Estimating Variation in Productivity Across State Medicaid Programs: Evidence from Dual-Eligibles (Part 2) *

Timothy Layton† Nicole Maestas‡ Daniel Prinz§ Mark Shepard¶ Boris Vabson‖

March 24, 2025

Abstract

Many social programs involve some level of local autonomy, with local governments making some program design decisions within a set of national guidelines, and with financing being shared between local and national governments. We study the extent to which this autonomy results in heterogeneity across U.S. states in the productivity of their primary social health insurance program: Medicaid. We do this for one of the most expensive groups of Medicaid beneficiaries—those also enrolled in Medicare, i.e. "dual-eligibles". Productivity is typically defined as the ratio of outputs to inputs. For duals in Medicaid, we define inputs as the fiscal cost of the Medicaid program. We define output as providing access to healthcare goods and services, specifically access to primary care, critical surgeries, and, most importantly, long-term services and supports not covered by the Medicare program. In Part 1 of this series, we leveraged duals who move across states to estimate state effects on fiscal costs, showing significant variation across states in these costs (a ratio of 3:1 for the highest to lowest spending states). In Part 2, we show similar variation in utilization of long-term care.

*The research reported herein was performed pursuant to grant RDR18000003 from the US Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA, any agency of the Federal Government, or NBER. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

†University of Virginia and NBER. Email: timothyjlayton@virginia.edu

‡Harvard University and NBER. Email: maestas@hcp.med.harvard.edu

§World Bank. Email: dprinz@worldbank.org

¶Harvard University and NBER. Email: Mark_Shepard@hks.harvard.edu

‖Harvard University. Email: vabson@hcp.med.harvard.edu

Social insurance programs account for an increasing share of current and future government spending for countries around the world. They also have a major impact on people's lives, affecting health, wealth, employment and labor supply, human capital investment, fertility, family structure, and overall levels of well-being. But while we know that the *existence* of these programs has important impacts, outside of the effects of overall levels of generosity, little is understood about how these programs are designed and the consequences of different design decisions.

In this paper, we continue our work studying heterogeneity in the productivity of social insurance programs. We focus on the largest social insurance program in the United States in terms of number of beneficiaries, a program where there is significant flexibility for states to design their own programs: Medicaid. As in our prior work, we focus our analysis on one of the most expensive groups of Medicaid enrollees: The elderly who are also eligible for Medicare. In prior work, we studied heterogeneity in Medicaid *spending* for duals across states. We leveraged duals moving across states to separate variation in spending due to differences in the composition (in terms of both characteristics and preferences) of duals across states from variation in spending due to *causal* state effects on spending. We showed that a significant portion of the overall variation in spending was indeed causal, suggesting significant differences in Medicaid program design across states.

Ultimately, however, we are interested in uncovering not heterogeneity in *spending* across states but heterogeneity in *productivity*. Some states that choose to spend less may do so by achieving efficiencies, i.e. they may produce similar output for less money. Such efficiencies are desirable and one would hope to spread these efficiencies across states. However, other states that choose to spend less may do so by producing less. Such spending reductions are less-desirable, as they do not reflect efficiencies but just lower levels of output. State preferences for lower spending and lower levels of output may be completely valid, but it may not be desirable for other states that wish to lower spending but maintain levels of production to mimic the program design of these states. Worse still, some states may spend *more* and achieve similar or *lower* levels of output. Such states could potentially improve efficiency by mimicking states with similar spending and higher levels of output.

Motivated by this, in this paper we move beyond estimating heterogeneity in spending across states and toward estimating heterogeneity in a key measure of output: Utilization of long-term care. As shown in our prior work, long-term care is responsible for a large share of Medicaid spending for duals (Layton et al. 2023). Long-term care comes in two forms: Institutional care (i.e., nursing homes) and non-institutional care (also known as home- and community-based services, or HCBS). We focus on institutional long-term care for three reasons: (1) according to our prior work it still makes up the majority of Medicaid long-term care spending in most states, (2) there is high-quality administrative data tracking utilization of institutional long-term care consistently across states, and (3) home- and community-based services are difficult to identify and measure consistently in claims data, the only source available with samples larger enough for the types of analyses we perform here.

We again start by documenting variation across states in utilization of long-term care. Similar to our work studying variation in spending, we document substantial variation in utilization of long-term care across states. In Connecticut and Indiana, almost 20% of duals have an episode of institutional long-term care in any given year, while fewer than 5% of duals have an episode in Oregon, Vermont, and Alaska.

The variation is even larger when incorporating the intensive margin as well as the extensive margin: In Connecticut and Indiana, the average dual uses around 60 days of institutional long-term care in a given year, while in Oregon, Vermont, and Alaska duals use fewer than 10 days. This variation persists when focusing on a set of states with high-quality Medicaid claims data, which is not surprising given that the utilization measures do not come from claims but instead from higher-quality administrative data (the Minimum Data Set, or MDS).

We then follow our prior work by decomposing this variation into variation due to differences in the composition of duals (in terms of characteristics and preferences) across states and variation due to *causal* state effects on use. We again do so by leveraging duals who move across states. We show that around 40-45% of the variation is causal, with the remaining variation stemming from differences in the composition of duals across states. This decomposition is similar for various measures of utilization, suggesting an important role for program design.

In the last part of the paper, we study state effects on spending and utilization simultaneously. Specifically, we ask whether states that causally increase spending also causally increase utilization of long-term care. Ultimately, we find this to be the case, though we find the relationship to be quite weak. We find that moving to a state with a 10% higher state spending effect leads to an increase in the probability of using any institutional long-term care during the year by 0.06 percentage points, or 0.7% of the baseline level of use. Similarly, we find that such a move would increase the number of days in long-term care by 0.18 days, or, again, around 0.8% of the baseline level of use.

These results suggest that while there is substantial heterogeneity across states in utilization of long-term care, only a portion of that variation is causal in nature. Further, while causal state effects on spending duals and duals' utilization of long-term care are positively correlated, that correlation is weak, with higher spending associated with only small gains in utilization. A potential implication of this result is that some states spend much more and get little in the form of increases in long-term care utilization for these higher levels of spending. Similarly, one could interpret these results to mean that states that are successful at reducing utilization of institutional long-term care (a common goal among Medicaid programs) are not seeing significant savings from these efforts. One potential explanation for this result is that states that decrease utilization of institutional long-term care do so by spending more on other forms of care, such as non-institutional forms of long-term care, and this spending almost fully offsets the savings achieved. Importantly, this shift to other forms of long-term care may be welfare-enhancing even if savings are trivial, as beneficiaries may prefer to remain at home and receive care there rather than being admitted to a nursing home. Indeed, another potential way to interpret our results could be that states are able to decrease reliance on nursing homes without 'breaking the bank' a consistent fear across state Medicaid programs looking to expand access to home-based services due to potential woodwork effects.

Overall, however, interpretation is difficult without understanding substitution between institutional and non-institutional long-term care. Indeed, in future papers in this series, we hope to estimate effects on measures of utilization of non-institutional long-term care. Such measurement is difficult, due to differences in tracking of these services across states in claims data. However, we hope to develop high-level measures nonetheless and build out a more complete picture of true 'productivity' differences across state Medicaid programs.

1 Data Description

We use national administrative enrollment and claims data from Medicare and Medicaid.

Enrollment Data: The primary data source is linked administrative data from the Centers for Medicare and Medicaid Services (CMS) covering the Medicare and Medicaid programs during the 2008-2015 time period. The different Medicaid and Medicare administrative files are all linkable to one other, through a standardized beneficiary identifier.

The Medicaid and Medicare data includes separate enrollment files covering each of the programs. The observation level in all of the enrollment files is at a beneficiary-year level. The Medicaid as well as Medicare enrollment data track demographic information such as age, gender, and birth date. Importantly, the data also include the beneficiary's zip code, which comes from the Social Security Administration and is updated when the beneficiary moves. They also track each beneficiary's basis for eligibility at a given point in time, based on standardized eligibility codes. The eligibility information can be used to determine which Medicaid and Medicare enrollees are eligible due to disability vs. some other eligibility pathway. Birth date information, meanwhile, allow us to observe when individuals in the treatment group should move from Medicaid to Medicare. Finally, the enrollment files track actual enrollment status in both Medicaid and Medicare on a month-by-month basis, as well as whether an individual is simultaneously enrolled in both Medicaid and Medicare.

Utilization Data: Utilization data for this project comes primarily from the Minimum Data Set (MDS). The MDS includes records for every nursing home stay, regardless of payer. The records are reported by the nursing homes themselves, and they include detailed information on the length of an episode as well as detailed assessment data including information on a variety of measures of resident health and functional status.

The MDS is widely regarded as more reliable than standard Medicare claims data for analyzing patient care in long-term care settings, particularly nursing homes. Unlike claims data, which primarily captures billing information related to services provided and reimbursement, the MDS offers detailed, clinically rich assessments of patients' physical, cognitive, and functional status. These assessments are conducted by healthcare professionals and updated at regular intervals, such as upon admission, quarterly, and when significant health changes occur, ensuring consistent and up-to-date information. Additionally, the MDS is subject to regulatory oversight by the Centers for Medicare Medicaid Services (CMS), which mandates accuracy and completeness as part of facility compliance. This level of detail, combined with standardized data collection across all Medicare- and Medicaid-certified facilities, makes the MDS a robust and reliable dataset for research focused on patient outcomes, quality of care, and the economics of long-term care, providing insights that claims data alone cannot offer.

The MDS underwent an overhaul in the middle of our sample period, as they shifted from MDS 2.0 to MDS 3.0. This overhaul involved some changes to the structure of the dataset that resulted in changes in the reported number of nursing home episodes. These effects were not always consistent across states. This can be seen in Figure 1, where Iowa (left panel) sees a significant increase in episodes between 2009 and 2011 (when the transition to MDS 3.0 occurred), while Idaho sees no such increase. We address this issue in our analyses by controlling for the version of the MDS interacted with state dummies.

State spending effects: For some analyses we will use estimates of state effects on spending for duals. These estimates come from our prior work in Layton et al. (2023).

2 Descriptive Statistics

We now describe the raw variation in long-term care utilization across states for duals. For these descriptive statistics, we focus on a sample of duals who are enrolled in both Medicaid and Medicare for at least 9 months of the year in 2012.[1] Panel A of Figure 2 presents for each state the portion of duals in our sample who have any institutional long-term care episode. Panel B presents for each state the average number of days spent in institutional long-term care for duals in our sample.

Differences in utilization of long-term care are large. In the highest utilization states (Connecticut, Indiana), almost 20% of duals spend time in a nursing home during any given year. In the lowest utilization states, however, fewer than 5% of duals have a nursing home episode in any given year. Even more dramatic variation is observed for days spent in a nursing home. In the highest utilization states, duals spend an average of around 60 days in a nursing home in any given year. In the lowest utilization states, however, duals only spend fewer than 10 days in a nursing home. Figure 3 restricts to a set of states with high-quality Medicaid claims data (identified in our prior work). These panels show that the variation persists even when restricting to this group, not surprising given that these measures of utilization do not come from claims.

These differences in utilization of institutional long-term care are large. For context, in our prior work we showed that raw per capita spending differences across states had a ratio of around 3:1 between the highest and lowest spending states. Here, the utilization ratio is more than 4:1 on the extensive margin and 6:1 when measuring the total number of days spent in nursing homes. Given that all of these beneficiaries have the same insurance coverage, Medicaid plus Medicare, this level of variation is stark.

While these descriptive results suggest that state program design is important for determining levels of utilization of institutional long-term care, they could also stem from differences in the composition of Medicaid enrollees across states, both in terms of characteristics and in terms of patient preferences. These compositional differences could arise due to differences in state-level Medicaid eligibility rules or due to natural variation in where different people live. In the next section, we leverage a quasi-experimental framework for differentiating between actual state effects on long-term care utilization versus compositional differences.

3 Estimating State Effects

So far, we have documented wide variation in utilization of institutional long-term care across states. However, the source of this variation remains unclear. Prior work focusing on geographic variation in spending in Medicare posits two potential sources of variation: differences in the composition of beneficiaries (in health status and preferences for medical care) and differences in physician and hospital practice patterns (Finkelstein et al. 2016, Cutler et al. 2019). As we emphasize in our prior work, however, variation in Medicaid spending and utilization can also stem from a third factor: variation across states in Medicaid program

[1] For those enrolled for fewer than 12 months, we construct 'annualized' spending measures by dividing their total annual spending by the number of months they're enrolled and multiplying by 12.

design. Indeed, our goal in this paper is to isolate this third factor.

In this section, we attempt to isolate the variation in utilization of institutional long-term care due to state Medicaid program design. We start by separating variation in long-term care utilization across states into 'person' and 'place' factors, using the migration of Medicaid beneficiaries across states and a two-way fixed effects (TWFE) design. Our approach follows our prior work in Layton et al. (2023), as well as prior work by Finkelstein et al. (2016) and work in labor economics that considers the role of firms in earnings inequality (Abowd et al. 1999, Card et al. 2013). This decomposition will allow us to identify the portion of variation in long-term care utilization across states that is due to 'state effects'.

3.1 Empirical Framework

Following our prior work (Layton et al. 2023), we start from a simple statistical model of long-term care utilization:

$$y_{ist} = \alpha_i + \gamma_s + \tau_t + x_{it}\beta + \epsilon_{ist} \tag{1}$$

where i indexes patients, s indexes states, t indexes years, y_{ist} is a measure of utilization, the α_i are patient fixed effects, the γ_s are state fixed effects, the τ_t are year effects, and x_{it} is a vector of observable characteristics. The patient and state effects in this model can be separately identified if some of the patients move across states. We thus construct a sample of moving dual eligibles to estimate this model. Specifics regarding the construction of this sample are discussed in Section 3.2. Again, the sample construction mimics our prior work.

This model can also be represented as an event study that allows us to graphically display our results and assess the validity of the identifying assumptions. Let $o(i)$ denote patient i's origin state and $d(i)$ denote her destination state. (For patients who do not move, $o(i) = d(i)$.) Let $r(i,t)$ denote years relative to the year patient i moves (e.g. $r(i,t) = -2$ two years before patient i moves), which we will refer to simply as "relative years". To transform Equation (1) into an event study, we define $\delta_i = \bar{y}_{d(i)} - \bar{y}_{o(i)}$, the difference in average spending between patient i's origin and destination states. Combining the origin place effect $\gamma_{o(i)}$ with the patient effect α_i from Equation (1) into a single patient effect, $\mu_i = \alpha_i + \gamma_{o(i)}$, we can write down an event study as

$$y_{it} = \mu_i + \theta_{r(i,t)}\hat{\delta}_i + \tau_t + x_{it}\beta + \epsilon_{it} \tag{2}$$

where i indexes patients, t indexes years, y_{it} is a measure of utilization, μ_i is a patient effect, $\hat{\delta}_i$ is the estimated difference between patient i's origin and destination states in average utilization, the $\theta_{r(i,t)}$ are relative-year-specific coefficients, the τ_t are year effects, and x_{it} is a vector of observable characteristics (including an interaction between state and the version of the MDS to capture changes in reporting as discussed above). The $\theta_{r(,t)}$ coefficients show where the average utilization of movers falls for each relative year compared to the average utilization levels in their origin and destination states. For example, if on average movers' utilization is the same as the average in their origin state, then $\theta_{r(i,t)} = 0$ and if on average it is the same as the average in their destination state, then $\theta_{r(i,t)} = 1$.

The model makes several important simplifying assumptions that we previously highlighted in our prior work but that we repeat here for convenience, modified to correspond to the specific analyses in this paper.[2] First, health shocks that coincide with the time of the move *and* are correlated with utilization in the origin and destination would bias our estimates of state effects. One potential violation would be if the need for a nursing home triggers a move by a dual, *and* makes them more likely to move to a state with higher nursing home use. We provide a few pieces of evidence that our analysis does not suffer from this type of problem. First, as in our prior work, we find no evidence in our event study analysis of pre-move trends in utilization, showing that movers with different origin-destination differences in utilization are on similar spending trends (see Figure 4). Second, in our prior work we provided evidence of little correlation between the pre-move utilization of movers and the "size" of the move (i.e., the difference between origin and destination levels of utilization), suggesting little selection on pre-move spending levels, with any differences being absorbed by patient fixed effects α_i.

Second, we assume that the patient effects α_i and state effects γ_s are additively separable. This rules out the possibility that different types of patients would behave differently within a state. For example, the causal effect of a place cannot be larger for high-utilization patients. We test this assumption by including specifications measuring the extensive margin of utilization and specifications measuring utilization in days (capturing both the intensive and extensive margins.

Third, because state effects are identified only by movers, for state effects to have a broader interpretation, we need to assume that state effects for movers and non-movers are similar. If state effects are different for movers, our estimated state effects would not be valid for entire population. As one piece of evidence that movers and non-movers may be relatively similar, in our prior work, we showed that movers and non-movers look fairly similar on age and gender, though they do differ on healthcare spending. Medicare spending is somewhat similar, but Medicaid spending differs, with movers spending less than non-movers.

Fourth, this model does not allow past utilization y_{it} to influence patient effects α_i, i.e. habit formation. Habit formation would mean that the estimated patient effect α_i could be partly driven by past state effects γ_j. We test for this possibility via the event studies of the effects of moves. Habit formation would appear as a gradual convergence of beneficiary spending to the spending level of their destination state. If effects appear immediately, on the other hand, habit formation is relatively unimportant. Ultimately, we find the latter.

3.2 Sample Definition

For all analyses, we use the same sample we used in our prior work (Layton et al. 2023). However, we repeat the sample description here for convenience. For this exercise we use the Medicare files from 2007 to 2015. Because we are interested in identifying movers, we drop beneficiaries only observed for one year. We define a patient's origin state as the first state that they appear in.[3] We define non-movers as beneficiaries

[2] For a more detailed discussion, see Finkelstein et al. (2016).

[3] We use the Medi*care* data to identify state of residence, not the Medi*caid* data. The Medicare data continuously measures residence using data from the Social Security Administration. When duals move across states, they typically experience a gap in their Medicaid coverage, as they need to sign up for coverage in their new state (it is not automatic). During that gap, we do not observe state of residence in the Medicaid data, but we do in the Medicare data.

who are always in the same state and movers as beneficiaries who are in two states during the timeframe. For our event study, we use only movers and drop non-movers. We define as the move year (relative year 0) as the first year that a patient is not in their origin state (the first state they appear in) but in their destination state (the second state they appear in). We limit to beneficiaries with no additional cross-state moves in the five years after the initial move, including back to the original source state or to a third state.

Further, we limit to beneficiaries who are enrolled in Medicare for all twelve months in relative years -2, -1, and 0 and in Medicaid in relative years -1 and 1. (This allows for some time to switch Medicaid enrollment from the origin state to the destination state.) We also drop years with incomplete (fewer than 12 months) Medicare or Medicaid enrollment. For analyses using outcomes from Medicare data, we construct a separate sample of beneficiaries never enrolled in Medicare Advantage. Finally, we require that movers have at least 75% of their spending in their source state pre-move and in their destination state post-move. This should eliminate beneficiaries with multiple homes, such as "snow birds."

In our prior work, we provide strong evidence that the beneficiaries we categorize as movers do in fact move. Specifically, we show that the location of almost 100% of their claims moves to their destination state after the move.

3.3 Event Study Results

Panel A of Figure 4 shows our main event study results for any nursing home utilization during the year. The move takes place at $t = 0$. Prior to the move, coefficients are essentially zero and not moving significantly. This suggests that before a move the utilization of beneficiaries moving to places with high levels of utilization is trending similarly to the utilization of beneficiaries moving to places with low levels of utilization. This is important for the validity of our empirical strategy, as it implies that we can infer causal effects of the move from any changes at the time of the move.

In the year of the move, utilization converges discontinuously. In year 0, the coefficient jumps to around 0.4 and remains at that level through year 5. This implies that when a beneficiary moves, within a year their utilization of long-term care converges 40% of the way from their origin place's level to their destination place's level of utilization. Importantly, it remains there in perpetuity, implying that there does not appear to be much habit formation/state dependence. Instead, place effects appear sharp and persistent.

Panel B of Figure 4 shows the same event study but with the outcome being the number of months with some time spent in the nursing home instead of a dummy variable for any time spent in a nursing home during the year. Results are similar, with pre-move coefficients steady around zero, and a sharp, discontinuous jump at the time of the move. Here, the coefficient continues to grow during the year after the move, implying a bit of state dependence. But the coefficient holds steady after year one, implying that any state dependence is short-lived. Again, the long-run coefficient is around 0.4.

Table 5 reports the θ coefficients from the version of the two-way fixed effects where we estimate a single θ for the entire post-move period. The upper panel represents the full set of states, while the lower panel restricts to states with high-quality Medicaid claims data. Each column represents a different outcome, including any nursing home stay, months with a day spent in a nursing home, quarters with a day spent in a nursing home, and days in a nursing home. Starting with the top panel, coefficients range from 0.29 to 0.38, implying between 30 and 40% of the geographic variation in utilization of institutional long-term care is due to state effects rather than differences in the composition of duals across states. The bottom panel reveals somewhat higher coefficients for the states with high-quality claims data, but coefficients are only modestly different, suggesting a similar story in those states.

Overall, these results can be interpreted to imply that around 30-40% of the geographic variation in utilization of institutional long-term care is due to state effects rather than differences in the composition of duals across states. We believe this suggests that state effects on utilization are quite significant.

3.4 Spending versus Utilization

In our prior work, we used these same methods to estimate state effects on Medicaid spending for duals. Ultimately, we are not interested just in state effects on spending or on utilization in isolation. Instead, we are interested in variation across states in the *productivity* of state Medicaid programs, or variation in output per unit of input.

Here, we do not estimate productivity, but we take a first step in that direction. Specifically, we ask, on average do states that spend more achieve higher levels of "output" as measured by institutional long-term care utilization. To do this, we modify our two-way fixed effects regression specification slightly. Specifically, we replace the $\hat{\delta}$ terms (the estimated difference between the utilization in the origin and destination states) with the (log) difference between the *estimated state spending effects* in the origin and destination states. Thus, the new θ coefficients will no longer tell us how much utilization converges from the origin state levels to the destination state levels, but instead the θs will tell us how much long-term care utilization increases when a beneficiary moves to a state with a 1% higher state effect on Medicaid spending.

The goal here is to assess whether states with higher causal spending effects also 'produce' more output in the form of higher utilization of institutional long-term care. Positive coefficients would indicate this relationship, and the size of the coefficient will give a sense of how strong the relationship is.

Table 6 presents the coefficients from these regressions. Again, the top panel includes all states and the bottom panel restricts to states with high-quality Medicaid claims data. Here, this restriction may actually be important, as the state spending effects were estimated using claims data, so the quality of that data may matter for the right-hand side variables (the spending effects). We start with the first column, where the outcome is a dummy variable for any nursing home stay during the year. The coefficient of 0.006 can be interpreted to suggest that a move to a state with a 10% larger state spending effect yields an increase in the probability of using any institutional long-term care during the year of 0.06 percentage points. Given that on average 9% of beneficiaries use institutional long-term care in any given year, this represents a relative increase of about 0.7%. We find similar results for other outcomes, with moves to states with 10% higher state spending effects yielding an increase of 0.003 months with a day spent in a nursing home (0.8%), an increase of 0.002 quarters with a day spent in a nursing home (0.8%), and an increase of 0.18 days (0.8%) spent in a nursing home in during the year. Results are almost identical when restricting to states with high-quality claims data (lower panel).

Overall, these results suggest that higher spending states do 'produce' more. However, the additional amount that they produce is quite modest, with only very small increases in utilization for fairly large increases in spending. This suggests that there are likely to be states that are much more ‘productive’ than others in terms of producing long-term care utilization for a given level of spending.

4 Conclusion

Social insurance programs around the world exhibit varying levels of local autonomy in program design. In this paper, we continue to study one of the largest social insurance programs in the world that provides a great deal of local autonomy: Medicaid. In prior work, we documented significant variation across states in per enrollee spending among program beneficiaries also eligible for Medicare (“duals”) who rely on this program for coverage of long-term services and supports and supplemental coverage for other services (Layton et al. 2023). We also leveraged beneficiaries moving between states to show that much of this variation is driven by causal state effects on spending rather than differences in the composition of Medicaid enrollees across states.

In this paper, we presented similar results for utilization of institutional long-term care, the primary driver of Medicaid spending for duals. We showed that there is significant variation in utilization of long-term care across states and that a significant portion of that variation is due to causal state effects rather than just differences in the composition of duals across states. Finally, we combined these estiamtes with the estimates of state spending effects from our prior work, and we showed that states that causally increase the spending of duals do indeed ‘produce’ more institutional long-term care. However, the amount of additional long-term care that they produce for the increase in spending is very small, suggesting that some states are likely lower productivity (higher spending but similar output) than others.

There are a number of potential ways to interpret these findings. First, the findings could imply that some states spend much more but get little (in terms of institutional long-term care utilization) for that additional spending. This would be seen as a negative result that suggests that some states are much less efficient/productive than others. However, other interpretations are also possible. Importantly, many states are actually trying to reduce their reliance on institutional long-term care under the assumption that institutional care costs more and is less desirable to beneficiaries than non-institutional care like home- and community-based services (HCBS). One possibility is that states that have lower levels of institutional long-term care utilization offset the lower spending from that lower utilization with increased spending on HCBS. This substitution could be one-for-one in that all beneficiaries who would have used institutional long-term care now use HCBS, or more than one-for-one in that some beneficiaries who would not have used institutional long-term care also take up HCBS. In either case, if beneficiaries prefer HCBS to institutional long-term care, states that produce less institutional long-term care and more HCBS could be ‘more productive’ than similar spending states that produce more institutional long-term care and less HCBS.

A full assessment of productivity thus necessitates estimation of state effects on the utilization of non-institutional long-term care in addition to institutional services. Estimation of these effects is difficult due to data quality issues in the measurement of HCBS. Specifically, there are major problems with consistent measurement of HCBS across states. However, in future work, we will attempt to produce these estimates and build a more complete picture of the productivity of state Medicaid programs.

References

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67* (2), 251–333. 5

Card, David, Jörg Heining, and Patrick Kline, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 2013, *128* (3), 967–1015. 5

Cutler, David, Jonathan S. Skinner, Ariel Dora Stern, and David Wennberg, "Physician Beliefs and Patient Preferences: A New Look at Regional Variation in Health Care Spending," *American Economic Journal: Economic Policy*, 2019, *11* (1), 192–221. 4

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Sources of Geographic Variation in Health Care: Evidence from Patient Migration," *Quarterly Journal of Economics*, 2016, *131* (4), 1681–1726. 4, 5, 6

Layton, Timothy, Nicole Maestas, Daniel Prinz, Mark Shepard, and Boris Vabson, "Estimating Variation in Productivity Across State Medicaid Programs: Evidence from Dual-Eligibles (Part 1)," 2023. NBER RDRC Working Paper. 1, 4, 5, 6, 9, 16

5 Figures

Figure 1: Trends in Nursing Home Episodes in Iowa and Idaho

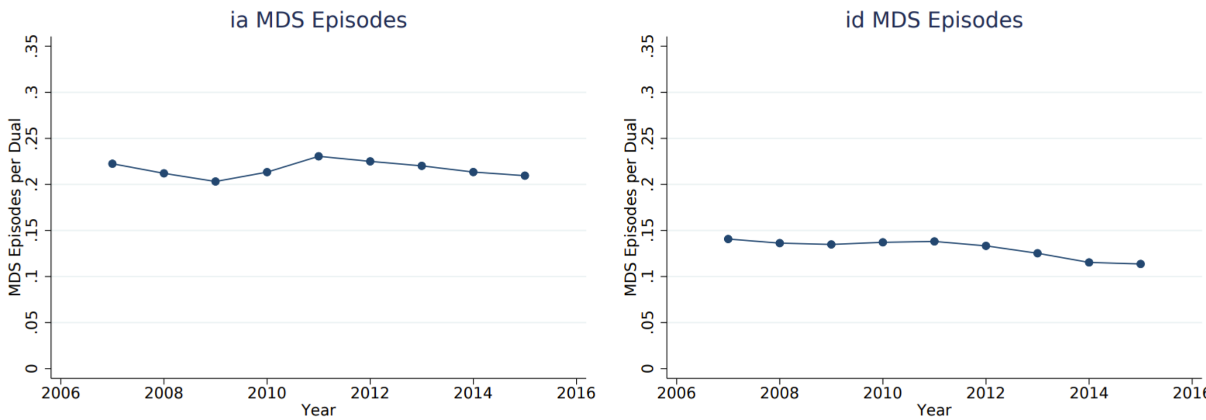

Note: The figure shows utilization of institutional long-term care over time as measured in the MDS. Both panels show the percent of dual-eligibles in a state that spent at least one day in a nursing home in each year. The left panel shows utilization in Iowa and the right panel shows utilization in Idaho. The purpose of the figure is to show how in some states (like Iowa), there was a shift in reporting with the shift from MDS 2.0 to MDS 3.0 in 2010-2011.

Figure 2: Variation across States in Institutional Long-Term Care Utilization among Duals - All States

(a) Any Nursing Home Episode

(b) Days Spent in Nursing Home

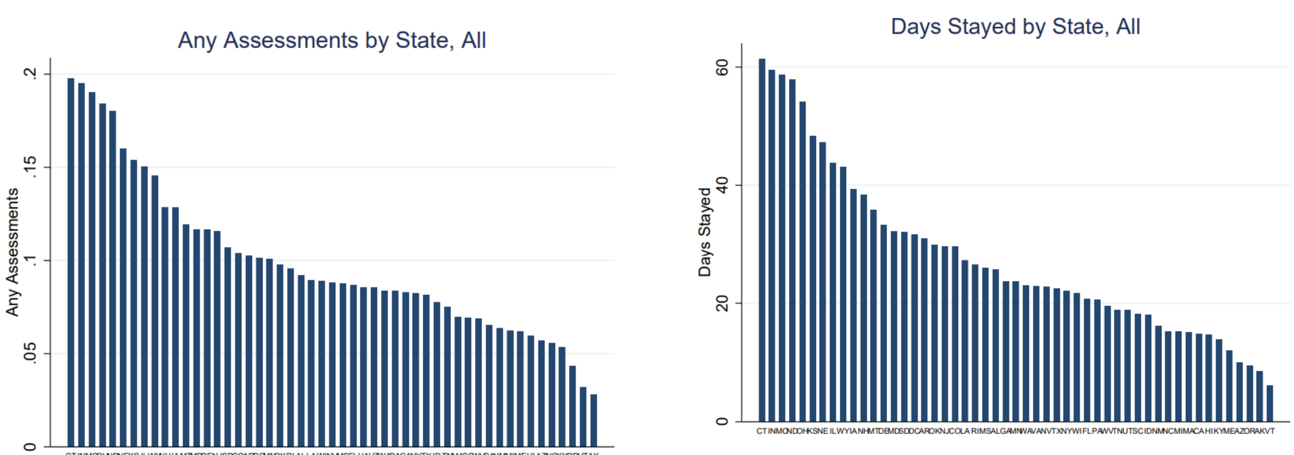

Note: The figure shows utilization of institutional long-term care among duals for each state. In the left panel, each bar represents a state and shows the portion of duals that state that spent at least one day in a nursing home during the year. In the right panel, each bar represents the average number of days spent in a nursing home per capita.

Figure 3: Variation across States in Institutional Long-Term Care Utilization among Duals - High Quality Data States

(a) Any Nursing Home Episode

(b) Days Spent in Nursing Home

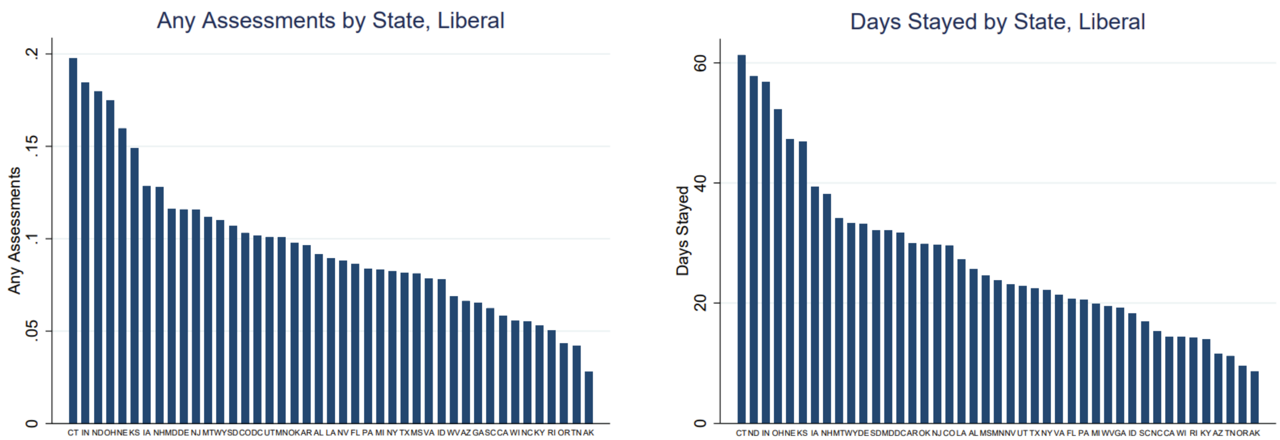

Note: The figure shows utilization of institutional long-term care among duals for each state, restricting to states with high quality Medicaid claims data. In the left panel, each bar represents a state and shows the portion of duals that state that spent at least one day in a nursing home during the year. In the right panel, each bar represents the average number of days spent in a nursing home per capita.

Figure 4: Event Study of Nursing Home Utilization Before Versus After Moves

(a) Any Nursing Home Episode

(b) Months with a Day Spent in Nursing Home

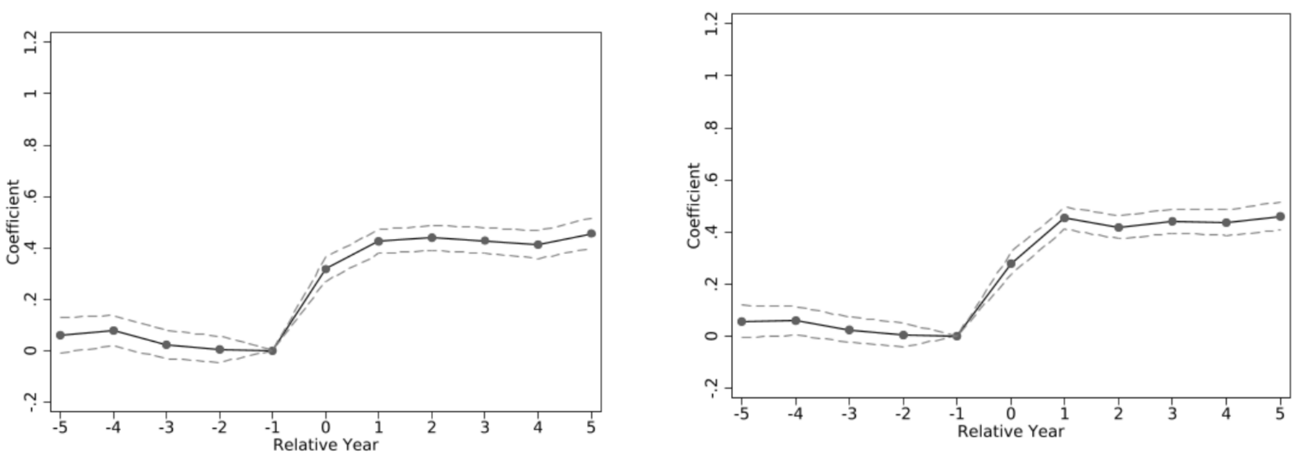

Note: The figure shows event studies of institutional long-term care utilization over time. Specifically, the figure shows coefficients from a two-way fixed effects model with person fixed effects and time fixed effects and a set of relative (to the time of move) time fixed effects interacted with a variable equal to the difference in the average level of utilization in the destination state versus the origin state.

Figure 5: Two-way Fixed Effect Regression Estimates of Effects of Moves to Higher Utilization States on Utilization

	(1) Assessments: Any	(2) Assessments: Months	(3) Assessments: Quarters	(4) Any Stay	(5) Days Stayed
All State Mean	0.382^{**}	0.383^{**}	0.326^{**}	0.380^{**}	0.290^{**}
	(0.013)	(0.012)	(0.010)	(0.013)	(0.010)
Observations	1.87e+06	1.87e+06	1.87e+06	1.87e+06	1.87e+06
Treatment Observations	70,984	70,984	70,984	70,984	70,984

Standard errors in parentheses
$^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

	(1) Assessments: Any	(2) Assessments: Months	(3) Assessments: Quarters	(4) Any Stay	(5) Days Stayed
Liberal State Mean	0.445^{**}	0.442^{**}	0.380^{**}	0.445^{**}	0.350^{**}
	(0.018)	(0.017)	(0.014)	(0.018)	(0.014)
Observations	1.55e+06	1.55e+06	1.55e+06	1.55e+06	1.55e+06
Treatment Observations	50,217	50,217	50,217	50,217	50,217

Standard errors in parentheses
$^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Note: The table shows coefficients from two-way fixed effect "value-added" models of the effects of moves on institutional long-term care utilization. The left-hand side variable varies by column. The right-hand side includes person fixed effects, time fixed effects, and the interaction between the destination-origin difference in average utilization and a dummy representing the "post-move" time periods.

Figure 6: Two-way Fixed Effect Regression Estimates of Effects of Moves to Higher Spending States on Utilization

	(1) Assessments: Any	(2) Assessments: Months	(3) Assessments: Quarters	(4) Any Stay	(5) Days Stayed
All Log Spending Effect	0.006**	0.028**	0.021**	0.006**	1.807**
	(0.000)	(0.002)	(0.001)	(0.000)	(0.116)
Mean	0.09	0.35	0.27	0.09	22.65
Observations	1,873,503	1,873,503	1,873,503	1,873,503	1,873,503
Treatment Observations	70,984	70,984	70,984	70,984	70,984

Standard errors in parentheses
$^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

	(1) Assessments: Any	(2) Assessments: Months	(3) Assessments: Quarters	(4) Any Stay	(5) Days Stayed
Liberal Log Spending Effect	0.007**	0.028**	0.022**	0.007**	1.831**
	(0.001)	(0.003)	(0.002)	(0.001)	(0.156)
Mean	0.08	0.33	0.26	0.08	21.92
Observations	1,547,532	1,547,532	1,547,532	1,547,532	1,547,532
Treatment Observations	50,217	50,217	50,217	50,217	50,217

Standard errors in parentheses
$^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Note: The table shows coefficients from two-way fixed effect models of the effects of moving to a higher-spending state on institutional long-term care utilization. The left-hand side variable varies by column. The right-hand side includes person fixed effects, time fixed effects, and the interaction between the destination-origin difference in state spending effects (as estimated in our prior work in Layton et al. (2023)) and a dummy representing the "post-move" time periods.