



Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Do consumers gamble to convexify?☆

Thomas F. Crossley^{a,*}, Hamish Low^b, Sarah Smith^c^a University of Essex and Institute for Fiscal Studies (IFS), United Kingdom^b University of Cambridge and IFS, United Kingdom^c University of Bristol and IFS, United Kingdom

ARTICLE INFO

Article history:

Received 8 April 2016

Received in revised form 20 July 2016

Accepted 28 July 2016

Available online 31 August 2016

JEL classifications:

E21

D12

D81

L83

C18

Keywords:

Lotteries

Income effects

Consumption

Durables

External validity

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. This idea has been demonstrated theoretically, but not explored empirically. We test this idea by focusing on a key implication: 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-constrained. This is consistent with credit-constrained, risk-averse agents gambling to convexify their budget set.

© 2016 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Why do risk-averse consumers sometimes gamble? One idea, first proposed by Ng (1965), is that discreteness in spending or in labor supply opportunities can induce local non-concavities in the value functions of risk-averse agents. This generates local risk-loving behavior and makes it rational to gamble in order to have a chance of crossing the threshold required to finance a lumpy purchase. A similar idea was advocated by Chetty and Szeidel (2007) in the context of committed consumption: when individuals are close to the region of needing to change their commitments, it may be optimal to gamble in order to cross the threshold of making the change. Bailey and Olson Wonnacott (1980) argued that access to credit markets made such gambling irrational, but Hartley and Farrell (2002) showed theoretically that rational gambling

☆ This work was supported in part by the ESRC-funded Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (grant number RES-544-28-5001.) For helpful comments and suggestions we thank Sule Alan, Helen Simpson and seminar participants at Royal Holloway University of London, McMaster University, Leicester University, Naples, University College London, the Lincoln College Applied Microeconometrics Conference at the University of Oxford, and the 2015 SOLE-EALE World Meetings. All errors are our own.

* Corresponding author at: Department of Economics, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, United Kingdom.
E-mail addresses: tcross@essex.ac.uk, tfcrossley@gmail.com (T.F. Crossley).

might still occur where borrowing and lending rates differ, where capital market imperfections exist, or if individuals' time preference rates differ from interest rates.¹

The first contribution of this paper is to use a model to develop the idea that households might gamble to cross purchase thresholds. Our analysis shows that this mechanism implies that lottery players will have systematically different income effects from nonplayers. The second contribution is to provide empirical evidence that some individuals do appear to play the lottery as a strategy for purchasing discrete goods. We do not suggest that financing discrete purchases provides the only – or even the main – motivation for gambling,² but we show that it may be important for credit-constrained households.³ In particular, it may help to explain infrequent lottery purchases, which amount to nearly 40% of the total ([Gambling Commission, 2014](#)).

There are several reasons why it is important to know whether, in practice, consumers gamble to convexify choice sets. First, there is broad interest in whether credit constraints and indivisibilities in consumption pose particular challenges for poorer households. If so, it is important to understand what strategies poor households use to overcome those challenges. For example, [Mullainathan and Shafir \(2009\)](#) discuss the role of lotteries in allowing poorer households to achieve “small to big transformations”.⁴

Second 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 Tranaes, 2005](#)). It is interesting to know whether this is simply a technical convenience, or if it in fact captures the way that individuals behave when faced with non-convexities.

Third, following [Imbens et al. \(2001\)](#), lotteries have been used to identify income effects across a wide range of spheres, including consumption and labor supply.⁵ As with all instrumental variable estimates that identify treatment effects among a sub-group, the external validity of these results is a crucial issue and many papers acknowledge that their estimated income effects are valid only for a subset of the population. However, if consumers gamble to convexify choice sets, then current gamblers are an endogenously selected group: consumers will be more likely to purchase lottery tickets when they have a desire to “convexify” and so these households will respond more strongly to income shocks than a random household receiving the same shock. A particular concern for external validity is that the selection process is directly related to the outcome of interest (probability of a durable purchase or a discrete change in labor supply). This is analogous to randomization bias in randomized trials ([Heckman and Smith, 1995](#)) if patients decide whether to subject themselves to randomization on the basis of their need the treatment, or the anticipated benefit of treatment in their particular circumstances.

To highlight the mechanisms at work and to motivate our empirical strategy, 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. The strength of the 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 gambling to convexify is an empirical question.

To look for evidence that consumers gamble to convexify we use data from the British Household Panel Survey and focus on purchases of consumer durables. Our empirical strategy is effectively a “difference-in-differences” design with household fixed effects, contrasting estimated within-household income effects for lottery windfalls with income effects for other windfalls (specifically inheritances) among households that are credit-constrained and households that are not. We use unconstrained households 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.

Our main result is that, among constrained households, purchases of consumer durable goods are much more responsive to a lottery win than to receipt of other windfall income: among the constrained, the income effect of a lottery win is five times greater than the income effect of a non-lottery windfall of the same size. By contrast, there is no difference in the estimated income effects of different windfalls for unconstrained households. We also show that there is no differential

¹ One natural question is why households do 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 For example, [Tufano \(2008\)](#) and [Kearney et al. \(2011\)](#) have emphasized the entertainment aspect of prize-linked savings products in explaining their potential attraction.

³ In the UK, this includes not only the National Lottery but also premium bonds, a government bond which have a lottery element.

⁴ Related to this is the 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 via a lottery. [Handa and Kirton \(1999\)](#) provide evidence from Jamaica that people use their allocation from the ROSCA to buy durable goods.

⁵ Other papers using lotteries to estimate income and wealth effects include: [Lindal \(2005\)](#), [Gardner and Oswald \(2007\)](#), [Apouey and Clark \(2015\)](#), [Hankins and Hoestra \(2011\)](#), [Kuhn et al. \(2011\)](#) and [Cesarini et al. \(2013\)](#).

effect of different types of windfall for spending on (a limited set of) non-durable items among constrained households. 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 (“players”). For the subset of these individuals who are constrained, purchases of consumer durable goods are more responsive to non-lottery windfall income than purchases by non-players. Our “small winnings test” implies that it is not the source of the money (lottery versus other windfall) that matters, but rather that lottery players are in different economic circumstances than non-players.

These findings highlight the importance of characterizing consumption opportunity sets in understanding consumer choices under uncertainty. They suggest that introducing wealth lotteries in structural models of discrete choice is not just a technical fix, but captures a genuine aspect of consumer behavior. And these results question the external validity of lotteries as an instrument for estimating income effects. We consider this as a specific example of the more general case for using insights from economic theory to shed light on the nature of external validity concerns associated with instrumental variable estimates.

The rest of the paper proceeds 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. Model of gambling to finance indivisible purchases

Our model is a one period model with two stages.⁶ 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 fair⁷: 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 $1 - q$.

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: $u(x_2 - dp, d) = u(x_2 - dp) + \eta d$; 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.⁸

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 $V_2^{d=1}(x_2) \geq V_2^{d=0}(x_2)$, i.e. $u(x_2 - p) + \eta \geq u(x_2)$.

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

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

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 Fig. 1. The degree of non-concavity will depend on the price of the durable good, p , and on the utility value of the durable good, η .

⁶ 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 to convexify transparent.

⁷ 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-concavity.

⁸ The additive separability assumed here is not necessary. It is however necessary to restrict the degree of substitutability between durable and non-durable consumption. We assume expected utility, although an extension to a non-expected utility framework may broaden the regions where nonconvexities occur.




Fig. 1. Durable Purchase Decision.




Fig. 2. Lottery and Durable Purchase Decisions.

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 $E[V_1^{l=1}(x_1)] - V_1^{l=0}(x_1) \geq 0$. Note that:

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

and

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

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

Proof. See Appendix (Supplementary material).

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 Fig. 2. In other words, the maximum wealth range over which lottery purchases will occur is given by the size of the winnings.

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, $x_1 \in [\underline{x}_1, \bar{x}_1]$, which contains x_2^* ($x_1 < x_2^* < \bar{x}_1$), in which the agent will purchase a lottery ticket.

Proof. See Appendix (Supplementary material).

From Result 2, we know that $x_2^* - (1 - q)/q \leq \underline{x}_1 < \bar{x}_1 \leq x_2^* + 1$. Within these bounds, the size of the region $x_1 \in [\underline{x}_1, \bar{x}_1]$ depends on parameter values (η , q and the curvature 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 $\underline{x}_1 < x_1 \leq \bar{x}_1$ in which the agent buys a lottery ticket and then buys a durable if and only if she wins the lottery; and a region $x_1 > \bar{x}_1$ in which the agent does not buy the lottery ticket but does buy the indivisible good. This is shown in Fig. 2.

This simple model illustrates that lottery players are likely to be close to the margin of a discrete decision. The implications of this theory for observable data are discussed in the next section.

3. Empirical implications

The model above assumes that a given person, at a given time, will react in the same way to an income windfall regardless of its source. However, because lotteries are played by particular consumers, at particular times, the average response to a lottery windfall may be very different to the average response to other windfalls. This suggests, first, that we can test the theory that consumers gamble to convexify by comparing responses to different kinds of windfalls. Second, research that uses lotteries to estimate income effects will not estimate a population average income effect. This section develops these implications and then we take them to data in the second half of the paper.

We contrast the endogenously chosen lottery with a random windfall. To be concrete, we imagine that a random fraction (λ) of the population has an elderly, spinster aunt, who with some probability, q , will die in the current period, leaving a windfall. In this thought experiment we hold the fraction of individuals with an elderly aunt equal to the fraction of consumers who chose to play the lottery. The key point of the thought experiment is that, while individuals chose to buy a lottery ticket, they do not choose to have an aunt. We think of the incidence and survival of aunts as random, so that the income effect from inheritances then approximates the population average income effect.

In taking these implications to data, there are two sets of issues to consider. First, there is the structure of the data and in particular whether or not we can identify potential recipients (lottery ticket purchasers, or consumers with an aunt) in the data. The second set of issues revolves around whether observed inheritances might differ from a random windfall, and how we might deal with this in our empirical strategy. We take up the data structure issue first.

3.1. Different data structures

We consider two cases, corresponding to two different data structures. In the first case, which resembles most of the empirical studies using lottery windfalls, income effects are estimated by comparing recipients (lottery winners) and potential recipients (i.e. people who play the lottery, but lose). In the second case, which more closely resembles our data, the comparison is between recipients and non-recipients. For lotteries, the latter includes both losers and non-players; for inheritances, the latter includes those without an aunt as well as those whose aunt survives. We show that in both cases the extra spending by lottery winners is a biased estimate of the population average income effect (the income effect arising from a random windfall) because of the desire to convexify among some individuals, some of the time.

3.1.1. Comparing recipients and nonrecipients among potential recipients

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 $P^c(d = 1|\text{recipient}) = 1$ lottery winner purchases the indivisible good is one:

We use the superscript “c” to indicate that playing the lottery was a choice taken by the individual. This, and subsequent probabilities are summarized in Table 1.

In the case of the inheritance, the distribution of aunts is random and so the potential for receipt is random. This is unlike the distribution of lottery tickets which results from a choice. However, among those with aunts, the actual receipt of inheritance is random and this is like the randomness in winning the lottery. We assume that the expected value of an

Table 1
Probabilities of purchase: chosen lotteries versus random inheritances.

	Windfall recipient $P(d = 1 recipient)$	Potential recipient, but no windfall (losing ticket, surviving aunt) $P(d = 1 nonrecipient, potential)$	Non-recipient* (losing ticket or no ticket, surviving aunt or no aunt) $P(d = 1 nonrecipient)$
Chosen Lottery	1	0	$\frac{1-F(\bar{x}_1)}{1-q\lambda}$
Random Inheritance	$1 - F\left(x_2^* - \frac{1-q}{q}\right)$	$1 - F(x_2^* + 1)$	$1 - F^a$
Difference	$F\left(x_2^* - \frac{1-q}{q}\right) \geq 0$	$F(x_2^* + 1) - 1 < 0$	$\frac{\lambda q(1-F)}{1-\lambda q} > 0^a$
Approx Bias: (Diff in diff)		$1 + F(x_2^* - \frac{1-q}{q}) - F(x_2^* + 1) > 0$	$F\left(x_2^* - \frac{1-q}{q}\right) - \lambda q \left(1 - \left(F\left(x_2^*\right) - F\left(x_2^* - \frac{1-q}{q}\right)\right)\right)$ $1-\lambda q$

^a The probabilities of a non-recipient 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-recipient when the lottery is chosen is given by: $(1 - F(\bar{x}_1)) / (1 - q\lambda)$ and the probability of purchase by a non-recipient with a random inheritance is given by: $\frac{(1-q)\lambda(1-F(x_2^*+1)) + (1-\lambda)(1-F(x_2^*))}{1-\lambda q}$. However, since \bar{x}_1 lies between $(x_2^* + 1)$ and x_2^* , evaluating each CDF at the same value of x_2^* is a reasonable approximation.

aunt is 0⁹: when she is alive, there is a per-period cost of 1, analogous to the ticket price of the lottery; a probability q of an inheritance being received, analogous to the probability of winning the lottery; and a windfall payout of an inheritance of 1/q. Net of cost, recipients receive $(1 - q)/q$ giving cash on hand of $x_2 = x_1 + (1 - q)/q$. They will purchase the indivisible good if $x_1 \geq x_2^* - (1 - q)/q$. We use the superscript “R” to denote that the inheritance is a random windfall. Thus

$$P^R(d = 1|recipient) = 1 - F\left(x_2^* - \frac{1-q}{q}\right)$$

This implies recipients (that is, winners) from the chosen lottery are more likely to purchase the indivisible good than recipients of the random inheritance (column 1 of Table 1):

$$P^C(d = 1|recipient) - P^R(d = 1|recipient) = F\left(x_2^* - \frac{1-q}{q}\right) \geq 0 \tag{3}$$

The intuition behind this result is straightforward. In the case of the random inheritance, some recipients will come from below the lower threshold and will not have enough cash on hand to buy the divisible good even if they receive the windfall. This difference tends to zero as q becomes increasingly small: if there is a windfall that is very large but with a very small probability of receipt, it is in the interest of everyone with cash-on-hand below x_2^* to gamble to convexify, and all recipients of a windfall will chose to buy the indivisible good.

For those that chose to play, but lost, the probability that they purchase the indivisible good is zero: $P^C(d = 1|non - recip, potential) = 0$. By comparison, among the those with aunts, and whose aunt survives, there are some consumers with cash on hand above the upper threshold $(x_2^* + 1)$. These consumers will have enough cash on hand to purchase the divisible good even though they do not receive an inheritance, i.e. $P^R(d = 1|non - recip, potential) = 1 - F(x_2^* + 1)$. The probability of purchase is therefore lower among losers of the chosen lottery (column 2 of Table 1):

$$P^C(d = 1|non - recip, potential) - P^R(d = 1|non - recip, potential) = F(x_2^* + 1) - 1 \leq 0 \tag{4}$$

Putting together the differences in purchase probabilities between recipients and the difference in purchase probabilities between potential recipients who do not receive, it is clear that the income effect in the case of the chosen lottery suffers from upward bias compared to income effect from the random inheritance. 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 column of Table 1. The size of the bias becomes smaller as the range in which tickets are bought becomes larger.

3.1.2. Comparing recipients to all non-recipients

Non-recipients comprise, in the case of the lottery, both non-players and losers, and in the case of the random inheritance, both those without an aunt and those whose aunt survives. Starting with the chosen lottery, those who choose not to play are those with $x_1 < \underline{x}_1$ or $x_1 > \bar{x}_1$ while losers are the fraction $1 - q$ of lottery players, all of whom have $\underline{x}_1 < x_1 \leq \bar{x}_1$. Of

⁹ This is just an innocuous normalization to aid the comparison with the lottery.

these non-recipients, only agents with cash on hand $x_1 > \bar{x}_1$ buy the indivisible good. Recall that λ is the fraction of lottery players: $\lambda = F(\bar{x}_1) - F(x_1)$, where $F(\cdot)$ is the cumulative distribution of cash on hand (x_1) in the population. Thus

$$P^C(d = 1 | non - recip) = \frac{1 - F(\bar{x}_1)}{1 - q\lambda} \quad (5)$$

The effect of winning the lottery (relative to non-recipients) on the probability of indivisible good purchase is therefore the difference (row 1, columns 1 and 3 of Table 1):

$$P^C(d = 1 | recip) - P^C(d = 1 | non - recip) = 1 - \frac{1 - F(\bar{x}_1)}{1 - q\lambda} = \frac{F(\bar{x}_1) - q\lambda}{1 - q\lambda} \quad (6)$$

With random inheritance those with an aunt that survives are fraction $(1 - q)\lambda$ of the population, and they purchase the durable if $x_1 \geq x_2^* + 1$. Those without an aunt are the fraction $(1 - \lambda)$ of the population, and they purchase the durable if $x_1 \geq x_2^*$. The overall fraction of the population that are non-recipients is, as with the lottery, $(1 - q)\lambda + (1 - \lambda) = 1 - q\lambda$. Thus the fraction of non-recipients who purchase the durable is:

$$\begin{aligned} P^R(d = 1 | non - recip) &= \frac{P^R(d = 1, non - recip)}{P^R(non - recip)} \\ &= \frac{(1 - q)\lambda(1 - F(x_2^* + 1)) + (1 - \lambda)(1 - F(x_2^*))}{1 - q\lambda} \end{aligned} \quad (7)$$

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

$$P^R(d = 1 | non - recip) \approx \frac{(1 - F(x_2^*))(1 - q\lambda)}{1 - q\lambda} = (1 - F(x_2^*)) \quad (8)$$

The denominator is the fraction of the population who are not recipients. The first part of the numerator, $1 - F(x_2^*)$, is the fraction of all individuals whose cash-on-hand means they would purchase the durable regardless of receipt. Some of these individuals will receive a random inheritance and so this fraction is multiplied by the fraction of the population that are not recipients. Given the approximation that the cost of an aunt is small, the probability of purchase for non-recipients of a random windfall is independent of the fraction of the population with an aunt.

By contrast, with a chosen lottery, none of those individuals who would purchase regardless of the lottery outcome actually choose to buy lottery tickets. This means these “always purchasers” are all non-recipients, and the probability of purchasing the durable among non-recipients is given by Eq. (5).

To aid interpretation of the difference between Eqs. (5) and (8), 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-recipients from the chosen lottery and random inheritance of (column 3 of Table 1):

$$P^C(d = 1 | non - recip) - P^R(d = 1 | non - recip) \approx \frac{\lambda q (1 - F(x_2^*))}{1 - q\lambda} \geq 0 \quad (9)$$

Note that λ enters this difference only through P^C and not through P^R . It is only the fraction of individuals who choose the lottery that matters and so we do not require the assumption that the fraction of aunts is equal to the fraction of lottery players.

The probability of purchasing the durable among the non-recipients from the chosen lottery is higher. This arises from a subtle composition effect because the group of non-recipients 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-recipients by choice (i.e. chose not to have a ticket because they would have purchased the durable anyway) do receive a windfall in the case of random inheritances. This then reduces the number of purchasers of the durable among those who were non-recipients. As q increases, this upward bias in the purchase rate of non-recipients (relative to a random inheritance) gets smaller. By contrast, we showed above that the upward bias in the purchase rate of winners (relative to a random inheritance) is increasing in q . This means that the net bias in income effects gets larger as q increases.

This can be illustrated by calculating numerically the size of the bias among recipients and among non-recipients in our simple model, at particular parameter values. We assume a log-normal distribution for cash-on-hand, log utility for consumption, and consider a high and low value for the utility of the durable (η). Fig. 3 shows the difference in the probability of purchase between the chosen lottery and random inheritance. 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 .

This discussion has highlighted the differing income effects that arise 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

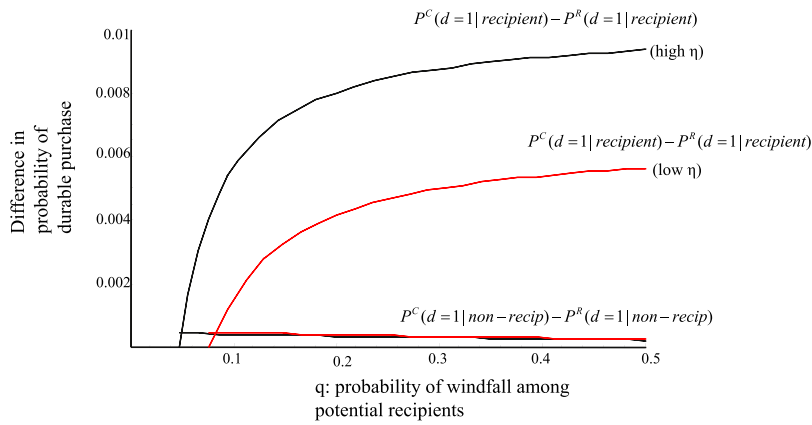


Fig. 3. Chosen Lottery vs Random Inheritance.

chosen to participate in because of gambling 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 windfalls (Bailey and Olson Wonnacott, 1980; Hartley and Farrell, 2002). The need to gamble to convexify is also diminished if there are multiple indivisible goods of different sizes so that the indivisibility is less “lumpy”, or if there are uninsurable income shocks which provide some convexification. This discussion highlights the large number of factors that affect the convexification decision and that would need to be specified for a realistically calibrated model. Instead of following this approach, we look directly for evidence of convexification in data on household choices. We now take up this empirical approach in greater detail, including how we deal with ways in which an inheritance may differ from random windfalls.

3.2. Empirical framework

We adopt a reduced form empirical approach directly motivated by our model. Our main empirical strategy is to estimate a difference-in-differences (DiD) in income effects. In particular, we compare the effect of lottery winnings on purchases of indivisible goods¹⁰ with the effect of inheritances – a different kind of windfall – and we compare these differences in income effects between households who are likely to be credit-constrained and those who are not. The latter is because we only expect a demand for convexification among the constrained. Unconstrained lottery players must be doing so for reasons outside our model (entertainment, for example) and so we would not expect them to be grouped below the threshold of a discrete purchase.

We estimate an empirical model along the following lines:

$$d_{it} = (\beta_1 + \beta_2 C_{it}) Lot_{it} + (\beta_3 + \beta_4 C_{it}) Inh_{it} + \alpha' X_{it} + u_{it} \tag{10}$$

where d_{it} is a measure of durable purchases by household i at time t ; $C_{it} = 1$ if the agent is 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, presence of constraints, employment status, financial expectations and year dummies. The error term, u_{it} , consists of a household-specific fixed effect and a random noise term, i.e. $u_{it} = \phi_i + \varepsilon_{it}$.

In the context of the model above, inheritances are intended to approximate a random windfall. 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 (including age) and individual fixed effects. Note that the critical value will vary in the population and over time for a given individual according to tastes and needs.

Previous empirical literature has shown that durables respond to unexpected windfalls (see Keeler and Abdel-Ghany, 1985), so we would expect $\beta_1 = \beta_3 \geq 0$. The theoretical considerations developed in the previous section suggest that, among constrained households, selection into playing the lottery will lead to differential responses to a lottery win compared to other windfalls. Under the convexification hypothesis, we expect durable purchases to respond more strongly to a lottery win than to an inheritance among constrained households, i.e. $(\beta_1 + \beta_2) > (\beta_3 + \beta_4)$, and hence $\beta_2 > \beta_4 > 0$.

To claim that consumers are gambling to convexify, we need to rule out alternative interpretations and we deal with this in a number of ways. First, we include household fixed effects to remove level differences: the time-invariant unob-

¹⁰ 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.

servable characteristics, such as risk or time preference that affect both lottery purchases and durable consumption. These characteristics include any permanent propensity or preference for durables that differs between inheritors and winners. The inclusion of household fixed effects means that we are comparing changes in durable purchases, and not levels, across subjects. Only a small fraction of our sample experienced both a lottery win and an inheritance and so identification is largely across rather than within subjects.¹¹

Second, our difference-in-differences strategy, comparing income effects across lottery wins and inheritances across constrained and unconstrained households, controls for general differences in income effects that affect both constrained and unconstrained households. As noted above, 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.

Differences between the two types of windfall may include the possibility that inheritances are anticipated, as discussed by Hurst and Lusardi (2004). As well as the DiD strategy, we present additional evidence showing that household financial expectations and consumption do not adjust in anticipation of an inheritance. This suggests that while individuals may know they have an elderly aunt, they do not know the timing and the amount of any inheritance.

An alternative explanation of different income effects is that the source of the money may affect what people feel that they can spend the money on. This idea was termed “emotional accounting” by Levav and McGraw (2009) and nicely summarized 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.”

We implement an additional empirical test (which we call the “small winnings test”), the basis of which is the following: among the people who received an inheritance there are likely to be some who were gambling to convexify, but who lost the (endogenously-chosen) 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 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 money, and thus the explanation must be a selection story like the convexification hypothesis.

The convexification hypothesis identifies a potential 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 constrained. It is still possible (if improbable) that there is an alternative, time-varying selection mechanism that operates only on gamblers who are constrained. We cannot conclusively eliminate this possibility, but we present additional strong evidence against there being such a selection mechanism in the form of a falsification test involving non-durable consumption.

4. Data

Our main analysis uses data taken from the British Household Panel Survey (BHPS) from 1997 to 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 5500 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.¹² We select single and two-adult households where the head is aged 20–70. Our analysis sample contains information on 6147 households (29,859 observations).

4.1. Consumer durables

We focus on durables that 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, tumble driers, 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.

¹¹ In the whole sample, 8% of households report receiving an inheritance and a lottery win, but we restrict our empirical work to a sample of inheritances and lottery wins within the range £100–£5000 where the fraction falls to 2%.

¹² 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.

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 (rent-to-own) companies also require a credit check and may charge high interest rates if the repayments are made over a long period. We think it is plausible that, compared to these alternatives, buying a lottery ticket may not be an unattractive option.¹³

4.2. Credit constraints

The BHPS does not have a question that asks directly about access to credit; we define constrained households as those with no (income from) savings or investments. This is a broad definition by which around half of all household-year observations are defined as constrained.¹⁴ Note that, with this broad definition, our estimates are likely to under-estimate the true convexity effect, compared to an approach where we could identify exactly which households face credit constraints. We also show results additionally excluding anyone with household income in the top two-third of the distribution. However, recent evidence from Kaplan et al. (2014) indicates that even high income households with illiquid but not liquid assets may face a hand-to-mouth existence.

4.3. 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).¹⁵

In the sample as a whole, 21 per cent of households reported at least one Lottery winning, while 5 per cent reported an inheritance. However, the average amounts received in the two cases are very different. The mean (median) Lottery winning was £290 (£40) compared to £29,949 (£5000) in the case of inheritances. This is not surprising given the structure of National Lottery payouts.¹⁶ 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 £5000. Anyone who receives a windfall of more than £5000 in any wave is dropped from the analysis and in our initial analysis we ignore small (<£100) lottery wins and inheritances. In this range, 13% of households report ever receiving a lottery win, 8% report ever receiving an inheritance and 2% report receiving both. 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 – £5000, 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 £1001–£5000.

Table 2 provides summary statistics for constrained and unconstrained households. Many unconstrained households receive windfalls from lottery wins, and indeed a higher proportion than among those who are constrained. In fact, the existence of lottery winners who are not constrained is necessary for the DiD strategy described above. 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. Fig. 4 shows that budget shares on gambling decline markedly with total expenditure, consistent with the need to gamble to convexify being concentrated

¹³ 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).

¹⁴ Young and Waldron (2006) 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. This information is not available in the BHPS.

¹⁵ 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.

¹⁶ The odds of winning £10 are 1:57, compared with odds of 1:1031 to win around £100, 1:55,490 to win around £1000, 1:2,330,636 to win around £100,000 and 1:13,983,817 to hit the jackpot.

Table 2
Descriptive statistics.

	Unconstrained					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	£2959	£3142	£2948	£3200	£2952	£2063	£2559	£2046	£2378	£2059
Degree (0/1)	0.217	0.131	0.223	0.220	0.217	0.106	0.066	0.108	0.174	0.106
Kids (0/1)	0.552	0.476	0.558	0.532	0.553	0.782	0.653	0.786	0.721	0.783
Couple (0/1)	0.762	0.859	0.756	0.821	0.760	0.599	0.738	0.594	0.705	0.598
Mean windfall		£514		£1984			£595		£1875	
Median windfall		£224		£1500			£250		£1200	
N	13,757	802	12,955	363	13,394	16,129	497	15,632	190	15,939

Notes: Number of durables refers to purchases made over the past 12 months of televisions, fridge/freezers, washing machines, tumblers/driers, dishwashers and home computers). Age, degree are for head of household. Income is household net monthly income. Winners and inheritors refer to those who receive lottery winnings and inheritances in the range £100–£5000. Constrained refers to no income from savings/dividends.




Fig. 4. Household Gambling Expenditure, 2007 EFS.

Budget Shares and% with Positive Expenditure.

among low income groups. Fig. 4 also shows that the fraction of households with positive gambling expenditure is around 40% across a wide range of incomes, again consistent with the idea of there being more than one motive for gambling.

Returning to Table 2, within the range we focus on (£100–£5000) there is no statistically significant difference in average windfall size between those who are potentially 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 constrained who have had a windfall due to a lottery win (but not for the 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. We model the number of durables purchased during the previous twelve months as the dependent variable. Below we show that the results are very similar when the dependent variable is a binary indicator for whether or not the household purchased any durables. We include lottery wins and inheritances in amounts (in £’00s). Below, we show results for a binary indicator for whether or not the household inherited/received lottery winnings.

Columns (1) and (2) of Table 3 are estimated using OLS. The results in column (1) indicate a stronger propensity to consume durables out of lottery winnings than out of an inheritance. In Column (2) we interact the windfall variables with a dummy variable indicating whether the household is constrained. This corresponds to Eq. (11) above and implements our main DiD test. The results in column (2) show that the stronger response to lottery winnings than to inheritances is driven just by those who are constrained, in line with our model. Columns (3) and (4) include household fixed effects to

Table 3

Main regression results. Dependent variable = Number of durables purchased in the last 12 months.

	(1) OLS	(2) OLS	(3) FE	(4) FE
β_1 Lot(£'00)	0.012** (0.003)	0.004 (0.003)	0.010** (0.004)	0.003 (0.004)
β_2 Lot(£'00)*C		0.017** (0.007)		0.016** (0.007)
β_3 Inh(£'00)	0.003** (0.001)	0.003* (0.002)	0.004** (0.002)	0.004* (0.002)
β_4 Inh(£'00)*C		0.000 (0.003)		−0.000 (0.004)
N	29,859	29,859	29,859	29,859
R ²	0.046	0.046	0.008	0.009
$\beta_1 = \beta_3$ [p-value]	[0.012]	[0.632]	[0.135]	[0.814]
$(\beta_1 + \beta_2) = (\beta_3 + \beta_4)$		[0.007]		[0.036]

Notes: Robust standard errors in brackets, clustered at the household level (6147 households). ** denotes statistically significant at the 5% level; * at the 10% level. Lot and Inh refer to lottery winnings and inheritances in the range £100–£5000. C = constrained = no income from savings/dividends. 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; financial expectations for next year, constrained; year dummies.

control for time-invariant unobservable characteristics, including time- and risk preferences that may affect both durable purchases and lottery participation. Column (4) presents the DiD test including household fixed effects. This allows for fixed effects in estimating income effects and double-differences income effects (across the constrained and unconstrained, and across inheritors and lottery winners), ruling out any alternative selection mechanism which operates on time-invariant unobservables and is not limited to the constrained. Column (4) is our preferred specification.

We find no evidence for any general “lottery winnings effect” – there is no significant difference in the response to lottery and inheritances among unconstrained households. However, among the constrained, the marginal propensity to consume durables out of (endogenously-selected) lottery winnings is nearly five times stronger than that out of an (exogenously determined) inheritance. To give some indication of how big these responses are, consider a typical medium-sized lottery win or inheritance of £500. This would result in a 0.095 increase in the number of durables purchased within the year among constrained households, which is an increase of 19% over the baseline purchase rate of those who are neither winners nor inheritors (0.49 from Table 2). The corresponding numbers for a £500 inheritance are a 0.020 increase in the number, which is a 4% increase over the baseline purchase rate. While this focuses on the constrained, they comprise nearly half of our sample. These numbers suggest that using lottery wins as an instrument is likely to do a poor job in estimating population average income effects.

5.2. Alternative specifications

Table 4 summarizes the results from a number of alternative specifications. To facilitate comparison, we include the results from our preferred specification (Table 3, column 4) in the first column of Table 4.

First, we impose common support on our sample. A possible limitation of regression adjustment is that, except in the special case of discrete independent variables and a fully-saturated model, it allows estimation of counterfactuals for treated units for whom there are no similar control units. To address this, we estimate propensity score models for the treatment group (constrained, lottery winners) versus each control group and impose common support in the probability, given characteristics, of being a constrained lottery winner (a propensity score). Given the similarity in characteristics among the groups (Table 2), imposing common support results in dropping relatively few observations and the estimates, shown in Column (2) of Table 4 are very similar.

Column (3) of Table 4 confirms that the results are also similar if we adopt a binary dependent variable and estimate the probability of durable purchase rather than the number of durables purchased.

There may be a concern that the relationship may be mis-specified since our windfall variables include a large number of zeroes. If we include lottery wins and inheritances as binary indicators (and include household fixed effects), we find similar estimated responses to lottery winnings and inheritances among constrained households (Column (5)). However, one issue with this specification is that the typical lottery win is much smaller than the typical inheritance making the effects hard to compare directly. Column (6) presents a further specification that includes both a binary indicator and the amount of the windfall and *p*-values for the test that lotteries and inheritances have the same effect on durables purchased using the mean of the two types of windfalls (£550 and £2000 respectively). The finding is the same – we find a stronger effect of lottery winnings than inheritances on durable purchases, but only among the constrained.

Finally, column (7) includes a tighter definition of constrained, including only those in the bottom third of the income distribution. The broader measure may well understate the importance of the bias induced by self-selection into lottery playing if the broader measure is treating some unconstrained individuals as constrained. Comparing column (7) with column (1), the difference between constrained lottery players and constrained inheritors is greater with the narrower definition of a constraint.

Table 4
Alternative specifications.

	(1) FE Main Results	(2) FE common support	(3) FE Discrete dependent variable	(4) OLS Discrete windfall variable	(5) FE Discrete windfall variable	(6) FE Discrete + continuous windfall	(7) FE Tighter definition of constrained
β_1 Lot(£'00)	0.003 (0.004)	0.003 (0.004)	0.003 (0.002)			0.004 (0.004)	0.007* (0.004)
β_2 Lot(£'00)*C	0.016** (0.007)	0.016** (0.008)	0.005* (0.003)			0.014 (0.009)	0.012 (0.008)
β_3 Inh(£'00)	0.004* (0.002)	0.004* (0.002)	0.004** (0.001)			0.001 (0.003)	0.005** (0.002)
β_4 Inh(£'00)*C	−0.000 (0.004)	−0.002 (0.004)	−0.002 (0.002)			0.009 (0.006)	−0.006 (0.005)
β_5 Lot(0/1)				0.031 (0.029)	0.008 (0.040)	−0.008 (0.047)	
β_6 Lot(0/1)*C				0.183** (0.055)	0.119* (0.066)	0.032 (0.081)	
β_7 Inh(0/1)				0.090** (0.044)	0.118** (0.055)	0.094 (0.087)	
β_8 Inh(0/1)*C				−0.062 (0.076)	−0.108 (0.103)	−0.273* (0.159)	
N	29,859	29,281	29859	29,859	29,859	29,859	29859
R ²	0.009	0.009	0.008	0.046	0.008	0.009	0.009
$\beta_1 = \beta_3$ [p-value]	[0.814]	[0.770]	[0.761]				[0.607]
$\beta_5 = \beta_7$				[0.258]	[0.106]		
$(\beta_1 + \beta_2) = (\beta_3 + \beta_4)$	[0.036]	[0.024]	[0.035]				[0.014]
$(\beta_5 + \beta_6) = (\beta_7 + \beta_8)$				[0.014]	[0.242]		
$(\beta_5 + \bar{x}\beta_1) = (\beta_7 + \bar{x}\beta_3)$, $\bar{x} = \text{£}550$						[0.293]	
$(\beta_5 + \bar{x}\beta_1) = (\beta_7 + \bar{x}\beta_3)$, $\bar{x} = \text{£}2000$						[0.553]	
$((\beta_5 + \bar{x}\beta_1) + (\beta_6 + \bar{x}\beta_2)) = ((\beta_7 + \bar{x}\beta_3) + (\beta_8 + \bar{x}\beta_4))$, $\bar{x} = \text{£}550$						[0.047]	
$((\beta_5 + \bar{x}\beta_1) + (\beta_6 + \bar{x}\beta_2)) = ((\beta_7 + \bar{x}\beta_3) + (\beta_8 + \bar{x}\beta_4))$, $\bar{x} = \text{£}2000$						[0.018]	

Notes: Robust standard errors in brackets, clustered at the household level. ** denotes statistically significant at the 5% level; * at the 10% level. Lot and Inh refer to lottery winnings and inheritances in the range £100–£5000. C = constrained = no income from savings/dividends, except specification (7) where C = constrained = no income from savings/dividends and in the bottom third of the income distribution. 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; financial expectations for next year, constrained; year dummies. In specification (3) the dependent variable is a discrete (0/1) measure of whether a durable was purchased during the previous 12 months. The tests for specification (6) are evaluated at the mean of both lottery winnings (£550) and inheritances (£2000).

Table 5

Are windfalls expected?

Fixed effects regression results.

Dependent variable: (0/1) whether household head expects financial situation to improve over the next 12 months

	Whole sample	Constrained	Unconstrained
Lot (0/1) _{t+1}	0.008 (0.016)	−0.009(0.028)	0.012 (0.020)
Inh (0/1) _{t+1}	0.018 (0.028)	0.034 (0.049)	−0.019(0.037)
Other (0/1) _{t+1}	0.035**(0.014)	0.014 (0.022)	0.058**(0.019)
N	27,410	14,508	12,884
R ²	0.001	0.000	0.001

Notes: Robust standard errors in brackets, clustered at the household level. ** denotes statistically significant at the 5% level; * at the 10% level. Lot and Inh refer to lottery winnings and inheritances in the range £100–£5000. “Other” refers to other windfalls and includes life insurance policy payments, pension lump-sums, redundancy payments, personal accident claims and “anything else”. Constrained = no income from savings/dividends.

Table 6

Small winnings test.

	Main results	Further test
β_1 Lot(£'00)	0.003 (0.004)	0.003 (0.003)
β_2 Lot(£'00)*C	0.016**(0.007)	0.016**(0.006)
β_3 Inh(£'00)	0.004* (0.002)	0.004* (0.002)
β_4 Inh(£'00)*C	−0.000(0.004)	−0.004(0.004)
δ_1 SmLot(0/1)*Inh(£'00)		−0.000(0.005)
δ_2 SmLot(0/1)*Inh(£'00)*C		0.019 (0.013)
$\beta_1 = \beta_3$ [p-value]	[0.814]	[0.780]
$(\beta_1 + \beta_2) = (\beta_3 + \beta_4)$	[0.036]	[0.008]
$(\beta_1 + \beta_2) = ((\beta_3 + \beta_4) + (\delta_1 + \delta_2))$		[.898]
N	29,859	29,859
R ²	0.009	0.009

Notes: Robust standard errors in brackets, clustered at the household level (6147 households). ** denotes statistically significant at the 5% level; * at the 10% level. Lot and Inh refer to lottery winnings and inheritances in the range £100–£5000. C = constrained = no income from savings/dividends; SmLot is an indicator if the household receives a lottery win of less than £100. Regressions include full set of controls as in Table 3.

5.3. Are inheritances anticipated?

As noted in section 3, one potential concern is that inheritances may differ from lottery wins in being reasonably well anticipated by the individual. 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–£5000). Only the coefficient on other windfalls is positive and significant; our financial expectation data do not contain any evidence that medium inheritances are anticipated. Consistent with this, sensitivity analysis (details available on request) that included lead terms in the durables regression to pick up the effect of any anticipated windfalls found no significant anticipation effects.

5.4. The small winnings test

We also perform what we call the “small winnings” test, by estimating the following empirical model:

$$d_{it} = (\beta_1 + \beta_2 C_{it}) Lot_{it} + (\beta_3 + \beta_4 C_{it}) Inh_{it} + (\delta_1 + \delta_2 C_{it}) SmLot_{it} \times Inh_{it} + \alpha' X_{it} + u_{it} \quad (11)$$

where $SmLot_{it} = 1$ if someone receives a lottery win of less than £100, and equals 0 otherwise. Our hypothesis is that, among 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 but rather will have the larger income responses of those who receive a medium-sized lottery win (i.e. $(\beta_1 + \beta_2) = (\beta_3 + \beta_4) + (\delta_1 + \delta_2)$)

The results in Table 6 show that this is exactly what we find in our data. Column (1) reproduces (from Column (4) of Table 3) the results from our main DiD specification with household fixed effects. In Column (2) we report estimates of Eq. (11) in which we interact the inheritance variables with a dummy indicating a small lottery win. We find that 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 responses of 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. Constrained gamblers have larger responses and this is consistent with the idea that they are a selected group: close to a purchase margin.

Table 7
Falsification tests.

	Number of durables		Food at home (£)	Meals out (£)
β_1 Lot (£'00)	0.003	(0.004)	-0.045 (0.102)	0.846**(0.307)
β_2 Lot (£'00)*C	0.016**	(0.0052)	0.195 (0.167)	0.349 (0.457)
β_3 Inh (£'00)	0.004*	(0.002)	0.047 (0.061)	0.169(0.1136)
β_4 Inh (£'00)*C	-0.000	(0.004)	-0.009 (0.083)	0.381*(0.217)
$\beta_1 = \beta_3$ [p-value]	[0.814]		[0.440]	[0.045]
$(\beta_1 + \beta_2) = (\beta_3 + \beta_4)$	[0.036]		[0.453]	[0.082]
N	29,859		28,859	29,859
R ²	0.009		0.109	0.072

Notes: Standard errors in brackets, clustered at the household level (6147 households). ** denotes statistically significant at the 5% level; * at the 10% level. Lot and Inh refer to lottery winnings and inheritances in the range £100–£5000. C = constrained = no income from savings/dividends. Other controls as in Table 3.

5.5. Falsification tests

Finally, in Table 7 we present the results from running our main specification but with measures of non-durable spending. The BHPS contains only a small number of these measures – we include weekly household spending on food for home consumption or, separately, food out (in restaurants) on the left-hand side.¹⁷ Since these are both divisible goods, these results provide a falsification test of the convexification hypothesis.

We find zero income effects for both lottery wins and inheritance receipts when we examine spending on food for home consumption. For meals out, we find a difference in the propensity to spend out of lottery winnings and inheritances. People are more likely to spend money on a meal out when they win on the lottery than when they inherit, consistent with an emotional accounting story. Crucially, however, this difference is common to both constrained and unconstrained households – both types react to a moderate win on the lottery by celebrating with a meal out. On this evidence, our finding of a differential response to lottery wins and inheritances for the constrained and not for the unconstrained is true only for durable purchases, consistent with our model of gambling to convexify.

6. Conclusion

Following an idea first proposed by Ng (1965) this paper shows that consumers are more likely to gamble when faced with a discrete decision and we illustrated how using windfalls from endogenously chosen lotteries could give rise to biased estimates of population income effects. The key point is that the group of lottery players is determined by a time-varying selection mechanism that is directly related to the outcome of interest.

We have presented convincing empirical support for the hypothesis that consumers sometimes gamble to convexify their choice sets. The purchase of durables responds more strongly to a lottery win than to another windfall among constrained households. Our empirical strategy – difference-in-differences with household fixed effects – rules out any alternative explanation for this finding that involves time-invariant unobservable characteristics of lottery players and/or that applies to all lottery players (constrained and un-constrained). It is hard to think of another selection mechanism that can explain this result. Our small winnings test and falsification test using items of non-durable spending provide further support for our preferred explanation.

Our findings are important for a number of reasons. They provide at least a part of the explanation for gambling among low-income households, and also for the popularity of prize-linked savings products amongst these households. Our finding complements the recent discussion by Mullainathan and Shafir (2009) that lotteries may play a role in the household finances of low-income households. 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 financial constraints they face.

Our findings demonstrate that assuming that individuals facing non-convex choice sets play wealth lotteries – as is often done in structural models with discrete choices – is not just a technical convenience. This modeling strategy captures an important aspect of how real individuals behave when faced with such non-convexities.

Our findings also highlight issues with using lottery winnings to instrument for income. The random success of winning a gamble would seem to make it a natural instrument for unanticipated income changes and has motivated the widespread use of lotteries in identifying income effects. However giving consideration to theoretical reasons for why people gamble is crucial for understanding exactly what is being estimated in this case. Our findings indicate that the degree of over-estimation is likely to be sizeable. Among constrained households, using lottery winnings leads to estimated income effects that are five times bigger than using other windfall income. Since the definition of constrained consists of half of all households in our sample, this suggests that using lottery wins as an instrument is likely to do a poor job in estimating population average income effects.

¹⁷ In the BHPS, the food data are banded and we take the mid-points.

We have defined constrained as those with no income from savings or investment. It is important to note that this might be a marker for perceived illiquidity rather than actual illiquidity: individuals with limited financial literacy may believe themselves to be constrained and so resort to inefficient financial instruments like lotteries. The policy implications of actual versus perceived illiquidity may differ: the former is an argument for making more instruments available, while the latter is about public awareness of existing instruments.

Playing lotteries may be part of a larger set of strategies where risks are taken to overcome thresholds. The presence of payday lenders, pawn shops and evidence of weak financial institutions are often found in areas of high crime and highly variable outcomes.

More generally, gambling data has been used to identify consumer preferences, beliefs on probabilities and wealth elasticities. A key example in the literature is the attempt to identify whether the under-purchase of short-odd gambles is due to risk loving preferences or due to probability misperception.¹⁸ Our analysis suggests that gambling is induced by a rational response to features of some individuals' consumption opportunity sets. Analyses that ignore these features of the consumption opportunity set (such as non-convexities) will mischaracterize preferences for risk and evidence of probability misperception. In some circumstances, gambling by fully-informed, risk-averse individuals is rational behavior borne out of necessity.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jebo.2016.07.023>.

References

- Apouey, B., Clark, A.E., 2015. Winning big but feeling no better? The effect of lottery prizes on physical and mental health. *Health Econ.* 24 (5), 516–538.
- Bailey, M.J., Olson Wonnacott, M.P., 1980. The marginal utility of income does not increase: borrowing, lending, and friedman-Savage gambles. *Am. Econ. Rev.* 70 (June (3)), 372–379.
- Besley, T., Coate, S., Loury, G., 1993. The economics of rotating savings and credit associations. *Am. Econ. Rev.* 83 (September (4)), 792–810.
- Cesarini, D., Lindqvist, M., Notowidigdo, M., Ostling, R., 2013. The Effect of Wealth on Household Labor Supply: Evidence from Swedish Lotteries. http://emlab.berkeley.edu/users/cle/e250a_f13/CLNO.pdf.
- Chetty, R., Szeidel, A., 2007. Consumption commitments and risk preferences. *Quart. J. Econ.* 122 (2), 831–877.
- Collard, S., Kempson, E., 2005. *Affordable Credit: The Way Forward*. Policy Press, Bristol.
- Epley, N., Gneezy, A., 2007. The framing of financial windfalls and implications for public policy. *J. Socio Econ.* 36, 36–47.
- Gambling Commission, 2014. *Gambling Participation: Activities and Mode of Access*. <http://www.gamblingcommission.gov.uk/pdf/Survey%20data%20on%20gambling%20participation%20year%20to%20March%202014.pdf>.
- Gardner, J., Oswald, A., 2007. Money and mental well-being: a longitudinal study of mediumsized lottery wins. *J. Health Econ.* 26 (1), 49–60.
- Ghandi, A., Serrano-Padial, R., 2014. Does belief heterogeneity explain asset prices: the case of the longshot bias. *Rev. Econ. Studies*.
- Handa, S., Kirton, C., 1999. The economics of rotating savings and credit associations: evidence from the Jamaican 'Partner'. *J. Dev. Econ.* 60, 173–194.
- Hankins, S., Hoestra, J., 2011. Lucky in life, unlucky in love? The effect of random income shocks on marriage and divorce. *J. Hum. Resour.* 46 (2), 403–426.
- Hansen, G., 1985. Indivisible labor and the business cycle. *J. Monetary Econ.* 16, 309–328.
- Hartley, R., Farrell, L., 2002. Can expected utility theory explain gambling? *Am. Econ. Rev.* 92 (3), 613–624.
- Heckman, J.J., Smith, J.A., 1995. Assessing the case for social experiments. *J. Econ. Perspect.* 9 (2), 85–110.
- Hurst, E., Lusardi, A., 2004. Liquidity constraints, household wealth and entrepreneurship. *J. Political Econ.* 112, 318–347.
- Imbens, G., Rubin, D.B., Sacerdote, B.I., 2001. Estimating the effect of unearned income on labor earnings, savings, and consumption: evidence from a survey of lottery players. *Am. Econ. Rev.* 91 (4), 778–794.
- Jappelli, T., 1990. Who is credit-constrained in the U.S. economy. *Quart. J. Econ.* 105 (1), 219–262.
- Jullien, B., Salanie, B., 2005. Empirical Evidence on the Preferences of Racetrack Bettors in: Efficiency of Sports and Lottery Markets. In Hausch, D., Ziemba, W. (Eds.) *Handbook in Finance*.
- Kaplan, Greg, Gianluca, Violante, Justin, Weidner, 2014. *The Wealthy Hand-To-Mouth*. Brookings Papers On Economic Activity.
- Kearney, M.S., Tufano, P., Guryan, J., Hurst, E., 2011. Making savers winners: an overview of prize-lined savings products. In: Mitchell, Olivia (Ed.), *Financial Literacy: Implications for Retirement Security and the Financial Marketplace*. Oxford University Press.
- Keeler, J., Abdel-Ghany, M., 1985. The relative size of windfall income and the permanent income hypothesis. *J. Bus. Econ. Stat.* 3 (3), 209–215.
- Kuhn, Peter, Kooreman, Peter, Soetevent, Adriaan, Kapteyn, Arie, 2011. The effects of lottery prizes on winners and their neighbors: evidence from the dutch postcode lottery. *Am. Econ. Rev.* 101 (5), 2226–2247.
- Landsberger, M., 1966. Windfall income and consumption: comment. *Am. Econ. Rev.* 56, 534–540.
- Lentz, R., Traaen, T., 2005. Job search and savings: wealth effects and duration dependence. *J. Labor Econ.* 23 (3), 467–490.
- Levav, J., McGraw, P., 2009. Emotional accounting: feelings about money and consumer choice. *J. Market. Res.* 46, 66–80.
- Lindal, M., 2005. Estimating the effect of income on health and mortality using lottery prizes as an exogenous source of variation in income. *J. Hum. Resour.* 40 (1 (Winter, 2005)), 144–168.
- Mullainathan, S., Shafir, E., 2009. Savings policy and decision-making in low-income households. In: Michael Barr, Rebecca Blank (Eds.), *Insufficient Funds: Savings, Assets, Credit and Banking Among Low-Income Households*. Russell Sage Foundation Press, pp. 121–145.
- Ng, Y., 1965. Why do people buy lottery tickets? choices involving risk and the indivisibility of expenditure. *J. Political Econ.* 73 (5), 530–535.
- Rogerson, R., 1988. Indivisible labor, lotteries and equilibrium. *J. Monetary Econ.* 21 (1), 3–16.
- Snowberg, E., Wolfers, J., 2010. Explaining the favorite-long shot bias: is it risk-love or misperceptions? *J. Political Econ.* 118 (August (4)), 723–746.
- Tufano, P., 2008. Savings whilst gambling: an empirical analysis of U.K. premium bonds. *Am. Econ. Econ. Rev. P&P* 98 (2), 321–326.
- Young, G., Waldron, M., 2006. The state of British household finances: results from the 2006 NMG Research Survey. *Bank Engl. Quart. Bull.*, 397–403.

¹⁸ e.g. Jullien and Salanie, (2005), Snowberg and Wolfers (2010), Ghandi and Serrano-Padial (2014).