

The Effects of Lottery Prizes on Winners and their Neighbors: Evidence from the Dutch Postcode Lottery

Peter Kuhn*

University of California at Santa Barbara

Peter Kooreman
Tilburg University

Adriaan R. Soetevent
University of Amsterdam (ASE).

Arie Kapteyn
RAND and Tilburg University

March 2009
Revised May 5, 2010

Each week, the Dutch Postcode Lottery (PCL) randomly selects a postal code, and distributes cash and a new BMW to lottery participants in that code. We study the effects of these shocks on lottery winners and their neighbors. Consistent with the life-cycle hypothesis, the effects on winners' consumption are largely confined to cars and other durables. Consistent with the theory of in-kind transfers, the vast majority of BMW winners liquidate their BMWs. We do, however, detect substantial social effects of lottery winnings: PCL nonparticipants who live next door to winners have significantly higher levels of car consumption than other nonparticipants.

JEL classifications: D12, C21. *Keywords:* social interactions, natural experiments.

* Corresponding author; pjkuhn@econ.ucsb.edu. We thank the editor, three anonymous referees, Tobias Klein, Andrew Oswald, Bruce Sacerdote, and Linda Waite, as well as many seminar participants for helpful comments. We also thank Herma van der Vleuten (University of Groningen), Kim Paulussen, Kamieke van de Riet, Peter Willems (GfK Benelux), and Jenny van Laar (Bridgis) for their contributions to the data collection process. The data collection was made possible through a grant from the Netherlands Organization for Scientific Research (NWO). Soetevent's research was supported by the MacArthur Research Network on Social Interactions and Economic Inequality and by the Netherlands Organization for Scientific Research under grant 457-07-010. Kuhn thanks UC Berkeley's Center for Labor Economics for their generous hospitality while working on this project. Supporting materials are available at <http://www.econ.ucsb.edu/~pjkuhn/Data/PCL/PCLIndex.html>.

1. Introduction

Economic theory offers a rich set of predictions concerning the effects of income shocks on household behavior. For example, the permanent income hypothesis argues that households should save the lion's share of any income shocks they receive. The classical theory of in-kind transfers predicts that households who receive such shocks in kind (for example as vouchers for food, school, or housing) should in most cases treat them as cash. According to the Easterlin hypothesis, positive shocks to one's neighbors' incomes should reduce one's happiness, while Veblen effects suggest that shocks to neighbors' incomes could also affect one's own consumption.

To date, empirical testing of all the above predictions has been hampered by the lack of credibly exogenous variation in either a household's own income, or in the income of its neighbors. Recently, however, progress regarding own income effects has been made by using a sample of lottery winners (Imbens, Rubin and Sacerdote 2001) and by exploiting the random timing of income tax rebates (Agarwal, Liu and Souleles 2007). Social effects of income shocks have recently been studied by Luttmer (2005), and Angelucci and De Giorgi (2009), though only the latter paper has access to randomized variation in neighbors' incomes. We are aware of no natural-experimental evidence concerning the effects of in-kind transfers on consumption decisions.

In this paper we study all the above hypotheses using data from the Dutch Postcode Lottery (PCL). Each week, this lottery allocates a prize to participants in a randomly chosen postcode (containing 19 households on average). More than one quarter of the Dutch population participates in the lottery. A participant wins €12,500 per ticket. In addition, one participating household in the winning postcode receives a new BMW. From an experimental design perspective, the lottery provides PCL participants in the winning code with an unexpected temporary income shock equal on average to about eight months of income, while leaving all other households' incomes unchanged.¹ Our survey data includes information on consumption and happiness for four groups of

¹ These income shocks are much larger than those studied by Agarwal et al. (2007); other differences between our study and theirs are that we have income information, we have considerable detail on the types of household expenditures, and the income shocks in our experiment were almost surely unexpected by their recipients. Compared to Imbens et al. (2001), we have a larger sample, more detail on types of consumption expenditures, smaller income shocks, and the income shocks we study are temporary.

households: lottery participants and nonparticipants in winning and in nearby non-winning postcodes. Given the inherent randomness in the prize draws and absent externalities between winning and non-winning postcodes, participants in non-winning postcodes constitute a valid counterfactual for participants in winning postcodes. This allows us to test for the effects of unexpected, temporary income shocks (both cash and in kind) on winning households' consumption and happiness under quite general conditions.

Similarly, under the above conditions nonparticipants in non-winning postcodes constitute a valid counterfactual for nonparticipants in winning postcodes. This allows for a clean test for social effects of income shocks on non-participating households' consumption and happiness. A noteworthy feature of our analysis of social effects is its partial-population design, in which a subset of the members of a fixed peer group receives an exogenous shock. Unlike what Moffitt (2001) calls group-changing interventions (where subjects are moved to a new peer group), partial population designs are not contaminated by the causal effects of mobility itself. Partial population designs have recently been used to estimate the extent of information dissemination and learning among neighbors and friends (Duflo and Saez 2003; Duflo, Kremer and Robinson 2004; Miguel and Kremer 2004) and peer effects in school participation (Bobonis and Finan 2009; Lalive and Cattaneo 2009).

One notion of an ideal partial population design starts with a sample of social groups (say villages) assumed to be isolated from one another (so there are no cross-village externalities). Next, a number of villages are randomly assigned to be 'treatment villages'. Finally, a random subset of the households in these treatment villages is treated. No households in the nontreatment villages are treated. In this design, between-village comparisons identify (a) the effect of village-level treatment on the mean outcome of the entire village, inclusive of all social interaction effects; (b) the effect of village-level treatment on treated households; and (c) the effect of village treatment on non-treated households. The latter provides a clean test for the existence of social effects since it should be zero if there are no externalities within villages.

Most, if not all, actual studies of social effects differ from the above ideal in some way. For example, in a well-known study of health-related interventions, Miguel and Kremer (2004) randomize treatment eligibility across groups (in their case Kenyan

schools), but cannot experimentally manipulate individuals' treatment takeup decisions within schools. Thus, while they can identify the effects of school-level treatment on the entire school (inclusive of social interactions between treated and nontreated children), they are forced to rely on nonexperimental methods to identify within-school social effects.

Like Miguel and Kremer, our paper differs from the ideal design in that, rather than being randomly assigned to treatment within social groups, individuals in our data choose whether to take up treatment or not. A crucial distinction, however, is that in our case, this takeup decision (participation in the Postcode Lottery) is made *before* the random selection of groups (in our case postcodes) to be treated is made. This has two consequences: First, in contrast to designs like Miguel and Kremer's, we observe *which* individuals in the nontreated postcodes would have been treated had those postcodes been randomly selected for treatment. Thus we can estimate separate effects of postcode treatment for PCL participants and nonparticipants.² Second, although lottery nonparticipants are not a random sample of the population, the difference in their outcomes between treated versus nontreated postcodes provides a valid test for the existence of within-postcode social effects on that *ex ante*-identifiable group. Thus, in contrast to Miguel and Kremer, we are able to test for within-group social effects directly from our experimental design.³

Most related to our study is Angelucci and De Giorgi (2009), which is to our knowledge the first study of social effects in household consumption using a partial-population design. Using data from a program that targets poor households in small rural communities in Mexico with bimonthly conditional grants to improve education, health

² In designs where agents self-select into treatment after eligibility is randomized, the treatment effect on-the-treated (TOT) is sometimes estimated by the difference in outcome between treatment and control groups, divided by the fraction of members f in the treatment group who take up the treatment. This procedure suffers from the potential defect that the treatment effects are derived by comparing a subsample of observations from the experimental group to the entire sample of control observations. See Burtless and Greenberg (1983) for an early discussion.

³ Our design does not, however, yield an experimental estimate of the effects of neighbors' lottery winnings on PCL *participants*. Heterogenous treatment effects between the treated and nontreated may however be less of a concern in our context than in cases where individual selection into treatment occurs after group randomization, since in the latter case we would expect persons who expect to benefit most from the treatment to select into it (and those who benefit most from spillovers, e.g. from *others'* immunizations, to select out of it). In our case, it is hard to think of a reason why neighbors' income shocks would have a larger effect on the consumption of PCL nonparticipants than on the consumption of participants.

and nutrition, they find strong evidence of positive program externalities on non-eligible households through informal insurance and credit markets. Treatments are randomized over villages. Each village contains eligible (poor) and non-eligible (non-poor) households, and eligibility is a nonrandom event based on a household's poverty status. Estimates of the average treatment effects (ATEs) and social or indirect treatment effects (ITEs) are obtained by comparing mean observed outcomes for the poor (nonpoor) in treatment and control villages.⁴ Thus, their design is analogous to ours, with pre-program income playing the same role as PCL participation in our study. Our study, however, relates to an urbanized, developed economy with broad participation in formal capital and insurance markets, where we might expect the permanent income hypothesis to be more applicable.

In addition to simple comparisons between winning and non-winning codes, we present estimates from a regression-based approach that accounts for differences in treatment intensity (i.e. *amounts* won, both in the household and in its vicinity), and explore the sensitivity of our results to a wide range of alternative specifications of a household's social comparison group, such as its two and four nearest neighbors and the set of other households within 25 meters. The regression framework also allows us --at the cost of some additional assumptions-- to identify some additional quantities (such as pure 'own' effects of lottery income shocks on participants) in addition to the above ATE and ITE.

Our estimated effects of lottery winnings on winners are consistent with theories of in-kind transfers (e.g. Moffitt 1984, 1989): we find that the vast majority of BMW winners do not own a BMW six months after the lottery. Consistent with a simple life-cycle model of consumption, we do not detect any effect of winning the postcode lottery on most components of winning households' expenditures, including food at home, transportation, and total monthly outlays. However, we do find effects on car expenditures and other durable expenditures of winners. These effects are consistent with a version of the life-cycle consumption model in which households adjust the timing of durables purchases to smooth consumption (Browning and Crossley 2009) or a model

⁴ As in Miguel-Kremer, individuals in Angelucci and De Giorgi's design who are eligible based on their income could in principle fail to sign up for the program. This turns out to be rare in practice, however.

with self-imposed borrowing constraints (Shefrin and Thaler 1988).

Turning to social effects, we detect statistically significant effects of lottery prizes on the car consumption of neighbors of winners. For example, having an immediate neighbor win the PCL raises the probability that a household will buy a car in the next six months by close to 7 percentage points and reduces the mean age of its main car at the survey date by half a year (about a 7 percent decline). Relative to the modest effects of the lottery prizes on the consumption choices of winning households, these effects on neighbors are large.

2. The Dutch Postcode Lottery and the Street Prize

The *Nationale Postcode Loterij* (PCL) is the second largest national lottery in the Netherlands, with a revenue market share of 26.6 percent in 2005. Contrary to the *Staatsloterij* (State Lottery), the largest Dutch lottery (market share 42.6 percent in 2005), the PCL is a charity lottery: A condition for its license is that at least 50 percent of revenues must be donated to approved charities. Since its inception in 1990 the PCL has grown steadily. In 2005 almost 30 percent of the Dutch population participated in the PCL, with an estimated annual expenditure per participant of €175.

In the PCL the lottery ticket number is the participant's 6-digit postcode. Thus, conditional on purchasing a ticket, a household's probability of winning the PCL in any given week is approximately equal to one divided by the number of postcodes in the Netherlands (about 430,000). The popularity of the PCL is sometimes attributed to its potential to induce regret among nonparticipants (Zeelenberg and Pieters 2004); nonparticipants living in a winning code know for sure that they would have won had they purchased a ticket. Moreover, the weekly award of prizes is widely publicized in the media, including – in most cases – a broadcast on national television around 10 pm on Sunday evenings. This five-minute program features happy winners and, occasionally, less happy non-winners.

During our sample period, participants paid from €6.25 to €6.75 per ticket (the price increased during the sample period), which is debited monthly from their bank account. There are no restrictions on the number of lottery tickets that can be purchased per participant. The PCL awards several prizes, ranging from very large ones (around €10

million, once or twice a year) to very small ones.⁵ In this paper we focus on one particular prize, the weekly Street Prize. If a postcode is selected as the winning code, a prize of €12,500 per lottery ticket is awarded to participants living in that postcode. Net of the 25% lottery tax, which is applied to all lottery prizes larger than €454, this amounts to €9,375 during our sample period. Because randomization is over postcodes instead of tickets, the number of tickets owned does not affect the probability of winning, only the amount won conditional on winning.

In addition to the monetary prizes, one of the Street-Prize winners wins a new BMW. The BMW winner is chosen by randomly selecting one of the winning lottery tickets. In contrast to the monetary prizes, the probability of winning the BMW does increase with the number of tickets held. BMW winners have the option of receiving the cash value of the BMW in lieu of the car itself. This however involves a substantial tax penalty, since the PCL authority pays the 25% lottery tax for winners who accept the BMW in kind, but not for those who choose the cash equivalent (about €25,000). Of course, winners also have the option of selling their new BMW, and incurring any associated transactions costs.⁶

3. Econometric Framework

As mentioned, our data is a sample of households living in winning postcodes, as well as in neighboring, nonwinning postcodes. Throughout our discussion, we refer to a winning postcode plus all the non-winning nearby codes associated with it as a “codegroup”. In addition to being part of a codegroup $CG(i)$ and a postcode $PC(i)$ (with $PC(i) \subseteq CG(i)$), we define for each household i its assumed social comparison group $N(i)$. A natural candidate is the other households in the postcode: $N(i) = PC(i) \setminus \{i\}$, but our empirical analysis also considers alternative definitions of $N(i)$, like a household’s two (four) immediate neighbors and the set of other households within 25 meters distance. The treatment consists of a partial-population experiment with lottery prizes being awarded to lottery participants ($P_i = 1$) in a randomly chosen set of winning postcodes ($W_i = 1$).

⁵ Eligibility for the large ‘jackpot’ prize requires households to pay an extra monthly fee (which increased from €1.25 to €1.50 during our sample period).

Effects on Winners

Define C_{1i} as the consumption of household i in a winning postcode ($W_i=1$) if a prize would fall in its postcode; similarly, C_{0i} is the consumption of a household if no prize would fall in its postcode. Observed consumption equals $C_i = C_{0i} + W_i(C_{1i} - C_{0i})$. We consider the effect of living in a winning postcode on the consumption of lottery participants. The average treatment effect on participants is given by:

$$(1) \quad \text{ATE} = E[C_{1i} | W_i=1, P_i=1] - E[C_{0i} | W_i=1, P_i=1]$$

The data do not reveal potential consumption in the absence of the treatment for participants in winning codes. We assume that the value of this expectation is the same as the potential consumption in the absence of prizes for participants in non-winning codes:

$$\text{Assumption 1: } E(C_{0i} | W_i=1, P_i=1) = E(C_{0i} | W_i=0, P_i=1).$$

Under Assumption 1, the difference

$$(2) \quad E[C_i | W_i=1, P_i=1] - E[C_i | W_i=0, P_i=1]$$

identifies the ATE. We note here that the ATE, or the effect of living in a winning code on PCL participants, includes (in addition to the ‘own’ effects of winning money) any social effects deriving from other winning households in the postcode; in this sense the ATE does not identify ‘pure’ own effects of lottery winnings (unless, of course, social effects are zero, which is the implicit assumption in essentially all other studies of own income shocks and consumption). As noted, Section 7 isolates pure own effects by imposing some additional structure.

Effects on Neighbors of Winners

Similarly, the average indirect effect of the lottery prize on non-participating households is given by:

$$(3) \quad \text{ITE} = E(C_{1i} | W_i=1, P_i=0) - E(C_{0i} | W_i=1, P_i=0)$$

Under random assignment of prizes, the expected consumption in the absence of a prize-win is equal for winning and non-winning postcodes, that is:

$$\text{Assumption 2: } E(C_{0i} | W_i=1, P_i=0) = E(C_{0i} | W_i=0, P_i=0)$$

Analogous to the case for participants, under Assumption 2, the difference

⁶ We have observed a small number of ads for BMWs won in the PCL on the Dutch equivalent of eBay.

$$(4) \quad E(C_i | W_i = 1, P_i = 0) - E(C_i | W_i = 0, P_i = 0)$$

identifies the ITE.⁷

4. Data

From September 2003 until July 2006 we sent out written surveys to all addresses in PCL-winning postcodes, six months after the prize was won. Moreover, for each winning postcode, we selected one or more neighboring postcodes as a control group, and sent out the same written survey to those postcodes.⁸ In all, we surveyed households in 419 postcodes. An average postcode contains 19 households; the smallest postcode surveyed contained 4 households and the largest 105. Very few Dutch postcodes (8/419 = 2.1% in our sample) contain more than 45 households. The survey contains questions on household composition, demographic variables (age, gender, ethnicity, family relationships and marital status), education, labor supply, happiness, car ownership, large expenditures, income, and lottery participation. For some of the questions respondents provided information on both current behavior (i.e. six months after the prize) and, retrospectively, behavior a year earlier (i.e. six months before the prize).

At the beginning of the survey, households were invited to participate in “scientific research on expenditures and income of Dutch households”, without any reference to social interactions or the PCL. Questions about lottery participation were asked after the consumption questions, near the end of the survey. Households were offered €7.50 (€10.00 in a small number of cases) to complete the survey. Respondents

⁷ Assumption 2 shows the two opposing objectives in selecting the non-winning postcodes. First, to control as well as possible for any unobserved household characteristics that vary smoothly over space, households in non-winning codes should not live too far away from households in winning postcodes. This maximizes the credibility of Assumption 1. On the other hand, nonparticipants in non-winning postcodes should on average live farther away from winners than do nonparticipants in winning postcodes to ensure that consumption by nonparticipants in non-winning codes is not affected by the prize draw. Importantly, it is likely that any indirect effects on nonparticipants in non-winning postcodes are of the same sign as for nonparticipants in winning postcodes. In that case, our estimated ITE can be interpreted as a lower bound for the true ITE.

⁸ Specifically, for each winning postcode, we examined a map of the area, and first selected the neighboring code that was closest to the winning code. We then added other adjacent postcodes up to the point where the number of non-winning addresses was at least twice the number of addresses in the winning code. If fewer than this number of addresses were available in adjacent codes, we settled for less. In the rare circumstances where we had to choose between multiple neighboring postcodes at a similar distance from the winning postcode, those postcodes with a size (in terms of number of addresses) similar to the winning postcode were included first.

could choose among a number of charities to receive this fee, or could provide any bank account number (including their own) to which the token was to be credited. If households did not respond within two weeks, a reminder was sent to households whose phone number was not known. Other non-respondents were called and asked to complete the questionnaire by telephone with the assistance of a survey agency employee. The response rate was 32.1 percent in winning postcodes, 33.0 percent in non-winning codes, and 32.7 percent overall. This overall rate is close to the average response percentage of Statistics Netherlands for similar surveys. Our final sample contains 2011 observations, 510 of which were completed by telephone.

Since our sampling frame is based on addresses six months after the lottery, our sample would be nonrepresentative if households' propensity to move out of a postcode depends on whether that code (or household) won the PCL. We examined this question directly using the Cadastre and Public Register Agency data on house sales, finding no significant difference in the number of home sales after the PCL prize draw between winning and non-winning addresses. These results are reported on the paper's website.

The street addresses of all respondents to our survey are known. Identification of social effects is however facilitated if we know the location of all lottery winners in the winning codes surveyed, including those winners who did not respond to our survey. The PCL administration provided us with the location (street addresses) of all winners, plus information on winnings for every winning postcode in our survey. From these addresses, and from geographical (x, y) coordinates obtained from municipal land registry offices, we compute two alternative classes of social comparison groups for a typical household (in addition to the postcode itself). The first defines a household's two immediate neighbors as those addresses in the same postcode with house numbers (or unit numbers in the case of multi-unit addresses) immediately below and above its own house number (houses in a postcode all share the same street name); its nearest four neighbors are defined analogously. A second class of proximity measures uses the (x, y) coordinates (in meters) for all addresses in the postcodes in our sample, plus those of all the lottery winners as reported by the PCL administrators. In this data, addresses within the same building have the same (x, y) coordinates. Thus, within buildings, our 'neighbor' variable based on the unit number might be a superior measure of proximity to

winners than that based on (x, y) coordinates. One might argue that physical visibility or social distance is more relevant to peer effects than Euclidian distance. Our choice of the latter is guided by data availability.

By definition, a postcode can only be a winning code if it contains at least one household that purchased a PCL ticket; non-winning codes do not have to satisfy this criterion to be in our sample. As a result, in our sample, the PCL participation rate in winning codes is significantly higher than in non-winning codes (0.32 versus 0.26; $p = .003$). This difference becomes smaller and statistically insignificant in a linear probability model that controls for codegroup fixed effects and respondent characteristics (see the paper's website). To check whether this phenomenon affects our results, we redid the full analysis restricting the sample to postcodes that contain at least 16 addresses (also on the paper's website). In these postcodes, the probability that no one participates in the PCL is smaller than 1 percent.⁹ Our main results remained unchanged.

Tables 1 and 2 report descriptive statistics for the households in our survey. In addition to providing a statistical portrait of our respondents before the 'treatment', Table 1 provides a test of the exogeneity of the lottery win by comparing lottery participants and nonparticipants in winning versus non-winning postcodes.¹⁰ If winning postcodes are randomly selected, if the decision to respond to our survey is independent of whether the code was a winning one, and if households' reporting of their own consumption behavior is not affected by living in a winning postcode, there should be no statistically significant differences between the two columns of Table 1. The table shows that participants in winning codes only differ from those in non-winning codes with regard to higher education level, age, and the amount spent on food away from home.¹¹ For nonparticipants, only the share of two-headed households is somewhat larger in winning codes. Overall, the message of Table 1 is one of consistency with random selection of winning codes and absence of response or recall bias. To test this more formally, we estimated linear probability models of living in a winning code as a function of all the

⁹ The PCL administration did not provide the number of participants in non-winning codes, in which case we could simply have dropped the non-winning codes without participants.

¹⁰ The statistics in Table 1 are also consistent with national means from Statistics Netherlands. See this paper's webpage for details.

variables in Table 1, separately for PCL participants and nonparticipants.¹² In neither case could we reject the hypothesis that all of the coefficients in the regression (excluding the constant term) were zero ($p=0.2771$ for participants, $p=0.9012$ for nonparticipants).

Table 2 presents descriptive statistics on awareness of and participation in the postcode lottery and on amounts won. In both winning and non-winning postcodes, all participating households say they remember that a PCL Street Prize was awarded in that code six months after the fact. Awareness among nonparticipants is significantly higher in winning postcodes than in non-winning codes. Both for participants and nonparticipants, those living in the winning postcode do recall the number of households who won, and report house numbers of at least some of these winners significantly more often than their counterparts in the neighboring non-winning codes do (for example, for nonparticipants the numbers are 42 vs. 17 percent and 20 vs. 4 percent, respectively). Such high awareness levels would seem to be conducive to social effects.

The next two sections of Table 2 provide information on PCL ticket ownership and winnings. The average PCL participant held about 1.8 tickets; the average amount won was €18,596. After the 25% lottery tax, these winnings correspond to about €13,947, or about 7 months of post-tax income for a typical Dutch household. 11.2 percent of ticket owners in winning codes won a BMW. Adding in the expected value of this BMW (we value the BMW at €25,000), the average amount won by a household in the PCL rises to $€13,947 + .112(€25,000) = €6,047$, or about 8 months of income for an average family in our sample.¹³

The remainder of Table 2 presents a variety of indicators of the amount of lottery winnings that took place in the geographical vicinity of a typical household in our data. According to our data, in winning codes, just under half of PCL nonparticipants lived next door to a PCL winner.¹⁴ The average amount won by a PCL nonparticipant's two

¹¹ The recalled expenditure amounts tend to be higher for the winners than for the participants in non-winning postcodes. Conceivably, the winnings induce respondents to recall higher amounts than they actually spent in the past.

¹² Education and age effects were interacted with whether the household was single-headed, to incorporate the fact that our indicators of these characteristics differ by household type.

¹³ The BMW awarded in the PCL is a model 116i. Between 2005 and 2007 we found advertised new prices for this vehicle ranging from €25,400 to €28,500.

¹⁴ In principle, a household in (but at the boundary of) a non-winning code could live next door to a PCL winner; our method of identifying neighbors will not capture these households. Households living in non-

immediate neighbors (combined) was €12,552. Both of these numbers were somewhat higher among PCL participants; this may reflect social interactions in ticket holdings. For an average non-winning household living in a winning code, 2.4 winning households lived in the same building (and were thus assigned the same (x, y) location) and 2.9 winning households lived within 25 meters. A typical non-winning household in a winning code shared that code with 7.8 winning households who won a total of €199,842 between them.

5. Effects of Lottery Prizes on Winners

Table 3 presents descriptive statistics for a list of outcome variables that are possibly affected either by winning the lottery, or by living close to lottery winners. The difference between the entries in columns 3 and 1 provides the empirical analogue to the ATE described by equation (2).

For PCL participants in winning and non-winning codes we find statistically significant post-lottery differences in expenditures on food away from home, other monthly items, and total monthly expenditures.¹⁵ We also find that participants in winning codes were 4.5 times as likely to initiate major exterior home renovations during this period and spent over €500 more on noncar durables than participants in non-winning codes.¹⁶ Note that these responses are consistent with liquidity-constrained versions of the life cycle consumption model (e.g. Browning and Crossley 2009, where households adjust the timing of durables purchases to smooth consumption), or with ‘mental accounting’ models with self-imposed borrowing constraints (Shefrin and Thaler 1988). In addition, Table 3 shows some significant differences in car ownership between winning and non-winning codes six months after the lottery date. For example, the main car of participating households in winning codes is on average thirteen months newer

winning codes but close to winners are, however, included in our Euclidean-distance-based measures of proximity to winners.

¹⁵ The effect for food away from home, however, seems to be driven by the fact that pre-lottery expenditures in that category were already higher among participants in winning codes; see Table 1. These retrospective consumption measures may however be biased among lottery winners.

¹⁶ The ‘occasional’ expenditure amounts (including vacations and non-car durables) in Table 3 are not directly comparable to the pre-lottery levels in Table 1 because the latter refer to the 18-month period preceding the lottery date in one’s codegroup.

than in non-winning households.¹⁷ PCL winners are also somewhat more likely than non-winners to donate the fee they receive for completing our survey to charity.

Since BMWs are a prize in the PCL, it is of interest to look specifically at BMW ownership six months after the lottery. The results for BMW ownership are clear: six months after the lottery, participating households in winning codes are statistically no more likely to own a BMW than participating households in non-winning codes. Thus it appears that most BMW winners either elected to receive the cash prize in lieu of the BMW or sold their BMWs shortly after they received them. To explore this result further, the bottom panel of Table 3 provides additional details on post-PCL BMW ownership. It shows that 25 BMW winners responded to our survey. Of these, only four, or 16 percent, still owned a BMW at the survey date.¹⁸ While this percentage of BMWs is more than one would expect in a random sample of Dutch households, overall the behavior of the BMW winners in our sample is remarkably consistent with simple models of in-kind transfers (see for example Moffitt 1984, 1989): whenever a gift in kind would induce a suboptimal consumption mix (as a new BMW is likely to do for the vast majority of Dutch households), that gift should, if possible, be converted into its cash equivalent, and either spent on other items or saved.

6. Effects of Lottery Prizes on Neighbors of Winners

The difference between the entries in columns 4 and 2 of Table 3 provides the empirical analogue to the ITE described by equation (4). We do not find significant differences in consumption for any of the monthly or non-car durable expenditures. Among nonparticipants, those who live in winning codes are less likely to play the PCL six months later. The regret aspect of the PCL combined with a “lightning never strikes (the same postcode) twice” misperception might provide an explanation: Having observed the “losers” in their neighboring winning postcode, nonparticipants in non-

¹⁷ To avoid the possibility that this is simply the mechanical consequence of the fact that cars (specifically BMWs) were a prize in the PCL, all BMW winners have been excluded from this sample.

¹⁸ Our survey collects information on a maximum of two cars per household. Could a significant number of BMW winners still own a BMW as their third, or higher-order car? For 18 of the 25 BMW winners, this is impossible, since they own either zero or one car at the survey date. Of the remaining 7 households, two report owning a BMW at the survey date. Thus at the very most, 5 BMWs could be missing from our sample for this reason. Since owning more than two cars is very rare in the Netherlands, the true number is likely much smaller.

winning codes may feel a strong urge to “insure” against non-winning through participation. Witnessing one’s postcode-mates win the PCL does not make non-winning households any less happy six months after the fact. However, Table 3 also shows that nonparticipants in winning codes are significantly more likely (24 versus 17 percent) to buy a car in the six months after the lottery date, to own more cars at the survey date, and to own more total car efficiency units at the survey date.¹⁹

Of all our consumption indicators, cars are (a) likely to be the most visible to one’s residential neighbors, and (b) durable. Unlike an expensive party or vacation, a household’s neighbors are continuously reminded of its new car after it has been purchased. This leads us to conduct a more detailed analysis of social effects of lottery winnings on car consumption, using simple comparisons that take the best possible advantage of the exogenous assignment of lottery winnings in our sample. Table 4 presents simple means of our four car consumption indicators for three subgroups of nonparticipants: those who live in non-winning codes, those who live in winning codes but do not live next door to a PCL winner (these households will be affected by any postcode-wide social effects on car consumption), and those who do live next door to a PCL winner (who are affected both by postcode-wide social effects and those stemming from their immediate neighbors). In addition to our car consumption measures, the second row of Table 4 asks – as a placebo test – whether a household acquired one of its currently-owned vehicles in the six months before (rather than after) the lottery. Clearly, all indicators of car consumption in Table 4 except row two are largest (smallest for car age) for households living next door to PCL winners. To the extent that, within postcodes, living next door to a PCL ticket holder is exogenous, Table 4 provides convincing evidence that (a) social effects of winning the PCL do exist, but (b) they are highly localized, restricted in large part to a household’s nearest neighbors.

Table 5 extends the analysis of Table 4 by addressing the possible endogeneity of living within one or two doors of a PCL winner. If living next door to a PCL winner is endogenous (which in our framework would require unobserved household characteristics to be correlated *within* postcodes at the extremely detailed level of next-

¹⁹ Our efficiency units measure combines information on both the number and quality of cars owned by the household at the survey date. A car that is less than one year old is defined as one efficiency unit; all other

door neighbors), Table 4’s estimates could be biased. To address this concern, Table 5 uses information from our survey to construct an indicator of whether a household lives next door to a PCL participant. As already noted, this indicator is necessarily incomplete because it is survey-derived. In particular, to be in the sample for columns 1 and 2 of Table 5, a survey household must have had at least one of its two nearest neighbors respond to the PCL survey *and* at least one of those neighbors must report holding a PCL ticket in our survey. That said, by construction, all the households in the Table 5 sample live next door to a known PCL ticket owner; the only variation in whether their neighbors have PCL winnings is generated by the random selection of winning codes. As the Table makes clear, all indicators of current car consumption with the exception of the ‘placebo’ measure in row 2 are higher (lower for car age) for those households who were exposed to the treatment of having a next-door neighbor win the PCL. Sample sizes are small, however, and only one of the comparisons is statistically significant, at the 10% level. In columns 3 and 4 we replicate this analysis, expanding the sample to nonparticipating households who live within two doors of a known PCL participant. Both sample size and statistical significance now increase.

7. Regression results

In this section we specify and estimate a simple regression model of consumption in the presence of lottery winnings. These regressions control for codegroup fixed effects and for any differences in observed pre-lottery household characteristics (including income) that are present despite randomization, thereby increasing statistical precision. The regression approach also facilitates using alternative definitions of a household’s social comparison group (see Table 2) and information on the intensity of treatment (amount of winnings). Our regression framework also generates estimates of the pure ‘own’ effects of lottery winnings on consumption, and relates these estimates to the ITEs and ATEs estimated in the preceding sections.

Specifically, we assume a household’s consumption is given by:

$$(5) \quad C_i = \alpha + \beta WIN_i + \gamma WIN_{N(i)} + \mu_i,$$

cars owned by the household are depreciated by 15% per year.

where WIN_i is the household's own lottery winnings, $WIN_{N(i)}$ measures the amount won by its comparison group, $N(i)$, and μ_i represents other (observed and unobserved) determinants of consumption.²⁰ We define β in (5) as the *own effect* of lottery winnings, and γ as the *social effect*. Both of these coefficients are related in a straightforward way to the ITE and ATE estimated in the previous section; specifically, the ITE is proportional to γ , while the ATE (the effect of winnings on winners) is a linear combination of γ and β (since winners are exposed to income shocks among their neighbors as well).²¹

Finally, the regression approach facilitates the testing of some simple hypotheses in the theory of consumption. For example, according to the life cycle model, households' consumption decisions should be more strongly linked to their own nonlottery income than to own lottery winnings, because the former contains a larger permanent component. The life-cycle model also predicts larger consumption responses to own lottery winnings among older, and among poorer households because in both cases the income shock is a higher fraction of remaining lifetime income. Models of conspicuous consumption predict that only *visible* goods should be subject to neighbors' income shocks. Finally, psychological models of adaptation predict that temporary income shocks will raise happiness more than permanent income shocks, while models where utility depends on relative income predict that positive shocks to neighbors'

²⁰ An alternative specification of (5) would have own consumption depend not on neighbors' income but on neighbors' *consumption* (i.e. that social effects are endogenous, not exogenous). In this paper we focus on the latter, for three reasons. Most obviously, shocks to neighbors' incomes are randomly assigned in our context; changes in their consumption are not. Second, households' consumption is multidimensional, which makes it far from clear how to best model consumption interdependencies. (Does my propensity to buy a car depend on your purchase of a specific model and quality of car, on your decision to purchase any car, on the fact that you recently made a visible purchase of any kind, or simply on the fact that you won some money?). Third (and related), effects of neighbors' consumption and income on a household's own consumption are, in general, not separately identified by a neighbor's-income instrument.

²¹ Since (ignoring taxes) each winning ticket pays 12.5 (thousand) Euro, $WIN_i = 12.5 \cdot T_i \cdot W_i$ and $WIN_{N(i)} = 12.5 \cdot T_{N(i)} \cdot W_i$, where T_i denotes the number of tickets held by the household. Substituting (5) into (2) and (4) reveals that, under Assumptions 1 and 2, the ATE and ITE can be expressed, respectively, as $ITE = 12.5\gamma\bar{T}_{N(i)}$ and $ATE = 12.5\beta\bar{T}_i + 12.5\gamma\bar{T}_{N(i)}$, where \bar{T}_i and $\bar{T}_{N(i)}$ respectively are the mean number of own and neighbors' tickets for a representative household in the sample. Note also that (5) implicitly assumes that β and γ are the same for PCL participants and nonparticipants; this is the additional assumption that identifies β in this context.

incomes should reduce happiness.²²

The regression approach also has some potential drawbacks. Note first that whereas W_i (living in a winning code) is randomly determined, individual ticket holdings and hence WIN_i may be correlated with unobserved household tastes or constraints that may affect its consumption of all items, including lottery purchases. For example, risk-loving households may buy more PCL tickets, and may also have different consumption patterns for other items than risk-averse households. Fortunately, we can account for this kind of bias by including controls for the number of tickets T_i held by the household at the time of the win. Specifically, we include in all of our regressions a quadratic in number of tickets purchased plus a fixed effect for participation in the PCL.²³

Unfortunately, a similar strategy is not available for $WIN_{N(i)}$, because we do not have data on lottery ticket ownership for households who did not respond to our survey. Thus, if (for example) $WIN_{N(i)}$ is the total amount won in the household's postcode, $WIN_{N(i)}$ will depend not only on whether the postcode was a winner, but on the number of households it contains and their propensity to own PCL tickets; these are nonrandom factors that might be correlated with a household's consumption decisions. We rely on codegroup fixed effects to control for such factors, but also remark that, as already noted, (a) our results are not sensitive to controls for postcode size, and (b) PCL ticket ownership in our sample is very similar between winning and nonwinning postcodes, especially within codegroups. We also show below that our results are not sensitive to codegroup fixed effects, suggesting that spatial correlation of unobservables does not play a large role in the relationships we identify here.²⁴

²² For recent reviews of both these literatures, see Kahneman, Krueger, Schkade, Schwarz and Stone (2006), and Clark, Frijters, and Shields (2008).

²³ The results are very similar if we include a dummy for each integer number of tickets purchased.

²⁴ Another remedy for possible unobserved differences in neighbors' ticket holdings is to time-difference our consumption measures within households; this is possible for some of our consumption indicators where we have (retrospective) information referring to the period before the lottery win. While recall bias is a concern with these estimates, they are largely similar to our main results. See the previous version of this paper (Kuhn et al, 2008).

Own Effects of Lottery Winnings

Row 1 of Table 6 reports our estimates of β , based on equation (5). Coefficients in this row represent the effects of winning €10,000 on the outcome variable in question; each column of the table corresponds to a different regression (the full specification is described in Table 6). According to Table 6, winning the PCL –now (in contrast to the ATE) controlling for being near other winners-- has few detectable direct effects on monthly and occasional expenditure categories, including food at home, transportation expenditures, other monthly expenditures, and total monthly expenditures. Own effects are, however, found for various aspects of car consumption. For example, winning €10,000 appears to reduce the average age of a household’s main car by about 0.4 years six months after the lottery date, and to raise expenditures on non-car durables by €10.²⁵ Winning the PCL has no effect on a household’s reported happiness six months after the event. Contrary to the simple results for winners in Table 3 (which reflect a combination of own and social effects), greater lottery winnings do not raise the likelihood that a household will donate its fee for completing our survey to charity.

Compared to lottery winnings – which are temporary and unexpected – cross-sectional income differences should have a substantially larger permanent component. In consequence, the life-cycle model of consumption predicts that cross-sectional income differentials should be more strongly related to current expenditures than are lottery winnings, at least for non-durables. While this is true for most of our point estimates, the only statistically significant differences between the effects of lottery and nonlottery income are for whether the household donated its survey fee to charity (at 10%) and for happiness (at 1%). To explore some other predictions of the life-cycle model, we re-estimated Table 6 separately for low- and high-education groups, and for households whose heads are above and below the median age of 50, to see if ‘own’ responses to lottery income were larger among the old and poor. While interesting, the results may reveal more about the distinction between luxuries and necessities and age-related

²⁵ Comparing these estimates to the ATEs in Table 3 suggests that own effects account for most, if not all, of those ATEs. In particular, according to Table 6 the estimated pure own effects of winning €18,000 (the mean amount won by PCL participants) on car age, car efficiency units and non-car durables are -.723, .069 and €58 respectively. These compare to ATEs of -1.114, .078, and €33 respectively. Thus, only a PCL winner’s car age (i.e. the size of its car quality ‘upgrade’) appears to be affected by social effects acting on PCL winners.

variation in a household's consumption priorities than the life cycle hypothesis *per se*: low-education households spend more of their lottery winnings on cars, and less on vacations than high-education households. Young households also spend more of their lottery winnings on cars than older households, but older households' non-car durables consumption is much more responsive to lottery income than younger households'.

The one outcome on which lottery and nonlottery income have the most dramatically different effects in Table 6 is happiness. Indeed, in contrast to our results for lottery winnings, and consistent with both Easterlin (1974) and Stevenson and Wolfers (2008), higher total income is very strongly associated with happiness in a cross-section of households. One interpretation of these contrasting results is that the six-month lag between the PCL win and the survey date is too long: lottery winnings could affect own happiness, but the effects are very transitory, as argued by both Easterlin and by Kahneman et al. (2006). Alternatively, the results of Gardner and Oswald (2007) suggest that six months might be too short: In their analysis of Britons who receive lottery wins between £1000 and £200,000 they find that in the year a prize is won, mental stress goes up, while in subsequent years lottery winners show less stress than non-winners. Finally, it is possible that happiness is simply more linked to long-run personal income than short-run fluctuations, both because permanent income differences enable the household to take more happiness-improving actions, and perhaps because long-term income differences are more likely to be seen as earned and thus "legitimate".²⁶

Social Effects of Lottery Winnings

We estimated the social effects of living near PCL winners using all the indicators of neighbors' winnings summarized in Table 2, but report only the results using four representative indicators in rows 3-6 of Table 6. As already noted, the estimate in row 3 is taken from the same regression as the coefficients in rows 1 and 2, which uses residence in a winning postcode (W_i) as our indicator of neighbors' winnings ($WIN_{N(i)}$). Each of the remaining rows re-runs this regression, *replacing* W_i by a different alternative

²⁶ Another interpretation of our own-happiness results is of course the possibility that the cross-sectional correlation between income and happiness is driven by reverse causation (happier people are more successful in the labor market). This interpretation would, however, be at odds with Stevenson and Wolfers' (2008) claim that economic growth promotes happiness.

indicator of neighbors' winnings. (The estimated coefficients on own winnings and own nonlottery income did not change much, and are not reported to save space.) Of these, row 4 is representative of our results when defining neighbors using Euclidean distance measures: none of these distance-based indicators of neighbors' winnings had statistically significant effects on consumption. We conjecture that these measures do not discriminate sufficiently among the very large share (81 percent) of Dutch households who live in multi-unit dwellings. (Households are defined as living in a multi-unit dwelling if their address shares a (x,y) -location with at least one other address.)

Excluding these distance-based measures, Table 6 provides some evidence that neighbors' winnings affect current PCL participation and disposition of the survey fee. For the vast majority of consumption indicators, however, no effects are found. An outcome of particular interest is happiness. According to column 14 in Table 6, living in a winning postcode (not winning oneself) has no effect on household's happiness; more precisely we say with 95% confidence that living in a winning postcode (but not winning oneself) reduces happiness by no more than .11 of a standard deviation, and that it raises happiness by no more than .07 of a standard deviation. This result contrasts with Luttmer's (2005). However, his result refers to effects of neighbors' *earnings* on happiness, which have a bigger permanent component than lottery winnings, and neighbors' earnings are not randomly assigned in his analysis. The absence of an effect of exogenous changes in neighbors' incomes on own happiness in our data is also consistent with Stevenson and Wolfers' (2008) claim that relative incomes do not have large effects on happiness.

Finally, consistent with the ITE estimates from the previous section, Table 6 suggests the presence of social effects for two aspects of consumption that are arguably most visible to one's neighbors: exterior home renovations and cars. While the evidence for the former is confined to one regression specification, the evidence for cars is more robust. Statistically significant effects are found for all four indicators of car consumption, and for three of our four measures of neighbors' winnings. These estimates of social effects on car consumption are substantial in size. For example, having an immediate neighbor win the PCL raises the probability that a household will buy a car in the next six months by close to 7 percentage points and reduces the mean age of its main

car at the survey date by half a year (about a 7 percent decline). For two car consumption indicators (total car efficiency units and the age of the main car), the estimated effects of an immediate neighbor winning the PCL are very similar in size to the estimated own effects of winning €10,000; for the incidence of car purchases in the past six months it is actually greater than the own effect. Thus, relative to the fact that the one-time lottery win has only a modest effect on households' own consumption choices (consistent with the life-cycle hypothesis), these effects are substantial in magnitude.

A final possible concern with the results in Table 6 is the possibility that our procedure for selecting 'control' postcodes for the winning codes may have induced some nonrandomness. (In most cases we simply included all adjoining codes, but occasionally some discretion was involved: see footnote 8.) Although the available evidence suggests that our randomization was successful, we also tested the robustness of our results to the selection of non-winning postcodes by removing the codegroup fixed effects from the regressions in Table 6. In regressions without fixed effects, we are comparing the behavior of households in winning codes to households in *all* non-winning postcodes in our sample. The results (reported on the paper's website) were virtually unchanged. In sum, our regression estimates of social effects confirm the nonparametric results reported in sections 5 and 6.

8. Conclusion

We have used the natural experiment associated with the Dutch postcode lottery (PCL) to study the own and social effects of a temporary, unexpected income shock equal to about eight months of income on households' consumption behavior and self-reported happiness. The natural experiment provided by the PCL has a number of advantages, including exogeneity of the income shock to a household's residential neighbors and the absence of direct causal effects of household mobility. According to our estimates, the effects of a lottery prize on winners are confined largely to cars and other consumer durables. This finding is consistent with a permanent income model in which households adjust the timing of their durables purchases to smooth consumption, or with mental accounting models in which households are reluctant to borrow from accounts viewed as assets. In addition, as predicted by simple models of transfers in kind, the vast majority

of households who exogenously receive a large, in-kind transfer (a new BMW) converted that prize into other goods or savings, despite the transactions cost or tax penalty associated with doing so.

We do find robust evidence for effects of lottery prizes on neighbors of winners, but only for one good --car consumption-- which is likely to be easily, and repeatedly, visible to a household's neighbors. Higher-income households in our data are significantly happier than other households, but lottery winnings do not make households happier, nor do they make neighboring households less happy. The latter two results, based on genuinely exogenous shocks to own and neighbors' incomes, would seem to present challenges for both relative income-based and for 'habituation' models of happiness.

What models of consumer behavior might explain the social effects estimated in our data? While it is tempting to interpret our estimates as reflective of a psychological need to "keep up with the van den Bergs"²⁷, we note that they could also be driven by other factors. For example, social spillovers in car consumption could be driven by winning households selling their used car to neighbors; or by households passing money to immediate neighbors, who might be family members.²⁸ In this regard it is noteworthy that intrafamily income transfers were indeed the main channel responsible for the social effects detected by Angelucci and De Giorgi (2009) in Mexican villages; while this seems less likely in Dutch postcodes we cannot rule it out. Also, it is worth re-emphasizing that our estimates do not distinguish 'imitative' consumption patterns, or what are sometimes more generally called endogenous social effects (I buy a car because you buy one) from exogenous social effects of neighbors' incomes on a household's consumption. Still, we find convincing evidence that households' consumption of visible, durable goods (and *only* such goods) is affected by genuinely exogenous shocks to their neighbors' incomes.

²⁷ Parallel to "Jones", van den Berg is the fourth most common surname in the Netherlands, according to Wikipedia. (De Jong is first).

²⁸ Another possibility, suggested by Grinblatt, Keloharju and Ikaheimo (2008), is information sharing about a specific car model. We think this is unlikely in the current context because most of the cars purchased by nonwinners after the lottery date are used vehicles (i.e. more than one year old). Note also that our estimated social effects for cars cannot be attributed to BMW winners selling an unwanted BMW to neighbors, since (among other reasons) none of the cars owned by nonwinners who bought a car since the lottery date are new BMWs.

Finally, we note that, despite the lack of detectable own spending responses for most consumption items, our results contain some encouraging news for fiscal policies such as unexpected tax rebates designed to stimulate consumer spending in developed economies: To the extent that such ‘stimulus’ policies aim specifically at durables (a well known recent example is the U.S.’s “cash for clunkers” subsidies for automobile purchases) our results suggest that they may have substantial own effects, as well as significant social multiplier effects (Glaeser, Sacerdote and Scheinkman 2003). These social multipliers are distinct from, and would presumably operate in addition to, the usual Keynesian multipliers that have been studied in this context.

References

- Agarwal, Sumit, Chunlin Liu, and Nicholas S. Souleles. (2007). "The Reaction of Consumer Spending and Debt to Tax Rebates—Evidence from Consumer Credit Data," *Journal of Political Economy* 115(6) (December): 986-1019.
- Angelucci, M., and G. De Giorgi (2009). "Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles' Consumption?", *American Economic Review*, vol. 99(1), pp. 486-508.
- Bobonis, G. J., and F. Finan (2009), "Neighborhood Peer Effects in Secondary School Enrollment Decisions", *Review of Economics and Statistics* 91 (4), pp. 695–716.
- Browning, M. and T. Crossley (2009), "Shocks, stocks and socks: smoothing consumption over a temporary income loss," *Journal of the European Economic Association* 7(6), pp. 1169-1192.
- Burtless, Gary and David Greenberg (1983) "Measuring the Impact of NIT Experiments on Work Effort" *Industrial and Labor Relations Review*. Vol. 36 (4). p 592-605. July 1983.
- Clark, Andrew E., Paul Frijters, and Michael A. Shields "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles" *Journal of Economic Literature* 2008, 46:1, 95–144.
- Duflo, E., M. Kremer, and J. Robinson (2004). "Understanding Technology Adoption: Fertilizer in Western Kenya – Preliminary Results from Field Experiments," Unpublished manuscript, Massachusetts Institute of Technology.
- Duflo, E., and E. Saez (2003). "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment," *Quarterly Journal of Economics*, 118(3), 815–842.
- Easterlin, Richard A. (1974). "Does Economic Growth Improve the Human Lot? Some Empirical Evidence". In *Nations and Households in Economic Growth: Essays in Honor of Moses Abramowitz*, Paul A. David and Melvin Reder, eds. New York: Academic Press.
- Gardner, Jonathan and Andrew J. Oswald (2007). "Money and mental wellbeing: A longitudinal study of medium-sized lottery wins" *Journal of Health Economics* 26 (2007) 49–60
- Glaeser, Edward L., Bruce I. Sacerdote and Jose A. Scheinkman (2003). "The Social Multiplier". *Journal of the European Economic Association* 1(2-3) (April-May): 345-53
- Grinblatt, Mark, Matti Keloharju and Seppo Ikaheimo (2008). "Social Influence and Consumption: Evidence from the Automobile Purchases of Neighbors", *Review of Economics and Statistics*, Vol. 90, No. 4: 735–753.
- Imbens, G., D. Rubin, and B. Sacerdote (2001). "Estimating the Effect of Unearned Income on Labor Earnings, Savings and Consumption: Evidence from a Survey of Lottery Players" *American Economic Review* 91(4) (September): 778-794.
- Kahneman, Daniel, Alan Krueger, David Schkade, Norbert Schwarz, and Arthur Stone (2006). "Would You Be Happier If You Were Richer? A Focusing Illusion", *Science*, 312, pp. 1908-1910.

- Kuhn, Peter J., Peter Kooreman, Adriaan R. Soetevent and Arie Kapteyn (2008). "The Own and Social Effects of an Unexpected Income Shock: Evidence from the Dutch Postcode Lottery" NBER working paper no. 14305 (May).
- Lalive, R., and Alejandra Cattaneo (2009). "Social Interactions and Schooling Decisions," *Review of Economics & Statistics*, vol. 91(3), pp. 457-477.
- Luttmer, Erzo. (2005). "Neighbors as Negatives: Relative Earnings and Well-Being" *Quarterly Journal of Economics*, 120(3) (August): 963-1002.
- Miguel, E. and Kremer, M. (2004). Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72(1): 159-217.
- Moffitt, R. A. (1984). The effects of grants-in-aid on state and local expenditures: The case of AFDC. *Journal of Public Economics*, 23, 279-305.
- Moffitt, R. A. (1989). Estimating the value of an in-kind transfer: The case of food stamps. *Econometrica*, 57(2), 385-409.
- Moffitt, R. A. (2001). "Policy Interventions, Low-Level Equilibria, and Social Interactions," in *Social Dynamics*, ed. by S. N. Durlauf, and H. P. Young, pp. 45-82. MIT Press.
- Shefrin, Hersch M. and Richard H. Thaler, Richard H. (1988). "The Behavioral Life-Cycle Hypothesis" *Economic Inquiry* 26(4) (October): 609-43
- Stevenson, Betsey and Justin Wolfers (2008). "Economic Growth and Happiness: Reassessing the Easterlin Paradox", *Brooking Papers on Economic Activity*, Spring, pp.1-87.
- Zeelenberg, M. and Pieters, F.G.M. (2004). "Consequences of regret aversion in real life: The case of the Dutch postcode lottery". *Organization Behavior and Human Decision Processes*, 93, 155-168.

Table 1: Descriptive Statistics, Pre-Lottery Characteristics

Permanent or pre-lottery characteristic:	Non-winning postcodes		Winning postcodes	
	(1) Participants	(2) Non- participants	(3) Participants	(4) Non- participants
Basic Demographics:				
Number of persons in household ¹	2.74	2.12	2.73	2.18
Two-headed household? ¹	0.767	0.581**	0.780	0.637**
Secondary education? ²	0.940	0.907	0.942	0.918
Higher vocational training or university? ²	0.326*	0.342	0.256*	0.333
Age ³	49.60**	50.01	52.49**	49.77
Number of children in household	0.664	0.461	0.668	0.507
For single-headed households:				
Head works?	0.514	0.418	0.510	0.468
Head's hours ⁴	14.94	14.19	19.22	15.90
For two-headed households:				
Husband works?	0.723	0.639	0.655	0.632
Husbands' hours ⁴	28.21	23.87	24.85	23.74
Wife works?	0.623	0.549	0.569	0.546
Wife's hours ⁴	17.46	15.32	15.57	15.37
Monthly expenditures: (euro, one year before survey date)				
Food at home	426.28	438.58	460.24	412.80
Food away from home	83.85***	100.74	119.63***	83.42
Transportation	177.43	178.03	178.90	186.47
Other monthly	239.07	240.10	290.82	245.78
Total monthly	936.46	951.79	1057.51	897.60
Occasional Expenditures:				
Exterior home renovations before lottery (any?) ⁵	0.090	0.040	0.063	0.042
Other home renovations before lottery (any?) ^{5,6}	0.183	0.167	0.148	0.149
Vacation expenditures before lottery (euro) ⁵	307.48	349.24	486.12	457.05
Non-car durables expenditures before lottery (euro) ⁵	1737.37	1574.20	1971.67	1519.74
Total Annual Household Income (euro, post tax, pre lottery)				
	25341	22679	27322	22490
Happiness⁷				
	6.87	6.68	6.92	6.78
Sample size⁸				
	301	878	223	477

Notes: Statistically significant differences between columns 1 and 3 and between columns 2 and 4 denoted by ***, **, and * for $p < .01$, $.05$, and $.10$ respectively

1. Refers to one year before the survey date.
2. For single-headed households, indicates whether the head has at least the level of education indicated. For two-headed households, indicates whether at least one head has this level of education.
3. For single-headed households, age in years. For two-headed households, a simple average of the husband's and wife's ages.
4. Usual weekly hours one year before the survey date. Equals zero for non-workers.
5. Refers to the 18 month period preceding the lottery date.
6. Includes interior renovations and those whose type could not be determined.
7. Self assessed on a scale from 1 to 10 (refers to one year before the survey date)
8. Gives the number of observations in winning and non-winning codes. Sample sizes vary across rows of the table due to missing values and sample restrictions.

Table 2: Descriptive Statistics, Lottery Characteristics

	Non-winning postcodes		Winning postcodes	
	(1) Participants	(2) Non- participants	(3) Participants	(4) Non- participants
Awareness of lottery:				
Remember PCL Street Prize?	1.000	0.534***	1.000	0.816***
Recall number of winning households	0.369***	0.170***	0.776***	0.421***
Recall winners' house numbers	0.086***	0.042***	0.363***	0.195***
Lottery participation and winnings:				
Number of tickets held	1.853	0	1.783	0
Amount of cash won (euros)	0***	0	18596***	0
Won BMW?	0***	0	0.112***	0
Indicators of neighbors' winnings:				
Had an immediate neighbor who won?	0***	0***	0.596***	0.488***
Number of immediate neighbors who won (0, 1, 2)	0***	0***	0.744***	0.583***
Total amount won by immediate neighbors (incl. BMW value)	0***	0***	15135***	12552***
Number of winning households:				
In same building	0.299***	0.473***	2.685***	2.362***
Within 25 meters	0.385***	0.536***	3.329***	2.859***
Within 100 meters	4.336***	3.991***	7.982***	6.581***
In your postcode	0***	0***	9.946***	7.815***
Total amount won:				
In same building	8056***	12414***	65315***	61789***
Within 25 meters	10257***	13875***	82095***	74500***
Within 100 meters	107558***	99513***	203266***	165921***
In your postcode	0***	0***	255518***	199842***

Notes: Statistically significant differences between columns 1 and 3 and between columns 2 and 4 denoted by ***, **, and * for $p < .01$, $.05$, and $.10$ respectively. See Table 1 for sample sizes.

Table 3: Descriptive Statistics, Outcome Variables

Post-lottery characteristic:	Non-winning postcodes		Winning postcodes	
	(1) Participants	(2) Non- participants	(3) Participants	(4) Non- participants
Monthly expenditures: (euro, at the survey date)				
Food at home	464.91	471.40	494.55	450.26
Food away from home	83.64***	97.87	124.50***	89.67
Transportation	189.89	192.34	213.87	211.81
Other monthly	254.50**	255.32	335.72**	279.53
Total monthly ¹	995.12**	1002.48	1180.88**	1025.55
Occasional expenditures:				
Exterior home renovations since lottery (any?)	0.0100**	0.0171	0.0448**	0.0231
Other home renovations since lottery (any?)	0.0831	0.0718	0.0717	0.0650
Vacation expenditures since lottery (euro) ²	449.67	195.31	482.96	167.61
Non-car durables expenditures since lottery (euro) ²	658.90**	805.69	1191.80**	744.85
Total annual household income (euro, post tax, pre lottery)	26662	23337	28444	24170
Other outcomes:				
Happiness ³	7.02	6.82	7.06	6.87
PCL participant at survey date?	0.924	0.131**	0.933	0.094**
Donate survey fee to charity?	0.442*	0.418	0.525*	0.440
Car variables (non-BMW winners only):				
Acquired car since lottery date? ⁴	0.156	0.173**	0.170	0.236**
Number of cars (up to 2)	1.203	0.929***	1.212	1.023***
Age of main car (years)	6.502***	7.139	5.388***	6.786
Total car efficiency units ⁵	0.484**	0.350***	0.562**	0.406***
BMW six months after lottery (incl. BMW winners):				
share respondents owning BMW	.037	.015	.031	.008
among cash winners (198 obs.)			.015	
among BMW winners (25 obs.)			.160	

Notes: Statistically significant differences between columns 1 and 3 and between columns 2 and 4 denoted by ***, **, and * for $p < .01$, $.05$, and $.10$ respectively. See Table 1 for sample sizes.

1. Sum of food at home, food away from home, transportation plus other expenditures.
2. Refers to the 6 month period between the lottery and survey dates.
3. Self assessed on a scale from 1 to 10 (refers to the survey date)
4. Equals one if the household acquired any of the autos it currently owns after the lottery date.
5. A car that is less than one year old counts as one unit. All other cars are depreciated at 15 percent per year.

Table 4: Car Consumption Indicators for PCL Nonparticipants

	(1)	(2)	(3)	(4)
	Non-Winning Codes	Winning Codes		
		More than 2 doors from a PCL Winner	Within 2 doors of a PCL Winner	Next door to a PCL Winner
Bought a car since lottery date?	.173 (646)	.242 (95)	.234** (274)	.265*** (189)
Bought a car between 6 and 12 months ago	.127 (647)	.147 (95)	.136 (273)	.133 (188)
Number of cars at survey date	.929 (877)	.920 (137)	1.065***†† (340)	1.073***†† (233)
Age of main car (years)	7.139 (660)	6.978 (93)	6.722 (281)	6.515* (196)
Car efficiency units	.350 (851)	.354 (131)	.428***† (328)	.444***†† (225)

*, **, ***: statistically different from column 1 at 10, 5 and 1 percent respectively

†, ††, †††: statistically different from column 2 at 10, 5 and 1 percent respectively

Sample sizes in parentheses.

Table 5: Car Consumption Indicators for PCL Nonparticipants who are Neighbors of Participants

	(1)	(2)	(3)	(4)
	Next-door neighbors of PCL Participants		Live within two doors of a PCL Participant	
	Living in Non-winning Codes	Living in Winning Codes	Living in Non-winning Codes	Living in Winning Codes
Bought a car since lottery date?	.136 (66)	.217 (46)	.140 (121)	.186 (86)
Bought a car between 6 and 12 months ago?	.152 (66)	.109 (46)	.148 (122)	.116 (86)
Number of cars at survey date	.955 (88)	1.073 (55)	.994 (157)	1.147** (102)
Age of main car (years)	7.386 (70)	6.298 (47)	7.133 (128)	6.500 (88)
Car efficiency units	.352 (83)	.451* (55)	.376 (152)	.485** (100)

*, **, ***: Differs from non-winning codes at 10%, 5% and 1% respectively.

Table 6: Own and Social Effects of Lottery Winnings

Regressor:	OUTCOME:															
	CAR CONSUMPTION				MONTHLY EXPENDITURES					OCCASIONAL EXPENDITURES				OTHER		
	Got Car	Number	Age	Total units	Food at Home	Food Away	Trans-Port	Other	Total	Exterior Renos	Other Renos	Vacation	Non-Car Durables	Happi-ness	PCL Now	Charity
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	
Lottery Winnings	.0208 (.0197)	.0165 (.0252)	-.4019* (.2086)	.0382* (.0210)	13.43 (16.03)	17.65** (7.93)	9.61 (11.81)	19.10 (18.49)	66.56 (41.30)	.0038 (.0032)	.0086 (.0091)	-3.87 (62.43)	310.23** (131.27)	-.0226 (.0500)	.0067 (.0083)	-.0135 (.0141)
Nonlottery Income	-.0002 (.0080)	.0334*** (.0082)	-.1170 (.0801)	.0245*** (.0065)	8.80 (11.29)	15.62** (6.06)	-7.96 (11.99)	15.62* (9.28)	43.43* (24.95)	.0007 (.0033)	.0075* (.0041)	91.56 (74.41)	200.51** (79.35)	.1243*** (.0249)	.0086* (.0046)	.0140* (.0077)
Winning Postcode?	.0460* (.0243)	.0003 (.0272)	-.3857 (.2454)	.0197 (.0163)	-33.17 (38.21)	-10.33 (12.91)	24.71 (21.62)	2.28 (25.77)	-12.49 (71.81)	.0088 (.0076)	-.0198* (.0120)	-54.06 (94.82)	-28.51 (176.44)	-.0323 (.0793)	-.0326** (.0138)	.0711*** (.0235)
ALTERNATIVE MEASURES OF NEIGHBORS' WINNINGS:																
# winners in postcode	.0041* (.0021)	.0040* (.0024)	-.0464** (.0180)	.0043*** (.0014)	-1.26 (2.70)	-.98 (1.03)	.98 (1.71)	-.23 (3.63)	-1.36 (6.54)	.0014** (.0006)	-.0008 (.0011)	9.36* (5.44)	-4.34 (18.90)	.0008 (.0087)	-.0022** (.0011)	.0082*** (.0018)
Winnings in 25 meters	.0018 (.0020)	-.0003 (.0024)	-.0244 (.0213)	.0001 (.0012)	-.98 (1.91)	.40 (.82)	.29 (1.72)	-.83 (2.61)	-1.51 (5.05)	-.0006 (.0005)	.0006 (.0011)	6.06 (5.66)	16.39 (12.50)	-.0069 (.0077)	-.0021** (.0010)	.0032 (.0025)
Neighbor won?	.0675** (.0310)	.0029 (.0323)	-.4996* (.2788)	.0407* (.0237)	-43.49 (37.47)	-10.52 (11.77)	47.54 (34.83)	-1.20 (40.01)	8.96 (97.34)	-.0078 (.0113)	.0084 (.0172)	-32.77 (98.41)	-139.84 (245.05)	.1729 (.1057)	.0122 (.0193)	.0843*** (.0302)

Notes: Rows 1-3 show coefficients on own winnings, own non-lottery income and a winning postcode dummy, all included in the same regression. Rows 4-6 show the coefficients on alternative measures of neighbors' winnings when they are substituted for the winning postcode dummy in the regression shown in rows 1-3 (coefficients on own winnings and nonlottery income do not change much when different indicators of neighbors' winnings are used). All specifications also include a fixed effect for lottery participation, a quadratic in the number of tickets purchased, and controls for the presence of a partner, number of children and its square, age and its square, education, and a full set of codegroup fixed effects. Standard errors clustered on postcodes. Total winnings (after tax) are measured in euros/10000 and include BMW values. "Got Car" means that one of the household's currently-owned cars was acquired in the six months since the lottery date. See previous tables for variable definitions. Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1

F-Tests fail to reject equality of the Lottery Winnings and NonLottery Income coefficients for all consumption items but two: Happiness ($p=.0097$) and Charity ($p=.0910$).