

Characteristics of the Front-Line Child Welfare Workforce

Frank Edwards¹, Christopher Wildeman²

Direct all correspondence to Frank Edwards: fedwards@cornell.edu

The analyses presented in this publication were based on restricted data from the National Child Abuse and Neglect Data System (NCANDS) Child File. These data were provided by the National Data Archive on Child Abuse and Neglect at Cornell University, and have been used with permission. The data were originally collected under the auspices of the Children's Bureau. Funding was provided by the Children's Bureau, Administration on Children, Youth and Families, Administration for Children and Families, U.S. Department of Health and Human Services. The collector of the original data, the funding agency, NDACAN, Cornell University, and the agents or employees of these institutions bear no responsibility for the analyses or interpretations presented here. The information and opinions expressed reflect solely the opinions of the authors. The authors are especially grateful to Casey Family Programs for funding this research project.

¹ Frank Edwards is a postdoctoral associate at the Bronfenbrenner Center for Translational Research at Cornell University and an incoming Assistant Professor of Criminal Justice at Rutgers University - Newark

² Christopher Wildeman is a Professor of Policy Analysis and Management (PAM) and Sociology (by courtesy) in the College of Human Ecology at Cornell University, where he is also co-director of the National Data Archive on Child Abuse and Neglect (NDACAN) and associate director of the Bronfenbrenner Center for Translational Research (BCTR). Since 2015, he has also been a Senior Researcher at the Rockwool Foundation Research Unit in Copenhagen, Denmark.

ABSTRACT

In this study, we provide new national- and state-level estimates of workload and workforce instability among child welfare agencies using previously unavailable data that includes unique identifiers for US child welfare caseworkers in 46 states and supervisors in 43 states. We identify and describe 139,921 unique caseworkers and 31,124 unique supervisors who were in the child welfare workforce between 2003 and 2015. The median caseworker handles 55 cases annually and is on the job for about 1.8 years. The median state has an annual front-line caseworker turnover rate of between 14 and 22 percent annually, and a supervisor turnover rate of about 20 percent annually.

1. INTRODUCTION

When a child welfare agency suspects that a child may be a victim of abuse or neglect, a front-line caseworker is tasked with making a series of critical, and often life-altering, decisions about that child's safety. Caseworkers and supervisors translate policy and law into interventions with children and families, doing the difficult work of protecting children from abuse and neglect and connecting vulnerable families to critical services.

Unfortunately, we lack basic comparable statistics on the organizational capacity of the child welfare workforce to respond to child maltreatment, the caseloads that workers handle, and the amount of turnover present in the workforce. While many studies have evaluated how and why front-line workers choose to leave the child welfare workforce (Benton, 2016; Fernandes, 2016; Griffiths, Royse, Culver, Piescher, & Zhang, 2017; Hyosu Kim & Kao, 2014; Truter, Fouché, & Theron, 2017), and many state agencies have commissioned studies of dynamics in their workforces (Hornby Zeller Associates, 2012; McKinsey and Company, 2008; Patel, McClure, Phillips, & Booker, 2017), few sources of data have been available to assess and compare child welfare workforces across jurisdictions and over time.

This study describes and makes available³ a new data source that researchers and policy makers can use to explore the effects of workforce characteristics on child and family outcomes, and to evaluate factors leading to changes in the workforce over time. We provide an exploratory analysis of a new set of geographically and longitudinally comparable measures of worker tenure, worker turnover, worker caseloads, and workforce capacity relative to the at-risk population for 46 states between 2003 and 2015.

³ We provide the state-level data and code used in this analysis at github.com/URL excluded for peer review

We rely on worker identifiers assigned to all screened-in cases recorded in the National Child Abuse and Neglect Data System (NCANDS) recently made available in restricted-access files. Using these identifiers, we provide annual state and national estimates of the capacity of front-line child welfare workforces, the ratio of cases per worker, the amount of time workers remain in their position, and the rate at which workers leave agencies.

1.2 Why the child welfare workforce matters

Front-line child welfare caseworkers are responsible for making decisions that directly affect child safety and family stability. However, an agency's ability to pursue both safety and permanency for children may be negatively affected when agencies are understaffed, when workers are overburdened with high caseloads, when workforces have high turnover, and when the workforce is inexperienced (Fluke, Corwin, Hollinshead, & Maher, 2016; Graef & Hill, 2000; Hansung Kim, 2011; Pietrowiak & Gambino, 2003; Strolin, McCarthy, & Caringi, 2006; N. J. Williams & Glisson, 2013). Recognizing these consequences, states have engaged in systematic efforts to evaluate and improve workforce retention (Casey Family Programs, 2017; Hornby Zeller Associates, 2012; McKinsey and Company, 2008; Patel et al., 2017).

Turnover is perhaps the most commonly used measure of workforce stability in studies of the child welfare workforce. Turnover captures the relative share of new and veteran workers within an agency over time. National estimates have suggested annual turnover rates as high as 22 percent for caseworkers and 12 percent for front-line supervisors (American Public Human Services Association, 2005), with some states reporting turnover rates for front-line caseworkers of over 30 percent (Patel et al., 2017).

This instability is consequential for children and families. High turnover is associated with poor outcomes for child safety and permanency (Pietrowiak & Gambino, 2003; Strolin-Goltzman, Kollar, & Trinkle, 2010), though effects may be moderated by the organizational culture of child

welfare agencies (N. J. Williams & Glisson, 2013). In states with over-taxed and unstable front-line child welfare workforces, investigations may linger incomplete, with potential consequences for child safety and family stability as cases are handed over to new workers (Child Welfare Information Gateway, 2016; Farber & Munson, 2010). More experienced workers tend to prefer a family preservation focused approach to casework, while workers with shorter tenures tend to prefer a child safety orientation to decision-making (Fluke et al., 2016). The frequency of worker-coordinated visits between family members for children in care may also decline, affecting the likelihood of successful reunification (Farber & Munson, 2010). High turnover also results in exceptional costs to agencies that must spend a substantial amount of time and money recruiting and training new workers (Casey Family Programs, 2017; Graef & Hill, 2000; Griffiths et al., 2017).

Workforce instability may be both a cause and consequence of excessive workloads. We know that in many places caseworkers face difficult working conditions (Cahalane & Sites, 2008; Griffiths et al., 2017; Hansung Kim, 2011; Lizano & Mor Barak, 2015; Scannapieco & Connell-Carrick, 2007; S. E. Williams, Nichols, Kirk, & Wilson, 2011), and are often assigned an unmanageably high number of cases (AFSCME, 2016; Children's Bureau, 2017; Hughes & Lay, 2012; McKinsey and Company, 2008; Patel et al., 2017; Sage, 2010). Prior research has shown that high caseloads are associated with worker burnout, lower quality relationships between workers and families, and poor outcomes for system-involved children (Casey Family Programs, 2017; Griffiths et al., 2017; Pietrowiak & Gambino, 2003; Strolin et al., 2006; Yamatani, Engel, & Spjeldnes, 2009), but studies of worker caseloads have historically been either cross-sectional in design, or longitudinal within a single jurisdiction.

While few studies have sought to compare the capacities of child welfare system workforces relative to the size of populations at-risk of child welfare system contact, we believe that new data offer an opportunity to explore how macro-level characteristics may affect worker and child outcomes. System-level dynamics related to the capacity and resource constraints facing agencies

may affect the likelihood that a victim of child abuse or neglect receives an agency response or receives services (Wulczyn & Halloran, 2017). Organizational capacity likely also helps to explain variation in worker caseloads and may be associated with child and family outcomes.

2. MEASURING CHARACTERISTICS OF THE CHILD WELFARE WORKFORCE

2.1 Data

This analysis provides a first national and longitudinal comparison of the front-line child welfare workforce using four standardized measures: staff tenure, staff turnover, caseloads, and workforce capacity. To do so, we rely on newly available restricted-access versions of the National Child Abuse and Neglect Data System (NCANDS) Child Files from 2003 through 2015 that include an anonymized unique identifier for both the caseworker and supervisor assigned to each allegation of child abuse or neglect that receives a formal response from an agency. We use these identifiers to create worker-level and state-level tables and provide a series of data visuals that allow for comparison of workforce characteristics across states.

The NCANDS child file captures information on all screened-in cases of reported child abuse and neglect at the report-child level. A single report may involve multiple children and a single child may be the subject of multiple reports. Front-line staff input case-level information into state data systems, which is then mapped onto the NCANDS submission format and reported to the US Children's Bureau for analysis and distribution. State participation is voluntary, but all fifty states have reported data between 2012 and 2016. The earliest year included in this analysis is 2003, in which 44 states participated.

In a collaboration with the Children's Bureau, the National Data Archive on Child Abuse and Neglect has made available a restricted-access version of the NCANDS child file that includes full information on the county in which an investigation was initiated, exact dates of case events, information on child deaths, and caseworker and supervisor identification numbers. In this study,

we focus on these worker identification numbers. These anonymized identification numbers are assigned in state data systems, but are not a documented element of the NCANDS child file⁴. The data contain no additional information on caseworker characteristics, such as demographics, education, or agency role.

For the purposes of this study, we assume that the assigned number identifies a unique caseworker and supervisor, that this number tracks a worker over time, and that this identifier is valid within states over time. As with most administrative data, the derived measures reported below are sensitive to the local organization and reporting practices of child welfare agencies. We detail the implications of these assumptions in the discussion.

Some states fail to record caseworker or supervisor identifiers in their NCANDS data for the full period used in this analysis. We fully lack data on New York, Ohio, Pennsylvania, and Vermont, and lack data on supervisors in Hawaii, Louisiana, and North Carolina. Some states, such as Georgia, Oregon and North Dakota only report data later in the time series. Additionally, many supervisor and worker identifiers do not appear to be valid. We describe our procedures for classifying identifiers as valid in section 2.3. From these data, we are able to identify and describe 139,921 unique caseworkers and 31,124 unique supervisors active in 46 states between 2003 and 2015. We provide a comparison of our workforce counts to state-reported figures in Appendix B.

2.2 Measures

2.2.1 Staff Tenure

We provide a measure of workforce tenure as the difference in years between the initial reporting date of a worker's last assigned case and their first assigned case. Because we lack

⁴ Our state-year counts of workers derived from these identifiers are similar to those published by the Children's Bureau in their annual Child Maltreatment reports (see, for example U.S. Department of Health and Human Services, 2017, p. 13). We compare our derived measures to these data in Appendix B.

reliable national data prior to 2003, we underestimate tenure for that group of workers whose first case date falls outside of the window included in our data. This likely induces a negative bias in our tenure estimates in the first and final years of the data. We also assume a workforce exit when a caseworker or supervisor stops carrying a caseload, though some of these workers likely remain employed within child welfare agencies. We are unable to identify workers and supervisors who remain in the workforce but are not assigned cases with the NCANDS data, which likely biases our estimates of tenure downward.

2.2.2 Staff Turnover

We define turnover as the proportion of workers who exit the focal workforce in a given year. Turnover is calculated for year t as the proportion of the set of worker ID numbers, $WORKERS$, from year $t - 1$ that are no longer present in the data in year t .

$$turnover_t = \frac{|WORKERS_{t-1} - WORKERS_t|}{|WORKERS_{t-1}|}$$

However, this turnover measure may provide too liberal an estimate of workforce instability, as some subset of caseworkers are likely to be promoted to supervisory positions year-to-year. Because we have no data on management-level child welfare workforce positions in these data, we cannot capture the proportion of supervisor exits that are internal promotions or transfers to supervisory roles that are not directly assigned to cases.⁵

To provide a conservative lower-bound on our caseworker turnover estimate, we construct an additional turnover measure that assumes exiting caseworkers must be treated as a promotion

⁵ This feature of the data likely leads to an upward bias in our estimates of supervisor turnover. Our estimates of median supervisor turnover are indeed larger than estimates previously published by the American Human Services Association (2005) and by Patel et al. in a study using Texas' child welfare workforce data (2017).

for each new supervisor ID. We count all unique IDs in the set of supervisor IDs, *SUPERVISORS*, in a state in year t that are not in the data in year $t - 1$, yielding a lower-bound on caseworker turnover as

$$\text{turnover lower bound}_t = \frac{|WORKERS_{t-1} - WORKERS_t| - |SUPERVISORS_t - SUPERVISORS_{t-1}|}{|WORKERS_{t-1}|}$$

2.2.3 Caseloads

We define caseloads as the average annual number of cases assigned to a caseworker or supervisor. We count all cases assigned to a caseworker or supervisor ID, then divide this count by the tenure (in years) of the ID number. NCANDS is structured at the report-child level; individual reports may have multiple child subjects and individual children may be subject to multiple reports. For the measures below, we reduce the data to the report-level, only including a single record for each unique report identifier. This approach may understate workloads when reports involve many children, but results in a simple and interpretable unit for measuring workloads.

2.2.4 Workforce Capacity

We define workforce capacity as the resources an organization has relative to the size of the population subject to intervention. In this case, we provide a measure of capacity as the number of front-line child welfare caseworkers in a state-year divided by the number of children in poverty in that state-year. We detail alternative approaches to measuring workforce capacity that include counts of screened-in reports as a denominator for caseworker capacity and a ratio of caseworkers to supervisors in Appendix A. More specifically, we define capacity as the count of the set of ID numbers assigned to cases, *WORKERS*, per 1,000 children in poverty in state s and year t . We obtain data on children in poverty from the US Census Bureau's Small Area Income and Poverty Estimates Program for each state and year between 2003 and 2015. This measure crudely captures

the size of the population of children likely to be at-risk of a child welfare investigation, enabling a comparison of staffing ratios as a function of the at-risk population across places. We detail alternative workforce capacity measures in Appendix A.

$$Capacity_{st} = 1000 \times \frac{|WORKERS_{st}|}{ChildPoverty_{st}}$$

2.3 Missing data and measurement error

As with all administrative data, these data both offer dramatic insight into organizational processes and present significant challenges for analysis. State data collection and reporting procedures for worker identifiers are not documented and vary in unknown ways. Our exploratory analysis suggests that some of these identifiers are not valid; some workers are assigned unbelievably high numbers of cases (suggesting that agencies may use a single identifier for all cases over a period of time), and some states report only a single ID number in a given year. Some ID numbers are also only assigned to a single case or a very small number of cases, or are assigned cases over a very short period of time. In an effort to include only those ID numbers that are assigned to actual workers and are not artifacts of the data collection process, we develop a series of exclusion criteria.

We identify a caseworker or supervisor identifier as erroneous and classify it as missing if any of the following conditions are met: the ID is the only number assigned within a state in any year, the ID is assigned cases for less than one month, the ID is assigned fewer than 5 total cases during its tenure⁶, or the ID is assigned an unrealistically high number of cases. We assume that any

⁶ This approach does likely remove actual workers with very short tenures, or who are assigned very few cases. IDs with very low tenure create instability in caseload measures, and IDs with very few cases may artificially bias caseload measures downward. While these thresholds are arbitrary, we

worker ID assigned more than 400 cases on average per year, or a supervisor ID assigned more than 4,000 cases on average per year is erroneous. Our prior beliefs about realistic caseloads come from previously published reports: the highest average annual caseload for caseworkers reported in official statistics from state agencies is about 200 (U.S. Department of Health and Human Services, 2017, p. 14). All missing ID numbers and ID numbers we classify as erroneous are excluded from the statistics and figures presented below.

3. FINDINGS

3.1 Capacity

The organizational capacity of a child welfare agency to respond to child maltreatment is, in part, described by the ratio of the size of its workforce to the size of the population that is at-risk of becoming a victim of child abuse or neglect (for alternative measures of workforce capacity, see Appendix A). As we show in Figures 1 and 2, the front-line capacity of child welfare agencies varies dramatically.

The median state has about 2.9 caseworkers per 1,000 children in poverty. States at the 25th percentile have about 2.2 caseworkers per children in poverty, and states at the 75th percentile have about 4.3 caseworkers per 1,000 children in poverty. The median state has about 0.8 supervisors per 1,000 children in poverty. A state at the 25th percentile has about 0.6 supervisors per 1,000 children in poverty, and a state at the 75th percentile has about 1.2 supervisors per 1,000 children in poverty.


Figure 1 displays our estimates of the capacity of front-line child welfare workforces by state for both caseworkers and supervisors. Points indicate the median value of workers per 1,000

believe that they strike a balance between mitigating measurement error and including real workers in the analysis.

children in poverty between 2003 and 2015 for each state in the data. There is tremendous variation in the capacity of child welfare workforces across states. Illinois, Rhode Island, and Louisiana all have a relatively low number of caseworkers per children in poverty at less than 1.5 caseworkers per 1,000 children in poverty. On the other hand, a small group of states including Connecticut, Wyoming, and Alaska each have high staffing ratios, with a median rate of more than 8 caseworkers per 1,000 children in poverty between 2003 and 2015. The ratio of front-line supervisors to children in poverty also varies dramatically across the states. Illinois, Rhode Island, Florida, Texas, and Kansas all have fewer than 0.5 supervisors per 1,000 children in poverty, while Wyoming, North Dakota, Alaska, Massachusetts, and Connecticut all have more than 2 supervisors per 1,000 children in poverty over this period.

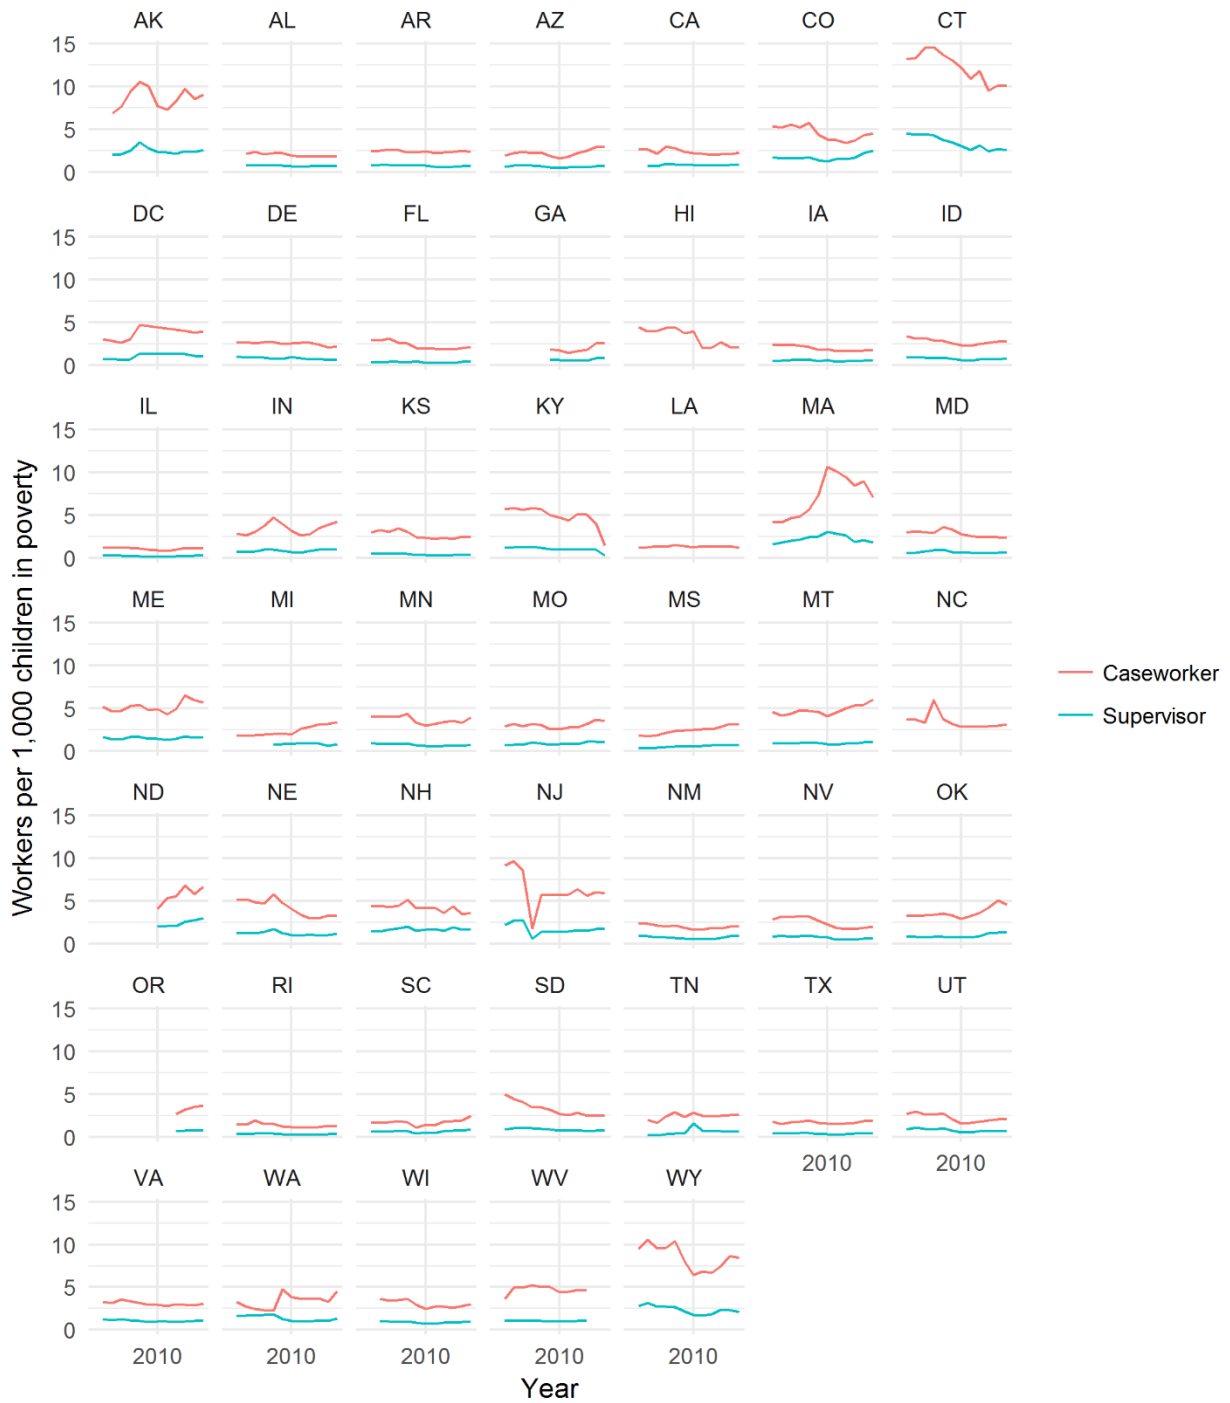
Figure 2 shows how the capacity of front-line state child welfare workforces has varied over time. The dramatic range of staffing ratios across states is again evident in this plot. Visual inspection does not reveal a clear trend in workforce capacity across states, though linear regression models including state intercepts and a linear national trend suggest a small negative national trend in workforce capacity for both caseworkers and supervisors. It is important to note that child poverty rates increased precipitously for most states during this period, which included the recession beginning in 2008. This change in the denominator would result in declining workforce capacity even in the context of a constant workforce.

Figure 1. Front-line caseworkers and supervisors per 1,000 children in poverty. Median values by state, 2003 – 2015.



Note: Hawaii and Louisiana did not report supervisor identifiers during this period

Figure 2. Front-line caseworkers and supervisors per 1,000 children in poverty. Annual values by state, 2003 – 2015.



Note: Hawaii and Louisiana did not report supervisor identifiers during this period

3.2 Caseloads

The median caseworker in the United States handles 55 cases per year. Caseworkers on the left-hand side of the distribution have very low annual caseloads, while many workers on the right-hand side of the distribution have exceptionally high annual caseloads. Workers at the 25th percentile of the distribution have an annual caseload of about 19, while caseworkers at the 75th percentile of the data have an annual caseload 5 times larger, at 97 cases per year. A small proportion of workers have exceptionally high caseloads; 10 percent of caseworkers handle more than 130 cases per year. Figure 3 shows the distribution of workloads for all workers and supervisors, and Figure 4 illustrates the distribution of median workloads across states.

The median supervisor in the US oversees 88 cases per year, and their workloads are far more variable than are the workloads for caseworkers. On the left-hand side of the distribution in Figure 3 (bottom panel), supervisors at the 25th percentile of the data handle about 25 cases per year, while on the right-hand side of the distribution, supervisors at the 75th percentile of the data handle at least 238 cases per year. At the 90th percentile, supervisors are handling over 435 cases per year.

Figure 4 shows the distribution of median worker caseloads across states between 2003 and 2015. A cluster of states have relatively low caseloads. Massachusetts caseworkers handled a median of 9 cases per year (50 percent interval [5, 21]), Connecticut workers handled 14 [8, 22] cases per year, and Wyoming workers handled a median of 16 [9, 28] cases per year. A handful of states had a median of more than 100 cases per year per caseworker. Illinois [87, 128], Florida [86, 136], Utah [65, 144], Rhode Island [84, 153], and Iowa [115, 147] all had exceptionally high caseloads over this period for caseworkers. For supervisors, Massachusetts, California, North Dakota, and Maine had a low median supervisor caseload, at below 30 cases per year. Rhode Island, Iowa, Illinois, and Florida had median supervisor caseloads of over 300 cases per year over this period, with Florida at the highest median value of 466 [282, 593] cases per year.

Figure 3. Cases per year assigned to caseworkers and supervisors, by worker ID 2003 – 2015.

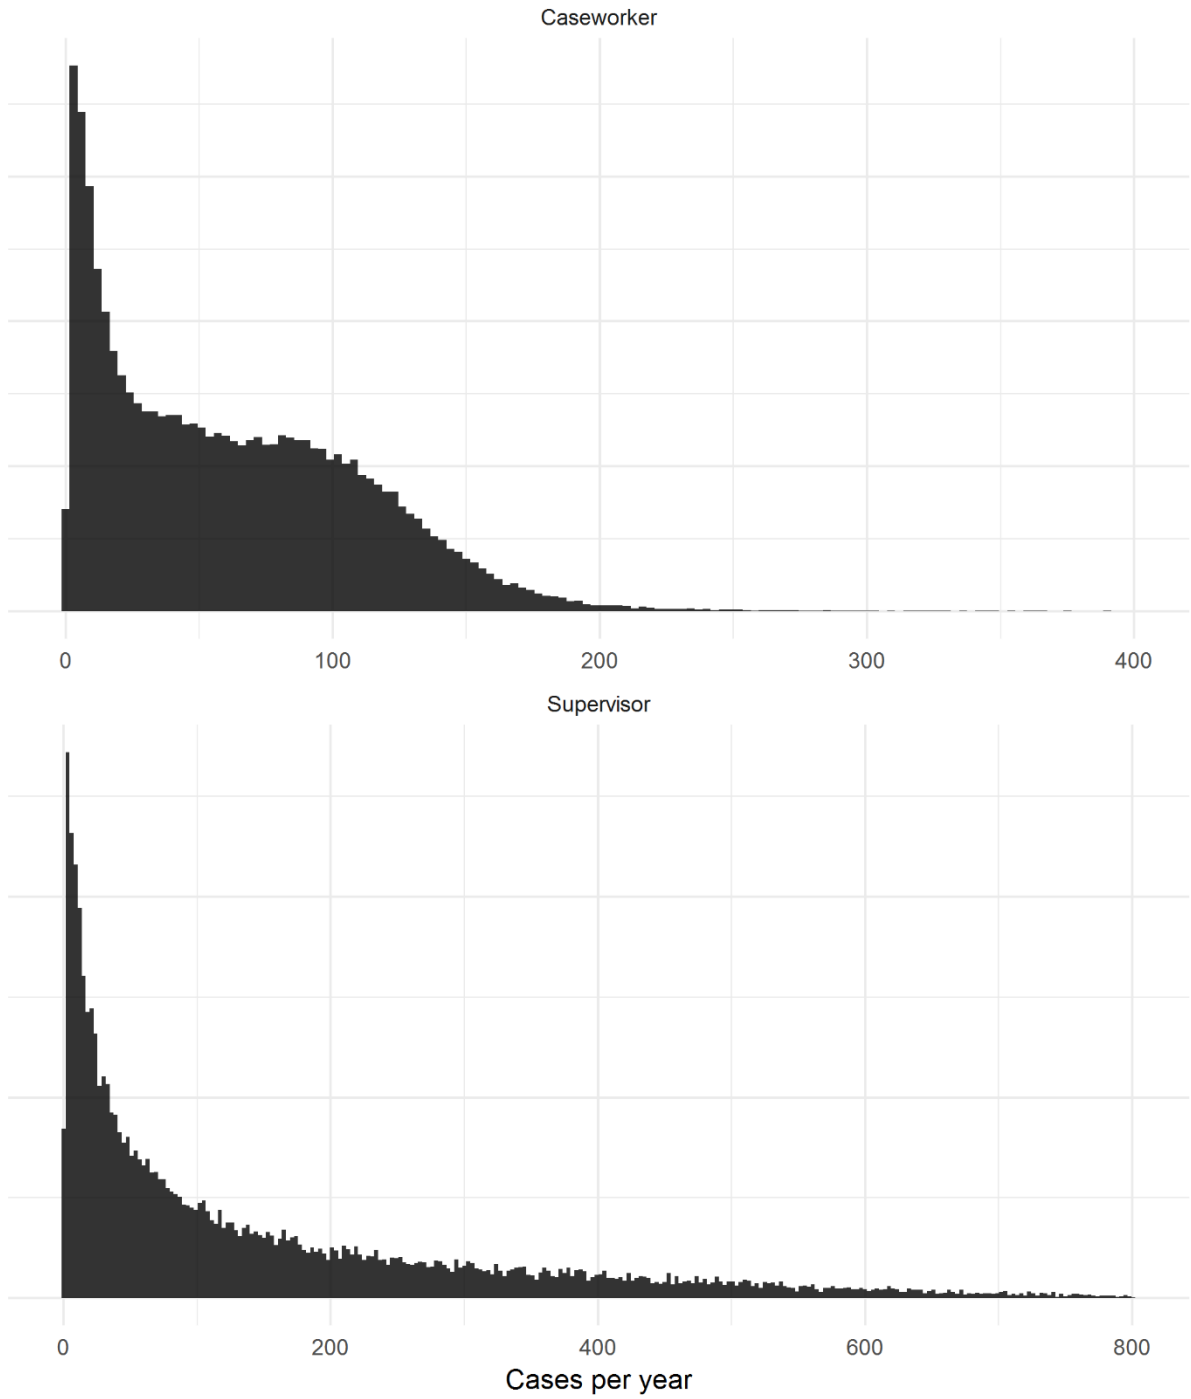



Figure 4. Median cases per year assigned to caseworkers and supervisors by state, 2003 – 2015.



Note: Hawaii and Louisiana did not report supervisor identifiers during this period

3.3 Tenure

Figure 5 displays the distribution of our estimates of worker tenure for all workers with a valid set of cases in the NCANDS data between 2003 and 2015. Note that supervisors in the data tend to have a longer tenure than do caseworkers. We find that the median caseworker in the US has about 1.8 years between their first and last assigned case, with workers at the 75th percentile having at least 4 years of experience, and workers at the 25th percentile having at least 0.8 years of experience. We find that the median supervisor in the US has about 2.5 years between their first and last assigned cases, with a fifty percent interval including long right tails [1.2, 5.3].

Figure 6 shows median worker tenure across all states with valid data. Only Florida has a median worker tenure of less than one year, though a large group of states have a relatively low median worker tenure. Illinois, Texas, North Carolina, Missouri, South Carolina, Tennessee, Georgia, Louisiana, Michigan, Oklahoma, Hawaii, Oregon, and West Virginia all have a median worker tenure of less than 1.8 years. Massachusetts, California, Connecticut, and Rhode Island all have a median caseworker tenure of greater than 4 years. Kansas, Illinois, Michigan, and Tennessee all have a median supervisor tenure of less than 1.5 years. On the other side of the distribution, the median supervisor in California, Connecticut, and Rhode Island has at least 8 years of tenure.

Figure 5. Length of time between first and last assigned case (tenure) for front-line workers, by worker ID 2003 – 2015.





Figure 6. Length of time between first and last assigned case (tenure) for front-line workers. Median values by state, 2003 – 2015.



Note: Hawaii and Louisiana did not report supervisor identifiers during this period

3.4 Turnover

Turnover, the proportion of workers in a prior year that persist in the workforce in a focal year, provides a direct measure of workforce stability. We illustrate rankings and trends in front-line child welfare workforce turnover in Figures 7 and 8.

Figure 7 displays the median proportion of workers who exited the ID pool year-to-year between 2004 and 2015 for supervisors and caseworkers. The dot indicates the total proportion of exiting workers (our upper bound estimate), and the line indicates bounds on turnover assuming some caseworkers were promoted to supervisors, given the count of new supervisors and the count of caseworker exits in year t . Note that this estimated range is substantial in some states, and in many state-years, a maximum promotion assumption results in an estimate of zero turnover. For those states where we lack data on supervisors, we are unable to estimate an interval based on promotion assumptions.


We estimate that a median state has an annual turnover rate of between 0.14 and 0.22 for its caseworkers, and about 0.20 for supervisors. As with other measures, however, states vary dramatically in the stability of their workforces. We find that Iowa, Rhode Island, Illinois, and Washington DC all have a median worker turnover rate of less than 0.15 (upper bound). We find that 17 states have annual worker turnovers (upper bound) of greater than 0.25. Nebraska and South Dakota have exceptionally high caseworker turnover, with an upper bound over 0.3 and lower bounds greater than 0.25. We find that Connecticut, Rhode Island, New Jersey, New Hampshire, and California have relatively stable supervisor workforces, with turnover of less than 0.15 annually. Texas⁷ and Georgia had unstable supervisory workforces over this period, with a median turnover greater than 0.3.

⁷ Though our estimates for Texas caseworker turnover are similar to estimates published by Patel et al. (2017, p. 11), based on data from Texas' Department of Family and Protective Services, our estimates of supervisor turnover in Texas diverge sharply. They estimate a 2015 supervisor turnover

Figure 8 displays changes over time in child welfare workforce turnover across states. Caseworker turnover includes our estimated upper and lower bounds for turnover based on supervisor promotion assumptions. We can see substantial heterogeneity in turnover across states and within states over time, but visual inspection does not reveal a consistent set of temporal trends across states. Ordinary least squares models of caseworker turnover with state intercepts and a national linear time trend suggest a small positive national trend in caseworker turnover rates, and a small positive national trend in supervisor turnover rates.

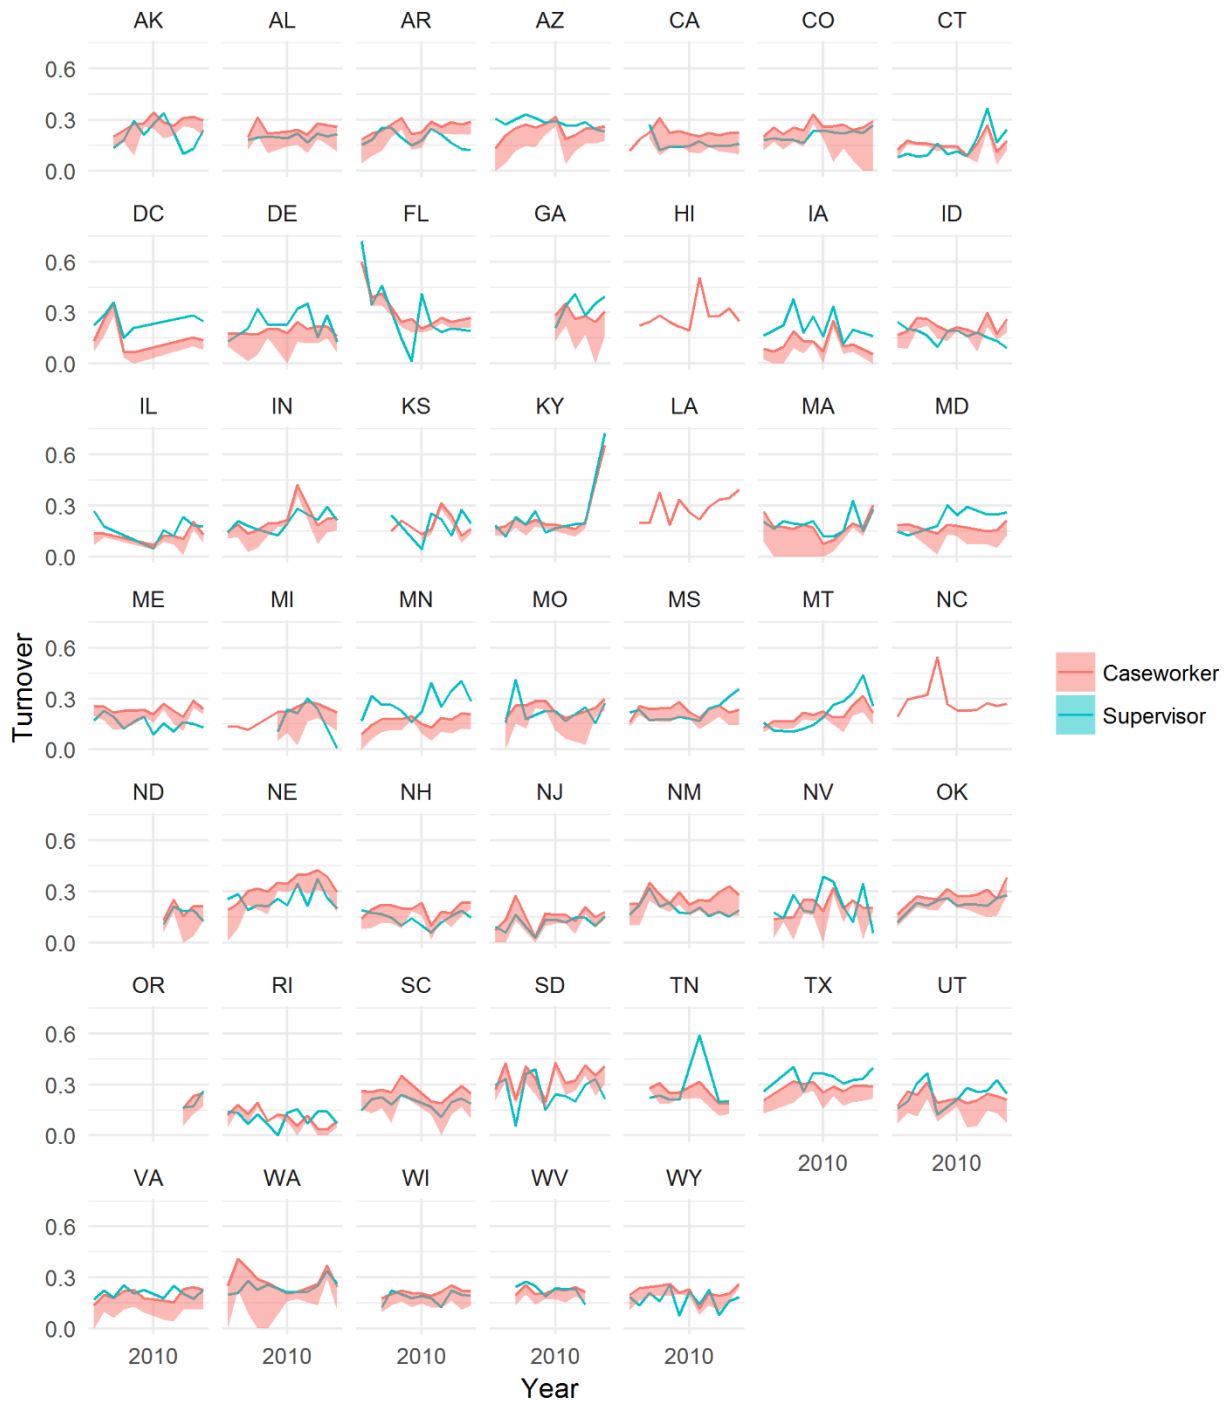
rate of about 0.1, while we estimate a supervisor turnover of nearly 0.4. This may be due to differences in the classification of supervisors, or the movement of supervisors away from investigations, at which point they exit our set of included supervisors.

Figure 7. Proportion of workers exiting the workforce annually (turnover). Points indicate all exits, lines indicate possible promotions to supervisors. Median values by state, 2004 – 2015.



Note: Hawaii and Louisiana did not report supervisor identifiers during this period

Figure 8. Proportion of workers exiting the workforce annually. Solid lines indicate all exits, shaded area indicates possible promotions to supervisors. Annual values by state, 2004 – 2015.



Note: Hawaii and Louisiana did not report supervisor identifiers during this period

4. DISCUSSION

We use previously unreleased data on caseworkers and supervisors from 46 states between 2003 and 2015 to identify variation in the capacity, experience, caseloads, and turnover rates of front-line child welfare workforces. We find that the median state has about 2.9 front-line caseworkers per 1,000 children in poverty. The median caseworker handles 55 cases per year, though this figure varies dramatically across workers; caseworkers above the 90th percentile of caseloads handle more than 130 cases per year. The median caseworker is on the job for about 1.8 years, and the median supervisor remains on the job for about 2.5 years. We estimate that the median state has an annual caseworker turnover rate of between 14 and 22 percent of its workforce. Prior point-in-time estimates have suggested turnover rates of around 20 percent (American Public Human Services Association, 2005).

On each of the dimensions we consider, the data show tremendous variation both across caseworkers and supervisors (within states) and between states in the composition and experiences of the front-line child welfare workforce. We find that some caseworkers and supervisors are handling exceptionally high caseloads. We identified over 13,000 caseworkers nationally (10 percent of those included in this study) who have handled, on average, more than 130 cases per year. We also find that the front-line child welfare workforce is unstable. Half of workers in this study were on the job for less than 1.8 years, and half of supervisors were on the job for less than 2.5 years. Turnover in the workforce is relatively high; about one in five workers and supervisors exit the front-line workforce year-to-year.

Although the findings from this specific analysis are informative, the possibilities the data offer for future research and policy evaluation are substantial. By offering this comparative description of the workforce data and providing state-level data for researchers and practitioners, we hope to spur new research on the causes of workforce instability and on the consequences of the workforce for child safety and family well-being.

For example, these data provide a new opportunity to explore whether workforce stability leads to higher or lower rates of repeated maltreatment, or whether workforce characteristics contribute to other indicators of child safety and well-being such as the death rate due to injuries for young children in a state. They make evaluations of the relationships between front-line workforce stability and case-level decision making (such as substantiation or service delivery) across states relatively straightforward. They also provide a compelling new opportunity to explore how and why workforce stability varies across places and within places over time. Using these data, policy makers and researchers can evaluate the efficacy of focused retention and recruitment efforts, how pay and benefits affect staff turnover, or how shocks to workloads affect workforce stability.

Of course, these data are imperfect. We lack information on salient demographic characteristics of the workforce, and lack information on the training and external experience of workers. Finer-grained data at the local level will be necessary to explore how these characteristics of the workforce matter for child protection. Because agency structures and policies (in addition to data-reporting practices) vary dramatically across states, there is likely geographic heterogeneity in the information included in worker identifiers. The kinds of workers assigned by identifier to cases in the NCANDS data may differ across places. The data may include investigations workers or workers overseeing ongoing services depending on agency practices. And while we have tried to identify invalid identifiers in exploratory analyses, more work is needed to validate these data (see Appendix B for comparisons to state-provided employment counts). Though professional organizations and unions have issued clear guidelines for worker caseloads (AFSCME, 2016; American Public Human Services Association, 2005), because NCANDS lacks data on the date of case closure or worker roles, we are unable to measure concurrent caseloads relative to worker duties. We are only able to provide cumulative caseloads per worker.

We urge caution in drawing conclusions about policy efficacy or needed changes in workforce practices directly from the descriptive analyses in this paper. However, these data do offer a new opportunity to rigorously analyze the causes and consequences of workforce instability both across jurisdictions and within agencies over time that are likely to have significant value for child welfare policy and practice. While the rank ordering of state workforce characteristics are informative, more work needs to be done to externally validate these data and to explore how these workforce characteristics relate to outcomes for children and youth.

Limitations aside, however, these data offer a unique opportunity for expanding knowledge on how the composition of front-line child welfare workers affects the health and wellbeing of both the children they seek to protect and the workers themselves.

REFERENCES

- AFSCME. (2016). *Cost Savings from Reasonable Child Welfare Workloads*. Washington, DC: AFSCME Department of Research and Collective Bargaining Services. Retrieved from <http://afscmestaff.org/wp-content/uploads/2016/12/Cost-Savings-from-Lower-Caseloads-FINAL.pdf>
- American Public Human Services Association. (2005). *Report from the 2004 child welfare workforce survey: State agency findings*. Washington, DC. Retrieved from <http://www.theprofessionalmatrix.com/docs/WorkforceReport2005.pdf>
- Benton, A. D. (2016). Understanding the diverging paths of stayers and leavers: An examination of factors predicting worker retention. *Children and Youth Services Review*, 65, 70–77. <https://doi.org/10.1016/j.childyouth.2016.04.006>
- Cahalane, H., & Sites, E. W. (2008). The climate of child welfare employee retention. *Child Welfare*, 87(1), 91–114.
- Casey Family Programs. (2017). *Healthy Organizations*. Seattle, WA: Casey Family Programs. Retrieved from https://caseyfamilypro-wpengine.netdna-ssl.com/media/HO_Turnover-Costs_and_Retention_Strategies.pdf
- Child Welfare Information Gateway. (2016). *Caseload and Workload Management*. Washington, DC: Children's Bureau. Retrieved from https://www.childwelfare.gov/pubPDFs/case_work_management.pdf
- Children's Bureau. (2017). *Child and Family Service Review: Round 3 Findings*. Washington, DC: Administration for Children and Families. Retrieved from https://www.acf.hhs.gov/sites/default/files/cb/cfsr_round3_findings_2015_2016.pdf

- Farber, J., & Munson, S. (2010). Strengthening the Child Welfare Workforce: Lessons from Litigation. *Journal of Public Child Welfare*, 4(2), 132–157.
<https://doi.org/10.1080/15548731003799340>
- Fernandes, G. M. (2016). Organizational climate and child welfare workers' degree of intent to leave the job: Evidence from New York. *Children and Youth Services Review*, 60, 80–87.
<https://doi.org/10.1016/j.chilyouth.2015.11.010>
- Fluke, J. D., Corwin, T. W., Hollinshead, D. M., & Maher, E. J. (2016). Family preservation or child safety? Associations between child welfare workers' experience, position, and perspectives. *Children and Youth Services Review*, 69, 210–218.
<https://doi.org/10.1016/j.chilyouth.2016.08.012>
- Graef, M. I., & Hill, E. L. (2000). Costing child protective services staff turnover. *Child Welfare*, 79(5), 517–533.
- Griffiths, A., Royse, D., Culver, K., Piescher, K., & Zhang, Y. (2017). Who stays, who goes, who knows? A state-wide survey of child welfare workers. *Children and Youth Services Review*, 77, 110–117. <https://doi.org/10.1016/j.chilyouth.2017.04.012>
- Hornby Zeller Associates. (2012). *Workload Study: Prepared for Alaska Office of Children's Services*. Troy, NY: Alaska Office of Children's Services. Retrieved from <http://dhss.alaska.gov/ocs/Documents/Publications/pdf/HZA-workload-study-2012.pdf>
- Hughes, S., & Lay, S. (2012). *Direct Services Workers' Recommendations for Child Welfare Financing Reform*. Washington, DC: Child Welfare League of America. Retrieved from <http://www.cwla.org/wp-content/uploads/2014/05/DirectServiceWEB.pdf>
- Kim, Hansung. (2011). Job conditions, unmet expectations, and burnout in public child welfare workers: How different from other social workers? *Children and Youth Services Review*, 33(2), 358–367. <https://doi.org/10.1016/j.chilyouth.2010.10.001>

- Kim, Hyosu, & Kao, D. (2014). A meta-analysis of turnover intention predictors among U.S. child welfare workers. *Children and Youth Services Review*, *47*, 214–223.
<https://doi.org/10.1016/j.chilyouth.2014.09.015>
- Lizano, E. L., & Mor Barak, M. (2015). Job burnout and affective wellbeing: A longitudinal study of burnout and job satisfaction among public child welfare workers. *Children and Youth Services Review*, *55*, 18–28. <https://doi.org/10.1016/j.chilyouth.2015.05.005>
- McKinsey and Company. (2008). *Child Welfare Staffing Study*. Salem, OR: State of Oregon Department of Human Services. Retrieved from
<http://www.thesoutherninstitute.org/docs/publications/ORChildWelfareStaffingStudy08.pdf>
- Patel, D., McClure, M., Phillips, S., & Booker, D. (2017). *Child Protective Services Workforce Analysis and Recommendations*. Dallas, TX: TexProtects: The Texas Association for the Protection of Children. Retrieved from
http://ncwwi.org/files/Retention/Child_Protective_Services_Workforce_Analysis_and_Recommendations.pdf
- Pietrowiak, D., & Gambino, J. (2003). *HHS Could Play a Greater Role in Helping Child Welfare Agencies Recruit and Retain Staff* (Child Welfare No. GAO-03-357). Washington, DC: United States General Accounting Office. Retrieved from
<https://www.gao.gov/assets/240/237373.pdf>
- Sage, M. D. (2010). *Child welfare workforce turnover: Frontline workers' experiences with organizational culture and climate, and implications for organizational practice* (Dissertation). Portland State University, Portland, OR.
- Scannapieco, M., & Connell-Carrick, K. (2007). Child welfare workplace: The state of the workforce and strategies to improve retention. *Child Welfare*, *86*(6), 31–52.

- Strolin, J. S., McCarthy, M., & Caringi, J. (2006). Causes and effects of child welfare workforce turnover: Current state of knowledge and future directions. *Journal of Public Child Welfare*, 1(2), 29–52.
- Strolin-Goltzman, J., Kollar, S., & Trinkle, J. (2010). Listening to the voices of children in foster care: Youths speak out about child welfare workforce turnover and selection. *Social Work*, 55(1), 47–53.
- Truter, E., Fouché, A., & Theron, L. (2017). The Resilience of Child Protection Social Workers: Are They at Risk and If So, How Do They Adjust? A Systematic Meta-Synthesis. *The British Journal of Social Work*, 47(3), 846–863. <https://doi.org/10.1093/bjsw/bcw048>
- U.S. Department of Health and Human Services. (2017). *Child Maltreatment 2015*. Washington, DC: U.S. Department of Health and Human Services, Administration for Children and Families, Administration on Children, Youth and Families, Children’s Bureau. Retrieved from <http://www.acf.hhs.gov/programs/cb/research-data-technology/statistics-research/child-maltreatment>
- Wildeman, Emanuel, Leventhal JM, Putnam-Hornstein E, Waldfogel J, & Lee, H. (2014). The prevalence of confirmed maltreatment among us children, 2004 to 2011. *JAMA Pediatrics*, 168(8), 706–713. <https://doi.org/10.1001/jamapediatrics.2014.410>
- Williams, N. J., & Glisson, C. (2013). Reducing turnover is not enough: The need for proficient organizational cultures to support positive youth outcomes in child welfare. *Children and Youth Services Review*, 35(11), 1871–1877. <https://doi.org/10.1016/j.chilyouth.2013.09.002>
- Williams, S. E., Nichols, Q. I., Kirk, A., & Wilson, T. (2011). A recent look at the factors influencing workforce retention in public child welfare. *Children and Youth Services Review*, 33(1), 157–160. <https://doi.org/10.1016/j.chilyouth.2010.08.028>

Wulczyn, F., & Halloran, J. (2017). Foster Care Dynamics and System Science: Implications for Research and Policy. *International Journal of Environmental Research and Public Health*, 14(10), 1181. <https://doi.org/10.3390/ijerph14101181>

Yamatani, H., Engel, R., & Spjeldnes, S. (2009). Child Welfare Worker Caseload: What's Just Right? *Social Work*, 54(4), 361–368. <https://doi.org/10.1093/sw/54.4.361>

APPENDIX A. ALTERNATIVE WORKFORCE CAPACITY MEASURES

Child poverty at the state-level is a crude measure of child well-being and risk of child maltreatment. The size of the population of children at-risk of becoming subject to a child abuse and neglect report is not easily obtainable from official statistics, though efforts have been made to quantify these risks through demographic models (Wildeman et al., 2014). One simple alternative metric to the measures presented in the main analysis – workers per 1,000 children in poverty – is to calculate ratios of workers relative to the volume of various child welfare indicators generated from the NCANDS Child File.

In Figure A1, we compare four caseworker workforce capacity measures: caseworkers per child poverty, caseworkers per screened-in maltreatment report, caseworkers per report (screened-out and screened-in), and caseworkers per confirmed maltreatment victim. With the exception of the victimization ratio, all measures are relatively strongly correlated ($r > 0.7$). The densities of each of these measures (provided on the diagonal) have similar distributions.

In Figure A2, we compare five supervisor workforce capacity measures: supervisors per child poverty, supervisors per screened-in maltreatment report, supervisors per report (screened-out and screened-in), supervisors per confirmed victim, and supervisors per caseworker. The measure reported in the main analysis, supervisors per children in poverty, is strongly correlated with supervisors per screened-in report, with a slightly weaker correlation with supervisors per report. Supervisors per child poverty is correlated at about $r = 0.5$ with supervisors per confirmed victim. The correlation between a supervisors per caseworker capacity measure and either poverty or maltreatment indicator based capacity measures are somewhat weaker ($0.35 \leq r \leq 0.45$).

We encourage analysts to explore these and other alternative workforce capacity measures when evaluating child welfare systems across places and within-places over time. For example, while we find that Alaska and Wyoming have high caseworker to child poverty ratios, the incredibly

low population density and large land-area of these states likely necessitate higher staffing ratios to provide services that may be equivalent in quality to places with higher population density.

Figure A1. Comparison of caseworker workforce capacity measures. Bivariate scatterplots below the diagonal, variable densities on the diagonal, bivariate correlations above the diagonal.

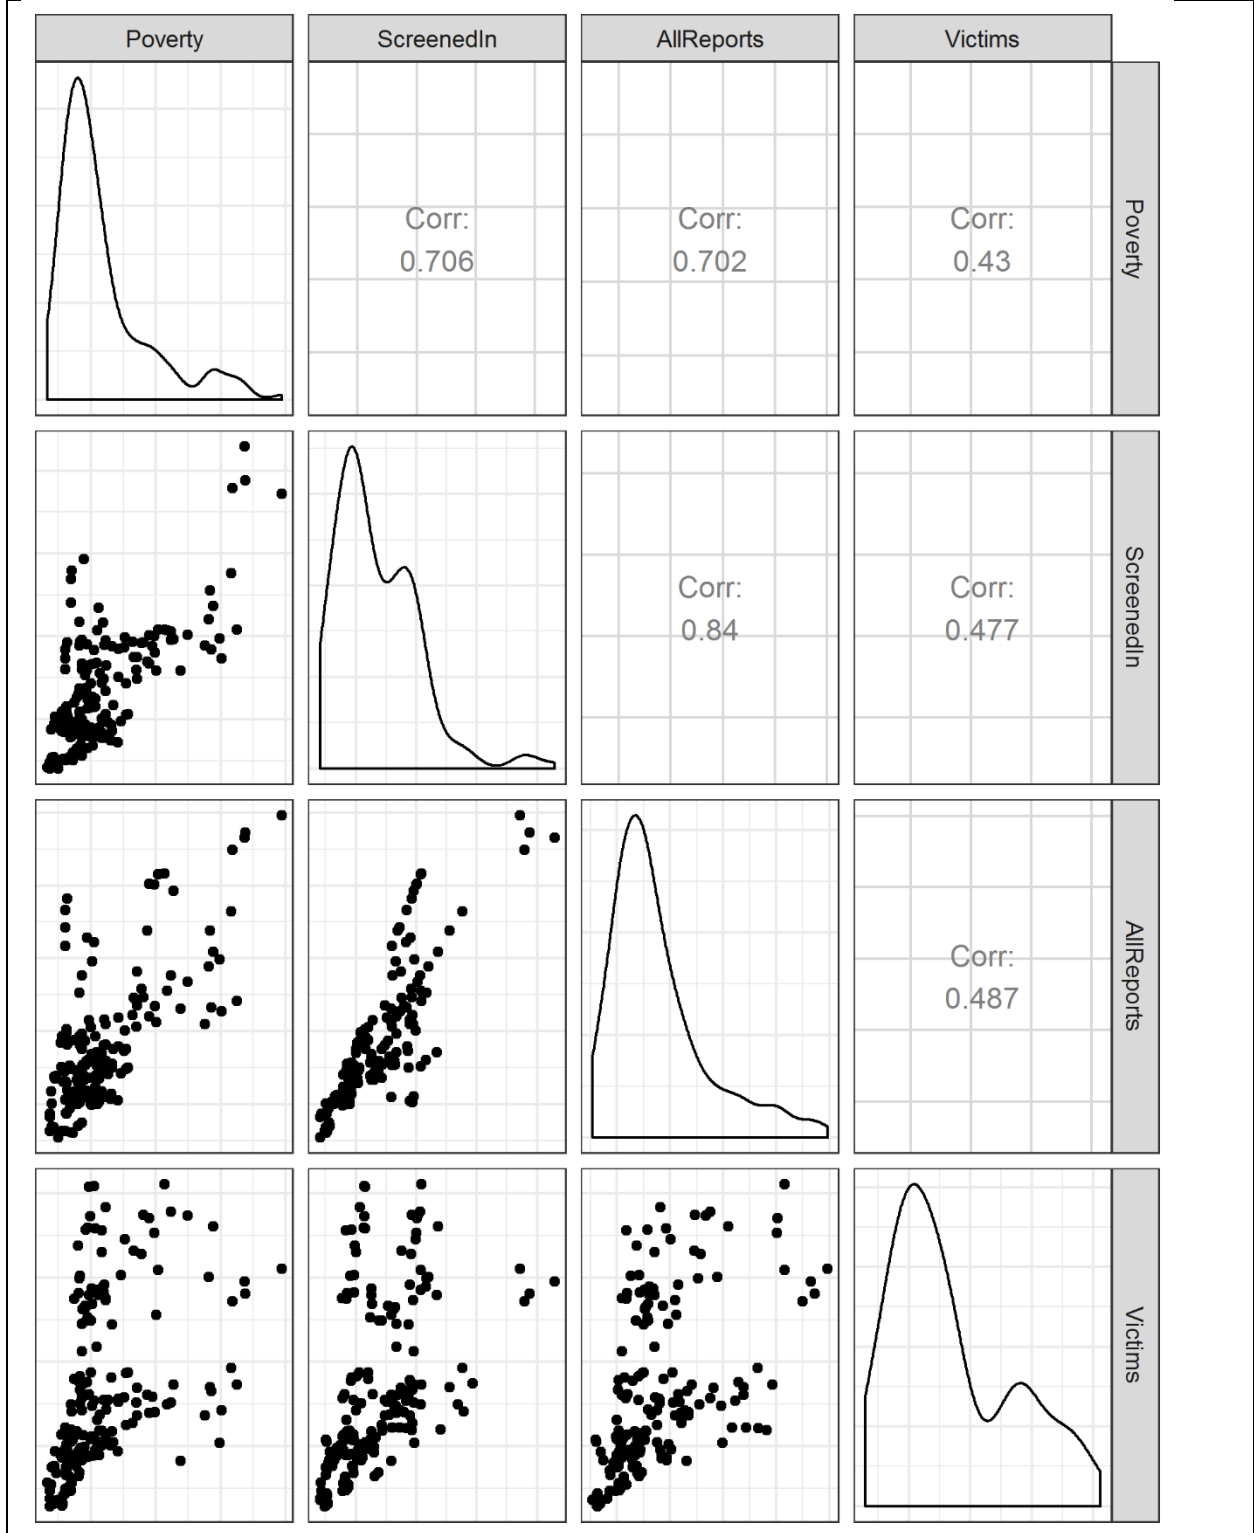
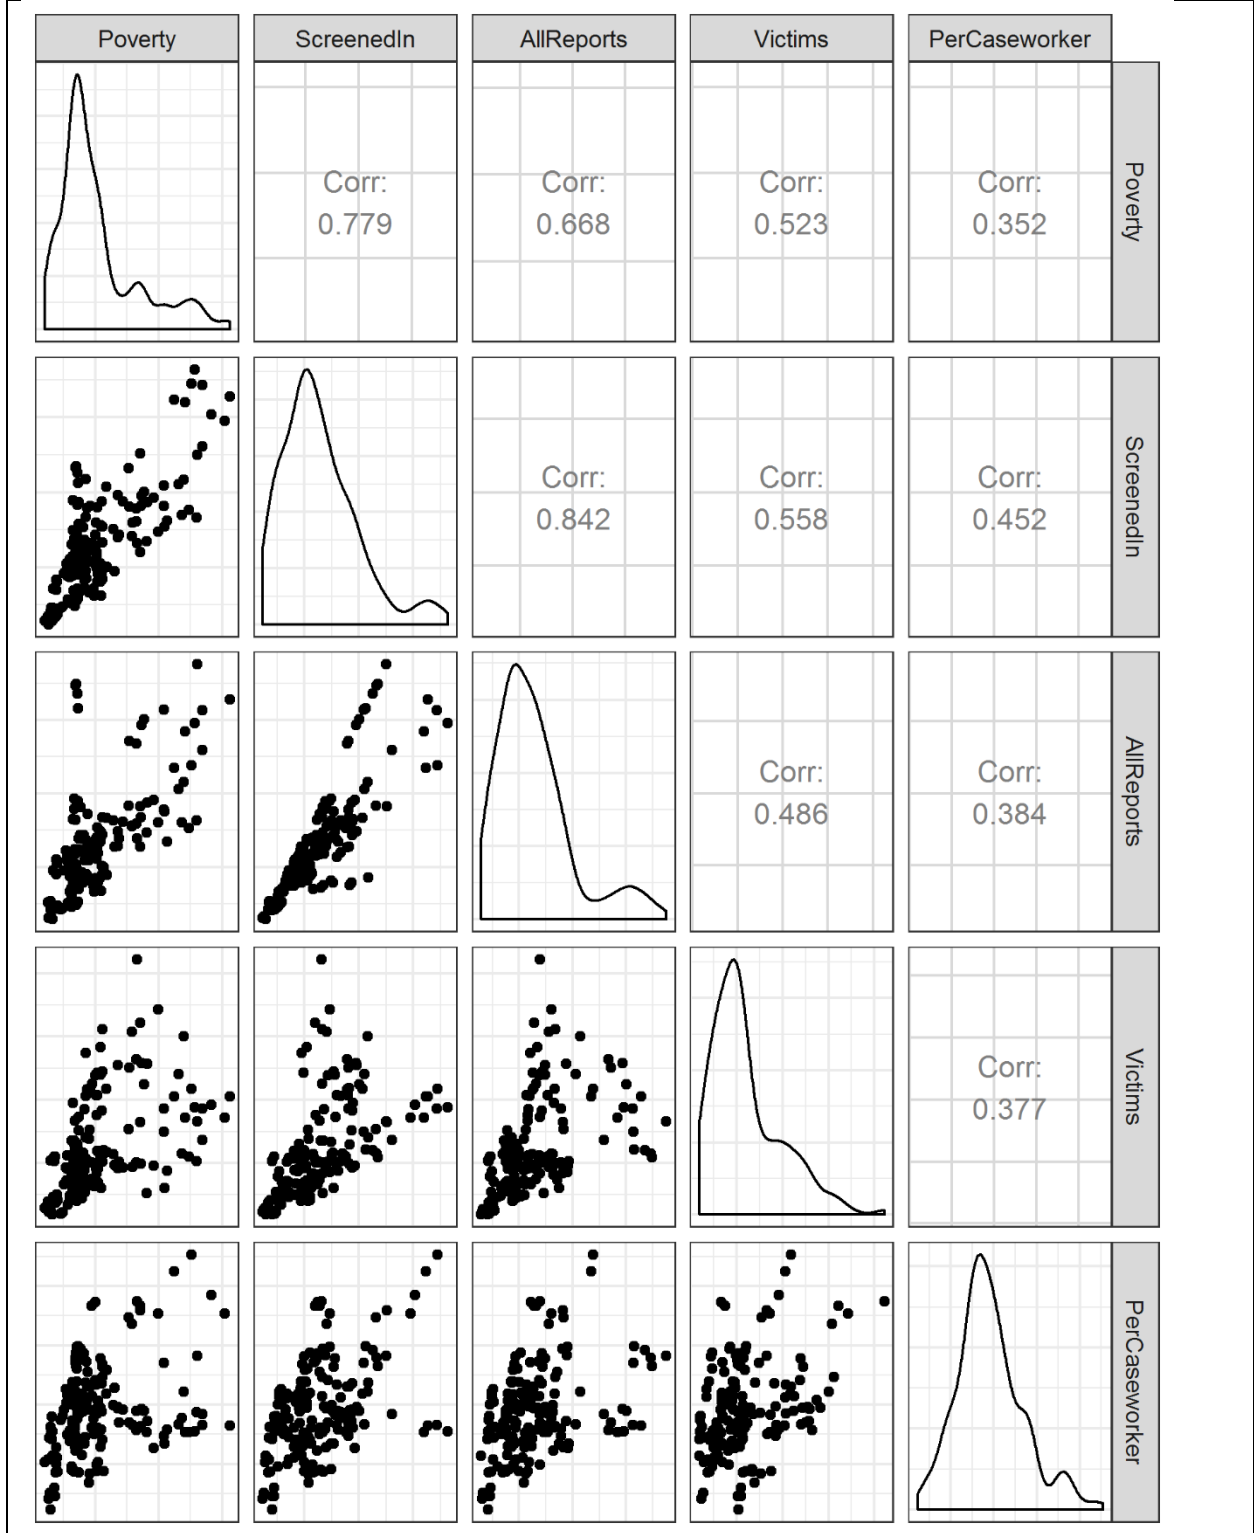


Figure A2. Comparison of supervisor workforce capacity measures. Bivariate scatterplots below the diagonal, variable densities on the diagonal, bivariate correlations above the diagonal.



APPENDIX B. COMPARISON TO NCANDS AGENCY FILE STAFFING DATA

While there are few available national statistics published on the child welfare workforce, one component of the NCANDS, the Agency File, includes some information on state-level staffing. This file includes two variables related to staffing: a count of the number of staff responsible for CPS functions (screening, intake, and investigation/assessment of reports) during the year, and a second variable that counts the number of staff responsible for the screening and intake of reports during the year. While these variables do not match perfectly onto our anonymous ID derived count – our variable includes only workers assigned to cases and our variable is inclusive of all workers assigned to a case during a year while these are likely point-in-time measures – they do offer an opportunity to validate our measures against alternative data that should be correlated with the population of interest.

In Figure B1, we plot the NCANDS Agency File state-year level measure of screening, intake, and investigations workers against state-year counts of caseworkers derived from our analysis of NCANDS Child File worker identifiers, along with a $y = x$ line to evaluate whether our measure reports more or fewer workers than the state-reported numbers. Of the 167 observations with valid data for both measures between 2012 and 2015, our data counts more caseworkers in 110 of the cases than are reported in the Agency File. In 57 cases, we identify fewer workers than are reported in the Agency File. The two measures are strongly correlated ($r = 0.93$). Our measure may be over-reporting workers relative to the Agency File in some places, as we provide a cumulative total of workers assigned to cases over the year, rather than a point-in-time employment count. Our measure may be under-reporting workers relative to the Agency File in some places because the universe of workers included in the Agency File count is inclusive of intake and screening workers, who presumably are not assigned by ID number as a caseworker in the NCANDS Child File.

Figure B1. Scatterplot of NCANDS identifier derived state-year caseworker employment counts and state-reported counts of screening, intake and investigation workers in the NCANDS Agency File, 2012 – 2015. Plotted line indicates equal values. Note log base 10 scale on axes.

