Child poverty and income mobility in the United States*

Tim Smith[a,†], Michael S. Delgado[a] Raymond J.G.M. Florax[a,b]

[a] *Department of Agricultural Economics, Purdue University, West Lafayette, IN, USA*
[b] *Department of Spatial Economics, Vrije Universiteit Amsterdam, and Tinbergen Institute, Amsterdam, The Netherlands*

October 4, 2018

Abstract

Childhood poverty has long-lasting effects, but much of the literature on human capital formation suggests that different kinds of poverty have different effects. We classify children's poverty spells according to four distinct features – duration, intensity, timing, and concentration – which allows us to differentiate between childhood poverty profiles, to assign roles to each feature, and to examine interactions between them. We find that the duration of childhood poverty has a compounding nonlinear effect on adulthood rank, and that the other features matter much less. These findings suggest that policy interventions aimed at children who have experienced a short to moderate spell of poverty, and are at risk of experiencing more, will be most efficient.

Keywords: intergenerational mobility, child poverty, human capital
JEL-classification: D31, I32, J62

* We have profited from comments and suggestions of Huan Li, Brigitte Waldorf, Tom Hertz, and participants in the Spatial, Health, and Population Economics (SHaPE) Seminar at Purdue. Any remaining errors are our own. This research has been partly financed through USDA grant 58-6000-5-0044. The authors note that the views expressed are those of the authors and should not be attributed to the Economic Research Service (ERS) or the United States Department of Agriculture (USDA).

† Corresponding author. Tel. 1 + (765) 543–4062. E-mail: smit1638@purdue.edu.

1 Introduction

Compared to the rest of the OECD, in the United States rates of child poverty remain high while intergenerational economic mobility remains low. According to OECD child poverty statistics, 20.2 percent of children in the United States were poor, compared to an average rate of 13.6 percent across the OECD; out of the other North American or Western European OECD countries, the only one with a higher rate of child poverty than the United States is Spain, with a rate of 22.7 percent (OECD, 2018).[1] Corak (2006) and Blanden (2013) find that the United States is substantially less mobile than other Western countries when mobility is measured with an intergenerational income elasticity that measures income similarity across generations; in both studies, the United States ranks second lowest in mobility among the countries studied. We explore these issues through an empirical model of individual income rank (as an adult) that is a function of a multi-dimensional index of childhood poverty that captures the age at which the child entered poverty, the intensity and duration of that poverty, and the concentration of individuals' time in poverty.[2] Our goal is to better understand the way in which childhood poverty affects adulthood outcomes: accounting for both poverty intensity and duration allows us to determine the extent to which adulthood outcomes are affected differently by poverty spells of different natures, accounting for age at poverty exposure allows us to investigate whether younger children are disproportionately affected (as adults) relative to older children, and incorporating the concentration of poverty lets us consider how these factors vary depending on the spacing of periods of time in poverty. This analysis is policy-relevant as our analysis informs policies that target childhood poverty by clarifying the context (nature and age of exposure) through which poverty affects individual socioeconomic outcomes; indeed, efficient use of social funds requires detailed knowledge of the processes that connect childhood inputs to adulthood outcomes.

Our work is motivated by the large literature that has investigated how different types of childhood interventions lead to positive socioeconomic outcomes for the individual as an adult. Several papers, for instance Cunha et al. (2010) and Agostinelli and Wiswall (2016), develop micro-theoretic foundations that link early-childhood human capital investments to socioeconomic outcomes as an adult, and provide supporting empirical evidence that parental investments made at an early childhood age compound over the remaining years of childhood. These papers build on the classic papers by Becker and Tomes (1979, 1986), focusing on the technology by which parents might influence an individual's economic outcomes. One alternative approach to focusing on the shape of a production function is to instead look for reduced-form (causal) evidence that children who experience a major socioeconomic adjustment (e.g., geographic relocation) subsequently experience better socioeconomic opportunities as adults (Chetty et al., 2016; Chyn, 2016), or to estimate intergenerational income elasticities or rank correlations (Chetty et al., 2014). The structural investment models predict large differ-

[1]The OECD child poverty rates in the United States are consistent with other recent United States sources; for instance, the Annie E. Casey Foundation and the National Center for Children in Poverty both report a rate of 19 percent in 2018, based on American Community Survey data (Koball and Jiang, 2018; Foundation, 2018).

[2]Concentration refers, in principle, to the degree to which someone's time in poverty occurred all at once, rather than across childhood. We discuss our measurements in detail in Section 3.2, but generally speaking, we want to capture the difference between a child who dipped into poverty at different points in childhood, and a child who experienced the same poverty, but across a contiguous stretch of years, such as ages three through six.

ences in the effect of socioeconomic interventions by age at which the intervention occurs, while much of the (causal) reduced form evidence suggests that intervention effects are largely homogeneous across ages, although some recent work finds that age thirteen is a kind of threshold, before which certain interventions are more effective (Chetty et al., 2016). These findings lead to a policy-relevant empirical puzzle in the sense that different empirical findings suggest different interventions: if the effects of childhood experiences on adulthood outcomes are (largely) constant across age at which the experience occurs, policymakers can rely on relatively broad policies such as housing vouchers, school desegregation (on income or racial lines), or improved primary education in poor neighborhoods improve an individual's lifetime opportunities. On the other hand, strongly age-differential effects suggest that policymakers take a much narrower focus on neonatal care and accessible pre-Kindergarten childcare and education.

Approaching childhood poverty as a treatment with several distinct features helps us to shed further light onto this policy-dilemma. In measuring poverty as a continuous, multivariate process that accounts for duration, intensity, timing, and concentration components, we posit that differences in childhood poverty spells (along these dimensions) produce different effects on an individual's position in the income distribution as an adult. By considering these distinct components of poverty, we are able to detect nonlinearities in the poverty-adult outcomes relationship that the child development literature suggests, while maintaining a broad focus on childhood poverty, rather than focusing directly on production functions for skills. Our intuition is that each of these measures correspond to a distinct feature of an individual's exposure to poverty: duration – defined as the number of years an individual spent in poverty as a child, captures how long an individual spent in poverty; intensity – defined as the average difference between one's household income and the poverty line when in poverty, in percentage points, reflects how serious that poverty was in income terms; and timing – the age at which each individual first experienced poverty, accounts for the developmental stage at which a person experienced poverty. We measure the concentration of childhood poverty with several measures, but the most important are the standard deviation of ages at which each individual was poor and the number of contiguous poverty spells each individual experienced as a child, both of which reflect the extent to which an individual's experience of poverty was spread across childhood, rather than being concentrated in one period. In our view, concentration is somewhat different from the other three treatments because we expect that it affects adulthood outcomes only through them, rather than mattering on its own, so while it enters our models in the same ways as the other treatments, it occupies a different position in our discussion. Accounting for concentration is very important, however, because it lets us hold this important, and often ignored, dimension of poverty constant as we model the effects of the components we expect to have first-order influence on adulthood rank. Drawing this distinction between timing and duration lets us isolate timing effects: in our sample, there are many values of duration for each value of age at first exposure, and vice versa. This approach lets us focus directly on the age heterogeneity issue central to the apparent differences between these literatures.

The economic foundation supporting our work comes from a particular subset of the human capital literature concerned with estimating elements of the technology that transforms inputs individuals receive in childhood into skills (e.g., Cunha and Heckman, 2010; Almond and Currie, 2011; Heckman and Mosso, 2014).[3] These skills, in turn, affect

[3] Almond and Currie (2011) review this literature and note in their abstract that "Child and family characteristics measured at school entry do as much to explain future outcomes as factors that labor

adulthood outcomes, such as income and educational attainment, but the magnitude of this effect decreases in the age at which the individual receives the investment. This age-decreasing relationship is driven by a phenomenon Cunha et al. (2010) refer to as 'dynamic complementarity': the extent to which investments in skill development in later childhood influences the individual's outcomes depends on previously acquired skills, and so being poor for an additional year will not only lead to lost skills for the year of poverty but will depress the absorption of skills acquired in subsequent years. Translated into our view of poverty – i.e., distinguishing between age, duration, intensity, and concentration – we expect that poverty duration and intensity have distinct effects on individual outcomes, and that these distinct effects are different by age of poverty exposure and the degree of poverty concentration. Duration corresponds to a sustained reduction in parental investments, while intensity corresponds to a dramatic reduction in parental investments, the interaction between duration and intensity corresponds to a sharp and sustained reduction in parental investments. If a sharp, sustained reduction in investment occurs during early childhood, particularly if the reduction is concentrated in this period rather than spread into later childhood, then the micro-theoretic literature would suggest that we would find a significantly negative effect on socioeconomic outcomes as an adult, and that holding the size of the reduction in parental investment constant, the effect would be disproportionately smaller if it occurred later in childhood.[4] Yet, the most striking estimates from the Moving to Opportunity experiment indicate that there was a strong positive effect of improved neighborhood quality for individuals who were younger than thirteen when they moved, but not for older children (e.g., Chetty et al., 2016).[5] Within this group of pre-thirteen year olds, however, Chetty et al. (2016) find that the effect of an additional year in a better neighborhood is linear, arguing that this piecewise linear relationship explains the difference in responses between younger and older children. Moving to a better neighborhood increases earnings (or other socioeconomic outcomes) linearly, but a fixed 'disruption cost' of moving decreases the effect of total exposure to the improved environment and so the treatment effect drops to zero for children who move later and receive fewer years in the better neighborhood. Translated into our model of poverty, these results imply that poverty spells do not have heterogeneous, age-differential

economists have more traditionally focused on...". Cunha and Heckman (2010) provide a complementary overview, also emphasizing the disproportionate benefits of early childhood investment.

[4]Similarly, Agostinelli and Wiswall (2016) find several sources of nonlinearity in the relationship between childhood investments and adulthood outcomes: the average treatment effect of a transfer of \$1000 drops substantially (60 percent or more) by the time a child is 11 or 12, compared to when she is 5 or 6, and that much of this drop occurs between the first period and the period when the child is 7 or 8. Further, the effects of the transfer on 5 and 6 year old children are approximately twice as large as the average treatment effect at the tenth percentile of family income, and that they decline convexly as income rises, consistent with both nonlinearity in intensity, and with an interaction between intensity and timing.

[5]The Moving to Opportunity experiment is an intervention funded and implemented by the U.S. Department of Housing and Urban Development (HUD) that involved 4,604 families living in high poverty census tracts between 1994 and 1998 in Baltimore, Boston, Chicago, Los Angeles and New York City. Each family was randomly assigned to a group that received either a voucher that subsidized private rents in tracts with poverty rates below 10 percent, a voucher that allowed them to move into subsidized housing without a location constraint, or into the control group. The study has followed individuals who moved since initial assignment, and early studies (e.g. Katz et al. (2001), Kling et al. (2007), Ludwig et al. (2013)), found no evidence for long-term benefits on earnings or employment, although they did find positive effects on mental health and schooling for females. People treated in early childhood reached adulthood relatively recently, however, so these studies are not able to model the effects Chetty et al. (2016) detect for this subgroup.

effects on individual outcomes, and that instead entrance into poverty at an early age corresponds to a constant (by year) drop in individual income rank with no significant effect if the poverty spell was to begin after age thirteen.

We proceed using data from the Panel Study of Income Dynamics (PSID). These data are a key feature of our study: though the data are not experimental, the rich set of longitudinal data is crucial in defining our more nuanced poverty measures. These measures let us distinguish between different aspects of poverty which allow us to consider the results of relocation studies and the structural work on the technology of human capital formation in a common framework. This common framework is necessary for understanding what features of childhood poverty might produce these parallel bodies of results, which is the key contribution of our paper. Our use of these observational data in a reduced form setting places our contribution between the studies relying on field experiments and the studies developing structural models. We start by benchmarking our PSID data to the results from Chetty et al. (2016), and then expand our focus to the nonlinear, multivariate framework introduced in Section 2; yet, we do not consider treatment as at fine a scale as the structural literature, opting instead to bridge their results to our coarser, but somewhat more tangible, set of treatment and outcome measures.

We find a strong and statistically significant association between adulthood income rank and poverty duration, with much larger marginal effects at lower values (between one and six years of childhood poverty) than at relatively large values (between eleven and seventeen years of childhood poverty). This relationship is not mediated by the other treatments, and none of them are significantly associated with adulthood income rank in our preferred specifications. These results reinforce the conclusions Cunha and Heckman (2010) and Agostinelli and Wiswall (2016) reach regarding the importance of early interventions, because they suggest that interventions targeting individuals with a small to moderate dose of poverty will have larger effects than targeting someone who has already experienced a long poverty spell. Our results also indicate that apparent differences between the Chetty et al. (2016) results and structural results are driven by complementarity between household poverty model and the neighborhood inputs they vary through their experimental design, as Agostinelli (2018) has argued.

2 Hypotheses and Modeling Strategy

2.1 Null Hypotheses

We formulate and test two null hypotheses related to our multi-dimensioned measure of childhood poverty:

Null Hypothesis 1 The intensity and duration of childhood poverty do not have distinct impacts on the economic status of an individual during adulthood.

Null Hypothesis 2 The effect of childhood poverty on economic outcomes at adulthood is linear in the age at which the individual was exposed to poverty.

The intuition we have presented in the introduction largely leads us to believe that empirical evidence leads to a rejection of these null hypotheses. Specifically, we expect that different aspects of poverty – captured through intensity and duration – have different effects on individual outcomes because they correspond to different natures of parental investments in the child. Likewise, a large majority of empirical evidence leads us to expect that the relationship between childhood poverty and individual economic outcomes are nonlinear in age of exposure,. Furthermore, we expect that these effects in turn vary in the concentration of childhood poverty, with greater concentration in early childhood corresponding to more negative and nonlinear effects than poverty that is otherwise equivalent, but spread across childhood.

2.2 Empirical Models

Our empirical model of income rank is

$$rank_i = g(poverty_i) + X_i\beta + \epsilon_i, \tag{1}$$

where $i = 1, 2, \ldots, n$ indexes individuals, $rank_i$ is the rank of individual i's income in the national income distribution during the time of adulthood, $poverty_i$ is some measure of poverty (defined below) during the time the individual was a child, $g(\cdot)$ is a smooth (differentiable) but otherwise unspecified function, X_i represents a vector of control variables with coefficient vector β, and ϵ is a mean-zero error term. This model assumes cross-sectional data because we average across years as a child and adult separately for each individual to generate a dataset that links each individual's economic status as an adult to his/her childhood economic status. At the same time, this model also captures alternative models from the literature; for example, in the relocation studies the measure of poverty exposure is implicitly 17 minus the age at which the individual moved out of poverty. The relationship between this model and models of childhood investment technology is less clear because childhood investment levels do not easily translate into poverty; and yet, that literature implies that $g(\cdot)$ in our model should be convex and decreasing in continuous poverty measures for a fixed age of exposure. In our case, we define poverty as a function of age of exposure, duration, and intensity, so that Equation 1 becomes

$$Rank_i = g(Age_i, Duration_i, Intensity_i, Concentration_i) + X_i\beta + \epsilon_i. \tag{2}$$

We estimate our regression model using several different estimation strategies that allow for both parametric (i.e., relatively restrictive) and nonparametric (i.e., relatively flexible) assumptions on $g(\cdot)$. The simplest model allows only for linear interactions between age, duration, intensity, and concentration, and can be estimated using ordinary least squares. This type of model is potentially quite restrictive by not allowing for the general types of nonlinearity that are shown to exist in the early-development literature. The most general specification allows for general nonlinearities and interactions and requires the use of advanced estimation methods. We present technical estimation details for these models in the appendix; see, also, Li and Racine (2007). We consider the range of structure for $g(\cdot)$ so that we can statistically test the more stringent specifications against the more relaxed specifications, thereby allowing us to understand the extent to which there are nonlinearities (and possibly what type of nonlinearities) in the childhood poverty-economic status relationship.

This scope of models allows us to detect nonlinearities and interactions without requiring distributional or functional form assumptions. Our approach is different from the structural models, which are built on microeconomic models with two or more discrete periods during childhood, and which relate investments to skills, and skills to outcomes. This difference is intentional: though we are studying the long-term effects of poverty, as do the relocation studies, we also want our models to be able to detect nonlinearities that the theoretical and empirical results of the structural literature predict, albeit indirectly, in the context of a model of poverty. In other words, our empirical model has the potential to provide a bridge between the relocation studies and the structural model literature.

3 Data and Measurements

3.1 Sample Description and Selection

Our data comes from the Panel Study of Income Dynamics (PSID).[6] The PSID includes a large number of parent-child pairs, as well as a rich set of variables capturing non-income inputs and outcomes across individuals' lives. Our final sample includes observations across many birth year cohorts; we summarize the distribution of our two samples (described below) via histograms in Figure 1. The mode (1968) is substantially more frequent than the second most frequent year; representation over the other cohorts is reasonably even, however. The frequencies in Figure 1 have standard deviations of about 33 and 27, respectively, which explains why the time period begins in 1963 when the PSID started in 1968.

Estimating intergenerational relationships in income ranks (or income) is plagued by various technical difficulties, many of which place constraints on the timing of observations during both childhood and adulthood, and consequently limit the sample. In our case, these issues are exacerbated by the fact that we need accurate measurements of our three primary treatment variables, all of which depend on different sample selection criteria, further limiting the size of our preferred sample. We impose requirements on the number of observations in both childhood and adulthood. In childhood, we require that each child be observed in each possible childhood year. In adulthood, we require that each individual is observed at least three times between the ages of 25 and 35. The adulthood requirement is necessary to reduce the influence of lifecycle bias, both from measuring income at inconsistent points in individuals' lives within the cross-section we use, and from measuring incomes in years in which it correlated less well with permanent income.[7]

Our second restriction – that we completely observe all years in childhood – is very important. Relaxing this restriction, so that we do not observe every year of each child's childhood, would substantially increase the risk of mismeasuring duration. If, for example, we relax the restriction and include anyone who is in the sample for ten or more years during childhood, and then measure duration as the percentage of observed years in which

[6]The PSID is a household survey that began in 1968 and continues today, which tracks income, employment, education, demographic characteristics, and a wide variety of additional information on individuals and households. One valuable feature of the PSID is the fact that it adds new members to the initial households (such as new children or partners), and continues to do this across generations so as to increase the sample size while maintaining high-quality links between individuals and households. These data have been used frequently in studies of intergenerational mobility (Solon et al. (1991) ,Solon (1992), Page and Solon (2003)) and of the relationship between childhood circumstances and adulthood outcomes (Brooks-Gunn et al. (1993), Duncan et al. (1998), Wagmiller et al. (2006), Smith (2009)).

[7]See Haider and Solon (2006) and Hertz (2007) for further discussion of these lifecycle bias issues in the context of intergenerational mobility measures.

the individual's household was poor,duration would be measured with error, because the distribution of poverty in the unobserved years is unknown.

One restriction which we do not impose in our main sample merits discussion as well: we do not restrict the sample to individuals who experience only a single contiguous poverty spell in our main analysis. This restriction ensures that each individual has as concentrated a poverty spell as possible, given their duration, which means that our estimates of $g(\cdot)$ are not affected by unobserved differences in concentration. This is helpful because it bolsters the credibility of our timing measure by ensuring that each child's time in poverty is spaced in the same way: no one leaves and re-enters, which could happen at different intervals that are difficult to hold constant. Since we know each person's entire duration is contiguous, and that every person is observed at every possible PSID interview in childhood, and we know how old each person was when she first entered poverty, we can map each timing and duration value to a specific age interval. In that case, age at first exposure does fully characterize the timing of total poverty exposure. By contrast, without this restriction age at first exposure reflects when poverty began, then rapidly loses precision as we move from that point in time. We do not use this sample in our primary results because of the cost in observations: we lose approximately 50 percent of our preferred sample if we impose this restriction. We do, however, repeat our analysis using this sample, which serves as one means of determining the impact of omitting concentration in our primary results, because it holds the level of concentration constant across individuals.

3.2 Key Measurements

Our outcome of interest is family income rank in adulthood, following Chetty et al. (2014). In the PSID, this variable is defined as the sum across all sources of income at the household level, and consequently includes transfer. In our preferred sample, the correlation between total household income and household income without transfers is 0.98 among childhood observations and approximately 0.996 among adulthood observations; thus, our results do not depend substantially on whether or not we include transfers.[8] We deflate total household income to real 2012 dollars before computing income ranks, and we report descriptive statistics of family income during both childhood years (ages 1-17) and adulthood years (ages 25-35), in Table 1. To provide a clearer impression of the distribution of these key variables, we also present scatter plots of age, duration, and intensity combinations in Figure 2, and of rank by age, duration, and intensity in Figure 3. Since the PSID is a sample of households and not the population, we approximate rank by computing each individual's percentile in the national distribution of family income ranks, available from the Integrated Public Use Micro Sample (IPUMS) database (Ruggles et al., 2018). Since the IPUMS data are only available decennially, we use the distribution that corresponds to the nearest census year.[9]

Our approach to measuring poverty is perhaps the most important aspect of our data preparation. From the PSID, we are able to capture three key features of childhood poverty exposure: the duration of exposure, the intensity of exposure, and the timing of exposure. We measure duration as the number of years an individual spent in a household

[8]These correlations are similar in the smaller subsample we discuss later in this section: 0.985 and 0.996, respectively.

[9]For example, we evaluate an individual's income rank in 1974 using the 1970 income distribution, while we use the 1980 income distribution for 1976.

with a total family income below the poverty threshold, both of which are recorded as survey responses in the PSID at the family level. This removes the need for assigning or imputing a poverty line measure from other variables and/or data sources, as the poverty line for each family, defined by their characteristics, is already recorded. Duration values can thus take any whole-number value between 0 and 17, because an individual is no longer a child at age eighteen. We define the intensity of poverty as the ratio of family income to the poverty line in years in which the family fell below the line, averaged over all the years of each individual's childhood. This means that intensity is continuous on $[0, 1]$, as averaging only over years in which the household was below the poverty line treats all years which would have ratios greater than one as zeroes in the summation. Finally, we capture the timing of poverty spells with the individual's age at his or her first exposure to poverty. Conditional on intensity and duration, this timing measure provides a straightforward measurement of when someone experienced poverty, which we then use to understand how the timing of poverty affects outcomes.

In addition to bounding the influence of concentration by holding it constant through sampling restrictions, we model the impact of concentration on rank and the relationship between rank and other treatments by including several measures of concentration in models of $g(\cdot)$. We expect that this factor only matters in the context of the other three treatments, and so we treat it as somewhat secondary. Our results, particularly regarding age at first exposure to poverty, may be biased by error in measurement of the timing of poverty. It is helpful to think of this in terms of concentration of poverty years: a child with three years of poverty that occur contiguously has experienced more concentrated poverty than a child who experienced the same amount of poverty in separate periods. This concentration is separate from timing, as two children who entered poverty at the same age, and stayed equally poor for the same number of years could still have different levels of concentration. It is, however, related to timing because greater concentration along with an early age implies, through dynamic complementarity, a more significant early source of deprivation for given levels of duration and intensity. Consequently, we expect that the relationship between rank and age (and perhaps the other relationships as well, particularly in models with interactions) may depend on this measurement.

Our primary measure of concentration is the standard deviation of the ages at which each individual was poor, divided by duration. Individuals with contiguous poverty spells will have the minimum standard deviation for duration, which will increase as poverty ages spread out further, but the standard deviation will also increase as duration increases, even if contiguity is preserved.[10] To illustrate this measure, consider the example of two children with contiguous poverty spells, one of 2 years and the other of 17 years, each entering at age 1. The standard deviation of (1,2), 0.71, is much lower than the standard deviation of (1,2,...,17), 5.05, despite the fact that both children experienced a 'concentrated' period of poverty in early childhood. It does not make sense to say the higher-duration child experienced less concentrated poverty because that concentrated spell continued. Dividing the standard deviation by duration removes the majority of this artificial relationship between duration and concentration. Under this normalization, concentration values for possible values of duration, given a starting age of 1 with contiguous poverty years, range from 0.297 to 0.354, compared to 0.71 to 5.05 without it. As an alternative measure of the concentration of poverty, we consider the number of poverty spells each individual experienced. To a certain extent, this measure reflects a feature of poverty similar to our standard deviation measure, but it differs in that it

[10]We use this measure because we can compute this value for our full sample.

focuses narrowly on moves in and out of poverty. This variable has the advantage of being more easily interpretable, and, of potentially capturing the impact of disruption associated with repeated moves in and out of poverty, however. Figure 4 shows the distribution of spells in our sample; the coverage is reasonable across different values: we have a fairly large number of individuals with two or three spells, and enough with four or five spells to conduct statistical inference about these groups. We also consider the variance of the intensity measures in children's poverty years and the variance in children's poverty status, i.e. a binary variable taking a value of one in years when a child was poor, and zero otherwise. These do not affect our results, however, so for the sake of space and focus, we do not discuss them further here or in our results and discussion.

The distinction we impose on this conceptual model of poverty exposure – considering duration, intensity, timing, and concentration as distinct elements that interact to determine outcomes – enjoys support in the literature, and the results and arguments found in the papers we build on inform our hypotheses regarding the treatment surface. Wagmiller et al. (2006) identifies several distinct strands of literature on child poverty and life chances: using proxies for permanent status, as in the case of intergenerational mobility models; measuring the effect of cumulative exposure to some circumstance; and capturing differences arising from the timing of a homogeneous experience – each of which include a diverse selection of papers going back decades. Duncan et al. (2010) build on this work, focusing on timing, find a different result based on a finer-grained dataset, consistent with strong timing effects. Both papers emphasize their measurement of poverty duration as a kind of 'dose' of treatment, in contrast to the more common approach of measuring average income or using a measurement at one time as a proxy for long-term income or deprivation, and find that long durations that occur early in life have a much larger effect than shorter ones. Brooks-Gunn and Duncan (1997) and Duncan et al. (1998) consider the intensity of poverty as well, distinguishing between different poverty line shortfalls, as well as making the duration and timing distinctions. Both studies find that more intense poverty is more closely associated with lower cognitive ability and academic achievement, as is earlier poverty and longer durations of poverty. Taken together, they establish a conceptual foundation consistent with the division of poverty attributes we have proposed. None of these papers, however, provide empirical evidence sufficient to identify the shape of a treatment surface for any of these components. Yet, in the context of Cunha et al. (2010), intensity and duration capture essentially the same thing, i.e., the degree of foregone investment that affects a given child due to a lack of financial resources. Our treatment of intensity and duration as distinct is, however, important, because Chetty et al. (2016) are only able to observe duration. Our ability to also account for intensity is one important difference between our data and theirs.

Our multidimensional approach is fairly straightforward, but even in this relatively simple framework, the breadth of our three measures lets us consider alternatives that Chetty et al. (2016) cannot test in their experimental design. Most importantly, the MTO design suppresses the distinction between duration and timing, because the only information available is age at the time of treatment assignment. This means that, for someone treated at 13, the duration of added time in a less poor neighborhood will always be 4, which need not be the case more generally. In the PSID, however, we observe many different duration values for individuals entering poverty at each possible age, and the inverse, which lets us disentangle the relationships between adulthood rank and each of these treatments, and to allow for heterogeneity in each of these relationships as the other value changes.

We summarize our key variables in Table 1. These summary statistics show that a large majority of the children in our sample (discussed in more detail in the next subsection) experience poverty starting fairly early in childhood, and experience short to medium-length poverty spells, both with a fairly large amount of variation based on the variables' respective standard deviations. Intensity is substantially less dispersed, however, with an interquartile range that extends over only about 20 percent of the support, compared to about 30 or 45 percent for age and duration respectively.

3.3 Control Variables

In most of our models, we control for an array of childhood factors which could otherwise bias the relationship between adulthood rank and childhood poverty. Given that we collapse the data into two time periods, childhood and adulthood, in most cases the control variables (summarized in Table 2) are averages of categorical variables describing the head of the individual's household during childhood; for example, 'Married' captures the percentage of the individual's childhood in which the head of her household was married. The exceptions to this definition are our education variables: we define the household head's education as the highest level of education of the household head observed during childhood, which we then collapse into high school, high school and some additional education, and a bachelor's degree or more (labeled 'college').

Our control variables all pertain to the individual's household during childhood, and fall into several categories: demographic variables, education variables, occupation variables, industry variables, employment, and residence. Head of household (HOH) race, the individual's sex, and parental marital status comprise our demographic variable group, and with the exception of the individual's sex, all are averages across each individual's childhood observations. In the case of parents' marital status, the interpretation is straightforward as it captures the percentage of childhood each individual lived in a household with married parents. However in the case of parental race/ethnicity, the measure is somewhat less intuitive. We measure race with a mean value of categorical variables defined based on the value for the head of household, so while most observations have values of zero or one, a decimal value is possible. We maintain this measure because the head of household can change, and we expect that important mechanisms through which head race *per se* could affect outcomes conditional on other observables we condition on include residential sorting, labor market discrimination, and social networks, and ignoring variation in the head of household by using only the modal category may omit this information. Our education variables capture the maximum attainment the head of household achieved at any point during each individual's childhood, so that if the head of household had completed high school but nothing more at a child's birth, but finished college when the child was older, our 'college' binary variable would take a value of one, and the other attainment variables would take values of zero. We code education in this way because we believe it reflects underlying parental human capital better than percentage of time in each category. While education clearly has economic consequences, those occur partly through labor supply and job type, for which we control with other variables.

We also measure occupation and industry variables as averages, and for completeness, we include all of the top-level categories for both, despite the fact that several categories (e.g. farm laborers) are very small both nationally and in our sample. These variables capture the kind of work each individual's head of household did across the individual's childhood. We expect job type to affect both economic security and the kinds of skills and norms that parents transmit to children. We measure employment and location in the same way, capturing parental labor supply with the percentage of the individual's childhood in which her head of household was employed, and location of residence with the percentage of childhood each individual spent in one of the four regions the PSID defines: South, Northeast, North-Central, or West. The majority of individuals in our sample never move, and the majority of nonmovers spent their childhoods in the south, but enough individuals moved between regions during childhood that we believe this more continuous measure of residence captures meaningful information that a binary measure, would not capture.

3.4 Benchmark Replication

In Appendix A, we present a replication exercise in which we estimate models analogous to the age models in Chetty et al. (2016). We present these models to demonstrate that our data and measurements are capable of producing results similar to estimates found in the relocation literature. As we discuss in more detail in Appendix A, our estimates of the Chetty et al. (2016) duration effect fall within the 95 percent confidence interval they compute. Our motivation is simple: if we cannot detect effects similar to the ones they detect using a version of their empirical model, it would be difficult to argue that we reach an explanation for some differences between linear and nonlinear poverty exposure effects in the context of our model. The evidence we present in Appendix A shows that our data are acceptable by this standard; through these regressions, we show that the baseline conclusions from Chetty et al. (2016) are borne out in our PSID sample.

4 Results

4.1 Duration, Intensity, and Age results

In presenting our results, we begin with a discussion of results from a model of rank as a function of duration, intensity, age, and controls only, and then expand that discussion by considering concentration as well. We separate the results in this manner because duration, intensity, and timing are directly relevant to rank, while concentration merits inclusion because it captures aspects of poverty that they omit, reinforcing estimates of the treatments we care about for their own sake.

We begin by fitting a series of parametric, linear in parameters models, which let us estimate $g(\cdot)$ from Equation 2 in a somewhat restrictive, but easily presented and interpreted, manner. In Table 3, we fit a simple additive model of $g(\cdot)$, a model with linear interactions, and an additive model with quadratic terms. In each of these models, duration is consistently significant, and the goodness of fit changes very little: based on the results of an F-test, we cannot reject the null hypothesis of no difference between Model 1 and each of Models 2 and 3. Across specifications, the marginal effect of duration is economically meaningful: an additional year of poverty costs approximately two percentiles of adulthood rank in the linear specifications. In the quadratic specification (Model 3), an additional year is associated with a decrease of approximately three rank percentiles at one year of duration, and approximately half a percentile at sixteen years of duration. When we add controls in Table 4, this pattern persists; duration coefficients are smaller in absolute value, and in the interaction specification (Model 2), the coefficient is only significant at the 10 percent level, but losing approximately one percentile per year is still meaningful, given the width of the income distribution. The persistence of this association across specifications, and its robustness to the inclusion of a wide variety of controls, suggests that poverty matters primarily through duration. Across specifications, the rank-duration relationship does not vary in the other treatments as we hypothesize, and the other treatments are not significantly associated with income rank on their own. Overall, the impression that emerges from these parametric models is that duration matters, but that its effects are linear and relatively small, although certainly not negligible; they are consistent with the Chetty et al. (2016) linearity result, and not consistent with results from the structural models that predict large and nonlinear effects.

We are primarily interested in the shape of $g(\cdot)$, however, and these models only go so far in letting us model that shape. When we plot a surface of fitted values across the support of our treatment variables, we see many predictions outside the permissible range of adulthood income rank (i.e., $[0, 1]$), primarily in regions with limited support, such as high intensity and low duration. In these regions, we are effectively extending a nonlinear function out of sample assuming that the shape of the function is constant, which, based on these non-credible fitted values, we believe to be too strong an assumption.

To produce estimates which respect these constraints on fitted values while allowing an unspecified form of nonlinearities and interactions, we fit a fully nonparametric model, regressing adulthood rank on an unspecified function of age, duration, and intensity.[11] When we fit this model, presented in Figure 5, both age and intensity are found to not influence income rank, which varies only in duration. In other words, the slope of duration does not vary in intensity, which we can see in Figure 5 because the duration curve is constant across the intensity axis, and the same pattern repeats as we evaluate the surface at different values of age. These estimates are significant at the 95 percent confidence level, computed by bootstrapping the mean of the nonparametric function.

In technical terms, this occurs because the optimized value of the bandwidth term, which governs the size of the neighborhoods for which the estimator fits local regressions, is very large, so the estimated conditional mean of rank is constant across these 'smoothed out' variables.[12] This result is not equivalent to a formal specification test, however, so we test the hypothesis that a nonparametric model of rank as a function of duration only is equivalent to a nonparametric model of rank as a function of age, duration, and intensity using standard nonparametric specification tests. We use two tests: a goodness-of-fit test comparable to a nonparametric F test, in which the restricted model is the model with only duration and the unrestricted model includes age, duration, and intensity, and a direct test for the irrelevance of a subset of variables in a nonparametric regression, originally developed by Lavergne and Vuong (2000).[13] Both tests fail to reject the null off no difference between specifications at the 10 percent confidence level, confirming the validity of smoothing these variables out in these regressions.[14]

These results are consistent with the parametric results in suggesting that age and

[11]We fit this model using bandwidths estimated by least squares cross validation, chosen for local constant least squares (LCLS) regression, which constrains the fitted values to fall on $[0, 1]$.

[12]See Hall et al. (2007) for a more detailed treatment of bandwidth estimators smoothing out irrelevant variables.

[13]For a detailed explanation of these tests, and of the nonparametric testing paradigm in general, see Chapter 6 of Henderson and Parmeter (2015)

[14]The p-value in the goodness of fit test is about 0.95, and about 0.12 in the Lavergne and Vuong (2000) irrelevance test, both of which constitute evidence that we cannot reject the null hypothesis of no difference between the models.

intensity do not matter apart from duration, but it does suggest substantial nonlinearity in the rank-duration relationship, as well as a stronger relationship overall. The nonparametric models predict that a child who is poor for six years, the average value of duration in our sample, will reach an adulthood position in the income distribution approximately thirteen percentiles lower than if she had spent only one year in poverty, in contrast to values between six and nine percentiles in the parametric models. At higher levels of duration, the slope flattens substantially, however. They are also more consistent with the structural literature than with the Chetty et al. (2016) linearity result, as they suggest that a disproportionate amount of the effects of poverty comes from the first few years. This makes sense in the context of dynamic complementarity, because the structural models (e.g. Cunha et al. (2010), Agostinelli and Wiswall (2016)) predict that the benefits of additional family resources will be substantially lower after having spent several years in poverty than after having received a small dose. The fact that the slope levels off above the mean, and especially at high values, may seem to contradict the structural results, but the fact that those papers consistently find weaker effects for investments in older poor children reflects a similar leveling off process. They make the argument in terms of timing, but intervening early in a poor child's life also implies intervening at a relatively low value of duration, and that is consistently what the structural literature recommends.

The nonparametric results let us detect nonlinearities without imposing functional form assumptions, while also constraining the fitted values to fall on $[0, 1]$, but they are not conditioned on control variables, which seem to be important based on the parametric results.[15] We build on the fully nonparametric results by fitting a series of partially linear models, in which $g(\cdot)$ is allowed to take an arbitrary form as in the nonparametric model, but it is conditioned on a parametric, additive, linear in parameters function of our control variables. We summarize this model in Figure 6, where we plot the fitted values against duration only, because as in the fully nonparametric models, the other features are not relevant. The same rank-duration pattern that we see in the fully nonparametric models emerges here, and while the parametric models produce similar results – the 95 percent confidence intervals, not pictured for readability, overlap – the slopes do vary locally, particularly at low to medium durations. These are our preferred estimates, because they retain attractive features of the nonparametric model while allowing us to control for the parental and household characteristics summarized in Table 2. This requires the assumption that duration, intensity, and age are additively separable from the control variables. Making this assumption is warranted because we are interested in the marginal effect of deprivation imposed by different features of poverty, above and beyond whatever the inputs provided by the parental characteristics we hold constant.

We provide an alternative visualization in Table 7, in which we plot the mean values of a nonparametric regression of rank on duration and a partially-linear regression of rank on age, duration, intensity, and controls by years of duration, with bootstrapped confidence intervals. The shape of the rank-duration surface is similar between the two models: in both panels, we see sharp nonlinearity at low values of duration, and less

[15]We do not include controls in the nonparametric models because fully nonparametric estimators effectively allow arbitrary interactions and nonlinearities for all variables, which makes isolating and interpreting $g(\cdot)$ difficult. Furthermore, adding additional variables to a nonparametric model dramatically increases the amount of data needed for any degree of precision, and with only 984 observations in our preferred sample, precisely estimating a nonparametric function of between thirty and forty variables is not feasible.

nonlinearities at high values, although in the conditional model, the tighter confidence interval suggests substantial nonlinearity even at high levels of duration. Taken together, these results indicate that the rank-duration relationship is nonlinear, and that the shape of this relationship is robust to the inclusion of our controls.

4.2 Concentration Results

We repeat much of this analysis accounting for concentration as well as duration, intensity, and age. First, we hold concentration constant by using only individuals who were poor in contiguous years only, so that no one moves in and out of poverty. Using this sample substantially reduces our sample size and the support of our treatment variables, particularly duration, so we also fit models using our preferred sample, in which we use the concentration measures discussed in Section 3.2 to control for concentration.

In the contiguous poverty sample, each individual has the lowest possible concentration of poverty years for their duration and age at first exposure, so unobserved differences in concentration cannot bias our estimates. The most substantial difference between the partially linear model results in this sample and the results from the main sample is that intensity is not smoothed out by the kernel selection procedure.[16] This result suggests that leaving concentration unobserved suppresses the relationship between adulthood rank and intensity; individuals with the most concentrated poverty appear to be affected by the intensity of that poverty more than the average individual in the main sample. In economic terms, this is consistent with the predictions of the structural literature: dynamic complementarity (Cunha et al., 2010) suggests that a deeper shock followed closely by additional shocks will have a larger long-term effect, as the deprivation in this period will compound more than it would had the child experienced the same duration, intensity, and starting age spread over several years. This interpretation is consistent with the fact that our estimated duration curves look similar to those from the main sample, because they are conditioned on intensity, and while it would make sense that a more concentrated poverty duration would also produce a compounding effect, that effect is mediated by intensity. Alternatively, this result may reflect the fact that intensity is better measured in this sample, because there is only one poverty spell. It is worth noting that the rank-duration curve remains nonlinear in this subsample, although less so than in the main sample.

We adjust our primary models by adding the standard-deviation based measure of concentration, discussed in Section 3.2, which increases as concentration decreases. Table 5 shows estimates from parametric models including this concentration measure. Adding concentration to these models does not affect the marginal effects of other treatment variables, and in many cases, it does not significantly change the model's fit, according to F-tests comparing the model with concentration to the base I-D-A modes; the same result holds in nonparametric and semiparametric models. Overall, these results suggest that, while the contiguous sample does display different behavior, the omission of concentration in our main results does not affect our findings.

We summarize the fitted values by years of poverty in Figure 10, by spell and years of duration in Figure 11, and by duration and intensity across different numbers of spells

[16] We fail to reject the null hypothesis of no difference between the age, duration, and intensity model and the duration only model in the Lavergne and Vuong (2000) irrelevance test (p = 0.516) at a 5% confidence level, but reject it at the 95 percent confidence level (p=0.001) based on the goodness of fit test. This is inconclusive, so we proceed with these intensity results.

in Figure 12. The averages in Figure 10 follow a more locally linear pattern than the averages in Figure 7, but they are estimated much less precisely, in part because this figure ignores the intensity and spells dimension to present the results simply. Our original result falls well within the 95 percent confidence interval, so while this figure suggests that including spells makes a difference, it does not undermine the original finding. In Figure 11, it is clear that the shape of the duration curve does vary across different numbers of spells, with children with more spells seeing a steeper drop before ten years of duration, but a similar flattening at high values. The lack of support in some duration-spell number bins, along with the width of the 95 percent confidence interval, makes this comparison tenuous, however. We reach a similarly ambiguous conclusion based on the semiparametric regressions presented in Figure 12, in which no clear pattern emerges across numbers of spells, although the general shape of the rank-duration curve persists, albeit with many deviations in the (wavy) surfaces. Taken together, these figures suggest that the number of spells matters in a way concentration does not, but their inclusion does not substantially alter the rank-duration curve.

5 Robustness Checks

5.1 Adjusting for Selection on Observables

Our first robustness assessment is to reduce, if not necessarily eliminate, sample selection bias using propensity score methods developed for continuous treatments, through the use of a Covariate Balancing Generalized Propensity Score (CBGPS) approach developed by Fong et al. (2018). In brief, this method is designed to reduce the bias in a regression of a continuous outcome on a continuous treatment and controls by minimizing the association between the treatments and the controls, and by reducing sensitivity to model misspecification, without making strong distributional assumptions. In our setting, this is helpful because many of the relatively large number of covariates we include as controls are correlated with our treatments. Furthermore, with multiple continuous treatments and many controls, the number of potentially credible specifications is high, so reducing bias due to misspecification also becomes particularly important. We are particularly concerned with the degree to which our estimates of the rank-duration slope are biased by covariate imbalance, as duration consistently has a strong relationship with rank, while the other two treatments do not. To this end, we apply the CBGPS procedure to balance covariates with regard to duration, and not the other treatments.[17]

We re-estimate models of rank as a function of duration adjusting for CBGPS-based weights, designed to consistently estimate the model after adjusting for the CBGPS itself, which we summarize in Figure 13 and Table 6. In Figure 13, we summarize the differences between correlations between each control variables and duration, before and after CBGPS weighting. The height of the second bin indicates that for eighteen of our control variables variables, weighting reduced their correlation with duration by between 0 and .1. In almost all cases, the decrease is positive in absolute value, indicating that the adjustment did, in fact, help to balance the covariates with respect to treatment. The relatively large number of covariates showing only a small difference may suggest the effect was limited, but in absolute terms, the decreases we see are much greater than

[17]To our knowledge, the procedure cannot be applied in a way that adjusts for balance over three covariates.

those in Fong et al. (2018), whose decreases were all between 0 and 0.1 in absolute value, so that they would have no frequency in bins beyond the second, corresponding to larger reductions in imbalance. This occurs primarily because our baseline imbalance was much larger than the imbalance in Fong et al. (2018).

In Table 6, we summarize parametric models equivalent to the parametric models with controls in Table 4, but adjusted for CBGPS weights. The duration coefficients become larger across all specifications, and the fit improves, suggesting that bias from covariate imbalance was, if anything, limiting the size of the rank-duration association. At the same time, however, these models produce implausibly large positive coefficients on intensity, and while this likely has to do with the fact that we can only adjust for imbalance with respect to duration, it does limit the robustness of the CBGPS specifications. Our semiparametric results thus remain our preferred specification, but the CBGPS results do reinforce the semiparametric results by showing that our duration results, if not necessarily the null results for the other treatments, are robust to adjustment for selection on observables.

Adjusting for the CBGPS changes the qualitative interpretation of our results only slightly. This exercise, does, however, serve to enhance the credibility of our baseline estimates of the rank-duration curve, because it establishes that our initial estimates were affected by bias arising from misspecification and imbalance in covariates given treatment only to a limited degree.

5.2 Adjusting for sample selection bias

In Appendix B, we address concerns with our sampling criteria and measurements which could bias our results. In Appendix B, we address the possibility that our use of a left-censored dataset, in which every observation has nonzero values for our treatment variables, biases our results by ignoring selection into zero or nonzero treatment. Our primary results rely on a sample with only treated individuals, but if individuals are sorted into zero and nonzero poverty groups endogenously, our results would be biased through the omission of modeling this process. To address this concern, we apply a Heckman correction, and find that our parametric results are very similar to regressions run using the full dataset.

6 Discussion

We make two related contributions: we show that children's poverty is closely related to adulthood rank through cumulative disadvantage imposed through duration, and we find that, contrary to our expectations, the timing and intensity of poverty have very little effect on the model given duration, either on their own or by shifting the duration slope. Our substantive finding is that poverty duration is, conditioning on a wide array of observables and using a sample designed to minimize bias from various sources of measurement error, strongly associated with a lower adulthood rank, and the rank cost of more duration grows very quickly at low values, but then tapers off. Our model predicts that, all other things being equal, someone who spent her whole childhood in poverty would achieve an adulthood income rank approximately 23 percentile ranks lower than she would have otherwise. The majority of this rank reduction occurs in the first six years of poverty (which need not occur in the first six years of life). After six years of poverty, there is a relatively flat period through eleven years of poverty, then a more or less linear reduction from eleven to seventeen years. To a certain extent, this shape is consistent with predictions from the structural literature (Heckman, 2008; Cunha et al., 2010; Agostinelli and Wiswall, 2016) regarding the timing of shocks, because absorbing a medium to long spell of childhood poverty is impossible if one enters poverty for the first time at an older age. Similarly, the relatively flat slope for children with a relatively long poverty duration suggests that a non-marginal reduction in poverty is necessary to improve their adulthood outcomes.

How, then, do our results help to bridge the structural, skill-based literature with the experimental treatment effect literature? Our results reinforce the policy implications of research that finds a particularly large effect of family resources on low-resource households. This research suggests that interventions, such as augmenting parental resources through the earned income tax credit (Dahl and Lochner 2012) or providing compensatory investments in skills (Cunha and Heckman (2010), Agostinelli and Wiswall (2016)), would be most effective for children who are likely to experience, a relatively low to moderate degree of poverty duration, because as our results show, this initial period is when the negative effects of poverty are most severe. Our result is weakly age-neutral, because a child in her fourth or fifth year of poverty could be any age, but the general principle of targeting relatively young and disadvantaged children, prior to any sustained exposure to poverty, holds.

On the other hand, our model provides some insight into why Chetty et al. (2016) find a linear age effect despite the early childhood literature. Treatment in the Chetty et al. (2016) case is movement away from poverty, which roughly corresponds to an end of "treatment" in our model – i.e., preventing a child from reaching poverty duration of 13 or more years. In this region, however, the rank-duration slope is more or less linear and is remarkably shallow: the confidence interval includes a straight line with a slope of approximately -0.0675 ranks per year, lower than the linear slope we estimate in our parametric model, and much lower than the slopes we estimate elsewhere in the semiparametric and nonparametric models. Though our estimated rank-duration relationship is nonlinear, we see a relatively flat relationship at longer poverty spells. In this region, however, the rank-duration slope is more or less linear and is remarkably shallow: the confidence interval includes a straight line with a slope of approximately -0.0675 ranks per year of poverty, lower than the linear slope we estimate in our parametric model, and much lower than the slopes we estimate elsewhere in the semiparametric and nonparametric models. The Chetty et al. (2016) interpretation - that a fixed disruption cost of moving overwhelms the sum of the benefits for these children – also fits our estimates. At the same time, we believe that the overall nonlinearity in the rank-duration relationship stems from the differences between neighborhood and household poverty. Page and Solon (2003) show that within-household income covariances for children are about twice as high as within-neighborhood covariances, suggesting that while neighborhoods matter, families matter more. If MTO treatment had a small to null effect on children's parents – which earlier results (e.g.Kling et al. 2007,Ludwig et al. 2013,Jacob et al. 2015) find is the case – treatment changes neighborhood inputs while leaving the household characteristics, which include the propensity to enter poverty, constant. The structural literature does not generally draw a distinction between neighborhood and household inputs, instead specifying general investment, which may occur through positive neighborhood inputs; yet this literature does predict that children who incur significant disadvantage will gain less from new investments.[18] Thus, in the Chetty et al. (2016) context, treatment pro-

[18]The notable exception to this characterization is Agostinelli (2018), which incorporates both house-

vides improved neighborhood inputs, but does not change the more important household inputs, and so the improved investment is not sufficient to take advantage of nonlinearities in returns to overall investment. Furthermore, our finding that the duration slope is constant across age and poverty intensity removes the objection that the MTO results conflate age and duration, as our results suggest that this conflation does not matter (although this is not equivalent to claim that the timing of poverty does not matter).

We control for a wide array of factors, but there are some parent and child characteristics we cannot control for well. To the extent that our controls account for the channels through which these factors affect adulthood rank, we can manage their influence partially. One example of such a variable is parental health. Our controls limit threat of bias from parental health variables, in part because of the breadth of poverty: in many cases, unobserved parental health variables will affect adult outcomes through our poverty variables (e.g. chronic conditions suppressing wages). In other cases, these shocks, e.g. a serious acute illness in a parent, could be realized in employment, which we capture with our employment and job type (i.e. industry and occupation) variables. Factors which reduce important parental factors such as income or time availability, but that do not affect one of our poverty measures, are a more serious concern. The persistence of the rank-duration relationship, across many specifications and conditional on many controls which should affect household resources, helps to mitigate it, however. This does not eliminate the possibility of bias from unobserved, sub-poverty disadvantage, but it would have to be the case that the poor households were consistently near poverty in non-poverty years for this to threaten our results. Doubling the poverty threshold and recalculating our poverty measures only slightly increases the average duration of poverty, however - shifting it from about 6 to about 7 years - so this does not appear to be the case in our sample. At the child level, most notable omission is some measure of baseline ability: microeconomic models of skill formation, from Becker and Tomes (1979) to more recent work by Cunha and Heckman (2010) and Agostinelli and Wiswall (2016) emphasize the importance of initial ability stocks, affected by genetics, prenatal health, and random chance. For this to cause large bias, however, it would have to be the case that children more likely to experience adverse peri-natal conditions or poor genetic endowments are more likely to be poor conditional on our observables, and while this in fact seems reasonable, we also expect that our controls - particularly education, marital status, and employment - mitigate this source of bias.

7 Conclusion

In this paper, we have explored the relationship between poverty and adulthood rank in the income distribution, by focusing on poverty duration, intensity, and timing. The literature on childhood skill suggests that these features are complementary, conferring increasingly large skill penalties as poverty increases. Consequently, we expect that separating these factors allows us to gain insight into the ways in which different facets of poverty affect individual economic outcomes. We find that the duration of poverty, is the only component of poverty that matters in the long term, and that the slope of the rank-duration model is nonlinear in the manner suggested by the structural literature.

hold and neighborhood effects, and finds that parental investments and peer effects are substitutes within a period, but that parental investments in one period improve both the quality of peers and the effect of those peers

This result reinforces the policy recommendations of the skill-development literature, i.e. that investments in children who have experienced relatively low amounts of poverty, but are at risk of experiencing more, are the most efficient. Our results also help to contextualize differences between Chetty et al. (2016) and this structural literature, as we find a nonlinear relationship between poverty duration and income rank for children with short to medium poverty spells, but an approximately linear and much flatter relationship for children with long poverty spells.

References

Agostinelli, F. (2018). Investing in children's skills: An equilibrium analysis of social interactions and parental investments. *Working Paper*.

Agostinelli, F. and M. Wiswall (2016). Estimating the technology of children's skill formation. *NBER Working Paper no. 22442*.

Almond, D. and J. Currie (2011). Human capital development before age five. *Handbook of Labor Economics 4*, 1315–1486.

Becker, G. and N. Tomes (1979). An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy 87*(6), 1153–1189.

Becker, G. and N. Tomes (1986). Human capital and the rise and fall of families. *Journal of Labor Economics 4*(3), 1–39.

Blanden, J. (2013). Cross-country rankings in intergenerational mobility: a comparison of approaches from economics an sociology. *Journal of Economic Surveys 27*, 38–73.

Brooks-Gunn, J. and G. Duncan (1997). The effects of poverty on children. *The Future of Children 7*(2), 55–71.

Brooks-Gunn, J., G. Duncan, P. K. Klebanov, and N. Sealand (1993). Do neighborhoods influence child and adolescent development? *American Journal of Sociology 99*(2), 353–395.

Chetty, R., N. Hendren, and L. Katz (2016). The effects of exposure to better neighorhoods on children: new evidence from the Moving to Opportunity experiment. *American Economic Review 106*(4), 855–902.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics 129*(4), 1553–1623.

Chetty, R., N. Hendren, P. Kline, E. Saez, and N. Turner (2014). Is the United States still a land of opportunity? Recent trends in intergenerational mobility. *No. w19844. National Bureau of Economic Research*.

Chyn, E. (2016). Moved to opportunity: The long-run effect of public housing demolition on labor market outcomes of children. *Working Paper*.

Corak, M. (2006). Do poor children become poor adults? Lessons from a cross country comparison of generational earnings mobility. *Institute for the Study of Labor Working Paper*.

Cunha, F. and J. Heckman (2010). Investing in our young people. *IZA Discussion Paper Series, No. 5050*.

Cunha, F., J. Heckman, and S. Schennach (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica 78*(3), 883–931.

Dahl, G. B. and L. Lochner (2012). The impact of family income on child achievement: evidence from the earned income tax credit. *American Economic Review 102*(5), 1927–1956.

Duncan, G., W. J. Yeung, J. Brooks-Gunn, and J. Smith (1998). How much does childhood poverty affect the life chances of children? *American Sociological Review 63*(3), 406–423.

Duncan, G., K. Ziol-Guest, and A. Kalil (2010). Early-childhood poverty and adult attainment, behavior, and health. *Child Development 81*(1), 306–325.

Fong, C., C. Hazlett, and K. Imai (2018). Covariate balancing propensity score for a continuous treatment: application to the efficacy of political advertisements. *Annals of Applied Statistics 12*(1), 156–177.

Foundation, A. E. C. (2018). Children in poverty. `https://datacenter.kidscount.org/data/tables/43-children-in-poverty-100-percent-poverty?loc=1&loct=1#detailed/1/any/false/870,573,869,36,868/any/321,322`.

Haider, S. and G. Solon (2006). Life-cycle variation in the association between current and lifetime earnings. *American Economic Review 96*(4), 1308–1320.

Hall, P., Q. Li, and J. Racine (2007). Nonparametric estimation of regression functions in the presence of irrelevant regressors. *Review of Economics and Statistics 89*(40), 784–789.

Heckman, J. (1979). Sample selection bias as a specification error. *Econometrica 47*(1), 153–161.

Heckman, J. (2008). Schools, skills, and synapses. *Economic Inquiry 46*(3), 289–352.

Heckman, J. and S. Mosso (2014). The economics of human development and social mobility. *Annual Review of Economics* (6), 689–733.

Henderson, D. and C. Parmeter (2015). *Applied Nonparametric Econometrics*. Cambridge University Press.

Hertz, T. (2007). Trends in the intergenerational elasticity of family income in the United States. *Industrial Relations: A Journal of Economy and Society 46*(1), 22–50.

Jacob, B., M. Kapustin, and J. Ludwig (2015). The impact of housing assistance on child outcomes: evidence from a randomized housing lottery. *Quarterly Journal of Economics 130*(1), 465–506.

Katz, L., J. Kling, and J. Liebman (2001). Moving to opportunity in Boston: early results of a randomized mobility experiment. *Quarterly Journal of Economics 116*(2), 607–654.

Kling, J., J. Liebman, and L. Katz (2007). Experimental analysis of neighborhood effects. *Econometrica 75*(1), 83–119.

Koball, H. and Y. Jiang (2018). Basic facts about low income children. `http://www.nccp.org/publications/pdf/text_1194.pdf`.

Lavergne, P. and Q. Vuong (2000). Nonparametric significance testing. *Econometric Theory 16*(4), 576–601.

Li, Q. and J. S. Racine (2007). *Nonparametric Econometrics: Theory and Practice.* Princeton University Press.

Ludwig, J., G. Duncan, L. Gennetian, L. Katz, R. Kessler, J. Kling, and L. Sanbonmatsu. (2013). Long-term neighborhood effects on low-income families: evidence from Moving to Opportunity. *American Economic Review 103*(3), 226–231.

Mazumder, B. (2005). Fortunate sons: new estimates of intergenerational mobility in the United States using social security earnings data. *Review of Economics and Statistics 87*(2), 235–255.

OECD (2018). OECD family database: child well-being module. `http://www.oecd.org/els/soc/oecdfamilydatabasechildwell-beingmodule.htm`.

Page, M. and G. Solon (2003). Correlations between brothers and neighboring boys in their adult earnings: the importance of being urban. *Journal of Labor Economics 21*(4), 831–855.

Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyer, J. Pacas, and M. Sobek (2018). IPUMS USA: Version 8.0 (dataset).

Smith, J. (2009). The impact of childhood health on adult labor market outcomes. *Review of Economics and Statistics 91*(3), 478–489.

Solon, G. (1992). Intergenerational income mobility in the United States. *American Economic Review 82*(3), 393–408.

Solon, G., M. Corcoran, R. Gordon, and D. Laren (1991). A longitudinal analysis of sibling correlations in economic status. *Journal of Human Resources 26*(3), 509–534.

Wagmiller, R., M. C. Lennon, L. Kuang, P. Alberti, and J. L. Aber (2006). The dynamics of economic disadvantage and children's life chances. *American Sociological Review 71*, 847–866.

A Replication Exercises

A.1 Replication Models

We begin by replicating measures of intergenerational mobility and the effect of poverty exposure analogous to those found in the literature, particularly in Chetty et al. (2014) and Chetty et al. (2016). This is necessary because we need to establish the extent to which we are able to replicate (at least qualitatively) established findings, if we are to argue that we can established results them using the broader, if less experimental, observational data we employ.

First, we replicate intergenerational elasticity results (Solon, 1992; Mazumder, 2005) and intergenerational rank-rank correlation results (Chetty et al. 2014 Chetty et al. 2014) results, using our data on household income (measured in real 2012 dollars, to maintain consistency with Chetty et al. 2016) and rank in the distribution of household incomes. We fit the standard regression equation

$$y_i = \alpha_0 + y_i^{\pi} + \epsilon, \tag{3}$$

where y refers to either logged incomes in the intergenerational elasticity model, or ranks in the rank-rank model, i indexes individuals, and the superscript π denotes a value in person the parental generation relative to person i.

Next, we replicate the regression that establishes the linearity result in Chetty et al. (2016). In that paper, the authors estimate the marginal effect of an additional year spent in a poor neighborhood using treatment dummies and the age at randomized treatment assignment, in the following equation:

$$y_i = \alpha + \beta_{E0} Exp_i + \beta_{S0} S8_i + \beta_{EA} Exp_i AgeRA_i + \beta_{SA} S8AgeRA_i + s_i\gamma + \epsilon_i. \tag{4}$$

In their notation, Exp and $S8$ denote treatment assignment, either to the full experimental treatment of vouchers to low-poverty neighborhoods, or to the 'Section 8' treatment, in which households were assigned vouchers which allowed them to move in to subsidized housing, often in higher poverty neighborhoods. This equation reproduces the Chetty et al. 2016 equation. $AgeRA$ denotes the individual's age at assignment into either group, y refers to a number of outcomes, including income at age 25, and i indexes individuals. The value of interest is β_{EA}, the marginal effect of being assigned to the experimental group a year later, which roughly reflects the cost of spending an additional year in a poor neighborhood, compared to a low poverty one. In their preferred specification, this coefficient has a value of -723.7, in a regression on adulthood income, so that it is interpretable in dollar terms. This means that the result implies that an additional year in a poor neighborhood costs about $724 in adulthood income, a substantial, but not overwhelming, value.

While we cannot replicate this regression precisely, we can produce an object with a similar interpretation, that is, the cost of an additional year of poverty (although we cannot control the alternative – in the Chetty et al. 2016 case a low poverty neighborhood – the way Chetty et al. 2016 do). In the following regression equation,

$$faminc_i = \alpha + \beta_1 Duration + \beta_2 Age + \epsilon_i \tag{5}$$

β_2 captures the marginal effect of an additional year of poverty on adulthood family income. We include age at first exposure as well, because, as we discuss above, the

Chetty et al. 2016 model measures age and duration simultaneously. To approximate their regression as faithfully as possible, we must control for this aspect of poverty exposure.

A.2 Replication Results

Our results are shown in Tables 7 and 8, where we present the standard intergenerational elasticity model applied to our data and a rank-rank mobility model following Chetty et al. (2014), and models following the formula in Equation 4, which produces a duration coefficient comparable to the age effect in Chetty et al. (2016). We report models with and without controls, and age-duration interactions. The results in Table 7 suggest a much lower degree of mobility than standard estimates – between 0.4 and 0.5 for the intergenerational elasticity, and approximately 0.35 for the rank-rank correlation – but neither is so large as to suggest that our data are incomparable with samples commonly used in this literature. This sample includes only children who experienced some poverty during childhood, which likely tightens the income distribution in these regressions, so that mobility is lower in this sample than in the more general sample most papers in this literature use. In Table 8, our estimates on the coefficients on poverty duration are also substantially higher than the Chetty et al. (2016) estimates of the effect of an additional year in a poor neighborhood ($\sim$ 724), but only in the first two models, which do not include controls.[19] The marginal effect of an additional year of duration includes the Chetty et al. estimate in its 95 percent confidence interval when we include controls, so while our point estimates are somewhat larger in absolute value than theirs, we believe they are qualitatively comparable.

Our objective in fitting these models is not to compete with the Chetty et al. (2016) or Chyn (2016) estimates, but rather to establish the extent to which our data match theirs. We see that our data leads us to similar conclusions, particularly in Models 3 and 4 in Table 8, where despite using different data sources, we obtain similar estimates of the effect of additional poverty exposure.

B Adjustment for Censored Treatment

For the results reported in Section 4, the dataset includes only individuals who never experienced poverty as children. In this section, we present models of duration fit to the whole sample, that includes PSID individuals who never experienced poverty. On the one hand, excluding individuals in the PSID that never experience poverty makes our definition of duration, intensity, and age clearer. On the other hand, it is possible that this sampling approach leads to a sample selection bias. To estimate our model using this full sample, while addressing the fact that many individuals who never experience poverty have a zero-value of treatment, we use the correction from Heckman 1979. As Table 9 shows, fitting regressions of rank on duration and controls yields a substantially different result in the full sample compared to the treated-only sample, but when we control for the Inverse Mills Ratio, we see that the estimates of the marginal effect of duration are very similar between the base dataset and the full dataset. This shows that our omission of zeroes in our preferred results in Section 4 does not bias our findings. Figure 15 shows

[19] In the relocation studies, the researchers obtain treatment exogeneity through their design. While we cannot claim to have achieved comparable bias reductions, controlling for a range of household covariates is the closest approximation available to us, because in those studies, household covariates are the confounders that randomization circumvents.

that this holds for the quadratic duration specification as well, because the fitted values of the regression on the full sample with the correction are extremely close to the fitted values from the treated sample when plotted against duration measured in years.

It is important to note that we use only our control variables as independent variables in the first stage of the Heckman procedure. We do this because we do not have a variable, or set of variables, that we can credibly exclude from the second stage of the model. This lack of an exclusion restriction does not, hinder the model's ability to adjust for selection into nonzero duration, provided that the Inverse Mills Ratio is sufficiently nonlinear over its support in the data. Figure 16 shows that the Inverse Mills Ratio is nonlinear, and so we expect identification to hold.

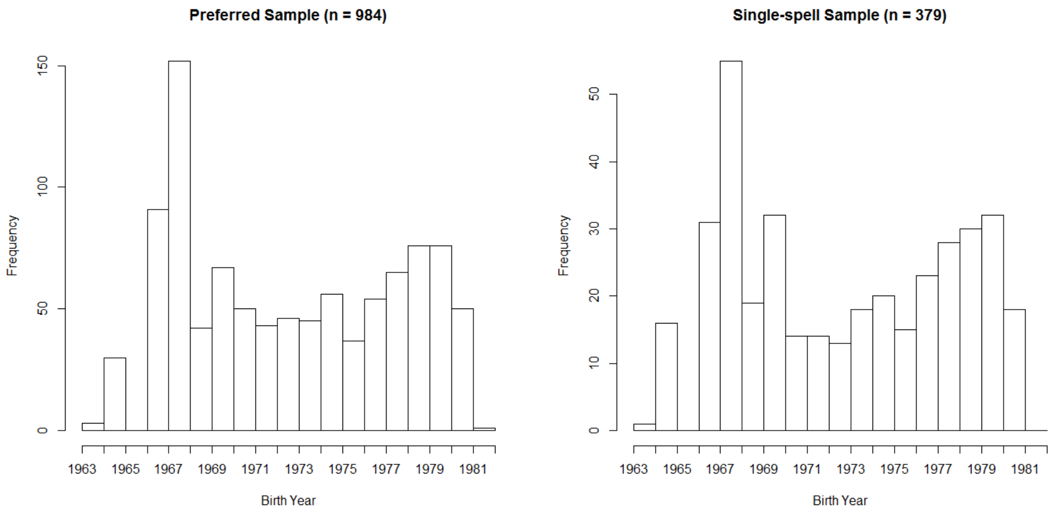

Figure 1: Frequencies of observations by birth year and sample

Figure 2: Bivariate Scatterplots of the three treatment variables

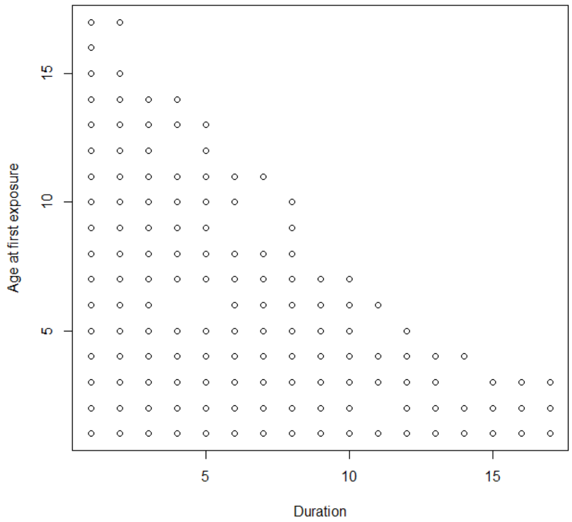

A. Duration and age

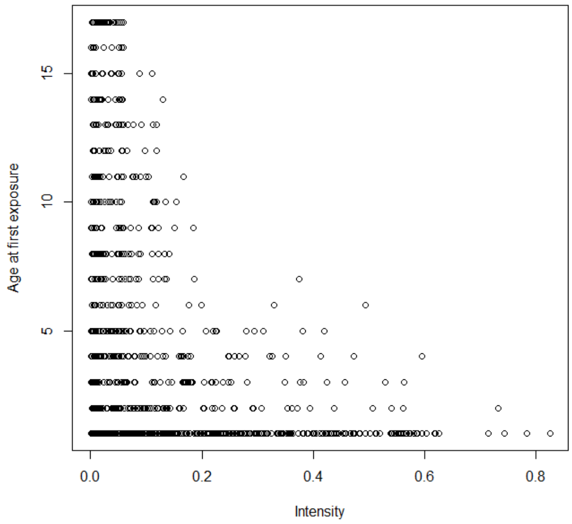

B. Intensity and age

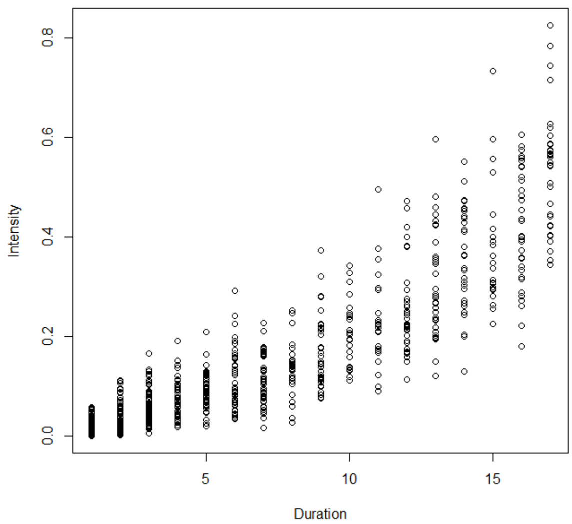

C. Duration and intensity

Figure 3: Bivariate scatterplots of rank by treatment variables

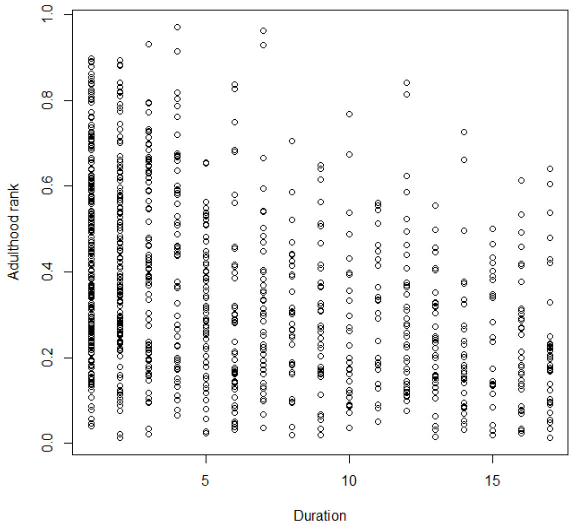

A. Rank and duration

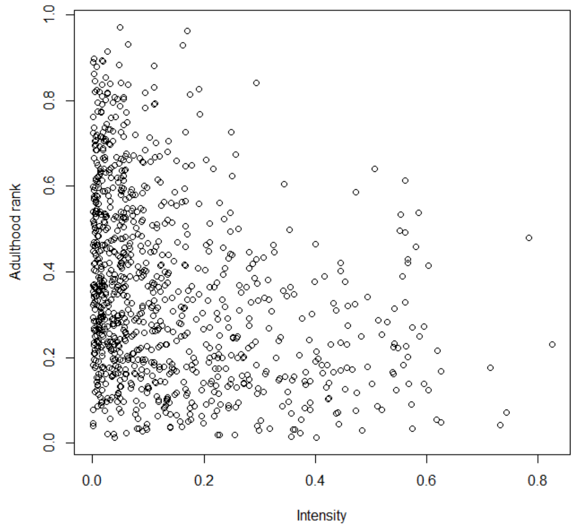

B. Rank and intensity

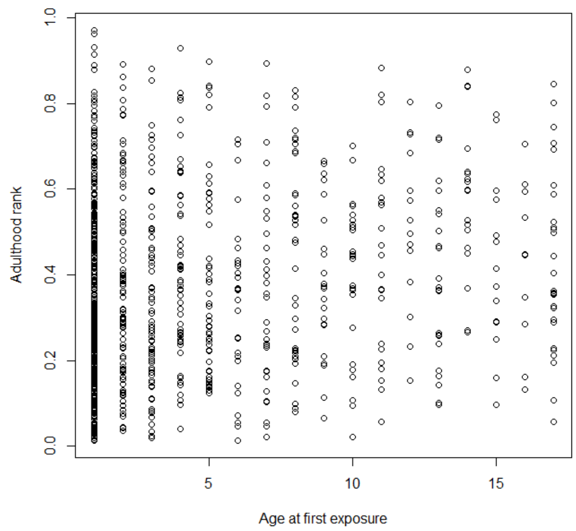

C. Rank and age

Figure 4: Frequency of numbers of poverty spells during childhood

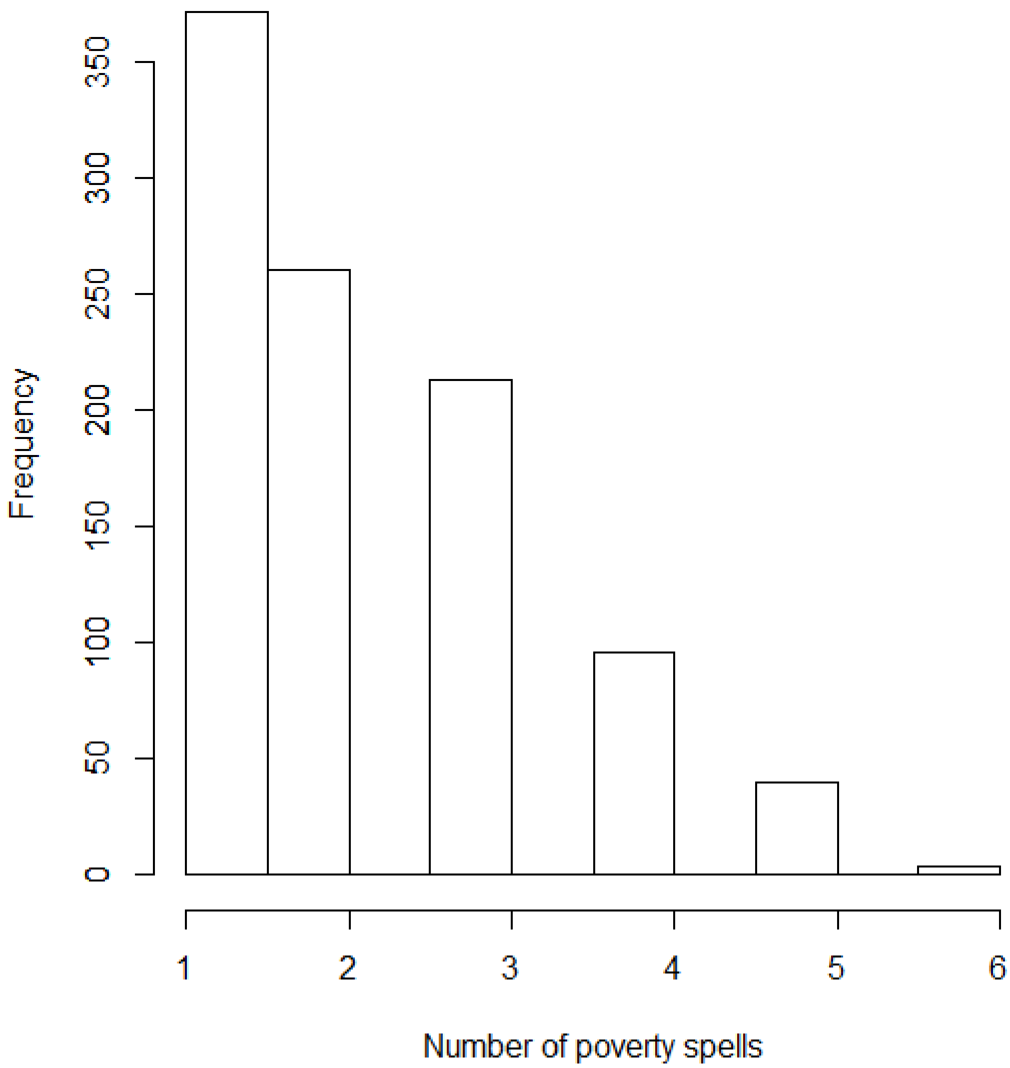

Figure 5: Fitted values by intensity and duration, from nonparametric regression of rank on treatments

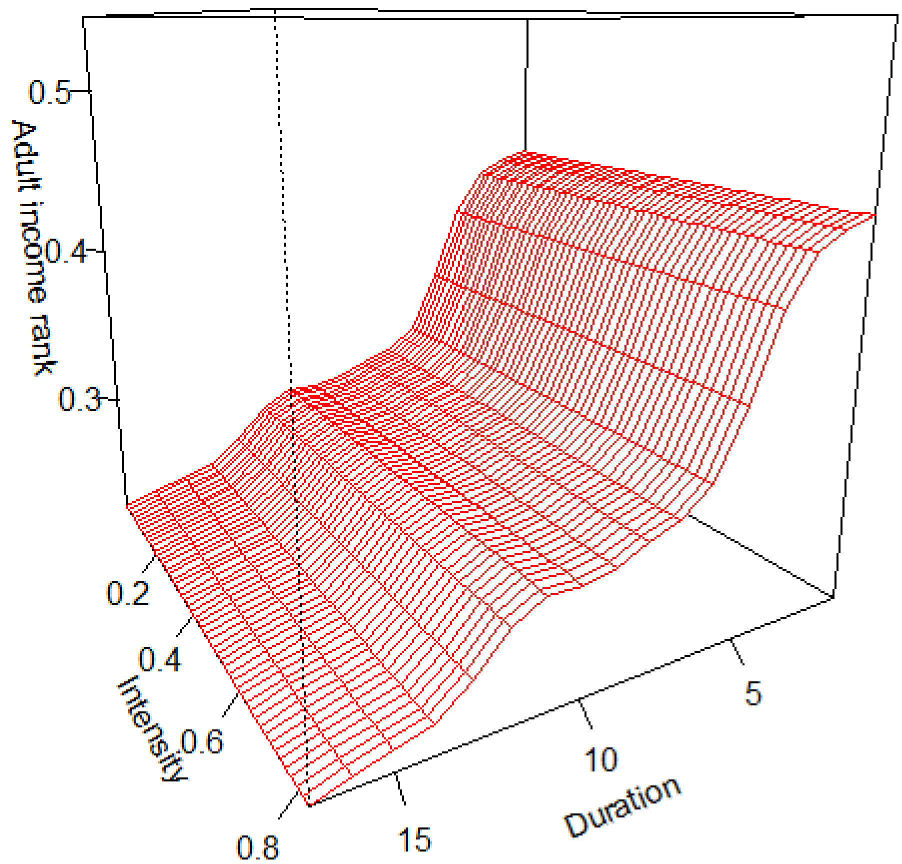

Figure 6: Conditional duration models by polynomial degree, overlayed with the partially linear model

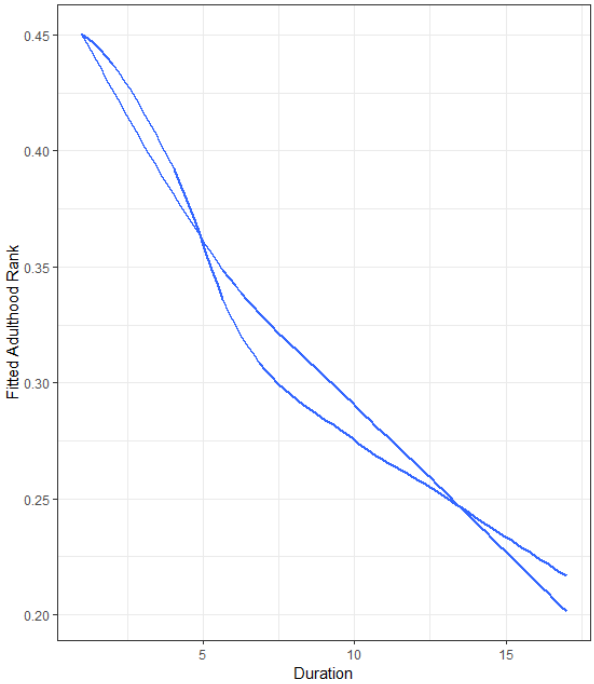

A. Linear parametric

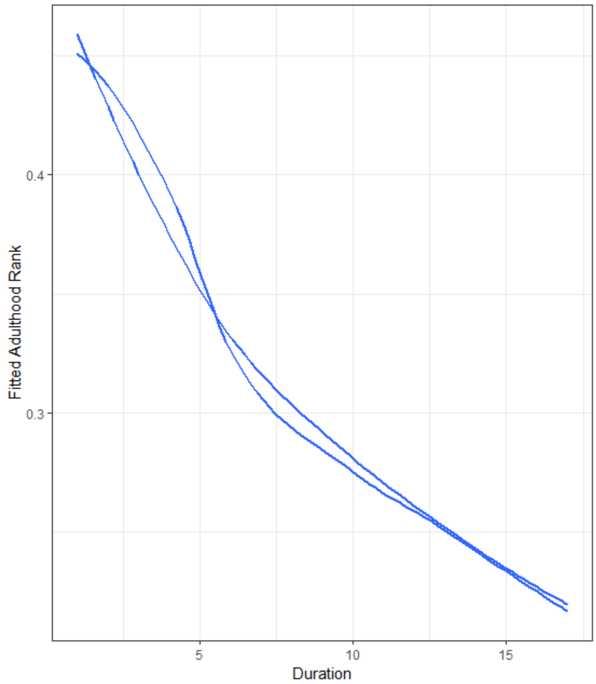

B. Quadratic parametric

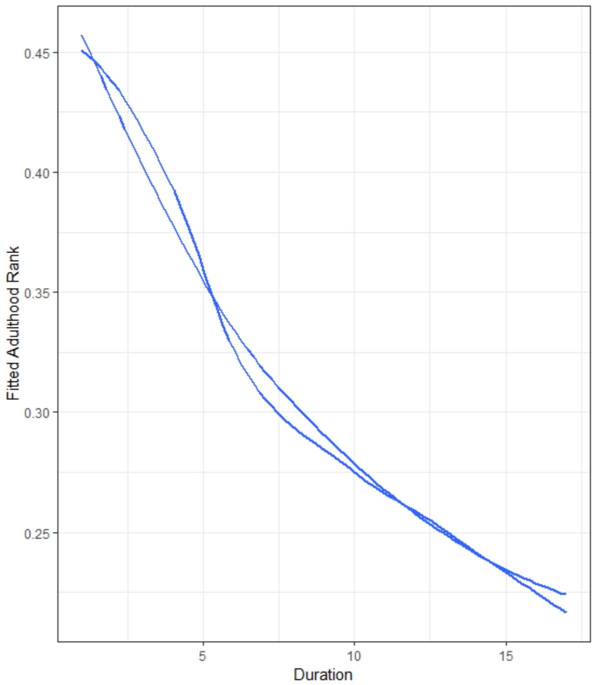

C. Cubic parametric

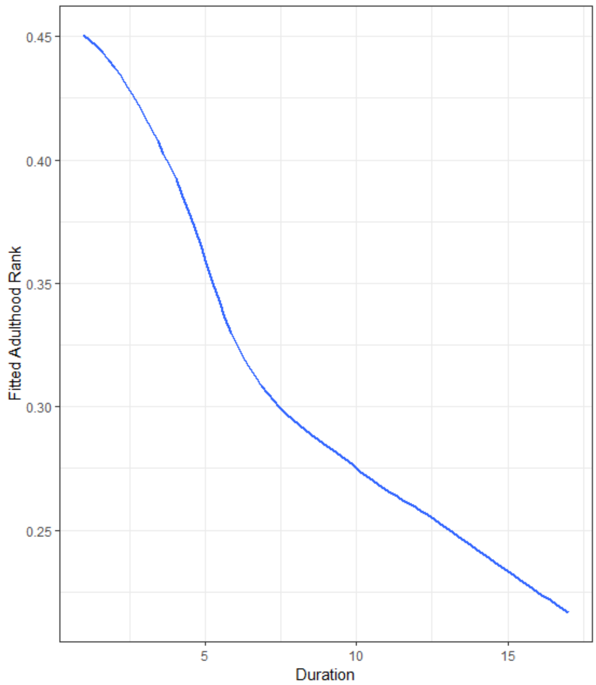

D. Partially linear model only

Figure 7: Mean fitted values by years of duration with bootstrapped confidence intervals

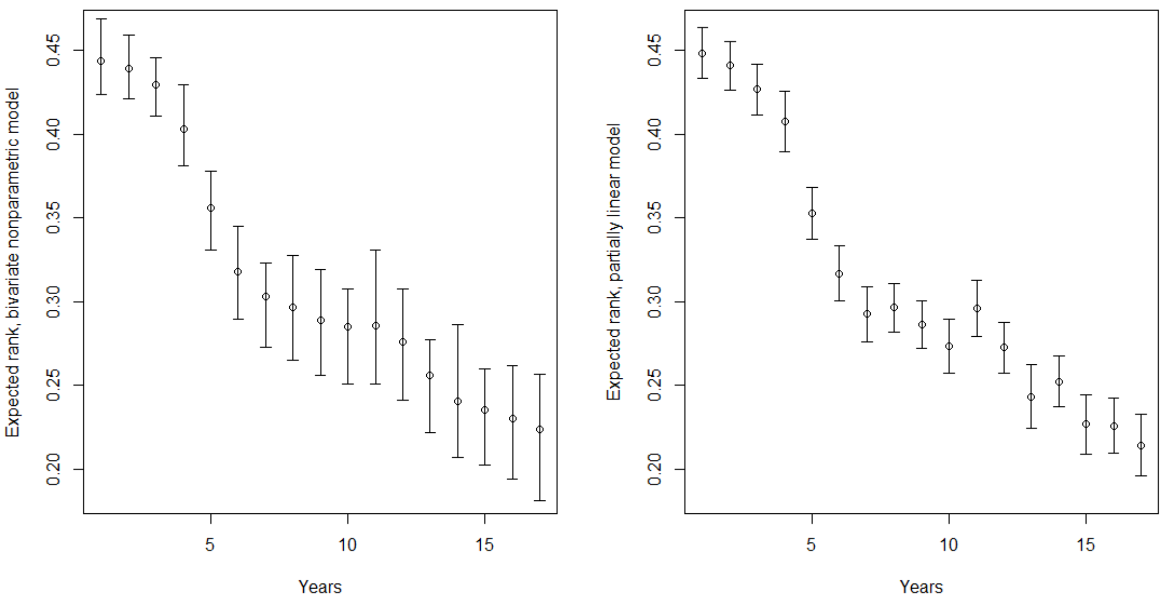

Bivariate nonparametric duration model Partially linear model

Figure 8: Conditional duration models by polynomial degree, overlayed with partially linear model (contiguous sample)

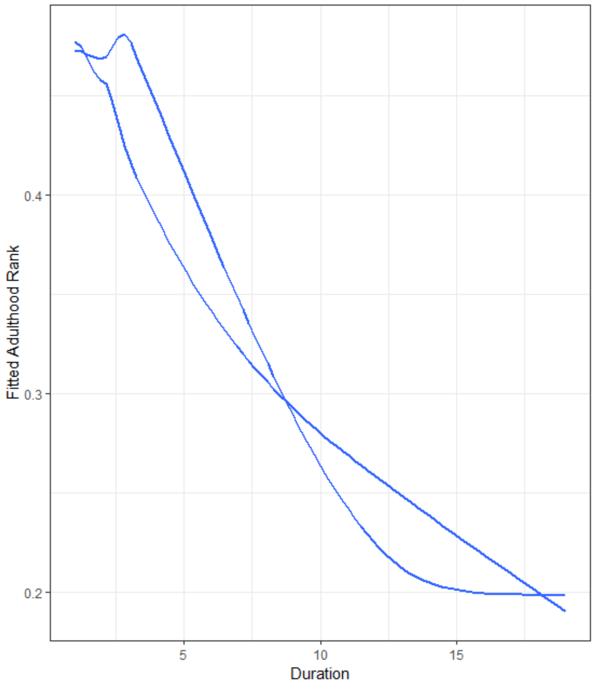

A

. Linear Parametric

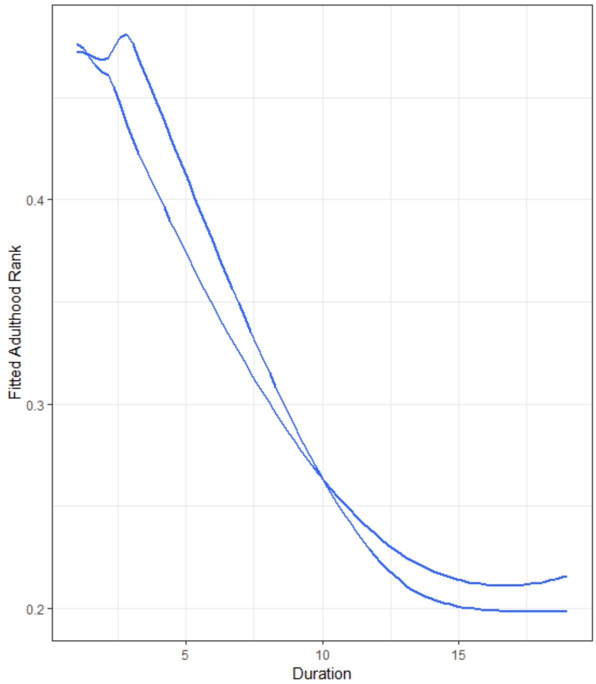

B. Quadratic Parametric

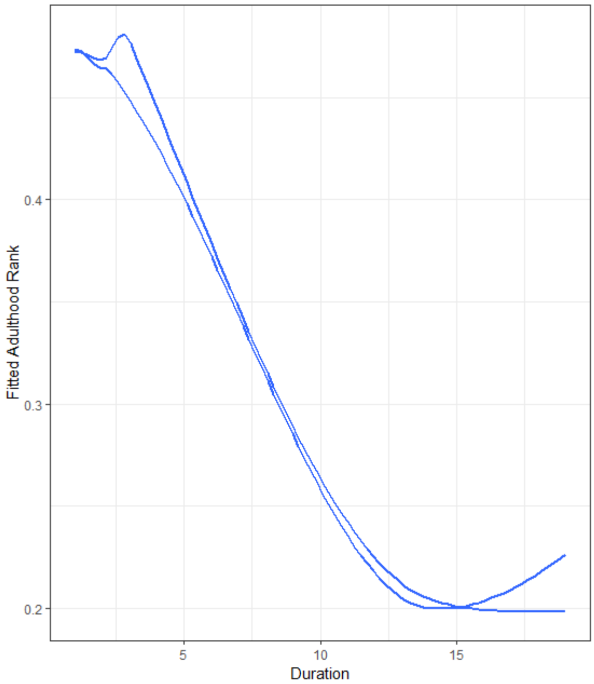

C. Cubic Parametric

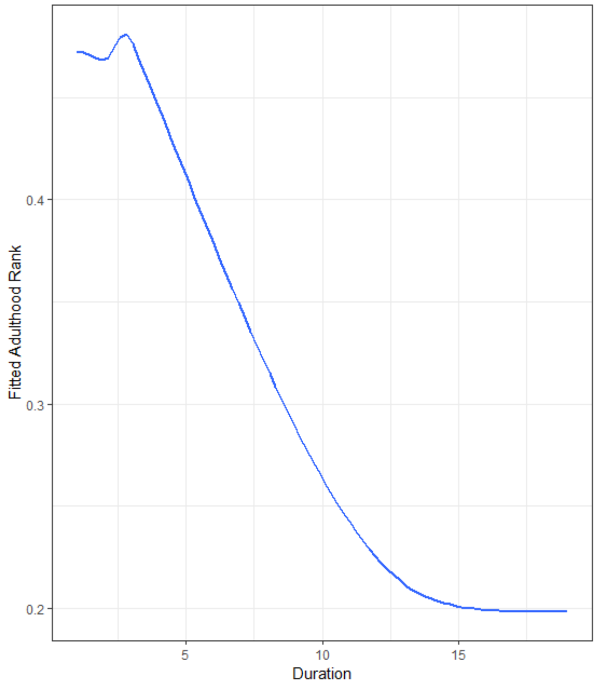

D. Partially Linear Model

Figure 9: Conditional intensity models by polynomial degree, overlayed with partially linear model (contiguous sample)

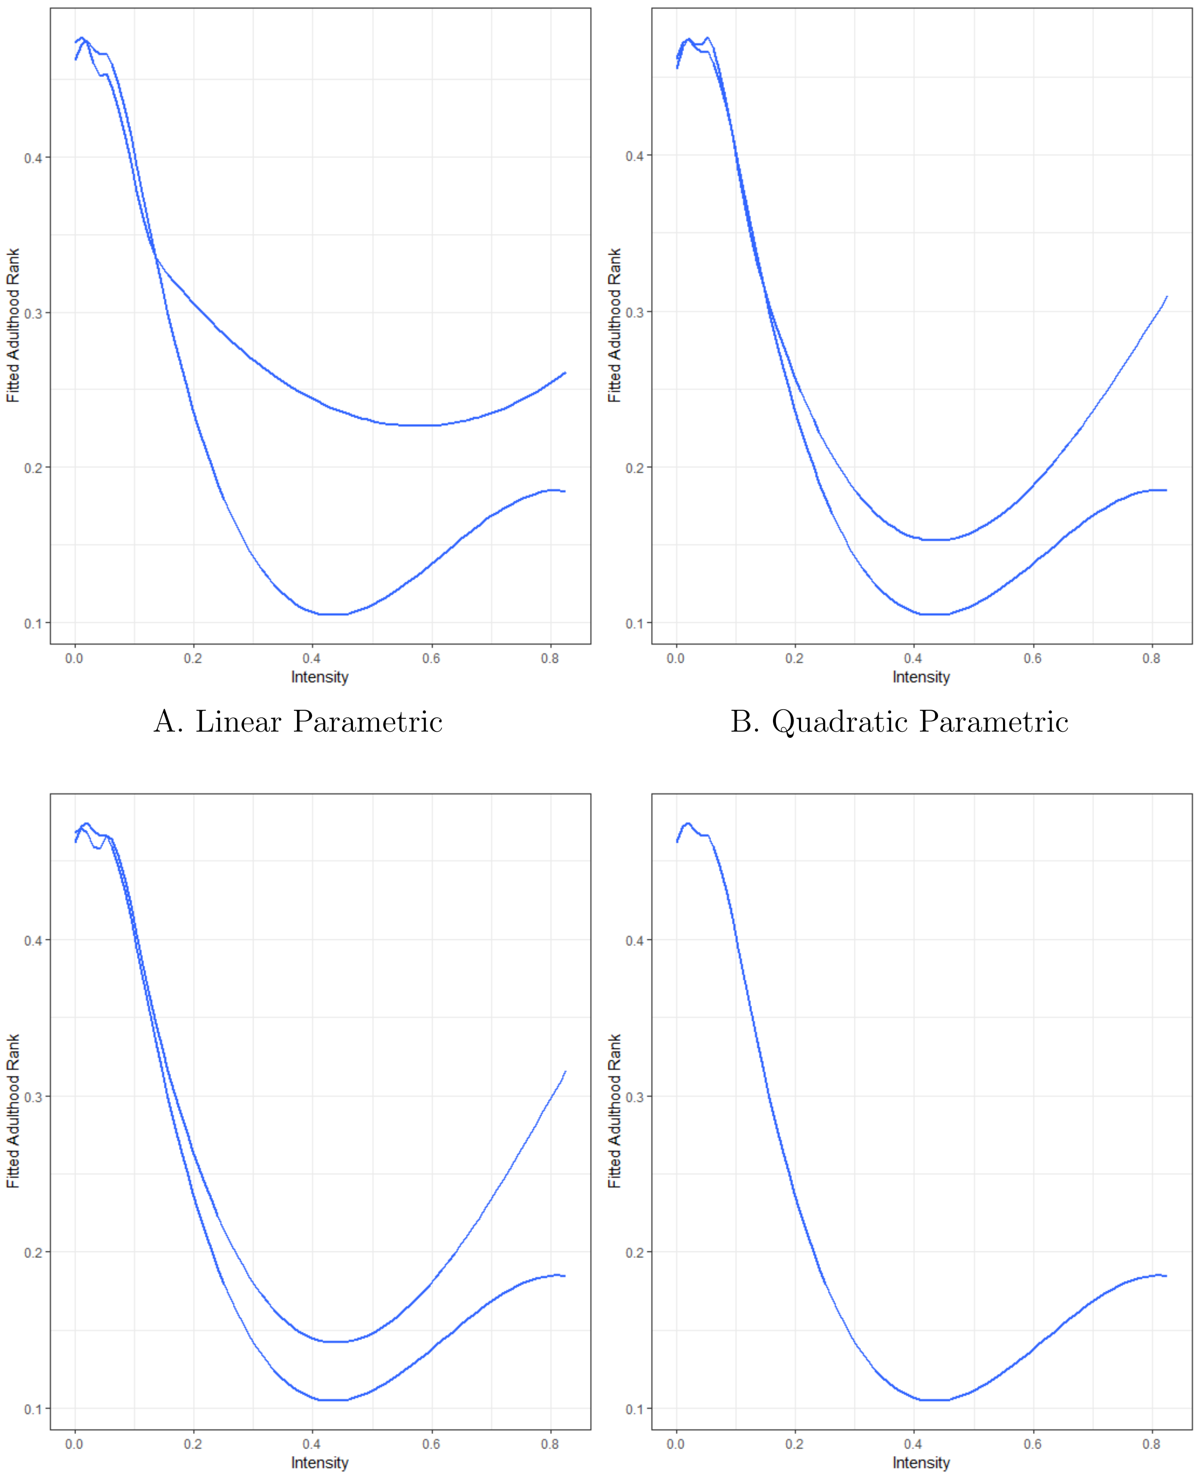

A. Linear Parametric

B. Quadratic Parametric

C. Cubic Parametric

D. Partially Linear Model

Figure 10: Partially linear model including spells, averaged by year

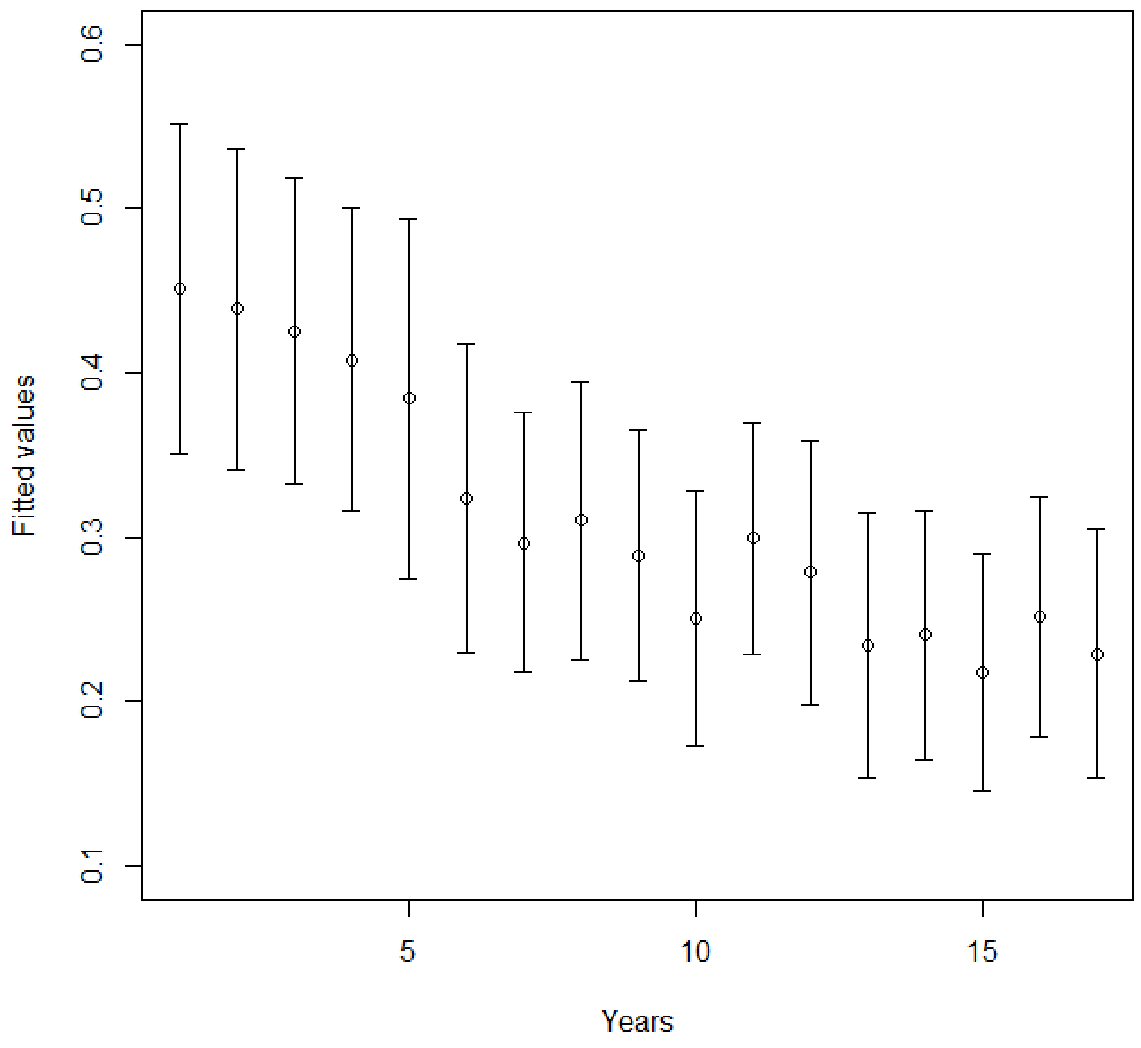

Figure 11: Partially linear model means, by spell and duration

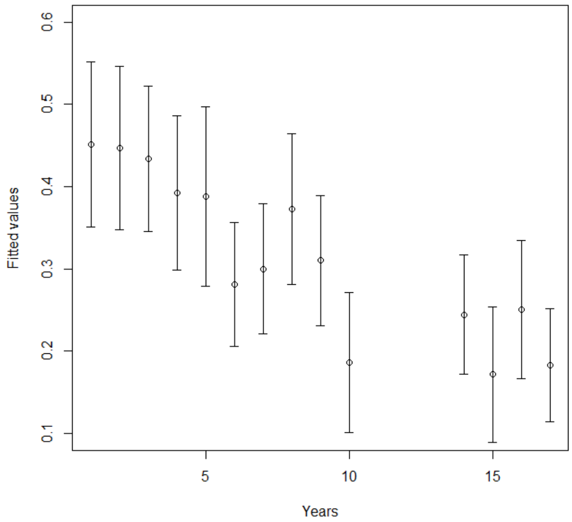

A. One spell

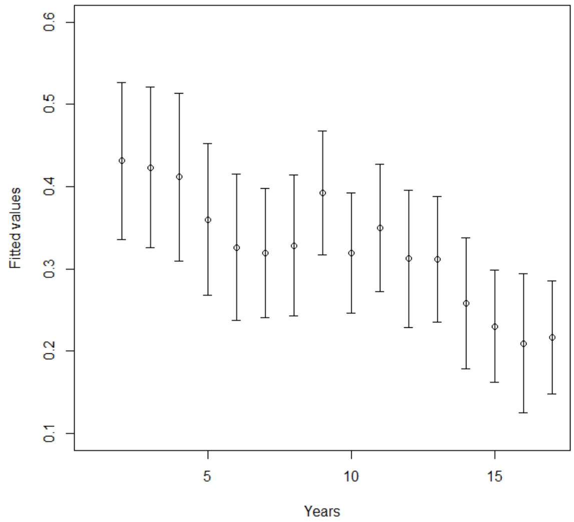

B. Two spells

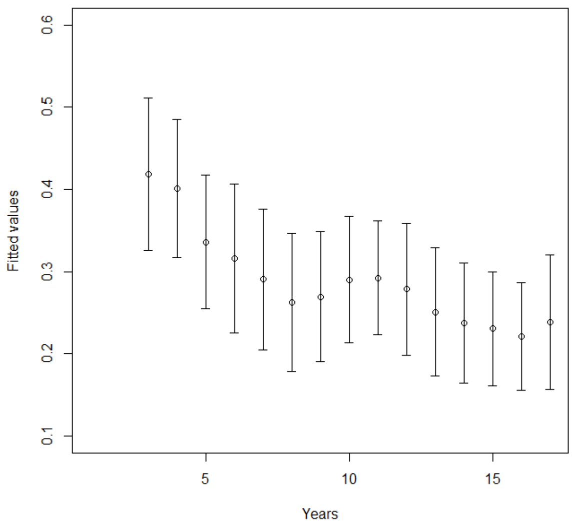

C. Three spells

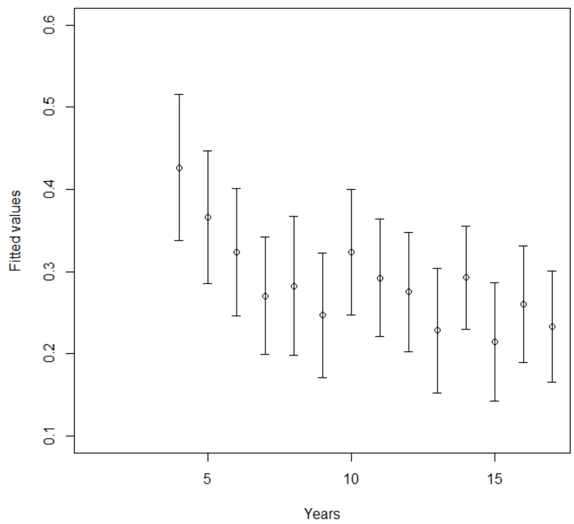

D. Four spells

Figure 12: Partially linear model intensity-duration surfaces, by spells

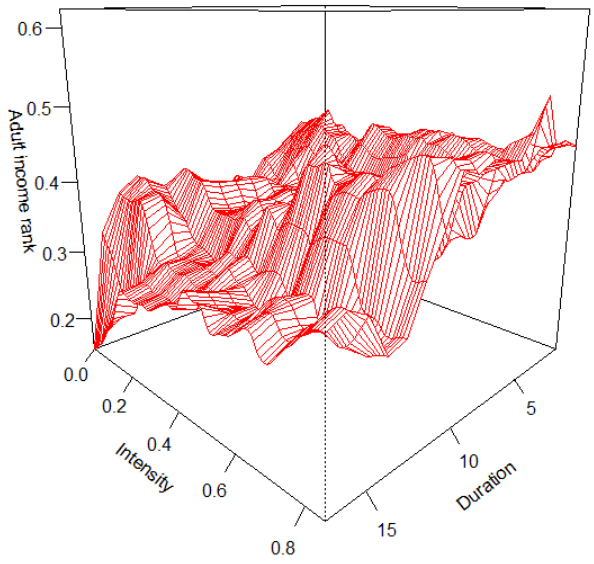

A. One spell

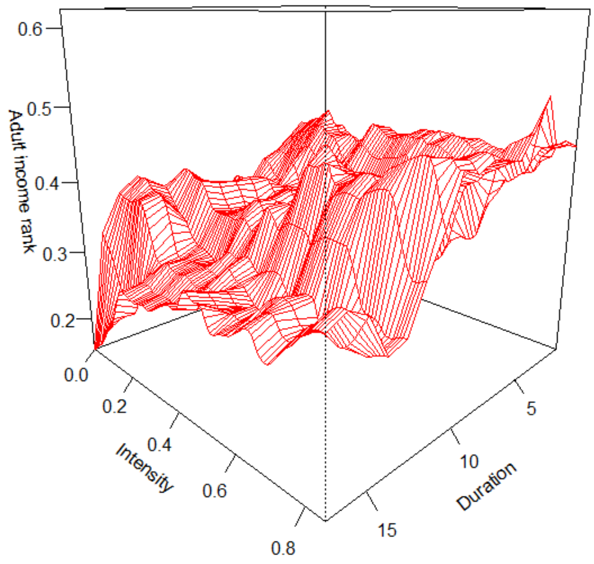

B. Two spells

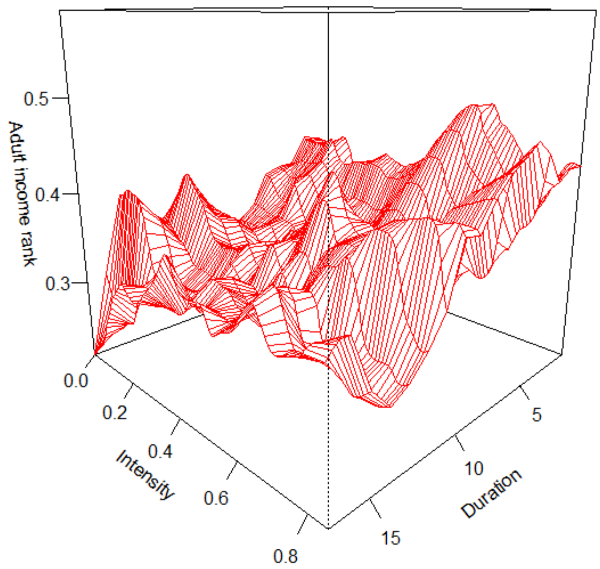

C. Three spells

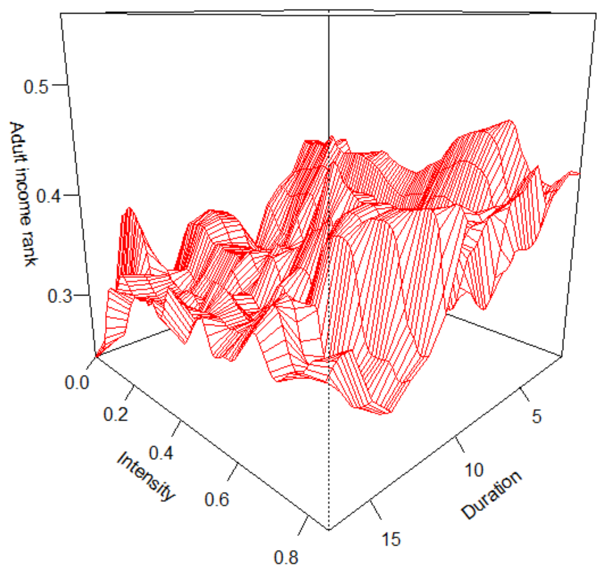

D. Four spells

Figure 13: Change before and after CBGPS balancing, in the correlation of rank and the control variables

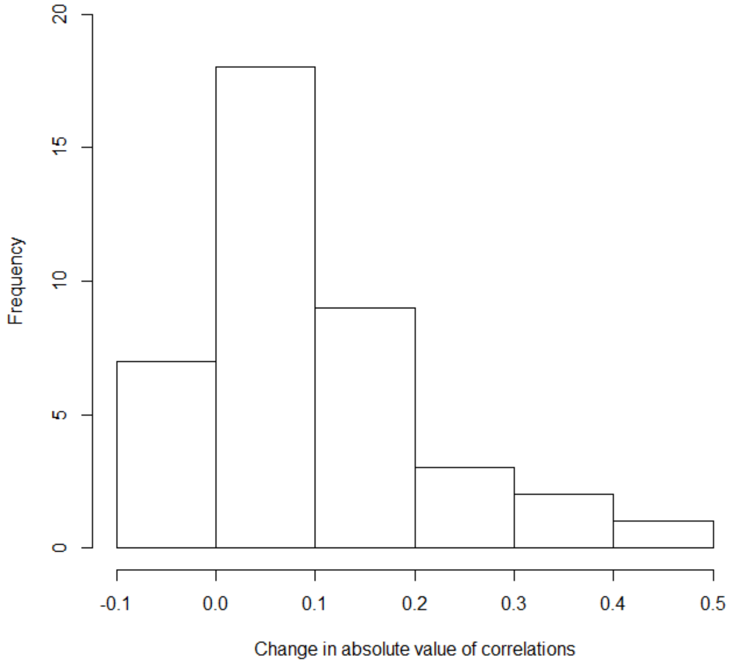

Figure 14: Regression of rank on parental rank with 95 percent CI and scatterplot

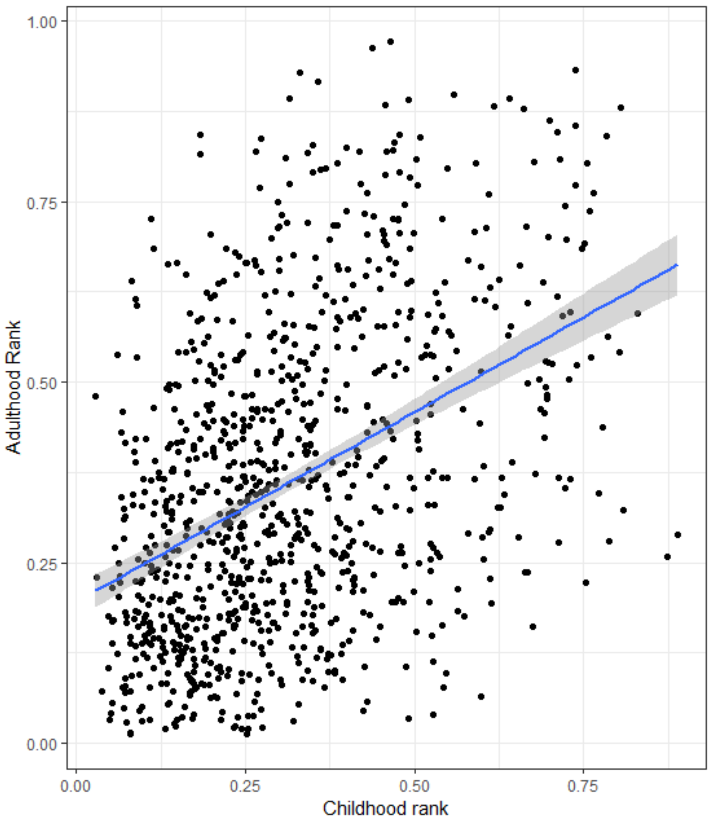

Figure 15: Fitted values for the censored sample (n=984, dashed) and full sample (n=2281, solid) with the Heckman correction

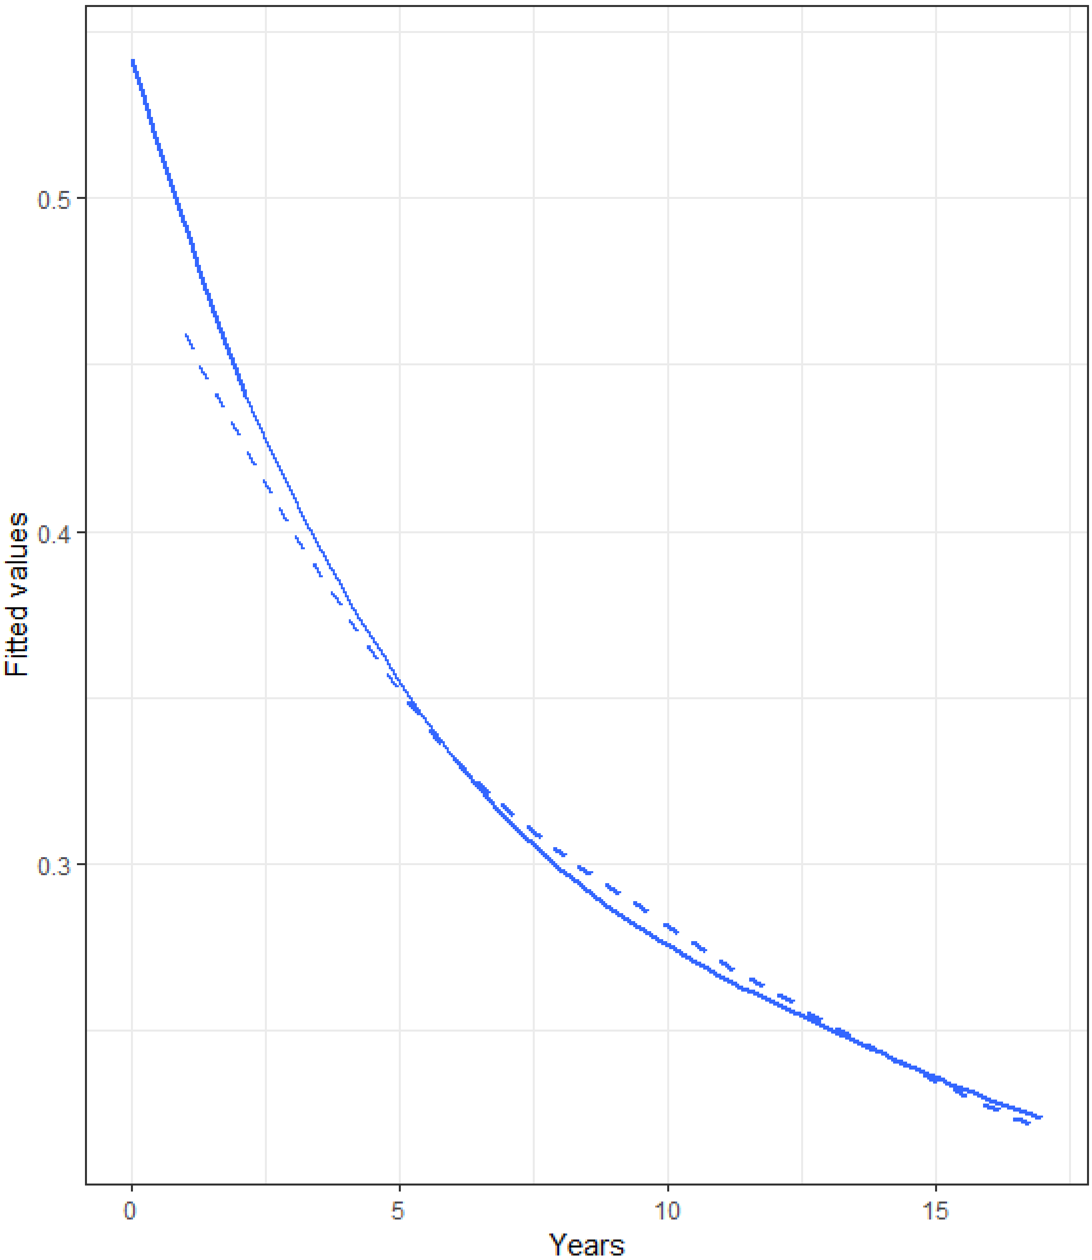

Figure 16: Inverse Mills Ratio vs. linear predictors ($x\hat{\beta}$) from first-stage probit model

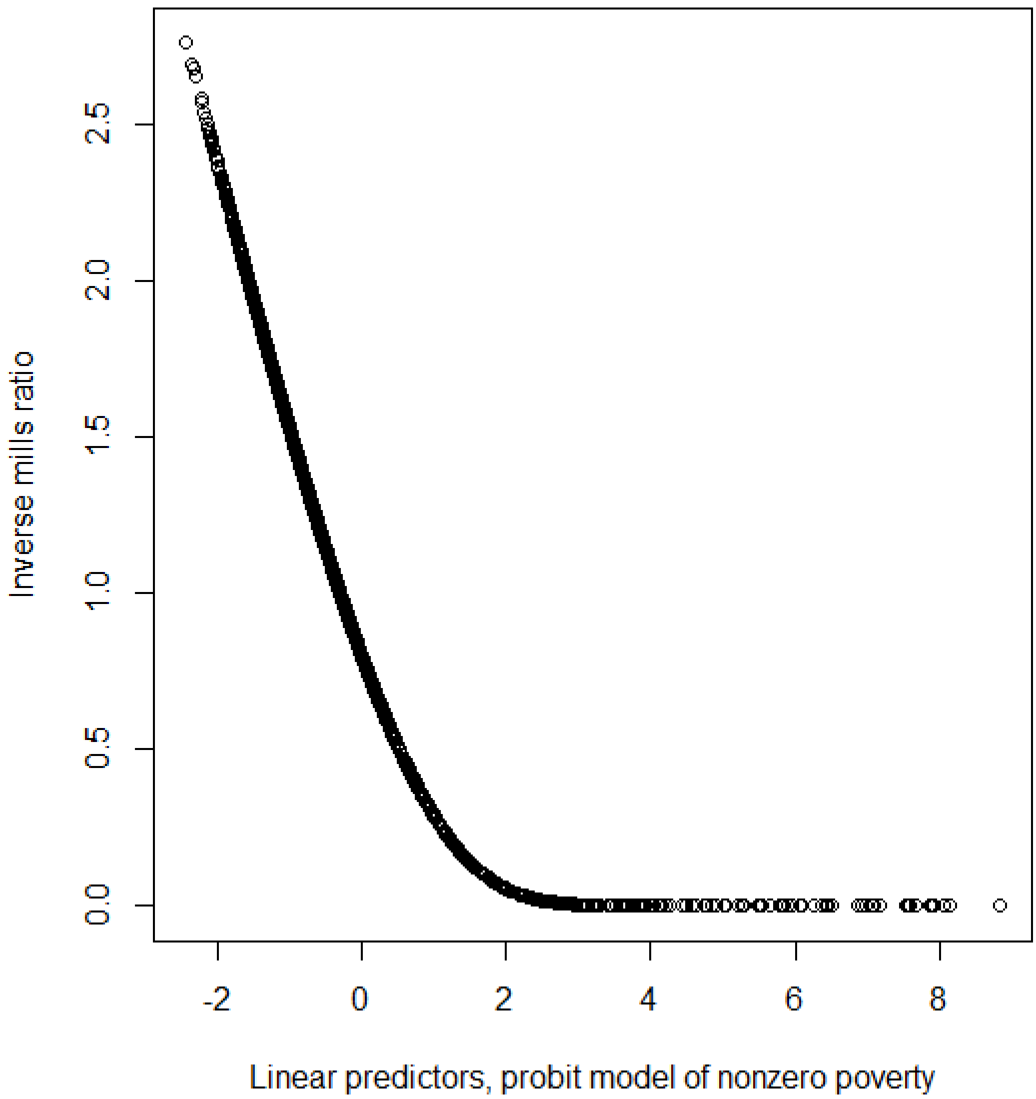

Table 1: Key Variable Descriptive Statistics

	Minimum	Mean	Maximum	St. Dev.
Age	1.00	4.27	17.00	4.30
Duration	1.00	6.07	17.00	5.12
Intensity	0.00	0.14	0.78	0.15
Family income (adulthood)	131.90	50120	568200	37845.18
Family income (childhoood)	4054	41300	235100	23528.14
Rank (adulthood)	0.01	0.22	0.99	0.16
Rank (childhood)	0.03	0.31	0.86	0.17

Table 2: Descriptive Statistics of the control variables

	Minimum	Mean	Maximum	SD
Demographic variables:				
Black Head of Household (HOH)	0.00	0.54	1.00	0.49
Other Race HOH	0.00	0.02	1.00	0.12
Hispanic HOH	0.00	0.03	1.00	0.14
Female	0.00	0.56	1.00	0.50
Married	0.00	0.63	1.00	0.38
Number of children in family	1.00	2.65	14.00	2.04
Individual's birth order	1.00	1.34	7.00	0.68
Education variables:				
High School Attainment	0.00	0.15	1.00	0.36
High School plus some additional Attainment	0.00	0.23	1.00	0.42
College Attainment	0.00	0.06	1.00	0.17
Occupation variables:				
Professional Occupation	0.00	0.08	1.00	0.18
Manager Occupation	0.00	0.03	1.00	0.10
Sales Occupation	0.00	0.09	1.00	0.19
Clerical Occupation	0.00	0.15	1.00	0.26
Craftsmen Occupation	0.00	0.16	1.00	0.26
Operatives Occupation	0.00	0.06	1.00	0.17
Transportation Occupation	0.00	0.08	1.00	0.20
Laborer Occupation	0.00	0.00	0.38	0.01
Farm Manager Occupation	0.00	0.02	1.00	0.12
Farm Labor Occupation	0.00	0.17	1.00	0.27
Service Occupation	0.00	0.03	1.00	0.13
Household Work Occupation	0.00	0.01	0.86	0.05
Industry Variables:				
Agricultural Industry	0.00	0.01	0.82	0.05
Mining Industry	0.00	0.06	1.00	0.17
Construction Industry	0.00	0.22	1.00	0.32
Manufacturing Industry	0.00	0.07	1.00	0.18
Transportation Industry	0.00	0.15	1.00	0.25
Finance Industry	0.00	0.02	1.00	0.09
Business Service Industry	0.00	0.04	1.00	0.14
Personal Service Industry	0.00	0.05	1.00	0.16
Entertainment Service Industry	0.00	0.01	0.73	0.05
Professional Service Industry	0.00	0.15	1.00	0.26
Employment and welfare variables:				
Employment	0.00	0.72	1.00	0.28
Sum of AFDC transfers during childhood ($)	0.00	309701	3316162	55879
Location Variables:				
Northeast Residence	0.00	0.11	1.00	0.29
North Central Residence	0.00	0.24	1.00	0.41
South Residence	0.00	0.53	1.00	0.47
West Residence	0.00	0.12	1.00	0.30

Table 3: Parametric linear in parameters income rank models

	Dependent variable:		
	Adulthood income rank		
	(1)	(2)	(3)
Age	0.0001	0.002	0.0009
	(0.002)	(0.003)	(0.005)
Duration	-0.021^{***}	-0.02^{***}	-0.035^{***}
	(0.003)	(0.004)	(0.007)
Intensity	0.204^{*}	0.151	0.321
	(0.092)	(0.282)	(0.223)
Age×Duration		-0.0005	
		(0.001)	
Age×Intensity		-0.03	
		(0.055)	
Duration×Intensity		0.002	
		(0.016)	
Age×Duration×Intensity		0.003	
		(0.005)	
Age^2			-0.00008
			(0.0003)
Duration^2			0.0009^{**}
			(0.0004)
Intensity^2			-0.31
			(0.3)
Constant	0.458^{***}	0.459^{***}	0.360^{***}
	(0.015)	(0.019)	(0.006)
Observations	984	984	984
R^2	0.149	0.151	0.154
Adjusted R^2	0.147	0.144	0.148
Note:	$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$		

Table 4: Parametric linear in parameters income rank models, with controls

	Dependent variable:		
	Adulthood income rank		
	(1)	(2)	(3)
Duration	-0.01^{***}	-0.0094^{*}	-0.015^{**}
	(0.003)	(0.0.004)	(0.0074)
Age	-0.0012	0.0004	-0.0041
	(0.0016)	(0.003)	(0.0052)
Intensity	0.082	-0.099	-0.0067
	(0.096)	(0.29)	(0.23)
Duration^2			0.0005
			(0.0004)
Age^2			0.00017
			(0.0003)
Intensity^2			0.117
			(0.305)
Duration$\times$Age		-0.0005	
		(0.001)	
Duration$\times$Intensity		0.009	
		(0.016)	
Age$\times$Intensity		-0.036	
		(0.055)	
Duration$\times$Age$\times$Intensity		0.0045	
		(0.0051)	
Constant	0.911^{***}	0.94^{***}	0.968^{***}
	(0.277)	(0.279)	(0.279)
Observations	984	984	984
Controls	YES	YES	YES
R^2	0.257	0.259	0.259
Adjusted R^2	0.223	0.222	0.223

Note: Controls include demographic, occupation, industry, employment, welfare, and locational characteristics, from Table 2.
$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table 5: Parametric models with concentration

	Dependent variable:		
	Adulthood Rank		
	(1)	(2)	(3)
Years	−0.021***	−0.019**	−0.010***
	(0.003)	(0.008)	(0.003)
Age	0.00002	0.002	−0.002
	(0.002)	(0.004)	(0.002)
Intensity	0.200**	0.353	0.084
	(0.092)	(0.439)	(0.095)
Concentration	−0.012	−0.027	−0.009
	(0.009)	(0.045)	(0.009)
Years×Age		0.001	
		(0.003)	
Years×Intensity		0.013	
		(0.046)	
Age×intensity.conditional		−0.146*	
		(0.079)	
Years×Concentration		0.003	
		(0.021)	
Age×Concentration		0.003	
		(0.016)	
Intensity×Concentration		0.353	
		(0.691)	
Years×Age×Intensity		0.019	
		(0.023)	
Years×Age×Concentration		−0.003	
		(0.008)	
Years×Intensity×Concentration		−0.110	
		(0.157)	
Age×Intensity×Concentration		0.269	
		(0.199)	
Years×Intensity×Age×Concentration		−0.046	
		(0.075)	
Controls	NO	NO	YES
F test p. val., vs. no concentration	-	.999	.999
Observations	984	984	984
R^2	0.151	0.164	0.256
Adjusted R^2	0.147	0.151	0.224
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 6: Parametric models with controls, weighted by CBGPS

	Dependent variable:		
	Adulthood income rank		
	(1)	(2)	(3)
Duration	-0.022^{***}	-0.014^{***}	-0.026^{***}
	(0.003)	(0.005)	(0.008)
Age	-0.004^{***}	0.001	-0.020^{***}
	(0.002)	(0.003)	(0.005)
Intensity	0.300^{**}	0.460	0.577^{**}
	(0.121)	(0.314)	(0.260)
Duration^2			0.0002
			(0.0005)
Age^2			0.001^{***}
			(0.0003)
Intensity^2			-0.548
			(0.433)
Duration$\times$Age		-0.003^{**}	
		(0.001)	
Duration$\times$Intensity		-0.025	
		(0.019)	
Age$\times$Intensity		0.019	
		(0.059)	
Duration$\times$Age$\times$Intensity		0.003	
		(0.006)	
Constant	1.006^{***}	1.023^{***}	1.095^{***}
	(0.311)	(0.311)	(0.311)
Observations	984	984	984
Controls	YES	YES	YES
R^2	0.318	0.324	0.328
Adjusted R^2	0.287	0.290	0.295
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table 7: Baseline IGE and Rank-Rank Regressions

	Dependent variable:	
	Log Family Income (adulthood)	Adulthood Rank
	(1)	(2)
Log Family Income (childhood)	0.553***	
	(0.040)	
Family Rank (childhood)		0.526***
		(0.036)
Constant	4.766***	0.196***
	(0.422)	(0.013)
Observations	984	984
R^2	0.162	0.182
Adjusted R^2	0.161	0.181
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 8: Duration and Age Baseline Regression

	Dependent variable:			
	Family Income ($)			
	(1)	(2)	(3)	(4)
Duration	$-2{,}351.426^{***}$	$-2{,}127.944^{***}$	$-1{,}099.100^{***}$	-902.837^{**}
	(229.586)	(288.553)	(328.760)	(374.375)
Age at first exposure	−132.492	130.431	−343.796	−118.859
	(258.287)	(330.159)	(256.885)	(328.841)
Duration×Age		−131.486		−111.802
		(102.897)		(102.056)
Constant	$64{,}168.240^{***}$	$63{,}673.270^{***}$	$239{,}288.900^{***}$	$243{,}587.900^{***}$
	(2,411.867)	(2,442.007)	(44,713.500)	(44,880.650)
Observations	984	984	984	984
Controls	NO	NO	YES	YES
R^2	0.118	0.120	0.243	0.244
Adjusted R^2	0.116	0.117	0.209	0.210
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 9: Coefficients of interest, adjusted by Heckman correction procedure

	Dependent variable:			
	Adulthood Rank			
	(1)	(2)	(3)	(4)
Years	-0.007^{***}	-0.010^{***}	-1.194^{***}	-2.158^{***}
	(0.002)	(0.002)	(0.307)	(0.352)
Years^2			0.328	0.423^{*}
			(0.200)	(0.225)
Inverse Mills Ratio		0.124^{***}		0.106^{***}
		(0.022)		(0.024)
Constant	0.900^{***}	1.067^{***}	0.868^{***}	1.020^{***}
	(0.276)	(0.166)	(0.275)	(0.166)
Observations	984	2,251	984	2,251
Controls	YES	YES	YES	YES
R^2	0.254	0.313	0.256	0.314
Adjusted R^2	0.224	0.301	0.226	0.302
Note:	*p<0.1; **p<0.05; ***p<0.01			