The Intergenerational Persistence of Child Maltreatment*

E. Jason Baron† Ezra G. Goldstein‡
Christopher Mills§ Emily Putnam-Hornstein¶

February 2026

PLEASE DO NOT CITE OR CIRCULATE

Abstract

Parental investments are crucial inputs for economic opportunity that have the potential to be transmitted across generations. We use novel linkages of California administrative data across nearly three decades to estimate the prevalence and intergenerational transmission of a leading costly disinvestment, child maltreatment, while documenting its implications for inequality and opportunity. While it is known that approximately one third of U.S. children are ever reported to Child Protective Services (CPS), we estimate that approximately 61 percent of children have indirect exposure through a *parent* who was reported to CPS as an alleged maltreatment victim during the parent's own childhood. We show that exposure to alleged maltreatment—indirect or direct—is disproportionately borne by low-income children: Among children in the bottom decile of the income distribution, approximately 80 percent have a parent with CPS contact (38 percent substantiated victim), and about 50 percent of children in the bottom decile are directly reported themselves (20 percent substantiated victim), compared to much lower rates in the top income deciles. Using parent-child linkages, we document that maltreated individuals are 64 percent more likely to have a maltreated child. A variety of empirical tests—including comparisons of siblings with different maltreatment exposure—suggest that lived experiences of maltreatment may themselves be causally transmitted, rather than simply reflecting a persistence of poverty, disproportionate surveillance, or genetics. Finally, we demonstrate that neighborhood concentration of CPS involvement is one of the strongest (negative) predictors of economic mobility. Together, these findings imply that child maltreatment may operate as a behaviorally transmitted barrier to intergenerational mobility, reinforcing disadvantage across generations.

*We thank Robert Collinson, Max Gross, Melissa Kearney, Kate Musen, and seminar participants at BYU, IU Indianapolis, Notre Dame, and PPCW 2025 for their helpful comments and suggestions. John Prindle, Himal Suthar, and the Children's Data Network provided crucial support for linking California administrative datasets featured in this paper. We acknowledge AI tool support from Anthropic, Google, and OpenAI for reviewing and refining elements of the paper. All errors are our own.

†Duke University and NBER. E-mail: jason.baron@duke.edu.
‡Georgia Institute of Technology. E-mail: eggoldstein@gatech.edu.
§University of Notre Dame. E-mail: chris.mills@nd.edu.
¶University of North Carolina Chapel Hill. E-mail: eph@unc.edu.

I Introduction

Parental investments are crucial inputs into children's well-being and long-run economic outcomes (Cunha and Heckman, 2007; Del Boca, Flinn and Wiswall, 2013; Heckman, 2006; Heckman and Masterov, 2007). The production of human capital, however, is not shaped solely by positive investments: it is also contingent on parental *disinvestments* and similar negative shocks that deplete a child's potential for social mobility. If such disinvestments function as barriers to opportunity, and if the propensity to disinvest is itself transmitted across generations, then they may play an important role in the intergenerational persistence of income and overall well-being.

Child abuse and neglect (hereafter maltreatment) represents one of the most severe and consequential forms of parental disinvestment. Maltreatment is both widespread and costly in the United States. By age 18, approximately one-third of children experience a formal investigation by Child Protective Services (CPS), and roughly five percent enter foster care (Kim et al., 2017; Wildeman and Emanuel, 2014). A substantial literature documents the long-run consequences of maltreatment, linking childhood abuse and neglect to lower educational attainment, reduced earnings, and increased involvement with the criminal justice system (Currie and Spatz Widom, 2010; Currie and Tekin, 2012). The aggregate economic burden is correspondingly large: Peterson, Florence and Klevens (2018) estimate that the lifetime cost of new cases of child abuse and neglect occurring in a single year may amount to as much as 10.9 percent of U.S. GDP.

A central concern is that child maltreatment may persist across generations. For example, CPS risk assessments routinely incorporate whether a parent was previously involved with CPS as a child, reflecting the widespread belief that individuals who experience maltreatment are more likely to maltreat their own children. Maltreatment is a particularly consequential form of childhood adversity because it directly disrupts the caregiving environment that underlies the formation of human capital, potentially attenuating the effectiveness of other opportunity-enhancing inputs such as schools, neighborhoods, or social programs. Because maltreatment is highly prevalent and disproportionately concentrated among families at the lower end of the income distribution, intergenerational transmission of maltreatment could constitute an important mechanism limiting social and economic mobility. Moreover, if maltreatment is transmitted across generations, commonly cited estimates of its prevalence and social costs—which typically focus on a single generation—may substantially understate its burden (Adermon, Lindahl and Palme, 2021; Finlay, Mueller-Smith and Street, 2023).

Despite the prominent belief of intergenerational transmission in policy discussions and

practice, rigorous evidence remains limited. Indeed, ex ante, neither the direction nor the magnitude of intergenerational transmission is obvious: individuals who experience maltreatment may actively avoid repeating similar behaviors, form partnerships that mitigate risk, or otherwise break cycles of harm, implying that observed persistence could be limited—or even absent—despite strong correlations in disadvantage.

Establishing whether maltreatment itself is transmitted across generations is challenging for two primary reasons. First, a rigorous analysis requires data that link childhood maltreatment to future fertility and to maltreatment outcomes among the next generation. Given the long time horizons involved and privacy constraints governing administrative records, such measurement has historically been infeasible. As a result, existing estimates largely rely on retrospective surveys or small, selective samples (Widom, Czaja and DuMont, 2015). Second, identifying causal transmission is complicated by usual selection concerns. Maltreatment risk differs systematically by socioeconomic status, race, neighborhood conditions, and many other factors. Consequently, observed correlations between parent and child CPS contact may reflect persistent economic disadvantage or shared environments rather than the behavioral transmission of maltreatment itself.

In this paper, we employ novel administrative data linkages in California to study children's cumulative exposure to child maltreatment—direct and indirect through parents—and the intergenerational transmission of child maltreatment from parent to child. We construct an intergenerational dataset linking the universe of CPS records to state birth certificates over a 26-year period. These linkages allow us to document the prevalence of maltreatment from birth through adulthood and to map childhood CPS involvement onto fertility decisions and maltreatment outcomes in the next generation. We observe the full lifecycle of CPS involvement—from initial referral to investigation, case opening, and foster care placement—allowing us to study a gradient of CPS contact and how exposure at different levels in one generation relates to CPS involvement in the next.

Before turning to estimates of intergenerational transmission, we first document the prevalence and distribution of CPS involvement using our linked administrative data. Tracking complete birth cohorts of children born in California from 1999-2001 through age 18, we find that contact with the child welfare system is remarkably common. At least 29 percent of children are referred to CPS, 26 percent are investigated, and 10 percent are substantiated as victims of abuse or neglect. More intensive forms of involvement are also nontrivial: six percent of children experience an open CPS case, and four percent are placed in foster care at least once during childhood. Consistent with previous studies (e.g., Doyle, Emanuel and Eyzaguirre (2025); Wildeman and Emanuel (2014); Yi, Edwards and Wildeman (2020)), these estimates

underscore that CPS involvement and child maltreatment are not rare events but a widespread feature of childhood in the U.S., with a meaningful share of children experiencing deep system contact.

We provide some of the first estimates of CPS involvement and child maltreatment stratified by socioeconomic standing. Using rich information from birth certificates, we show that 49 percent of children born into the bottom decile of the predicted income distribution are referred to CPS by age 18, nearly 43 percent are investigated, and 19 percent are substantiated victims of abuse or neglect. In contrast, among children in the top income decile, only about 9 percent are referred to CPS. Similar gradients emerge when stratifying by maternal education, Medicaid coverage at birth, and household instability. We further document that these socioeconomic gradients persist across racial groups and are mirrored in objective outcomes such as potential-maltreatment-related child mortality, suggesting that observed disparities in CPS contact reflect meaningful differences in underlying risk rather than differential surveillance alone. Taken together, these patterns highlight that child maltreatment is both widespread and sharply stratified by early-life disadvantage—conditions that make intergenerational transmission a potentially important mechanism for the persistence of economic mobility.

Our prevalence patterns raise a core insight: A child has two biological parents, and these parents may *each* have experienced maltreatment that can be passed on to a child. Using parent-child linkages in administrative data, we conduct a backward-looking analysis that asks: among children who come into contact with CPS early in life, how common is a parental history of CPS involvement? Focusing on the universe of children born in California in 2022, we link early-life CPS contact for children to their parents' own childhood CPS histories. We estimate that intergenerational prevalence is strikingly high. Approximately 81 percent of children with CPS contact in their first three years of life (and 61 percent of all children) have a parent with a history of CPS involvement.

We then turn to a forward-looking analysis that examines whether childhood maltreatment increases the likelihood of maltreatment in the next generation. We follow cohorts of girls born in California between 1987 and 2001 from childhood into adulthood, linking their CPS histories as children to CPS outcomes for their own children. Descriptively, women with a substantiated maltreatment allegation during childhood are roughly 64 percent more likely to have a child who is themselves substantiated for maltreated, relative to women with no substantiated childhood history. These differences are evident across multiple measures of system involvement, including referrals, substantiations, case openings, and foster care placements among their children.

While these descriptive intergenerational correlations are large, they are unlikely to be causal. Mothers who experienced substantiated maltreatment as children differ from other mothers along many observed and unobserved dimensions, including socioeconomic background, neighborhood conditions, and family instability. To address these concerns, we employ two complementary empirical strategies. First, we estimate selection-on-observables models with rich controls and focus on the intensive margin of CPS involvement, comparing girls who were all reported to CPS but differed in whether maltreatment was substantiated. Second, we implement a sibling fixed effects design that compares sisters listed on the same childhood maltreatment referral but who experienced different levels of substantiation, thereby differencing out shared family background, parental characteristics, and other time-invariant household factors.

Across specifications, we find consistent evidence of intergenerational transmission. In our preferred sibling fixed effects specification, a maternal history of substantiated maltreatment increases the expected number of times her children are substantiated by approximately 0.128 additional substantiations—an increase of roughly 22 percent relative to the sample average. We find similar or larger relative increases for other CPS outcomes, including referrals, case openings, and foster care placements. Importantly, these effects condition on fertility, indicating that intergenerational transmission is not driven solely by differences in childbearing behavior. Estimates are qualitatively similar using an alternative approach linking maltreated children to future perpetrator records that do not rely on fertility and can be estimated for either girls or boys. Although underpowered, we also employ two supplemental exercises that yield similar conclusions: within-sibling variation in maltreatment from an unrelated perpetrator outside the home, and within-sibling differences in years exposed to a maltreatment-prone home environment using age variation across sibling groups removed to foster care and thereafter no longer maltreated.

We next examine the mechanisms behind this persistence. The evidence points to behavioral transmission rather than purely economic or genetic channels. We find strong type-specific patterns in intergenerational maltreatment. Mothers who experienced physical abuse are most likely to have children who experience physical abuse; those who experienced sexual abuse are most likely to have children who experience sexual abuse; and those who experienced neglect show especially strong persistence in neglect. These patterns hold even when we control for other forms of maltreatment. Although cross-type associations are positive—indicating broader persistence in family dysfunction—the strongest relationship is between the same type of maltreatment across generations. This structure is difficult to explain by persistent poverty, general disadvantage, or increased monitoring alone. Such factors would likely raise overall CPS involvement across categories, but not selectively reproduce the same type of

harm. Instead, the evidence is consistent with the intergenerational transmission of specific caregiving behaviors learned in childhood, suggesting that maltreatment may be passed down through behavioral channels, not only through shared economic conditions.

Moreover, transmission rates are similar regardless of whether the mother herself was maltreated by a biological parent or by a non-biological caregiver. If shared genetics were the primary mechanism, persistence should be substantially stronger when maltreatment was perpetrated by a biological parent. The similarity of these magnitudes instead points to lived experience and environmental exposure as the primary channels of transmission.

Finally, we examine the relevance of maltreatment for economic opportunity and conduct a descriptive exercise relating local child maltreatment rates to measures of intergenerational mobility from Chetty et al. (2014). We document a strong negative correlation: neighborhoods with higher per-capita rates of CPS referrals exhibit substantially lower levels of upward income mobility for children born to low-income families. Neighborhood rates of CPS referrals remain a leading correlate of low upward mobility even after accounting for a rich set of socioeconomic and demographic characteristics—including many of the neighborhood factors previously shown to predict mobility. Although purely descriptive, these findings reinforce that child maltreatment is not only concentrated among disadvantaged families and persistent across generations, but may restrict opportunities for upward mobility.

Our study contributes to the literature documenting the intergenerational persistence of maltreatment in the U.S. by providing population-level estimates of transmission from administrative data. Although the theory that maltreatment transmits intergenerationally is pervasive, estimates of its prevalence have largely relied on self-reporting measures and retrospective reports which are frequently limited by recall bias, small sample sizes, and a lack of appropriate comparison groups (Berlin, Appleyard and Dodge, 2011; Thornberry, Knight and Lovegrove, 2012; Widom, 1989; Widom, Czaja and DuMont, 2015). An exception is Font et al. (2020), which leverages administrative data from Wisconsin and finds that, by age 25, between one-quarter and one-third of individuals with out-of-home care histories in adolescence experienced some form of parent or perpetrator CPS involvement. Despite this growing descriptive evidence, few studies attempt to disentangle whether intergenerational transmission reflects a causal effect of maltreatment exposure or selection, and even fewer investigate the mechanisms underlying observed persistence. We add to this literature by leveraging linked administrative records from California, measuring prevalence across socioeconomic backgrounds. Furthermore, we attempt to move beyond correlational evidence by employing a sibling fixed effects design. This approach mitigates concerns about selection by enabling comparisons of maltreatment histories within households, and also allows us to

explore potential mechanisms underlying intergenerational transmission.

Second, we contribute to the literature on the determinants of child maltreatment in the United States. Prior work documents that external conditions and policy interventions—such as extreme heat (Evans, Gazzé and Schaller, 2025), Early Intervention programs (Aizer and Brito Rebolledo, 2025), opioid exposure (Evans, Harris and Kessler, 2022; Gihleb, Giuntella and Zhang, 2022), and changes in economic resources and employment (Lindo, Schaller and Hansen, 2018; Rittenhouse, 2025)—shape maltreatment risk. We add to this literature by highlighting a previously underexplored determinant: parental exposure to maltreatment during childhood. We show that a parent's own maltreatment history is a strong predictor of CPS involvement in the next generation, even within families. As a result, interventions that reduce maltreatment today may yield benefits that extend beyond the directly treated child. In particular, evaluations of prevention programs—such as Early Intervention—may understate these programs' social returns if they do not account for intergenerational spillovers in maltreatment risk.

II CPS Involvement as a Measure of Maltreatment

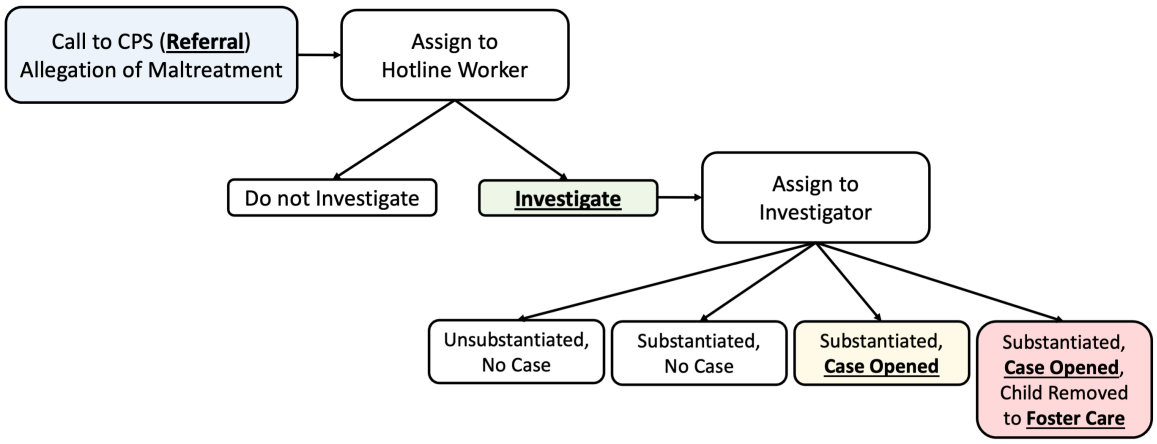
Before turning to our empirical estimates, we provide institutional context for how children come into contact with CPS and clarify how we interpret different forms of system involvement as signals of maltreatment. This background is important for understanding terms used throughout the paper, such as "investigation," "substantiation," "case opening," and "foster care," and for motivating our use of increasingly serious forms of CPS involvement as proxies for more severe or sustained exposure to maltreatment. While referrals mark a family's entry point into the child welfare system, deeper forms of involvement often reflect greater agency concern and greater expenditure of resources to address perceived risk.

The objective of CPS is to keep children safe from abuse and neglect, intervene to support families and children when the risk of ongoing maltreatment is present, and, when necessary, arrange alternative caregiving arrangements. In California, the setting of this study, CPS is administered at the county level; county agencies operate under the oversight of the California Department of Social Services (CDSS).

Figure 1 provides a simplified overview of the CPS system. A child enters the system through a report—or "referral"—from a party concerned that the child may be experiencing maltreatment. Referrals can originate from a wide range of reporters, including medical professionals, police officers, teachers, neighbors, relatives, or CPS staff themselves. These reports are typically made through a centralized hotline and routed to county-level call centers.

The first major decision occurs at the hotline stage: whether to "screen in" the referral for formal investigation, or to "screen out" the call if it does not meet statutory criteria for CPS intervention.

Figure 1: The CPS Process Simplified



Notes: This figure illustrates the core stages of Child Protective Services (CPS) involvement referenced in this paper, beginning with a referral and proceeding through potential investigation, substantiation, case opening, and placement in foster care. While exact procedures vary across states and counties, this diagram represents a canonical pathway through the system.

When a referral is screened in, it is assigned to a CPS investigator who typically conducts an unannounced home visit to assess the child's living environment. The investigator evaluates the allegations in the referral and determines whether the report is: substantiated, meaning there is sufficient evidence that abuse or neglect occurred based on state-defined criteria; unfounded, indicating there is not enough evidence to support the allegation and maltreatment is deemed not to have occurred; or inconclusive, where the evidence is unclear and does not allow a definitive determination either way.

If the investigator substantiates the allegation, they next must determine the appropriate level of intervention. There are three main outcomes. The first is no further action, typically used in lower-risk cases where maltreatment was minimal, situational, or appears resolved. In such a case, the investigator may suggest voluntary services aimed at reducing the risk of subsequent harm, but the agency establishes no formal ongoing action with the family. The second is a case opening for in-home services, applied when the family presents a moderate to high risk of future harm, but the child can remain safely at home with support (e.g., parenting classes, behavioral health therapy, substance use treatment, regular monitoring). The third

is an out-of-home placement (foster care), reserved for high-risk cases where remaining in the home poses immediate and ongoing danger. This outcome usually requires court approval and may result in temporary or extended separation from the parents. Depending on the circumstances, children may reunify with their biological families, be adopted (often by a relative or foster caregiver), or age out of the system.

Because these decisions require increasing levels of agency resources and legal involvement, we interpret deeper levels of CPS engagement as stronger signals of actual maltreatment. We place particular emphasis on substantiation as a measure of maltreatment. Unlike referrals or investigations—which may reflect suspicion, precautionary screening, or information gathering—substantiation represents a distinct determination by the agency that abuse or neglect occurred, based on statutory criteria and a supervisor's affirmation. Although substantiation decisions are necessarily discretionary, they have an important role in child welfare practice; substantiated findings can trigger eligibility for formal services and inform future risk assessments. For these reasons, we interpret substantiated allegation decisions as a particularly informative signal of exposure to abuse or neglect.

It is also important to clarify what CPS considers to be maltreatment. Under the federal Child Abuse Prevention and Treatment Act (CAPTA), maltreatment is broadly defined to include physical abuse, sexual abuse, emotional abuse, and neglect. Neglect is typically understood as a caregiver's failure to provide for a child's basic needs, such as food, shelter, medical care, supervision, or education, and is often linked to issues such as substance use disorders (Goldstein and Font, 2025; Palmer et al., 2022).

Although neglect can be affected by household resources (Rittenhouse, 2025), it is also known by practitioners that neglect is often distinct from poverty (Herd et al., 2022). States retain considerable discretion in interpreting and applying definitions of abuse and neglect, and most do not classify involuntary neglect due to material hardship, such as lack of housing or food arising from poverty, as maltreatment under state law. In California, this distinction is codified in statute. The California Welfare and Institutions Code § 300(b)(2) provides that a child shall not be found to be neglected solely due to "homelessness," "poverty," or "indigence or other conditions of financial difficulty."[1] In practice, this means that children experiencing significant deprivation may not be considered victims of maltreatment unless CPS determines that the caregiver is capable of providing adequate care but fails to do so. These distinctions are important for interpreting our estimates, which reflect CPS-defined maltreatment, not the broader set of adversities children may experience.

[1]See also the California Department of Social Services, *Child Welfare Services Manual*, Section 31-125, available at https://www.cdss.ca.gov/ord/entres/getinfo/pdf/cws2.pdf.

Using CPS records as a proxy for maltreatment has important strengths and limitations. On the one hand, CPS data are likely to understate the true prevalence of maltreatment, since they only capture situations that are reported to and formally acted on by the child protection system. Many incidents, particularly those involving emotional abuse or neglect in hard-to-detect forms, may go unreported, and survey-based studies often yield higher estimates of maltreatment exposure. On the other hand, CPS data may also overstate maltreatment prevalence, insofar as some reports are unfounded, driven by interpersonal disputes, or reflect minor concerns that do not rise to the level of significant risk. It is difficult to systematically assess the magnitude of such false positives. To address this uncertainty, our analysis leverages a gradient of system involvement to distinguish between relatively minor and more serious or sustained CPS contact.

III Data Sources and Record Linkage

Our analyses of both the prevalence of child maltreatment and its intergenerational transmission rely on administrative data from the State of California that allow us to observe CPS involvement over long horizons and to link children to their biological parents.

We draw on administrative records from two primary sources: the Child Welfare Services/Case Management System (CWS/CMS) and birth certificate records from the California Department of Public Health–Vital Records (CDPH-VR). These datasets are linked using a two-stage probabilistic record linkage procedure designed to accommodate missing, inconsistent, or error-prone identifiers common in administrative data.

In the first stage, there is a blocking procedure to restrict attention to plausible record pairs. Blocking requires an exact match on birth month and birth year, combined with fuzzy matching on first or last names. In the second stage, record pairs within a block are assigned match probabilities based on a rich set of features. These include exact and fuzzy matches on Social Security number (when available), full date of birth, child and parent names, parental identifiers, and geographic information. Match probabilities are generated using a machine learning algorithm following the probabilistic record linkage framework developed in Prindle, Suthar and Putnam-Hornstein (2023). Record pairs are then classified as matches if their predicted match probability exceeds a threshold of 0.825, following Prindle, Suthar and Putnam-Hornstein (2023).[2]

The CWS/CMS database contains statewide records for all children reported to CPS from

[2]Generally, the record linking algorithm performs well. In a simulation analysis, the procedure identified 94 percent of true matches, and just 1.7 percent of non-matches were labeled as matching (Prindle, Suthar and Putnam-Hornstein, 2023).

1998 through 2024 and captures the full lifecycle of CPS involvement, beginning with hotline referrals and extending through investigations, case openings, and foster care placements. Referral records include information on alleged maltreatment, the reporter's relationship to the child, and demographic characteristics of children and household members, as well as the hotline worker's screening decision. For screened-in referrals, the data track investigation outcomes and agency decision-making, including whether an individual allegation was substantiated, resulted in a case opening for ongoing in-home services, and whether the child was placed into foster care. The database also contains foster care placement histories, including placement settings, entry and exit dates, and reasons for entry and exit into and out of placement.

Birth certificate records cover the universe of in-state births from 1999 through 2023 and provide a rich set of covariates not observed in CPS data, including maternal education, parental age, race and ethnicity, prenatal care utilization, household structure at birth, and infant health at birth (e.g., birthweight and gestational age). Together, these linked data allow us to define birth cohorts, stratify CPS involvement by early-life socioeconomic characteristics, and examine both the prevalence and intergenerational persistence of CPS contact.

IV Prevalence of CPS Involvement and Child Maltreatment

Estimates of child maltreatment prevalence in the United States have traditionally relied on synthetic cohort methods applied to nationwide administrative datasets, including the National Child Abuse and Neglect Data System (NCANDS) and the Adoption and Foster Care Analysis and Reporting System (AFCARS). These systems track CPS investigations and foster care placements using anonymized child-level identifiers, allowing researchers to construct longitudinal event histories. Seminal work by Kim et al. (2017) using this approach estimated that approximately 37 percent of U.S. children experience a CPS investigation before age 18. Related studies estimate that 12.5 percent of children are ever substantiated victims of maltreatment, and up to 6 percent are placed in foster care during childhood (Wildeman and Emanuel, 2014; Wildeman et al., 2014). These widely cited figures underscore the large scale of child welfare involvement nationwide.

However, synthetic cohort methods rest on important assumptions—namely, that age-specific rates of involvement remain stable over time and that administrative data accurately capture each child's first contact with the system. Recent research by Doyle, Emanuel and Eyzaguirre (2025) evaluates these assumptions using a direct birth cohort approach, constructing life

tables from over two decades of AFCARS data. Tracking children born between 2000 and 2002, they find that 5 percent entered foster care by age 18, suggesting that synthetic cohort estimates may modestly overstate lifetime risk—particularly during adolescence.

While NCANDS and AFCARS offer rich national data on system contact, their lack of linkage to external information on children and families makes it difficult to analyze how CPS involvement varies by other factors. For instance, researchers may want to understand how CPS contact varies across the income distribution to better understand the role of maltreatment in shaping inequality and economic mobility. As a result, there is still much to learn about the distribution of CPS involvement.[3]

This section builds on this literature by producing new birth cohort-based estimates of child welfare involvement using linked administrative data from California. Our dataset combines detailed CPS records with birth and death certificate data, enabling us to follow the full population of children born in 1999-2001 through age 18. This linkage allows for disaggregation of CPS involvement by a wide range of characteristics observable at birth, including race, predicted parental income, and parental education.

We estimate the cumulative prevalence of multiple forms of child welfare contact—including referrals, investigations, substantiations, case openings, and foster care placements. Importantly, the inclusion of death certificate data also allows us to examine patterns of child mortality plausibly related to maltreatment, helping to distinguish between true differences in underlying risk and disparities that may reflect differential surveillance. Our aim is to provide a comprehensive, descriptive account of early-life contact with the CPS system in a large, diverse state, shedding light on both the magnitude and distribution of child welfare involvement and laying the groundwork for our investigation of the importance of child maltreatment in shaping economic opportunity and social mobility in the U.S.

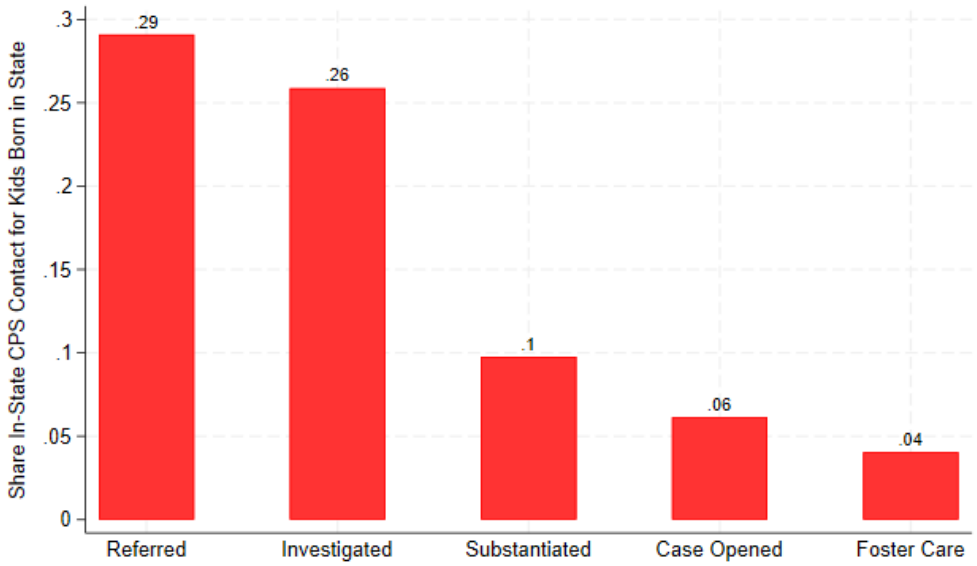
IVA Overall Prevalence

Figure 2 summarizes the cumulative prevalence of CPS involvement for the 1999-2001 birth cohorts in California. For each stage of the child welfare system—referral, investigation, substantiation, case opening, and foster care placement—the figure reports the share of children who experienced that form of contact at least once by age 18. We find that approximately 30 percent of children were referred to CPS, 27 percent were the subject of an investigation, and 10 percent had at least one substantiated allegation of maltreatment. Of the full cohort, 6 percent had an open case, and 4 percent were placed in foster care at least

[3] A few notable exceptions have leveraged linked administrative data to begin addressing these gaps. For examples, see Putnam-Hornstein et al. (2021) and Ryan et al. (2018).

once before adulthood. These results underscore that CPS contact is far from rare: nearly one in three children is referred to CPS, and a nontrivial share are subject to deeper involvement, including out-of-home placement.

Figure 2: Overall Prevalence of Childhood CPS Contact in California



Notes: This figure reports the cumulative share of children in the 1999-2001 California birth cohorts who experienced each type of CPS involvement before their 18th birthday. Involvement is defined in non–mutually exclusive terms. For example, children who were placed in foster care are also included in the categories of referral, investigation, substantiation, and case opening. All estimates are based on linked administrative records from California birth certificates and CPS data.

Note that these estimates likely represent a conservative lower bound on the true prevalence of maltreatment exposure among children born in California for two reasons. Because we rely on state administrative records, our data capture only those CPS contacts that occur within California. Children who move out of state during childhood are no longer observable in the CPS system, even if they are subsequently investigated elsewhere. As a result, children who were born in California but exited the state before experiencing CPS involvement are treated as never having had contact with the system. Moreover, our administrative data captures child maltreatment referrals, and the true prevalence may be higher if maltreatment is under-reported. These limitations suggest that the true cumulative risk of CPS involvement for this cohort may be higher than what we observe.

IVB Prevalence by Race

The overall rates of CPS involvement in California mask substantial disparities by race and ethnicity. Figure A.1 shows the cumulative prevalence of each type of CPS contact by age 18, disaggregated by race and ethnicity. The differences are striking: among Black children born in California in the 1999-2001 cohorts, 50 percent were referred to CPS, 47 percent were investigated, 20 percent were the subject of at least one substantiated allegation, 15 percent had an open case, and 11 percent were placed in foster care. For Hispanic children, the corresponding rates were 32 percent, 29 percent, 10 percent, 5 percent, and 3 percent. Non-Hispanic white children experienced lower rates of involvement across virtually all categories: 26 percent referred, 22 percent investigated, 9 percent substantiated, 4 percent with a case open, and 3 percent placed in foster care. CPS contact rates were even lower for Asian and Pacific Islander children across all stages of involvement.

IVC Prevalence by Socioeconomic Status at Birth

To explore heterogeneity in CPS involvement by SES, we take advantage of the fact that we observe a rich set of covariates for the mother and the father of the child at the time of birth from the birth certificate data. Because direct income measures are not available on the birth certificate, we impute family income deciles and quartiles using a linear prediction model trained on data from the American Community Survey (ACS). Online Appendix BA describes the model in more detail.

Figure 3 presents the cumulative prevalence of CPS contact by age 18 across five stages—referral, investigation, substantiation, case opening, and foster care placement—disaggregated by these income and SES indicators. The differences across the income distribution are stark. Panel (a) shows that, among children in the bottom income decile, 54 percent are referred to CPS, 43 percent are investigated, 19 percent are substantiated, 13 percent have a case opened, and 10 percent are placed in foster care. In contrast, for children in deciles 2 through 10, the respective rates fall to 27 percent, 24 percent, 9 percent, 5 percent, and 3 percent.

When looking at heterogeneity by income quartile in Panel (b) instead, the gradient is similarly pronounced. In the bottom quartile, 46 percent of children are referred to CPS, nearly 20 percent have a substantiated investigation, and 9 percent are placed in foster care. These rates decline steadily across quartiles. In the top quartile, only 12 percent of children are referred, 2 percent are substantiated victims, and fewer than 1 percent placed in foster care. Notably, however, CPS involvement is far from negligible even in the third quartile: 25 percent

of children are referred to CPS. This highlights that while CPS contact is significantly more prevalent at the bottom of the income distribution, it remains surprisingly common even in the upper-middle segments of the distribution

Because our income measure is based on prediction, we also construct a supplementary "SES Index" using a related set of observable characteristics: Medicaid coverage at birth, maternal education, whether a father is listed on the birth record, paternal education (if available), and an imputed measure of neighborhood income based on birth hospital.[4] Each component variable is standardized to have mean zero and unit variance, with higher values indicating a more traditionally advantageous outcome. Missing values are imputed with the variable's mean. The SES Index is then calculated as the sum of these standardized components and re-standardized to have mean zero and variance one.

In Panel (c), we show CPS involvement by SES quartile. Among children in the bottom quartile, about 43 percent are referred to CPS, 40 percent are investigated, and about 17 percent are substantiated victims. By contrast, among those in the top quartile, roughly 11 percent are referred, 9 percent investigated, and just 2 percent substantiated. These patterns underscore the strong socioeconomic gradient in early-life CPS contact.

The patterns are similarly stark across other proxies for disadvantage. At least 40 percent of children whose mothers had less than a high school degree (Panel (d)), or whose births were covered by Medicaid (Panel (e)), are referred to CPS—compared to around 10 percent among children of more-educated mothers (with at least a bachelor's degree) or around 20 percent for those with private coverage. Around 55 percent of children whose fathers are not listed on the birth certificate—a potential signal of household instability—are referred to CPS (Panel (f)), and more than 25 percent become a substantiated victim of abuse or neglect.

IVD Prevalence by SES-by-Race

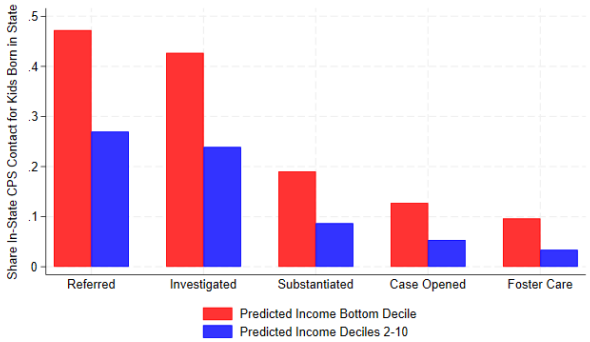
Having documented disparities in CPS involvement by both race and SES, we next examine their interaction. Understanding how CPS contact varies across combinations of race and early-life SES may reveal patterns that are masked when considering these dimensions alone.

To reduce dimensionality, we focus on two key outcomes: referrals—the most common form of contact—and substantiations—the main treatment variable in our subsequent analyses. We disaggregate each outcome by race and SES, using three different measures of socioeconomic status: predicted parental income at birth, our constructed SES index, and maternal

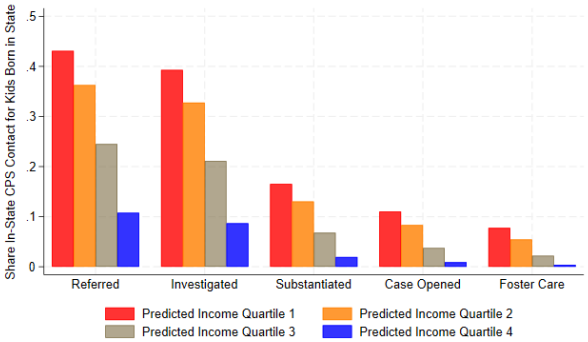
[4]We create this variable by using the mom's ZIP code of residence—available from 2018 onward—to estimate the average ZIP-code-level median household income for families delivering at each hospital in 2018–2019.

Figure 3: CPS Contact by Socioeconomic Characteristics at Birth

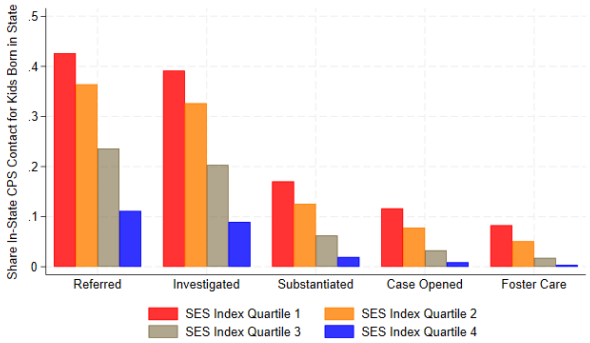
(a) By Estimated Income Decile



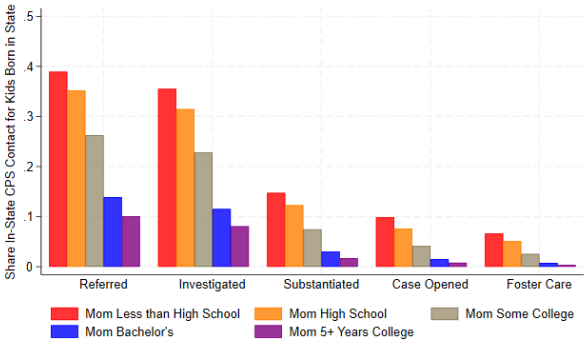
(b) By Estimated Income Quartile



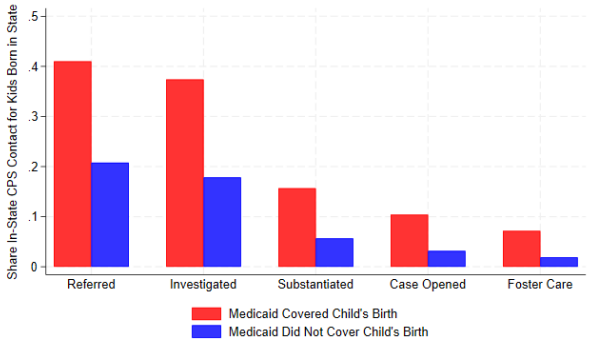
(c) By SES Index Quartile



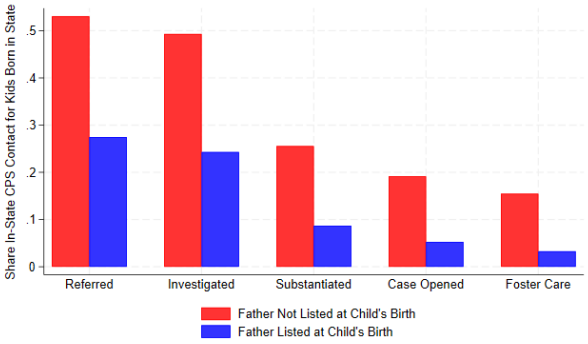
(d) Mother's Education at Birth



(e) Medicaid Coverage at Birth



(f) Father Listed on Birth Certificate



Notes: Each panel shows the cumulative prevalence of CPS contact by age 18 across five stages of involvement: referral, investigation, substantiation, case opening, and foster care placement. Each prevalence measure is disaggregated by a socioeconomic characteristic observed at birth. Panel (a) shows prevalence by estimated household income decile; Panel (b) shows prevalence by estimated household income quartile; Panel (c) shows prevalence by a standardized index of socioeconomic status; Panel (d) compares children by mother's level of education; Panel (e) shows differences by whether the child's birth was covered by Medicaid, which is a proxy for lower income; and Panel (f) disaggregates by whether a child's father was listed on the birth certificate. All estimates are based on linked California birth certificates and CPS records for the 1999-2001 birth cohorts.

education. Each panel in Figure 4 shows CPS contact rates by race within deciles of predicted income and SES index, and by mothers' education.

The differences across groups are substantial. Among children born to Black mothers in the bottom decile of the predicted income distribution, over 80 percent are referred to CPS by age 18, and more than 40 percent are substantiated victims of maltreatment. Rates are also extremely high for children born to White mothers in the same income decile: more than 70 percent are referred and over 30 percent are substantiated. These findings suggest that while Black children overall experience higher rates of CPS involvement, White children in the lowest-income households also face substantial risk of deep system contact.

We find similar—though slightly attenuated—patterns when using our SES index. Roughly 70 percent of children born to either Black or White mothers in the bottom SES index decile are referred to CPS, with about 30 percent experiencing a substantiated allegation. Notably, conditional on being in the bottom SES decile, racial disparities in CPS contact are relatively small.

The same pattern holds when looking at patterns by maternal education. Among children whose mothers did not complete high school, CPS involvement is again widespread. Over 70 percent of children born to Black mothers and more than 60 percent born to White mothers in this group are referred to CPS. Roughly 30 percent and 25 percent, respectively, are substantiated victims.

Taken together, these results highlight two key points. First, CPS involvement is extremely common among children growing up in the most socioeconomically disadvantaged households, regardless of race. Second, while racial disparities in overall CPS contact are large, those disparities are much smaller when comparing children within the same low-SES strata.

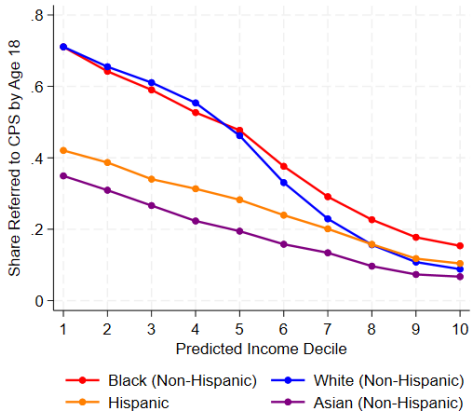
IVE Prevalence of Child Mortality

One potential concern with using CPS contact as a measure of child maltreatment is that system involvement may reflect differential surveillance or reporting biases rather than underlying risk. For example, children from disadvantaged backgrounds may come into contact with CPS not because maltreatment is more common, but because they are more likely to be observed by mandated reporters or reside in communities with higher CPS presence. While our analysis incorporates a gradient of CPS involvement—from referrals through foster care placements—to help differentiate between minor and more severe forms of contact, one may still be concerned about differential detection and discretionary system responses.

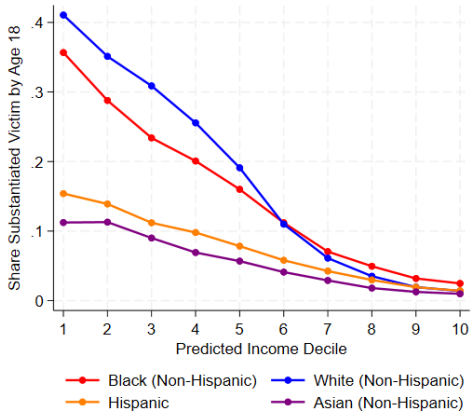
To further probe whether observed disparities in CPS contact reflect real differences in child

Figure 4: CPS Contact by Socioeconomic Status and Race/Ethnicity

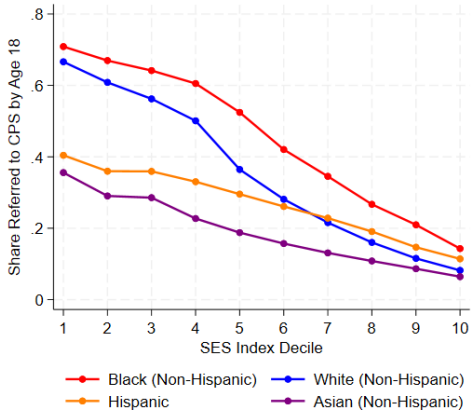
(a) Frac Referred to CPS, by Estimated Income



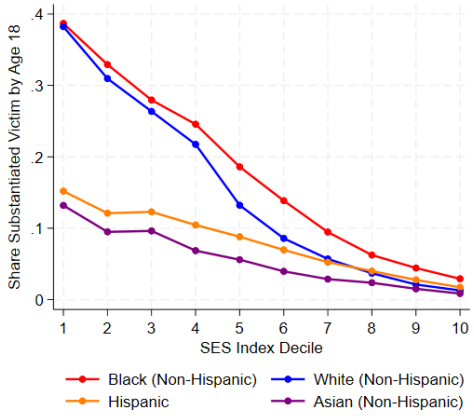
(b) Frac Subst'd Maltreatment, by Est Income



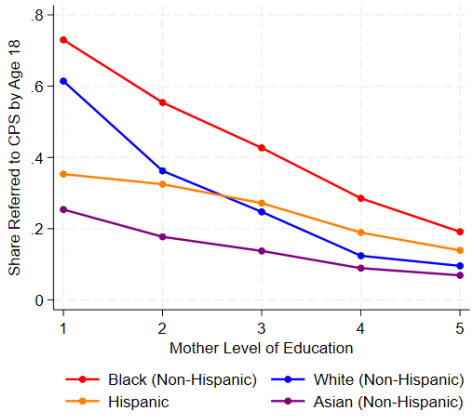
(c) Frac Referred to CPS, by SES Index



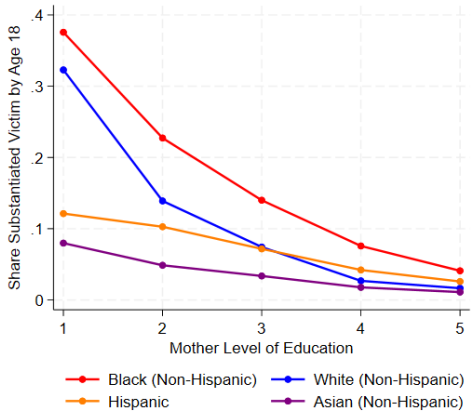
(d) Frac Subst'd Maltreatment, by SES Index



(e) Frac Referred to CPS, by Mom Education



(f) Frac Subst'd Maltreatment, by Mom Educ

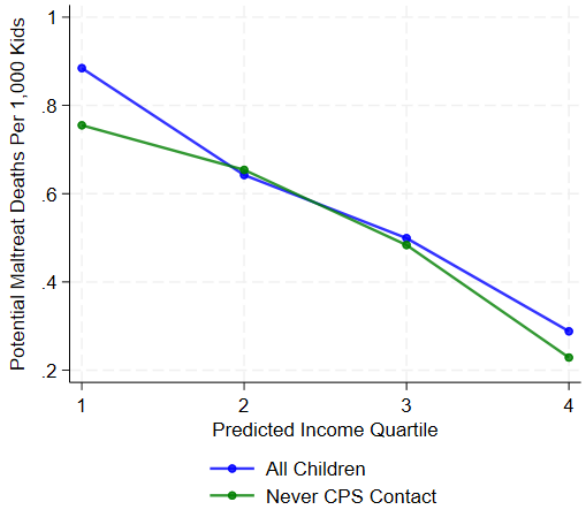


Notes: This figure shows the cumulative prevalence of CPS referrals (left panels) and substantiated maltreatment (right panels) by age 18, disaggregated by race and socioeconomic status (SES) for children born in California from 1999-2001. Each row reflects a different SES measure: predicted household income decile (top), SES index decile (middle), and maternal education level at birth (bottom). All SES measures are based on variables observable at birth. Maternal education is reported in five categories: (1) less than high school, (2) high school diploma or GED, (3) some college, (4) bachelor's degree, and (5) five or more years of postsecondary education.

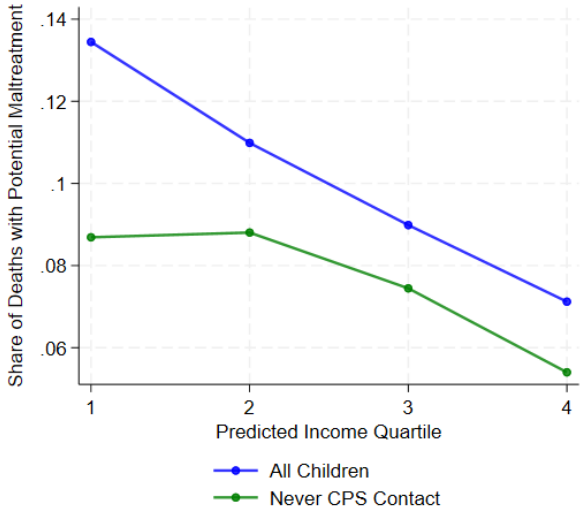
risk, we turn to a more objective and high-stakes outcome: child mortality. Our linked administrative data include death certificate records for children born in California from 1999-2001, which allow us to examine how early-life mortality—especially deaths plausibly related to maltreatment—varies across the income distribution.

Figure 5: Child Deaths from Possible Maltreatment Categories, by Socioeconomic Status

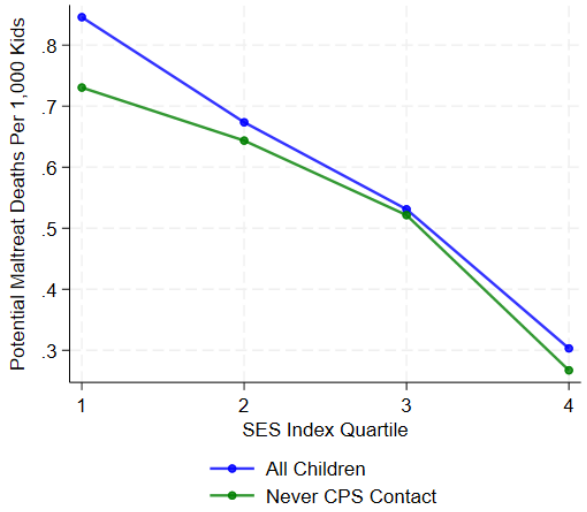
(a) Deaths per Capita, by Estimated Income



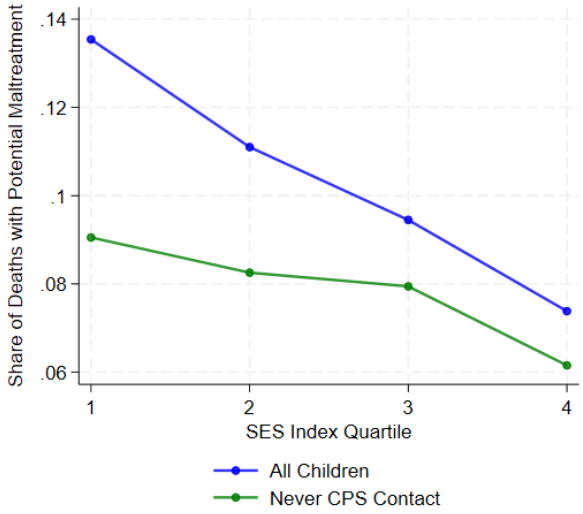
(b) Share of Deaths, by Estimated Income



(c) Deaths per Capita, by SES Index



(d) Share of Deaths, by SES Index



Notes: This figure shows the cumulative prevalence of child deaths before age 15 – a ground-truth measure of child risk – disaggregated by SES for children born in California from 1999-2001. The left column shows both deaths from potential-maltreatment-related categories (i) per 1,000 births of all children (in blue), (ii) and among never-CPS-involved children only (thus not influenced by CPS surveillance; in green), respectively. The right column shows the share of potential-maltreatment-related deaths by age 15 as a *fraction of all child deaths* for the given group. Each row reflects a different SES measure: predicted income decile (top) and SES index decile (bottom). All SES measures are based on variables observable at birth.

We begin by examining child mortality rates that are plausibly related to maltreatment, disaggregated by predicted income quartile. In Figure 5, Panel (a), we plot the number of maltreatment-related deaths per 1,000 children, by income quartile, for the full birth cohort. We identify deaths as maltreatment-related based on ICD codes and cause-of-death classifications commonly used in prior work.[5] The results reveal a stark socioeconomic gradient: among children in the lowest income quartile, the rate of maltreatment-related deaths exceeds 0.8 per 1,000 children. This rate declines steadily across the income distribution and falls to roughly 0.3 per 1,000 in the top quartile.

One might ask whether this mortality gradient is itself driven by CPS involvement—for example, among children placed in foster care. To address this, we replicate the analysis separately for children with no recorded CPS contact. As shown in the same panel, there is still a strong income gradient in maltreatment-related mortality when restricting to children who were never referred to CPS, suggesting that differences in mortality risk are not driven by system involvement.

Another concern is that children in lower-income households may simply face higher mortality risk overall, unrelated to maltreatment. To address this, Panel (b) of Figure 5 shows the share of all child deaths that are classified as plausibly maltreatment-related, again by income quartile. Among children in the lowest income quartile, roughly 16 percent of all deaths fall into this category. This share declines with income and falls below 8 percent in the highest quartile. Importantly, we observe a similar gradient even among children who were never referred to CPS, reinforcing the notion that the observed disparities reflect real underlying differences in risk—not just differential detection or surveillance.

Taken together, these findings offer additional evidence that early-life conditions associated with socioeconomic disadvantage are linked not only to higher rates of CPS involvement but also to elevated risk of severe harm. The persistence of the mortality gradient among children with no CPS contact suggests that the disparities we observe in administrative records reflect more than just differential system exposure—they track with meaningful differences in child well-being.

V Parents' CPS Involvement and Maltreatment

Having described the prevalence of child maltreatment and CPS involvement by socioeconomic status, we now turn to describing children's indirect exposure to maltreatment through their parents, and estimating the direct transmission of maltreatment from parent to child. As

[5] More detail on this coding and classification appears in Section BC.

mentioned above, while practitioners speak of intergenerational cycles of trauma passed from parents to children, this persistence could be confounded by surveillance or other omitted variables such as persistent poverty. In the following sections, we describe new statistics on children's indirect exposure to maltreatment through their parents' experiences, and the persistence of maltreatment across generations. We then provide suggestive estimates that maltreatment is causally transmitted.

VA Backward-looking Analysis

Starting with the full population of children born in California from 2022, we employ a backward-looking analysis to estimate the share of parents who were themselves CPS-involved. Table 1 presents descriptive statistics for the California backward-looking sample, which describes the percent of CPS-involved children (G2) who had a parent (G1) who was themselves CPS involved. The sample includes all children in California's 2022 birth cohort and their maltreatment history prior to turning age three. Column 1 presents means for children with no CPS referral prior to age three. Column 2 presents means for children with at least one CPS referral before age three. G1 CPS contact statistics are estimated using data available from 1998 onward.

From this table, we see that 68 percent of children with CPS contact had a parent (either a mother or a father) who also appears in the CPS records during the observable period (1998–2024), compared to 40 percent among children without CPS contact. The relative disparity is even greater for more intense system outcomes, such as investigation (63 percent vs. 35 percent) and removal (20 percent vs. 7 percent). These raw comparisons, however, understate the true extent of intergenerational CPS involvement because parental CPS histories are only partially observed. CPS records begin in 1998, and many G1 parents were born before this date or migrated into the state after early childhood, leading to unobserved CPS contact in their earlier years. As shown in Figure A.3, most linked G1 parents were born prior to the start of the administrative data, implying that backward-looking descriptive estimates miss a substantial share of parental CPS exposure. To address this limitation, we construct imputed estimates that account for incomplete observation of maternal CPS histories.

Figure 6 presents rates of mother's (G1) California CPS contact, split by child's (G2) history of CPS contact for children in the 2022 California birth cohort, across a variety of imputation assumptions. The first panel presents results identical to the raw data in Table 1: Among children born in 2022 who were referred to California CPS before age three, 58 percent had a matched birth certificate mother who was herself referred to California CPS (measured

Table 1: Summary Statistics for Backward-Looking Sample

	(1) Child No CPS Record	(2) Child Has CPS Record
Child CPS		
Child Ever Investigated	.	0.87
Child Ever Substant'd Victim	.	0.37
Child Ever Case Opened	.	0.28
Child Ever Removed	.	0.21
Bio Mom Ever Perpetrator	.	0.82
Only Unrelated/Unknown Perpetrator(s)	.	0.063
Either Parent CPS		
Either Parent Has CPS Record	0.40	0.68
Either Parent Ever Investigated	0.35	0.63
Either Parent Ever Substant'd Victim	0.14	0.33
Either Parent Ever Case Opened	0.079	0.21
Either Parent Ever Removed	0.066	0.20
Mother CPS		
Mom Has CPS Record	0.28	0.58
Mom Ever Investigated	0.24	0.52
Mom Ever Substant'd Victim	0.090	0.26
Mom Ever Case Opened	0.047	0.16
Mom Ever Removed	0.036	0.15
Father CPS (if Father Known)		
Dad Has CPS Record	0.23	0.43
Dad Ever Investigated	0.20	0.39
Dad Ever Substant'd Victim	0.069	0.15
Dad Ever Case Opened	0.037	0.092
Dad Ever Removed	0.034	0.092
Birth Characteristics		
Mother's Birth Year	1990.7	1993.0
Mom High School or Less	0.34	0.63
Low Birthweight	0.067	0.12
Labor Delivery Medicaid	0.38	0.75
Birth Order/Parity	2.27	3.49
Father Missing from Birth Cert	0.053	0.24
Father Age	33.6	31.6
N Child Observations	398,580	21,813

Notes: The table presents descriptive statistics for the backward-looking sample, which describes the percent of CPS-involved children (G2) who had a mother (G1) who was herself CPS involved. The sample includes all children in California's 2022 birth cohorts and their history prior to turning three years old. Column 1 presents means for children with no CPS referral prior to age 3. Column 2 presents means for children with at least one CPS referral before age 3. Mother and father CPS contact statistics are raw estimates using linked CPS data available from 1998 onward.

from January 1998 onward), whereas only 28 percent of never-referred children in the 2022 cohort have a mother matched to a CPS record. The second set of estimates imputes mother CPS contact based on mothers' potential omitted years of contact between their birth date and the earliest year with matched CPS records (1998), using the procedure described in Figure A.5. This procedure uses observed post-1997 data to estimate CPS contact before a certain age X, conditional on *not* having any CPS contact at or after age X.[6] The third set of estimates accounts for undercounting maternal CPS contact due to migration into California by additionally restricting the sample to mothers (G1) who themselves were born in California, which is identifiable from G2 child birth certificate records. Finally, the fourth set of estimates further restricts to children in the 2022 birth cohort who either had a substantiated allegation of maltreatment before age three, versus those who had no CPS referral.

Our estimates show an extremely high prevalence of family history of suspected maltreatment for CPS-involved children. The final two panels of Figure 6 show that approximately 71 percent of CPS-involved (referred) young children had a mother who herself was CPS-involved, and 74 percent of maltreated young children had a mother who was CPS involved. The latter is nearly 30 percentage points higher than the rate of maternal involvement for non-CPS-contacted young children.[7]

Figure 7 presents a complementary perspective on intergenerational CPS involvement by distinguishing parental CPS histories by parent type. Restricting attention to parents who were themselves born in California, the figure reports the share of children with CPS contact who had either parent—mother or father—with a recorded history of CPS involvement during childhood. This approach reduces undercounting due to parental migration. Exposure to parents with prior CPS contact is substantial. Among children with any CPS contact before age three, approximately 80 percent have at least one parent with a recorded childhood CPS history, compared to roughly 60 percent among children with no CPS contact.[8]

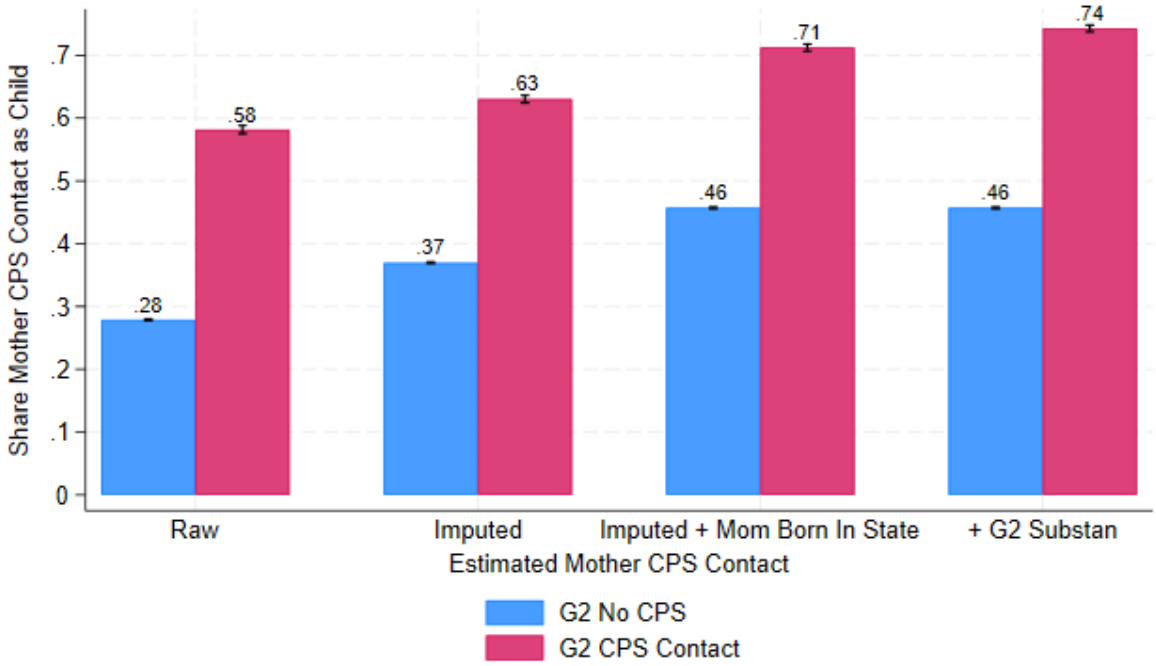
Panel A of Figure A.9 examines heterogeneity in backward-looking intergenerational CPS involvement across the income distribution. The figure plots, by predicted parental income

[6] To address the undercounting of maltreatment caused by CPS records only beginning in 1998, we impute missing history for mothers whose early childhoods are unobserved due to data left-censoring. We calculate the probability of unobserved early-life CPS contact using younger, fully observed birth cohorts to estimate the likelihood of contact before a specific age, conditional on having no later contact. These probabilities are then applied to mothers in earlier birth cohorts, correcting for censored maltreatment that occurred prior to the start of the dataset. We describe the procedure further in Figure A.5

[7] Figure A.6 provides more granularity by type of contact between children, showing that nearly 80 percent of young children with two or more substantiations, or those placed in foster care, had a mother who herself had CPS contact.

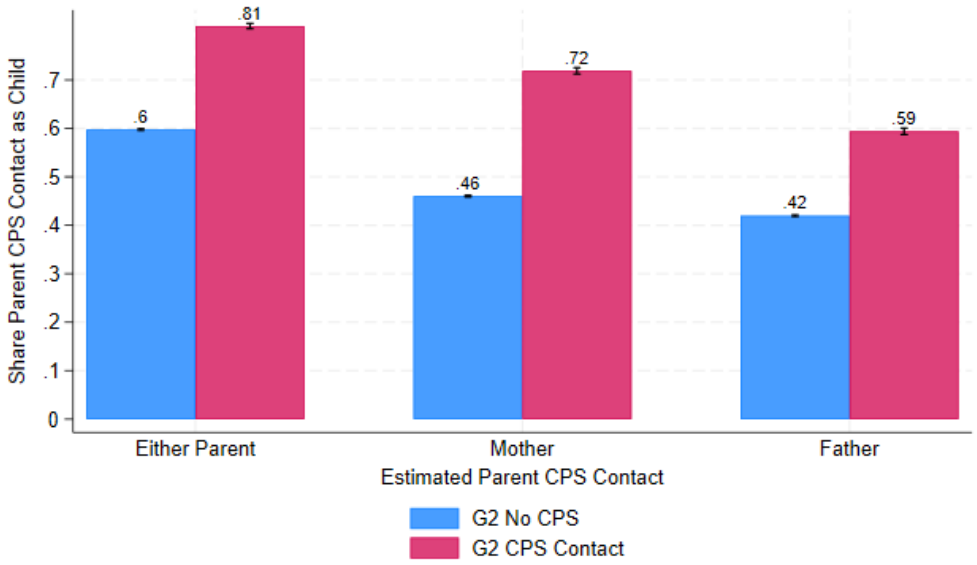
[8] Figure A.8 presents estimates of comparable magnitude using raw data, which understate parent exposure due to missing CPS contact prior to 1998, but require no imputation.

Figure 6: Backward-Looking: Estimated Mother's (G1) CPS Contact by Child's (G2) CPS Contact



Notes: This figure presents rates of mother's (G1) California CPS contact, split by child's (G2) history of CPS contact for children in the 2022 California birth cohort. For example, the first set of estimates show that among children born in 2022 who were referred to California CPS before age three, 58 percent had a matched birth certificate mother who was herself referred to California CPS (measured from January 1998 through December 2024), whereas only 28 percent of never-referred children in the 2022 cohort have a mother matched to a CPS record. The second set of estimates imputes mother CPS contact based on mothers' potential omitted years of contact between their birth date and the earliest year of available CPS records (1998), using the procedure described in Figure A.5. The third set of estimates accounts for undercounting mother's CPS contact due to mothers potentially growing up outside of California by additionally restricting the sample to mothers (G1) who themselves were born in California, which is identifiable from G2 child birth records. Finally, the fourth set of estimates further restrict to children in the 2022 birth cohort who had a substantiated allegation of maltreatment before age three, versus those who had no CPS referral.

Figure 7: Backward-Looking: G1 Parent Inferred CPS Contact by G2 Contact



Notes: This figure shows imputes rates of CPS contact for parents (G1), split by child (G2) CPS contact. Panel 1 shows the percentage of children (G2) in the 2022 California birth cohort who have at least one parent listed on the birth record who was referred to CPS. Panel 2 shows rates of contact for mothers, and Panel 3 shows rates of contact for fathers (second listed parent on birth record). The sample is restricted to children born in 2022 to California-born mothers. Parental CPS histories are measured using raw administrative records with imputation for potential contact prior to 1998 (See Figure A.5). Father estimates are likely understated due to unlisted fathers on some birth records (24 percent unlisted for CPS-involved children and 5 percent for non-CPS-involved children).

decile, the share of children who have at least one parent with a recorded history of CPS contact during childhood, separately for children with and without CPS contact themselves. The sample is restricted to children born in 2022 to California-born mothers.[9]

Two patterns stand out. First, parental CPS involvement is common across the income distribution but is substantially more prevalent at the bottom. Among children born into the lowest income decile, nearly 80 percent of children with no CPS contact nonetheless have a parent with a CPS history, while the corresponding share exceeds 85 percent for children who themselves experience CPS contact. Second, the income gradient is downward sloping for both groups, and at nearly every point in the income distribution, children with CPS contact are more likely to have a parent with prior CPS involvement than children without CPS contact. This gap narrows with income and largely disappears only at the very top of the income distribution, where the two series converge.[10] Together, these patterns indicate that intergenerational CPS involvement is most pronounced among economically disadvantaged families but that disparities in intergenerational involvement persist across the income distribution.

Altogether, the backward-looking estimates highlight (i) the high prevalence of parent CPS history for CPS-involved children, and (ii) the disparities in family history for CPS vs. non-CPS involved children. In Section VIA, we use a forward-looking analysis to provide suggestive evidence that maltreatment may be causally transmitted from parent to child across generations.

VI Intergenerational Transmission of Maltreatment

The previous section showed that a majority of children have indirect exposure to maltreatment through their parents, and that parents' exposure is greater for children who themselves are CPS-involved and confirmed maltreatment victims. In this section, we provide

[9] Panels C and D of Figure A.9 repeat the analysis for the same sample where CPS histories are calculated using our imputation.

[10] Panel B replicates the analysis, but focuses on a more intensive measure of parental exposure: whether either parent was ever a confirmed victim of maltreatment. The figure plots, by predicted parental income decile, the share of children with at least one parent who had a substantiated CPS allegation during childhood, separately for children with and without CPS contact themselves. The same qualitative patterns emerge, though at lower levels. Among children born into the lowest income decile, approximately 35 percent of children with no CPS contact have a parent who was a confirmed victim of maltreatment, compared to roughly 48 percent among children who themselves experience CPS contact. The prevalence of parental substantiation declines sharply with income for both groups, producing a steep downward-sloping gradient across the income distribution. At each income decile, children with CPS contact are more likely to have a parent with a substantiated maltreatment history than children without CPS contact. Although the two series converge at the top of the income distribution, a gap remains even among the highest-income families.

evidence consistent with a causal transmission of maltreatment across generations. We also provide evidence on mechanisms through which maltreatment may be transmitted.

VIA Forward-looking Analysis

To examine the effects of child maltreatment on the maltreatment of one's future children, we focus on a sample of girls born in California between 1987-2001. We measure CPS involvement for girls in the California birth cohorts from ages 0 to 14. We refer to these girls as generation 1 ("G1") for the remaining analyses, and their children identified on birth certificate records as generation 2 ("G2"). We restrict the benchmark samples to girls because maternal linkages are more reliably observed than paternal ones, as fathers are substantially less likely to appear on birth certificate records.[11]

Our sample has some limitations. In California, birth records are only available for the 1999 birth cohort onward, implying that G1-to-G2 birth-to-birth linkages can only be measured until age 24. To alleviate this constraint, we instead use a base sample of *all girls referred to CPS* in the 1987-2001 birth cohorts and compare the *intensive* margin of their CPS contact, namely whether a girl was a substantiated (confirmed) victim of maltreatment. While the estimates do not include children who were never referred, the comparison is relevant because (i) conditioning on any CPS contact attenuates concerns about differential surveillance and (ii) the share of ever-CPS-contacted children still comprise a large share of the child population (outlined in Section IV). Finally, the CPS data begin in 1998, which limits the range of tracking to, for instance, ages 4 and up for a girl born in 1994. While this means we do not observe CPS contact from birth for all girls, we deliberately cap the observation window at age 14 to maintain a consistent upper bound of maltreatment experienced prior to the potential birth of any children.[12]

VIA.1 Empirical Strategy

To estimate how maltreatment experienced during childhood affects maltreatment for one's own future children, we compare G1 girls with different levels of CPS involvement, linked to their own children's (G2) CPS records. Naturally, comparing children with different

[11] In Table A.3, we verify the robustness of our maltreatment transmission strategy for both girls and boys by examining whether a maltreated G1 individual is differentially likely to become a future perpetrator of maltreatment via being listed on a later CPS report, rather than examining whether their child appears as a victim. An advantage of this supplemental strategy is that it does not depend on fertility or other match rates to birth records.

[12] We include birth-year-by-county fixed effects in all specifications to account for differential windows of observation. A negligible fraction of the sample gave birth before age 15. To the extent that this happens, we define G2 births as only those births occurring ages 15 and up.

levels of CPS involvement raises concerns about selection bias: G1 children exposed to more severe maltreatment may differ in both observed and unobserved ways from those with less intensive (or without) CPS involvement, and such differences may affect the likelihood of maltreatment for G2 children. To address this, we employ two complementary empirical approaches. The first approach compares girls who were all reported to CPS but differed in whether a maltreatment allegation was confirmed (substantiated), conditioning on a rich set of observable characteristics such as zip code, family size, and reporter type. The second strategy leverages sibling comparisons, where all siblings were reported, but vary along substantiation decisions. For each approach, we assess the extensive and intensive margins of contact, and include a battery of robustness checks. Finally, we introduce three suggestive, supplemental approaches: (i) a test for transmission of maltreatment type (physical abuse, sexual abuse, or neglect), (ii) isolating idiosyncratic non-household perpetrator incidents where one sibling is maltreated, and (iii) isolating years of exposure to a maltreatment-prone home environment by using age variation among siblings who are removed to foster care.

Selection-on-Observables Our first approach leverages our linked administrative records to control for observable differences across children. While this strategy does not guarantee conditional independence, it allows us to adjust for substantial observed heterogeneity and isolate variation in maltreatment among otherwise similar children who have already come to the attention of CPS.

We consider two key measures: (1) whether a child was ever substantiated by CPS as a maltreatment victim through age 14, and (2) a child's total number of substantiations through age 14. Within this sample, differences in substantiation reflect variation in the assessed credibility of maltreatment. Moreover, because child maltreatment is often chronic, repeated substantiations plausibly capture persistent agency concern and sustained exposure to potentially unsafe environments.

We estimate the following models:

$$Y_i = \beta_1 NumSubstan_i + \theta_1 X_i + \varepsilon_i \tag{1}$$

$$Y_i = \beta_2 AnySubstan_i + \theta_2 X_i + \varepsilon_i \tag{2}$$

where i indexes a G1 girl and Y_i is a G2-related outcome, such as the number of times G2 children are substantiated maltreatment victims. $AnySubstan_i$ is an indicator for whether the G1 girl ever experienced a substantiated CPS investigation between ages 0 and 14, and $NumSubstan_i$ is the total number of observed substantiations during this period.

The covariate vector X_i includes G1 race and ethnicity, age at first observed CPS investigation, ZIP code of residence at first referral, birth-year-by-county fixed effects, and the reporter type for the child's first referral (e.g., medical, school, family friend, law enforcement, etc.). The standard control set also includes child birth-year-by-month fixed effects, number of children listed on a child's first referral, and ages of the youngest and oldest children listed.

G2-related outcomes include indicators for counts of CPS reports, investigations, substantiations, case openings, and removals to foster care. For all estimates, the intensive-margin outcome "number" captures the total number of events across all of G1's children. To illustrate, suppose a G1 mother has two G2 children, Anna and Ben. If Anna is substantiated on one report and Ben is not, and both children are substantiated on a subsequent report, the total number of G2 substantiations assigned to the G1 mother is three (two for Anna, one for Ben). We primarily focus on the intensive margin because the chronicity of CPS involvement and substantiations is considered by practitioners to be a strong signal of maltreatment severity.

Sibling Fixed Effects Design Our second empirical approach uses within-family comparisons to isolate the association between maltreatment severity and later outcomes. Specifically, we compare sisters who were both reported to CPS but experienced different investigation outcomes: For example, instances where one sister was substantiated while another was not, or where sisters appeared on the same report together but experienced different numbers of lifetime substantiations. The sibling fixed effects design allows us to account for shared household, parental, and neighborhood characteristics that are common across siblings, allowing identification to come from differential maltreatment exposure within the same family.

The sibling fixed effects specification is given by:

$$Y_{is} = \beta_1 NumSubstan_{is} + \theta_1 X_{is} + \lambda_{r(s)} + \varepsilon_{is} \tag{3}$$

$$Y_{is} = \beta_2 AnySubstan_{is} + \theta_2 X_{is} + \lambda_{r(s)} + \varepsilon_{is} \tag{4}$$

where i indexes a G1 girl and s indexes her sibling group. The term λ_s denotes a sibling-group, with siblings identified by a shared referral $r(s)$ fixed effect that absorbs all characteristics shared across sisters, including family background and neighborhood conditions. The vector X_{is} includes only covariates that vary within sibling groups, such as child's birth year and their age at the time of the referral. Equations 3 and 4 are estimated on a restricted sample of families with at least two sisters who were investigated by CPS. Identification is driven by differences in substantiation outcomes across sisters within the same family. Standard errors

are clustered at the CPS referral level.

The sample for the sibling fixed effect design comprises: (i) referrals with two or more siblings who each satisfy the standard sample criteria (see beginning of Section VIA); (ii) the referral has at least one alleged parent perpetrator—biological or step-parent—and no alleged sibling perpetrators, and (iii) all children listed on the report have the same zip code of residence.[13] A final requirement for the sibling fixed effects design is (iv) for a given referral, variation among the siblings in terms of either (a) total observed substantiations from ages 0-14 (intensive margin comparison) or (b) any substantiation from ages 0-14 (extensive margin comparison). For example, suppose that Anna and Brianna are siblings and have no other siblings or co-resident children. Suppose Anna is a substantiated victim on one occasion by age 14 and Brianna is a substantiated victim on two occasions. Then, referrals with Anna and Brianna would not influence the extensive margin sibling fixed effects estimates (because both have at least one lifetime substantiation), but they *would* influence the intensive margin estimates (because they have different numbers of lifetime substantiations). To accommodate flexible sibling groups, a child can appear on multiple qualifying referrals over time. Every child is weighted equally: specifications weigh child observations by the inverse of the number of times a child appears in the sample.[14]

While the sibling fixed effects design substantially reduces bias from unobserved differences across households, our estimates may still be biased due to unobserved factors that vary across children within the same family. The direction of this bias is ambiguous ex ante. For instance, if a sibling without a substantiated investigation was nevertheless maltreated but went undetected, the referral fixed effects estimates likely understate the true effects of maltreatment. Conversely, if substantiation reflects child-specific traits that are themselves predictive of later outcomes, such as behavioral challenges or differential

[13]This restriction, together with a biological parent perpetrator, is an attempt to ascertain true co-resident siblings. The CPS-involved population includes children in blended families and children who move between homes. We validate our sibling identification approach as follows: We measure the share of 1999-2004 inferred sibling pairs who have the same birth mother on matched birth records from 1999-onward. We confirm that the vast majority of sibling pairs inferred using this procedure share the same biological mother (98 percent of all child-referrals and 92 percent of child-referrals with variation in child lifetime substantiations).

[14]A strength of the primary sibling fixed effects design is that it flexibly handles adjustments to sibling groups (pairwise referral fixed effects) while controlling for common background, and estimates remain interpretable (the association of outcomes with one additional substantiation, or ever being substantiated). However, it is true that exposure to maltreatment on other referrals could still be endogenous to child characteristics. To partially address this concern, in Appendix Table A.1 we estimate a more restrictive sibling fixed effects design that only uses referrals where at least one sibling is a victim and one sibling is not on the *given* referral. This design's advantage is that it is easy to explain—it uses variation on a given referral at a given point in time. However, it is more difficult to interpret: when one child is substantiated on a given referral, this might impact the likelihood that they are substantiated on other referrals. Thus, interpreting the coefficient of the restrictive design is more challenging. Still, the results are qualitatively similar.

parental perceptions, the resulting estimates could be biased in either direction. Moreover, differences in maltreatment exposure may coincide with differences in parental investments across children, complicating interpretation. While a sibling fixed effects design may not fully eliminate concerns about unobserved confounders, it provides valuable complementary evidence of the impact of substantiated maltreatment (and a corresponding bundle of inputs) and helps strengthen the case that maltreatment is associated with meaningful differences in intergenerational outcomes.

VIA.2 Results

Table 2 reports descriptive statistics for the samples used in the forward-looking analysis. Column (1) summarizes characteristics of the "standard sample", used for the first empirical strategy, which includes all G1 girls with at least one CPS referral in childhood. The average G1 child in this sample was born in the mid-1990s, and the sample is racially and ethnically diverse, with Hispanic girls comprising the largest race/ethnicity share, followed by non-Hispanic White and Black girls. Within this sample, an individual experienced on average 2.7 referrals during childhood. Approximately 31 percent experienced substantiated maltreatment, with an average of 0.47 substantiations during childhood. By early adulthood, 39 percent of G1 girls had given birth, with an average of 0.87 children and a mean age at first birth of 21. Socioeconomic conditions at the time of birth indicate substantial disadvantage: average predicted household income is approximately $46,000, two-thirds of births are Medicaid-covered, and college attendance is relatively low.

Column (2) summarizes the sibling sample. As in the standard sample, cohorts are concentrated in the mid-1990s and the sample is predominantly Hispanic, with substantial representation of non-Hispanic White and Black girls. On average, a G1 girl in the sibling sample experienced 3.55 referrals in childhood. 40 percent experienced at least one substantiated investigation, with an average of 0.67 substantiations. Just under half of girls in the sibling sample gave birth during the observation period, with an average of 1.01 children and a mean age at first birth of 20.9. Measures of socioeconomic status at birth indicate high levels of disadvantage, including low predicted household income, high Medicaid coverage at delivery, and low parental educational attainment.

Conditional on fertility, children (G2) of mothers in both samples experience substantial CPS involvement. In the standard sample, 16 percent of mothers have at least one child with a CPS referral and 6.7 percent have at least one child with a substantiation; on average, mothers have 1.04 total G2 referrals and 0.13 total G2 substantiations. In the sibling sample, these rates are higher: 19 percent have a child with a referral, and 8 percent have a child with a

Table 2: Summary Statistics: Forward-Looking Analysis Sample

	(1) Standard Sample	(2) Sibling Sample
G1 Demographics		
Female	1	1
G1 Birth Year	1994.9	1994.6
Black (Non-Hispanic)	0.12	0.13
White (Non-Hispanic)	0.32	0.31
Hispanic	0.45	0.50
Asian (Non-Hispanic)	0.043	0.042
G1 Childhood CPS Contact		
G1 Any Referral as a Child	1	1
G1 Total Referrals as a Child	2.68	3.55
G1 Any Substantiation as a Child	0.31	0.40
G1 Total Substantiations as a Child	0.47	0.67
Fertility and G1 SES at Birth		
Gave Birth	0.39	0.45
Number of Children	0.87	1.01
Age at First Birth	21.0	20.9
Predicted Household Income	46296.8	44690.8
Birth Mom Any College	0.30	0.28
Birth Dad Any College	0.27	0.26
No Partner Listed	0.16	0.16
Medicaid Covered Delivery	0.66	0.69
Hospital Mean Household Income	56161.3	55313.6
Child (G2) Outcomes, Conditional on Fertility		
Any G2 Child with a Referral	0.16	0.19
Total G2 Child-Referrals	1.04	1.27
Any G2 Child with a Substantiation	0.067	0.080
Total G2 Child-Substantiations	0.13	0.22
Number of G1 Observations	1,177,273	926,192
Number of Unique G1 Children	1,177,273	413,422

Notes: This table reports summary statistics for two California samples used in the forward-looking analyses. The first (*standard*) sample includes all girls (G1) born between 1987 and 2001 who had any CPS contact between ages 0 and 14. The second (*sibling*) sample consists of all referrals from girls in the first sample that also include at least one other girl from the first sample (a sibling pair), and where there is variation among siblings in terms of observed cumulative substantiations through age 14. Referrals in the sibling sample are restricted to those listing at least one alleged parent perpetrator (biological or otherwise) and no alleged sibling perpetrators.

substantiation, with averages of 1.27 referrals and 0.22 substantiations per mother.

Table 3 presents descriptive evidence on forward-looking intergenerational CPS involvement from G1 to G2. Panel A reports extensive-margin outcomes—whether any G2 child is ever referred or substantiated—shown both for all G1 girls and conditional on fertility. Among all G1 girls, 20 percent of those who experienced substantiated maltreatment in childhood have a G2 child with a CPS referral, and 9.2 percent have a G2 child with a substantiation, compared to 14 percent and 5.6 percent, respectively, among G1 girls who were never substantiated. Conditioning on fertility increases these probabilities mechanically: among G1 mothers, 46 percent of those with a childhood substantiation have a child who is referred to CPS, and 21 percent have a child who is substantiated, compared to 39 percent and 15 percent among mothers who were never substantiated.

Panel B reports intensive-margin outcomes, capturing the number of G2 referrals and substantiations. Among all G1 girls, those with childhood substantiation have an average of 1.41 G2 referrals and 0.25 G2 substantiations, compared to 0.87 referrals and 0.14 substantiations among those never substantiated. Conditional on fertility, the corresponding averages are 3.22 referrals and 0.57 substantiations for previously substantiated G1 mothers, compared to 2.31 referrals and 0.38 substantiations for those never substantiated. Overall, confirmed maltreated G1 girls are 64 percent more likely to have a maltreated G2 child (40 percent conditional on fertility) and have on average 79 percent more child substantiation incidents (50 percent conditional on fertility). Taken together, these descriptive patterns demonstrate persistence on both the extensive and intensive margins and motivate the analysis that follows, which estimates these conditional relationships while accounting for observed characteristics and within-family variation.

Table 3: Conditional Probabilities of Intergenerational CPS Involvement

Panel A: Extensive Margin

	All		Conditional on Fertility	
	G2 Ref	G2 Substan	G2 Ref	G2 Substan
G1 Childhood Substantiation	0.20	0.092	0.46	0.21
G1 Never Substantiated	0.14	0.056	0.39	0.15

Panel B: Intensive Margin

	All		Conditional on Fertility	
	N G2 Refs	N G2 Substan	N G2 Refs	N G2 Substan
G1 Childhood Substantiation	1.41	0.25	3.22	0.57
G1 Never Substantiated	0.87	0.14	2.31	0.38

Notes: This table reports conditional probabilities of CPS contact and substantiated maltreatment among children (G2) in California, by their mothers' (G1) own childhood CPS histories. Panel A shows the extensive margin of G2 outcomes (a binary measure of any child referral or substantiation, respectively), and Panel B shows a continuous measure of G2 outcomes (number of child-referrals or child-substantiations, respectively).

Forward-looking transmission Table 4 reports estimates of the effect of G1's childhood substantiated maltreatment on the number of substantiated maltreatment incidents experienced by her G2 children. We focus our forward-looking analysis on the intensive margin—specifically, the number of G2 substantiations—because this outcome most directly captures repeated or sustained exposure to confirmed maltreatment. We interpret this margin as focusing on variation in agency-verified maltreatment that is plausibly associated with sustained, chronic maltreatment risks. Panel A uses the number of G1 childhood substantiations as the key independent variable, while Panel B uses an indicator for any G1 childhood substantiation. Columns 1-4 present estimates in the standard sample under increasingly restrictive specifications: Column 1 is unconditional, Column 2 adds our set of standard controls (including county-by-birth-year fixed effects), Column 3 additionally controls for characteristics measured at the birth of G1's first child (G2), such as G1's educational attainment and father/second-parent presence, which amounts to conservatively overcontrolling for any direct effects of maltreatment on G1's adult socioeconomic status. Column 4 excludes all G2 CPS contact events for which a non-parent is listed as the alleged perpetrator in order to mitigate any mechanical multigenerational links, such as a grandparent (G0) who maltreats both their child (G1) and grandchild (G2). Columns 5–6 report corresponding estimates in the sibling fixed effects sample, with Column 6 adding referral fixed effects to account for across-family variation.

Table 4: Intergenerational Transmission: Child Substantiated Incidents

Outcome: G2 Child Substantiations

	Standard Sample				Sibling FE Sample	
	(1) No Controls	(2) Standard Controls	(3) G2 Birth Controls	(4) No Multi-Gen	(5) No Controls	(6) Standard Controls
Panel A: Number of G1 Childhood Substantiations						
Total G1 Childhood Substantiations	0.139***	0.108***	0.077***	0.066***	0.133***	0.104***
	(0.003)	(0.003)	(0.003)	(0.002)	(0.004)	(0.012)
Panel B: Any G1 Childhood Substantiation						
G1 Any Substantiation as a Child	0.186***	0.131***	0.089***	0.078***	0.194***	0.128***
	(0.005)	(0.005)	(0.005)	(0.004)	(0.007)	(0.021)
County x Birth Cohort		×	×	×		×
Standard Controls		×	×	×		
Referral Fixed Effects +						×
G2 Birth Controls			×			
Mean of Dependent Var	0.444	0.444	0.448	0.287	0.539	0.539
Observations	464,177	464,177	441,058	442,906	284,723	284,723

Notes: This table presents estimates of the intergenerational transmission of child maltreatment from mother (G1) to child (G2). The outcome is the number of substantiated maltreatment incidents for G2 children of the G1 mother. All specifications condition on having at least one child. Columns 1–4 report results from the selection-on-observables specification (without referral fixed effects), while Columns 5–6 report estimates from the sibling design, where Column 6 includes referral-level fixed effects. Panel A uses the mother's *number* of childhood substantiations as the key independent variable, and Panel B compares future mothers with no substantiations to those with at least one. Column 1 presents unconditional estimates from equations 1 and 2. Column 2 adds the full set of standard controls described in the text. Column 3 further adds birth characteristics of G1's first child (G2)—which are technically endogenous to G1's childhood history and could attenuate estimates toward zero—to demonstrate robustness to over-controlling. Column 4 excludes any G2 CPS contact in which a non-parent is listed as the alleged perpetrator, mitigating mechanical links arising from potential grandparent (G0) maltreatment of both G1 and G2. Columns 5 and 6 report results from the sibling fixed effects sample. Column 5 reports unconditional estimates. Column 6 includes referral fixed effects and additionally includes G2 birth-year and G1 age-at-referral fixed effects. Robust standard errors are reported in parentheses for Columns 1–4, and standard errors are clustered at the referral level for Columns 5–6. Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

Across specifications, the estimates indicate substantial persistence from G1 substantiated maltreatment to G2 substantiated maltreatment. In the standard-sample specification (Column 2), an additional G1 childhood substantiation is associated with 0.108 (24 percent) additional G2 substantiations (Panel A), and having any substantiated maltreatment as

a child is associated with 0.131 (30 percent) additional G2 substantiations (Panel B), relative to a baseline mean of 0.444. These magnitudes remain relatively precise and economically meaningful when adding potentially "over-controlling" G2 birth characteristics (Column 3: 0.077 per substantiation; 0.089 for any substantiation), and when excluding cases with non-parent alleged perpetrators (Column 4: 0.066 per substantiation; 0.078 for any substantiation), suggesting the results are not driven purely by observable birth conditions or mechanical multigenerational reporting links (e.g., maltreatment of both G1 and G2 involving a grandparent, G0).

In the referral fixed effects sample, the estimates are similarly large. Column 6 implies that each additional G1 childhood substantiation is associated with 0.104 (19 percent) additional G2 substantiations, and that having any G1 substantiation is associated with 0.128 (24 percent) additional G2 substantiations (Panels A and B, respectively), relative to the sibling-sample mean of 0.539 for mothers with children. Taken together, the pattern across both the selection-on-observables and within-family specifications suggests that intergenerational persistence is substantial.

Table 5 extends the sibling fixed effects design to a set of additional outcomes.[15] Forward-looking measures of G2 maltreatment mechanically depend on whether G1 has children; childhood maltreatment could influence intergenerational transmission both through parenting behavior and through selection into having children. If G1 childhood substantiation affects the timing or likelihood of having children, then part of the observed intergenerational persistence in CPS involvement may operate through differential fertility. Column (1) shows that G1 substantiated maltreatment is associated with a higher likelihood of fertility. An additional childhood substantiation is associated with a 2.6 percentage point increase in the probability of any birth (Panel A), while having any substantiation is associated with a 3.7 percentage point increase (Panel B), relative to a mean fertility rate of 0.467. These estimates are relatively modest and therefore unlikely to explain the magnitude of the intergenerational persistence in CPS involvement that we document. Our results are also robust to alternative specifications that are less sensitive to fertility.[16]

[15] In Table A.2, we additionally extend the selection-on-observables design to the same set of outcomes.

[16] We probe the sensitivity of the intergenerational transmission estimates to conditioning on fertility in the following ways. First, we examine an alternative measure that does not depend on a biological parent-child link: whether G1 was ever listed as an alleged perpetrator. This alternative measure captures intergenerational involvement with the child welfare system through alleged perpetration rather than solely through substantiation of one's own child. Table A.3 shows that childhood substantiated maltreatment predicts later involvement with CPS as an alleged or confirmed perpetrator, regardless of fertility. For girls, the effects are sizable and precisely estimated, particularly for guardian perpetration, while estimates for non-guardian roles are smaller but remain positive. For boys, the associations are more modest but generally positive, indicating that persistence in CPS involvement extends beyond biological parent–child links and reflects broader patterns of later-life system contact. Furthermore, Table A.4 shows that the findings are robust

Columns (2)-(6) examine alternative measures of G2 child welfare involvement, all estimated within the referral fixed effects design. Across outcomes, G1 childhood substantiated maltreatment is positively and precisely associated with more intensive CPS contact among the next generation. Each additional G1 childhood substantiation is associated with 0.445 (12 percent) additional G2 referrals, 0.342 (13 percent) investigations, 0.102 (16 percent) substantiations, 0.065 (18 percent) case openings, and 0.052 (20 percent) foster care removals (Panel A). Along the extensive margin (Panel B), mothers with any childhood substantiation experience 0.61 (16 percent) additional G2 referrals, 0.46 (18 percent) investigations, 0.12 (19 percent) substantiations, 0.072 (20 percent) case openings, and 0.056 (22 percent) foster placements relative to those with none.

Table 5: Impacts of G1 Maltreatment on G1 Fertility and Measures of G2 Maltreatment
Sibling Fixed Effects Design

	(1) Any Birth	(2) Tot G2 Referral	(3) Tot G2 Investigation	(4) Tot G2 Substan	(5) Tot G2 Case	(6) Tot G2 Foster
Panel A: Number of G1 Childhood Substantiations						
Total G1 Childhood Substantiations	0.026***	0.445***	0.342***	0.102***	0.065***	0.052***
	(0.002)	(0.052)	(0.037)	(0.012)	(0.008)	(0.007)
Panel B: Any G1 Childhood Substantiation						
G1 Any Substantiation as a Child	0.037***	0.61***	0.46***	0.12***	0.072***	0.056***
	(0.0032)	(0.094)	(0.065)	(0.022)	(0.014)	(0.011)
Referral Fixed Effects +	×	×	×	×	×	×
Mean of Dep Var	0.467	3.737	2.609	0.635	0.360	0.261
N Observations	926,192	284,723	284,723	284,723	284,723	284,723

Notes: This table extends the referral fixed effects specification from Column 6 of Table 4. Each column presents results from separate regressions in which the dependent variable corresponds to a different outcome. Panel A uses the number of G1 childhood substantiations as the explanatory variable, while Panel B compares future G1 mothers with at least one substantiation to those with none. Column 1 reports the effect of G1 maltreatment on the likelihood of ever having at least one child. Columns 2–6 condition on having at least one child and report effects on various measures of G2 child welfare involvement: (2) number of G2 child CPS referrals, (3) investigations, (4) substantiations, (5) case openings, and (6) removals to foster care. Standard errors are clustered at the referral level. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The pattern of results across progressively more severe stages of CPS involvement suggests that intergenerational persistence is not limited to greater reporting or surveillance. Instead, the

to controlling for the intensive margin of fertility and cumulative G2-child-years of exposure to potential CPS contact.

associations extend to opening an ongoing case and removals to foster care—decisions that require higher evidentiary thresholds and reflect more serious system involvement. Taken together, we interpret these findings as indicating that childhood maltreatment is associated with elevated downstream risk at multiple points in the child welfare process, adding to the broader evidence of forward-looking transmission of maltreatment beyond differences in reporting.

Additional Robustness Exercises. We provide two additional robustness exercises for the intergenerational transmission of maltreatment. The first exercise isolates within-sibling variation in maltreatment that comes from an unrelated perpetrator from outside the home. The second exercise uses a design that leverages years of exposure to a maltreatment-prone home environment by comparing the ages of siblings who are removed from their home due to abuse or neglect, are placed in foster care, and are never subsequently substantiated victims of maltreatment again. The idea behind this design is to consider, for example, two siblings of different ages. The older sibling is exposed to maltreatment for a longer period of time than the younger sibling, until the situation is resolved by placing both children in foster care. We hypothesize that the length of time spent in an abusive or neglectful home may be associated with the likelihood of maltreating one's own children later in life. Although imprecise, the estimates from these exercises are consistent with the causal transmission of child maltreatment across generations.

Table A.5 presents estimates the association of maltreatment from a perpetrator outside the home with a G1 child's future outcomes, compared to a G1 sibling who was not maltreated. An example of this scenario is the abuse or neglect of one sibling by a community member when outside the home, such as teachers, coaches, or religious figures. This scenario is designed to extrapolate away from potential within-sibling-group differences that could arise from parental reinforcement or compensation associated with selective maltreatment of one sibling. The table presents sibling fixed effects estimates of one sibling being maltreated by an outside perpetrator.

The "outside perpetrator" sample is limited to G1 children born in 1999-2004, for whom their own birth records are available, in order to isolate sibling groups (children sharing birth mother) and in order to estimate cumulative exposure to maltreatment over the life course.

Column 1 of Table A.5 shows that the outside perpetrator sibling victim had approximately one additional substantiation during their childhood than their sibling, which is consistent with the premise of the design. Column 2 shows that fertility for the maltreated sibling is greater (economically but not statistically significant), their G2 children are more likely to

be maltreatment victims (economically but not statistically significant), and the maltreated sibling is more likely to become a confirmed perpetrator regardless of fertility (economically significant and marginally statistically significant).

Table A.6 compares the time-exposure of siblings to an abusive or neglectful home environment. The sample is limited to G1 children born in 1999-2004, for whom their own birth records are available, to isolate sibling groups and to estimate cumulative exposure to maltreatment and foster care over the life course. The sample is additionally restricted to sibling groups who (i) each experience their final removal to foster care for reasons related to maltreatment at the exact same point in time, (ii) all siblings are younger than age 14 at the time of final removal, and (iii) no further confirmed maltreatment occurs for any sibling after the final removal, regardless of whether they return home from foster care or not. As a result, the analysis sample is restrictive and point estimates are only suggestive. Specifications control for sibling fixed effects by gender, and year of birth by gender. The independent variable is a child's age at the time of final removal.

Column 1 of Table A.6 shows that for every year older a child is at the time of removal, they had 0.08 more lifetime maltreatment substantiations, which is consistent with the older child facing greater cumulative maltreatment exposure on average. More years of potential exposure are associated with greater fertility (column 2), more transmission of maltreatment to G2 children conditional on fertility (Column 3), and—suggestively—a greater frequency of being a confirmed perpetrator regardless of fertility (Columns 4 and 5; not statistically significant).

Altogether, the estimates across four different designs, each with multiple permutations, are consistent with a causal interpretation of the transmission of maltreatment from parent to child. In the following section, we provide additional evidence of the forward-looking transmission of maltreatment and mechanisms behind transmission.

Mechanisms of Intergenerational Transmission. Table 6 provides further support for the causal interpretation of maltreatment persistence by showing persistence in the *type* of maltreatment the parent and child experienced. Because the specifications in this section are demanding from a statistical power perspective, we revert to the standard sample.

Each column regresses the standardized number of G2 substantiations of a given type (physical abuse, sexual abuse, or neglect) on the standardized number of G1 childhood substantiations of each type. The diagonal entries therefore capture same-type transmission (e.g., G1 physical abuse predicting G2 physical abuse), while off-diagonal entries capture cross-type associations. For example, row 1 Column 2 shows the regression estimate on total G1 childhood physical

Table 6: Type-Type Intergenerational Association of Substantiated Maltreatment *by Maltreatment Type*

	Standardized Number of G2 Substantiations		
	(1) Physical Abuse	(2) Sexual Abuse	(3) Neglect
Standardized Number G1 Physical Abuse	0.022***	0.004**	0.026***
	(0.002)	(0.002)	(0.002)
Standardized Number G1 Sexual Abuse	0.011***	0.014***	0.018***
	(0.002)	(0.002)	(0.002)
Standardized Number G1 Neglect	0.011***	0.007***	0.035***
	(0.002)	(0.002)	(0.002)
Standard Controls	×	×	×
Mean of Dep Var	-0.000	0.000	-0.000
N Observations	464,246	464,246	464,246

Notes: This table presents estimates of the intergenerational (mother–child) association in substantiated maltreatment by type of allegation, measured at the intensive margin. All independent and dependent variables are standardized (mean 0, variance 1) for comparability within columns. Each regression estimates the effect of the number of times a mother (G1) was a substantiated victim of each type of maltreatment on the count of how many times her children (G2) were substantiated for a given maltreatment type. For example, Column 1 reports the coefficient from a regression of the number of child G2 physical-abuse substantiations on the number of mother G1's own childhood substantiations as a victim of physical abuse (row 1), sexual abuse (row 2), and neglect (row 3). The diagonal entries capture same-type persistence (e.g., G1 neglect associated with G2 neglect), while the off-diagonal entries capture cross-type associations. Standard sample control variables are included in all regressions. All specifications condition on having at least one child. Inference is based on robust standard errors. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

abuse substantiations from a regression of G2 sexual abuse substantiations on G1 physical abuse substantiations, G1 sexual abuse substantiations, and G1 neglect substantiations. Therefore, all estimates control for number of substantiations of other types.

The estimates show consistent same-type persistence. G1 physical abuse is most strongly associated with G2 physical abuse, G1 sexual abuse is most strongly associated with G2 sexual abuse, and G1 neglect is most strongly associated with G2 neglect. While cross-type associations are also positive and statistically significant, the diagonal terms are generally larger, indicating that intergenerational transmission is not purely a general increase in CPS involvement (for example, endogenous surveillance) but exhibits meaningful *type*-specific transmission. We interpret these findings as an indication of learned behavior or the reproduction of experiences in childhood.

We also test whether transmission operates through the expected perpetrator. Table A.7 uses a reverse-regression framework to examine whether a mother's (G1) own childhood maltreatment predicts the perpetrator type involved in her children's (G2) substantiations. When examined separately, G2 substantiations involving a biological mother perpetrator, biological father perpetrator, and unrelated adult perpetrator are each positively associated with G1's own childhood substantiation. However, when all perpetrator types are included jointly (Column 4), only substantiations involving the biological mother remain positive and statistically significant. This pattern suggests that the strongest transmission channel operates through the mother herself, rather than through assortative mating or general household instability.

Finally, we attempt to distinguish between whether the transmission of maltreatment reflects lived experiences or underlying characteristics. Table A.8 compares transmission based on who perpetrated G1's own childhood maltreatment. If underlying characteristics, such as genetic endowments, were the primary mechanism, we would expect stronger transmission when G1 was maltreated by her biological mother. Instead, the coefficients are similar in magnitude: experiencing maltreatment from a biological mother is associated with 0.135 additional G2 substantiations, while maltreatment from a non-biological caregiver is associated with 0.188 (or 0.152 when jointly included). The similarity of these magnitudes suggests that transmission is not driven primarily by genetic predisposition, but rather by lived experience and environmental exposure. In other words, G1 children who are maltreated themselves are more likely to have G2 children who are maltreated, regardless of whether or not the G1 perpetrator was a biological parent.

Altogether, the estimates presented in this section point to meaningful, behavioral intergenerational transmission of maltreatment rather than a purely mechanical or

surveillance-driven pattern. The type-to-type persistence and the concentration of transmission through the mother as perpetrator suggest that lived experience and internalized caregiving dynamics play an important role. Moreover, the similar transmission magnitudes for biological and non-biological perpetrators are consistent with a view that maltreatment persists across generations largely through experiential and behavioral channels.

VII Local Maltreatment Rates and Intergenerational Mobility

A central motivation for our study of the intergenerational transmission of child maltreatment is its potential relevance for economic mobility. We have shown that child maltreatment—as proxied by CPS involvement—is highly prevalent, sharply concentrated among socioeconomically disadvantaged families, and intergenerationally persistent within families. Coupled with a large literature documenting the long-run consequences of maltreatment for educational attainment, earnings, health, and involvement with the criminal justice system (Currie and Spatz Widom, 2010; Currie and Tekin, 2012), these patterns suggest that maltreatment may represent an important mechanism linking early-life disadvantage to low intergenerational mobility.

To explore this possibility, we conduct a descriptive analysis that relates local rates of child maltreatment to measures of intergenerational mobility from Chetty et al. (2014). While this exercise does not identify causal effects, it provides suggestive evidence on whether places with higher exposure to maltreatment are also places where children experience systematically worse economic trajectories in adulthood.

Figure A.10 illustrates the relationship between ZIP-code-level rates of child maltreatment investigations and absolute upward mobility, measured as the mean adult income rank of children born to parents at the 25th percentile of the national income distribution, using the mobility estimates developed by Chetty et al. (2014) and updated in Chetty et al. (2022). The relationship is strongly negative: ZIP codes with higher rates of CPS involvement exhibit substantially lower levels of upward mobility. Children growing up in high-maltreatment areas are far less likely to climb the income distribution than those raised in otherwise comparable low-maltreatment areas.

This association does not appear to be merely a reflection of broader neighborhood disadvantage. Figure 8 and Table A.9 show that the negative relationship between maltreatment rates and upward mobility persists even after controlling for a rich set of ZIP-code-level characteristics that prior work in Chetty et al. (2014) has identified as key

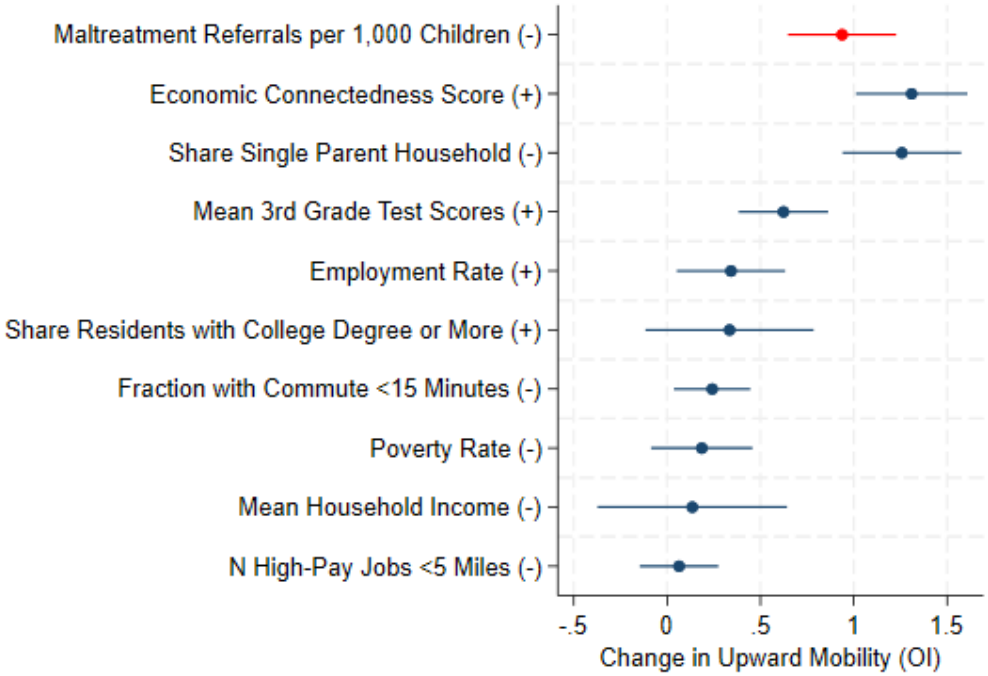
correlates of mobility. These include measures of family structure, racial composition, poverty, income inequality, residential segregation, and local labor market conditions. In fact, across specifications local child maltreatment referral rates remain among the strongest predictors of mobility outcomes. In standardized terms, variation in CPS referral rates is at least as predictive of upward mobility as many commonly emphasized neighborhood characteristics, including the share of single-parent households.

We emphasize that these results are descriptive and do not establish a causal relationship between child maltreatment and intergenerational mobility. Local CPS referral rates may proxy for many adverse conditions not fully captured by available controls, and mobility itself may influence system involvement. Nevertheless, the strength and robustness of the association suggest that maltreatment exposure may be an important component of the local environment shaping children's long-run economic prospects.

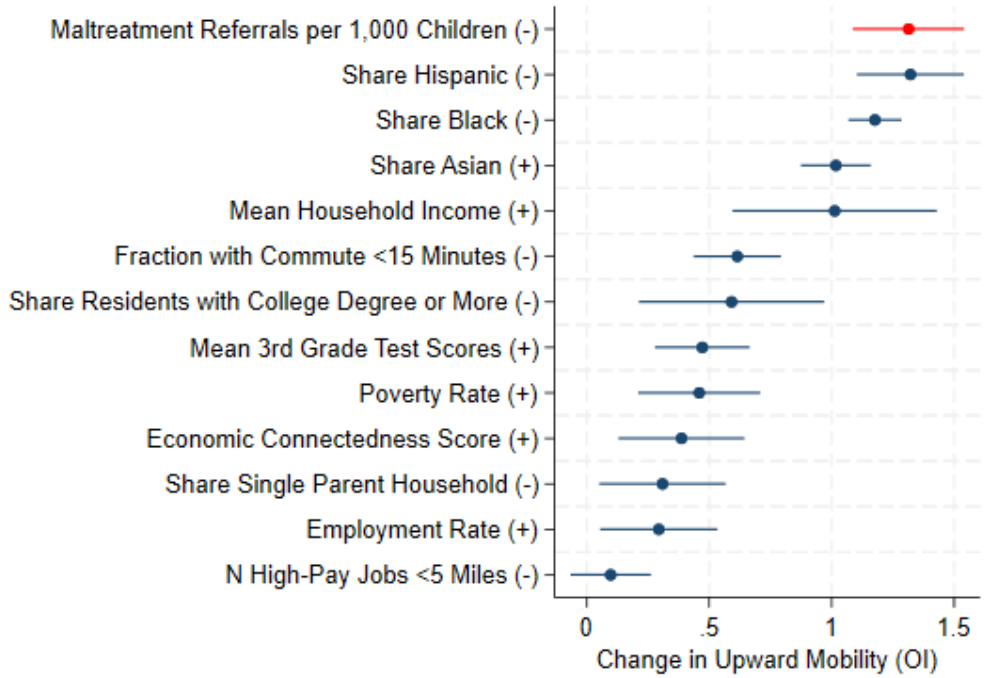
Taken together, these patterns motivate future research that directly links individual-level maltreatment exposure to earnings and mobility outcomes across generations. Rich administrative data that connect CPS records to tax data—and eventually to the outcomes of individuals' own children—would allow researchers to move beyond the correlations presented here and quantify the role of maltreatment as a mechanism for the persistence of economic disadvantage.

Figure 8: Maltreatment Referrals vs. Other Correlates with Economic Mobility

(a) Correlates with ZIP-Level Upward Mobility
All Feature Variables are Standardized



(b) Correlates with ZIP-Level Upward Mobility, including Race/Ethnicity
All Feature Variables are Standardized



Notes: This figure presents a coefficient plot of ZIP-code level upward mobility regressed on other ZIP code level

characteristics. Upward mobility is measured as the mean national income rank in adulthood for children born to parents at the 25th percentile of the national income distribution using Census tract level estimates from (Chetty et al., 2018) re-weighted by the fraction of each ZIP code's population belonging to a given tract. ZIP code characteristics are derived using tract-level characteristics from Chetty et al. (2018). Economic connectedness is defined at the ZIP code level in Chetty et al. (2022). In this figure, upward mobility is redefined on a 0-100 scale (not 0-1) for ease of interpreting magnitudes as percentile points. For comparability, all characteristics are standardized across ZIP codes with mean 0 and variance 1. The estimates are derived from regressing the mobility measure on all of the standardized characteristics simultaneously in a multivariate regression. Panel (a) shows estimates without controlling for ZIP-level demographic shares. Panel (b) shows estimates that include regressors for ZIP-level demographic shares. Estimates are ordered by magnitude, with a negative "(-)" sign indicating a negative association between the characteristic and the mobility measure.

VIII Conclusion

This paper documents the prevalence and intergenerational persistence of child maltreatment using linked administrative data from California spanning nearly three decades. We show that contact with the child welfare system is common: nearly one-third of children are referred to CPS by age 18, and a meaningful share experience substantiated maltreatment, case openings, or foster care placement. Exposure is sharply stratified by socioeconomic status at birth, with children born into the lowest-income households facing dramatically higher rates of CPS involvement than those born into more advantaged families. At the same time, CPS contact is not confined to the poorest families; it remains prevalent across much of the income distribution.

We further document that indirect exposure to maltreatment through parents is widespread. A substantial share of young children—particularly those with early CPS contact—have a parent who was themselves involved with CPS during childhood. We estimate that the majority of young children have a parent with prior CPS contact.

Using two complementary empirical strategies, including a sibling fixed effects design, we provide evidence consistent with meaningful intergenerational transmission of maltreatment. Girls who experienced substantiated maltreatment in childhood are significantly more likely to have children who themselves experience substantiated maltreatment. These associations persist when conditioning on rich observable characteristics and when comparing sisters within the same family. The magnitudes are economically meaningful across multiple measures of CPS involvement—including referrals, substantiations, case openings, and foster care placements—and are not explained by differences in fertility. Additional evidence on

type-specific transmission and perpetrator patterns suggests that persistence operates in part through behavioral channels.

Finally, we show that neighborhood-level rates of child maltreatment are strongly negatively associated with upward economic mobility. Although descriptive, this relationship underscores the possibility that maltreatment constitutes an important determinant of opportunity.

Taken together, our findings suggest that child maltreatment may operate as a behaviorally transmitted barrier to intergenerational mobility. Because maltreatment is both common and concentrated among disadvantaged families, its persistence across generations has the potential to reinforce disadvantage over time. Interventions that reduce maltreatment or mitigate its long-run consequences may therefore yield benefits that extend beyond directly affected children, attenuating the transmission of disadvantage across generations.

References

Adermon, Adrian, Mikael Lindahl, and Mårten Palme. 2021. "Dynastic Human Capital, Inequality, and Intergenerational Mobility." *American Economic Review*, 111(5): 1523–48.

Aizer, Anna, and Emilia Brito Rebolledo. 2025. "Can Early Intervention Reduce Future Child Maltreatment?" *NBER Working Paper Series*, , (33341).

Berlin, Lisa J, Karen Appleyard, and Kenneth A Dodge. 2011. "Intergenerational continuity in child maltreatment: Mediating mechanisms and implications for prevention." *Child Development*, 82(1): 162–176.

Chetty, Raj, John N Friedman, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter. 2018. "The opportunity atlas: Mapping the childhood roots of social mobility." National Bureau of Economic Research.

Chetty, Raj, Matthew O Jackson, Theresa Kuchler, Johannes Stroebel, Nathaniel Hendren, Robert B Fluegge, Sara Gong, Federico Gonzalez, Armelle Grondin, Matthew Jacob, et al. 2022. "Social capital II: determinants of economic connectedness." *Nature*, 608(7921): 122–134.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where is the land of opportunity? The geography of intergenerational mobility in the United States." *The Quarterly Journal of Economics*, 129(4): 1553–1623.

Cunha, Flávio, and James J Heckman. 2007. "The technology of skill formation." *American Economic Review*, 97(2): 31–47.

Currie, Janet, and Cathy Spatz Widom. 2010. "Long-Term Consequences of Child Abuse and Neglect on Adult Economic Well-Being." *Child Maltreatment*, 15(2): 111–120.

Currie, Janet, and Erdal Tekin. 2012. "Understanding the Cycle: Childhood Maltreatment and Future Crime." *Journal of Human Resources*, 47(2): 509–549.

Del Boca, Daniela, Christopher Flinn, and Matthew Wiswall. 2013. "Household Choices and Child Development." *The Review of Economic Studies*, 81(1): 137–185.

Doyle, Joseph J., Natalia Emanuel, and Raimundo Eyzaguirre. 2025. "Cumulative Risks of Foster Care Placement for U.S. Children: A Research Note Comparing Birth Cohort and Synthetic Cohort Approaches." *Working paper*.

Evans, Mary F, Ludovica Gazzé, and Jessamyn Schaller. 2025. "Temperature and Maltreatment of Young Children." *The Review of Economics and Statistics*, 1–37.

Evans, Mary F., Matthew C. Harris, and Lawrence M. Kessler. 2022. "The Hazards of Unwinding the Prescription Opioid Epidemic: Implications for Child Maltreatment." *American Economic Journal: Economic Policy*, 14(4): 192–231.

Finlay, Keith, Michael Mueller-Smith, and Brittany Street. 2023. "Children's Indirect Exposure to the U.S. Justice System: Evidence From Longitudinal Links between Survey and Administrative Data." *The Quarterly Journal of Economics*, 138(4): 2181–2224.

Font, Sarah, Maria Cancian, Lawrence M Berger, and Anna DiGiovanni. 2020. "Patterns of intergenerational child protective services involvement." *Child Abuse & Neglect*, 99: 104247.

Gihleb, Rania, Osea Giuntella, and Ning Zhang. 2022. "The Effect of Mandatory-Access Prescription Drug Monitoring Programs on Foster Care Admissions." *Journal of Human Resources*, 57(1): 217–240.

Goldstein, Ezra G., and Sarah A. Font. 2025. "Prevalence and Treatment of Maternal Substance Use Disorder in Child Welfare." *JAMA Health Forum*, 6(3): e250054.

Heckman, James J. 2006. "Skill formation and the economics of investing in disadvantaged children." *Science*, 312(5782): 1900–1902.

Heckman, James J, and Dimitriy V Masterov. 2007. "The productivity argument for investing in young children." *National Bureau of Economic Research Working Paper Series*, , (13016).

Herd, Toria, Sarah Font, Lindsey Palmer, Claire Selin, Stacey Shipe, Anneke Olson, Catherine Diercks, Kayla Brown, Casey Mullins, Melissa Jonson-Reid, Sheree L. Toth, and Jennie G. Noll. 2022. "Disentangling Neglect from Poverty." Policy Brief.

Kim, Hyunil, Christopher Wildeman, Melissa Jonson-Reid, and Brett Drake. 2017. "Lifetime Prevalence of Investigating Child Maltreatment Among US Children." *American Journal of Public Health*, 107(2): 274–280.

Lindo, Jason M., Jessamyn Schaller, and Benjamin Hansen. 2018. "Caution! Men not at work: Gender-specific labor market conditions and child maltreatment." *Journal of Public Economics*, 163: 77–98.

Palmer, Lindsey, Sarah Font, Andrea Lane Eastman, Lillie Guo, and Emily Putnam-Hornstein. 2022. "What Does Child Protective Services Investigate as Neglect? A Population-Based Study." *Child Maltreatment*.

Peterson, Cora, Curtis Florence, and Joanne Klevens. 2018. "The Economic Burden of Child Maltreatment in the United States, 2015." *Child abuse & neglect*, 86: 178–183.

Prindle, John, Himal Suthar, and Emily Putnam-Hornstein. 2023. "An open-source probabilistic record linkage process for records with family-level information: Simulation study and applied analysis." *Plos one*, 18(10): e0291581.

Putnam-Hornstein, Emily, Eunhye Ahn, John Prindle, Joseph Magruder, Daniel Webster, and Christopher Wildeman. 2021. "Cumulative rates of child protection involvement and terminations of parental rights in a California birth cohort, 1999–2017." *American journal of public health*, 111(6): 1157–1163.

Rittenhouse, Katherine. 2025. "Income and Child Maltreatment: Evidence from a Discontinuity in Tax Benefits." *The Review of Economics and Statistics*, 1–46.

Ryan, Joseph P, Brian A Jacob, Max Gross, Brian E Perron, Andrew Moore, and Sharlyn Ferguson. 2018. "Early exposure to child maltreatment and academic outcomes." *Child maltreatment*, 23(4): 365–375.

Thornberry, Terence P, Kelly E Knight, and Peter J Lovegrove. 2012. "Does maltreatment beget maltreatment? A systematic review of the intergenerational literature." *Trauma, Violence, & Abuse*, 13(3): 135–152.

Widom, Cathy Spatz. 1989. "The cycle of violence." *Science*, 244(4901): 160–166.

Widom, Cathy Spatz, Sally J Czaja, and Kimberly A DuMont. 2015. "Intergenerational transmission of child abuse and neglect: Real or detection bias?" *Science*, 347(6229): 1480–1485.

Wildeman, Christopher, and Natalia Emanuel. 2014. "Cumulative Risks of Foster Care Placement by Age 18 for US Children, 2000–2011." *PloS one*, 9(3): e92785.

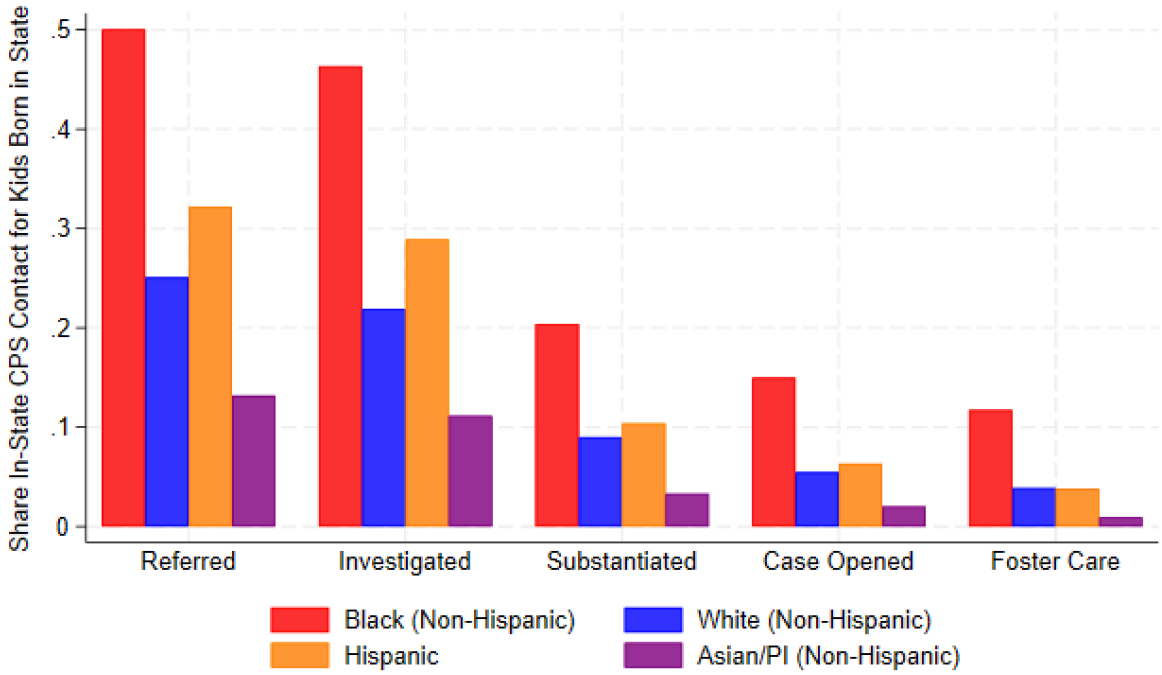
Wildeman, Christopher, Natalia Emanuel, John M Leventhal, Emily Putnam-Hornstein, Jane Waldfogel, and Hedwig Lee. 2014. "The prevalence of confirmed maltreatment among US children, 2004 to 2011." *JAMA pediatrics*, 168(8): 706–713.

Yi, Youngmin, Frank R Edwards, and Christopher Wildeman. 2020. "Cumulative prevalence of confirmed maltreatment and foster care placement for US children by race/ethnicity, 2011–2016." *American journal of public health*, 110(5): 704–709.

Appendix A

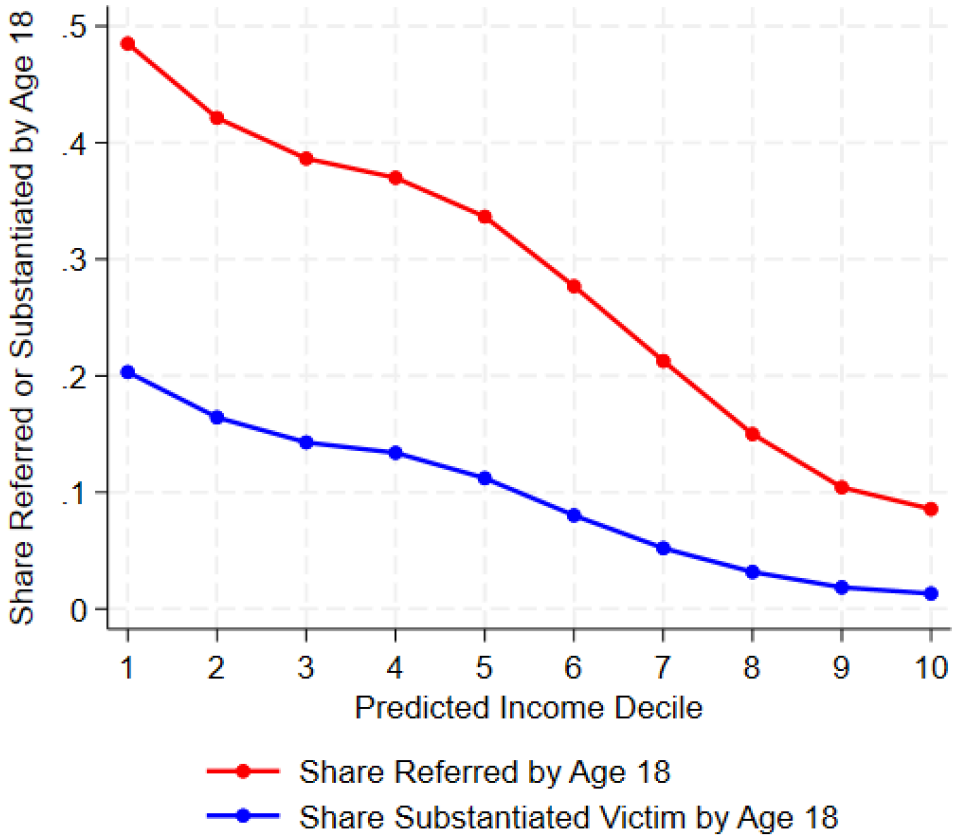
A Additional Figures

Figure A.1: Prevalence of Childhood CPS Contact, by Race/Ethnicity



Notes: This figure reports the cumulative share of children in the 1999-2001 California birth cohorts who experienced each type of CPS involvement before their 18th birthday, disaggregated by mother's race and ethnicity as recorded on the child's birth certificate. Involvement is defined in non–mutually exclusive terms. For example, children who were placed in foster care are also included in the categories of referral, investigation, substantiation, and case opening. All estimates are based on linked administrative records from California birth certificates and California CPS data, implying that estimates are a lower bound.

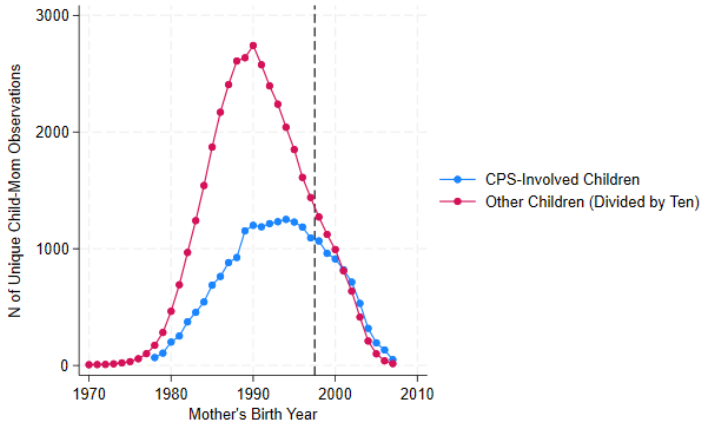
Figure A.2: Prevalence of Childhood CPS Contact and Substantiation, by Income Decile



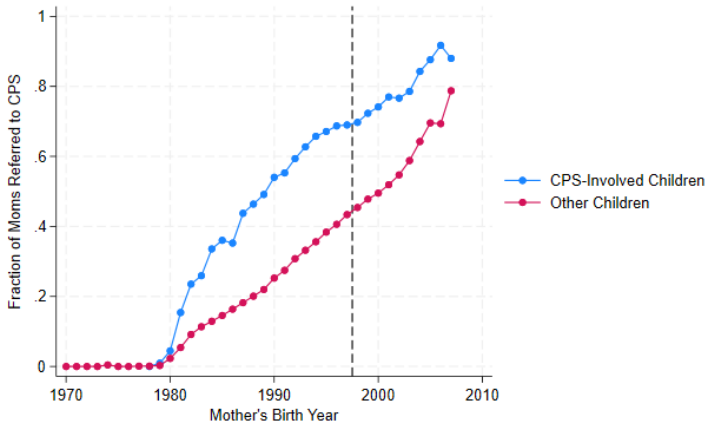
Notes: This figure reports the cumulative share of children in the 1999-2001 California birth cohorts who experienced each type of CPS involvement before their 18th birthday, disaggregated by estimated income decile at birth. Involvement is defined in non–mutually exclusive terms. For example, children who are substantiated victims are also included in the referral category. All estimates are based on linked administrative records from California birth certificates and California CPS data, implying that the estimates are a lower bound.

Figure A.3: IG Backward-Looking: Mother Age and CPS Contact in Raw Data

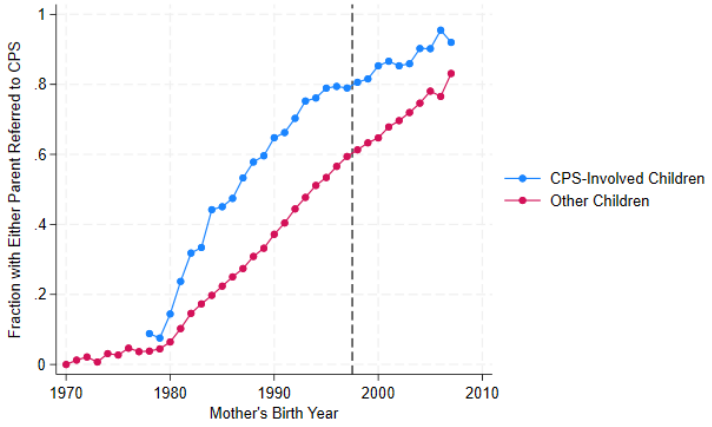
(a) Distribution of Mom Ages, by Child CPS Contact



(b) Mother CPS Contact (1998-), by Child CPS Contact



(c) Either Parent CPS Contact (1998-), by Child CPS Contact

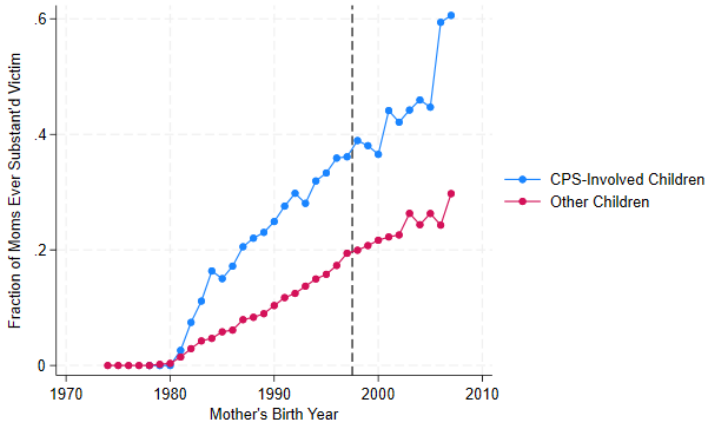


Notes: This figure presents the distribution of mom (G1) birth-years relative to the year of first available

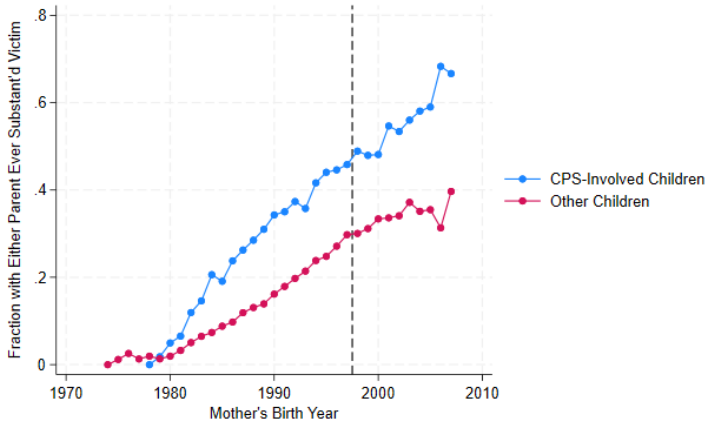
CPS records (1998) for the California backward-looking sample with G2 children born in 2022. Panel (a) shows the distribution of mother (G1) birth years – derived from mother's age on the G2 child's birth record – for CPS-involved G2 children and never-CPS-involved G2 children, respectively. Most G1 mothers with the current data linkage were born prior to the availability of CPS records in 1998, requiring imputation for the backward-looking descriptive estimates. Panel (b) shows mother's rate of CPS contact during her childhood in the raw data, measured using records from 1998-onward, by mother's birth year and child's CPS contact status. Mothers of CPS-involved children themselves had higher rates of CPS contact as children, even conditional on birth year to account for systematic age-at-birth differences for CPS-involved families. Panel (c) presents rates of either parents' (Mother or Father/second parent if listed on G2 birth record) CPS contact in the raw data, disaggregated by mother birth year and G2 child contact.

Figure A.4: IG Backward-Looking: Mother Age and Parent Substantiations in the Raw Data

(a) Mother Confirmed Maltreatment Victim (1998-), by Child CPS Contact

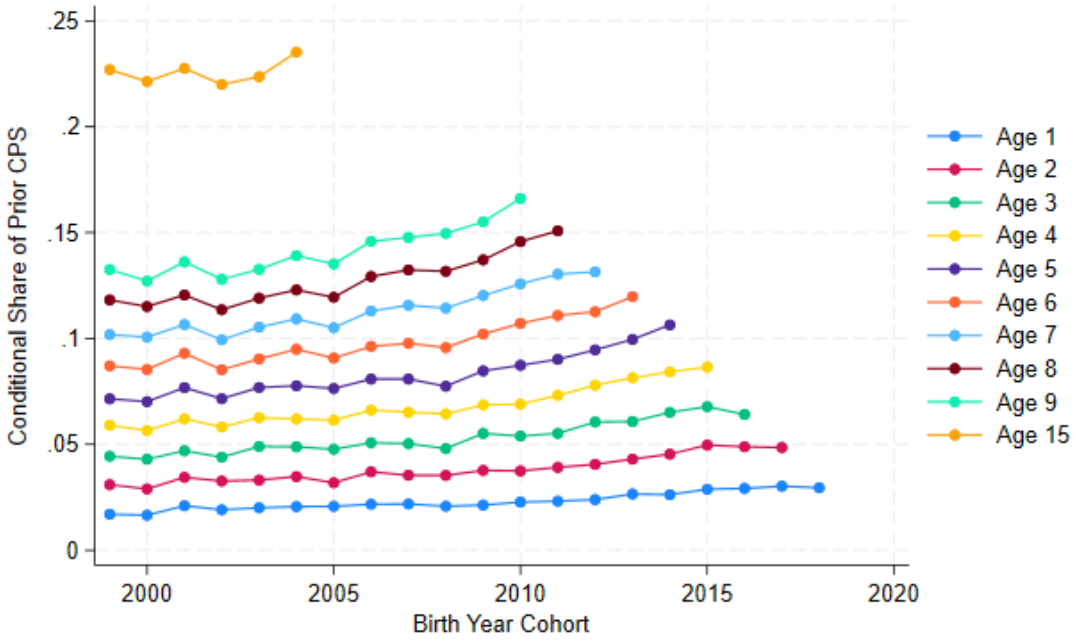


(b) Either Parent Confirmed Maltreatment Victim (1998-), by Child CPS Contact



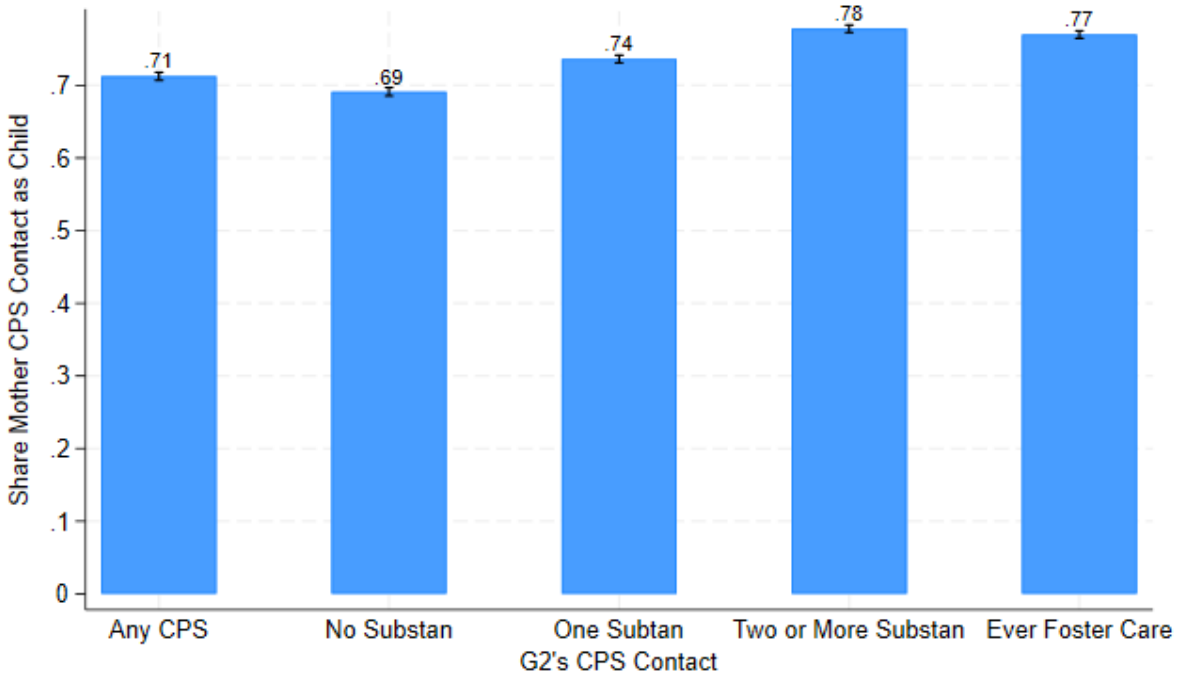
Notes: Note: This figure presents rates of maltreatment victimization in raw CPS data (1998-2024) for parents of children born in California in 2022. The procedure is identical to that in Figure A.3 but for parent substantiation instead of parent referral.

Figure A.5: Cohort Rates of Prior CPS Contact, Conditional on No Contact After Age X



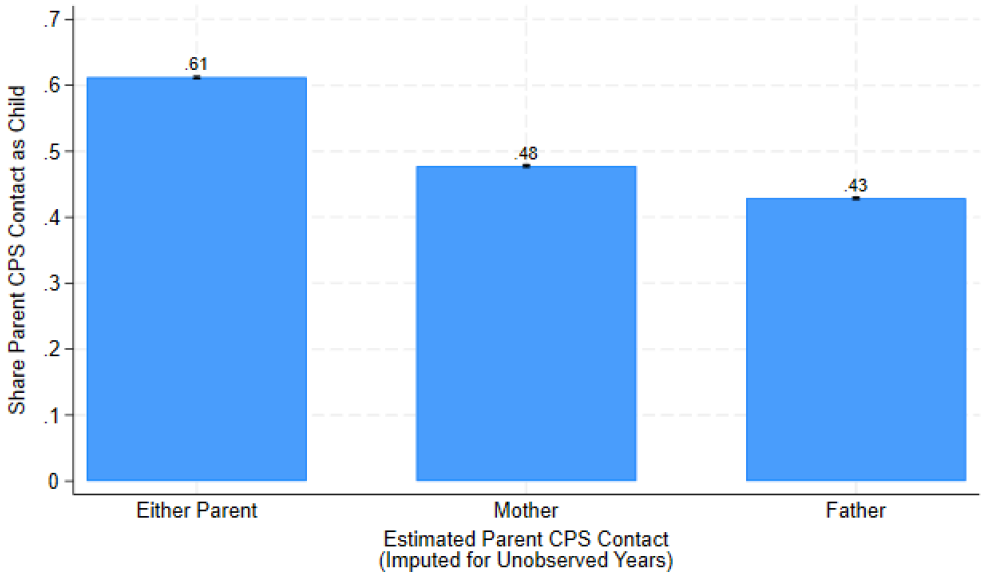
Notes: Because CPS data are unavailable prior to 1998, estimating rates of parents' CPS contact requires imputing unobserved CPS contact for those earlier years. This figure shows the share of California girls who had at least one CPS referral before age X, conditional on having no referral at or after age X, by birth cohort and for selected ages ($X \in \{[1, 9], 15\}$). For example, for the 1999 birth cohort, approximately 5 percent of children who did not have any referral from ages 3-18 did have a prior CPS referral from ages 0-2. The stability of these conditional rates across cohorts provides support for using the earliest observed empirical estimates at each age to impute mothers' likely CPS contact prior to the start of matched CPS data (1998). An analogous imputation procedure is used for men/fathers (not shown).

Figure A.6: Backward-Looking: G1 (Future Mother) CPS Contact by Intensity of G2 Contact



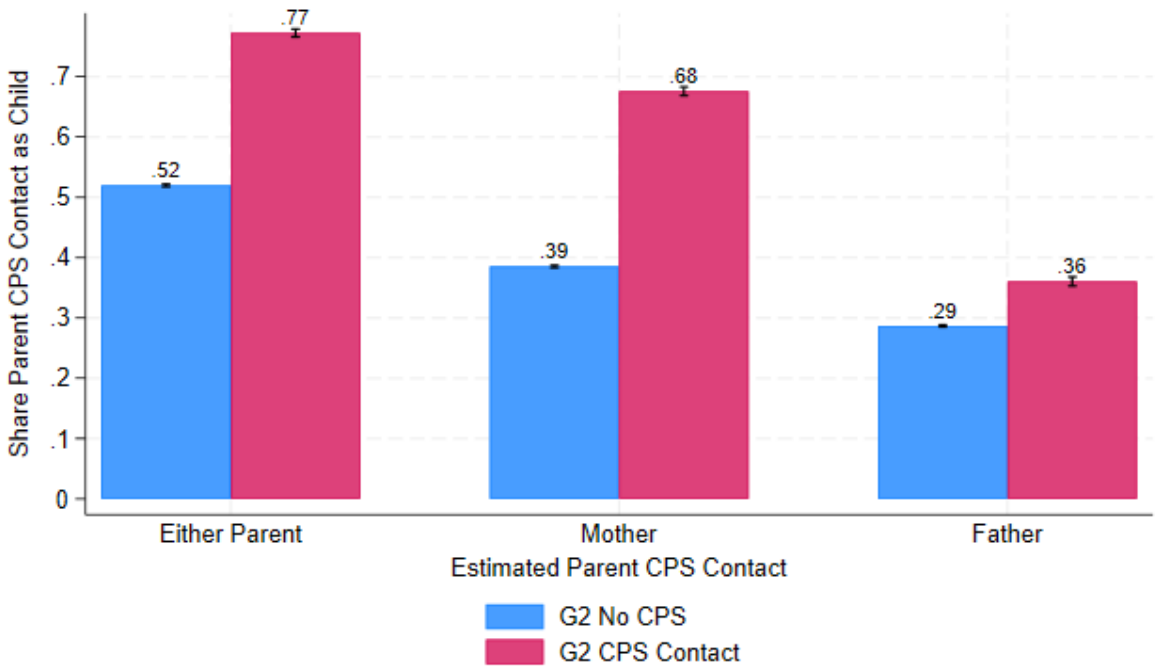
Notes: This figure presents rates of mother's (G1) imputed California CPS contact, split by child's (G2) intensity of CPS contact for children in the 2022 California birth cohort. The estimates use mothers' imputed contact for all mothers themselves born in state. The first estimate includes all children referred to CPS before age 3. The second estimate uses the subset of ever-referred children who were never substantiated victims of maltreatment before age 3; the third estimate uses the subset of ever-referred children who had exactly one substantiation before age 3, and so on.

Figure A.7: Backward-Looking: Overall Rates of Parent CPS Contact, Inferred



Notes: This figure shows imputed rates of CPS contact for both parents (G1). Panel 1 shows the percentage of children (G2) in the 2022 California birth cohort who have at least one parent listed on the birth record who was referred to CPS. Panel 2 shows rates of contact for mothers, and Panel 3 shows rates of contact for fathers (second listed parent on birth record). The sample is restricted to children born in 2022 to California-born mothers. Parental CPS histories are measured using raw administrative records for 1998-onward, and with imputation based on conditional contact prior to 1998 (See Figure A.5). Father estimates are likely understated due to unlisted fathers on some birth records (24 percent unlisted for CPS-involved children and 5 percent for non-CPS-involved children).

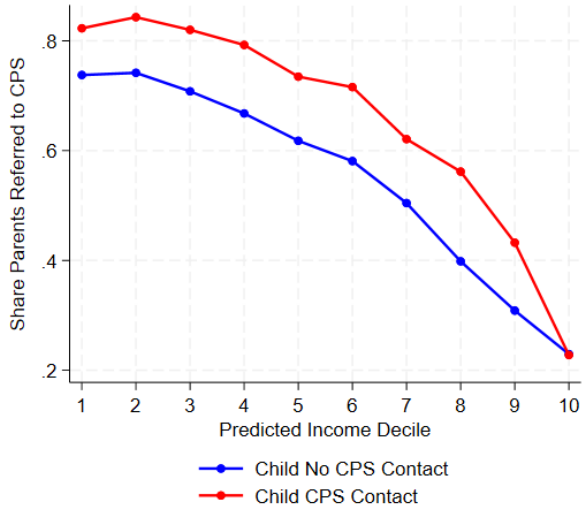
Figure A.8: Backward-Looking: G1 Parent CPS Contact by Intensity of G2 Contact (Raw Data)



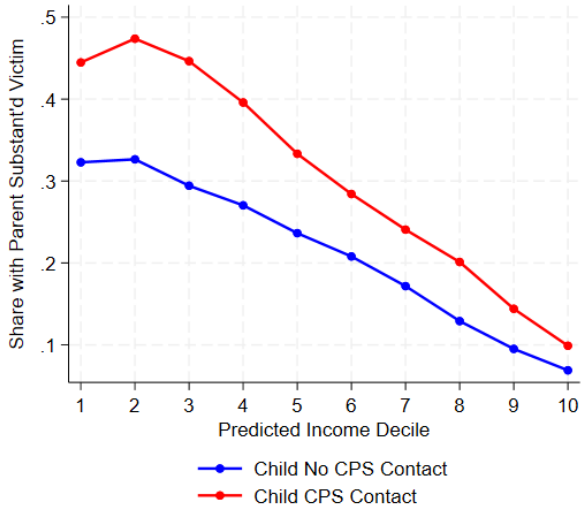
Notes: This figure shows raw rates of CPS contact for both parents (G1), split by child (G2) CPS contact. Panel 1 shows the percentage of children (G2) in the 2022 California birth cohort who have at least one parent listed on the birth record who was referred to CPS. Panel 2 shows rates of contact for mothers, and Panel 3 shows rates of contact for fathers (second listed parent on birth record). The sample is restricted to children born in 2022 to California-born mothers. Parental CPS histories are measured using raw administrative records with no imputation, and therefore may understate true exposure due to left-censoring prior to 1998. Father estimates are likely understated due to unlisted fathers on some birth records (24 percent unlisted for CPS-involved children and 5 percent for non-CPS-involved children).

Figure A.9: Backward-Looking: Inferred G1 Parent CPS Contact by Income Decile

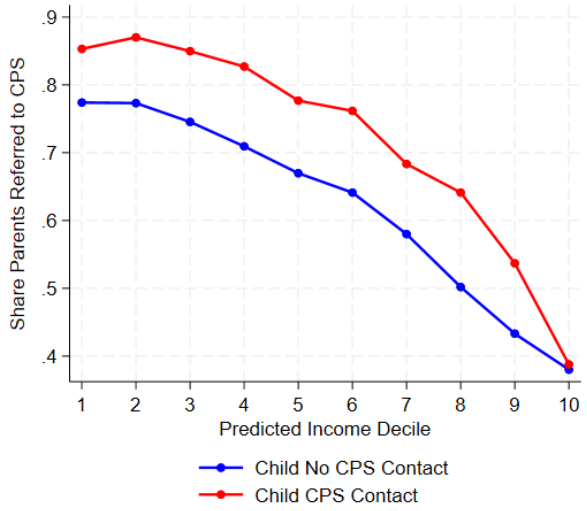
(a) Raw Parent CPS Contact



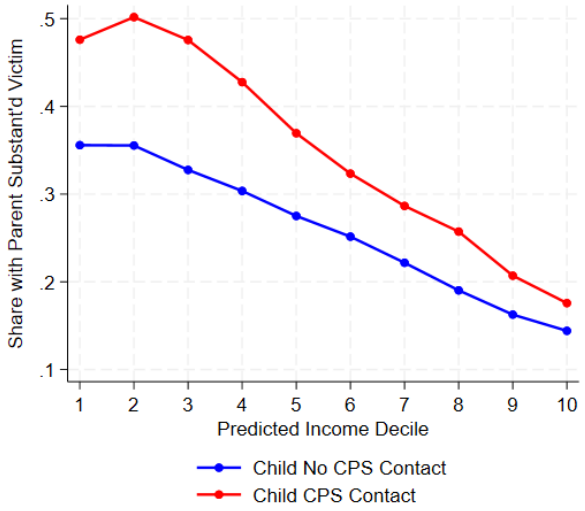
(b) Raw Parent Substantiated Victim



(c) Imputed Parent CPS Contact

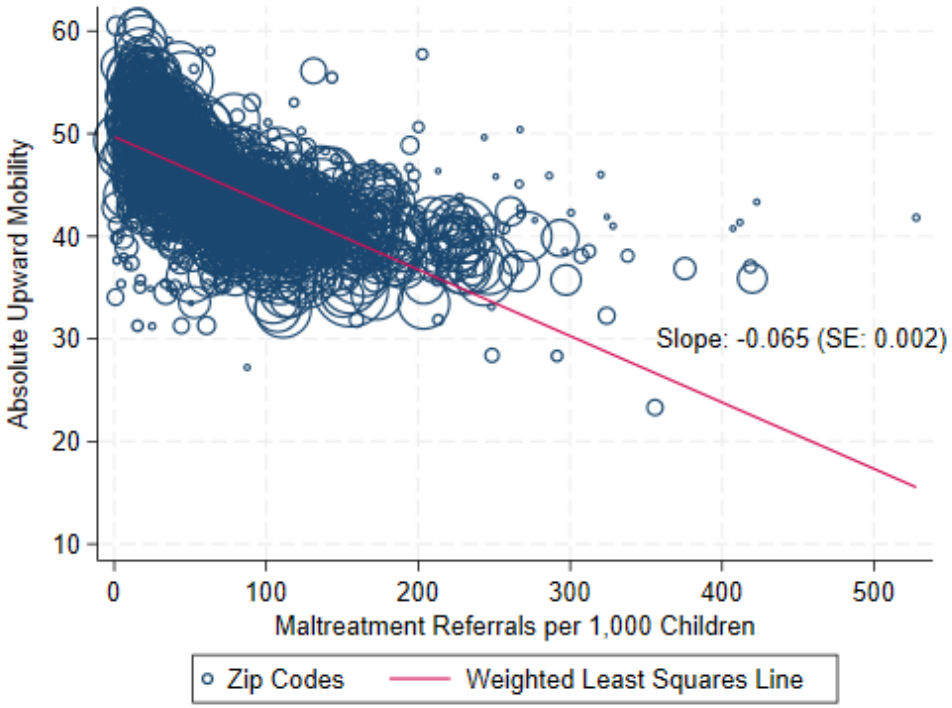


(d) Imputed Parent Substantiated Victim



Notes: This figure shows, by estimated parental income decile, the share of children who have at least one parent (mother or father) with a recorded history of CPS involvement (left column) or confirmed maltreatment (right column) during childhood. Subfigures present separate lines are shown for children with no CPS contact before age three (blue) and children with any CPS contact before age three (red). The sample is restricted to children born in 2022 to California-born mothers. Parental CPS histories are measured using raw administrative records for 1998-onward in panels (a) and (b), and additionally with imputation based on conditional contact prior to 1998 in panels (c) and (d) (See Figure A.5). Predicted income is estimated using birth certificate characteristics.

Figure A.10: Neighborhood Maltreatment Investigations and Upward Mobility



Notes: This figure shows the association between neighborhood (ZIP code) rates of alleged maltreatment and upward economic mobility. ZIP-code-level child maltreatment referral rates are calculated using our administrative dataset for California over the ten-year period 2000-2009. Absolute upward mobility is defined as the average national income rank in adulthood of children born to parents in the 25th percentile of the income distribution (Chetty et al., 2014). Mobility estimates and ZIP-code-level characteristics are obtained from 2000 when available.

B Additional Tables

Table A.1: Robustness to Differences in Maltreatment Victimization on a Given Referral

Sibling Fixed Effects Robustness

	(1) Any Birth	(2) Tot G2 Referral	(3) Tot G2 Investigation	(4) Tot G2 Substan	(5) Tot G2 Case	(6) Tot G2 Foster
Additional G1 Substantiation on a Shared Referral	0.020***	0.849***	0.631***	0.200***	0.107***	0.082***
	(0.005)	(0.175)	(0.121)	(0.038)	(0.024)	(0.020)
Referral Fixed Effects +	×	×	×	×	×	×
Mean of Dep Var	0.510	3.963	2.739	0.663	0.380	0.274
N Observations	29,629	10,444	10,444	10,444	10,444	10,444

Notes: This table presents robustness checks using a restrictive within-referral variation sample. Unlike the primary analysis, which leverages variation across all childhood referrals, this design identifies the association using only referrals where at least one sibling is a victim and one sibling is not, on the given referral. This isolates variation at a single point in time but changes the interpretation of the coefficient to reflect substantiation on a specific discordant referral rather than cumulative childhood exposure. Standard errors are clustered at the referral level. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Impacts of G1 Maltreatment on G1 Fertility and Measures of G2 Maltreatment

Selection-on-Observables Design

	(1) Any Birth	(2) Tot G2 Referral	(3) Tot G2 Investigation	(4) Tot G2 Substan	(5) Tot G2 Case	(6) Tot G2 Foster
Panel A: Number of G1 Childhood Substantiations						
Total G1 Childhood Substantiations	0.028***	0.522***	0.367***	0.108***	0.072***	0.062***
	(0.001)	(0.015)	(0.011)	(0.003)	(0.002)	(0.002)
Panel B: Any G1 Childhood Substantiation						
G1 Any Substantiation as a Child	0.040***	0.643***	0.452***	0.131***	0.089***	0.077***
	(0.001)	(0.022)	(0.015)	(0.005)	(0.003)	(0.003)
Standard Controls	×	×	×	×	×	×
Mean of Dep Var	0.395	2.629	1.814	0.444	0.252	0.173
N Observations	1,175,272	464,246	464,246	464,246	464,246	464,246

Notes: This table extends the selection-on-observables specification from Column 2 of Table 4. Each column presents results from separate regressions in which the dependent variable corresponds to a different outcome. Panel A uses the number of G1 childhood substantiations as the explanatory variable, while Panel B compares future G1 mothers with at least one substantiation to those with none. Column 1 reports the effect of G1 maltreatment on the likelihood of ever having at least one child. Columns 2–6 condition on having at least one child and report effects on various measures of G2 child welfare involvement: (2) number of G2 child CPS referrals, (3) investigations, (4) substantiations, (5) case openings, and (6) removals to foster care. Standard errors are corrected for heteroskedasticity. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Child Victims Become Future Perpetrators, Regardless of Fertility

Panel A: Girl Victims

	(1) Alleged Perp Role: Any Guardian	(2) Confirmed Perp Role: Any Guardian	(3) Alleged Perp Role: Non-Guardian	(4) Confirmed Perp Role: Non-Guardian
Panel A: Number of G1 Childhood Substantiations				
Total G1 Childhood Substantiations	0.033***	0.018***	0.002***	0.001***
	(0.001)	(0.001)	(0.001)	(0.000)
Panel B: Any G1 Childhood Substantiation				
G1 Any Substantiation as a Child	0.048***	0.022***	0.0019**	0.00086**
	(0.0027)	(0.0019)	(0.00089)	(0.00035)
Referral Fixed Effects +	×	×	×	×
Mean of Dep Var	0.192	0.077	0.015	0.002
N Observations	926,192	926,192	926,192	926,192

Panel B: Boy Victims

	(1) Alleged Perp Role: Any Guardian	(2) Confirmed Perp Role: Any Guardian	(3) Alleged Perp Role: Non-Guardian	(4) Confirmed Perp Role: Non-Guardian
Panel A: Number of G1 Childhood Substantiations				
Total G1 Childhood Substantiations	0.007***	0.001	0.004***	0.001*
	(0.001)	(0.001)	(0.001)	(0.000)
Panel B: Any G1 Childhood Substantiation				
G1 Any Substantiation as a Child	0.0080***	0.0021*	0.0041***	0.00083
	(0.0020)	(0.0013)	(0.0013)	(0.00065)
Referral Fixed Effects +	×	×	×	×
Mean of Dep Var	0.085	0.030	0.034	0.007
N Observations	966,357	966,357	966,357	966,357

Notes: This table reports the association of an individual's child maltreatment substantiations from ages 0-14 with their future appearances on child maltreatment reports as an older adolescent or adult perpetrator (ages 15+), regardless of the individual's own fertility. Estimates are reported using a sibling fixed effects design, where two siblings appearing on the same referral vary in their maltreatment victimization exposure. Panel A reports estimates for girls/women. Panel B reports estimates for boys/men. Outcomes for columns 1 and 2 include future incidents as a guardian perpetrator (bio parent, non-bio (step-)parent, or other guardian). Outcomes for columns 3-4 include incidents as a non-guardian and non-sibling perpetrator. Outcomes for columns 1 and 3 include all referrals, regardless of substantiation. The dependent variable in columns 2 and 4 includes only reports where sufficient evidence was gathered to confirm the individual as a perpetrator of substantiated maltreatment. Standard errors are clustered at the referral level. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4: Robustness of Intergenerational Transmission to Intensive Margin of Fertility

	Standard Sample		Sibling FE Sample	
	(1) G2 Child-Years Observed	(2) G2 Child Substan Per Year	(3) G2 Child-Years Observed	(4) G2 Child Substan Per Year
G1 Any Substantiation as a Child	0.913***	0.005***	0.881***	0.006***
	(0.026)	(0.000)	(0.100)	(0.001)
Standard Controls	×	×		
Referral Fixed Effects +			×	×
Mean of Dep Var	7.451	0.019	8.751	0.021
N Observations	1,175,272	464,177	926,192	284,658

Notes: This table addresses possible differences between maltreated and non-maltreated women not just on the extensive margin of fertility (any child), but also the intensive margin (differences in child-years of exposure to potential maltreatment reports due to differences in mother age-at-birth and total number of children). Columns 1 and 3 show the additional number of child-years for maltreated girls. Columns 2 and 4 scale G2 child substantiations by the window of observability (from child birth through Q4 2024, added together for all G2 children born to a given mother). The estimates show that the IG transmission result is robust to differences in fertility on the intensive margin. Columns 1-2 present estimates using the standard specification with a rich set of selection-on-observables control variables. Columns 3-4 present estimates from the sibling fixed effects specification. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Transmission of Maltreatment Robustness Check: One Sibling Maltreated by a Perpetrator from Outside the Home

	(1) Total G1 Substantiations	(2) Any Birth	(3) Total G2 Substantiations	(4) N Times Became Perpetrator
Additional G1 Substantiation from Perp Outside Home	0.982***	0.025	0.040	0.088*
	(0.068)	(0.045)	(0.034)	(0.046)
Sibling-group-by-gender FEs	×	×	×	×
Birth-cohort-by-gender FEs	×	×	×	×
Mean of Dep Var	1.220	0.208	0.041	0.057
N Observations	245	245	245	245

Notes: This table presents sibling fixed effects estimates of one sibling being maltreated by an outside perpetrator. These estimates provide one example of maltreatment that may be more independent of any potential parental reinforcement of – or compensation for – sibling differences. The sample is limited to G1 children born in 1999-2004, for whom their own birth records are available, in order to isolate sibling groups with high precision and in order to estimate cumulative exposure to maltreatment over the life course with high precision. As a result, the analysis sample is restrictive and point estimates are only suggestive. Column 1 shows that the outside perpetrator sibling victim had approximately one additional substantiation during their childhood than their sibling, which is consistent with the premise of the design. Column 2 shows that fertility for the maltreated sibling is greater (not significant), their G2 children are more likely to be maltreatment victims (economically but not statistically significant), and the maltreated sibling is more likely to become a confirmed perpetrator regardless of fertility (economically significant and marginally statistically significant). Standard errors are clustered at the referral level. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Transmission of Maltreatment Robustness Check: Sibling Years of Maltreatment Prior to Removal to Foster Care

	(1) Total G1 Substantiations	(2) Any Birth	(3) Tot G2 Substantiations	(4) N Times Bio Parent Perpetrator	(5) N Times Other Perpetrator
G1 Age at Final Foster Care Removal	0.075***	0.045**	0.235*	0.001	0.005
	(0.025)	(0.018)	(0.140)	(0.010)	(0.005)
Sibling-group-by-gender FEs	×	×	×	×	×
Birth-cohort-by-gender FEs	×	×	×	×	×
Fertility-by-gender FEs			×	×	×
Mean of Dep Var	1.793	0.144	0.173	0.012	0.004
N Observations	8,875	8,875	346	8,875	8,875

Notes: This table presents an alternative design to estimate the impact of maltreatment: comparing the time-exposure of siblings to an abusive or neglectful home environment. The sample is limited to G1 children born in 1999-2004, for whom their own birth records are available, in order to isolate sibling groups with high precision and in order to estimate cumulative exposure to maltreatment and foster care over the life course with high precision. The sample is additionally restricted to sibling groups who (i) each experience their final removal to foster care for reasons related to maltreatment at the exact same point in time, (ii) all siblings are younger than age 14 at the time of final removal, and (iii) no further confirmed maltreatment occurs for any sibling after the final removal, regardless of whether they return home from foster care or not. As a result, the analysis sample is restrictive and point estimates are only suggestive. Specifications control for sibling fixed effects by gender, and year of birth by gender. The independent variable is a child's age at the time of final removal. For example, Column 1 shows that for every year older a child is at the time of removal, they had 0.08 more lifetime maltreatment substantiations, which is consistent with the older child facing greater cumulative maltreatment exposure on average. More years of potential exposure is associated with greater fertility (column 2), more transmission of maltreatment to G2 children conditional on fertility (column 3), and – suggestively – a greater frequency of being a confirmed perpetrator regardless of fertility (columns 4 and 5; not statistically significant). Standard errors are clustered at the referral level. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Intergenerational Transmission: Transmission through the Expected Perpetrator (Prior Victim)

	Mother: Any Substan as Child			
	(1)	(2)	(3)	(4)
Tot G2 Substan: Perp Bio Mom	0.015***			0.017***
	(0.001)			(0.001)
Tot G2 Substan: Perp Bio Dad		0.014***		-0.003**
		(0.001)		(0.001)
Tot G2 Substan: Unrelated			0.016***	-0.002
			(0.002)	(0.002)
Standard Controls	×	×	×	×
Mean of Dep Var	0.350	0.350	0.350	0.350
N Observations	464,246	464,246	464,246	464,246

Notes: This table presents reverse-regression estimates of the intergenerational association between parents' (G1) childhood maltreatment histories and the perpetrator types involved in their children's (G2) substantiated maltreatment. The dependent variable is an indicator for whether the G1 mother experienced any substantiated maltreatment as a child. The key independent variables are counts for how often any G2 child was substantiated with a perpetrator of a given type (biological mother, father, or unrelated adult). Columns 1–3 report pairwise regressions for each perpetrator type, while Column 4 includes all three categories jointly. All models include the standard set of controls. Inference is based on robust standard errors. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.8: Maltreatment from Kin *and* Non-Kin Perpetrators is Transmitted
Test of Genetics vs. Lived Experience

	Tot G2 Substantiations		
	(1)	(2)	(3)
Any G1 Substan: Bio Mom	0.135***		0.133***
	(0.006)		(0.006)
Any G1 Substan: Non-Bio Mom		0.188***	0.152***
		(0.026)	(0.026)
Standard Controls	×	×	×
Mean of Dep Var	0.444	0.444	0.444
N Observations	464,246	464,246	464,246

Notes: This table presents a test of whether the intergenerational transmission of maltreatment reflects genetic inheritance or lived experience, using information on perpetrator type for a mother's (G1) own childhood maltreatment. Each column reports results from regressions of a count of G2 child substantiations on indicators for whether their G1 mother was ever substantiated as a victim of maltreatment by a given perpetrator type (e.g., her biological mother, or a non-biological (step)mother). If transmission was largely based on genetic predisposition, stronger G1-G2 associations would be expected for G1 maltreatment from G0 biological parents. In this instance, however, coefficients for transmission of maltreatment experienced by non-biological perpetrators are of similar magnitude. All outcomes include the standard control set. Inference is based on robust standard errors. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: ZIP-Level CPS Contact and Upward Mobility

	No Race Controls	Race Controls
	(1)	(2)
Maltreatment Referrals per 1,000 Children	-0.937***	-1.314***
	(0.148)	(0.116)
Economic Connectedness Score	1.310***	0.387***
	(0.152)	(0.131)
Share Single Parent Household	-1.258***	-0.309**
	(0.162)	(0.132)
Mean 3rd Grade Test Scores	0.623***	0.471***
	(0.122)	(0.0987)
Employment Rate	0.342**	0.294**
	(0.148)	(0.122)
Share Residents with College Degree or More	0.335	-0.591***
	(0.230)	(0.193)
Fraction with Commute <15 Minutes	-0.242**	-0.614***
	(0.104)	(0.0910)
Poverty Rate	-0.186	0.459***
	(0.138)	(0.127)
Mean Household Income	-0.135	1.012***
	(0.259)	(0.213)
N High-Pay Jobs <5 Miles	-0.0643	-0.0972
	(0.107)	(0.0839)
Share Hispanic		-1.322***
		(0.111)
Share Black		-1.177***
		(0.0553)
Share Asian		1.017***
		(0.0727)
Outcome Mean	44.540	44.540
Observations	1,595	1,595

Notes: The dependent variable is absolute upward mobility, defined as the mean national income rank in adulthood of children born to parents at the 25th percentile of the income distribution using Census tract level estimates from (Chetty et al., 2018) re-weighted by the fraction of each ZIP code's population belonging to a given tract. ZIP code characteristics are derived using tract-level characteristics from Chetty et al. (2018). Economic connectedness is defined at the ZIP code level in Chetty et al. (2022). In this table, upward mobility is redefined on a 0-100 scale (not 0-1) for ease of interpreting magnitudes as percentile points. All right-hand-side variables are standardized to have mean 0 and variance 1. Robust standard errors are reported in parentheses.

B Supplemental Materials

BA Estimated Household Income at Birth

Although California birth certificates do not directly report parent income, they contain a rich set of socioeconomic characteristics that can be used to predict household income. We impute household income at the time of birth using a regression model trained on California household data from the American Community Survey (ACS). There is an overlapping set of covariates between the birth records and ACS records, allowing us to predict income on birth records by applying the weights from the ACS prediction model. Our approach follows a two-stage procedure akin to Rittenhouse (2025).

BA.1 Training Sample from the American Community Survey

Household income is inferred using one-year IPUMS ACS files from 2000–2019. To approximate the California birth population, the ACS sample is restricted to households in California with at least one young children (ages 0-2) where the mother is present in the household. For each child, we record the following mother characteristics: race/ethnicity (non-Hispanic Black, non-Hispanic White, Hispanic, and non-Hispanic Asian/Pacific Islander), age group (15-19, 20-24, 25-29, 30-34, 35-39, 40-49, 50+), education level (less than high school, high school, some college, bachelor's degree, five or more years of postsecondary), and number of own children in the household (1, 2, 3+). In addition to mother characteristics, features found in the ACS that correspond with birth record variables include: family county of residence and whether the child is covered by Medicaid.[17] Finally, father education level is measured in the same ways as for mothers. An indicator for whether a father is present is included due to the substantial share of children with no co-resident father. When the father is absent, father-specific variables are set to a missing category.

The predicted variables from the ACS estimation are household income (in 2010 dollars), family income (expressed as a percentage of the federal poverty line), and mother's and father's personal income.

To predict household income at birth and other outcomes of interest, we estimate a linear regression model using the ACS training sample at the child level:

$$Y_i = \mathbf{X}_i\boldsymbol{\beta} + \boldsymbol{\gamma}_c + \epsilon_i \tag{5}$$

[17]The corresponding birth record variable, delivery covered by Medicaid, is slightly more generous but again captures a lower-income segment of the population.

where Y_i is household socioeconomic status (e.g., household income in 2010 dollars) for child i in the ACS, $\mathbf{X}_i$ includes the full set of harmonized covariates as well as all pairwise interactions among the demographic variables , and $\boldsymbol{\gamma}_c$ denotes county fixed effects. The model is estimated using ACS person weights. This flexible pairwise specification allows the relationship between covariates and income to vary across demographic subgroups. For example, the income penalty associated with low maternal education may differ by race/ethnicity or by whether a father is present.

BA.2 Matching Variables

Birth record and ACS variables have been harmonized for a set of common predictors. For birth records, we observe: mother's age at delivery (grouped into age bins corresponding to those defined for the ACS sample), mother's race and ethnicity (corresponding to those in the ACS sample), mother's level of education (corresponding when possible to the binned ACS categories), child's birth parity (corresponding to number of children in the home in the ACS), father's presence on the birth record (corresponding to father's presence in the household in the ACS), father's education (when available), county of birth (corresponding to county of residence in the ACS), and whether the delivery was covered by Medicaid (corresponding to whether the child is covered by Medicaid in the ACS).

Although birth records and ACS data share many common variables categories, matching between ACS and birth record variables is imperfect. For example, the ACS measures the number of own children currently in the household, whereas birth certificate records the number of prior live births. Similarly, mother's age in the ACS reflects age at survey, while the birth records indicate mother's age at delivery. These discrepancies may introduce modest prediction error but are unlikely to substantively affect the SES gradients documented in this paper.

BB Out-of-Sample Prediction

We apply the estimated coefficients from the ACS prediction model out of sample to generate predicted household income for each child birth record. Specifically we estimate:

$$\widehat{Y}_i^{\text{birth}} = \mathbf{X}_i^{\text{birth}} \hat{\boldsymbol{\beta}} + \hat{\boldsymbol{\gamma}}_{c(i)}$$

for each child i observed on a California birth certificate, where $\mathbf{X}_i^{\text{birth}}$ contains harmonized covariates from the birth record and $c(i)$ denotes a child's county of birth.

Other Outcomes and Supplemental SES Index

Using predicted household income, we can assign each child to a predicted income decile and quartile based on the within-birth-record distribution.

We construct a standardized index of socioeconomic status (SES) as a weight-agnostic complement to predicted household income. The SES index combines multiple indicators observable at birth: mother's education level (1-5 scale), father's education level (1-5 scale, if available), child delivery Medicaid coverage, father listed on birth certificate, and hospital neighborhood income. Hospital neighborhood income is calculated using the average of zip code median household income for births from 2018-2019, for whom zip code of residence is available.

Each component of the SES index is standardized with mean zero and variance of one. Missing values are imputed with the component mean prior to standardization. Components reflecting lower SES (Medicaid coverage, father absent) are reverse-coded so that higher values uniformly indicate more advantaged backgrounds. The SES index is computed as the sum of the five standardized components and is itself re-standardized to have mean zero and variance one. Children are then assigned to SES index quartiles and deciles based on this composite measure.

Finally, for some descriptive tables and figures, we present mother's level of education (1-5 scale) at birth as an alternate measure when stratifying by socioeconomic status. This variable has the drawback of being relatively coarse, but the advantage of being a highly transparent measure of child SES.

BC Death Records and Potential-Maltreatment Deaths

Classification of Potential Maltreatment Deaths

This appendix documents the procedure used to classify child deaths as potential maltreatment-related fatalities. Death records from the California Death Statistical Master File (DSMF) are linked to children in the study sample. We identify potential maltreatment deaths using the underlying cause of death as recorded by ICD-10 codes, supplemented by sample restrictions on age and timing of death. Our approach is intentionally broad, capturing deaths where maltreatment, neglect, or inadequate supervision may have been a contributing factor, rather than limiting the definition to deaths explicitly coded as abuse or neglect.

ICD-10 Cause-of-Death Codes

A death is flagged as a potential maltreatment death if the underlying cause of death falls within any of the ICD-10 code categories listed in Table B.1. These categories span direct indicators of abuse and neglect (e.g., T74 for confirmed maltreatment syndromes, T76 for suspected maltreatment), injuries and trauma consistent with physical abuse (e.g., all S-codes for bodily injuries, assault codes X85–Y09), causes suggesting supervisory neglect (e.g., drowning codes W65–W74, suffocation codes W75–W83, burns X10–X19), medical neglect indicators (e.g., severe malnutrition E43–E46, dehydration E86–E87, sepsis A40–A41, pneumonia J18, and gastroenteritis K52.9), and sudden infant death syndrome (R95), which is included because a subset of SIDS deaths may be attributable to unsafe sleep environments or co-sleeping practices associated with inadequate supervision.

Additional Sample Restrictions

In addition to the cause-of-death classification, we impose two restrictions that shape the analysis of child mortality outcomes.

Age restriction (age $\leq$ 14). The sample of first-generation (G1) individuals is restricted to those who were first referred to child protective services between ages 0 and 14. This restriction ensures that G1 individuals are observed at ages young enough to be consistent across birth cohorts in the data, while also being young enough that they are unlikely to have begun having children of their own during the observation window. This constraint is important for cleanly separating G1 childhood maltreatment exposure from the onset of G1 parenting and the measurement of second-generation (G2) outcomes.

Neonatal timing restriction (30-day threshold). Deaths among second-generation children are classified by timing relative to a 30-day postnatal threshold. Deaths occurring within the first 30 days of life are categorized separately from those occurring at 30 days or later. This distinction is particularly relevant for SIDS (ICD-10 code R95). Deaths in the immediate neonatal period are more likely attributable to congenital abnormalities, complications of prematurity, or perinatal conditions rather than to parenting behaviors or environmental risk factors. By contrast, SIDS deaths occurring after the first 30 days are more plausibly linked to modifiable environmental factors such as unsafe sleep positioning, co-sleeping, or soft bedding—conditions that may reflect supervisory adequacy or parenting practices. The 30-day split thus allows the analysis to separate deaths with a stronger potential connection to parenting and household environment from those more likely driven by medical or biological causes.

Table B.1: ICD-10 Codes Used to Classify Potential Maltreatment Deaths

ICD-10 Code(s)	Category	Description
Injuries and Trauma		
S (all)	Injuries/Trauma	Injuries to specific body regions
T20–T69	Injuries/Trauma	Burns, corrosion, frostbite, and poisoning
T71	Injuries/Trauma	Asphyxiation
T73	Injuries/Trauma	Other deprivation (e.g., neglect, starvation)
T74	Injuries/Trauma	Maltreatment syndromes, confirmed
T75	Injuries/Trauma	Other external causes
T76	Injuries/Trauma	Maltreatment syndromes, suspected
Assault		
X85–X89	Assault	Assault by various means
X90–X99	Assault	Assault (continued)
Y00–Y09	Assault	Assault (continued)
Burns and Scalds		
X10–X19	Burns/Scalds	Exposure to smoke, fire, flames, or contact with heat/hot substances
Suffocation and Entrapment		
W75–W79	Suffocation	Accidental suffocation/strangulation; aspiration of food or vomit
W83	Suffocation	Other specified threats to breathing (e.g., plastic bag entrapment)
Drowning		
W65–W66	Drowning	Drowning in bathtub
W74	Drowning	Unspecified drowning
Infection / Medical Neglect		
A40–A41	Infection	Sepsis (streptococcal and other)
J18	Infection	Pneumonia, unspecified organism
K52.9	Infection	Noninfective gastroenteritis and colitis, unspecified
Malnutrition and Dehydration		
E43–E46	Malnutrition	Severe and moderate protein-calorie malnutrition
E86–E87	Dehydration	Dehydration and electrolyte imbalance
Sudden Infant Death		
R95	SIDS	Sudden infant death syndrome

Data Processing Notes

Death records are drawn from the California DSMF for the years 1999 through 2023 and are linked to children in the study sample using individual identifiers. Only injury-related external cause codes (primarily in the X, Y, W, and S/T ranges) are explicitly recorded in the California death data; the broader medical codes (A, E, J, K, R series) are included to capture deaths where neglect or inadequate care may have been a contributing factor even if the manner of death is classified as natural. The final classification yields a binary indicator for potential maltreatment death that is merged to the analytical sample and further disaggregated by timing (pre- vs. post-30 days) for the mortality analyses.