Independent Review of Abt Associates Needs Assessment Report

Steven Raphael
Professor
Goldman School of Public Policy
University of California, Berkeley
2607 Hearst Avenue
Berkeley, CA 94720-7320
Tel: (510) 643-0536
E-mail: stevenraphael@berkeley.edu

Prepared for

Terry Hill, MD.
CEO, Medical Services
California Prison Receivership

January 16, 2009

This memo provides an overview and evaluation of the Abt Associates research report titled "Chronic and Long-Term Care in California Prisons: Needs Assessment." The memo also includes a discussion of the June 13, 2008 update to the original report that takes into account new population projections by the California Department of Corrections and Rehabilitation (CDCR) as well as the potential impact of the early release of approximately 20,000 inmates. I first provide an overview of the goals, methodological approach and conclusions of the original assessment and the update. I then provide a critical assessment of the final estimates.

To summarize my conclusions, the sampling strategy employed by Abt is certainly sensible and likely captures inmates most in need of long term care outside of the general population. The procedures used to impute long term care need among general population inmates from a survey of medical beds are likely to include some biases, although the identifiable biases are often opposite in sign and are likely to cancel one another out. While there is substantial uncertainty in any estimate based on probabilistic imputation, the methods employed in the Abt study are certainly sound and, given the practical constraints on the research estimate, represent a sensible set of compromises.

Regarding the June 2008 update, while it does not make much of a difference to the overall numbers, I prefer the estimates based on the CDCR population projections to the original estimates based on the linear growth interpolations calculated by Abt. My assessment regarding the effect of an early release of 20,000 inmates is that in the medium to long run the impact of such a release on long term care needs is likely to be immaterial. Absent fundamental change in sentencing policy, the impact of this release on the overall prison population will be quickly undone. In light of this fact, I believe it would be a mistake to reduce estimated future needs based on a proposed one-time early release.

1. Overview of the original assessment and June 13, 2008 Update

The main objectives of the Abt studies reviewed here are to estimate:

- the total long-term-care (LTC) needs of the CDCR inmate population,
- the distribution of LTC inmates across alternative level of care (LOC) categories that would (or should) be housed outside of the general population, and

- the future LTC and LOC needs of the CDCR inmate population given projected changes in the overall inmate population as well as changes in the internal age composition of California's state prison inmates.

Combined with a census of existing LTC resources in California's 33 state prisons, the total needs assessment provides a gauge of the degree to which current and future LTC needs exceed existing capacity, and thus provides estimates as to how many new LTC beds should be added to the system to prepare for the future.

Data gathering

The estimates of current overall needs as well as future projected needs are based on two data gathering efforts undertaken by Abt with the aid of Lumetra, the CDCR medical staff, and CDCR corrections officers. The first data gathering effort was a complete census of medical beds in the system's 33 facilities. The medical bed census involved counts of all inmates in medical beds being treated for physical ailments (i.e., not for mental illness) either within the prison itself or outsourced to a local medical facility. For each inmate, information was collected on basic demographics (from prison administrative records), detailed medical histories and current diagnoses (from medical charts), information on activities of daily life (ADL) as well as prison activities of daily life (PADL) limitations from proxy interviews with medical service providers, and finally an assessment from the CDCR medical staff regarding whether the inmate required long term care and if so, the appropriate level of care.

The second data gathering effort involved a stratified random sample of the general population (GP) of inmates in nine of the 33 state prison facilities. The purpose of this second data set was to estimate LTC needs among those inmates not currently occupying a medical bed. The sampling frame for this data set was designed to maximize the likelihood of capturing GP inmates with LTC needs. Given that a relatively small proportion of GP inmates are in need of long term care (perhaps less than 1.5 percent), Abt constructed a sampling frame that over-sampled inmates with a higher likelihood of LTC needs. The basic sampling frame worked as follows. First, using data on those with LTC needs from the medical bed census as well as administrative data for all inmates in the state's 33 prisons, Abt first estimated a simple model where the probability of needing LTC was modeled as a function of variables observable in

administrative records (age, prior hospitalizations, etc). Abt then used this model to fashion a probabilistic prediction for each GP inmate in the system regarding the likelihood of needed LTC. Those inmates with relatively high predicted probabilities were deemed to be at "high risk" of needing LTC while all other inmates where deemed to be "low-risk."

Next, correctional officers in each of the nine facilities were asked to nominate GP inmates that they believed to be in need of LTC outside of general population facilities. The use of correctional officer nominations was based on the experience of an earlier data gathering effort by Lumetra that found correctional officers to be quite knowledgeable about the unmet needs of current GP inmates, and willing to make such nominations. The nomination process coupled with the pre-classification of inmates into high- and low-risk groups then defines four strata of inmates from which to sample:

- high-risk inmates nominated by correctional officers,
- high-risk inmates not nominated by correctional officers,
- low-risk inmates nominated by correctional officers, and
- low-risk inmates not nominated by correctional officers.

The sampling frame employed allocated observations across these strata in a manner designed to minimize the sampling variance of the LTC population needs estimate. This involves over-sampling stratum where the incidence of unmet LTC needs is likely to be the highest and under-sampling stratum where the incidence of unmet LTC needs is likely to be low. To maximize precision of estimate of LTC needs among the riskiest groups, Abt chose to sample from the top three strata only and employed an indirect method to estimate LTC needs emanating from the lowest risk group (the last stratum listed above).

The survey of the general population gathered much of the same data that was gathered for inmates captured by the census of medical beds, although there are some key important differences that figure prominently in the methodological imputations employed by the study. The GP survey includes information on demographics from CDCR administrative records. The survey also gathered information on medical histories and current diagnoses from medical charts. The survey also collected information on ADL, PADL and cognitive limitations from proxy

interviews (although here proxy interviews were conducted with correctional officers rather than nurses).

The main difference between the data gathered in the GP survey and that gathered in the medical bed census concerns the evaluation of LTC needs. While correctional officers were asked whether they believed the person would eventually require care outside of the general population, the GP survey included neither a LTC nor a LOC needs assessment by a medical professional.

Estimating LTC Needs Among GP Inmates

The study combines information from the medical bed census along with information from the GP survey to estimate LTC needs among the surveyed GP inmates in the nine prisons as well as the likely LOC needs of inmates with some positive probability of needing long term care. Since the nine prisons surveyed were chosen to be representative of the 33 state facilities, and since within each stratum inmates were randomly sampled, once LTC and LOC needs are determined, total LTC and LOC needs can be estimated by multiplying each LTC and LOC probability determination by the inverse of the sampling probabilities for each GP inmate included in the survey. In other words, the sampling frame permits a determination of the number of state inmates that each inmate included in the GP survey "represents." Thus, for example, if the likelihood that a given inmate in the sample will need LTC is 0.01 and if this inmate is representative of 10,000 inmates statewide, than the contribution of inmates such as the sampled inmate to overall LTC needs is estimated to be 0.01x10,000 = 100.

The sampling probabilities (and thus the number of inmates represented by each surveyed inmate) were under the control of the researchers and thus represented the easier of the two pieces of information needed to generate the LTC population estimate. Estimating the probability of LTC need and the accompanying LOC designation for each sampled inmate is somewhat more complex and required some additional modeling. Since much of my evaluation below focuses on these imputations, here I'll discuss these modeling methods in some detail.

To impute LTC need among GP inmates, the study exploits three facts. First, the information sets of the GP survey and the medical bed census overlap quite a bit, with both including information on ADL, PADL, and cognitive limitations as well as information from medical charts. Second, among inmates in the medical bed census, not all inmates required long

term care. Third, inmates in the medical bed census received LTC and LOC assessments from medical professionals (while the GP survey inmates did not). The Abt study first uses the medical bed census data to estimate a model where the dependent variable is whether one needs long term care and the explanatory variables include all of the variables on physical and cognitive limitations as well as medical chart information that are included in both data sets. The parameters of this model are then selectively applied to high risk (to be discussed below) general population inmates to generate a LTC probability prediction.

The second probabilistic model estimated with the medical bed census data is a multinomial logit where the dependent variable is the level of care designation (specialized general population, low acuity, high acuity) among those deemed in need of LTC and the explanatory variables include all of the variables in the common information set for the two data sets. The multinomial logit model assigned a probability for each individual for each level of care, where the sum of the probabilities across the three possible levels must equal one. This second model is used to apportion all those with some positive probability of needing LTC across LOC categories. Again, the imputation involves first estimating the model with medical bed census data and then applying the model to all GP inmates with some positive predicted probability of LTC need. In conjunction with the weighting scheme derived from the sampling frame, the second stage permits a breakdown of current and future long term care projections by level of care categories.

Generating future needs projections

In addition to estimates of the current LTC and LOC needs, the Abt study and the follow up report present estimates of the likely future LTC needs. The estimates in the original report are based on linear interpolation of population growth rates and age-specific estimates of LTC needs from the assessment discussed above. The June 2008 update substitutes long-term population projections provided by the CDCR research staff for the population estimated calculated by the Abt researchers in the original report.

Summary of findings

Both the original and updated estimates place LTC bed needs for the year 2007 at more than 2,900. This exceeds current medical bed capacity of the 33-prison system. Both the original and updated estimates project LTC bed needs exceeding 5,000 by the year 2018.

2. Critical Evaluation of the Needs Assessment Methodology

With unlimited time and resources, the ideal manner of gauging LTC need would have been to medically assess each inmate in the system and to have a uniform assessment of (1) whether LTC is needed, and (2) the needed level of care. Of course, this would have been prohibitively expensive. The next best method would have been to draw random samples of the inmate population and have the same medical professionals that made LTC and LOC determinations for the inmates in medical beds deliver similar assessments for the sample of general population inmates. My guess is that given the vulnerability of the population in question, the human subjects hurdles associated with selective medical assessments of inmates would also have been prohibitive. Such assessments for those in medical beds was of course facilitated by the developed familiarity that the attending medical staff already possessed and thus, further examination was not necessary. The same would not apply, however, to a random sample of inmates not currently under medical supervision.

This being said, the research team probably had little choice other than to attempt to model the relationship between observable medical characteristics on the need for LTC. Moreover, since the researchers only had LTC assessments for those in medical beds, they were indeed restricted to using the medical bed census data to model this relationship and impute LTC needs for the general population of inmates. My comments on the LTC estimates for the general population focus on three aspects of this imputation. First, I discuss the likely directions of bias associated with applying the parameters of a LTC model estimated with the medical census data to the general population. Second, I discuss the multinomial logit apportioning of LTC shares across LOC categories. Finally, I discuss issues associated with the imputation of LTC need among the low-risk, not nominated strata of GP inmates that were not sampled.

Imputing LTC needs with a model estimated with medical census data

The main tool for imputing LTC needs to the general population is a logistic regression model of the form

Prob (LTC is needed) = F(ADL limitation, PADL limitation, cognitive impairment, age, ….)

where "Prob (LTC is needed)" refers to the likelihood that a given inmate needs long term care (with the probability bounded to being between 0 and 1), F(…) is the cumulative distribution function for a variable with a logistic distribution, and the variables within the function F(.) are the pre-specified variables that are likely to impact the need for long term care. When the model is estimated with data from the medical bed census, the estimated impact of any one variable on the LTC probability is roughly related the strength of the correlation between the variable in question and a medical assessment of LTC need after accounting for the relationship between the given explanatory variable and all other explanatory variables in the model. In most such models, a constant term is included in the specification that in this application can roughly be thought of as representing the base LTC probability of an inmate in a medical bed without any of the exacerbating or mitigating factors captured by the explanatory variables included in the analysis. Once this model is estimated, the model can be used to calculated a probability of needing LTC for any inmate so long as the researcher is able to observe values for each of the explanatory variables used in the analysis. In this instance, since the information sets of the GP survey and the medical bed census overlap, the researchers are able to apply the model estimated with the medical bed census data to the GP survey data.

The key issue of concern in making such an imputation is whether it is appropriate to apply a model estimated on a group of inmates already receiving medical care in beds housed outside of the general population to a group of inmate not receiving medical care and residing in the general population. One would suspect that even for two inmates with identical observable characteristics in terms of the variables included in the two data sets, the inmate in the medical bed is likely to be much less healthy (and perhaps more in need of LTC) than the inmate remaining in the general population. This would be the case if there are dimensions of health that are not captured by the surveys but that are readily observable by medical staff. If this is indeed the case, applying the model fit to medical census data to everyone in the GP population will likely over-estimate the LTC needs of the system.

This potential selection problem is most evident in the base probability of LTC need predicted by the Abt logistic regression. The logisitic regression predicts that someone with no

cognitive impairment, major ADL limitation or permanent PADL limitation has a base rate need of LTC of 27 percent. Given the conclusion of the study that under 2 percent of the 2007 inmate population was in need of LTC, applying such a based rate to the general population of inmates is clearly too high.

The authors of this study are well aware of this problem and take steps to attempt to mitigate the upward bias. Specifically, the authors apply the probabilistic prediction only to those sample inmates who have at least one major ADL or PADL limitation or a cognitive problem. All other sampled inmates are assumed to have an LTC probability of zero.

Note, this imputation strategy imparts two biases to the estimate of LTC needs that are of opposite sign. First, while those inmates in the general population with an observable limitation are certainly likely to have a higher base level LTC probability relative to those without such an impediment, my guess would be that it is still likely to be below the base level for those presenting in medical beds. Of course, the implied upward bias to the LTC probability estimates is smaller than it would be if the logistic model were applied to the population more generally. Nonetheless, my prior expectation would be that for the impaired among the general population, the predicted LTC probability is still biased upwards.

Second, the strategy imputes a LTC probability of zero for all inmates in the three relatively high risk strata that do not have an ADL, PADL, or cognitive impairment. Clearly, this will bias the overall needs assessment downwards. Moreover, as this population is likely to be numerically larger than the impaired GP population, this bias may be particularly significant.

Absent further information, it is impossible to gauge which bias is likely to be more severe. My guess however, would be that these biases are largely offsetting and that the imputation strategy is likely to generate estimates of LTC need among the sampled strata of the general population that is in the ballpark of the true need.

Apportioning LTC need across LOC categories

The second stage in the needs projection involves modeling the likelihood that an inmate who needs LTC requires a specific level of care. In other words, conditional on needing LTC, what level of care is one likely to require? The empirical model defines three levels and uses those deemed needing LTC among those in medical beds to model these probabilities. When applied for prediction purposes, the multinomial logit model will spit out the probability that a

given person will land in each of the three levels, with the sum of the probabilities equal to one. When aggregated across a target population, the sums of these probabilities provide estimates of level of care need.

This is a quite standard tool for forecasting discrete outcomes where there are more than two possibilities (often used for transportation planning and modeling of commute mode choices). The use in the current application is entirely appropriate and unlikely to suffer from the potential selection problems encountered in the LTC imputation.

Imputing LTC need among GP inmates in the unsampled strata

Recall, the sampling strategy involves only sampling inmates who are either deemed high risk by the pre-sampling logistic regression or who are nominated by correction officers as in need of LTC. Those inmates deemed low-risk who are not nominated are not sampled. Thus, absent a further correction, such inmates would not contribute to the estimate of the long term care needs assessment.

Of course, ignoring this population could not be correct. By the authors own estimates, perhaps as much as 33 of the 164 inmates in medical beds in the nine surveyed prisons come from the general population inmates who were not sampled. Thus, the counts based on the three sampled strata must be adjusted upwards to account for the much more populous, yet un-sampled fourth stratum.

The study makes this upwards adjustment in the following manner. Assume that the ratio of low-risk to high-risk inmates in the general population needing LTC that is captured by the employed sampling strategy is equal to the comparable proportion of inmates that one would have measured in medical beds had the sampling frame been applied to all inmates inclusive of those currently receiving medical treatment outside of the GP. In the nine sample prisons, there are 164 inmates in medical beds, 33 of which are among the predefined low-risk group. The authors estimate that the correction officers' nominations correctly identify 24 percent of those inmates in need of LTC. Hence, of the 33 inmates in the low risk group in medical beds, roughly (1-.24)*33 = 25.1 inmates would have been missed by the sampling procedure. Thus, the sampling procedure would have correctly identified 84 percent of the inmates in medical beds [(164-25.1)/164*100], and to derive a correct total for medical bed inmates, we would need to

divide the total identified (164-25.1) by 0.84. This is equivalent to multiplying the count by an adjustment factor of 1.19.

Applying this correction to the overall count essentially assumes that the ratio un-sampled to sampled inmates among those in hospital beds is equal to the comparable ratio among those in the general population who require LTC. This is an impossible identifying assumption to validate, though I can offer speculations as to its reasonableness. Absent further information, my initial instinct would be that the proportion from the un-sampled among those with LTC needs is likely to be lower among the general population than among those in medical beds. I would imagine that the inmates in medical beds from this lowest risk group are perhaps more likely to be receiving treatment due to traumatic injury or the onset of unanticipated serious illness, suggesting incidence levels of conditions requiring LTC among the general population of this group that is lower than that observed among the higher risk stratum.

Nonetheless, this un-sampled stratum clearly contributes something to LTC needs among inmates. Moreover, the number in need of LTC for 2007 attributable to this adjustment factor is only 417 beds (out of a total of roughly 2,900). If we suppose that the true needed upwards adjustment is only half this size (209 beds needed for the un-sampled GP inmates) this bias would only amount to 7 percent of projected needs for that year. Thus, while it is difficult to evaluate the direction of bias created by this assumption, the magnitude of any such bias is likely to be numerically benign.

Summary assessment of the LTC imputations

This section three alternative biases to the LTC assessment, two of which I believe to be positive, one of which is unambiguously negative. While it is impossible to infer the magnitude of these biases, my best assessment would be that they are likely offsetting and that the LTC needs assessment presented in this report are close to the truth.

3. Assessing long term needs assessments

In June of 2008, the Abt research team presented a set of updated needs assessment estimates that made two principal changes. First, the research team recalculated future needs projections using long term population projections calculated by the CDCR. Second, the research team considered the impact of a one time early release of roughly 20,000 inmates.

The new population projections are generally lower than those resulting from the linear interpolation of growth rates presented in the earlier report. However, the two projections are still quite similar to one another and thus this particular update makes very little difference. Nonetheless, it seems advisable to base projected needs on the population projections estimates by the CDCR research staff.

The update also presents a series of long term needs assessments assuming a one-time release of 20,000 prisoners. The authors present subsequent estimates with and without recidivism by the released inmates. My thinking with regards to the impact of a one time release on the state's prison population is that in the absence of meaningful reforms to sentencing and parole policy, the temporary reduction will be undone quite quickly by the return of released inmates on parole violations and by prison admissions via new court commitments. This being said, I believe that it would be a mistake to project lower needs based on a one time release. An overly simple example will serve to illustrate this point.

Imagine a fictional society where the rate at which people break the law and the rate at which they are prosecuted and sent to prison generate 10,000 prison admissions per month. Suppose further that we begin with no prison inmates and that each new prisoner admitted to the system will serve exactly one year. In the first month, 10,000 inmates will be admitted to prison while nobody will be released (since by assumption there is nobody to release), giving us a prison population of 10,000. In month 2, we would add another 10,000, but still nobody would be released, since the prior month's admissions are serving the second month of a twelve month sentence. Thus, our prison population would climb to 20,000. This process would continue until month 12 when the prison population reaches 120,000 inmates. In month 13, however, the first 10,000 prisoners admitted in month 1 will be released offsetting the 10,000 new admissions. In fact, for each month afterwards, the flow out of prison (the 10,000 inmates whose sentences expired) will be exactly matched by the flow into prison (the new 10,000 inmates) and our prison population will stabilize at 120,000.

Now suppose that from this stable point we decide to attempt to reduce the prison population through a one-time early release of all inmates who are scheduled to be released over the next six months. This will indeed reduce the prison population in the month of the early release by 50,000 inmates (the reduction is not 60,000 since the releases are offset by 10,000 new admissions). In the subsequent five months no new releases will occur, since those

previously scheduled to be released during these months were released in the one time acceleration. However, in the absence of sentencing and parole reforms, we will still experience new admissions on the order of 10,000 per month. In fact, new admissions may rise above this level, as many of those whom we released early will violate parole and be sent back to prison. In short order, new admissions during the five month period following the early release will drive the inmate population back up to 120,000. Thus, the early release bought only a temporary reduction in the prison population.

Permanently reducing the prison population requires either (a) permanently reducing the rate at which people enter the system or (b) permanently reducing the amount of time that inmates serve. In more technical terms, the "steady-state" prison population will depend on these factors alone. Any one-time releases will certainly lower the prison population in the short run, but such policy shocks will have no medium to long run effect on the state's incarceration rate.