Caseload Reductions in New York City – An Impact Assessment

Fred Wulczyn, Sara Feldman, and Scott Huhr

Center for State Child Welfare Data
Chapin Hall at the University of Chicago
April 2019

Executive Summary

As part of its Waiver strategy, New York City reduced caseworker caseloads within the network of private agencies that provide foster care services on behalf of New York's Administration for Children's Services (ACS), with the expectation that doing so would expedite permanency.

For the evaluation, we asked whether the rate of exit to permanency increased for children whose time in care coincided with when private agencies reached the new caseload target.

In sum, we found that exit rates increased by 9 percent in the years following implementation of the caseload reduction. We also found that when pre-implementation admissions are compared with post-implementation admissions, the median length of stay declined from 525 days to 475 days.

To understand what would happen if the increased rate of exit was reversed because caseloads returned to their pre-implementation levels, we projected the number of foster children five years out following the reversal. We found that the number of children in care would increase by 425 children at a cost of $90 million over 5 years, all else being equal.

1 Waiver Strategy

As part of its Waiver strategy, New York City implemented caseload reductions within the network of private agencies that provide foster care services on behalf of New York's Administration for Children's Services (ACS). The strategy and underlying theory of change are straightforward: to increase the rate of exit, ACS reduced the number of children on each caseworker's caseload. Advocates often target caseloads as an indicator of service quality – too many cases on a worker's caseload limits the amount of time each case receives. Too little time stretches out the work needed to reach a permanency outcome.

As a strategy, caseload reduction doesn't per se change the type of work or how the work is done. Rather, productivity improves simply because there are more people doing the work that needs to be done.[1]

2 Evaluation

Regarding the evaluation, we asked whether the rate of exit to permanency increased for children whose time in care coincided with when private agencies reached the new caseload target. We did this using a unique file that allowed us to measure child to worker ratios in each agency over time. Then, because ACS implemented the caseload reduction system-wide, we used historical placement data to compare exit rates when caseloads were high to exit rates after caseloads reached their target levels. We expected to see a higher rate of exit for (1) children already in care when the changes went into effect and (2) children admitted to care after the changes went into effect. Technical details of our approach to these questions are found in Section 4 – Technical Details.

3 Impact

To establish whether caseload ratios reached target levels, for each private agency we measured the number of cases per worker on a monthly basis for calendar years 2012 through 2016, inclusive. Reductions in caseloads were authorized in 2014. By the start of 2015, average caseloads across all agencies reached the target level of 12 children per worker (Technical Details - Figure 1).

We then examined whether exit rates improved after controlling for characteristics of children (e.g., age, race, and gender) and their placement history (e.g., when did they enter care, how long had they been in care, how many placements had they experienced, and which provider agency provided care?). Detailed findings are included Section 6 – Model Details. In sum, we found:

- Exit rates increased by 9 percent in the post-implementation years when compared to pre-implementation periods.
- Median length of stay for children admitted to care *after* the caseload reduction was 475 days; median duration for children admitted to care *before* the caseload reduction was 525 days. The pre/post difference is approximately 9 percent.

[1] We routinely encounter similar strategies in everyday life. For example, grocery stores open checkout lanes during times of peak activity. TSA opens more security lines at the airport when travel is heaviest. We don't mean to suggest that casework is simply a matter of moving children through a checkpoint. Nevertheless, because of legal requirements regarding the process and quality of care (e.g., reasonable efforts), children require a certain amount of effort to move them to permanency. The work involved requires an expenditure of time. The additional capacity allows for faster processing. In the case of clerks and TSA agents, and caseworkers, during times when capacity increases, workers on the frontline are not doing anything differently necessarily. Rather, they are able to process more cases simply because, with fewer cases, they have more time to devote to the work each child requires.

To understand what would happen if caseworker caseloads returned to pre-Waiver levels, we built a projection model that forecasts the number of children in foster care, the number of children admitted to care over five years, and length of stay. With the length of stay projection, we adjust the exit rates to their pre-waiver level. In effect, we assume that a caseload increase will slow exit rates down for the same reason the caseload reduction resulted in an exit rate increase: with fewer workers to do the work that needs to be done, the time needed to complete the work stretches out.

With the projection model, we estimate a baseline result. The baseline is an estimate of the number of children in foster care at the beginning of each year for the next 5 years, assuming that admission and exit rates remain unchanged over that period. Then, for comparison with the baseline, we adjust the exit rate to mimic what would happen if caseloads increase and the waiver-induced gains were reversed. We measure the impact as an increase in the foster care population and the aggregate cost of foster care using a $85 per day as the blended average cost per day.

The results of the projection model suggest that an exit-rate slowdown would result in an increase in the foster care population, after 5 years, of 425 children, from 8,804 to 9,229. Based on an increase in the number of foster children, because discharge rates slow down by 9 percent, the cost of providing care would increase by $90.46 million. The project model results are found in Section 6.a – Monetizing the Length of Stay Changes.

4 Technical Details

The evaluation uses an interrupted time series model. We elected this approach because implementation of the reductions was system-wide and simultaneous. As a consequence, there is no natural contemporaneous counterfactual. That said, this is a typical situation from the perspective of change initiatives in child welfare systems. Leadership often contemplates a strategic shift in resources that affects the system as a whole. When this happens, the impact assessment has to look for other, reasonable strategies for assessing whether the changes are having their intended impact.

In NYC, the caseload reductions took place over calendar year 2014. In its simplest form, contract agencies adjusted staffing patterns so that the average number of children per worker dropped from about 15 (pre-2014) to just under 12 by 2015.

To observe worker caseloads, rather the rely on self-reports of agencies or workers, we used a unique link between agencies, caseworkers and children served to assess on a monthly basis the number of workers working in an agency, the number of children served by those workers at those agencies, and the *monthly* worker/child ratio, separately for each unique agency in the City network. In this way, we were able to ask whether in a given month the standard had been met and whether children served in months when the standard was met were more likely to leave care.

The results from this linked caseworker/child file are found in Figure 1. As depicted, between the start of 2012 and the start of 2014, average caseloads drifted upward from 13 children per worker to about 15 children per worker before starting downward over calendar year 2014, as the caseload reduction was put in place. By the beginning of 2015, the number of children per worker dipped below 12 children per worker where it has remained since.

Figure 1: Average Number of Children Per Worker

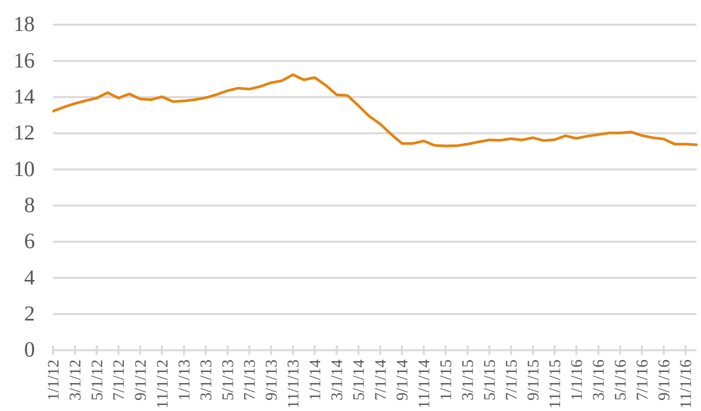

5 WAIVER IMPACT

The interrupted time series analysis was done as follows. We marked time in months before the implementation of the caseload standard (-10 to -1) and after implementation (1 to 10). In the statistical model, we used this indicator to assess the time trend. We then looked at who was in care at the beginning of each month and asked whether they left care during the month. We also noted whether the month in question fell within the period when caseloads met the standard. This provides for a more nuanced understanding of the effect insofar as some of the children and young people admitted in 2012-2014 would have been in care when the caseload reductions went into effect. By looking more closely at exposure – i.e., who was in care when the caseload reductions went into effect, while bearing in mind when the admission to care happened (i.e., how long they had been in care already), we give ourselves a better chance of seeing the impact.

These data, found in Figure 2, show a persistent drop in the likelihood a child would leave care in the months leading up to the reduction. For example, D – 1 refers to the first 30 days of a placement. In the months leading up to the caseload reduction (-10 to -1), the likelihood of leaving care dropped from about 6.6% to just under 6%. In the months following, the change (1 to 10), the likelihood increased from about

6.6% to just above 7 percent. The data show a similar pattern across person-periods. The children in care during the second person-period, regardless of when they entered, were less likely to leave care during that person-period as time went on. After the caseload reduction, the rate (or likelihood) increased. It is worth noting that the results presented in Figure 2 hold up after differences in the population served are taken into consideration. As mentioned, this provides a more detailed view of length of stay changes relative to the caseload reduction.

Figure 2: Change in the Likelihood of Leaving Care by Person-Period
Relative to When the Caseload Reduction Started

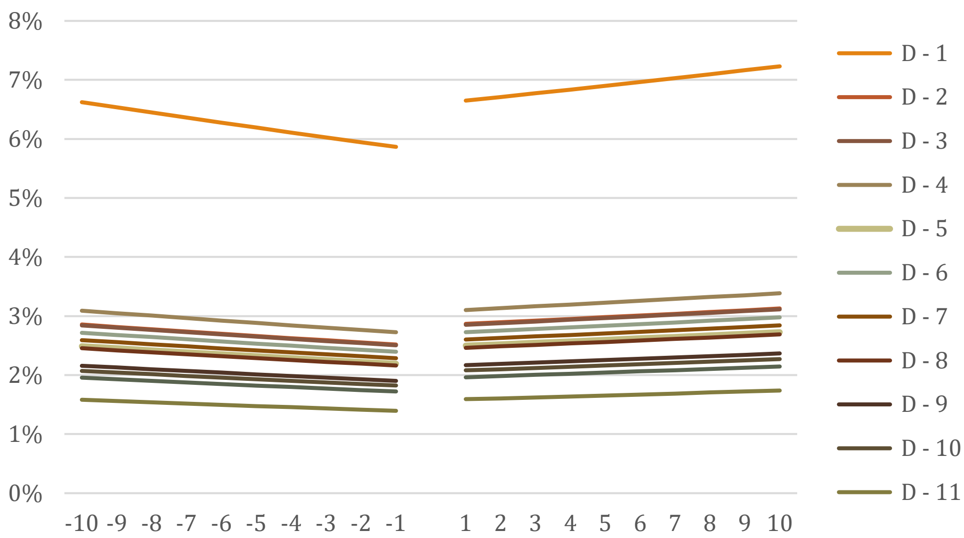

6 Model Details

Details of the model used to understand the Waiver effects associated with the caseload reduction are found in Table 1. Our principal interest is in the Post Implementation x Trend interaction effect, as that represents the treatment effect. Simply put, as a general matter, the trend in NYC indicates the over the period of observation – 2012 through 2016 – exits rates were slowing (see Time trend in Table 1). Post-implementation (in the months after the caseload reduction), exit rates were increasing. When the post-implementation person-periods (i.e., placement months) are adjusted for the time trend, the rates of exit show a significant increase.

The random effects nature of the model accounts for the fact that children placed with some agencies leave care faster because of the agency's own performance. By controlling for the agency effect, we increase the validity of the results.

Other factors in the model account for demographic attributes of the children (age, gender, race/ethnicity) and history of placement. A child placed with an agency today may have been placed with another agency at some earlier time. Adjusting for the number of prior agency spells, as we call them, accounts for the fact children change placements. In an indirect way, this adjustment controls for the mix of reasons children leave placement and what those placement changes mean for when children leave care.

Finally, we control for placement month. The likelihood of leaving care changes with the how long the child has been in care. In this case, we compare exit rates in subsequent months with the exit rate in the first month of placement. Generally, exit rates in the first month are highest, as indicated by the top line in Figure 1.

Table 1: Random Effects Interrupted Time Series Model Using Discrete Time

Effect	Estimate	Standard Error	t value	Prob.
Intercept	-3.7385	0.08659	-43.17	.0001
Time trend	-0.01894	0.00181	-10.46	.0001
Post implementation	0.08023	0.04463	1.8	0.0723
Post imp. x Trend	0.03692	0.002065	17.88	.0001
Females	Reference			
Males	-0.01294	0.02109	-0.61	0.5393
Whites	Reference			
Blacks	0.15	0.06861	2.19	0.0288
Hispanics	0.1636	0.0702	2.33	0.0198
Other	1.2229	0.06812	17.95	.0001
1st agency placement	Reference			
2nd placement	-0.03555	0.02677	-1.33	0.1842
3rd placement	-0.2017	0.04792	-4.21	.0001
4th placement	-0.5471	0.06972	-7.85	.0001
Infants	Reference			
1 to 5-year olds	0.3721	0.02682	13.87	.0001
6 to 13-year olds	0.4931	0.02901	17	.0001
14 and above	-0.1706	0.0416	-4.1	.0001
Placement month 1	Reference			
Placement month 2	-0.8226	0.04966	-16.57	.0001
Placement month 3	-0.8097	0.0506	-16	.0001
Placement month 4	-0.7152	0.05027	-14.23	.0001
Placement month 5	-0.9295	0.05538	-16.78	.0001
Placement month 6	-0.811	0.05432	-14.93	.0001
Placement month 7	-0.831	0.05606	-14.82	.0001
Placement month 8	-0.9124	0.05897	-15.47	.0001
Placement month 9	-1.0377	0.06309	-16.45	.0001
Placement month 10	-1.0909	0.06554	-16.65	.0001
Placement month 11	-1.3438	0.07362	-18.25	.0001
Placement month 12	-1.1144	0.06832	-16.31	.0001
Placement month 13	-1.0937	0.069	-15.85	.0001
Placement month 14	-1.1361	0.07162	-15.86	.0001
Placement month 15	-1.4116	0.08194	-17.23	.0001
Placement month 16	-1.4642	0.0853	-17.16	.0001
Placement month 17	-1.4591	0.08657	-16.86	.0001
Placement month 18	-1.2098	0.07919	-15.28	.0001
Placement month 19	-1.2329	0.08153	-15.12	.0001
Placement month 20	-1.3441	0.08697	-15.46	.0001
Placement month 21	-1.3998	0.09053	-15.46	.0001

Effect	Estimate	Standard Error	t value	Prob.
Placement month 22	-1.2444	0.08608	-14.46	.0001
Placement month 23	-1.4544	0.09608	-15.14	.0001
Placement month 24	-1.3098	0.09157	-14.3	.0001
Placement month 25	-1.3733	0.09583	-14.33	.0001
Placement month 26	-1.1415	0.08832	-12.92	.0001
Placement month 27	-1.0575	0.08686	-12.18	.0001
Placement month 28	-1.2913	0.09781	-13.2	.0001
Placement month 29	-1.118	0.09247	-12.09	.0001
Placement month 30	-1.1608	0.09613	-12.08	.0001
Placement month 31	-1.2041	0.09993	-12.05	.0001
Placement month 32	-1.2708	0.1047	-12.13	.0001
Placement month 33	-1.1633	0.1017	-11.44	.0001
Placement month 34	-1.2796	0.1088	-11.76	.0001
Placement month 35	-1.2684	0.1105	-11.47	.0001
Placement month 36	-1.3035	0.1141	-11.42	.0001

This person period format is important in the context of the model structure. Children placed *after* the caseload reductions when into effect would have experienced any potential benefit over the entirety of their placement trajectory. Child placed prior to the caseload reduction would have only experienced the potential benefit for those months in care that coincided with the policy change. That might have been the 2nd or the 32nd month of their placement. By the same token, children who enter and leave care before the caseload changes take effect remain in the analysis. By noting whether the person-period/placement month overlaps with the caseload reduction, within the interrupted time series framework, we can directly compare exit rates, with and without exposure to the treatment, after adjusting for the effect of time in care on exit rates.

6.a Monetizing the length of stay changes

In this section, we use a projection model to estimate the population and budgetary impact of an increase in caseloads that would happen if the positive gains of the reduced caseloads were lost (i.e., the caseload ratio returns to pre-2014 levels). In practical terms, were this to happen, it means that children in care when caseloads return to prior levels will stay longer (i.e., their likelihood of leaving will go down) and children admitted after the policy goes into effect will also be less likely to leave care. This is because the amount of work needed to discharge a child remains more or less constant but the number of people doing the work goes down. Effectively, the change in work force without some other increase in productivity causes a slowdown in movement through the system measured as an increase in the number of children in care because exits fall.

In Table 2, we report the following. Scenario 1 assumes a 9 percent change in the likelihood of leaving foster care.[2] The 9 percent reduction in the likelihood of leaving care is based on the change in exit rates observed during the time when the caseload reduction was in effect (see Figure 1). Basically, the results from the interrupted time series indicate that exits rates improved by about 9 percent as a result of the caseload reduction. In the model, we simply reverse the gain going forward. If caseloads return to levels observed prior to the policy shift (see Figure 1), we assume the loss of capacity results in a slowdown of exit rates, which we translate into a lower probability of leaving care.

The adjustment applied in Scenario 1 is a one-time adjustment applied to children in care at the beginning of the fiscal year (population in care). If children in that legacy population are still in care at the beginning of the next year, there is a similar reduction in the likelihood of leaving over the course of the next year. This adjustment is applied for 5 years. We apply a similar change to children admitted each year following the policy change (i.e., the increase in caseloads). In the case of Scenario 1, the magnitude of this adjustment is 9 percent. That is, we reduce the likelihood of exit by 9 percent. Again, this is based on the observed improvement in exit rates illustrated in Figure 2.

In Scenario 2, we maintain the 9 percent reduction in the likelihood of exit in the initial year (i.e., the first fiscal year of the change). However, in subsequent years, we apply an additional 1% per year to the downward adjustment of exit rates. Scenario 3 replicates Scenario 1 but with a slightly larger downward adjustment (15%) applied across the board. Scenario 4 replicates Scenario 3 but with the same 1% additional adjustment applied each year. This is intended to give a range of results. For example, if the impact is larger than anticipated, Scenarios 2, 3, and 4 provide an upward bound of the impact of increasing caseloads.

Scenario 1, which merely reverses the gains achieved by lowering the caseworker caseloads, is the most conservative estimate. A slowdown in exit rates due to an increase in caseloads – a constant amount of work to be done by fewer people – would increase the number of children in foster care by 425 children

[2] The estimates of impact – i.e., the results presented in Table 2 – were generated by comparing the results from a no change projection with results derived by changing the assumptions in the projection model. The projection model uses 6 years of historical data to estimate the number of children entering care and how long children will stay in care. The model generates estimates by groups of children stratified on the basis of admission volume and length of stay. The no change projection asks what would happen if history repeats itself – that is there is no change to the underlying trends. For example, if admissions are rising, the baseline projection model relies on an assumption as to whether the trend will continue or abate. The baseline is then compared with a second model built around a different set of assumptions. The second set of assumptions may be grounded in effects tied to policy and/or practice changes or in effects that are the result of external forces (e.g., an increase maltreatment rates that results in higher foster care admissions).

over 5 years at a 5-year cumulative cost of $90.46 million.[3] In Scenario 2, the additional downward adjustments to the exit rate over subsequent years increases the estimated population by 503 children (an additional 78 children) and an additional $7.4 million in spending.

Scenario 3 and 4 point to population increases of 734 and 812 children, with additional costs running to about $153.5 million and $160.4 million, respectively. Again, these estimates are intended to provide an upper bound on what could happen, if for some reason the impact of rising caseloads is larger than expected.

Table 2: Population and Budgetary Changes Resulting from an Increase in Caseloads

Summary of expected changes	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Summary of expected changes:				
Number of admissions	None	None	None	None
Changes in exit rates for:				
Population in care (average)	9%	9%[+]	15%	15%[+]
Children admitted (average)	9%	9%[+]	15%	15%[+]
Number of children in care - increase	425	503	734	812
Total budget impact (5 years)	-$90,461,898	-$97,405,163	-$153,472,506	-$160,396,890

[+] In Scenarios 2 and 4, in addition to the basic exit rate adjustment of 9% and 15%, further reductions of 1% per year were added in years 2 through 5 of the projection.

[3] These increases have baked into them a projected increase in the foster care population under a business as usual assumption. The baseline increase is a function of an increase in admissions as observed in recent years. The admission increase is held constant over the simulation. That is, projected admissions are held at the level of the most recent year.