The Center for

State Child Welfare Data

Risk Groups and the Average Cost of Paid Foster Care in Texas

Fred Wulczyn, Kristen Hislop, Scott Huhr

March 2019

CHAPIN HALL
AT THE UNIVERSITY OF CHICAGO

1313 E. 60th Street
Chicago, IL 60637
(773) 256.5100
fcda.chapinhall.org

1 Introduction

In Texas, under the legacy system of reimbursement, providers of out-of-home care receive from the Department of Family and Protective Services (DFPS) a per diem payment for each day of care provided. The rate paid (i.e., the per diem rate) depends on the placement type and the service level. The differential rates established under the legacy method reflect differences in the care provided, with higher rates tied to the cost of providing more specialized or intense care based on what children need. For reasons described below, DFPS has modified this approach. Rather than differentiated rates that rely on the placement type and service level to determine the rate paid, the new approach is based on a blended rate that applies to each day of care regardless of the child's placement type or service level. The single, blended rate has been utilized for the past 6 years in catchment areas operating under the terms and conditions of a Single Source Continuum Contract (SSCC) both through Foster Care Redesign and now Community-Based Care. That experience has prompted a review of how the blended rate is established. In past years, the blended rate used the level-specific number of days and the level-specific per diem rates for the state and age strata for the catchment area to calculate a blended (or weighted) rate that was then applied to each day of care actually provided. In more recent years, with the absence of service levels in the Community Based Care model, it has become increasingly difficult to predict the actual case mix of children inherent in the blended rate in each catchment area.

In this report, to address these concerns, we propose an approach the identifies groups of children distinguished by a set of characteristics (e.g., their age at placement, their reason for entering care, and their placement history) that are strongly correlated with the average daily rate paid in previous fiscal years but is less reliant on the service level system. The results of this model are then used by DFPS to generate the blended per diem rate. Details of how this was accomplished follow.

2 The Blended Per Diem Rate

For the past decade, the Texas Department of Family and Protective Services has been involved with a major restructuring of its foster care program. A major part of this effort involves shifting the approach to reimbursement for foster care services. Rather than rely on per diem rates that vary by service level and type of care, DFPS has shifted to a blended rate that combines multiple rates into a single rate. The shift to a blended rate is meant to deemphasize the service level system so that providers are focused on aligning the care provided with the real-time needs of the child. This approach moves the department closer to value-based purchasing strategies that are increasingly common in health care.

To address changes in payment methods DFPS is seeking a more robust methodology for stratifying the foster care population as the basis for its blended rate methodology. In the (recent) past, DFPS relied on statewide service level utilization to generate a blended per diem rate that would, in theory, generate the revenue an agency (an SSCC in this case) would have received under the legacy model. However, when the statewide blended rate is applied to catchment area specific utilization, the result will fall above or below the actual cost of providing care if the mix of days provided within the catchment area is different than the projected mix of days used to blend the statewide rate. These differences may arise for various reasons but the approach to stratification is important. With several years of experience,

practice suggests that this is what has happened with use of the blended rate. Specifically, in the 3B catchment area, the SSCC charged with managing paid foster care services (ACH Child and Family Services) has found the statewide blended rate to be too low because the actual case mix of days at the various levels is different than the proportion found in the statewide data. As a consequence, OCOK has experienced revenue shortfalls, which DFPS has worked to address through reconciliation.

Ordinarily, building a catchment area-specific blended rate would not pose technical challenges over and above those encountered when undertaking the task at a statewide level. However, the challenge in the Community Based Care (CBC) catchment areas is compounded by the fact that, under the new foster care model the state no longer requires the use of or tracks service levels for children through IMPACT, the state's case management system. This decision was made to de-emphasize reimbursement levels to placements and to better align SSCC performance in a manner that supports and incentivizes improvements in child well-being as demonstrated by a reduction in service need. However, this change means that DFPS has no way of identifying the service level used, the children placed in each of the various service levels, or the time spent in each service level. Insofar as each of those pieces of information play a vital role in stratifying the population for purposes of the blended rate calculation, DFPS has found it increasingly difficult to establish a blended rate on a prospective basis.

This problem has created other challenges. First, as each catchment area moves further away from the legacy system, the *potential* gap between the projected blended rate and the actual blended looms larger, which raises fiscal uncertainty. That same uncertainty may affect how new SSCCs prepare to take over foster care operations under the CBC model. Third, the aforementioned challenges highlight the need for a formal reconciliation process that addresses differences between projected and actual blended rates that arise.

The methodology sought by DFPS would address this interrelated set of problems by building a robust stratification of the foster care population that is less dependent on the legacy service level system. Ideally, children stratified through the application of such an algorithm, for lack of a better term, would render a projected blended rate developed by DFPS that is a more accurate reflection of what it will likely cost to provide care going forward.

Modeling the Average Cost per Child per Day

Our approach to stratifying the population is straightforward: we set out to isolate a set of factors correlated with the average cost per child per day over the course of their entire placement history.[1] In doing so, we followed three basic steps. First, we used descriptive data to isolate factors that were good candidates for including in the model. Second, we tested various factors, one-by-one and in various combinations to find the factors with the strongest association with the cost per child per day. When comparing results, we considered the variation explained (i.e., the R-squared) and the effect size. Among the variables tested, we included these in the models:

- Age at removal: under 1 year old (reference group), 1 to 5 years old, 6 to 13 years old, and 14 years old & above
- Parent reason which is an indicator that tracks parental risk factors associated with why children were removed and placed into substitute care. We tested parental risks in two ways – as individual reasons and as a count of how many different reasons were associated with each child spell
- Child reason which is an indicator that tracks the reason for placement. We tested placement reason in two ways – as individual reasons and as a counter of how many different reasons were associated with each child spell
- A spell counter for the number of child spells the child experienced (see Table 1)
- A paid spell sequence counter that tracks the number of paid conservatorship episodes experienced by the child (see Table 1)
- A rate span counter that tracks the number of unique rate spans experienced by the child (see Table 1)
- Various interaction terms, which tested whether the effect of one variable was dependent on the value of another variable, were tried. For example, the child spell/rate span interaction term tested whether the rate span effects the average cost per child per day were dependent on whether the rate span occurred during the first child spell or some subsequent child spell
- A catchment area indicator was used to detect differences in the average cost per child per day by catchment area

The models run were based on ordinary least squares regression (OLS). Two variants of the basic OLS model were tested: fixed effects and random intercepts. Although both approaches are essentially the same, the random intercept model is regarded as the most appropriate choice in analyses of clustered

[1] For a detailed explanation of how the cost per child per day was calculated, please see the explanation provided in the Appendix section entitled: How was the history of conservatorship used for the analysis? It is also important to note here that the average cost per child per day is based on the cost of care provided to a child over the entirety of their time in care, without regard for the fiscal year in which the care is provided. A child who spends a portion of two or more fiscal years in care will have the days provided in a given fiscal year reflected in the cost per day for the given fiscal year. Days of care that fall in other fiscal years will be included in the corresponding fiscal year. When the days included in a cost per day calculation are limited to the days that fall within a fiscal year, the results may differ from average cost per child per day. Our goal in this exercise is to identify strata of children, where the between strata average costs per child per day are large. The strata identified in this way is then used in the Department's blended rate methodology.

data as is the case in this example (i.e., the catchment areas serve as the cluster variable).[2]

Table 1: Factors Affecting the Cost Per Child Per Day

Variable	Coefficient	P-value
Intercept	$49.55	0.00
Age		
Infants	Reference	
1 to 5	$0.72	0.04
6 to 13	$13.89	0.00
14 & above	$36.93	0.00
Placement History		
1st child spell	Reference	
2nd child spell	$1.77	0.00
3rd child spell	$0.81	0.35
Rate Counter	$8.56	0.00
Reason for Placement		
Missing	$-3.06	0.00
Refuse parental resp.	$15.69	0.00
Other conservatorship*	$5.86	0.00
Medical neglect	$9.47	0.00
All others	$-8.78	0.00
Catchment Area		
1	$0.17	0.8313
10	$-1.23	0.2155
11A	$-1.18	0.1518
11B	$-4.68	0.0001
2	$-0.19	0.8312
3A	$0.54	0.5186
3B	$-4.63	0.0001
3C	$-2.54	0.0006
4	$1.23	0.1164
5	$-0.58	0.5009
6A	$-0.92	0.1971
6B	$3.33	0.0001
7A	$3.18	0.0001
7B	$4.73	0.0001
8A	$0.19	0.7927
8B	$-0.03	0.972
9	$2.60	0.0025

* - Other conservatorship refers to a placement reason as defined by DFPS.

Although we produced and reviewed several dozen models, in the interest of brevity, the results presented in Table 1 are from the penultimate conditional random effects model. Briefly, the results

[2] The random intercept models differ slightly from ordinary regression in that it controls for the amount of information provided by each catchment area. Specifically, if catchment areas vary in size (which they do) then large catchment areas contribute more to the statewide average than small catchment areas. As an example, if one were to compute an average length of stay using California and Vermont, the average would be more reflective of Ca. because there are so many more cases from Ca. contributing to the average. Although we have not discussed this specific issue, this is another reason why the statewide per diem rate is biased.

show the following:

- Age is the most important predictor of the cost of care. For example, compared to infants (i.e., age at admission), youth between the ages of 14 and 17 at the time of admission will cost an additional $36.93 cents per child per day.[3]
- Placement history is important. Each time the child moves between levels of care, the cost per day increases by $8.56 per child per day on average.
- Costs per child per day also increase as the number of child spells rises, but these differences are generally smaller than the change associated with type/service level changes.
- The cost per child per day does depend on the reason for removal but not the parental risk factors (which were dropped from the final model). Children placed because the parent refused parental responsibility for their child cost an extra $15.69 per day, on top of the other applicable child characteristics and associated costs. For example, a teenager placed because the parents didn't accept parental responsibility would cost $49.55 (the base or intercept) plus $36.93 plus $15.69, given the reason for placement, and other measured characteristics of the child that were included in the model. Another example would be, if the child was in the midst of their third rate span an additional $25.68 would be added to the expected cost of care (i.e., 3 x $8.56)
- There is important variation between catchment areas in the cost of care. These cost differences arise because the use of different levels of service varies between catchment area. The use of different levels of service varies because, even though Youth for Tomorrow makes service-level determinations, the supply of services is in important determinant that affects the service level selected. The catchment areas with statistically different costs per child per day are shaded in Table 1.

3 Stratifying the Population

As a second step in the process, we went back to the source data and created groups of cases based on the unique combination of the factors in the model. We call these groups strata. Each stratum represents a unique combination of the factors taken from the model, with some exceptions that we explain below. For example, we combined each value of age group with every reason for placement, the rate counter, and child spell. Naturally, this process produced a large number of unique sub-populations which we reduced as follows. Because cost per child per day differences are small when 1 to 5-year olds are compared with children under the age of 1, we combined them into a single group of infants to 5-year olds, inclusive. Regarding placement history, we dropped the number of child spells from the model, again because the differences were small. We also grouped the rate span variable into two groups: children who experienced 1 or 2 rate spans and children who experienced 3 or more.[4] We chose this particular grouping because the majority of children have either one or two unique rate spans within a spell of paid placement. Finally, we recategorized the reason for placement, based on the fact children may have more than one reason listed. When children are placed because parents refuse to

[3] Note that the intercept in the conditional model is $49.55. In the conditional model, the intercept represents cost for an average child, which in this case is an infant in their first child spell plus other descriptors of the child: catchment area, rate counter, and reason for placement.

[4] See the description of how the data were used, found in the Appendix, for an illustration of how we constructed the rate spans.

accept parental responsibility, the average cost per child per day is significantly higher than it is for children placed for other reasons, so we left them as a separate group. Children placed for reasons tied to 'other conservatorship' and medical neglect were linked together because both are associated with lower average cost per day.[5] Children with multiple reasons were grouped together for purposes of stratification as were children with no reason listed and children with other placement reasons that did not fit into the other categories already described.[6] The results of this effort to reduce the number of categories used are found in Table 2 and 3.

Table 2: Basic Summary Data for Unique Strata - Children Placed in Foster Care: 2011 - 2015

Age, Placement Reason & Rate Spans	Total Spending	Total Days	PC - PDR	Number of Obs.
Total	$1,989,882,100	26,515,282	$75.05	81,497
Age				
Under 6	$754,224,341	13,124,817	$57.47	41,552
6 to 13	$843,165,082	9,910,296	$85.08	26,520
14 & Above	$392,492,678	3,480,169	$112.78	13,425
Placement Reason				
Parental Res.	$153,652,836	1,230,929	$124.83	3,889
Other Cons + Med Nglct.	$54,145,387	582,237	$93.00	1,394
Multiple	$140,873,659	1,621,983	$86.85	4,481
None	$58,169,013	749,031	$77.66	1,987
Other	$1,583,041,206	22,331,102	$70.89	69,746
Number of Rate Spans				
Under 3	$1,100,030,401	18,077,219	$60.85	70,075
3 or More	$889,851,699	8,438,063	$105.46	11,422

PC – PDR: Per Child Per Day Rate. Placement reasons: Parental Res. = Refusal to accept parental responsibility; Other Cons + Medical Neglect = other conservatorship (as defined by DFPS or medical neglect; Multiple reasons = more than one reason given; none = nothing listed; Other = all other reasons not separately listed. Dollar figures derived by applying 2017 type of care / service level specific rates to the days of care provided.

For purposes of stratifying the population, the most critical column is PC – PDR, which references the per child per diem rate for each unique group of children. The PC - PDR is calculated by summing over all the days of care provided to each of the 81,497 children using the paid cost of care, regardless of type / service level, experienced by those children during their spells of paid foster care. The total count of days and the total spending for those days is summed over all the members of the strata to generate a cost per child per day within the strata.

From this data, we know that the expected cost per child per day for children under age 6 is on average $57.47 per day. This reflects the fact that children in this age group spend most of their time at a basic service level, although some care provided to this group is specialized or intense, which have higher per

[5] Other conservatorship is a reason for placement found in the IMPACT data used for the analysis.
[6] The other placement reason category includes the following reasons: abandonment, emotional abuse, infant risk, neglectful supervision, physical abuse, physical neglect, and sexual abuse. In advance of producing the final model, we tested each reason for placement individually. From those results, we made judgments as to whether the reason contributed to significant differences in the average rate paid. That process led to the specification shown in Tables 1 and 2.

diems. The more expensive care is less common with this age group but relative to the rate for basic care alone, the daily rate is elevated slightly because the mix includes all types and service levels. By way of comparison, young people 14 years of age and older cost more per day ($112.78) because they are more likely to be placed in congregate care. There are relatively fewer children in this age group (13,425) but their daily cost is higher.

Table 2 also shows the rate differences for children grouped according to their reason for placement and the number of rates spans. As already noted, children placed because their parents refused to take responsibility cost considerably more per day ($124.83) than children placed for any other reason. Similarly, children who experience different types of care (i.e., their time in paid placement is spent at different service levels) cost more per day because, by definition, some of their time is spent in specialized or intense care. Specifically, children who experience 3 or more changes in service level cost $105.46 per day on average.

The next step in the process involved grouping children together based on their *combined* characteristics (i.e., their age, the reason for placement, *and* the number rate spans) to produce unique, mutually exclusive groups of children – i.e., the strata. Table 3 shows the results of this step after dropping rate span from the stratification approach.[7] The strategy for reducing the number of strata balanced parsimony and clinical relevance. Parsimony reflects a preference for relatively few unique strata; clinical relevance reflects the idea that the remaining strata ought to provide at least some guidance from a case management perspective. For example, the cost of caring for older children who are placed because the reason for placement lists the parent's refusal to accept responsibility is well above average. For case managers, the tendency to use higher service levels within this group should, perhaps, be reflected in the care plan. We also considered population size. For some strata, once all the unique combinations of age, reason, and rate span were assembled, there were fewer than 200 children per strata over five years of data. At a catchment area level, the count of children in these groups would be very small, which means those strata exert little if any influence on the overall blended per diem rate.[8]

In the end, the reduced number of strata, as shown in Table 3, produces robust results when the expected PDR is compared with the observed PDR. The observed PDR was generated from the actual data; the expected PDR was calculated from the results in Table 2. Generally, the expected PDR is close to the observed PDR. The two exceptions are children age 6 to 13, and children above the age of 14 who were placed because their parents refused to accept responsibility (6 to 13-year olds) and children admitted for reasons having to do with other conservatorship or medical neglect (young people above age 14). In the case of the older group, the number of days provided to children admitted for reasons having to do with other conservatorship or medical neglect is small. In both cases, the relatively large difference between the expected and observed PDR are due to variation in the underlying utilization of

[7] In the final analysis, the rate span variable was dropped from the model because at the time of admission, the number of rates spans is not known whereas age and placement reason are.
[8] It is worth noting here that that for children who fall outside the blended foster care rate, the department uses an exceptional care rate as the basis for reimbursement.

care.

Table 3: Reduced Strata Based on Age and Reason for Placement

Age and Reason	Total Spending	Total Days	Observed - PDR	Expected - PDR	Difference
Under 6					
Cons + Med Nglct	$21,331,908	276,929	$77.03	$75.38	($1.65)
All Other Reasons	$732,892,433	12,847,888	$57.04	$57.03	($0.01)
6 to 13					
Parental Res.	$55,607,049	442,245	$125.74	$115.81	($9.93)
Cons + Med Nglct	$21,898,393	210,636	$103.96	$103.66	($0.30)
All Other Reasons	$765,659,640	9,257,415	$82.71	$83.04	$0.33
14 & Above					
Parental Res.	$91,213,471	685,317	$133.10	$133.37	$0.27
Cons + Med Nglct	$10,915,087	94,672	$115.29	$121.50	$6.21
Multiple Reasons	$34,953,985	287,918	$121.40	$121.37	($0.03)
None w/ Other	$255,410,135	2,412,262	$105.88	$105.19	($0.69)

As the final step in the stratification process, we reproduced the analysis found in Table 1 using age, reason for placement, and catchment areas as statistical controls. These results are found in Table 4. As in Table 1, the results point to significant differences in the average cost per child per day. In this model, the intercept, which represents the average cost per day for a child between birth and 5 years of age at the time of admission, is higher than it is in the original model (see Table 1) because the effects of rate spans and placement history have been dropped. Indeed, all of the results in Table 4 are somewhat higher for that reason. For the catchment areas, the results again show the upward or downward adjustment to the average cost per child per day associated with children from those catchment areas. These differences arise because the blend of care used in each catchment area differs. Moreover, the effect of removing the number of rates spans from the model (i.e., comparing the catchment area results from Table 1 with those in Table 4) is further indication of how the utilization of care differs from one catchment area to the next, differences that arise for a variety of reasons including the supply of care (or service capacity) within catchment areas. Finally, the results in Table 4 were data transmitted to DFPS for use by the Department inside the Department's rate blending algorithms.

Table 4: Final Model for Linking Case Characteristics to the Average Cost Per Child Per Day

Variable	Coefficient	P-value
Intercept	$66.21	0.00
Age		
Birth to 5	Reference	
6 to 13	$26.20	0.00
14 & above	$48.69	0.00
Reason for Placement		
Refuse parental resp.	$20.15	0.00
Other conservatorship* + Medical neglect	$9.96	0.00
All others	-$9.80	0.00
Catchment Area		
1	$3.47	0.00
10	$2.35	0.05
11A	$1.41	0.17
11B	-$6.30	0.00
2	-$0.08	0.94
3A	-$0.12	0.91
3B	-$8.32	0.00
3C	-$3.17	0.00
4	$0.76	0.45
5	-$3.20	0.00
6A	-$2.20	0.02
6B	$3.78	0.00
7A	$0.97	0.33
7B	$2.30	0.03
8A	$2.22	0.02
8B	$2.34	0.02
9	$3.76	0.00

* - Other conservatorship refers to a placement reason as defined by DFPS.

4 Summary with Implications

Our aim in this analysis was to identify subgroups of children with distinct average costs per child per day. As shown in Table 3, when the results of the model are used to generate an expected per diem rate (i.e., the projected average cost per child per day) for each stratum, the expected per diem rate comes very close to the observed per diem rates for those same children, at that specific point in time. For that reason, DFPS should expect more robust results from the methodology it uses to generate the blended rate that serves in turn as the basis for reimbursing SSCCs.[9]

With that said, there is a small handful of implications to consider when using blended rates as a mechanism for reimbursing foster care providers, especially as stakeholders press on with the implementation of CBC, which is already underway. Among those issues, reconciliation, monitoring the use of foster care, risk-sharing, and the continuum of care are or should be top of mind. A brief

[9] It is important to note here that the model described does not generate the blended rate DFPS intends to use. Rather, the results describe the strata DFPS intends to use with its blended rate methodology.

summary of each issue follows.

4.a Reconciliation

Although a more robust stratification approach will improve the accuracy of the blended rate used for reimbursement, even a perfectly blended rate will not eliminate the need to reconcile the difference between the expected blended rate (based on forecasted case mix of children) and the observed blended rate (based on the actual case mix of children served). The need for reconciliation is predicated on the reasons why blended rates are used in the first place. Blended rates are designed to stimulate changes in the pattern of care on the part of care providers. If providers respond to the incentives and reduce the use of high-end care, the gap between the observed and expected rates will grow in the intended direction. The reverse may also be true – the observed blended rate may exceed the expected blended rate – for any number of reasons including those that are linked to the risk-sharing plan DFPS needs to develop in conjunction with stakeholders. Using the results described in Table 4 as the basis for an example, if there is a spike in the number of adolescents admitted because their parents refuse to accept responsibility then the gap between the observed and expected per diem will grow, even if the expected per diem was generated with perfect information. Whether the expected rate is too low or too high relative to what was observed, only a thoughtful reconciliation framework will resolve the differences to everyone's satisfaction.

4.b Monitoring foster care utilization and the blended rate

A thoughtful reconciliation process is highly dependent on a vigorous process of monitoring the utilization of care and outcomes. In the final analysis, the blended rate simply reflects the average cost of care across the population of children used to blend the rate. By definition, it reflects a number of factors including the number and proportion of children belonging to each of the identified strata, their length of stay, their outcomes, and cost of care provided. As an estimate of future costs, the accuracy of the blended rate will be affected by whether any of these inputs change going forward. On the one hand, if nothing changes, the blended rate will perform as expected. Even an increase in the number of children served will be absorbed with the blended rate provided the underlying mix of children *does not* change. On the other hand, if the dynamics underlying the blended rate change between the time the rate is established and when it is applied, a gap between the blended rate and the observed rate *will* materialize. If the foster care system is accountable for the safety, permanency, and well-being of foster children, then it is incumbent upon everyone involved to understand and respond to changes in the population served, either programmatically, fiscally, or through policy, as necessary. How those changes are identified is first and foremost a function of monitoring. As a starting point, each SSCC and DFPS should understand how the strata developed for the blended rate are reflected in the local caseload at the start of each fiscal year. Then, changes in the population should be monitored with the expectation that stakeholders will be in a position to articulate with documentary evidence why the changes occurred. Then and only then can the risk sharing model be applied equitably.

4.c Risk-sharing and incentives

Risk-sharing and incentives are an *inevitable* by-product of reimbursement. Indeed, the nature of risk-sharing and incentives under traditional per diem approach is driving human service systems, including child protective services, in the direction of alternatives that emphasize outcomes rather than services

provided. Nonetheless, these new approaches are not without their own risks and corresponding incentives that could adversely affect the services children receive and the outcomes they experience. For example, per diem systems are linked to incentives that increase the number of services provided relative to need (i.e., over serving) whereas blended rates may lead to fewer services being provided relative to need (i.e., under serving). For these reasons, it is important to be clear about risk, the incentives risk produce, how risk will be shared, and under what circumstances.

Regardless of how services are reimbursed, either prospectively through some sort of blended rate or through more traditional fee-for-service payments, risk-sharing begins with understanding the source of risk. Generally, risk is tied to the volume of services provided, measured as the number of admissions and the average number of units provided (or length of stay in the case of foster care), and the rate paid for each unit of service (i.e., the service level). A proper risk-sharing plan defines the risk, how much of the risk will be shared and by whom. A list of considerations includes but is not limited to:

- An increase in the demand for out-of-home care changes (i.e., an increase in admissions)
- An increase in cost due to an increase in care days used because of longer lengths of stay – i. e., the time needed to achieve positive outcomes
- An increase in cost due to increases in the unit cost of care due to the services needed to address the well-being of children
- A decline in the demand for out-of-home care (i.e., fewer admissions)
- A decline in costs that arise because the volume of care days drops due shorter lengths of stay
- A decline in costs due to a drop in the unit cost of care because the use of high-end care drops
- A change in the populations that contribute to the case mix tied to the volume and the unit cost of care due to other reasons such as those listed above

With these listed possibilities in mind, the risk-sharing plan considers how over- and under-payment will be managed. Alternative risk-sharing options include hold harmless provisions that are applied across the board or in reference to specific risks. For example, DFPS could hold providers harmless against an *increase* in cost due to an increase in admissions. This is how risk is shared under the legacy system. From the provider's perspective, a *decrease* in admissions is a risk they absorb, as under the legacy system. The point is this: each risk (volume, duration, and unit cost) has to be evaluated from the DFPS perspective and the provider's perspective. Only then will a reasonable risk-sharing plan that highlights the incentives and disincentives emerge. Risk corridors are another design option available to DFPS and its stakeholders. Basically, a risk corridor describes the point or points at which the risk sharing plan changes. The risk corridor is sometimes referred to as a stop loss provision in the case of downside risks. Risk corridors may also be applied to upside risks. In the case of an upside risk corridor, DFPS shares in the benefits of a more efficient *and* effective service system. Finally, the risk sharing plan may involve exclusions of some sub-populations because their costs are unusually high and/or unpredictable. For example, in our analysis of the blended rate, we did not include children served under the exceptional care rate in the CBC model.

4.d The continuum of care

In the run up to the current reform initiatives, the service level attracted considerable attention because stakeholders worried that differences in rates paid weighed on placement decisions, rendering the decision to move a child from one service level to another a fiscal matter as opposed to a clinical decision that aligns the service provided with the child's need. To the extent those dynamics were in place, the move to a blended rate made sense because blended rates de-emphasize the differences in reimbursement associated with service levels.

With that said, DFPS and its stakeholders should revisit what it costs to provide care at different points along the continuum of care. Because the mix of services along the continuum should reflect what is needed to meet the needs of children placed in those settings, stakeholders should expect the cost of providing clinically appropriate care, in addition to board and maintenance, to change.

An effort to understand the care provided and the underlying costs would have at least two direct benefits. First, the analysis would bring the blended rate closer to the cost of providing care along the continuum, all else being equal. Components of the underlying costs – property costs, professional services, and the efficacy of clinical interventions – change over time and there has to be a mechanism in place that recognizes these changes in the process by which rates are set. Second, the effort to pinpoint the cost of care would clarify the clinical rationale that undergirds the distinction between, say, treatment foster care and regular foster care (as one example). Specifically, what is the distinction between these two types of care and how does each type of care relate to the needs of children who would benefit from the services offered. In turn, clarity as to the clinical purpose of a particular type of care would help DFPS understand whether children are getting the care needed when they need it, questions that lie at the heart of contract monitoring. The answer to those questions depends on clearly differentiated service populations and a clear understanding of how the services offered in different placement settings are supposed to affect safety, the likelihood of permanency and the well-being of children targeted for those services.

5 Conclusion

By their nature, child protection systems are complicated: raising the children of others, even if only for a short time, has never been easy. When DFPS, together with its stakeholders, set out to improve the legacy system, it did so with a set of core principles in mind: success should be judged on whether outcomes improve; private sector partners should be given greater latitude over decision-making but only in exchange for greater accountability; and, the fiscal model underpinning reimbursement should align policy and practice with fiscal incentives by encouraging the use of family-based care whenever possible. Wisely, DFPS adopted a stage-in, long term approach to implementation. Starting with the transfer of responsibility using a blended rate and concluding with a more diversified, comprehensive risk sharing model, the plan recognizes the magnitude of the changes required from everyone. Deep structural change requires time, patience, *and* a learn-as-you-go approach. The addition of a Community-based Care model to the original Foster Care Redesign is but one example of an adaptive strategy at work. The shift to a more robust method of building a blended per diem rate is another. Put simply, improvement is a continuous, long-term undertaking informed by the answer to a single question: are the outcomes for children improving?

6 Appendix

To solve the stratification problem, in the absence of service level data, we tried three different approaches, all of which relied on historical data to varying degrees. The solutions tested include a population projection that predicts the expected mix of children in the population going forward; a predictive risk model that assigns children to levels of care based on a set of characteristics; and, a factor model that uses characteristics of children to project the average cost per child per day rather than the level of care. As the last of the models tested was the model preferred by DFPS, we discuss the details of the data that informed the model building process. Details of the other models we tested are available.

6.a How was the history of conservatorship used for the analysis?

Although the idea of placing a child in paid conservatorship (also known as paid foster care) is relatively simple, the reality is quite different when measured as the placement history of each child placed in out-of-home care. Children may enter paid foster care, leave paid foster care and then return. When in paid foster care, children may change service levels. Changes in service level may happen once or many times and in any order. In addition, children may experience time in unpaid conservatorship as in the case of a placement with unverified kin. The methods used to blend rates have to capture this complexity because the most expensive care is often linked to children who have the most disrupted patterns of care. If the methods adopted for the analysis systematically miss this complexity, the rates set will be downwardly biased.

To build the analytical file needed to capture each child's full placement history, we devised three interrelated units of analysis. The core unit of analysis is the *spell of paid conservatorship*. The spell of paid conservatorship is sometimes referred to as a *spell of paid foster care*. They are identical from a definitional perspective, so we use the terms interchangeably. The *spell of paid conservatorship* refers to a *continuous period of time spent in paid conservatorship* (i.e., there is a start and a stop date). Paid conservatorships start with placement into some type of setting that is *reimbursed by DFPS*. Paid conservatorship ends when the child is placed in unpaid conservatorship (i.e., unpaid foster care) or the child leaves conservatorship altogether. There are rules, previously established by DFPS, that determine how periods of paid conservatorship separated by unpaid conservatorship are linked. Specifically, depending on the reason and on how long the period of unpaid conservatorship lasted, the two spells of paid conservatorship may either be combined into one longer spell of paid conservatorship or two separate spells of paid conservatorship. When building the file, we followed those rules.[10]

The second unit of analysis is the *rate span*. Rate spans are periods of time (i.e., there is a stop and a start date) that always fall within spells of paid conservatorship during which the cost of care for each day is the same. Rate spans end when the rate changes or the child leaves paid conservatorship. Rate spans are tied to the type and level of care provided. Historical data captures both the type care and the service level.

The third unit of analysis is the *child spell*. The child spell is a continuous period of time during which the

[10] The details of those rules are available upon request.

child is in conservatorship, regardless of whether the conservatorship is paid or unpaid (i.e., paid or unpaid foster care). Child spells may consist of unpaid conservatorship only, paid conservatorship only, or both. Paid and unpaid conservatorships, when both are present, may occur in any order.

Although the three units of analysis are defined separately, they are closely connected, as shown in Table 1. The main link between the units of analysis lies in the fact that rate spans are nested within a spell of paid conservatorship and spells of paid conservatorship are nested within child spells. The examples in Table 1 portray this structure. The unit of analysis is found in column one; the counter in column 2 shows the hierarchical or nested structure of the data.

Depicted is the (mock) history for Child A who experienced four separate child spells. The first, second, and fourth child spells all involved spells of paid conservatorship. The third child spell involved unpaid conservatorship only whereas the second spell involved paid conservatorship only. There were four spells of paid conservatorship, one each in child spells one and two. Child spell four consisted of two spells of paid conservatorship separated by a spell of unpaid conservatorship. Each spell of paid conservatorship consisted of at least two distinct rate spans. Two spells of paid conservatorship (during child spells one and four) had three distinct rate spans. Again, rate spans refer to the periods of time during which every day of care was reimbursed at a specific rate based on both the type and level of care. Of course, the care history of most children is not so complicated, but for the children whose history is complicated, the structure depicted in Table 1 captures all the detail in an organized fashion suited the analysis undertaken.

For the analysis we did, we generated a variety of summary data from this file:

- Total length of time in the spell of paid conservatorship (i.e., LOS) regardless of fiscal year
- Total spending during the spell of paid conservatorship regardless of fiscal year
- Number of different rate spans during the spell of paid conservatorship
- Number of prior child spells
- Child specific, weighted daily rate derived by dividing the total spending for the child during the spell of paid conservatorship by the service level during the spell of paid conservatorship. This is a weighted daily rate because we use the total days and the total spend during each spell of paid conservatorship. An unweighted daily rate would be the simple average of the rates associated with each distinct rate span. For example, the unweighted rate for a child who spent time in care at $100 per day and $50 per day would be $75 day. The weighted rate includes the number of days at each rate level. Thus, if this same child spent 20 days in care $50 per day and 10 days in care at $75 per day, the weighted rate for the total time in care would be $58.33 ($1750/30 days). The calculation of the rate in this way is equivalent to the blended daily rate (BDR) for that child.

Table 1: History of Conservatorship Expressed as Child Spells, Paid Conservatorship Spells, and Rate Spans

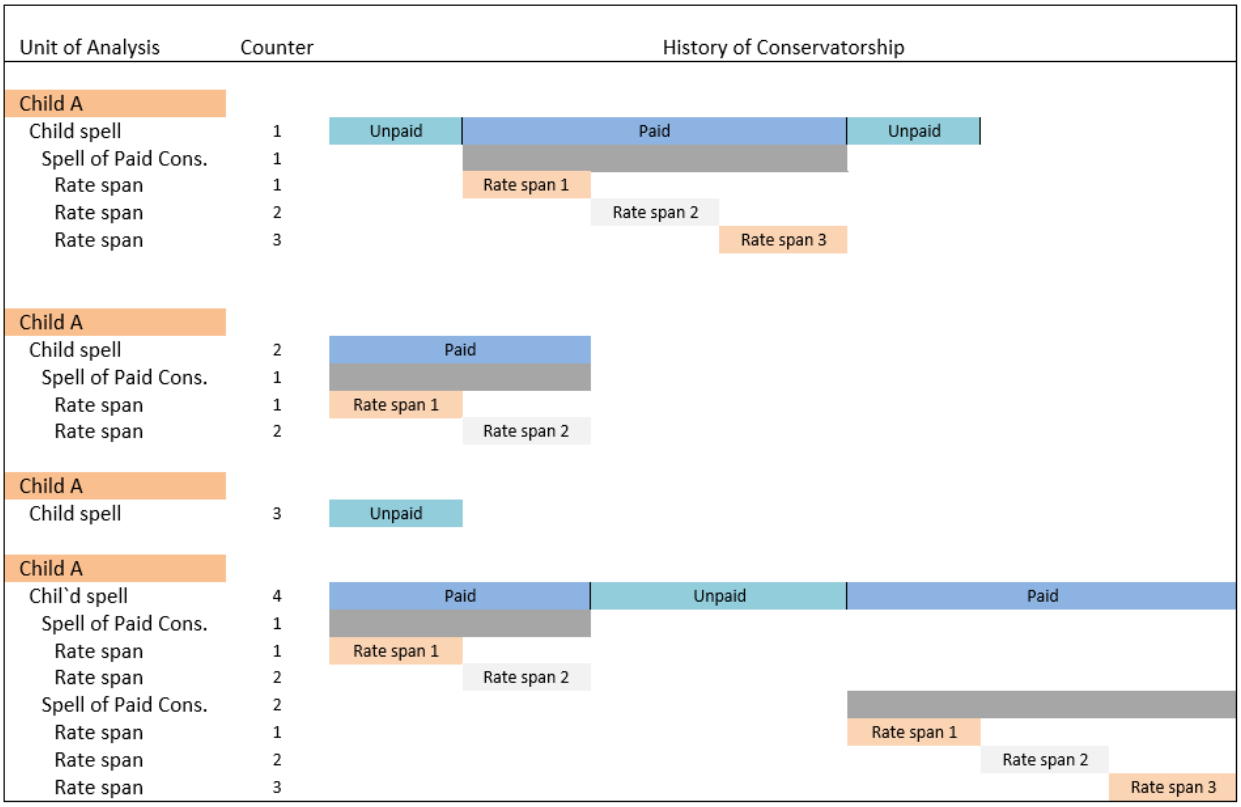

In addition to the summary information described above, we were able to summarize the data in other ways. For example, to predict the level of care at the start of the very first child spell, we selected the very first rate span from the very first child spell. To predict the level of care conditional on the first type of care, we selected the rate span in the child spell that followed the original rate span.