Is being an only child harmful to psychological health?: Evidence from an instrumental variable analysis of China's One-Child Policy

Shuxi Zeng Fan Li Peng Ding [1]

ABSTRACT

This paper evaluates the effects of being an only child in a family on psychological health, leveraging data on the One-Child Policy in China. We use an instrumental variable approach to address the potential unmeasured confounding between the fertility decision and psychological health, where the instrumental variable is an index on the intensity of the implementation of the One-Child Policy. We establish an analytical link between the local instrumental variable approach and principal stratification to accommodate the continuous instrumental variable. Within the principal stratification framework, we postulate a Bayesian hierarchical model to infer various causal estimands of policy interest while adjusting for the clustering data structure. We apply the method to the data from the China Family Panel Studies and find small but statistically significant negative effects of being an only child on self-reported psychological health for some subpopulations. Our analysis reveals treatment effect heterogeneity with respect to both observed and unobserved characteristics. In particular, urban males suffer the most from being only children, and the negative effect has larger magnitude if the families were more resistant to the One-Child Policy. We also conduct sensitivity analysis to assess the key instrumental variable assumption. KEY WORDS: Causal inference, marginal treatment effect, principal stratification, sensitivity analysis, heterogeneous treatment effect

[1]Shuxi Zeng is Ph.D. student, Department of Statistical Science, Duke University, Durham, NC, 27705 (email: shuxi.zeng@duke.edu); Fan Li is Associate Professor, Department of Statistical Science, Duke University, Durham, NC, 27705 (email: fl35@duke.edu); Peng Ding is Assistant Professor, Department of Statistics, University of California, Berkeley, CA 94720 (email: pengdingpku@berkeley.edu).

1 Introduction

The One-Child Policy (OCP) was a birth planning policy in China to control the rapid population growth during the mid 20th century. From late 1979 to 2015, the Chinese government enforced strict regulations to limit the number of children that each family could have. In most cases, each family was allowed to have only one child. A range of penalties were imposed on families who violated the OCP, including hefty financial fines, restriction to education, and demotion of parents working in the public sector. The OCP formally ended in 2015 when each family was allowed to have two children. As the world's most aggressive policy of population planning, the OCP had far-reaching and transformative influence on every facet of the Chinese society. Although the OCP was effective in controlling the population growth, its social, economic and cultural impact remains controversial. Various aspects of the OCP have been studied in the literature. For example, a stream of research focused on its demographic influence, such as the artificial selection of the gender of newborns or "man-made" twins (e.g. Ebenstein, 2010; Huang et al., 2016), marriage distortion, and welfare loss (Huang and Zhou, 2015). Another strand of research studied its impact on human capital accumulation by inducing more families to have a single child, e.g., whether the OCP leads to higher educational attainment or better physical conditions of the children (e.g. Rosenzweig and Zhang, 2009; Qian, 2009; Liu, 2014; Huang et al., 2020).

This paper investigates the causal effects of being an only child on subjective psychological well-being. This is an important social topic because a prevalent yet unsubstantiated perception in China is that the only children are more selfish, insecure and immature compared to the children with siblings. Indeed, the only children following the OCP in China are often called the generation of "little emperors" (Cameron et al., 2013). The effects of being an only child have been studied in different disciplines from complementary perspectives. Research in psychology has suggested that sibling companion contributes positively to the development of psychological health (Dunn, 1988; Brody, 2004; McHale et al., 2012). On the other hand, in labor economics, Becker and Lewis (1973) proposed the influential theoretical model of "quantity-quality interaction" for fertility choice, which explains the relationship between the size of a family (i.e. "quantity") and the well-being (i.e. "quality", such as physical health and education attainment) of each child. Some empirical studies found a negative association between the family size and school achievement (e.g., Leibowitz, 1974; Hanushek et al., 1992), but such an association was not observed in other quasi-experiments (e.g., Black et al., 2005; Angrist et al., 2010). Mogstad and Wiswall (2016) revealed a more complex and heterogeneous effect of the number of children on children's well-being. These studies found heterogeneous patterns across different populations. It is therefore of interest to extend the research to the Chinese population given its different cultural background.

Evaluating the effects of being an only child on psychological health is challenging because some unmeasured confounding factors may simultaneously affect the fertility decision and psychological health outcomes of the children within a family. The families with only one child may differ from the families with more than one child in systematic but unobserved ways (e.g. nurturing environment). A credible causal comparison requires proper adjustment of both measured and unmeasured confounding. A number of studies leveraged the OCP as an exogenous shock to address confounding. Cameron et al. (2013) adopted an instrumental variable (IV) approach with the birth year as an IV, leveraging a natural experiment on individuals who were born just before and just after the implementation of the OCP. Based on measures of subjective well-being, they concluded that only children are less trustworthy, more pessimistic and risk taking. Another approach is to use the implementation intensity of the OCP as an IV. More specifically, although the OCP was a national policy, it was implemented with different intensities across regions and time periods. Some researchers adopted the fine rate as a proxy for the intensity (Ebenstein, 2010; Huang et al., 2016). Attane (2002) developed the Indicator of Family Planning Policy Resistance (IFPPR)—a continuous positive index—to characterize the OCP implementation intensity. From a design perspective, the variation in the implementation intensity provides a natural experiment on fertility choice. Along this line of thoughts, Wu (2014) used the IFPPR as an IV to analyze a subsample of the China Family Panel Studies (CFPS) (Xie and Hu, 2014), and found that being an only child had a negative effect on self-reported psychological health. Wu (2014) employed a three-stage-least-squares method (Heckman, 1978), assuming a homogeneous treatment effect. Wu's analysis left three important statistical challenges open. First, the three-stage-least-squares model lacks a formal causal interpretation when the strong assumption of homogeneous treatment effect does not hold. Second, the sample units are clustered within provinces; importantly, the IV—the OCP implementation density—is measured at the cluster level defined by the province of birth. Third, the key IV assumptions such as exclusion restriction may be violated to a certain degree and its impact to the analysis remains unknown.

Motivated by these limitations in Wu (2014), we propose a set of new methods for continuous IV analysis. Specifically, under the potential outcomes framework to causal inference, we extend the local IV method (Heckman and Vytlacil, 1999, 2001), capitalizing on an intrinsic link between local IV and principal stratification (Frangakis and Rubin, 2002). This link was implicitly implied in Heckman and Pinto (2018) but has not been used in statistical analysis previously. Based on this link, we translate the causal estimands in the local IV method into the causal parameters under principal stratification. Within the principal stratification framework, we propose a flexible Bayesian model that allows for heterogeneous treatment effects and accommodates the clustering structure and ordinal outcomes, as well as provides straightforward posterior inference of the causal estimands. Our method extends the binary IV approach for randomized experiments with noncompliance (Angrist et al., 1996), and the principal stratification approach to clustered randomized trials (Frangakis et al., 2002; Jo et al., 2008; Forastiere et al., 2016). Moreover, we propose a sensitivity analysis method to assess the potential impact of the violation of the key exclusion restriction assumption, which supplements the existing literature on IV sensitivity analysis that has largely focused on linear models with constant effects (e.g. Small, 2007; Wang et al., 2018). Although our methods are originally motivated by the empirical application of OCP, they are applicable to a wide range of studies with continuous IV.

Our local IV analysis stratifies on two key background covariates, leading to four subpopulations defined by male or female and urban or rural areas. Across three self-reported measures of confidence, anxiety and desperation, being an only child does not affect the rural populations significantly but does exert small yet significant negative effects on the urban subpopulations. Thanks to the local IV method, we also detect that urban males suffered the most from being only children especially for those from families which were more resistant to the OCP. We offer possible explanations to the treatment effect heterogeneity and also provide several sensitivity checks and simulation evaluations. The data and programming code of this paper is available at:

https://github.com/zengshx777/OCP_LIV_JRSSA.

2 The Causal Inference Framework

2.1 Basic setup and assumptions

Consider a sample of N units from a population of interest, where the ith ($i = 1, 2, \cdots, N$) individual belongs to a cluster C_i ($C_i = 1, 2, \cdots, G$) defined by the province of birth. For individual i, we observe a set of pretreatment covariates X_i, for example, the demographic information and family background. Let T_i be the treatment indicator, with $T_i = 1$ if individual i is an only child and $T_i = 0$ otherwise. We observe a response variable Y_i and a continuous IV Z_i bounded between $z_{\min}$ and $z_{\max}$. In our application, Y_i is an ordered self-reported measurement of psychological health with a larger value representing a better condition, and Z_i is the IFPPR index, with a larger value indicating a higher policy implementation intensity (i.e., it is decreasing in the original index created by Attane (2002)).

We proceed under the potential outcomes framework. We invoke the Stable Unit Treatment Value Assumption (SUTVA) (Rubin, 1980). Specific to our application, SUTVA implies two components. First, whether individual i is an only child depends on the IFPPR in the corresponding province, but does not depend on the IFPPR in other provinces. Second, the IFPPR and the fertility decision for one family do not affect the psychological health of children from other families. Both assumptions are reasonable in our application. SUTVA allows us to write $T_i(z)$ as the potential treatment status had the IV of unit i taken the value z, and $Y_i(z, t)$ as the potential outcome had unit i been exposed to the IV value z and the treatment status t. For each individual i, we observe the treatment status $T_i = T_i(Z_i)$ and the outcome $Y_i = Y_i(Z_i, T_i(Z_i))$.

We now formally introduce the IV assumptions.

Assumption 1 $Z_i \perp\!\!\!\perp \{Y_i(z, t), T_i(z) : t = 0, 1, z \in [z_{\min}, z_{\max}]\} \mid X_i$.

Assumption 1 requires that the IV is randomly assigned with respect to the potential treatment status and outcomes, conditioning on the covariates X_i. This assumption is reasonable in our application because the intensity of implementing OCP largely depended on the specific province, and which province a family resided in is random conditional on the family background.

Assumption 2 *The probability* $\Pr(T_i = 1 \mid Z_i = z, X_i)$ *is a non-degenerate function of* z *for all* X_i.

Assumption 2 requires that the IV has non-zero effects on the treatment assignment, that is, the probability of being an only child varies with the intensity of the OCP conditioning on covariates. This assumption is testable by checking the empirical distribution of the treatment conditional on the covariates. We will present such evidence in Section 4.

Assumption 3 (Monotonicity) *For any* $z < z'$, $T_i(z) \leq T_i(z')$ *for all* $i = 1, ..., N$.

Assumption 3 requires that the IV monotonically affects the potential treatment status $T_i(z)$; it extends the monotonicity assumption on a binary IV in Angrist et al. (1996). In our application, monotonicity assumes that increasing the intensity of implementing OCP will not increase the number of children a family had, which is reasonable given the severe financial and social punishment incurred from violating the OCP.

Assumption 4 (Exclusion Restriction) *The potential outcomes is solely a function of* t*, that is, for any unit* i *and for any* z, z', $Y_i(z, t) = Y_i(z', t)$, *for* $t = 0$ *and* 1.

Assumption 4 requires that the IV affects the outcomes only through its effects on the treatment. In our application, this means that the intensity of the OCP only affects whether a child is an only child or not, but does not directly affect psychological health. Namely, given the same treatment status, different intensity of implementing OCP has no influence on the outcomes. Under exclusion restriction, we can use the single-index notation for the potential outcome, $Y_i(t)$, instead of the double-index $Y_i(z, t)$, which we adopt hereafter. Exclusion restriction is a crucial identification assumption in the IV analysis but might be questionable in practice. Intensity of the implementation of the OCP might have other channels to affect psychological health. For instance, it might change the divorce rate, which further affects psychological health of children. This could lead to violation of exclusion restriction, although we find the divorce rate was low throughout the study period. In Section 4.4, we will perform a sensitivity analysis to examine the consequences of potential violation to exclusion restriction in our application.

2.2 Causal estimands

We now introduce three causal estimands in the context of a continuous IV. The first estimand is the standard average treatment effect on the treated (ATT):

$$\tau^{\text{ATT}} = E\{Y_i(1) - Y_i(0) \mid T_i = 1\}. \tag{1}$$

The second estimand is the policy-relevant treatment effect (PRTE) (Heckman and Vytlacil, 2001):

$$\tau^{\text{PRTE}} = E\{Y_i(1) - Y_i(0) \mid T_i(z_{\min}) = 0, T_i(z_{\max}) = 1\}. \tag{2}$$

This estimand measures the causal effect for children from families who would give birth to more than one child at the lowest policy intensity but only one child at the highest policy intensity. These families changed their fertility decisions because of the OCP. The estimand τ^{PRTE} quantifies the causal effect for a subpopulation defined by the joint potential treatment status $T_i(z)$, where $T_i(z)$ measures the inclination to receive the treatment.

The third estimand is the marginal treatment effect (MTE):

$$\tau^{\text{MTE}}(z) = E\{Y_i(1) - Y_i(0) \mid T_i(z) = 1 \text{ and } T_i(z') = 0 \text{ for any } z' < z\}. \tag{3}$$

It measures the causal effect for the units at the margin of receiving the treatment at a given IV value z under the monotonicity assumption. Similar to τ^{PRTE}, it is also a "local" causal effect on the subpopulations partitioned by the joint potential treatment status under different IV values. In our application, $\tau^{\text{MTE}}(z)$ is the causal effect on the children from the families who would have only one child at a given policy intensity z but would have more children with less intensity. In other words, for each given intensity level z, $\tau^{\text{MTE}}(z)$ is the causal effect on the children from the families who would change their fertility decision to obey the OCP just at that level. The $\tau^{\text{MTE}}(z)$ over the range of the IV give a complete picture of the heterogeneous treatment effects in the study population.

2.3 Local instrumental variable and principal stratification

Heckman and Vytlacil (1999) derived nonparametric identification formulas for the three causal estimands, which involve estimating the partial derivative of the conditional expectation of the outcome given covariates and the estimated propensity scores. Carneiro et al. (2011); Carneiro et al. (2017) used the sample analogues of the partial derivative in the estimation. However, it is difficult to quantify the uncertainty of these estimators as well as to accommodate clustered data. Below we adopt a selection model representation of the problem, which naturally allows for flexible Bayesian inference based on hierarchical models. The key to this representation is an intrinsic connection between the local IV approach (Heckman and Vytlacil, 1999) and principal stratification (Frangakis and Rubin, 2002), as illustrated below.

Principal stratification is a general framework for adjusting post-treatment intermediate variables in causal inference, which extends the IV approach to noncompliance by Angrist et al. (1996). In the context of IV, the intermediate variable is the treatment status, and a principal stratum is defined as the joint potential values of the treatment under all possible values of the IV: $\mathcal{T}_i = \{T_i(z) : z \in [z_{\min}, z_{\max}]\}$. The key insight is that the principal stratum, by definition, is not affected by the observed value of the IV and thus can be viewed as a latent pretreatment variable (Frangakis and Rubin, 2002). Therefore, one can define the principal causal effects as comparisons of potential outcomes conditioning on one principal stratum or combination of several principal strata. Both PRTE and MTE are special cases of principal causal effects. Most applications of principal stratification focused on a binary IV with a few exceptions (e.g. Jin and Rubin, 2008; Bartolucci and Grilli, 2011; Schwartz et al., 2011). The main challenge to a continuous IV is that there can be infinitely many possible principal stratum $\mathcal{T}_i$, rendering modelling and estimation difficult.

Fortunately, we can reduce the principal stratum to a scalar. Under Assumptions 1–3, Vytlacil (2002) showed that the potential treatment status $T_i(z)$ follows a latent selection model. Namely, Assumptions 1–3 are equivalent to the following assumption.

Assumption 5 *There exists a random variable S_i and a monotone non-degenerate function $v(\cdot)$ such* *that the intermediate treatment status* $\{T_i(z) : z \in [z_{\min}, z_{\max}]\}$ *can be written as* $T_i(z) = \mathbf{1}_{v(z) \geq S_i}$*, with* $Z_i \perp\!\!\!\perp S_i | X_i$.

In Assumption 5, S_i is an unobserved threshold that determines the treatment status of individual i. In our application, it describes the latent "utility" of violating the OCP for individual i. A larger value of S_i means that individual i needs a larger encouragement $v(z)$ to obey the OCP. The selection model representation in Assumption 5 allows for flexible modeling and inference strategies. See Kline and Walters (2019) for more discussions on the numerical equivalence between Assumption 1–3 and Assumption 5.

In Assumption 5, $v(z)$ is monotone in z. Without loss of generality, we set $v(z) = z$ because we can apply a monotone transformation of $v(z)$ and S_i simultaneously. See Vytlacil (2002) for more technical discussions. Therefore, $T_i(z) = \mathbf{1}_{z \geq S_i}$ is a step function with respect of z and its shape is determined by the random variable S_i. We can then use S_i to characterize the entire vector of the potential values of the treatment $\mathcal{T}_i$, and thus will call S_i the principal stratum hereafter. Based on the selection model representation in Assumption 5, the MTE estimand in (3) is equivalent to the following principal causal effect:

$$\tau^{\text{MTE}}(s) \quad = \quad E\{Y_i(1) - Y_i(0) \mid S_i = s\}. \tag{4}$$

We can express the ATT and PRTE estimands as weighted averages of $\tau^{\text{MTE}}(s)$ over a range of principal strata. Specifically, averaging $\tau^{\text{MTE}}(s)$ over the distribution of S_i for treated units leads to τ^{ATT}:

$$\tau^{\text{ATT}} = \int_{-\infty}^{\infty} \tau^{\text{MTE}}(s) F_S(\mathrm{d}s \mid T = 1); \tag{5}$$

averaging $\tau^{\text{MTE}}(s)$ over the distribution of S_i between $[z_{\min}, z_{\max}]$ leads to τ^{PRTE}:

$$\tau^{\text{PRTE}} = \int_{z_{\min}}^{z_{\max}} \tau^{\text{MTE}}(s) F_S(\mathrm{d}s \mid z_{\min} \leq S \leq z_{\max}). \tag{6}$$

In our context, the treated units are those with principal stratum S_i smaller than the observed policy intensity. Importantly, (5) involves the $\tau^{\text{MTE}}(s)$ values for the principal strata below the minimal intensity $z_{\min}$. However, because families with $S_i < z_{\min}$ or $S_i > z_{\max}$ would not change their fertility decisions regardless of the policy intensity (these are called the always-takers and never-takers by Angrist et al. (1996)), the principal stratum S_i of these families is not unique even if we know the whole joint potential values $\mathcal{T}_i$. We mitigate this complication by imposing a model for principal strata S_i conditioning on covariates, based on which we can impute the individual principal stratum membership S_i for the always-takers or never-takers. A similar strategy was adopted in Glickman and Normand (2000).

3 Bayesian hierarchical selection and outcome models

3.1 General structure of Bayesian inference

The above connection between the local IV approach and principal stratification allows us to employ a flexible Bayesian modelling strategy to infer the causal estimands. Specifically, for each unit i, the complete data are $\{Y_i(1), Y_i(0), S_i, X_i, Z_i\}$, where S_i is equivalent to $\mathcal{T}_i$ under Assumptions 1–4. The causal estimands are functions of the complete data, and thus inferring the causal effects depends on the complete data likelihood. Let $\Pr\{Y_i(1), Y_i(0), S_i, X_i, Z_i \mid \theta\}$ denote the joint probability density function of the random variables for unit i governed by parameters $\theta = (\zeta, \phi, \psi)$. We factorize the complete-data likelihood into three parts:

$$\begin{aligned}&\Pr\{S_i, Y_i(1), Y_i(0), X_i, Z_i \mid \theta\}\\ &\qquad = \Pr\{Y_i(1), Y_i(0) \mid S_i, X_i, \zeta\} \times \Pr(S_i \mid X_i, \phi) \times \Pr(X_i, Z_i \mid \psi). \quad (7)\end{aligned}$$

In (7), ζ denotes the parameters for the model of the potential outcomes, ϕ denotes the parameters for the model of the principal strata, and ψ denotes the parameters for the distribution of the covariates and the IV. Following the common practice in the literature, we assume the three sets of parameters (ζ, ϕ, ψ) are distinct and *a priori* independent.

Based on the factorization in (7), we can further refine the definition of MTE, allowing it to be conditional on X_i:

$$\tau^{\text{MTE}}(s, x) \quad = \quad E\{Y_i(1) - Y_i(0) \mid S_i = s, X_i = x\}. \quad (8)$$

Averaging $\tau^{\text{MTE}}(s, x)$ over the distribution of X_i conditional on $S_i = s$ gives the $\tau^{\text{MTE}}(s)$ in (4).

The potential outcome $Y_i^{\text{mis}} = Y_i(1 - T_i)$ and the principal stratum S_i are not observed for any unit. From a Bayesian perspective, these unobserved values are no different from unknown model parameters, both of which are unobserved random variables that we need to draw posterior inference on (Rubin, 1978). Specifically, we will simulate the posterior distribution of θ, and impute the missing values $(Y_i^{\text{mis}}, S_i)_{i=1}^N$ from their posterior predictive distributions conditional on the observed data and θ. We can then perform posterior inference of the causal effects based on the posterior samples of θ and $(Y_i^{\text{mis}}, S_i)_{i=1}^N$.

3.2 Models and posterior inference

The factorization in (7) suggests that to infer causal effects we need to specify three models: (a) a model for principal strata conditional on the covariates $\Pr(S_i \mid X_i, \phi)$, (b) a model for the potential outcomes conditional on the covariates and principal stratum $\Pr(Y_i \mid S_i, X_i, \phi)$, and (c) the joint distribution of the observed covariates and IV $\Pr(X_i, Z_i \mid \psi)$. In Bayesian inference, we usually condition on the empirical distribution of the covariates instead of modelling the joint distribution (Ding and Li, 2018), and thus below we will focus on the first two models.

First, for the continuous principal stratum S_i, we postulate a hierarchical model that accounts for clustering:

$$\begin{aligned} S_i &\sim \mathcal{N}(\beta_S' X_i + r_{C_i}, \sigma^2), \\ T_i(z) &= \mathbf{1}_{z \geq S_i}, \end{aligned} \tag{9}$$

where $r_{C_i} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \tau_S^2)$ for $C_i = 1, 2, \cdots, G$, are the random effects capturing the correlation structure of the principal stratum membership S_i of the units within the same province C_i. In our application, r_{C_i} can be interpreted as the latent resistance to the OCP in province C_i. This model implies the correlation structure for the treatment assignment mechanism $\Pr(T_i = 1 \mid Z_i, X_i)$, which is equivalent to a Probit model with Normally distributed random effects:

$$\Pr(T_i = 1 \mid Z_i, X_i) = \Phi\{(Z_i - \beta_S' X_i - r_{C_i})/\sigma\}, \tag{10}$$

where $\Phi(\cdot)$ is the cumulative distribution function for the standard Normal distribution. Here we restrict the coefficient of Z_i to be 1 but allow for an unknown variance of ε_i. This parametrization differs from the standard Probit model but follows more closely the latent index representation in Assumption 5.

Second, for the ordinal potential outcomes, we postulate a proportional odds model with a cumulative logit link (Agresti, 2003): for $t = 0, 1$ and $k = 1, 2, \cdots K - 1$,

$$\text{logit}\{\Pr(Y_i(t) \leq k \mid X_i, S_i)\} = \alpha_{t,k} + \beta_t' X_i + \gamma_t S_i + \nu_{t,C_i}, \quad \text{with } \alpha_{t,k} < \alpha_{t,k+1}, \tag{11}$$

where $\nu_{t,C_i} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \tau_t^2)$ for $C_i = 1, 2, \cdots, G$, capture the correlation structure of the potential outcomes within a province. In our application, ν_{t,C_i} can be interpreted as the latent psychological characteristics in province C_i. Model (11) assumes that each potential outcome has its own increasing intercepts $\alpha_{t,k}$'s for $t = 0, 1$. The outcome model (11) differs from the classical selection model (Heckman, 1979): The former models the conditional distribution of $Y_i(t)$ given S_i, whereas the latter models the joint distribution of the treatment assignment and the outcome model. The parametrization in (11) offers more convenient inference for the MTE. We do not impose a joint model for $Y_i(1)$ and $Y_i(0)$ because the data contain no information about their association. This does not pose a problem for inferring the τ estimands because they depend only on the marginal distributions of the potential outcomes.

We impose standard weakly-informative priors for the parameters. For the regression coefficients β, α and γ, we impose the diffuse Normal priors $\beta_d \sim \mathcal{N}(0, 100 \times I_p)$, $\beta_t \sim \mathcal{N}(0, 100 \times I_p)$ and $\alpha_{tk} \sim \mathcal{N}(0, 100 \times I_{K-1})$, and $\gamma_t \sim \mathcal{N}(0, 100)$; for the variance of the Probit model, we impose a flat prior $\pi(\sigma^2) \propto 1/\sigma^2$; for the standard deviations of the random effects τ_S and τ_t, we impose the half-Cauchy priors $\pi(\tau_S) \propto \{1 + (\tau_S/A)^2\}^{-1}\mathbf{1}_{\tau_S \geq 0}$ and $\pi(\tau_t) \propto \{1 + (\tau_t/A)^2\}^{-1}\mathbf{1}_{\tau_t \geq 0}$ with the scale parameter $A = 25$ (Gelman et al., 2006). Given the models and the prior distributions of the model parameters, we can obtain the posterior distribution of $\zeta = \{\alpha_{t,k}, \beta_t, \gamma_t, \tau_t^2 : t = 0, 1, k = 1, \ldots K-1\}$. We then derive the posterior distributions of $\tau^{\text{MTE}}(s, x)$. These are all functions of ζ based on the hierarchical outcome model. Finally, we average these conditional effects over the empirical distribution of $\{S_i, X_i, C_i\}$ to obtain the unconditional effects. For example, given a posterior draw of the parameters ζ, the province random effect v_{t,C_i}, and principal stratum S_i, we can obtain a posterior draw of τ^{ATT} and τ^{PRTE} from

$$\tau^{\text{ATT}} = \sum_{i=1}^{N} T_i \tau^{\text{MTE}}(S_i, X_i; \zeta, v_{t,C_i}) \Big/ \sum_{i=1}^{N} T_i. \tag{12}$$

$$\tau^{\text{PRTE}} = \sum_{i=1}^{N} \delta_i \tau^{\text{MTE}}(S_i, X_i; \zeta, v_{t,C_i}) \Big/ \sum_{i=1}^{N} \delta_i, \text{ with } \delta_i = \mathbf{1}_{z_{\min} \le S_i \le z_{\max}}. \tag{13}$$

This yields the posterior distributions of τ^{ATT} and τ^{PRTE}. In (12) and (13), we emphasize that the conditional effects depend on the parameter ζ and province random effects v_{t,C_i}.

4 Empirical application

4.1 The data and preliminary analysis

We now provide more information about the data. The CFPS is a comprensive longitudinal household survey representing 95% population in China. We only analyze the data from the first national wave in 2010, so that each household only appears once in the questionaire. The survey sample covers approximately 14,000 households from 25 provinces. Following Wu (2014), we exclude the samples from provinces such as Ninxia, Gansu, Xinjiang, and Yunnan, which include a large proportion of minority ethnicity and hence avoid the OCP restriction.

From the 2010 wave of CFPS we obtain the subsample of children born after 1979, the year when the OCP was first imposed, with age ranging from 16 to 31. Pretreatment covariates include age, ethnicity (1 for the Han ethnicity, 0 otherwise), paternal and maternal education attainment measured in years, family income, parents marriage status (whether divorced or not), and parental age at the child birth. We stratify on the area (urban or rural) and gender (female or male) and split the sample into four subgroups: rural females, rural males, urban females, and urban males, with a sample size of 1747, 1708, 407, and 416, respectively. Within these subgroups, the numbers of only children are 147, 283, 254, and 296, respectively. For families violating the OCP, we only keep the oldest child within a household. Clearly, the urban area has a much higher proportion of only children than the rural areas: the proportions of the households with only children are 12.46% and 66.89% in rural and urban areas, respectively. We use three psychological measures as outcome variables: self-confidence, degree of anxiety and desperation. During the survey, investigators asked the interviewees for the frequency of experiencing these emotions, and then transformed the frequency into the Likert scale (Likert, 1932). The outcomes take discrete values from 1 to 5, with a larger value indicating a better psychological condition.

In Table 1, the summary statistics of the observed covariates between the families with one child (i.e. treated) and the families with more than one child (i.e. control) reveal notable differences in the pretreatment variables between the treatment and control groups. In general, parents of only children had more educational attainment, but had a higher divorce rate. Also, only children were younger on average, which is partly due to our choice of keeping only the oldest child from a family with multiple children. This choice is to eliminate the difference in the birth order in the study sample, which may have a lasting effect on personality (Rohrer et al., 2015). Moreover, families with only children on average had higher household income. The Han ethnicity had a higher proportion of only children than minor ethnicities; this is expected because the OCP was not enforced among the ethnical minorities. In addition, we control for the number of siblings in the households with more than one child. For females, we also check whether she has a younger brother, which could potentially influence her psychological health (Chu et al., 2007). Specifically, we create a dummy variable that equals zero for the girls who were only children or only had female siblings and equals one for the girls who had a younger brother. However, simple t-tests do not reveal any significant differences in the outcomes between the treatment and control groups.

Our IV is the implementation intensity of the OCP, measured by the Indicator of Family Planning Policy Resistance (IFPPR) (Attane, 2002), ranging from 0 to 140. This index is an aggregated measure for the implementation intensity of OCP in the period of 1980s and hence does not vary with time. As discussed in Section 2, Assumption 1 (randomization of the IV) and Assumption 3 (monotonicity) is reasonable in our application. To examine Assumption 2, Figure 1 shows the fitted probabilities of being an only child in a family as a function of the quantile of the IV, which are predicted from the hierarchical model (9) with covariates fixed at their means and the random effects fixed at zero.

Table 1: Baseline characteristics of the whole sample, households with only children, and households with multiple children. Standard deviations are presented in the parentheses.

	All samples	Only child	With siblings
	$N = 4278$	$N_1 = 980$	$N_0 = 3298$
Family background			
Maternal education (years)	5.05 (4.57)	7.93 (4.28)	4.19 (4.29)
Paternal education (years)	6.92 (4.39)	8.71 (3.99)	6.39 (4.37)
Maternal age at birth (years)	25.5 (4.34)	25.1 (3.44)	25.6 (4.57)
Paternal age at birth (years)	27.4 (4.92)	26.8 (3.82)	27.6 (5.20)
Family annual income (CNY[2])	45100 (165790)	59025 (87050)	40965 (182570)
Divorce	1.49% (12.1%)	2.95% (16.9%)	1.06% (10.2%)
Individual information			
Age (years)	26.5 (3.49)	25.0 (3.38)	27.2 (3.52)
Number of siblings	1.35 (1.17)	0 (0)	1.74 (1.04)
Proportion of majority ethnicity	90.6% (29.1%)	95.6% (20.4%)	89.1% (31.1%)
Outcome information			
Confidence measure	4.00 (0.95)	3.97 (0.91)	4.01 (0.97)
Anxiety measure	4.16 (1.06)	4.14 (1.06)	4.17 (1.07)
Desperate measure	4.26 (1.03)	4.22 (1.03)	4.27 (1.03)

[†] Annual income in the Chinese currency (CNY) is for the year when the household was surveyed.

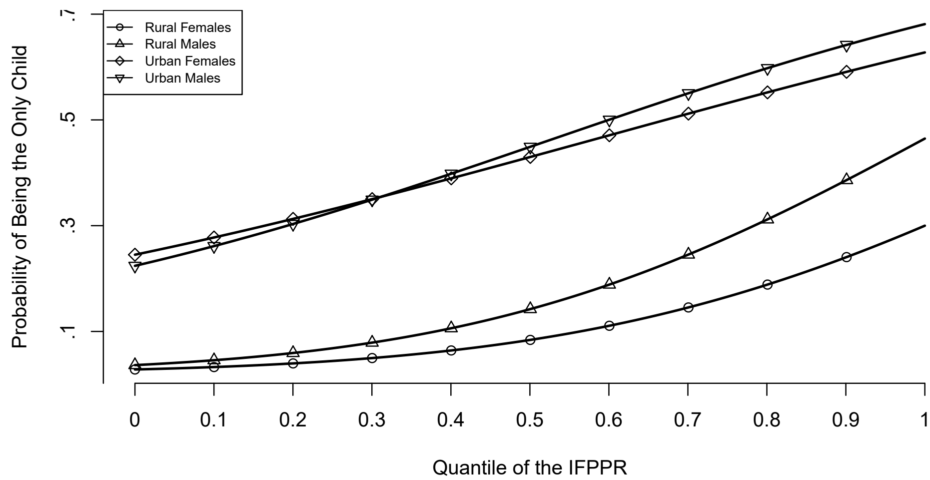

Figure 1: Probability of being treated (i.e. being an only child) as a function of the IFPPR quantiles

Figure 1 shows a clear increasing trend of being an only child as the IFPPR goes up in all four subpopulations, bolstering the plausibility of Assumption 2. We also see the heterogeneity in the treatment assignment mechanism across the four subpopulations. For both females and males, the individuals from the urban areas have a larger probability of being only children compared to those from the rural areas. This pattern can be attributed to the fact that urban residents were more inclined to obey government regulations because (a) the consequences of violating the OCP were usually more severe in the urban areas, and (b) urban residents were less likely to hold the traditional thinking of "more children are better" that was quite common in the rural areas. There is also a marked difference between females and males, especially in the rural areas. Girls had a lower probability of being only children. This pattern could be attributed to the traditional perception in China that "boys are superior to girls", namely, the families who had a girl as their first child would attempt to have a second child in hope of having a boy, even at the cost of violating the OCP. We observe a larger gender gap in the rural areas where such a perception was particularly prevalent. Another explanation for the discrepancy between urban and rural areas is the policy change. In mid 1980s, the Chinese government gradually relaxed the restriction in rural ares, allowing families to have a second child if the first one was a daughter (Banister, 1991; Huang et al., 2020). As a result, the proportion of female only children was lower in rural areas.

4.2 Marginal and average treatment effects

We applied our local IV approach with the Bayesian models (9) and (11) and priors in Section 3 to the CFPS data. We conducted separate analyses within each of the four subpopulations. We used `JAGS` (Plummer et al., 2003) to simulate the posterior distributions. For each analysis, we simulated three Markov chains with different starting values and 50,000 iterations for each chain, discarding the first half as burn-in and thinning the chains for every 25 iterations. The remaining 3,000 draws were used to approximate the posterior distributions. The Gelman–Rubin statistic (Gelman and Rubin, 1992) for all parameters are below 1.1 indicating the good mixing of the Markov chains.

Figure 2 shows the posterior means and credible intervals of the marginal treatment effects $\tau^{\text{MTE}}(s, x)$ against the principal stratum S_i, for the three outcomes. To compare across different subgroups and outcomes, we transform the principal stratum value into the quantile of its posterior distribution based on the imputed S_i in the sample. We use a formula from Ju and Geng (2010, page 133) to obtain the analytic expression of $\tau^{\text{MTE}}(s, x)$ from the hierarchical outcome model (11) as follows:

$$\begin{aligned}\tau_i(s, x) &= E\{Y_i(1) - Y_i(0) \mid X_i = x, S_i = s\} \\ &= \sum_{k=1}^{K-1} \{\text{sig}(\alpha_{0,k} + \beta_0' x + \gamma_0 s + v_{0,C_i}) - \text{sig}(\alpha_{1,k} + \beta_1' x + \gamma_1 s + v_{1,C_i})\},\end{aligned}$$

where $\text{sig}(x) = 1/\{1 + \exp(-x)\}$ is the sigmoid function. We then evaluate the posterior means and credible intervals at a grid of values of the quantiles of the principal stratum S_i, again setting the covariates at their means and the random effects at zero.

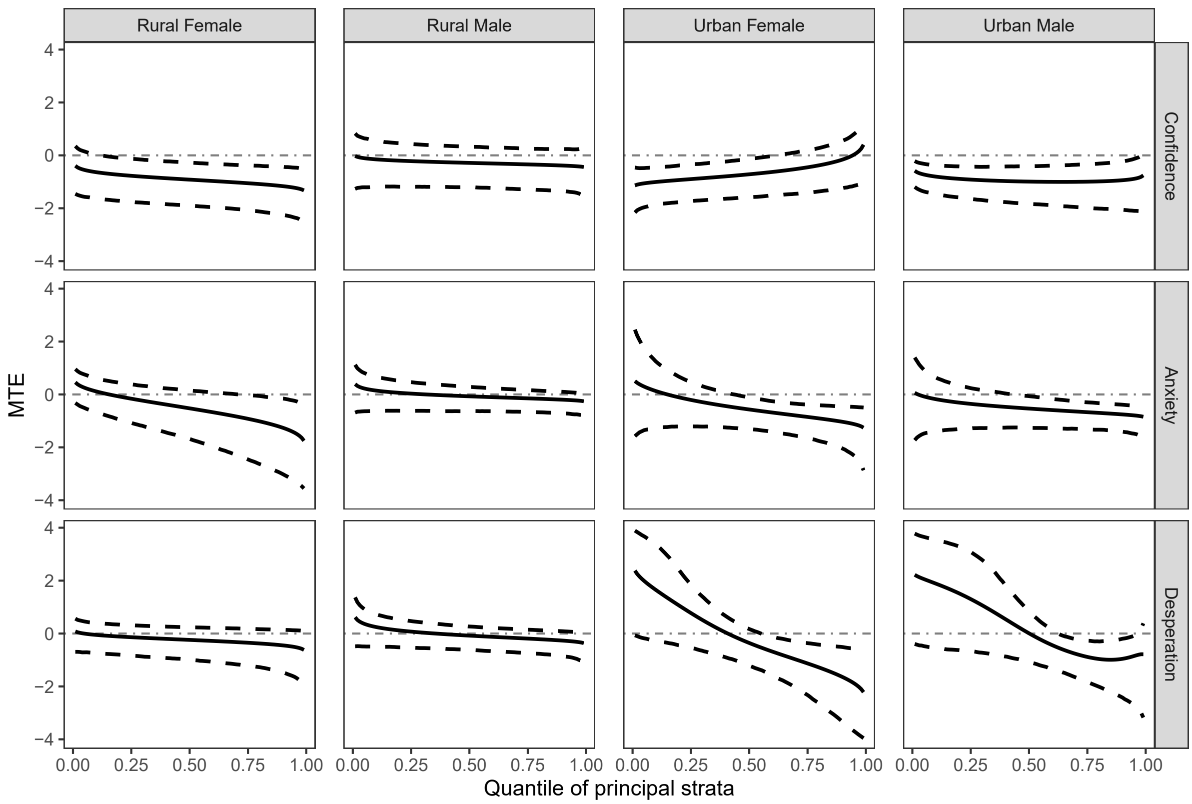

Figure 2: Posterior means and 95% credible intervals of $\tau^{\text{MTE}}(s, x)$ against principal strata s. X-axis represents the quantile of the principal strata S_i after normalization. A larger value in the x-axis means that the individual is from a family that is more resistant to the OCP.

For the confidence measure, a range of $\tau^{\text{MTE}}(s, x)$ lie below zero for the subgroups of rural females, urban females, and urban males, indicating negative treatment effects across principal strata. The males from urban areas are more likely to have a worse confidence measure, compared with males from rural areas. The patterns of $\tau^{\text{MTE}}(x, s)$ are similar for the anxiety and desperation measures. We observe a decreasing trend in the treatment effect as the principal stratification S_i increases. This trend is more apparent for the urban subgroups. Recall that $\tau^{\text{MTE}}(x, s)$ represents how the treatment effect of being an only child varies across principal strata, which quantify families's reluctancy of bearing only one child.

That is, for those from families that were more resistant to the OCP, the effect of being an only child is more negative. For instance, in the bottom right corner of Figure 2, the effect on desperation measure for urban males flips its sign from approximately 2 to -1 as the reluctance to obey the policy increases, taking up about 40% and -20% compared to the range of the outcome. A possible explanation is that the families with stronger preference for more children did not prepare well to bear only one child and thus provided a less nurturing environment for their children. In contrast, the decreasing trend is less noticeable for the rural subgroups, such as the desperation measure for rural males in the bottom left of Figure 2. Indeed, in the rural subgroups the effects of being an only child vary little across the families' degree of resistance to the OCP.

Based on formulas (12) and (13), we can infer τ^{ATT} and τ^{PRTE} based on their posterior means and 95% credible intervals, which are displayed in Figure 3. Overall, τ^{ATT} and τ^{PRTE} are very similar for each subgroup and outcome, with τ^{ATT} being slightly smaller. For all three outcome measures, in urban areas only children have on average $0.3 - 0.5$ (under the outcome scale of 1–5) smaller values than the children with siblings, translating into a statistically significant $7.5\% - 12.5\%$ decrease in self-reported psychological health. The effects on the rural individuals are much smaller and inconclusive. This pattern is consistent with the results from the marginal effects τ^{MTE}. One explanation is that the children grew up in rural areas had less restriction to communicate with their peers in the same community, such as a village. The companion of other children in the same cohort might substitute the effect of siblings, while the children in urban area might not have such opportunities.

It is worth noting that τ^{ATT} averages over all families with only children regardless of the OCP, whereas τ^{PRTE} averages over families who had only children as the consequence of the OCP. The similarity between τ^{ATT} and τ^{PRTE} is due to the large overlap of the target populations of τ^{ATT} and τ^{PRTE}. Namely, most families having only one child, corresponding to the target population of τ^{ATT}, would not have only one child at the lowest policy intensity, which belong to the target population of τ^{PRTE}. Indeed we found that the imputed principal strata of the majority (between 55% to 75%) of the families lie in the range of the observed values of the IV in all four subgroups. Therefore, the policy affected most families, in the sense that most families would have more than one child at the lowest policy intensity but have only

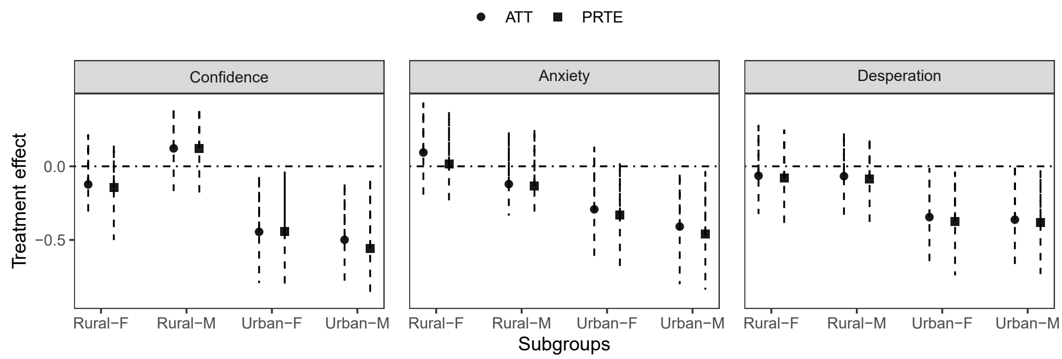

Figure 3: Posterior means and 95% credible Intervals of τ^{PRTE} and τ^{ATT} . The $\circ$'s are the posterior means of the treated effects on the treated and the $\square$'s are the posterior means of the policy related treatment effects. The vertical dashed lines are the credible intervals.

one child at the highest policy intensity.

In summary, our analysis suggests that being an only child has a negative impact on psychological health, especially for those from the urban areas. Specifically, being an only child decreases approximately 12%, 7.5% and 10% in the confidence, anxiety and desperation measures, respectively. One possible explanation of the negative effect is that only children might be subject to more pressure from parents, which influences psychological health, giving rise to an individual being less confident and more likely to feel anxious and desperate. The individuals from urban areas, especially urban males, might receive higher expectations from their families compared with those from rural areas. This leads to a worse impact of being an only child on the psychological health of the urban subgroup.

4.3 Comparison with alternative methods

We now compare the proposed local IV method with two alternative standard methods. The first method is direct regression adjustment based on ordinary least squares (OLS), where we fit a regression model of the outcome on the treatment and centered pretreatment variables with cluster-robust standard errors, $E(Y_i \mid T_i, X_i) = \beta_0 + \beta_t T_i + \beta_x X_i + \beta_{tx} T_i X_i$, and take the estimated coefficient β_t to estimate the ATT.

However, OLS cannot estimate the MTE or PRTE which are local effects defined by the IV. Moreover, OLS relies on the unconfoundedness assumption, that is, there is no unmeasured confounding beyond the observed pretreatment variables, which is unlikely to hold given the limited covariate information in the CPFS data.

The second method is the standard two-stage least squares (2SLS) for IV analysis. In the first stage, we use a regression model of T_i on X_i and Z_i, and in the second stage, we regress the outcome Y_i on X_i and the fitted values of T_i from the first stage to obtain the coefficient of the fitted T as an estimate of the ATT. We use the standard errors robust to the cluster structure in the two stages. The 2SLS method is more comparable to the local IV method. However, while the 2SLS estimate has a clear causal interpretation similar to the ATT when the treatment effects are homogeneous, its interpretation becomes ambiguous when treatment effects are heterogeneous (Heckman and Vytlacil, 2005). In contrast, the local IV method can provide a comprehensive picture of the potentially heterogeneous treatment effects via the marginal treatment effects curves shown in Figure 2.

Figure 4 shows the point estimates and 95% confidence intervals of the ATT estimated by OLS, 2SLS and local IV. The OLS estimates have much smaller standard errors than both IV methods, but the estimates are concentrated around zero and fail to detect significant effect in any outcomes and subgroups. Local IV and 2SLS generally agree on the signs of the effects, but local IV detects more significant effects with smaller standard errors than 2SLS. More importantly, as shown in the MTE curves from local IV, there is significant heterogeneity in treatment effects among urban children, particularly in anxiety and depression; reporting only the average effects as 2SLS and OLS would not capture the full picture of the heterogeneous effects.

4.4 Sensitivity analysis

Among the causal assumptions we made, exclusion restriction in Assumption 4 is the least justifiable by subjective knowledge. Therefore, in this subsection we conduct sensitivity analysis to examine the impact of the potential violation to exclusion restriction. To define the sensitivity parameter, we decompose the effect of the IV on the outcome into the component acting through the treatment and the

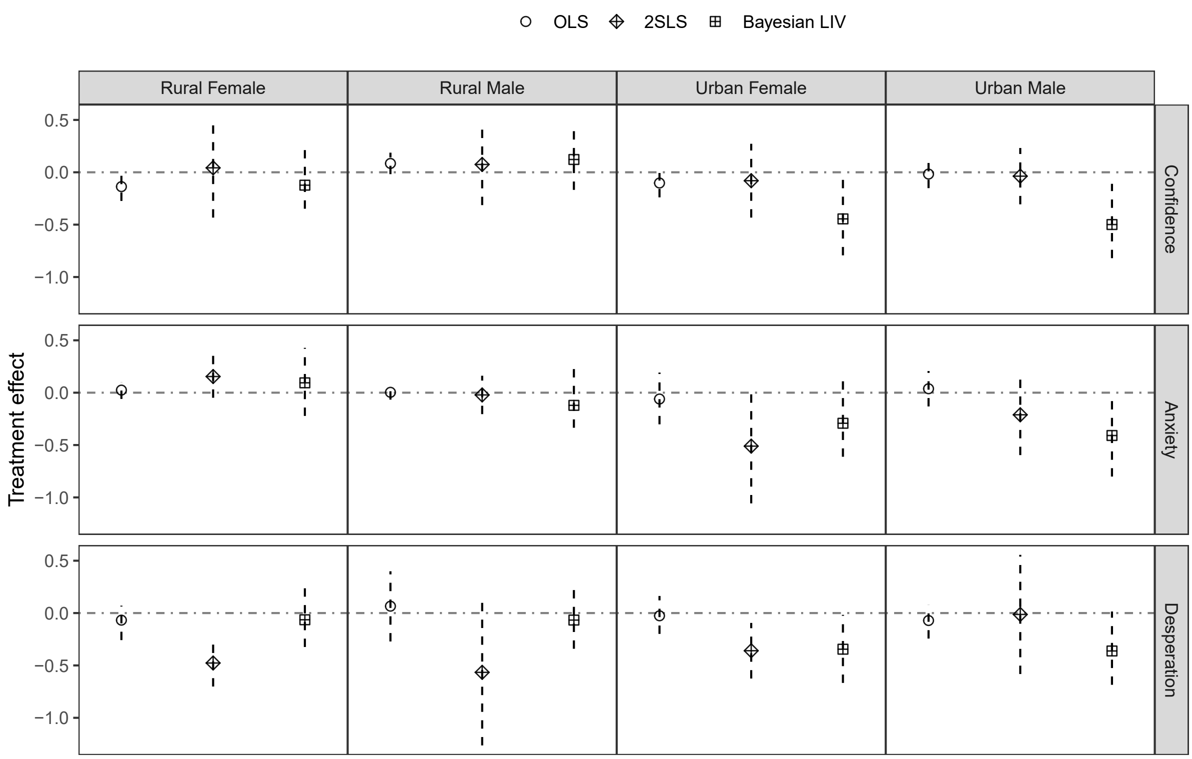

Figure 4: Comparison with other methods: dashed lines represent confidence or credible intervals.

remaining component:

$$\tau^{\text{direct}} = E\{Y_i(z_{\text{max}}, 0) - Y_i(z_{\text{min}}, 0)\},$$

$$\tau^{\text{indirect}} = E\{Y_i(z_{\text{max}}, 1) - Y_i(z_{\text{max}}, 0)\},$$

and define the sum of the two as:

$$\tau^{\text{total}} = E\{Y_i(z_{\text{max}}, 1) - Y_i(z_{\text{min}}, 0)\} = \tau^{\text{direct}} + \tau^{\text{indirect}}.$$

We call these three causal estimands as the direct, indirect and total effect for convenience, acknowledging that they are different from the ones in mediation analysis (Baron and Kenny, 1986; Imai et al., 2010). If Assumption 4 holds, namely $Y(z,t) = Y(t)$, then the direct effect of the IV equals zero and thus $\tau^{\text{total}} = \tau^{\text{indirect}}$. Therefore, we use the ratio of τ^{direct} and τ^{total} as the sensitivity parameter $r = \tau^{\text{direct}}/\tau^{\text{total}}$. When Assumption 4 holds, $r = 0$ and $|r|$ increases as the degree of violation to Assumption 4 increases. If we assume the scale of the direct effect $|\tau^{\text{direct}}|$ never exceeds that of the total effect $|\tau^{\text{total}}|$, the maximum value of $|r|$ equals one.

To conduct sensitivity analysis with respect to exclusion restriction, we add an IV term into the outcome model (11),

$$\text{logit}\{\Pr(Y_i(z,t) \leq k \mid X_i, S_i, C_i)\} = \alpha_{t,k} + \beta_t' X_i + \gamma_t S_i + \delta z + \nu_{t,C_i},$$

where δ measures how the IV directly affects the outcome. The coefficient δ determines the sensitivity parameter r based on the formula

$$r = \frac{\sum_{i=1}^{N}\sum_{k=1}^{K-1}\{\text{sig}(\alpha_{0,k} + \beta_0' X_i + \gamma_0 S_i + \delta z_{\text{min}}) - \text{sig}(\alpha_{0,k} + \beta_0' X_i + \gamma_0 S_i + \delta z_{\text{max}})\}}{\sum_{i=1}^{N}\sum_{k=1}^{K-1}\{\text{sig}(\alpha_{0,k} + \beta_0' X_i + \gamma_0 S_i + \delta z_{\text{min}}) - \text{sig}(\alpha_{1,k} + \beta_1' X_i + \gamma_1 S_i + \delta z_{\text{max}})\}}.$$

With a fixed value of δ and thus r, we can follow exactly the same procedure as before to estimate τ^{ATT} and τ^{PRTE}.

We plot the estimated τ^{ATT} and τ^{PRTE} against the sensitivity parameter r in Figure 5, with r on a grid ranging from -0.5 to 0.5. As a result, we only examine the case where the direct effect from the IV takes up at most 50% of total effect, which is a very large proportion in our application. For illustration purposes, we only display the effect on confidence measure of the urban males, which exhibits a strong negative effect under exclusion restriction. Figure 5 shows an increasing trend between the treatment effect and sensitivity parameter r. This pattern is reasonable as the sensitivity parameter characterizes how much of the effect can be explained by the influence of the IV acting directly on the outcome without affecting the treatment. If r is close to 1, a large proportion of the difference can be explained by the direct impact of the IV, which leads the indirect effect through the treatment changing from negative to zero. Therefore, our previous conclusions are more likely to incorrectly identify the negative effect when the sensitivity parameter r is larger.

We examine the threshold below which the indirect effect remains significantly negative, and thus quantify the robustness of our previous estimates. The estimated effect remains significantly negative

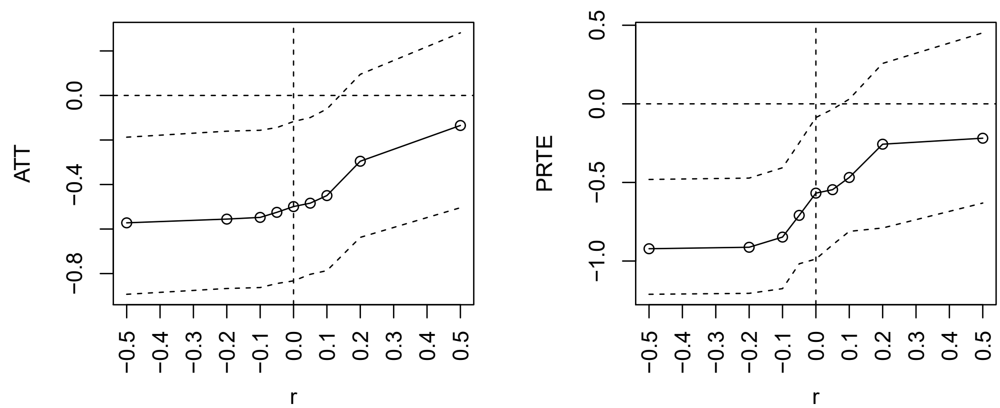

Figure 5: Posterior means and 95% credible intervals of τ^{PRTE} and τ^{ATT} against the sensitivity parameter r, for confidence measure and urban males; $r = 0$ when exclusion restriction holds.

.

when r is below 0.1. Namely, as long as the IV takes up less than 10% change in the variation of the outcomes, the impact of being an only child on the confidence measure remains negative for urban males. In Figure 5, the sensitivity of τ^{ATT} and τ^{PRTE} with respect to Assumption 4 share nearly identical trend against the sensitivity parameter. We relegate the sensitivity analysis results for the other measures and subgroups to the supplementary document.

5 Simulations

We further examine the operating characteristics of the proposed local IV method and compare it with the OLS and 2SLS methods in simulated scenarios. We independently simulate a covariate $X_i \sim \mathcal{N}(0,1)$ and an IV $Z_i \sim \mathcal{N}(0,1)$. We posit the following latent threshold model for the treatment

T_i:

$$S_i = 0.5 - 0.5X_i + \varepsilon_i,$$
$$T_i = \mathbf{1}_{Z_i \geq S_i},$$

where $\varepsilon_i \sim \mathcal{N}(0,1)$. The latent threshold S_i is a function of the covariate, and an individual is treated when Z_i exceeds the threshold value. For the potential outcomes, we posit linear models:

$$Y_i(0) = b_{00} + b_{01}X_i + e_i,$$
$$Y_i(1) = b_{10} + b_{11}X_i + e_i + l_i,$$

We assume that $(e_i, l_i, \varepsilon_i)$ follows multivariate Normal with mean zero, variance one, and the following correlation structure

$$\begin{pmatrix} 1 & 0 & p \\ 0 & 1 & h \\ p & h & 1 \end{pmatrix}.$$

The above model allows us to separate the correlation between treatment T_i and potential outcomes $Y_i(1), Y_i(0)$ and the correlation between T_i and treatment effect $Y_i(1) - Y_i(0)$. Specifically, the correlation p between errors e_i and ε_i controls the degree of confounding between the treatment and the outcome, and the correlation h between l_i and ε_i controls the degree of treatment effect heterogeneity with respect to the threshold value ε_i. A larger $|h|$ corresponds to larger heterogeneity, namely, the marginal treatment effect curve τ^{MTE} has larger variation across the principal stratum S_i.

To be consistent with the application, we focus on the estimands τ^{ATT} and τ^{PRTE}, the finite sample version of which in the simulated dataset b are calculated as follows,

$$\tau_b^{\text{ATT}} = \sum_{i=1}^{N} T_i\{Y_i(1) - Y_i(0)\} \Big/ \sum_{i=1}^{N} T_i,$$
$$\tau_b^{\text{PRTE}} = \sum_{i=1}^{N} \delta_i\{Y_i(1) - Y_i(0)\} \Big/ \sum_{i=1}^{N} \delta_i, \text{ with } \delta_i = \mathbf{1}_{z_{\min} \leq S_i \leq z_{\max}}.$$

We then apply the local IV method described in Section 3 and OLS and 2SLS described in Section 4.3.

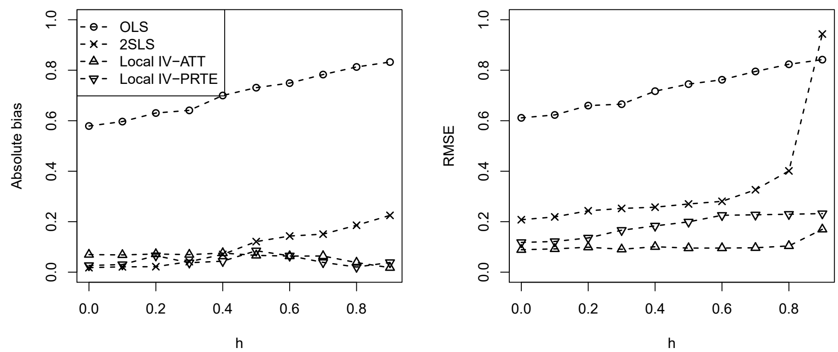

Figure 6: The absolute bias (left) and RMSE (right) of local IV (estimating τ^{ATT} and τ^{PRTE}), OLS and 2SLS estimates under different degree of treatment effect heterogeneity, h, in the simulation studies.

We fix the parameters $b_{00} = 0, b_{01} = 1, b_{10} = 3, b_{11} = 2, p = 0.5$ and sample size $N = 1000$, and vary the degree of heterogeneity h. For each setting, we simulate 1000 replicates and apply local IV, OLS and 2SLS to each replicate. We calculate the absolute bias and squared root of the mean squared error (RMSE) of τ^{ATT} and τ^{PRTE} of each method. Figure 6 displays the absolute bias and RMSE, respectively, between the local IV (for τ^{ATT} and τ^{PRTE}), OLS, and 2SLS estimates under different degree of heterogeneity h. We can see that the local IV approach obtains lower bias and RMSE compared with OLS and 2SLS in estimating both τ^{ATT} and τ^{PRTE}. Moreover, as the treatment effects become more heterogeneous between different principal strata, the advantage of local IV over OLS and 2SLS in terms of bias and RMSE becomes larger. This supports the use of the local IV method in settings with heterogeneous treatment effects, which is likely to be the case in the OCP application.

6 Conclusion

Leveraging the different implementation intensity of the OCP as a natural experiment, we evaluated the causal effects of being an only child on self-reported psychological health measures using data from the China Family Panel Studies. We found small but significant negative effects of being an only child. We also found two sources of treatment effect heterogeneity. First, the negative effects are more pronounced among those from urban areas. Second, the effects decrease with the latent degree of family resistance to the policy, characterized by a selection model of the family decision. Our results support the importance of sibship in psychological development. In the context of quantity-quality trade-off, our results suggest that the decreasing family size or the lack of sibship is not necessarily beneficial to individuals' psychological development. Namely, the "quality" of child, measured by the psychological health in our application, does not improve as the "quantity" decreases. Future work would investigate the possible trade-off between family size or the number of children and the well-being of children from other aspects.

From a methodological perspective, we made several extensions of the local IV method for continuous IV analysis (Heckman and Vytlacil, 1999). We elucidate an intrinsic connection between local IV and principal stratification, which allows us to employ Bayesian hierarchical models to accommodate complex data structure such as clustering. Within the same framework, one could also adopt more flexible Bayesian semiparametric or nonparametric models. For simplicity, we focused on average causal estimands for ordinal outcomes, while acknowledging the literature on ordinal outcome in the presence of noncompliance (Cheng, 2009; Baker, 2011) and 2SLS models for ordinal outcomes (Miranda and Rabe-Hesketh, 2006). We have also explored nonadditive estimands as the ones proposed in Lu et al. (2018) and found similar patterns as the average effects.

Acknowledgements

The authors are grateful to the editor of the special issue and two anonymous reviewers for constructive comments that help improve the clarity and exposition of the article. We also benefited from discussions with Liyun Chen, Chaoran Yu, Wentao Xiong, Zhichao Jiang, Wei Li, Yi Zhou, Wei Huang, and Yu Xie. Peng Ding was supported by the U.S. National Science Foundation (grants # 1713152 and # 1945136).

References

Agresti, A. (2003), *Categorical Data Analysis*, John Wiley & Sons, New York.

Angrist, J., Lavy, V., and Schlosser, A. (2010), "Multiple experiments for the causal link between the quantity and quality of children," *Journal of Labor Economics*, 28, 773–824.

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996), "Identification of causal effects using instrumental variables (with discussion)," *Journal of the American statistical Association*, 91, 444–455.

Attane, I. (2002), "China's family planning policy: an overview of its past and future," *Studies in Family Planning*, 33, 103–113.

Baker, S. G. (2011), "Estimation and inference for the causal effect of receiving treatment on a multinomial outcome: an alternative approach," *Biometrics*, 67, 319–323.

Banister, J. (1991), *China's changing population*, Stanford University Press.

Baron, R. M. and Kenny, D. A. (1986), "The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations." *Journal of Personality and Social Psychology*, 51, 1173.

Bartolucci, F. and Grilli, L. (2011), "Modeling partial compliance through copulas in a principal stratification framework," *Journal of the American Statistical Association*, 106, 469–479.

Becker, G. S. and Lewis, H. G. (1973), "On the Interaction between the Quantity and Quality of Children," *Journal of Political Economy*, 81, S279–S288.

Black, S. E., Devereux, P. J., and Salvanes, K. G. (2005), "The more the merrier? The effect of family size and birth order on children's education," *The Quarterly Journal of Economics*, 120, 669–700.

Brody, G. H. (2004), "Siblings' direct and indirect contributions to child development," *Current Directions in Psychological Science*, 13, 124–126.

Cameron, L., Erkal, N., Gangadharan, L., and Meng, X. (2013), "Little emperors: behavioral impacts of China's One-Child Policy," *Science*, 339, 953–957.

Carneiro, P., Heckman, J. J., and Vytlacil, E. J. (2011), "Estimating marginal returns to education," *American Economic Review*, 101, 2754–2781.

Carneiro, P., Lokshin, M., and Umapathi, N. (2017), "Average and marginal returns to upper secondary schooling in Indonesia," *Journal of Applied Econometrics*, 32, 16–36.

Cheng, J. (2009), "Estimation and inference for the causal effect of receiving treatment on a multinomial outcome," *Biometrics*, 65, 96–103.

Chu, C. C., Xie, Y., and Yu, R.-r. (2007), "Effects of sibship structure revisited: Evidence from intrafamily resource transfer in Taiwan," *Sociology of Education*, 80, 91–113.

Ding, P. and Li, F. (2018), "Causal inference: A missing data perspective," *Statistical Science*, 33, 214–237.

Dunn, J. (1988), "Sibling influences on childhood development," *Journal of Child Psychology and Psychiatry*, 29, 119–127.

Ebenstein, A. (2010), "The "missing girls" of China and the unintended consequences of the one child policy," *Journal of Human Resources*, 45, 87–115.

Forastiere, L., Mealli, F., and VanderWeele, T. J. (2016), "Identification and estimation of causal mechanisms in clustered encouragement designs: Disentangling bed nets using Bayesian principal stratification," *Journal of the American Statistical Association*, 111, 510–525.

Frangakis, C. E. and Rubin, D. B. (2002), "Principal stratification in causal inference," *Biometrics*, 58, 21–29.

Frangakis, C. E., Rubin, D. B., and Zhou, X.-H. (2002), "Clustered encouragement designs with individual noncompliance: Bayesian inference with randomization, and application to advance directive forms," *Biostatistics*, 3, 147–164.

Gelman, A. and Rubin, D. B. (1992), "Inference from iterative simulation using multiple sequences (with discussion)," *Statistical Science*, 7, 457–472.

Gelman, A. et al. (2006), "Prior distributions for variance parameters in hierarchical models (Comment on article by Browne and Draper)," *Bayesian Analysis*, 1, 515–534.

Glickman, M. E. and Normand, S.-L. T. (2000), "The derivation of a latent threshold instrumental variables model," *Statistica Sinica*, 517–544.

Hanushek, E. A., Rivkin, S. G., and Jamison, D. T. (1992), "Improving educational outcomes while controlling costs," in *Carnegie-Rochester Conference Series on Public Policy*, volume 37, Elsevier, 205–238.

Heckman, J. J. (1978), "Dummy endogenous variables in a simultaneous equation system," *Econometrica*, 46, 931–959.

— (1979), "Sample selection bias as a specification error," *Econometrica*, 47, 153–161.

Heckman, J. J. and Pinto, R. (2018), "Unordered monotonicity," *Econometrica*, 86, 1–35.

Heckman, J. J. and Vytlacil, E. (2001), "Policy-relevant treatment effects," *American Economic Review*, 91, 107–111.

— (2005), "Structural equations, treatment effects, and econometric policy evaluation," *Econometrica*, 73, 669–738.

Heckman, J. J. and Vytlacil, E. J. (1999), "Local instrumental variables and latent variable models for identifying and bounding treatment effects," *Proceedings of the National Academy of Sciences*, 96, 4730–4734.

Huang, W., Lei, X., and Sun, A. (2020), "Fertility Restrictions and Life-Cycle Outcomes: Evidence from the One-Child Policy in China," *Review of Economics and Statistics*, in press.

Huang, W., Lei, X., and Zhao, Y. (2016), “One-child policy and the rise of man-made twins,” *Review of Economics and Statistics*, 98, 467–476.

Huang, W. and Zhou, Y. (2015), “One-child policy, marriage distortion, and welfare loss,” .

Imai, K., Keele, L., and Tingley, D. (2010), “A general approach to causal mediation analysis.” *Psychological Methods*, 15, 309–334.

Jin, H. and Rubin, D. B. (2008), “Principal stratification for causal inference with extended partial compliance,” *Journal of the American Statistical Association*, 103, 101–111.

Jo, B., Asparouhov, T., Muthén, B. O., Ialongo, N. S., and Brown, C. H. (2008), “Cluster randomized trials with treatment noncompliance.” *Psychological Methods*, 13, 1–18.

Ju, C. and Geng, Z. (2010), “Criteria for surrogate end points based on causal distributions,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72, 129–142.

Kline, P. and Walters, C. R. (2019), “On Heckits, LATE, and Numerical Equivalence,” *Econometrica*, 87, 677–696.

Leibowitz, A. (1974), “Home investments in children,” *Journal of Political Economy*, 82, S111–S131.

Likert, R. (1932), “A technique for the measurement of attitudes.” *Archives of Psychology*, 140, 1–55.

Liu, H. (2014), “The quality–quantity trade-off: evidence from the relaxation of China’s one-child policy,” *Journal of Population Economics*, 27, 565–602.

Lu, J., Ding, P., and Dasgupta, T. (2018), “Treatment Effects on Ordinal Outcomes: Causal Estimands and Sharp Bounds,” *Journal of Educational and Behavioral Statistics*, 43, 540–567.

McHale, S. M., Updegraff, K. A., and Whiteman, S. D. (2012), “Sibling relationships and influences in childhood and adolescence,” *Journal of Marriage and Family*, 74, 913–930.

Miranda, A. and Rabe-Hesketh, S. (2006), “Maximum likelihood estimation of endogenous switching and sample selection models for binary, ordinal, and count variables,” *The Stata Journal*, 6, 285–308.

Mogstad, M. and Wiswall, M. (2016), “Testing the quantity-quality model of fertility: Estimation using unrestricted family size models,” *Quantitative Economics*, 7, 157–192.

Plummer, M. et al. (2003), “JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling,” in *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, volume 124, Vienna, Austria.

Qian, N. (2009), “Quantity-quality and the one child policy: The only-child disadvantage in school enrollment in rural China,” , NBER Working Paper 14973.

Rohrer, J. M., Egloff, B., and Schmukle, S. C. (2015), “Examining the effects of birth order on personality,” *Proceedings of the National Academy of Sciences*, 112, 14224–14229.

Rosenzweig, M. R. and Zhang, J. (2009), “Do population control policies induce more human capital investment? Twins, birth weight and China’s “one-child” policy,” *The Review of Economic Studies*, 76, 1149–1174.

Rubin, D. B. (1978), “Bayesian Inference for Causal Effects: The Role of Randomization,” *The Annals of Statistics*, 6, 34–58.

— (1980), “Comment on “Randomization analysis of experimental data: The Fisher randomization test” by D. Basu,” *Journal of the American Statistical Association*, 75, 591–593.

Schwartz, S. L., Li, F., and Mealli, F. (2011), “A Bayesian semiparametric approach to intermediate variables in causal inference,” *Journal of the American Statistical Association*, 106, 1331–1344.

Small, D. S. (2007), “Sensitivity analysis for instrumental variables regression with overidentifying restrictions,” *Journal of the American Statistical Association*, 102, 1049–1058.

Vytlacil, E. (2002), “Independence, monotonicity, and latent index models: An equivalence result,” *Econometrica*, 70, 331–341.

Wang, X., Jiang, Y., Zhang, N. R., and Small, D. S. (2018), “Sensitivity analysis and power for instrumental variable studies,” *Biometrics*, 74, 1150–1160.

Wu, L. (2014), *Are Only Children Worse Off on Subjective Well-Being?: Evidence from China's One-Child Policy*, Ph.D. thesis, Hong Kong University of Science and Technology.

Xie, Y. and Hu, J. (2014), "An introduction to the China family panel studies (CFPS)," *Chinese Sociological Review*, 47, 3–29.