



Stop Predicting the Breaking Point

A QIC Take on Workforce Topics in Child Welfare

Why AI People Analytics is Solving the Wrong Problem in Child Welfare

Can a machine-learning model tell us which caseworker is about to quit before they know it themselves? And if it could, would knowing actually help? I keep running into that question at the intersection of two trends I work on closely: the spread of AI into public human services, and the slow-burn workforce crisis in child welfare. The answer I keep arriving at is not the one the vendors are selling.

Public child welfare agencies lose a sizable share of their frontline workforce every year, with turnover driven by a well-documented mix of caseload, supervision, organizational factors, and compensation ([Elgin et al., 2025](#)). Faced with those numbers, it is entirely reasonable that agency leaders are looking at every tool in reach. And the pitch from AI vendors is genuinely alluring: ingest the HR and engagement data, train a model, flag the caseworkers most likely to leave, and intervene before they hand in their notice.

I want to argue, carefully, that this is the wrong place to point the technology. The tools are not the problem. Their target is. We are building algorithmic early-warning systems for workforce trauma while leaving the conditions that produce that trauma untouched.

The Techno-Solutionism Trap

There is real enthusiasm for bringing predictive modeling into social-work HR. Researchers are fitting models on engagement and exit-survey data to identify “predictive factors” of caseworker turnover, and a growing number of state agencies are piloting people-analytics platforms with AI layers bolted on top. These efforts sound innovative on paper.

The difficulty is that the peer-reviewed literature has already told us, with a fair amount of consistency, what drives child welfare turnover. Kim and Kao’s meta-analysis ([Kim & Kao, 2014](#)) pointed squarely at caseload, supervisory support, organizational commitment, and compensation. Barbee and colleagues ([Barbee et al., 2018](#)) reached compatible conclusions using intent-to-leave measures. We do not have a prediction problem. We have a resourcing problem.

Focusing AI on predicting who will leave shifts the burden of attrition onto the individual worker rather than the institution. If an agency needs a model to tell it that a caseworker carrying thirty high-acuity trauma cases is a flight risk, the agency does not need better analytics. It needs fewer cases per worker. People analytics in this context risks becoming a tool for symptom management — offering the illusion of control while systemic underfunding, low pay, and unmanageable caseloads remain untouched.

Quantifying the Unquantifiable

Corporate people-analytics tools typically track measurable digital outputs: emails sent, calendar density, keystroke cadence, task-completion rates. When those tools are ported into human services, the reliance on “interaction analytics” becomes actively distorting.

Child welfare is a deeply relational, emotionally taxing field. You cannot capture the labor of sitting with a grieving foster child or de-escalating a domestic crisis at a family’s kitchen table, through keystrokes and case-note turnaround times. And there is a well-documented observer effect at play: Bernstein’s [transparency paradox](#) research shows that workers visibly reshape their behavior to match what observational systems measure — and that the measurement itself changes the underlying work. In human services, that dynamic creates perverse incentives. An efficiency dashboard may flag the caseworker who spent two unlogged hours sitting with a traumatized teenager as “unproductive.” The system ends up algorithmically punishing the exact human-centered behaviors the profession depends on.

Systemic Error, Surveillance, and the Erosion of Professional Autonomy

There is also the question of what these systems do to the people they watch. Leicht-Deobald and colleagues ([2019](#)) lay out the mechanisms in careful detail: people-analytics platforms can encode historical bias, narrow the definition of good work, and shift power toward whoever controls the dashboard. Ajunwa, Crawford, and Schultz ([2017](#)) made the parallel legal argument for workplace surveillance more broadly. Child welfare should not have to learn these lessons twice.

In fact, the field has already seen what happens when algorithmic risk scoring enters consequential professional judgment. Chouldechova and colleagues’ analysis of the Allegheny Family Screening Tool ([Chouldechova et al., 2018](#)), and Saxena and colleagues’ qualitative work with caseworkers using predictive tools ([Saxena et al., 2020](#)), document a consistent pattern: the context thins, the score hardens, and the worker’s discretion contracts. If we train the same kinds of systems on workers rather than families, the dynamics will not be different. If a model learns that staff who request mental health days or fall behind on data entry are the ones who eventually quit, the rational managerial response is to write those workers off, not to support them.

A Paradigm Shift: From Worker Surveillance to Structural Diagnostics

This is not a call to banish AI from human services. I help lead a federally funded quality improvement center whose entire mission is to strengthen workforce analytics in public child welfare, and the tools we build — hiring-pipeline visualizations, turnover-cost calculators, onboarding and exit-survey instruments — use the same underlying techniques I am critiquing here. The distinction that matters is where the analytics point. Tools that illuminate institutional conditions — caseload balance, overtime distribution, hiring-funnel bottlenecks, onboarding completion — augment supervisory judgment. Tools that score individual workers on their likelihood of quitting quietly replace it.

So the proposal is narrow: redirect the gaze. Instead of using AI to predict which worker will crack under the pressure, use it to map the pressures themselves. Imagine analytics that flag a county court system for generating unsustainable administrative burden, rather than flagging the caseworker drowning in it. Imagine models that audit caseload balance in real time, that measure the downstream effects of after-hours emergency calls on team stability, or that automate the bureaucratic paperwork now stealing hours from families.

If public child welfare agencies want to address the retention crisis, they need to stop buying technology designed to monitor the breaking points of their staff. It is time to point the analytics upward — at the structural conditions leaders can actually change — rather than downward, at the workers trying to hold the system together.



Authored by Robert Blagg, PhD, an applied social welfare researcher dedicated to empowering stakeholders to use data to make timely and sound program and policy decisions.

References

- Ajunwa, I., Crawford, K., & Schultz, J. (2017). Limitless worker surveillance. *California Law Review*, 105(3), 735–776. <https://www.californialawreview.org/print/limitless-worker-surveillance>
- Barbee, A. P., Rice, K., Antle, B. F., Henry, K., & Cunningham, M. R. (2018). Factors affecting turnover intention of public child welfare workers: Comparing worker self-report versus supervisor perspective. *Journal of Public Child Welfare*, 12(5), 542–562. <https://doi.org/10.1080/15548732.2018.1436107>
- Bernstein, E. S. (2012). The transparency paradox: A role for privacy in organizational learning and operational control. *Administrative Science Quarterly*, 57(2), 181–216. <https://doi.org/10.1177/0001839212453028>
- Chouldechova, A., Benavides-Prado, D., Fialko, O., & Vaithianathan, R. (2018). A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. *Proceedings of Machine Learning Research*, 81, 134–148. <https://proceedings.mlr.press/v81/chouldechova18a.html>
- Elgin, D. J., Barbee, A. P., McCarthy, M. L., Kluckman, M., Ringeisen, H., & Dolan, M. (2025). Child welfare workforce onboarding, training, and professional development from 2021 to 2022 (OPRE Report #2025-078). U.S. Department of Health and Human Services, Administration for Children and Families, Office of Planning, Research, and Evaluation. <https://acf.gov/sites/default/files/documents/opre/opre-child-welfare-workforce-onboarding-aug25.pdf>
- Kim, H., & 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.childyouth.2014.01.015>
- Leicht-Deobald, U., Busch, T., Schank, C., Weibel, A., Schafheitle, S., Wildhaber, I., & Kasper, G. (2019). The challenges of algorithm-based HR decision-making for personal integrity. *Journal of Business Ethics*, 160(2), 377–392. <https://doi.org/10.1007/s10551-019-04204-w>
- Saxena, D., Badillo-Urquiola, K., Wisniewski, P. J., & Guha, S. (2020). A human-centered review of algorithms used within the U.S. child welfare system. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2), Article 181. <https://doi.org/10.1145/3392878>