NBER WORKING PAPER SERIES

FORTUNATE FAMILIES?
THE EFFECTS OF WEALTH ON MARRIAGE AND FERTILITY

David Cesarini
Erik Lindqvist
Robert Östling
Anastasia Terskaya

Working Paper 31039
http://www.nber.org/papers/w31039

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2023

The study was supported by the Swedish Research Council (421-2011-2139 and 2022-02686), the Handelsbanken Research Foundation (P2011:0032:1), Riksbankens Jubileumsfond (P15-0514:1) and the Ragnar Söderberg Foundation (E4/17). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Fortunate Families? The Effects of Wealth on Marriage and Fertility
David Cesarini, Erik Lindqvist, Robert Östling, and Anastasia Terskaya
NBER Working Paper No. 31039
March 2023
JEL No. D01,J12,J13

ABSTRACT

We estimate the effects of large, positive wealth shocks on marriage and fertility in a sample of Swedish lottery players. For male winners, wealth increases marriage formation and reduces divorce risk, suggesting wealth increases men's attractiveness as prospective and current partners. Wealth also increases male fertility. The only discernible effect on female winners is that wealth increases their short-run (but not long-run) divorce risk. Our results for divorce are consistent with a model where the wealthier spouse retains most of his/her wealth in divorce. In support of this assumption, we show divorce settlements in Sweden often favor the richer spouse.

David Cesarini
Department of Economics
New York University
19 West 4th Street, Room 511
New York, NY 10012
and NBER
dac12@nyu.edu

Erik Lindqvist
Swedish Institute for Social Research (SOFI)
Stockholm University
SE-106 91 Stockholm
SWEDEN
erik.lindqvist@sofi.su.se

Robert Östling
Department of Economics
Stockholm School of Economics
P.O. Box 6501
SE-113 83
Stockholm
Sweden
robert.ostling@hhs.se

Anastasia Terskaya
Department of Economics and IEB
Universidad de Barcelona
Carrer de John Maynard Keynes, 1, 11
Barcelona 08034
Spain
anastasiaterskaya@gmail.com

A data appendix is available at http://www.nber.org/data-appendix/w31039

1 Introduction

In a series of landmark papers, Becker famously argued that the basic tools of consumer theory could be used to understand choices in domains traditionally assumed to lie outside the scope of economic theory, such as fertility (Becker, 1960) and marriage (Becker, 1973, 1974; Becker et al., 1977). Becker argued that price and income changes were essential for understanding many of the findings from the observational literature. While his work helped spawn an enormous literature on family economics Hotz et al. (1997), it has proven notoriously challenging to stringently test some of the core predictions of the models in question. The basic difficulty is that the models make quantitative predictions about the effects of exogenous changes in prices, wages and unearned income. It is widely understood that identifying such variation in observational data remains a major challenge for efforts to design and implement credible tests of theoretical predictions (or credibly pinning down key elasticities). For example, Hotz et al. (1997) conclude a comprehensive review of economic models of fertility by remarking that “theory and econometric methods are...much better developed than the empirical literature” and proceed to call for intensified efforts to identify “plausibly exogenous variation in proxies for the price and income concepts appearing in the theories”. Lundberg and Pollak (1996) and Burstein (2007) emphasize that similar difficulties arise when testing economic models of marriage and divorce.

In this paper, we seek to make progress on the question of causality by providing credible estimates of the direction and magnitude of the long-run effects of unearned wealth on marital and fertility outcomes. Specifically, we leverage the randomized assignment of lottery prizes to analyze how large, positive wealth shocks impact marriage formation, divorce risk and fertility up to 10 years after a lottery windfall. Methodologically, our work is most closely related to previous studies that, following Imbens et al. (2011), leverage lottery wins to try to estimate the causal effects of unearned wealth. A small subset of these have considered one or more outcomes resembling the family-formation variables that we focus on here (Bulman et al., 2022; Tsai et al., 2022; Golosov et al., 2023; Bleakley and Ferrie, 2016; Boertien, 2012; Hankins and Hoekstra, 2011). The paper most closely related to ours is a a contemporaneous study of American lottery players by

Bulman et al. (2022) which reports a rich and credible set of analyses of how unearned wealth impacts marriage and fertility outcomes measured up to five years after the event.

More generally, our study is broadly related to a quasi-experimental literature concerned with estimating income effects on fertility and marriage outcomes. Previous studies have relied on a range of identification strategies to try to isolate plausibly exogenous variation in income. These include leveraging income loss from job displacement (Lindo, 2010; Amialchuk, 2013; Huttunen and Kellokumpu, 2016), asset price fluctuations (Schultz, 1985; Black et al., 2013; Lovenheim and Mumford, 2013; Klein, 2017), gender-specific components of labor demand shocks (Schaller, 2016; Kearney and Wilson, 2018; Autor et al., 2019), changes in income taxes (Groeneveld et al., 1980; Cain and Wissoker, 1990; Weiss and Willis, 1997; Baughman and Dickert-Conlin, 2003, 2009; Azmat and González, 2010), or to the welfare system (Moffitt, 1990; Rosenzweig, 1999; Hu, 2003; Bitler et al., 2004; Cohen et al., 2013; Berniell et al., 2020; Yonzan et al., 2020).

Our study has several methodological strengths that allow us to go beyond this earlier work. A first is that our data allow us to classify players into groups within which the prize amount was randomly assigned under the rules of the lottery, effectively replicating the conditions of a randomized control trial. Our subsequent causal inferences are based exclusively on comparisons of the outcomes of players who were in the same group but were randomly assigned different prize amounts. Second, the prize pool is almost $265 million, allowing us to estimate treatment effects with high precision, both in the overall sample and in many interesting subsamples. Third, the rich registry data allows us to observe outcomes realized many years after the lottery, in a sample virtually free of attrition. A systematic and comprehensive comparison of our findings to those of previous lottery studies of family-formation outcomes identifies several dimensions along which our work compares favorably and helps advance the literature in substantively important ways. A final contribution is that we try to go beyond earlier work in designing and reporting numerous analyses to identify potential mechanisms underlying our findings and inform theories of the family.

In our main analyses, we study marriage (for players unmarried at the time of winning), divorce (for married players) and fertility (for all players, regardless of marital status) over three event windows: the short run (up to two years after the lottery), the medium run (five years) and the long run (ten years). We conduct our analyses in a sample of female and male players, and in the combined sample. In our pooled sample, we find that lottery wealth increases the marriage probability and fertility rate in the short and median run, while the effect on divorce risk is statistically insignificant. Our separate analyses of male and female winners suggests that the effects in the pooled sample often mask interesting heterogeneity by gender. Among unmarried men, our point estimate suggests that a one-million SEK lottery win (≈\$140,000) increases the probability of marriage within a five-year horizon by 4.7 percentage points, a 30% increase. We also report results which suggest that when married men win the lottery, the windfall has a tendency to stabilize their existing marriages. According to our estimates, the risk of a divorce event in the ten years following the lottery falls by 6.0 percentage points in this sample, a 40% reduction. Lottery wealth also increases male fertility at all time horizons. Ten years after the lottery draw, male winners have 0.056 children more per million SEK, equal to a 13.5 percent increase in fertility during this period. In contrast, the only exception to the pattern of null results for female winners is that lottery wealth almost doubles their short-run probability of divorce. One interpretation of the absence of a discernible long-run increase in divorce risk among these women is that wealth accelerates the dissolution of marriages that were already underway.

Both in our combined and sex-stratified analyses, a striking finding is that the sign and magnitude of our estimated treatment effects often track the cross-sectional income gradients for the outcomes surprisingly closely. For example, Our results on marriage formation are consistent with the literature on gender differences in partner selection (Fisman et al., 2006; Bertrand et al., 2015; Almås et al., 2023), which has found that wealth appears to improve the marriage market prospects of men more than women, on average. The "gendered" treatment effects we report for divorce are directionally consistent with evidence that the association between husbands' and wives' incomes and divorce risk have opposite signs (Berniell et al., 2020; Boheim and Ermisch, 2001; Burstein, 2007; Doiron and Mendolia, 2012; Folke and Rickne, 2020; Killewald, 2016; Weiss and Willis, 1997). Our follow-up analyses suggest the improvements are not uniform over the income distribution: lottery wins cause the largest increase in marriage rates and the biggest reduction in divorce rates among unmarried men with low incomes.

To provide some further insight into potential mechanisms underlying the gendered effect of wealth on divorce, we examine under what conditions our results align with economic theories of marriage. The predictions of divorce-threat models of marriage often hinge critically on assumptions about how wealth is split in the event of divorce. If wealth is split equally between spouses, each spouse's single-state utility is independent of the identity of the winning spouse. Gender differences in the effects of lottery wealth can then only arise if the utility of remaining married depends on whether the husband or wife wins the lottery. Such differential impacts could arise for a variety of reasons, for example because there is a strong social norm idealizing a male breadwinner (Bertrand et al., 2015). If the winner instead retains most of the lottery prize in divorce, an alternative explanation for the gendered divorce pattern is that the wife has greater marginal utility of consumption in the single-state than in the married state, and the husband greater marginal utility of consumption in a married state than in the single state. Intuitively, a lottery win may give a discontent wife economic opportunity to leave the marriage, while men use the prize money in a way that increases the gains from marriage.

To make some progress on distinguishing between these broad classes of explanations, we conduct a number of follow-up analyses. Under Swedish marital law, the default rule is that all assets are shared equally between spouses in the event of divorce. We show that actual divorce settlements in Sweden are, in fact, often deviate from this default. While the richer spouse redistributes some assets to the poorer spouse in the typical settlement, the amount of redistribution is smaller than the amount required to ensure an equal split of the assets. This empirical finding suggests that the spouses in our samples may anticipate retaining a larger share of the prize money in the event of divorce. Such non-equal splits probably reflect a combination of social norms, nuptial agreements and a stronger bargaining position in divorce settlements for the richer spouse.

There are additional patterns in the data which are easier to reconcile in a world of non-equal splits and gender-based preferences. First, the winning spouse retains most of the lottery wealth within marriage and increases consumption of leisure more (Cesarini et al., 2017). Because a stronger outside option also implies greater bargaining power also inside the marriage, larger consumption increases for the winning spouse follows naturally from a model with non-equal divorce splits. Other explanations appear more far-fetched.

For example, a model with a male breadwinner norm and equal divorce-splits generates a larger increase in the winner's consumption only if the wife espouses the breadwinner norm more strongly than the husband. Second, the spike in divorces when wives win the lottery is only present for couples where the woman earns relatively little, consistent with the notion some women find the financial means to pursue a life as singles if they win the lottery (but not if their husbands do).

Our results also inform a debate about how fertility choices depend on income. Following Becker (1960), children are often introduced into economic models as normal durable goods. Studies on the cross-sectional relationship between income and fertility have reached different conclusions about both the sign and magnitude of any association (Anderson, 2008; Black et al., 2013; Kolk, 2019; Jones and Tertilt, 2008). Overall, our results suggest children are a normal good, even though income effects are not large, and may be stronger for men than for women. Because wealth increases men's marriage rate and lowers their divorce rate, the effect on male fertility might be partly mediated by a higher marriage rate (a back-of-the-envelope calculation suggests 20-40% of the effect). Suppose, further, that the reason wealth impacts men's marriage rate is that wealth makes men more attractive as spouses compared to other men. Then it may be misleading to extrapolate our estimates to a setting where everyone enjoys a large windfall gain. Under such a scenario, any effect of wealth on attractiveness is likely to be much smaller, and the wealth shock may have very little impact on marital choices that in turn could affect fertility.

The article is structured as follows. Section 2 reviews the lottery and the register data and describes how our estimation sample was constructed. Section 3 describes our identification strategy and provides evidence in support of our key identifying assumption. In Section 4 we report the results from our main analyses and compare our estimates to cross-sectional gradients and other studies using lottery data. Section 6 discusses how our results fit into previous theoretical and empirical literature. Section 7 concludes.

2 Data

Following Cesarini et al. (2016; 2017), we construct our estimation sample by matching three samples of lottery players, and their family members, to population-wide registers with annual information about labor market outcomes and demographic characteristics. Our basic identification strategy leverages our institutional knowledge of how prizes were awarded in each lottery to assign players to groups within which the lottery prize amounts are randomly assigned under the rules of the lottery. The group construction is almost identical to Cesarini et al. (2016), but to clarify and motivate our approach we provide a summary overview of the lottery data below.[1] We then explain the process by which we arrived at our final estimation sample and provide some summary information about the distribution of lottery prizes in this sample.

2.1 Kombi

Our first sample consists of about half a million individuals who participated in a monthly ticket-subscription lottery called Kombilotteriet ("Kombi"). The proceeds from Kombi go to the Swedish Social Democratic Party, Sweden's main political party during the post-war era. Subscribers choose their desired number of subscription tickets and are billed monthly. Our data set contains information about all draws conducted between 1998 and 2011. For each subscriber, the data contain information about the number of tickets held in each draw and information about prizes exceeding a million SEK. The Kombi rules are simple: two individuals who purchase the same number of tickets in a given draw face the same probability of winning a large prize. To construct the Kombi group identifiers, we therefore match each large-prize winner to (up to) 100 non-winning players who held the same number of tickets in the month of the draw and are similar to the winner in sex and age. The winning player and the non-winning matched "controls" are then assigned to the same group.

[1] Our procedure for generating the group identifiers is identical to that in Cesarini et al. (2016), except that all main analyses in this paper are restricted to a sample of individuals aged 18-44 at the time of the lottery event. A detailed description of the institutional features of the lottery samples, the primary sources of lottery data processing, data quality control, and group identifiers construction is provided in the Online Appendix to Cesarini et al. (2016).

2.2 Triss

Triss is a scratch-ticket lottery run by Svenska Spel, the Swedish government-owned gaming operator. Our Triss data set consists of winners of two types of prizes, here denoted Triss-Lumpsum and Triss-Monthly. Both types of winners were invited to a daily TV show. At the show, winners of Triss-Lumpsum (1994-2011) draw a prize ranging from 50,000 to 5 million SEK. Triss-Monthly winners (1997-2011) participate in the same TV show, but instead win a monthly installment where the size (10,000 to 50,000 SEK) and duration (10 to 50 years) are determined by two separately drawn tickets. We convert the Triss-Monthly prizes to present value by using a 2 percent annual discount rate.

Since we do not have information about non-winning Triss players, we rely on an identification strategy that compares players who won prizes of different magnitudes for reasons that are plausibly random. In Triss-Lumpsum, two players share a group identifier if and only if they won exactly one lump-sum prize in the same year and under the same prize plan. In Triss-Monthly, two players share a group identifier if and only if they won exactly one prize paid in monthly installments in the same year and under the same prize plan. We exclude a small number of players who won two prizes of the same type in a single year and under the same prize plan, since it is difficult to identify suitable controls for such winners. We also exclude shared prizes.

2.3 PLS

Our final lottery sample contains information about Swedish individuals with "prize-linked savings" (PLS) accounts. These accounts are automatically enrolled into lotteries where they have a chance of winning substantial monetary prizes. In most systems, banks fund the lottery prizes by reducing the interest paid on PLS account balances: PLS account owners effectively choose to forego some interest in return for the thrill of participating in the lotteries. The Swedish system is unusual in that the lottery prizes were subsidized by the government, making it lucrative to hold PLS accounts.

Under the Swedish PLS system, each account was assigned one ticket per 100 SEK in account balance, and automatically enrolled in a lottery that was held monthly for most of our study period. Each lottery ticket had the same chance of winning a prize, so a higher account balance increased the chance of winning. PLS account holders could win two types of prizes: fixed prizes and odds prizes. Fixed prizes were regular lump-sum lottery prizes that varied between 1000 and 2 million SEK. Odds prizes, on the other hand, paid a multiple of 1, 10, or 100 times the account balance to the winner (with the prize capped at 1 million SEK during most of the sample period). To define our PLS groups, we rely on one approach for fixed prizes and another one for odds prizes.

For fixed-prize winners, our identification strategy exploits the fact that the total prize amount is independent of the account balance among players who won the same number of fixed prizes in a given draw. We therefore assign two winners to the same group if they won the same number of fixed prizes in the same draw. Notice that this procedure excludes non-winning accounts from the sample. In practice, this is of little consequence, since the overwhelming majority of fixed prizes won are tiny compared to the largest prizes. Therefore, our approach effectively boils down to selecting our control group for winners from a sample of players who won very small prizes, rather than players who did not win any prizes at all. The main advantage of constructing a control group from small-prize winners rather than non-winners is that our information about prize-winning accounts is of very high quality. For the period 1986 through 1994, we observe all prize-winning accounts and can identify the owner of over 98% of them.

Among odds-prizes winners, there is no compelling justification for assuming that the total prize amount is randomly assigned among players who won the same number of odds prizes. On the contrary, since accounts with larger balances will, on average, win larger odds prizes, it is plausible that the owner of the account with larger balance differs along unobservable dimensions from the owner of the account with lower balance. For odds-prizes, we therefore proceed by matching each individual who won exactly one odds prize to an individual who also won exactly one prize (odds or fixed) in the same draw and whose account balances was similar to the winner's. A fixed-prize winner who is successfully matched to an odds-prize winner is hence assigned to the new odds-prize group and removed from the original fixed-prize group.

2.4 Estimation Sample

In this section, we summarize the key steps through which we arrived at our final estimation sample. We begin by pooling our three lottery datasets and matching players in this pooled sample to population-wide registers. This matching was performed by Statistics Sweden, using information about players' personal identification numbers (PINs) supplied by us. A complete description of how we processed the lottery data shared with Statistics Sweden and the quality-control filters we subsequently applied to the data set they delivered to us is available in the Online Appendix of Cesarini et al. (2016).

In the merged data set, the unit of observation is a lottery event: a player who wins on two occasions will therefore appear twice as two separate observations. For each lottery event, we have detailed information about any prize won (date, amount, type of prize, etc) and the characteristics of the lottery player. Since changes in marital status and fertility are strongly age-dependent (see Figure A.1), we restrict our main analyses to individuals aged 18-44 at the time of the lottery. Imposing the age restriction leaves 86,768 observations. Next, we drop observations with a missing marital status in the year prior to the lottery, which leaves us with 86,180 observations – 84,015 PLS prizes, 2,131 Triss prizes and 34 Kombi prizes. We then proceed by creating group identifiers for PLS and Triss observations in this sample, according to the previously described procedures.

In PLS, we drop observations assigned to groups with zero within-group variance in the magnitude of the prize (i.e. all players won the same prize amount) and observations assigned to an odds-prize group in which the total prize pool is below 100,000 SEK. These restrictions leave us with 83,199 PLS observations assigned to 543 groups (164 and 379 fixed-prize and odds-prize groups, respectively). In Triss, we drop 277 prizes where we have reason to believe the winning ticket was jointly owned or the information on prize share was missing, and one group with only one observation, leaving 1,854 Triss observations assigned to 38 distinct groups (19 for each type of lottery).

Next, we augment the data set with controls for each of the large-prize winners in Kombi. To select each winner's controls, we first identified all non-winning Kombi players who had purchased the same number of tickets as the winner in the month of win. For most winners, the number of exact matches exceeded 100, so we chose the 100 most

Table 1: Overview of Identification Strategy

				Group Identifiers	
Lottery	Period	Type	N_p	Construction	N_G
PLS Fixed	1986-2003	Lumpsum	80,253	Draw × # Fixed Prizes	164
PLS Odds	1986-1994	Lumpsum	2,617	Draw × # Prizes × # Tickets	379
Kombi	1998-2011	Lumpsum	3,410	Draw × # Tickets	33
Triss-Lumpsum	1994-2011	Lumpsum	1,580	Year × Prize Plan	19
Triss-Monthly	1997-2011	Monthly	266	Year × Prize Plan	19

Notes. This table provides a summary overview of our identification strategy which assigns players to different groups defined by group identifiers. Groups are defined so that the magnitude of each prize won is randomly assigned within groups under the lottery. We assign two players to the same group identifier if and only if they share the characteristics listed in the column labelled Construction. N_p is number of prizes and N_G is number of unique group identifiers in our main estimation sample composed of players aged 18 through 44 at the time of the lottery. PLS Fixed and PLS Odds are for players assigned to PLS fixed groups and to PLS odds groups.

similar in sex and age as the controls. When the number of exact matches was below 100, we retained them all as controls. Adding all Kombi controls to the sample expands the number of Kombi observations in the sample from 34 to 3,426 (34 large-prize winners and 3,392 controls) assigned to 33 groups. In one draw, two large-prize winners had identical ticket balances and therefore we assigned them to the same group. Each of 33 players were matched to 100 controls, one player was matched to 96 controls, and 4 controls (each from different cell) were eliminated because of missing pre-lottery marital status. These restrictions leave us with 88,479 observations.

Finally, we eliminate 353 observations where all outcome variables are missing. Table 1 reports the number of observations, the number of groups, and the study period for each lottery sample. Our final estimation sample consists of 88,126 observations – 82,879 from PLS, 1,846 from Triss and 3,410 from Kombi. These observations correspond to 76,866 unique individuals.

2.5 Lottery Prize Distribution

To help interpret our treatment-effect parameter, Table 2 provides some basic information about the prize distribution in our final estimation sample, both overall and separately by lottery. The prizes are expressed in 2010 SEK (deflated by the consumer price index) and net of taxes. To convey a sense of the magnitudes, the median annual disposable income of an adult Swede in 2010 was 178,100 SEK.

The bottom row of Table 2 reports the share of identifying variation contributed by each lottery to our final sample. Triss-Lumpsum and Triss-Monthly contribute the largest share of identifying variation (39.7% and 34.4%, respectively), followed by PLS fixed and odds prizes (10.6% and 12.3%). Kombi prizes contribute the lowest share of identifying variation. Clearly, the share of observations contributed to the final sample by each lottery is a very misleading estimate of the lottery's contribution to the final estimates. A related observation is that even though most prizes are small – Kombi controls and PLS small-prize winners (<10,000 SEK) account for 90% of the observations – it does not follow that our estimates are mostly informative about the treatment effects of small or modest changes in wealth. On the contrary, our estimates as primarily informative about the causal effects of positive windfall gains that are large, even from a life-cycle point. When we drop the 156 prizes above 2.5 million SEK, the total amount of identifying variation falls by 35% even though the sample size only drops by 0.2%.

Table 2: Distribution of Prizes

	(1)	(2)	(3)	(4)	(5)	(6)
				Triss...		
	PLS Fixed	PLS Odds	Kombi	Lumpsum	Monthly	All
0	0	0	3,376	0	0	3,376
1K to 10K	74,034	2,005	0	0	0	76,039
10K to 100K	5,737	222	0	527	0	6,486
100K to 500K	398	222	0	923	0	1,543
500K to 1M	18	52	2	65	0	137
1M to 2.5M	66	116	31	32	144	389
2.5M to 5M	0	0	1	19	90	110
>5M	0	0	0	14	32	46
N	80,253	2,617	3,410	1,580	266	88,126
Prize Sum (100M SEK)	3.68	2.69	0.41	4.38	7.68	18.84
% Treatment Variation	10.62	12.26	3.06	39.71	34.35	100.00

Notes. This table reports the distribution of lottery prizes for the pooled sample and the lottery subsamples. A lottery's share of treatment variation is calculated in a two-step process. First, we subtract each prize by its group-level mean and square the demeaned variable. Second, we calculate the sum of the squared variables for the lottery and divided by the sum of squares for all lotteries. *PLS Fixed* and *PLS Odds* are for players assigned to *PLS fixed* groups and to *PLS odds* groups.

3 Empirical Framework

3.1 Estimating Equation

We estimate the effect of unearned wealth by ordinary least squares (OLS), using the following estimating equation:

$$Y_{i,t} \quad = \quad \beta_0 + \beta_t L_{i,0} + X_{i,0}\delta_t + Z_{i,-1}\gamma_t + \epsilon_{i,t}, \tag{1}$$

where $Y_{i,t}$ is an outcome of player i measured t periods after the win. $L_{i,0}$ denotes the lottery prize won (measured in units of year-2010 millions of SEK), and $X_{i,0}$ is a vector of group identifiers. The term $Z_{i,-1}$ is a vector of baseline characteristics measured at year-end in $t = -1$. It includes indicator variables for (i) sex, (ii) college completion, (iii) being born in a Nordic country (iv) marital status as well as a third-order degree polynomial in age-at-win and a discrete numerical variable equal to i's number of children (see Section B of the Appendix for additional details).

Our key identifying assumption is that among players assigned the same group identifier, $L_{0,t}$ is independent of potential outcomes. If this assumption holds, we expect prize amount to be uncorrelated with $Z_{i,-1}$ conditional on the group-identifier fixed effects, and we expect OLS estimates of the β_t parameter to be unbiased both with and without the controls in $Z_{i,-1}$. We control for the pre-lottery characteristics since they absorb some of the residual variance, improving the precision of our estimates. In our main analyses we estimate the effect of the lottery wealth shock on fertility, marriage, and divorce $t = 2, 5, 10$ years after the lottery. We refer to these event windows as the short ($t = 2$), medium ($t = 5$) and long run ($t = 10$).

3.2 Outcomes and Estimation Sample

Our outcome variables are derived from data in two government registers: the *Longitudinal Integrated Database for Health Insurance and Labour Market Studies* (usually referred to by its Swedish acronym, "LISA") and the *Total Population Registry* ("RTB") (Statistics Sweden, 2017). These registries contain annual information about the three domains of family-formation that are the focus of this paper: marriage, divorce and fertility. For each domain, we analyze three outcome variables, one for each of the event horizons.

To analyze marriage formation, we generate three binary variables, one for each event horizon, in the subsample of players who were unmarried at year-end in $t = -1$. In this subsample ($N = 54,020$), each variable takes the value 1 if there is at least one marriage event recorded for the player in question over the relevant time horizon (and 0 otherwise). The variable is set to missing if marital status is missing at least once over the relevant time horizon.[2] Table A.1 in the Appendix shows that the short-run ($t = 2$) marriage rate is 9% among previously unmarried players and rises to 27% in the long run ($t = 10$). Our divorce outcomes are defined analogously, except that the event of interest now a marital dissolution and that the outcomes are defined only for players who were married at year-end in $t = -1$ ($N = 34,106$). Among previously married players, the probability of a divorce is 4% in the short run, rising to 14% in the long run. Finally, we analyze fertility by calculating the total number of post-lottery children born within 2, 5 and 10 years of the lottery event in the full estimation sample. We classify a child as "post-lottery" if the month of conception is the month of the lottery event or a later month. The variable is set to missing only for individuals who were not registered in Sweden in t and in all posterior years because, in this case, it is possible that some children born t years after the lottery event will not appear in the *Total Population Registry*. Table A.1 in the Appendix shows that the average number of children born within 2, 5 and 10 years of the lottery event is, respectively, 0.07, 0.21 and 0.40 in our sample.

Table A.1 shows how the size of the estimation sample varies for each of our nine outcomes. For instance, for long-run fertility outcomes, the estimation sample is marginally diminished ($N = 86,109$), due to a small number of players for whom we do not have ten complete years of post-lottery fertility information and/or who migrated from Sweden during the study period.

3.3 Inference

Throughout, we report two sets of p-values. The first are conventional p-values derived from analytical standard errors that have been clustered at the individual level (to adjust for non-independence that could potentially arise across observations when a

[2] An individual's marital status is missing if they are not registered in Sweden in a given year.

player appears multiple times). We refer to these as analytical p-values. We also report permutation-based p-values constructed by simulating the distribution of the relevant test statistic under the null hypothesis that the treatment effect is zero. In the majority of cases, the two sets of p-values are close in magnitude. In the instances that we observe meaningful differences, it is typically the case that the estimation sample is small or the outcome is a binary outcome with low prevalence (or both). In such instances, we favor the permutation-based p-values, which are valid under much weaker assumptions. In particular, analytical standard errors rely on asymptotic approximations of the sampling distribution that can be misleading, especially in small samples.[3] Appendix Section D.1 contains a detailed description of how we generate the permutation-based p-values.

In our primary analyses, we analyze the effect of unearned wealth on three outcomes, measured over three event horizons, in three different samples (men only, women only and both sexes pooled). To address the concern that we conducted many statistical tests, we report a false discovery rate (FDR) adjusted p-value for each of the $3 \times 3 \times 3 = 27$ null hypotheses of a zero treatment effect. Our method for calculating FDR-adjusted p-values follows the two-stage procedure proposed by Benjamini et al. (2006). Briefly, the procedure ensures the expected proportion of rejections that are incorrect – the false-discovery rate (FDR) – is bounded above by q. We follow the convention of setting $q = 0.10$ (Efron, 2010). In other words, the decision rule we adopt to declare each of the 27 individual hypotheses tested to be either significant or insignificant ensures that the (expected) proportion of significant test results that correctly reject the null is at least 90%. Appendix Section D.2 contains a detailed description of our procedures for multiple-hypothesis adjustment.

[3]The results of the Monte Carlo simulations described on pp. 93-96 in the Online Appendix of Cesarini et al. (2016) suggest that for binary outcomes, meaningful biases in analytical p-values are a concern in small estimation samples where the dependent variables has low prevalence. A recent work of Young (2019), that relies on a sample of approximately 50 experimental papers, shows that permutation tests find 13 to 22 percent fewer statistically significant results than are found using conventional methods.

3.4 Generalizability

Internal Validity

Our study's key identifying assumption is that within the groups described in Table 1, prize amounts are independent of potential outcomes. An implication of this assumption is that demographic characteristics determined before the lottery should not predict, individually or jointly, within-group differences in the lottery outcome. To test this assumption, we regressed the lottery outcome on group-identifier fixed effects and several pre-lottery characteristics measured at year-end in $t = -1$. The results of these balancing tests are shown in Columns (1-5) of Table A.2 in the Appendix. Consistent with our identifying assumption, we find none of the baseline characteristics are significantly associated with prize amount, either individually or jointly, once we condition on the group-identifier fixed effects. Columns (1) through (4) show that this conclusion holds in each of the individual lotteries and Column (5) shows that it holds in the combined sample. Column (6) shows that without group-identifier fixed effects, it *is* possible to predict prize amounts from pre-lottery characteristics. The fact that we are able to construct and use group identifiers in our analyses is thus more than a theoretical curiosity. The results in Column (6) suggest that the group identifiers matter in practice.

External Validity

A common concern about studies of lottery players is that individuals who choose to participate in lotteries differ in important ways from the general population, making it difficult to generalize any findings to the general population. Even though the concern itself often derives from inaccurate beliefs about the type of individuals who tend to participate in lotteries (Kaplan, 1987), it may have merit nonetheless. In designing our study, we took a number of steps to address the concern as comprehensively as possible.

In Table 3 we compare the distribution of baseline characteristics in each lottery sample, and the pooled lottery sample, to a representative sample that has been reweighted to match the sex- and age- distribution of our pooled sample. Baseline characteristics are always measured the year before the lottery. Since most of the identifying variation comes from large prize winners, we report summary statistics weighted by the prize amount won. In general, the differences are small. The overall proportion of players who are married (35%) in our pooled sample is similar to the population proportion (38%), as is the fraction of players attended college (25% compared to 29%). Finally, the average player's number of children in the year of win is similar to the population average (1.2 compared to 1.3).

A different type of concern is that the effect of lottery wealth may not be informative about other types of shocks to wealth or income. However, our previous work on Swedish lottery winners give little reason to think lottery prizes differ significantly from other types of wealth shocks. Winners spend their prize money over an extended period of time (Cesarini et al., 2016), are more satisfied with their personal finances even a decade after winning (Lindqvist et al., 2020), and mainly invest in safe assets (Briggs et al., 2021). In line with a standard model, winning the lottery leads to an immediate but modest reduction in labor supply that is quite stable over time (Cesarini et al., 2017). We further estimate a positive but statistically insignificant effect on self-assessed mental health (Lindqvist et al., 2020), a modest reduction in consumption of prescriptions drugs related to mental health (Cesarini et al., 2016) and no statistically significant effect on alcohol consumption (Östling et al., 2020).

Table 3: Summary Statistics for Baseline Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
			Triss...		Pooled	Reweighted
	PLS	Kombi	Lumpsum	Monthly	Lottery Sample	Representative Sample
Age-at-Win	34.38	37.16	33.83	34.46	34.35	34.35
1 if Female	0.49	0.29	0.47	0.45	0.46	0.46
1 if Nordic Born	0.97	0.95	0.89	0.92	0.93	0.92
# Children	1.13	0.99	1.36	1.11	1.17	1.30
1 if College	0.25	0.27	0.24	0.24	0.25	0.29
1 if Married	0.44	0.28	0.32	0.30	0.35	0.38
N	82,870	3,410	1,580	266	88,126	844,443

Notes. Sample averages are reported. Baseline characteristics are measured the year before the lottery event. *# Children* is the number of pre-lottery children (children born or conceived before the lottery event). Summary statistics shown in Columns (1-5) correspond to the lottery sample weighted by the prize amount won. The summary statistics shown in Column (6) correspond to a representative sample of Swedish adults aged 18-44, reweighted to match the sex- and age-distribution of the combined lottery sample reported in Column (5).

4 Results

We turn now to our results. We begin by reporting estimates from our core analyses of how windfall gains impact each of our outcomes over our three time horizons, before we turn to analyses related to the robustness of our findings, heterogeneity and additional outcomes.

4.1 Main Findings

The results from our main analyses of how wealth impacts marriage, divorce and fertility in the short ($t = 2$), medium ($t = 5$) and long run ($t = 10$) are summarized in Figure 1, which depicts our key parameter estimates along with their 95% confidence intervals. A more detailed summary of the results for marriage, divorce and fertility are shown in panels A, B and C of Table 4. For interpretational ease, the coefficients in Panels A and B are expressed as percentage-point changes per unit of 1M SEK won, whereas the results for fertility in Panel C are scaled so that a value of 100 would imply that a 1M SEK (about $150,000) increases average post-lottery fertility by one child over the relevant time period. We also report each coefficient estimate after it has been normalized by the mean of the dependent variable (relative risk). Overall, the FDR-adjusted p-values in Table 4 show that we can reject the null hypothesis of a zero treatment in 9 out of the 27 cases (at $q = 0.10$). Below, we briefly discuss the results in each panel in greater detail.

Overall, the estimates in Panel A suggest that unmarried lottery players who unexpectedly receive a substantial windfall are more likely to get married. All nine coefficient estimates are positive, suggesting that wealth encourages marital formation. Our estimates from the pooled sample imply a 1M SEK windfall increases the probability of getting married within $t = 2, 5, 10$ years by 2.29 (SE = 1.09), 3.25 (SE = 1.20) and 2.52 (SE = 1.51) percentage points, respectively. Expressed as relative risks, these coefficients are quite sizable, corresponding to a 25% increase in the short-run probability of getting married, a 20% increase in the medium run, and an 9% increase in the long-run. Columns (2), (5) and (8) show that the estimated wealth effects in women are consistently smaller, but never statistically distinguishable from the estimates for men. Irrespective of the time horizon, a standard F-test of the null hypothesis that wealth effects are the same

Table 4: Wealth Effects on Marriage, Divorce and Fertility

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	$t = 2$			$t = 5$			$t = 10$		
	All	By Sex F	By Sex M	All	By Sex F	By Sex M	All	By Sex F	By Sex M
Panel A: Got Married by t (Unmarried at $t = -1$)									
Effect × 100	2.287	1.167	2.890	3.249	0.695	4.684	2.520	1.122	3.275
SE	(1.086)	(1.442)	(1.322)	(1.197)	(1.632)	(1.673)	(1.513)	(2.352)	(2.113)
p (analytical)	[0.035]	[0.418]	[0.029]	[0.007]	[0.670]	[0.005]	[0.096]	[0.634]	[0.121]
p (resampling)	[0.013]	[0.469]	[0.026]	[0.004]	[0.718]	[0.005]	[0.078]	[0.599]	[0.110]
p (FDR)	[0.051]	[0.633]	[0.078]	[0.045]	[0.819]	[0.045]	[0.162]	[0.736]	[0.198]
Baseline Mean	0.093	0.103	0.085	0.165	0.179	0.155	0.269	0.287	0.257
Relative Risk	0.247	0.113	0.340	0.197	0.039	0.302	0.094	0.039	0.128
N	53,805	22,636	31,169	53,191	22,378	30,813	51,867	21,863	30,004
Heterogeneity p		[0.378]			[0.088]			[0.496]	
Panel B: Got Divorced by t (Married at $t = -1$)									
Effect × 100	1.284	3.701	-1.903	-0.441	2.297	-4.084	-2.065	-0.270	-6.030
SE	(1.145)	(1.811)	(1.064)	(1.446)	(2.305)	(1.498)	(1.568)	(2.380)	(2.081)
p (analytical)	[0.262]	[0.041]	[0.074]	[0.760]	[0.319]	[0.006]	[0.188]	[0.910]	[0.004]
p (resampling)	[0.154]	[0.002]	[0.331]	[0.758]	[0.251]	[0.106]	[0.283]	[0.893]	[0.064]
p (FDR)	[0.260]	[0.045]	[0.470]	[0.819]	[0.399]	[0.198]	[0.425]	[0.894]	[0.144]
Baseline Mean	0.042	0.040	0.044	0.084	0.080	0.089	0.145	0.138	0.152
Relative Risk	0.307	0.927	-0.430	-0.052	0.286	-0.460	-0.143	-0.020	-0.396
N	33,994	18,750	15,244	33,740	18,617	15,123	33,094	18,320	14,774
Heterogeneity p		[0.008]			[0.020]			[0.068]	
Panel C: #Post-Lottery Children by t (All)									
Effect × 100	1.209	0.214	2.048	2.915	0.865	4.278	3.310	0.777	5.583
SE	(0.650)	(0.767)	(0.993)	(1.252)	(1.645)	(1.847)	(1.660)	(1.886)	(2.659)
p (analytical)	[0.063]	[0.780]	[0.039]	[0.020]	[0.599]	[0.021]	[0.046]	[0.680]	[0.036]
p (resampling)	[0.048]	[0.819]	[0.008]	[0.010]	[0.598]	[0.012]	[0.043]	[0.727]	[0.020]
p (FDR)	[0.118]	[0.851]	[0.051]	[0.051]	[0.736]	[0.051]	[0.116]	[0.819]	[0.068]
Baseline Mean	0.068	0.067	0.070	0.208	0.197	0.218	0.393	0.369	0.415
Relative Risk	0.177	0.032	0.294	0.140	0.044	0.196	0.084	0.021	0.135
N	88,113	41,539	46,574	87,635	41,319	46,316	86,109	40,700	45,409
Heterogeneity p		[0.144]			[0.168]			[0.140]	

Notes. This table reports the estimated treatment effect of lottery wealth on the probability to get married for unmarried at $t = -1$ players, to get divorced for married at $t = -1$ players, and on the number of post-lottery children by year-end of $t = 2$, $t = 5$ and $t = 10$ for the pooled lottery sample and by gender. All specification control for baseline controls measured at $t = -1$ and group-identifier fixed effects. Standard errors are clustered at the individual level. Baseline mean is defined as the mean of the dependent variable of small-prize winners ($< 10,000$ SEK). The resampling based p-values are obtained from the resampling distribution of coefficients from 1,000 Monte Carlo simulations. $p(FDR)$ correspond to the false discovery rate adjusted resampling p-values computed using (Benjamini et al., 2006) and (Anderson, 2008) procedure. The heterogeneity p-value is from a two-sided t-test of the null hypothesis that the treatment-effect parameters are identical in the subsamples. Age-at-win: 18-44.

in men and women fails to reject at the 5% level.

Panel B shows the results from our analyses of divorce, which were conducted in the subsample of players who were married at the time of the lottery. In the pooled sample, all estimates are small in magnitude and statistically indistinguishable from zero. The results from the sex-stratified analyses nevertheless yield some intriguing evidence of gender heterogeneity. In the subsample of male winners, all point estimates suggest that windfalls stabilize marriages. The $t = 10$ effect suggests a 1M SEK windfall reduces the probability of divorce by 6.03 percentage points (SE = 2.08), equivalent to a 40% reduction in relative risk. In the female subsample, the pattern is strikingly different.

Figure 1: Wealth and Family Formation

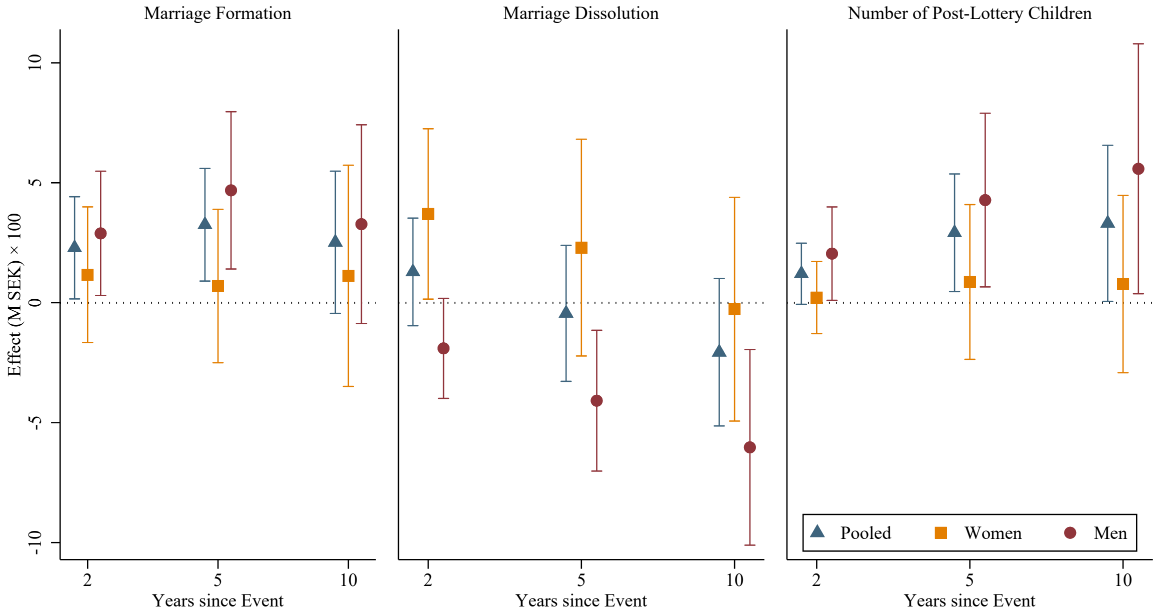

Notes. This figure reports the estimated treatment effects of a million SEK on the probability to get married for unmarried players, on the probability to get divorced for married players, and on the number of post-lottery children, all measured at year-end in 2, 5, and 10 years after the lottery. The results are reported in the pooled sample and by gender of the winner. The estimates are multiplied by 100. All specifications control for baseline controls measured at $t = -1$ and group-identifier fixed effects. Standard errors are clustered by individual, and the error bar corresponds to 95 percent analytical confidence intervals. Age-at-win: 18-44.

Whereas the long-run estimate is close to zero, married women who win the lottery are *more* likely to get divorced in the short run. In the female subsample, we estimate that a 1M SEK windfall increases short- and medium-run risks of divorce by 3.70 (SE = 1.81) and 2.30 (SE = 2.30) percentage points, respectively. Rescaling by the baseline probability (4%), our estimate suggests a 1M SEK windfall effectively doubles the short-run risk of divorce. Our results thus suggest that among married women, wealth produces a large spike in short-run divorce risk that subsequently dissipates and may be close to zero in the long run. One interpretation of this pattern is that unearned wealth accelerates the termination of marriages whose dissolution were already underway. While only two of our nine coefficient estimates for divorce are statistically distinguishable from zero, conventional F-tests reject the null that male and female coefficients are identical at $t = 2$ ($p = 0.008$), $t = 5$ ($p = 0.020$) and $t = 10$ ($p = 0.068$).

Panel C reports our results for fertility, which were conducted in the full estimation sample. For each of the three horizons, the outcome variable is the number of post-lottery children born over the relevant time horizon. The point estimates in Panel C are positive, consistent with children being a normal good. In our pooled sample, we estimate that a 1M SEK windfall increases expected number of children by 0.012 (SE = 0.007), 0.029 (SE = 0.012) and 0.033 (SE = 0.017) children in the short-, medium- and long run, respectively. Expressed as relative risks, the $t = 5$ estimate suggests that a 1M SEK windfall gain increases medium-run fertility by approximately 14%. In our sex-stratified analyses, we find evidence that children are a normal good for men. For all event windows, the estimated effects in the male subsample are positive and economically meaningful. For example, the $t = 10$ effect of a 1M SEK windfall is estimated to be 0.056 (SE = 0.027), which is equivalent to a 13.5% increase in long-run fertility. In the female subsample, the coefficients are also positive, but smaller. Since none of the three female estimates is statistically distinguishable from zero, we conclude that the evidence that children are normal goods is considerably weaker in the female subsample than in the male or pooled sample. We also note that our estimates are too imprecise to reject equal treatment effects in men and women, so the differences between men and women are only suggestive.

4.2 Sensitivity Analyses

To probe the robustness of the findings in Table 4, we first examined if our main results are sensitive to alternative sample inclusion criteria with respect to age. Our main analyses were all conducted in a sample limited to players aged 18-44 at the time of the lottery event. For our fertility outcomes, these cutoffs are well within the range typically used in the literature (e.g. Hotz et al., 1997). While the exact values are arbitrary, Panel C of Figure A.1 provides a compelling graphical justification for an upper limit somewhere in the mid 40s. The probability of having at least one more child is strongly related to age, reaching a peak around the age of 25. By age 45, it is effectively zero. Panels A and B of Figure A.1 show that for marriage and divorce outcomes, there are qualitatively similar patterns by age although marriage and divorce probabilities do not decline to zero at older age. For example, in a random sample of unmarried (married) Swedes aged between 45 and 64, we find that 10.3% (6.7%) get married (divorced) within ten years. We therefore reran our main analyses of marriage and divorce (but not fertility) in an estimation sample constructed using a more liberal age-at-win restriction of 18-64. Appendix Table A.3 reports the results, which are similar to those from our primary specification.

Second, we looked for evidence of nonlinear wealth effects by rerunning our original analyses omitting either small (<10,000 SEK for PLS winners) or large (>4 million SEK) prizes. If the marginal effects of wealth were everywhere diminishing or increasing, we expect the coefficient estimates to systematically move in opposite directions relative to the baseline estimate in Table 4. For example, under diminishing marginal returns, the specification with large prizes omitted is expected produce coefficient estimates closer to zero, whereas the specification with small prizes omitted should produce coefficients further away from zero. Under increasing returns, we expect the opposite pattern. Empirically, the results in Appendix Table A.4 provide little evidence of nonlinear effects: we fail to detect any systematic changes of coefficient estimates and the changes that we do observe are well within the range expected from sampling variation alone. Overall, the results do not provide evidence against the hypothesis that the marginal effects of wealth are roughly constant over the range of prizes in our data.

Third, we attempt to reassure that the heterogeneity by gender we document is not driven by differences in characteristics between couples where the wife played the lottery and couples where the husband played the lottery. In Tables A.5 and A.6 in the Appendix we reweight couples where the wife or the husband won in order to match the first and the second moments of pre-lottery characteristics in the pooled sample. That is, we construct the weights so that the reweighted sample of the couples where the wife won is similar to the reweighted sample of the couples where the husband won in terms of husband's and wife's age, income, the number of children, the year they played the lottery, and the type of the lottery they played (*i.e.*, Kombi, PLS, Triss). We use entropy balancing technique (Hainmueller, 2012) for the construction of the weights (the details are provided in Section E of the Appendix). The estimates of the lottery effect on the probability of marriage dissolution in the reweighted samples (reported in Table A.6 in the Appendix) confirm our main findings – we can reject equal treatment effect in the couples where wives win and where husbands win, even when the estimates are adjusted for differences in characteristics between these subsamples.

4.3 Heterogeneous Effects

We now investigate heterogeneity along dimensions other than gender. For each of our three outcomes, we analyze whether the lottery effect varies in the pooled sample by the winner's (i) pre-lottery income (above/below the sample median), (ii) age-at-win (above/below age 35) and (iii) pre-lottery parental status (no children/some children). For fertility we also test whether effects differ between married and unmarried players. We estimate equation (1) for each subsample separately and report the results in Table A.7 (marriage formation), Table A.8 (divorce) and Table A.9 (fertility) in the Appendix.

Overall, we find little evidence of heterogeneous effects, but there are a few notable exceptions. First, the positive effect on marriage formation is statistically distinguishable from zero only among low-income players. In our analyses by gender above, we only find an effects on marriage formation of men. In a further exploratory analysis, we split the sample by both income and gender. Though statistical power goes down with the number of subgroups, Figure A.2 shows the effect of lottery wealth on marriage formation is positive for low-income men ($\hat{\beta}_5 = 7.84, SE = 2.29$) and it is not statistically distinguishable from zero for high-income men and women.

Second, the effect of lottery wealth on fertility is larger for winners who are young (when we split the sample by age) and unmarried (when we split by marital status), though we cannot always reject identical effects. The stronger effect for young winners is unsurprising given that the fertility rate declines with age (see Figure A.1). The stronger effect for unmarried players suggest the fertility response is driven by new marriages of unmarried players. Figure A.2 shows the results when we further split the age- and marital-status subsamples by gender. Notably, the gender difference in the fertility response is larger for older winners, suggesting declining female fertility may attenuate the female fertility response. Figure A.2 also shows estimated effects are close to zero for women regardless of their marital status. For men, the effects on fertility are also similar between unmarried and married players.

4.4 Additional Outcomes

We consider two additional outcomes: cohabitation and spousal quality. We observe cohabitation for unmarried couples with joint children. Panel A of Table A.10 in the Appendix shows the increase in unmarried men's probability of marriage is robust to excluding cohabiting men from the sample, suggesting the effect on marriage probability is not simply due to cohabiting couples changing their relationship status. Panel B shows the effect of lottery winnings on cohabitation for "single" winners (unmarried and not cohabiting with children prior to winning). The estimated effect is positive for women and negative for men, but neither effect is statistically different from zero. Consequently, the increase in men's marriage rates likely reflects an increased propensity to marry rather than a general increase in couple formation.

We next consider whether changes in the marriage rate coincide with changes in spousal quality. To avoid selection bias, we regress indicator variables equal to one in case an individual married a spouse with certain characteristics. This way of defining the outcome variable thus combines the "extensive" margin (whether an individual is married) and the "intensive" margin (spousal quality conditional on marriage). We consider two proxies for spousal quality: age and education. Figure A.3 in the Appendix suggests the likelihood of marrying a younger spouse does not increase in lottery wealth for female and male winners. In unreported analyses, we do not find that men or women are more likely to marry a partner with a college education.

5 Benchmarking Lottery Estimates

In this section, we compare our lottery estimates to two different benchmarks: income gradients and estimates from previously published studies of lottery players.

5.1 Income Gradients

We estimate cross-sectional income gradients using a representative sample of Swedish adults that has been weighted to match the sex-, age- and event-year distribution in the lottery sample. Our reweighting procedure ensures that in the representative sample, the outcomes are defined over the same distribution of event windows as in the lottery sample. To estimate the income gradients, we estimate the equation

$$Y_{i,t} = \alpha_0 + \alpha_t I_{0,i} + Z_{i,-1}\theta_t + B_{0,i}\psi_t + u_{i,t} \tag{2}$$

where $Y_{i,t}$ is individual i's outcome measured t years after the event year, $Z_{i,-1}$ is the set of baseline covariates defined in the previous section, and $B_{0,i}$ is a vector of event-year fixed effects. To obtain a long-run measure of economic status purged of most transitory year-to-year fluctuations in income, $I_{0,i}$ is defined as the five-year average of all annual disposable incomes observed up to five years before the event year.

Our lottery estimates are not on a scale that easily permits comparisons to income gradients. To enable comparisons, we follow our previous research (Cesarini et al., 2016; Lindqvist et al., 2020) and convert the lump-sum prizes to annuity payouts, assuming a 20-year period and interest rate of 2%. For point of reference, a lump-sum prize of \$100,000 translates into an increase in permanent annual after-tax income of \$5,996.

Figure 2 and Table A.11 compares each of our annuity-rescaled lottery estimates to its corresponding income gradient. The lottery estimates are generally similar in magnitude to the income gradients. Even more strikingly, the patterns of heterogeneity by gender are often similar when we compare the experimental lottery estimates to the non-experimental gradients. For purposes of illustration, consider the gradient between short-run divorce risk and income. Panel A shows both the gradient and the lottery-estimate are positive in women and overall, but negative for men. Our comparison between lottery-based estimates and gradients thus bolster the credibility of previous studies that rely on observational data.

Figure 2: Comparison of Annuity-Rescaled Lottery Estimates to Income Gradients

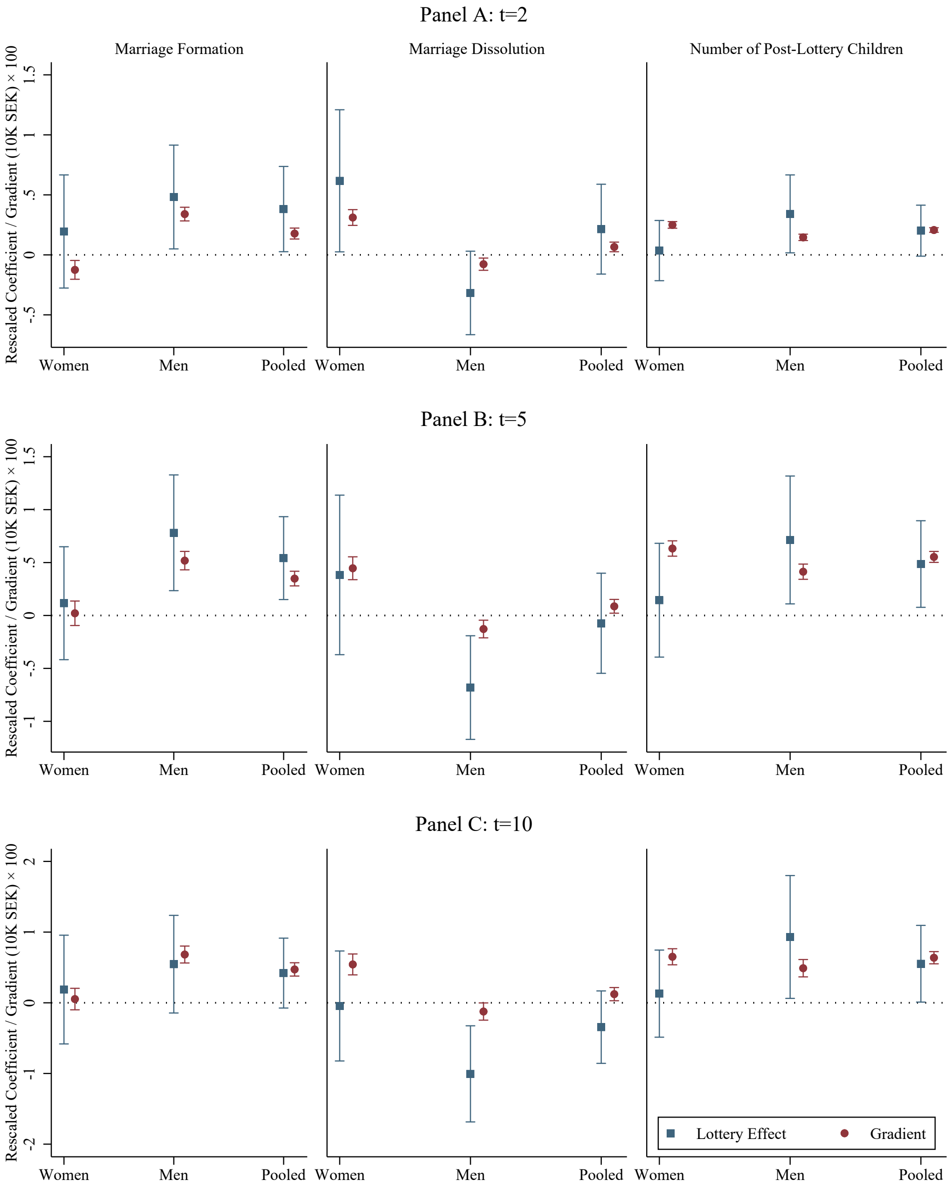

Notes. This figure reports the estimated treatment effect of annuitized lottery prize and annual disposable income gradient (in 100,000 SEK) on the probability to get married for unmarried at $t = -1$ (Column 1), on the probability to get divorced for married at $t = -1$ (Column 2), and on the number of post-lottery children (Column 3), all measured at year-end in 2 (Panel A), 5 (Panel B), and 10 (Panel C) years after the event. The results are reported in the pooled sample and by gender of the winner. The rescaled lottery effects are estimated using equation (1), and the income gradients are estimated using equation (2). The estimates are multiplied by 100. All specifications control for baseline controls measured at $t = -1$ and group-identifier fixed effects. Standard errors are clustered by individual, and the error bar corresponds to 95 percent analytical confidence intervals. Age-at-win: 18-44.

5.2 Comparison to Other Lottery Studies

We identified six quasi-experimental studies of lottery players that report estimates of wealth effects on outcomes similar to those in our primary analyses. Table A.12 provides summary information about each of these along some key dimensions, including identification strategy, sample, event-study windows, outcomes analyzed, and basic information about relevant estimates of reported in the original study.

We emphasize two dimensions along which our study compares favorably to this other work. First, our identification strategy compares the post-lottery outcomes of individuals *within* groups of ex ante identical participants from the same lottery. Within each group, we know that the magnitude of prizes won were randomly assigned under the rules of the lottery. This feature of our study distinguishes it from the remaining six studies, all of which make stronger identifying assumptions. Second, we can track our post-lottery outcomes for long periods after the win, allowing us to study effects for event-study windows as wide as a decade. This feature of our study allows us to go beyond earlier work in characterizing how wealth effects evolve over time.

Of the six studies, one analyzed lottery players in modern Taiwan (Tsai et al., 2022), one used data from the British Household Panel Survey (Boertien, 2012), one reported estimates of the effects of winning a land-lottery in early-19th century Georgia on long-run fertility (Bleakley and Ferrie, 2016), and the remaining three studied lottery winners in contemporary United States (Bulman et al., 2022; Golosov et al., 2023; Hankins and Hoekstra, 2011). Bulman et al. (2022) and Golosov et al. (2023) use multiple years of federal tax records for the entire US population to construct a panel with information about the winners of state lotteries. Since the studies use tax data from similar periods (2000-2019 and 1999-2016, respectively) the sample overlap is obviously very high. However, the papers differ substantially in their focus. While Golosov et al. (2023) analyze some marital outcomes in their full sample, their primary focus is on labor-market outcomes.

By contrast, Bulman et al. (2022) conduct a rich set of analyses of fertility and marital outcomes, many of which are directly relevant for and comparable to our study's findings. Their primary analyses are all conducted both in a full sample, and in men and women separately, applying several sample inclusion criteria that facilitate comparability with our results. For example, they impose an age restriction (25-44) similar to ours (18-44) and their paper's main specification excludes prizes greater than $500K. This restriction on prizes ensures that the identifying variation comes from prizes in a range similar to the range in our sample (Table 2). For these reasons, we generally prioritize comparisons of our results to Bulman et al. (2022) relative to those reported by Golosov et al. (2023).

The third US lottery study (Hankins and Hoekstra, 2011) analyzed data on state-lottery prizes won over a six-year period in a region of Florida. For several reasons, some of which are discussed below, we view the results in Hankins and Hoekstra (2011), and also Boertien (2012), as substantially less informative than those in the other four studies, and assign less weight to their results in our discussion below.

To facilitate comparisons across studies, we rescaled each of the originally reported effect-sizes to make them interpretable as estimates of the effect of a windfall measured in units of $100K. Below, any coefficients reported are harmonized coefficients, except where explicitly stated otherwise. See Appendix C for details on how the harmonized coefficients were generated.

Marriage

Four of the studies report estimates wealth effects on marriage (Bulman et al., 2022; Tsai et al., 2022; Golosov et al., 2023; Hankins and Hoekstra, 2011). For players unmarried at the time of win, Bulman et al. (2022) estimate that a $100,000 windfall increases the probability of being married two and five years after the lottery by 2.43 ($SE = 0.43$) and by 1.18 ($SE = 0.51$) percentage points, respectively. Since their data does not allow them to directly observe marriage or divorce events, they define an individual as married or not based on the marital status listed in the year-t tax filing (whereas we define the person as married if we observe at least one marriage event between the lottery and the end of year t). Rerunning our baseline model for $t = 2$ and $t = 5$ using Bulman *et al.*'s (2022) alternative definition of marriage, our harmonized $t = 2$ and $t = 5$ estimates were 1.67 ($SE = 0.78$) and 2.40 ($SE = 0.85$).[4] Tsai et al. (2022) also analyze wealth effects on marriage among players unmarried at the time of win. Their harmonized estimates

[4]These estimates are nearly identical to harmonized estimates from our main specification with marriage defined as described in Section 3.2, we conclude that the use of the alternative measure of marriage used by Bulman et al. (2022) is not a barrier to comparability.

for $t = 2$ and $t = 5$ respectively, are 1.16 ($SE = 0.34$) and 1.00 ($SE = 0.38$).

Overall, the results provide compelling evidence that unearned wealth increases the probability that single winners get married. For $t = 2$, the harmonized estimates range from 1.16 to 2.43, and an inverse-variance weighted average of the three coefficient estimates is 1.63, with a standard error of 0.25. For $t = 5$ the coefficients range from 1.00 to 2.40, now with an inverse-variance weighted average of 1.22 ($SE = 0.29$).

Next, we examine how Bulman *et al.*'s (2022) results broken down by gender compare to ours. In our sample, the wealth effects on marriage in the full sample appear to be primarily driven by unmarried men, especially those with low incomes. Bulman *et al.*'s (2022) harmonized estimate for $t = 5$ is 1.49 ($SE = 0.69$) among unmarried men, compared to 0.66 ($SE = 0.76$) for unmarried women. For point of reference, the analogous estimates in our sample are 3.47 ($SE = 1.15$) and 0.27 ($SE = 1.18$). Additionally, Bulman *et al.*'s (2022) follow-up analyses suggest that the effects are strongest among low-income men. Overall, Bulman *et al.*'s (2022) results on gender and marriage are thus qualitatively consistent with our main findings. However, we note that the magnitude of the estimated Swedish gender gaps are generally larger than those reported Bulman *et al.* (2022) and that both gender gaps are imprecisely estimated. Overall, a plausible interpretation of the evidence is that wealth effects on marriage may well be stronger among unmarried than among unmarried women, but that it is likely that the gender gaps we report are overestimates. An alternative possibility is that the actual gender gap are substantially larger in Sweden, but this strikes us as a less plausible first-order explanation.

Overall, the evidence from the three studies on how wealth impacts marriage rates up to five years after the lottery is congruent. There is clear and compelling evidence that the true harmonized wealth parameters are in the range 1-2 percentage points. There is also some evidence suggesting that wealth effects are larger for men. Finally, the evidence suggests that wealth effects on both unmarried men and women tend to be strongest among those with lower incomes.

The results reported in the fourth study on marriage, by Hankins and Hoekstra (2011), may seem to contradict our main conclusions about the sign and magnitude of wealth effects. The study's point estimates for marriage consistently go in the "wrong" direction, and the authors report a negative and statistically significant wealth effect on women that is highlighted in the abstract and interpreted as evidence that additional income makes single women less likely to marry. However, our examination of the results concludes that such conclusions may not be justified. In Online Appendix C.3, we show that Hankins and Hoekstra's harmonized estimates have large standard errors. Even under assumptions about the true effect sizes we consider optimistic, the study's statistical power was too low to provide meaningful information about the wealth effects of interest. Given this low power, we show that it is not uncommon, conditional on finding a statistically significant result, for the coefficient to have the "wrong" sign and for its magnitude to be implausibly large.

Divorce

Our literature review identified four studies of lottery players that estimate wealth effects on marital dissolutions (Bulman et al., 2022; Golosov et al., 2023; Boertien, 2012; Hankins and Hoekstra, 2011). Two of them (Hankins and Hoekstra, 2011; Boertien, 2012) are vulnerable to concerns about low power similar to those voiced about Hankins and Hoekstra's (2011) marriage analyses in the previous section (see also C.3). For example, Boertien (2012) uses data from the British Household Panel Survey and finds that lottery wealth causes a statistically significant increase in the divorce risk for men, but not women. However, the small number of winners in the sample, the small prize amounts (average prize £402) and the use of an empirical specification with prizes expressed in logarithms makes us skeptical that a detailed comparison would be instructive. Therefore, we focus on relating our findings to those in Bulman et al. (2022).

In our full sample of players married at the time of the lottery, our harmonized estimates are consistently small and statistically indistinguishable from zero. For example, our harmonized estimate for $t = 5$ implies a 0.32 percentage-point *reduction* ($SE = 1.04$) in the probability that a player married at win files for divorce within five years of the lottery. The evidence reported by Bulman et al. (2022) is based on analyses of how wealth impacts the probability that players remain married. Their harmonized $t = 5$ estimate suggests that players married at $t = 0$ are 0.98 percentage points ($SE = 0.48$) *less* likely to be married five years later. Even though the harmonized estimate reported has the opposite sign of ours, both estimates are close to zero and we cannot reject the null hypothesis that the parameters are identical. Overall, the evidence is consistent with wealth effects on divorce risk being small overall.

Several findings suggest that wealth may impact the divorce risk of married men and women differently. For example, our harmonized $t = 5$ estimates imply that married women are 1.87 ($SE = 1.62$) percentage points *less* likely to remain married, whereas men in our sample are 3.11 ($SE = 0.89$) percentage points *more* likely to remain married. The analogous estimates reported by Bulman et al. (2022) are -1.57 ($SE = 0.78$) for women and -0.58 ($SE = 0.78$) for men. The gender gap implied by these estimates is approximately one percentage point. While the sign of the gap matches what we find, the magnitude of the gap is substantially smaller. Finally, Bulman *et al.* (p. 21 2022) report evidence that wealth causes larger increases in divorce risk among women with low baseline earnings, a finding consistent with the results from follow-up analyses described in Section 6.2 below.

Fertility

Finally, we identified three lottery studies examining fertility: Bulman et al. (2022), Tsai et al. (2022), and Bleakley and Ferrie (2016). For ease of interpretation, the harmonized wealth effects are multiplied by a factor of 100, so that an effect of 1.00 means that a $100K windfall increases the expected total number of children over the relevant event horizon by 0.01 (or equivalently, that a windfall of $100K increases the probability of having one more child during the relevant time period by one percentage point).

Bulman et al. (2022) analyze fertility outcomes up to five years after the lottery event. While they find clear evidence of a modest increase in the likelihood of having a child shortly after the win, they find no evidence that the wealth impacts cumulative births up to five years after the lottery. Their harmonized estimate for cumulative fertility at $t = 5$ is 0.07 ($SE = 0.47$). In our sample, the analogous estimate is 2.1 ($SE = 0.90$), whereas Tsai et al. (2022) report an estimate of 1.36 ($SE = 0.40$) on the same scale. An inverse-variance weighted meta-analysis of the three coefficients yields an estimate of 0.95 ($SE = 0.29$), providing strong evidence against the null that the effect is everywhere zero.

Finally, Bleakley and Ferrie (2016) exploit a land lottery that took place in Georgia in 1832 to study, inter alia, how wealth impacts long-run fertility outcomes. The obvious differences in time and place clearly hamper comparison to work based on more recent data and we make no effort to try to generate harmonized estimates for this study. That said, the authors find that winners – who, on average, were assigned a plot of land worth the approximate equivalent of five years salaries for an unskilled laborer at the time – had 0.18 more children than nonwinners and demonstrate that the positive fertility effect persisted over the entire 18-year post-lottery period analyzed.

Overall, the weight of the evidence suggests to us that in most populations, unearned wealth has a small to modest positive effect effect on fertility.

6 Discussion

In this section, we discuss how our results fit into previous theoretical and empirical literature.

6.1 Marriage Formation

We find that lottery wealth has a positive effect on men's marriage formation, while the effect for women is close to zero. Splitting the sample by both gender and pre-win income, we further only find a positive effect for low-income men. How could these results be understood?

First, the fact that we see no effect of lottery wealth on the marriage rates of low-income women speaks against theories suggesting that young adults may delay marriage until they reach economic stability (Oppenheimer, 1988). Although classical economic models struggle to explain the documented gender heterogeneity, it aligns with the concept of hypergamy, i.e. the tendency for husbands to have a higher economic rank in the overall male population than their wives do in the overall female population (see Almås et al. 2023 for recent evidence from Norway). In the presence of hypergamy, high-income men have a larger pool of potential partners than their lower-income counterparts, which suggests that wealth may increase their probability of finding a match. Additionally, our results are consistent with earlier research findings indicating that there is a stronger relationship between income or other human-capital related factors and the propensity to marry in men than in women (Burgess et al., 2003; Ellwood and Jencks, 2004; Burstein, 2007; Bertrand et al., 2015; Killewald, 2016; Autor et al., 2019). One of the explanations for this pattern is that men and women may give different weights to different qualities when choosing a partner. This explanation fits with literature indicating that women value earning potential, intelligence, and social status of potential partner more than men do (Fisman et al. 2006; Hitsch et al. 2010b,a; Regan et al. 2000; Almås et al. 2023). A parsimonious interpretation of our results is therefore that lottery wealth increases the attractiveness of low-status men who would otherwise have a hard time finding a spouse.

6.2 Marriage Dissolution

We find that lottery wealth increases the short-run divorce probability in the couples where wives win, and reduces long-term divorce probability in couples where the husband wins. These gendered treatment effects are consistent with a large body of empirical evidence showing higher husband's earnings or employment stabilize marriages, while wife's income or employment have an opposite effect (Berniell et al., 2020; Boheim and Ermisch, 2001; Burstein, 2007; Doiron and Mendolia, 2012; Folke and Rickne, 2020; Killewald, 2016; Weiss and Willis, 1997), though there are several exceptions (Hoffman and Duncan, 1995). Yet while gender differences in the effect of earnings on marriage dissolution can be explained by the differential effect of earnings of the husband and wife on the gains from specialization according to comparative advantage, gender differences in the effect of lottery wealth cannot be explained by this mechanism.

To interpret our results, in Section F of the Appendix we considera symmetric cooperative bargaining model with singlehood as the threat point, as in Manser and Brown 1980 and McElroy and Horney 1981. A key assumption in the model is how wealth is split in a divorce settlement.[5] If lottery wealth is split equally, single-state utilities are independent of which spouse won the lottery. In this case, husbands' and wives' lottery wealth have different effects on marriage dissolution only if husbands' and wives' wealth affect married-state utilities differently. One such potential mechanism is a male

[5] In Section F.1 of the Appendix, we show the key role of wealth splits in divorce holds also in a more general theoretical framework (as in Weiss 1997) without assumptions regarding the bargaining process.

breadwinner norm (Bertrand et al., 2015). A male breadwinner norm implies gains from marriage increase when husbands win the lottery (decreasing divorce risk) but decreases when wives win (increasing divorce risk).

If the winner instead retains most of the lottery prize, an alternative explanation for the gendered divorce pattern is that the wife has greater marginal utility of consumption in the single-state than in the married state and the husband has greater marginal utility of consumption in the married state than in the single state. In this case, the gains from marriage increase in the husband's lottery wealth and decrease in the wife's lottery wealth. Intuitively, a lottery win may give a discontent wife economic opportunity to leave the marriage, while men use the prize money in a way that increases the gains from marriage.

We now turn to a discussion of which mechanism best fits patterns in the data. We first present descriptive evidence that divorce settlements in Sweden are often unequal and favor the richer spouse. We then present two additional empirical regularities. First, the winning spouse retains most of the lottery wealth and increases consumption more than the non-winning spouse. Second, both the increase in divorce risk when wives win and the decrease when the husband win are stronger in couples where the winner earns relatively little. We argue these regularities are easier to reconcile with our second mechanism, i.e. gender differences in preferences and a stronger bargaining position for the winning spouse, though other explanations (including a breadwinner norm) cannot be completely ruled out.[6]

The Division of Assets in Divorce

According to the Swedish Marriage Code, the default rule is that all assets (including inherited assets and pre-marital assets) should be split equally after divorce, unless spouses have a nuptial agreement or agree on another division of assets. As of today, 11% of Swedish married couples have a nuptial agreement.[7] In Appendix G.2, we present an investigation of a random sample of 997 Swedish nuptial agreements registered in 2013.

[6] Employing a method similar to Bertrand et al. (2015), Hederos and Stenberg (2022) find only weak evidence that Swedish couples adhere to a male breadwinner norm.

[7] The figure for nuptial agreements comes from e-mail correspondence with Henrik Bondesson of the Swedish credit reference agency *UC* on January 23rd, 2023.

Our investigation reveals one third of nuptial agreements state that all property is private property, implying there is no marital property to split in case of divorce. More than half (54%) of nuptial agreements specify that certain types of assets, typically real estate, inheritances and stocks in closely held corporations, are to be individual property. Notably, such contracts are often asymmetric in the sense that only one of the spouse's property is excluded from marital property.

Even absent a nuptial agreement, the identity of the lottery winner is likely to matter for bargaining power in divorce. Before a divorce settlement is reached, the prize money is controlled by the winning spouse unless a voluntary transfer is made, allowing the winner considerable freedom how to spend the money. And because a legal process that ensures the non-winning spouse his or her stipulated share in a divorce settlement is costly, bargaining theory suggest the winning spouse should be able to keep more than half the prize amount in divorces settled outside of court. Survey evidence from Brattström (2011) indicate only one fifth of divorcing couples in Sweden hired some kind of legal assistance to help with their divorce settlement. The survey also indicates spouses often agree on unequal divisions of assets even in the absence of nuptial agreements. Out of the thirty-four percent of respondents who said the division of assets was unequal, 53% said they "agreed to it" while 35% referred to a nuptial agreement.

In order to quantify the significance of non-equal splits, we rely on descriptive evidence from the Swedish Wealth Registry 1999-2007. We restrict the sample to couples with a total net wealth of at least 100,000 SEK, who have been a couple for at least three years, and where both partners are between 25 and 44 years old. We differentiate between three different types of couples: married with or without children, and cohabiting couples with joint children. We include cohabiting couples in our analysis because the rules governing asset division of separating cohabitants does not impose an equal split of all assets.[8] Whether divorcing married couples split assets more equally than separating cohabitants is thus an indication of the extent to which the institution of marriage impact sharing of wealth between partners.

[8]According to Swedish law, common property for cohabiting couples only include dwellings and associated property (e.g., furniture) procured for the purpose of living together. Property procured prior to cohabitation is not common property. As for married couples, cohabiting couples can regulate what is to be considered common property in a contract.

Figure 3 shows the wealth share of the wealthiest partner (at each point in time) the year prior to (panel A) and the year after (panel B) separation. Panel A shows within-couple wealth inequality prior to separation is often high. For instance, in 30% of married couples with children the wealthiest spouse holds at least 80% of total wealth, and inequality is higher for couples who are married without children or cohabit. Panel B shows measured wealth inequality is even higher after separation. Perhaps surprisingly, married couples do not seem to split their assets substantially more equally than cohabiting couples.

Figure 3: CDF of Wealth Shares of Wealthiest Partners

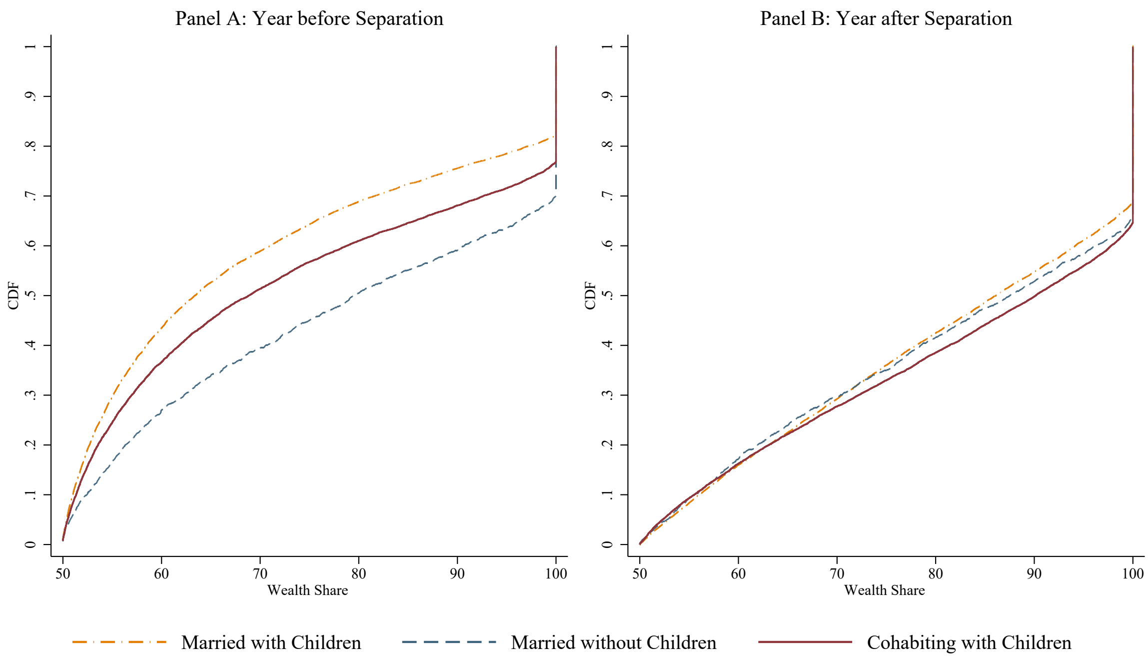

Notes. This figure shows the estimated cumulative distribution function (CDF) of the wealth share of the wealthiest partners in a couple, for the year before and after separation. The data is for couples where both partners are between the ages of 25 and 44 at the time of separation, and where the total combined wealth of the couple is above 100,000 SEK.

Because the identity of the wealthiest partner in a couple may change between Panel A (pre-separation) and B (post-separation), Figure 3 does not rule out redistribution of wealth from the wealthier partner at the time of separation. To estimate the extent of such redistribution, we regress each partner's post-separation share of assets on the pre-separation share. A coefficient of 1 implies the richer partner keeps his or her full share of assets post-divorce. A coefficient of 0 is consistent with an equal split of assets, but also with any other settlement where the post-separation wealth share is independent of the pre-separation wealth level. Table A.13 shows the coefficients from this regression are 0.62 for married couples with children, 0.68 for married couples without children and 0.75 for cohabiting couples. For each type of couple, we reject both equal splits (a coefficient of 0) and complete lack of redistribution from the richer to the poorer partner (a coefficient of 1). Although redistribution is higher for married couples, the small difference compared to cohabiting couples indicates marriage has a modest impact on the division of assets in divorce.

Appendix G.3 discusses the robustness of the Wealth Registry analyses discussed above. In particular, we argue idiosyncratic measurement error in asset values implies Figure 3 Panel B likely overstates wealth inequality after separation, but that such measurement error has limited implications for measured wealth inequality pre-separation and the estimated redistribution from the richer to the poorer spouse.

The Allocation of Lottery Wealth Within Marriage

Our empirical results suggest which spouse wins the lottery matters also for the allocation of consumption within the marriage. Table A.14 shows the winning spouse retains about two thirds of the lottery wealth and Cesarini et al. (2017) document stronger negative labor supply responses of the winning spouse. As shown in Section F.2, these empirical results are consistent with a divorce-threat model where the winner retains most of the lottery wealth in case of divorce.[9] The reason is the stronger outside option of the winning spouse increase bargaining power also inside the marriage.[10]

In contrast, a model with a male breadwinner norm and equal divorce-splits generates a larger increase in the winner's consumption only if the wife espouse the breadwinner norm more strongly than the husband. To see why, first note the spouse for which the norm is stronger will see a larger increase in utility when the husband wins, and a smaller increase (or a decrease) when the wife wins, for a given change in consumption. The norm-abiding spouse must thus receive a larger increase in consumption when the

[9] As shown in Section F.2, how large the winners' share of lottery wealth in divorce has to be in order for the winner to always consume more also within the marriage depends on spouses' marginal utilities of consumption in the single and married state.

[10] As pointed out by Cesarini et al. (2017), the non-equal division of lottery wealth might also be consistent with models with threat points internal to the marriage.

wife wins and a smaller increase when the husband wins for the marital surplus to be shared equally (which in turn follows from the assumption of equal splits of property in divorce). If the wife is more norm-abiding, her consumption must thus increase more when she wins the lottery, but less when her husband wins. If instead the husband is more norm-abiding, the breadwinner norm implies the consumption of the *non-winning* spouse should increase more, a prediction that contrasts sharply to the pattern in the data.

Exploratory Heterogeneity Analyses

We run a number of exploratory heterogeneity analyses to better understand the differential effects on divorce risk. Figure A.4 shows that the short-run effect of the wife's lottery wealth on the divorce probability is positive and statistically distinguishable from zero only in the couples where the wife earns less than the husband (Panel A) or below the sample median (Panel B), suggesting the treatment effect is stronger for women who are more likely to be financially dependent on their husband. These results thus support the notion that wives' lottery wins may increase short-term divorce risk because it gives some women the financial means to pursue a life as singles. Panel B of Figure A.4 also shows the reduction in divorce risk from husbands' lottery wealth is larger in couples where the husband earns below the sample median ($\hat{\beta}_5 = -6.36, SE = 2.57$) compared to couples where the husband earns above the sample median ($\hat{\beta}_5 = -1.93, SE = 1.77$), consistent with low economic status being detrimental to men's attractiveness as marital partners.

Beyond the divorce-threat model we discuss above, a general prediction of many marriage market models is that lottery wealth will facilitate or delay termination of marriages whose dissolution was already underway, while the effect on "stable" marriages will be negligible. To test this hypothesis, we test whether the effect of wife's wealth on marriage dissolution is stronger in (i) couples where the match quality is poor than in better matched couples and (ii) in newly married couples than in longer-term married couples. In order to analyze how the effect on wife's wealth on marriage dissolution varies with the match quality, we construct a proxy for the match quality defined as the absolute value of the difference in spouses age and education (larger distance suggesting lower match quality).[11] Panels C of Figure A.4 shows the estimated short-run effect on divorce when women win is positive in couples with match quality below the sample median but negative in couples with match quality above the sample median. Panel D of A.4 shows the effect of the wife's lottery wealth on the divorce probability by marriage duration, indicating that the short and the medium-run lottery effect is bigger in newly married couples (married for 3 years or less at the moment of win), though we cannot reject that the effects are identical. We do not find that the impact of a husband's wealth on marriage dissolution varies with the match quality or the length of the marriage

6.3 Fertility

Following the seminal work of Becker (1960), children are often introduced into economic models as normal durable goods. However, the assumption that children are normal goods is seemingly at odds with the fact that economic growth has coincided with a transition to lower fertility in many countries (Doepke, 2004; Clark, 2005; Galor, 2005; Bar and Leukhina, 2010; Ager et al., 2020). Cross-sectional studies have not reached a consensus on the effect of income on fertility, with some studies suggesting that income and fertility are positively correlated (Anderson, 2008; Black et al., 2013; Kolk, 2019) and others showing that there is a negative association (Jones and Tertilt, 2008).

One possible explanation for the observed negative relationship between income and fertility is that wage increases may have two offsetting effects for parental decisions to have children. On the one hand, the income effect is expected to increase fertility if children are "normal" goods. On the other hand, higher wages imply higher opportunity cost of time and an increase in the "shadow price" of raising children, which may reduce the optimal level of fertility (the substitution effect).

Testing whether the income effect is positive is challenging. As noted by Hotz et al. (1997), labor market and fertility decisions are taken simultaneously and, therefore, the income-fertility association cannot be interpreted as a causal relationship. Moreover,

[11] Becker (1973) predicts that there is positive assortative mating with respect to complementary traits (e.g., education or physical attractiveness), while the negative assortative mating with respect to substitutionary traits, such as earning power, would be optimal. Therefore, the distance in complementary traits can be informative about the match quality. In our sample of pre-lottery married players, husband's and wife's education are positively correlated ($\hat{\rho} = 0.46$) while the correlation between husband's and wife income is small and negative ($\hat{\rho} = -0.04$).

studies which rely on exogenous variation in wages generally cannot separate the income effect from the substitution effect. In order to confront these challenges, some studies estimate the income effect on fertility exploiting income shocks generated by husband's job displacement (Amialchuk, 2013; Lindo, 2010) or other exogenous shocks that affect men's labor income (Black et al., 2013; Schaller, 2016; Kearney and Wilson, 2018; Autor et al., 2019), assuming that men do not contribute to home production and therefore the effect on fertility is due to the income effect rather than to the changes in the "shadow price" of raising children. The evidence usually indicate that fertility increases with men's income, though there are exceptions (Huttunen and Kellokumpu, 2016). Further studies have found increases in housing wealth (Lovenheim and Mumford, 2013) and lottery winnings (see the discussion in Section 5.2) to increase fertility.

Our results are consistent with children being normal goods, as we find clear evidence that lottery wealth increases fertility in the pooled sample and in the subsample of male winners. However, a back-of-the-envelope calculation suggests about 20-40 percent of male winner's fertility response can be accounted for by their higher marriage rate.[12] Moreover, if men's improved marriage prospects are mainly explained by them becoming relatively more attractive partners compared to other men, a general increase in income may have no effect on marriage rates.

The gendered fertility effects may suggest men and women have different fertility preferences. The agreement between partners is an important determinant of fertility (Doepke and Kindermann, 2019), and, as discussed in Section 6.2, lottery wealth may increase the bargaining power of the winner. In this case, our results would be consistent with men preferring to have more children than women. However, according to the OECD, the av-

[12]We calculate the change in fertility t years after the lottery, $\triangle C_t$, as

$$\triangle C_t = \hat{\beta}_t^m \triangle C_t^m (1 - s^m) + \hat{\beta}_t^d \triangle C_t^d s^m$$

where $\hat{\beta}_t^m$ and $\hat{\beta}_t^d$ are the estimated effects of lottery wealth on marriage and divorce probabilities t years after the lottery event, $\triangle C_t^m$ and $\triangle C_t^d$ denote the difference in the number of post-lottery children between pre-win unmarried men who are married at t and pre-win married men who divorce at t, and s^mdenotes the share of men who are married at the time of the lottery event. Comparing $\triangle C_t$ to the estimated effect of lottery wealth on male winners' fertility indicates 21.5% ($t = 2$), 42.9% ($t = 5$) and 29.8% ($t = 10$) of the fertility response can be accounted for by the effect of lottery wealth on marital formation and dissolution. The validity of the calculation above hinges upon the strong assumption that the effects of lottery wealth on men's marriage formation and marital stability do not reflect an effect on desired fertility.

erage ideal number of children in Sweden is similar for men and women, suggesting there is no gender differences in preferences for children.[13] Moreover, to the extent the effect on male fertility is explained by increased marriage formation and stability, the fertility response reflects a relaxed constraint rather than preferences for children. In addition, the larger gender differences in the fertility response we found for older winners could be due to declining female fertility with age. In sum, the gendered effects we document may reflect the different constraints on fertility facing men and women rather than differences in preferences for children.

7 Conclusion

Our study leverages the randomized assignment of lottery prizes to estimate the effects of wealth on three important family outcomes – marriage formation, marriage dissolution, and fertility. Our estimates have strong internal validity and advance understanding of the broader question of how wealth affects family outcomes by providing credible and precise estimates for a large sample of Swedish lottery players. We find that lottery wealth increases the short- and medium-run probabilities of marriage. The overall wealth effect on marriage formation is driven by male winners. These gender asymmetries are consistent with previous literature that documents gender differences in partner selection, and indicates that women put higher weight on economic characteristics of potential partners than men do (Fisman et al., 2006; Hitsch et al., 2010b,a).

While the overall average treatment effect on marriage dissolution is not statistically distinguishable from zero, there is a consistent pattern of divergence between the estimated effects for husbands and wives. Specifically, when the winning player is a married woman, our estimates suggest that a 1M-SEK windfall almost doubles the baseline short-run divorce rate. This estimated effect appears to fade away in the long-run. We speculate that the positive wealth shock accelerates the exit from marriages whose dissolution was already underway. In contrast, long-term divorce risk goes down when

[13] According to the OECD, the average personal ideal number of children for Swedish women aged 15-64 is 2.41 while it is 2.33 for men. This difference is similar to the difference in other OECD countries. See OECD Family Database available at `http://www.oecd.org/els/family/database.htm`. The estimates are based on data from Eurobarometer 2011: Fertility and Social Climate.

husbands win the lottery. These results are compatible with previous empirical evidence showing that a higher husband's income or employment stabilizes marriages, while an increase in wives' income or employment has the opposite effect (Berniell et al., 2020; Boheim and Ermisch, 2001; Burstein, 2007; Doiron and Mendolia, 2012; Folke and Rickne, 2020; Killewald, 2016; Weiss and Willis, 1997). We show a divorce-threat model where the winning spouse retains the bulk of the lottery prize in case of divorce can rationalize our findings, and present suggestive empirical evidence in support of this explanation. In particular, we show divorce settlements in Sweden are often unequal and favor the richer spouse.

Finally, consistent with theoretical models that introduce children as normal goods (Becker, 1960), we find that the evidence from quasi-experimental studies of lottery players as a whole Tsai et al. (2022); Bulman et al. (2022); Bleakley and Ferrie (2016) is consistent with small but positive wealth effects on completed fertility. This conclusion is also fits with the conclusions from the broader quasi-experimental literature (Black et al., 2013; Lindo, 2010; Lovenheim and Mumford, 2013). The positive effect on fertility in our sample appears to be driven by male winners and appears to be partly explained by the increase in marriage formation for men. If the increase in men's fertility is due to increased attractiveness in the marriage market relative to other men, the positive wealth effects we estimate are likely upper bound for the effects one should expect in response to a more general upward shift in overall living standards and wealth.

References

Ager, P., B. Herz, and M. Brueckner (2020): "Structural Change and the Fertility Transition," *Review of Economics and Statistics*, 102, 806–822.

Almås, I., A. Kotsadam, E. R. Moen, and K. Røed (2023): "The Economics of Hypergamy," *Journal of Human Resources*, 58, 260–281.

Amialchuk, A. (2013): "The Effect of Husband's Job Displacement on the Timing and Spacing of Births in the United States," *Contemporary Economic Policy*, 31, 73–93.

Anderson, M. L. (2008): "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of theAbecedarian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association*, 103, 1481–1495.

Autor, D., D. Dorn, and G. Hanson (2019): "When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men," *American Economic Review: Insights*, 1, 161–78.

Azmat, G. and L. González (2010): "Targeting Fertility and Female Participation Through the Income Tax," *Labour Economics*, 17, 487–502.

Bar, M. and O. Leukhina (2010): "Demographic transition and industrial revolution: A macroeconomic investigation," *Review of Economic Dynamics*, 13, 424–451.

Baughman, R. and S. Dickert-Conlin (2003): "Did Expanding the EITC Promote Motherhood?" *American Economic Review*, 93, 247–251.

——— (2009): "The Earned Income Tax Credit and Fertility," *Journal of Population Economics*, 22, 537–563.

Becker, G. S. (1960): "An Economic Analysis of Fertility," in *Demographic and Economic Change in Developed Countries*, Columbia University Press, 209–240.

——— (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81, 813–846.

——— (1974): "A Theory of Marriage: Part II," *Journal of Political Economy*, 82, S11–S26.

Becker, G. S., E. M. Landes, and R. T. Michael (1977): "An Economic Analysis of Marital Instability," *Journal of Political Economy*, 85, 1141–1187.

Benjamini, Y., A. M. Krieger, and D. Yekutieli (2006): "Adaptive Linear Step-Up Procedures That Control the False Discovery Rate," *Biometrika*, 93, 491–507.

Berniell, I., D. de la Mata, and M. P. Machado (2020): "The Impact of a Permanent Income Schock on the Situation of Women in the Household: The Case of Pension Reform in Argentina." *Economic Development and Cultural Change*, 68, 1295–1343.

Bertrand, M., E. Kamenica, and J. Pan (2015): "Gender Identity and Relative Income Within Households," *Quarterly Journal of Economics*, 130, 571–614.

Bitler, M. P., J. B. Gelbach, H. W. Hoynes, and M. Zavodny (2004): "The Impact of Welfare Reform on Marriage and Divorce," *Demography*, 41, 213–236.

Black, D. A., N. Kolesnikova, S. G. Sanders, and L. J. Taylor (2013): "Are Children "Normal" ?" *The Review of Economics and Statistics*, 95, 21–33.

Bleakley, H. and J. Ferrie (2016): "Shocking Behavior: Random Wealth in Antebellum Georgia and Human Capital Across Generations," *Quarterly Journal of Economics*, 131, 1455–1495.

Boertien, D. (2012): "Jackpot? Gender Differences in the Effects of Lottery Wins on Separation," *Journal of Marriage and Family*, 74, 1038–1053.

Boheim, R. and J. Ermisch (2001): "Partnership Dissolution in the UK - The Role of Economic Circumstances," *Oxford Bulletin of Economics and Statistics*, 63, 197–208.

Brattström, M. (2011): "Bodelning mellan makar–verklighetens betydelse för framtidens regelutformning?–Delrapport 1 inom ramen för prosjektet Rättvis delning? Frågor om delningsrättens täckningsområde vid bodelning mellan makar eller sambor," *Tidsskrift for familierett, arverett og barnevernrettslige spørsmål*, 9, 56–98.

Briggs, J., D. Cesarini, E. Lindqvist, and R. Östling (2021): "Windfall gains and stock market participation," *Journal of Financial Economics*, 139, 57–83.

Bulman, G., S. Goodman, and A. Isen (2022): "The Effect of Financial Resources on Homeownership, Marriage, and Fertility: Evidence from State Lotteries," Tech. rep., NBER Working Paper No. 30743.

Burgess, S., C. Propper, and A. Aassve (2003): "The Role of Income in Marriage and Divorce Transitions Among Young Americans," *Journal of Population Economics*, 16, 455–475.

Burstein, N. R. (2007): "Economic Influences on Marriage and Divorce," *Journal of Policy Analysis and Management*, 26, 387–429.

Cain, G. G. and D. A. Wissoker (1990): "A Reanalyis of Marital Stability in the Seattle-Denver Income-Maintenance Experiment," *American Journal of Sociology*, 95, 1235–1269.

Cesarini, D., E. Lindqvist, M. J. Notowidigdo, and R. Östling (2017): "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries," *American Economic Review*, 107, 3917–46.

Cesarini, D., E. Lindqvist, R. Östling, and B. Wallace (2016): "Wealth, Health, and Child Development: Evidence from Administrative Data on Swedish Lottery Players," *Quarterly Journal of Economics*, 131, 687–738.

Clark, G. (2005): "Human Capital, Fertility, and the Industrial Revolution," *Journal of the European Economic Association*, 3, 505–515.

Cohen, A., R. Dehejia, and D. Romanov (2013): "Financial Incentives and Fertility," *Review of Economics and Statistics*, 95, 1–20.

Doepke, M. (2004): "Accounting for Fertility Decline During the Transition to Growth," *Journal of Economic Growth*, 9, 347–383.

Doepke, M. and F. Kindermann (2019): "Bargaining over Babies: Theory, Evidence, and Policy Implications," *American Economic Review*, 109, 3264–3306.

Doiron, D. and S. Mendolia (2012): "The Impact of Job Loss on Family Dissolution," *Journal of Population Economics*, 25, 367–398.

Efron, B. (2010): *Large-Scale Inference: Empirical Bayes Methods for Estimation,Testing and Prediction*, Cambridge University Press.

Ellwood, D. T. and C. Jencks (2004): "The Uneven Spread of Single-Parent Families: What do We Know? Where to We Look for Answers?" *Social Inequality*, 1, 3–77.

Fisman, R., S. S. Iyengar, E. Kamenica, and I. Simonson (2006): "Gender Differences in Mate Selection: Evidence From a Speed Dating Experiment," *Quarterly Journal of Economics*, 121, 673–697.

Folke, O. and J. Rickne (2020): "All the Single Ladies: Job Promotions and the Durability of Marriage," *American Economic Journal: Applied Economics*, 12, 260–87.

Galor, O. (2005): "The Demographic Transition and the Emergence of Sustained Economic Growth," *Journal of the European Economic Association*, 3, 494–504.

Golosov, M., M. Graber, M. Mogstad, and D. Novgorodsky (2023): "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income," *Quarterly Journal of Economics, in press.*

Groeneveld, L. P., N. B. Tuma, and M. T. Hannan (1980): "The Effects of Negative Income Tax Programs on Marital Dissolution," *Journal of Human Resources*, 654–674.

Hainmueller, J. (2012): "Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies," *Political Analysis*, 20, 25–46.

Hankins, S. and M. Hoekstra (2011): "Lucky in Life, Unlucky in Love? The Effect of Random Income Shocks on Marriage and Divorce," *Journal of Human Resources*, 46, 403–426.

Hederos, K. and A. Stenberg (2022): "Gender identity and relative income within households: Evidence from Sweden," *The Scandinavian Journal of Economics*, 124, 744–772.

Hitsch, G. J., A. Hortaçsu, and D. Ariely (2010a): "Matching and Sorting in Online Dating," *American Economic Review*, 100, 130–63.

Hitsch, G. J., A. Hortacsu, and D. Ariely (2010b): "What Makes You Click? - Mate Preferences and Matching Outcomes in Online Dating," *Quantitative Marketing and Economics*, 8, 393–427.

Hoffman, S. D. and G. J. Duncan (1995): "The Effect of Incomes, Wages, and AFDC Benefits on Marital Disruption," *Journal of Human Resources*, 30, 19–41.

Hotz, V. J., J. A. Klerman, and R. J. Willis (1997): "The Economics of Fertility in Developed Countries," *Handbook of Population and Family Economics*, 1, 275–347.

Hu, W.-Y. (2003): "Marriage and Economic Incentives Evidence from a Welfare Experiment," *Journal of Human Resources*, 38, 942–963.

Huttunen, K. and J. Kellokumpu (2016): "The Effect of Job Displacement on Couples' Fertility Decisions," *Journal of Labor Economics*, 34, 403–442.

Imbens, G. W., D. R. Rubin, and B. I. Sacerdote (2011): "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players," *American Economic Review*, 91, 778–794.

Jones, L. E. and M. Tertilt (2008): "Chapter 5: An Economic History of Fertility in the United States: 1826-1960," in *Frontiers of Family Economics*, ed. by P. Rupert, Emerald Group Publishing Limited, 165–230.

Kaplan, H. R. (1987): "Lottery Winners: The Myth and Reality," *Journal of Gambling Behavior*, 3, 168–178.

Kearney, M. S. and R. Wilson (2018): "Male Earnings, Marriageable Men, and Nonmarital Fertility: Evidence From the Fracking Boom," *Review of Economics and Statistics*, 100, 678–690.

Killewald, A. (2016): "Money, Work, and Marital Stability: Assessing Change in the Gendered Determinants of Divorce," *American Sociological Review*, 81, 696–719.

Klein, J. (2017): "House Price Shocks and Individual Divorce Risk in the United States," *Journal of Family and Economic Issues*, 38, 628–649.

Kolk, M. (2019): "The Relationship Between Lifecourse Accumulated Income and Childbearing of Swedish Men and Women born 1940-1970," *Stockholm Research Reports in Demography | no 2019:19.*

Lindo, J. M. (2010): "Are Children Really Inferior Goods? Evidence from Displacement-Driven Income Shocks," *Journal of Human Resources*, 45, 301–327.

Lindqvist, E., R. Östling, and D. Cesarini (2020): "Long-run Effects of Lottery Wealth on Psychological Well-Being," *The Review of Economic Studies*, 87, 2703–2726.

Lovenheim, M. F. and K. J. Mumford (2013): "Do Family Wealth Shocks Affect Fertility Choices? Evidence from the Housing Market Boom and Bust," *Review of Economics and Statistics*, 95, 464–475.

Lundberg, S. J. and R. A. Pollak (1996): "Bargaining and Distribution in Marriage," *Journal of Economic Perspectives*, 10, 139–158.

Manser, M. and M. Brown (1980): "Marriage and Household Decision-Making: A Bargaining Analysis," *International Economic Review*, 31–44.

McElroy, M. B. and M. J. Horney (1981): "Nash-Bargained Household Decisions: Toward a Generalization of the Theory of Demand," *International Economic Review*, 333–349.

Moffitt, R. (1990): "The Effect of the US Welfare System on Marital Status," *Journal of Public Economics*, 41, 101–124.

Oppenheimer, V. K. (1988): "A Theory of Marriage Timing," *American Journal of Sociology*, 94, 563–591.

Östling, R., D. Cesarini, and E. Lindqvist (2020): "Association Between Lottery Prize Size and Self-reported Health Habits in Swedish Lottery Players," *JAMA Network Open*, 3, e1919713–e1919713.

Regan, P. C., L. Levin, S. Sprecher, F. S. Christopher, and R. Gate (2000): "Partner PreferencesWhat Characteristics Do Men and Women Desire In Their Short-Term Sexual and Long-Term Romantic Partners?" *Journal of Psychology & Human Sexuality*, 12, 1–21.

Rosenzweig, M. R. (1999): "Welfare, Marital Prospects, and Nonmarital Childbearing," *Journal of Political Economy*, 107, S3–S32.

Schaller, J. (2016): "Booms, Busts, and Fertility Testing the Becker Model Using Gender-Specific Labor Demand," *Journal of Human Resources*, 51, 1–29.

Schultz, T. P. (1985): "Changing world prices, women's wages, and the fertility transition:Sweden, 1860-1910," *Journal of Political Economy*, 93, 1126–1154.

Statistics Sweden (2017): "Description of the Register: Total Population Register," Accessed on 2021-01-04 at www.scb.se/contentassets/8f66bcf5abc34d0b98afa4fcbfc0e060/rtb-bar-2016-eng.pdf.

Tsai, Y.-Y., H.-W. Han, K.-T. Lo, and T.-T. Yang (2022): "The Effect of Financial Resources on Fertility: Evidence from Administrative Data on Lottery Winners," *arXiv preprint arXiv:2212.06223*.

Weiss, Y. (1997): "The formation and dissolution of families: Why marry? Who marries whom? And what happens upon divorce," *Handbook of Population and Family Economics*, 1, 81–123.

Weiss, Y. and R. J. Willis (1997): "Match Quality, New Information, and Marital Dissolution," *Journal of Labor Economics*, 15, S293–S329.

Yonzan, N., L. Timilsina, and I. R. Kelly (2020): "Economic Incentives Surrounding Fertility: Evidence from Alaska's Permanent Fund Dividend," Tech. rep., National Bureau of Economic Research.

Young, A. (2019): "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results," *Quarterly Journal of Economics*, 134, 557–598.