than 15% of the income effect of a gambling win of the same size.

These numbers suggest using lottery wins to estimate income effects may get the population average income effect wildly wrong. External validity is very weak because selection into the "experiment" is directly related to the outcome of interest.

5.2. Are Inheritances Anticipated?

As noted in the previous section, one potential concern is that inheritances may differ from lottery wins in being reasonably well anticipated by the individual. However, there is no evidence of households adjusting either their financial expectations or their durable purchases ahead of receiving an inheritance. Table 5 reports the results of a fixed effects regression of a binary indicator for whether the (head of the) household expects their financial situation to improve over the next 12 months on a set of indicators for whether or not the household does in fact receive a lottery win, an inheritance or one of the other financial windfalls (life insurance payment, pension lump sum, personal accident claim, redundancy payment) over the following 12 months, focusing on medium-sized windfalls (between £100 - £5,000). Only the coefficient on other windfalls is positive and significant; medium inheritances do not appear to be anticipated. Consistent with this, sensitivity analysis (details available on request) that included lead terms in the regression analysis to pick up the effect of any anticipated windfalls found no significant effects.

[Table 5 about here]

5.3. The Small Winnings Test

Section 4 proposed a second test of the convexification hypothesis based on the idea that individuals who receive an inheritance but whom we know to have been gambling (because they report very small winnings) should behave like lottery winners (rather than like inheritors). The results in Table 6 show that this is exactly what we find in our data. Column (1) reproduces (from Column (3) of Table 3) our main difference-in-difference with individual fixed effects results, corresponding to equation (12). In Column (2) we report estimates of equation (13) in which we interact the inheritances variables with a dummy indicating whether the agent reported a small lottery win. We

find that credit-constrained inheritors that we know to have been gambling exhibit much larger income effects than other inheritors. In fact, their responses are not statistically different from the lottery winners. This test provides further confirmation that our findings in the previous section were not driven by differences in the way individuals respond to lottery winnings compared to inheritances. Instead, it is the characteristics and situation of the person who receives the money that matters. Credit-constrained gamblers have larger responses and this is consistent with the idea that they are a selected group: close to a purchase margin.

[Table 6 about here]

5.4. Falsification Test

Finally, in Table 7 we provide a falsification test based on the fact that the convexification hypothesis should generate differences in income effects only for indivisible goods. We therefore run the same regression but with weekly household spending on food for home consumption or, separately, food out (in restaurants) on the left-hand side.[15] Food at home is a necessity, while food in restaurants (like many of the durable goods we study) is a luxury. However, both are divisible goods.

Consistent with the convexification hypothesis and with the predictions of a standard life-cycle model, we find zero income effects for both lottery wins and inheritance receipts when we examine spending on food for home consumption. For food at home we do find differences between inheritances and lotteries, with a significant income effect only for the latter. This is consistent with the two sources being treated differently – the source of the money matters. Crucially, however, this difference is common to the credit-constrained and to the unconstrained, as would be expected for a good that is divisible. As described above, our difference-in-difference strategy deliberately removes such effects. Thus we find differential responses of the credit-constrained only for indivisible goods, consistent with our model of gambling to convexify.

[Table 7 about here]

Overall, our empirical results are fully consistent with the theory presented in Section 2 and lend strong support to the idea that consumers gamble to convexify.

---

[15] In the BHPS, the food data are banded and we take the mid-points.

6. Conclusion

This paper sheds light on why risk averse individuals gamble. We find that purchases of durables are more responsive to lottery winnings than to inheritances and argue that this difference in durable purchase is consistent with individuals gambling to convexify. We find this difference while controlling for individual fixed effects; we find it only for households who are credit-constrained who might have an incentive to gamble; and we find it only in expenditures on indivisible goods. Moreover, we find that an inheritance has a larger effect when received by individuals who were playing the lottery: it is the economic circumstances of the recipient, rather than the source of the money, that matters. These multiple lines of evidence are all consistent with gambling to convexify: the larger income effects associated with lottery winnings result from the fact that it is exactly those who are close to a purchase threshold who choose to gamble. It is hard to think of another plausible explanation that is consistent with all these multiple lines of evidence.

Our findings are important for a number of reasons. First, our findings provide at least a part of the explanation for gambling among low-income households, and also the popularity of so-called prize-linked savings products amongst these households. Of course, there are other reasons for people to gamble besides convexification. Recently, it has been suggested that prize-linked savings products could be used to promote saving through the "excitement factor" and this may be an important factor for many. However, our research points to another potential reason why such products might appeal to low-income households; they may allow consumers to overcome indivisibilities potentially more quickly than conventional savings products. On the other hand, given the poor return to playing lotteries, our evidence that individuals are gambling to finance indivisible purchases highlights the lack of financing options available to poor households, and the severity of the credit constraints they face.

Second, our finding highlights the difficulty of finding appropriate instruments – in this case for income. The random success of winning a gamble would seem to make it a natural instrument for unanticipated income changes that should allow for the identification of income effects, as in Imbens et al. (2001). In their conclusion, Imbens et

al. (2001) note the caveat that they have "no direct evidence concerning the difference of responses to lottery income versus other sources of unearned income" (page 793). We provide both evidence of and economic motivation for such differences. The decision to play lotteries is itself an economic decision that is not independent of the income effect on winning: those who choose to play lotteries will display larger income effects than the rest of the population. In our data the difference in magnitudes is very substantial: among the credit-constrained, the income effect of an inheritance is less than 15% of the income effect of a gambling win of the same size.

Finally, our evidence on how individuals deal with non-convexities provides guidance on how we should model the behaviour of those individuals in structural models. In particular, the appropriate way to model optimisation by individuals in the presence of non-convexities is to allow those individuals to play lotteries.

REFERENCES

Bailey, M. J., M. Olson, and P. Wonnacott, (1980). "The Marginal Utility of Income Does Not Increase: Borrowing, Lending, and Friedman-Savage Gambles," *The American Economic Review*, vol. 70, no. 3, June 1980, pages 372-379.

Besley, T., S. Coate, and G. Loury, (1993). "The Economics of Rotating Savings and Credit Associations," *The American Economic Review*, vol. 83, no. 4, September 1993, pages 792-810.

Chetty, R. and A. Szeidel, (2007), "Consumption Commitments and Risk Preferences", *The Quarterly Journal of Economics*, Vol. 122, no. 2, pages 831-877.

Collard, S. and E.Kempson, (2005), *Affordable Credit: The Way Forward*, Bristol: Policy Press

Epley, N., and A. Gneezy, (2007), "The framing of financial windfalls and implications for public policy." *Journal of Socio-economics*, 36, 36-47

Handa, S. and C. Kirton, (1999), "The economics of rotating savings and credit associations: Evidence from the Jamaican 'Partner'" *Journal of development Economics*, 60:173–194.

Hansen, G., (1985), "Indivisible labor and the business cycle" *Journal of Monetary Economics* 16: pp. 309—328

Hartley, R., and L. Farrell, (2002), "Can Expected Utility Theory Explain Gambling?" *American Economic Review*, vol 92(3) pp. 613-624.

Heckman, J.J. and J.A. Smith, (1995), "Assessing the Case for Social Experiements."*Journal of Economic Perspectives*, 9(2):85-110

Hurst, E. and A. Lusardi, (2004), "Liquidity constaints, household wealth and entrepreneurship," *Journal of Political Economy*, vol 112, pp. 318-47

Imbens, G., D.B. Rubin and B.I. Sacedote, (2001), "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players," *The American Economic Review*, 91(4):778-794

Jappelli, T., (1990), "Who is Credit-constrained in the U.S Economy," *The Quarterly Journal of Economics*, Vol. 105(1), pp. 219-262.

Keeler, J., and M. Abdel-Ghany, (1985), "The relative size of windfall income and the permanent income hypothesis", *Journal of Business and Economic Statistics*, vol 3, no. 3, pp. 209-215.

Kearney, M.S., P. Tufano, J. Guryan, and E. Hurst, (2010). "Making Savers Winners: An Overview of Prize-lined Savings Products." NBER Working Paper 16433.

Landsberger, M. (1966), "Windfall Income and Consumption: Comment," *The American Economic Review*, 56: 534-40

Levav, J. and P. McGraw (2009) "Emotional Accounting: Feelings About Money and Consumer Choice." *Journal of Marketing Research*, 46:66-80

Lentz, R. and T. Tranaes (2005), "Job Search and Savings: Wealth Effects and Duration Dependence", *Journal of Labor Economics*, 23(3), 467-490.

Ng, Y., (1965), "Why Do People Buy Lottery Tickets? Choices Involving Risk and the Indivisibility of Expenditure," *The Journal of Political Economy*, vol. 73, no. 5, , pages 530-535.

Rogerson, R., (1988) "Indivisible labor, lotteries and equilibrium," *Journal of Monetary Economics*, vol. 21(1), pages 3-16.

Tufano, P., (2008), "Savings Whilst Gambling: An Empirical Analysis of U.K. Premium Bonds." *American Economic Economic, Review P&P,* 98(2): 321-26.

Young, G. and M. Waldron (2006) "The state of British household finances: results from the 2006 NMG Research Survey", Bank of England Quarterly Bulletin, pp.397-403

Table 1: Probabilities of Purchase: Chosen Lotteries versus Random Lotteries

	Lottery Winner $\text{Prob}(d=1\|winner)$	Lottery Loser (had ticket, lost) $\text{Prob}(d=1\|lost)$	Non-winner* (losers + non-holders) $\text{Prob}(d=1\|non\,winner)$
Lottery Chosen	**1**	**0**	$\dfrac{1-F}{1-q\lambda}$
Random compulsory lottery	$1-F\left(x_2^*-\dfrac{1-q}{q}\right)$	$1-F(x_2^*+1)$	$\dfrac{(1-\lambda q)(1-F)}{1-\lambda q}$
Difference	$F\left(x_2^*-\dfrac{1-q}{q}\right)\geq 0$	$F(x_2^*+1)-1\leq 0$	$\dfrac{\lambda q(1-F)}{1-\lambda q}>0$
Approx Bias: (Diff in diff)	$1+F\left(x_2^*-\dfrac{1-q}{q}\right)-F(x_2^*+1)>0$		$\dfrac{F\left(x_2^*-\frac{1-q}{q}\right)-\lambda q\left(1-\left(F(x_2^*)-F\left(x_2^*-\frac{1-q}{q}\right)\right)\right)}{1-\lambda q}$

* The probabilities of a non-winner purchasing the durable good are approximations to the actual probabilities because the exact CDF's are calculated at different points. Hence the probability of purchase by a non-winner when the lottery is chosen is given by: $\left(1-F(\bar{x}_1)\right)/(1-q\lambda)$, and the probability of purchase by a non-winner when the lottery is random is given by:

$$\frac{(1-q)\lambda\left(1-F\left(x_2^*+1\right)\right)+(1-\lambda)\left(1-F\left(x_2^*\right)\right)}{1-\lambda q}.$$

However, since $\bar{x}_1$ lies between $\left(x_2^*+1\right)$ and x_2^*, evaluating each CDF at the same value of x_2^* is a reasonable approximation.

Table 2: Descriptive Statistics

	Unconstrained					Credit-constrained				
	All	Winners	Non-winners	Inheritors	Non-inheritors	All	Winners	Non-winners	Inheritors	Non-inheritors
Number of durables	0.516	0.544	0.514	0.640	0.512	0.493	0.714	0.486	0.574	0.492
Age	45.5	46.0	45.5	45.6	42.7	41.7	42.4	41.7	38.6	41.7
Income	£2,959	£3,142	£2,948	£3,200	£2,952	£2,063	£2,559	£2,046	£2,378	£2,059
Degree	0.217	0.131	0.223	0.220	0.217	0.106	0.066	0.108	0.174	0.106
Mean kids	0.552	0.476	0.558	0.532	0.553	0.782	0.653	0.786	0.721	0.783
Couple	0.762	0.859	0.756	0.821	0.760	0.599	0.738	0.594	0.705	0.598
Mean windfall		£514		£1,984			£595		£1,875	
Median windfall		£224		£1,500			£250		£1,200	
N	13,757	802	12,955	363	13,394	16,129	497	15,632	190	15,939

Table 3: Main regression results

	Number of durables purchased in last 12 months			
	(1)	(2)	(3)	(4)
	OLS	OLS	FE	FE
				common support
Winner $_{it}$ (β_1)	.0121**	.0052*	.0038	.0034
	(.0023)	(.0031)	(.0036)	(.0037)
Winner $_{it}$ * Constrained$_i$ (β_2)		.0160**	.0145**	.0150**
		(.0047)	(.0052)	(.0055)
Inheritor$_{it}$ (γ_1)	.0029	.0032	.0044**	.0043**
	(.0014)	(.0017)	(.0018)	(.0019)
Inheritor$_{it}$ * Constrained$_i$ (γ_2)		-.0008	-.0007	-.0022
		(.0029)	(.0032)	(.0033)
Tests (p-values)				
$\beta_1 = \gamma_1$	.0007	.5717	.8895	.8402
$(\beta_1 + \beta_2) = (\gamma_1 + \gamma_2)$		.0000	.0021	.0009

Standard errors in brackets. ** denotes statistically significant at the 5% level; * at the 10% level. Credit-constrained = no income from savings/ dividends; Lottery win (Winner) and Inheritance (Inheritor) are in £'00s. Other controls: Age of head of household and age squared; couple; indicators for number of children; home-owner; head of household is unemployed, retired, other non-work; spending on food (permanent income), financial expectations for next year, year dummies.

Table 4: Implications

		Effect of £500 Windfall		
		(1)	(2)	(3)
		Lotteries	Inheritances	Difference
Total Effect for Credit-constrained	Number of Durables	0.092	0.010	0.082[1]
	Percent of non-windfall number purchased	18.8%	2.1%	16.7%[2]
Effect of Credit-constrained – Effect for Unconstrained	Number of Durables	0.075	-0.011	0.086[3]
	Percent of non-windfall number purchased	15.3%	-2.2%	17.5%[4]

Implications are calculated using estimates from column 4 in table 3 (i.e. the Fixed Effect estimates, imposing common support).

[1] $[(\beta_1 + \beta_2) - (\gamma_1 + \gamma_2)]*5$; [2] $[(\beta_1 + \beta_2) - (\gamma_1 + \gamma_2)]*5/0.49*100$;

[3] $[\beta_2 - \gamma_2]*5$; [4] $[\beta_2 - \gamma_2]*5/0.49*100$

able 5: Are Windfalls Expected?

Dependent variable: (0/1), whether household head expects financial situation to improve over the next 12 months (Results of fixed effects regression)

	Whole sample	Credit-constrained	Not credit-constrained
Winner $_{t+1}$	0.0084 (0.0152)	-0.0099 (0.0254)	0.0122 (0.0202)
Inheritance $_{t+1}$	0.0181 (0.0253)	0.0340 (0.0440)	-0.0199 (0.0330)
Other windfall $_{t+1}$	0.0349** (0.0130)	0.0142 (0.0211)	0.0579** (0.0178)
N	27,410	14,508	12,902

Standard errors in brackets. ** denotes statistically significant at the 5% level. "Other windfalls" include life insurance policy payments, pension lump-sums, redundancy payments, personal accident claims and "anything else"

Table 6: Small Winnings Test

	Main results	Further test
Winner $_{it}$ (β_1)	.0038 (.0036)	.0038 (.0036)
Winner $_{it}$ * Constrained$_i$ (β_2)	.0145** (.0052)	.0145** (.0052)
Inheritor$_{it}$ (γ_1)	.0044** (.0018)	.0045** (.0021)
Inheritor$_{it}$ * Constrained$_i$ (γ_2)	-.0007 (.0032)	-.0027 (.0034)
Small Win * Inheritor (δ_1)		-.0006 (.0042)
Small Win * Inheritor * Constrained$_i$ (δ_2)		.0187** (.0096)
Tests (p-values)		
$\beta_1 = \gamma_1$	.8895	.8705
$(\beta_1 + \beta_2) = (\gamma_1 + \gamma_2)$	.0021	.0006
$(\delta_1 + \delta_2) = 0$		.0393
$(\beta_1 + \beta_2) = ((\gamma_1 + \gamma_2) + (\delta_1 + \delta_2))$		.8575

** denotes statistically significant at the 5% level
Credit-constrained = no income from savings/ dividends; Lottery win Winner) and Inheritance (Inheritor) are in £'00s. Small Win is an indicator if the household receives a lottery win of less than £100. Other controls as in table 3.

Table 7: Falsification Tests

	Number of durables	Spending on food	Spending on meals out
Winner $_{it}$ (β_1)	.0031	-.0005	.0077**
	(.0036)	(.0013)	(.0029)
Winner $_{it}$ * Constrained$_i$ (β_2)	.0106**	.0018	-.0041
	(.0053)	(.0020)	(.0043)
Inheritor$_{it}$ (γ_1)	.0041**	.0008	.0014
	(.0018)	(.0007)	(.0015)
Inheritor$_{it}$ * Constrained$_i$ (γ_2)	-.0008	-.0007	.0046
	(.0032)	(.0011)	(.0026)
<u>Tests (p-values)</u>			
$\beta_1 = \gamma_1$	.7837	.3834	.0533
$(\beta_1 + \beta_2) = (\gamma_1 + \gamma_2)$	.0346	.4710	.5474

** denotes statistically significant at the 5% level
Credit-constrained = no income from savings/ dividends; Lottery win (Winner) and Inheritance (Inheritor) are in £'00s. Other controls as in table 3.

Figure 1: Durable Purchase Decision

# Figure 2: Lottery and Durable Purchase Decisions

Figure 3: Chosen Lottery versus Random Lottery

Difference in probability of purchase

$Prob(d=1 \mid winner, choice) - Prob(d=1 \mid winner, rand)$
(high η)

$Prob(d=1 \mid winner, choice) - Prob(d=1 \mid winner, rand)$
(low η)

$Prob(d=1 \mid non-winner, choice) - Prob(d=1 \mid non-winner, rand)$

q: probability of lottery win

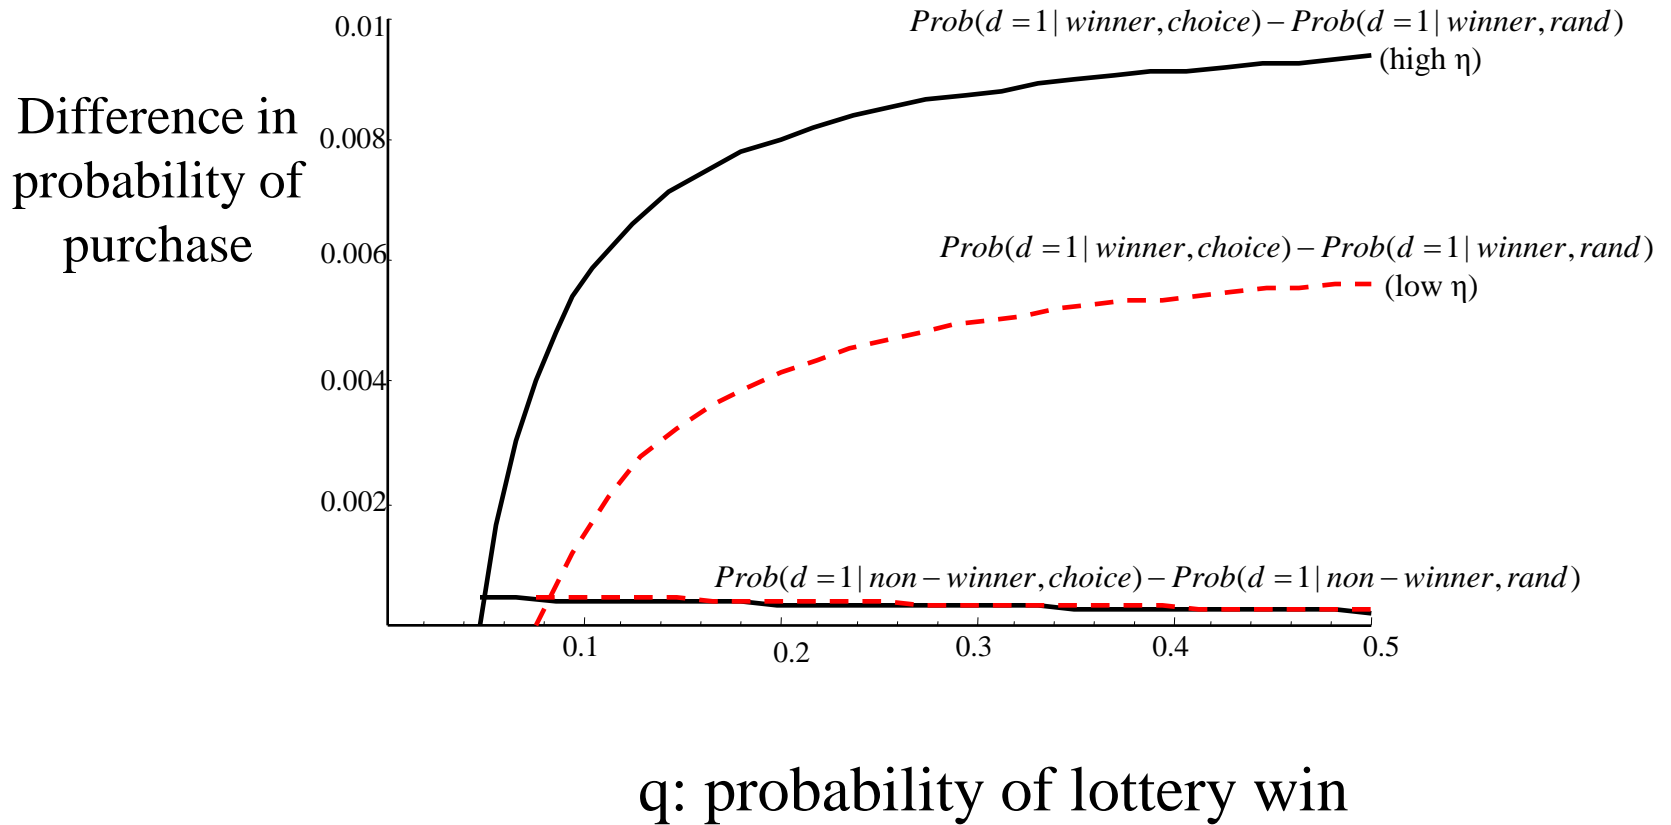# Figure 4: Household Gambling Expenditure, 2007 EFS
Budget Shares and % with Positive Expenditure

Appendix: Proofs of Results 2 and 3

Proof of Result 2

Result 2: *Lottery tickets are not purchased outside the interval* $x_1 \in \left[x_2^* - \frac{1-q}{q}, x_2^* + 1 \right]$.

Proof:

The value functions for not buying a lottery ticket is given by:

$$V_1^{l=0} = \begin{cases} u(x_1) & \text{if } x_1 < x_2^* \\ u(x_1 - p) + \eta & \text{if } x_1 \geq x_2^* \end{cases}$$

The expected value function for buying a ticket is given by:

$$E\left[V_1^{l=1}\right] = q \begin{cases} u\left(x_1 + \frac{1-q}{q}\right) & \text{if } x_1 < x_2^* - \frac{1-q}{q} \\ u(x_1 + \frac{1-q}{q} - p) + \eta & \text{if } x_1 \geq x_2^* - \frac{1-q}{q} \end{cases} + (1-q) \begin{cases} u(x_1 - 1) & \text{if } x_1 < x_2^* + 1 \\ u(x_1 - 1 - p) + \ln \eta & \text{if } x_1 \geq x_2^* + 1 \end{cases}$$

Now consider separately the incentive to buy a lottery ticket when cash-on-hand is below the interval and above the interval.

1) When

$$x_{1^{.}} < x_2^* - (1-q)/q,$$

cash-on-hand in period 2 will be sufficiently low that even if the lottery is won, $x_2 < x_2^*$, and so the household does not buy the indivisible good, regardless of the lottery outcome. Thus, the expected value of buying a lottery ticket becomes:

$$E\left[V_1^{l=1}\right] = qu\left(x_1 + \frac{1-q}{q}\right) + (1-q)u(x_1 - 1)$$

The value of not buying becomes: $V_1^{l=0} = u(x_1)$. Since the gamble is actuarially fair and utility, u, is concave, the value of not buying a lottery ticket is always greater than the expected value of buying the lottery ticket:

$$V_1^{l=0} = u(x_1)$$

$$\geq qu\left(x_1 + \frac{1-q}{q}\right) + (1-q)u(x_1 - 1) = E\left[V_1^{l=1}\right].$$

2)) When

$$x_{1\cdot} > x_2^* + 1,$$

cash-on-hand in period 2 will be sufficiently high that even if the lottery is lost, $x_2 > x_2^*$, and so the household buys the indivisible good regardless of the lottery outcome. Thus, the expected value of buying a lottery ticket becomes:

$$E\left[V_1^{l=1}\right] = qu\left(x_1 + \frac{1-q}{q} - p\right) + (1-q)u(x_1 - 1 - p) + \eta$$

And the value of not buying becomes:

$$V_1^{l=0} = u(x_1 - p) + \eta$$

Since the gamble is actuarially fair and utility, u, is concave, the value of not buying the lottery ticket is always greater than the value of buying the ticket.

$$V_1^{l=0} = u(x_1 - p) + \eta$$

$$\geq qu\left(x_1 - p + \frac{1-q}{q}\right) + (1-q)u(x_1 - p - 1) + \eta = E\left[V_1^{l=1}\right],$$

Proof of Result 3

Result 3: *There exists a region,* $x_1 \in \left[\underline{x_1}, \overline{x_1}\right]$, *which contains* x_2^* $\left(\underline{x_1} < x_2^* < \overline{x_1}\right)$, *in which the agent will purchase a lottery ticket.*

Proof:

We consider the incentive to buy a lottery ticket in the region of x_2^*. Define the difference in utility from purchasing the indivisible good and not purchasing it as

$$\delta = u(x_2 - p) + \eta - u(x_2)$$

We consider separately the incentive ε above and ε below x_2^*.

1) **Below** x_2^*: When

$$x_2 = x_2^* - \varepsilon$$
$$and \; so \; \delta < 0$$

we can write the expected value of buying a lottery ticket as:

$$E\left[V_1^{l=1}\right] = q\left(u\left(x_2^* - \varepsilon + \frac{1-q}{q} - p\right) + \eta\right) + (1-q)u\left(x_2^* - \varepsilon - 1\right)$$

And the value of not buying a ticket as:

$$V_1^{l=0} = u\left(x_2^* - \varepsilon\right)$$
$$= q\left(u\left(x_2^* - \varepsilon - p\right) + \eta\right) - q\delta + (1-q)u\left(x_2^* - \varepsilon\right)$$

$$E\left[V_1^{l=1} - V_1^{l=0}\right] = q\left(u\left(x_2^* - \varepsilon + \frac{1-q}{q} - p\right) + \eta - u\left(x_2^* - \varepsilon - p\right) - \eta\right)$$
$$+ q\delta$$
$$+ (1-q)\left(u\left(x_2^* - \varepsilon - 1\right) - u\left(x_2^* - \varepsilon\right)\right)$$

This is approximately equal to:

$$E\left[V_1^{l=1} - V_1^{l=0}\right] = q\left(u'(x_2 - \varepsilon - p)\left(\frac{1-q}{q}\right)\right)$$
$$+ q\delta$$
$$+ (1-q)\left(-u'(x_2 - \varepsilon)\right)$$

$$E\left[V_1^{l=1} - V_1^{l=0}\right] = (1-q)\left(u'(x_2 - \varepsilon - p) - u'(x_2 - \varepsilon)\right) + q\delta$$
$$= -p(1-q)u''(x_2 - \varepsilon) + q\delta$$

As

$$\varepsilon \to 0, \; x_2 \uparrow x_2^*, \delta \uparrow 0$$

and

$$E\left[V_1^{l=1} - V_1^{l=0}\right] > 0$$

2) **Above x_2^*:** When

$$x_2 = x_2^* + \varepsilon$$
$$and \; so \; \delta > 0$$

$$E\left[V_1^{l=1}\right] = q\left(u\left(x_2^* + \varepsilon + \frac{1-q}{q} - p\right) + \eta\right) + (1-q)u\left(x_2^* + \varepsilon - 1\right)$$

The value of not buying a ticket is:

$$V_1^{l=0} = u\left(x_2^* + \varepsilon - p\right) + \eta$$
$$= q\left(u\left(x_2^* + \varepsilon - p\right) + \eta\right) + (1-q)\left(u\left(x_2^* + \varepsilon - p\right) + \eta\right)$$
$$= q\left(u\left(x_2^* + \varepsilon - p\right) + \eta\right) + (1-q)u\left(x_2^* + \varepsilon\right) + (1-q)\delta$$

$$E\left[V_1^{l=1} - V_1^{l=0}\right] = q\left(u\left(x_2^* + \varepsilon + \frac{1-q}{q} - p\right) + \eta - u\left(x_2^* + \varepsilon - p\right) - \eta\right)$$
$$+ (1-q)\left(u\left(x_2^* + \varepsilon - 1\right) - u\left(x_2^* + \varepsilon\right)\right)$$
$$- (1-q)\delta$$

$$E\left[V_1^{l=1} - V_1^{l=0}\right] = q\left(u'(x_2 + \varepsilon - p)\left(\frac{1-q}{q}\right)\right)$$
$$+ (1-q)\left(-u'(x_2 + \varepsilon)\right)$$
$$- (1-q)\delta$$

$$E\left[V_1^{l=1} - V_1^{l=0}\right] = (1-q)\left(u'(x_2 + \varepsilon - p) - u'(x_2 + \varepsilon)\right) - (1-q)\delta$$

$$= -p(1-q)u''(x_2 + \varepsilon) - (1-q)\delta$$

As

$$\varepsilon \to 0, \ x_2 \downarrow x_2^*, \delta \downarrow 0$$

and

$$E\left[V_1^{l=1} - V_1^{l=0}\right] > 0$$