Using Machine Learning Methods to Predict Physician-Hospital Integration

Stuart Craig, Matthew Grennan, Joseph Martinez, and Ashley Swanson

August 2020

1 Introduction

In this document, we propose a new method that combines high-dimensional data with machine learning methods to predict physician-hospital integration. We compare the performance of this method with alternative approaches used in the healthcare economics literature for a large validated sample, finding that it outperforms previous methods by a substantial margin. We also compare the static predictions of this model to a large validated sample recently made available by the Agency for Healthcare Research and Quality (AHRQ) and again document a high degree of accuracy. Finally, we briefly summarize the implications of our method for the growth in physician-hospital integration over the years 2008-2016.

2 Input Data and Cleaning

The goal of this exercise is to take each physician in the Centers for Medicare and Medicaid Services' (CMS) Medicare Data on Provider Practice and Specialty (MD-PPAS) data, determine whether or not she is integrated with a hospital or hospital system, and assign her to the correct hospital/system as appropriate.

We rely on the following data sources, each of which is readily available for purchase by researchers:

- MD-PPAS data (2008-2016): contains each physician's national provider identifier (NPI) and specialty, up to two taxpayer identification numbers (TINs) under which the physician bills Medicare, total allowed amounts billed under each TIN, and CBSA code
- American Hospital Association (AHA) annual survey data (2008-2016): survey data on hospitals and hospital systems, cleaned up as a complete panel for years in which facilities are open as in Cooper et al. (2018)
- SK&A data (2008-2016): SK&A polls U.S. office-based providers, including physicians and non-physician medical professionals; the data contain NPI, practice address, specialty, and self-reported system and hospital ownership (assigned a unique 3-digit code and name by SK&A)[1]
- CMS Physician Compare data (2014-2016): physician NPI, practice address, specialty, and self-reported hospital affiliation for all clinicians enrolled in Medicare
- CMS National Plan and Provider Enumeration System (NPPES) data (2008,2010,2018): universe of physician NPIs, practice addresses, and specialties; infrequently updated after 2008 (when signing up for an NPI was mandatory)

With these raw data in hand, we performed several additional processing steps:

- We assigned geographic coordinates (latitude and longitude) based on physician practice addresses from SK&A, physician compare, and NPPES. When practice addresses were unavailable, we use the coordinates of the CBSA centroid.

[1] Only 65.5% of physician-years in MD-PPAS are present in SK&A data. However, 94.6% of TIN-years contain at least one physician that is present is SK&A in the same year. These TIN-years account for 98.2% of all physician-years in MD-PPAS.

- For each NPI-year, we keep only the TIN with the highest allowed amount. 82% of NPI-years have a single TIN. After removing secondary TINs, 96% of total allowed amounts remain in the sample.
- We matched each combination of self-reported SK&A hospital (system) ownership code, year, and hospital referral region (HRR) to a hospital (system) observation in the AHA survey in the same year.[2] We used the following match procedure:
  - We manually matched each SK&A system ownership code to its appropriate AHA system.
  - For each SK&A hospital ownership code, we matched to AHA using geography and string matching (Jaro-Winkler) on names. Candidate matches were generated by iteratively matching and relaxing string and geographic distances. Unique SK&A-AHA matches generated by this procedure were accepted without further validation if (a) the median physician-hospital distance (across physicians in the TIN) is less than 50 miles; or b) the hospital is geographically closest and has string distance less than 0.15. All other matches were validated manually.

For each combination of NPI, TIN, year, and hospital/system from the AHA, we construct the following variables:

- TIN legal name and AHA name similarity metrics (Jaro-Winkler and restricted edit distances)
- Geographic distance between physician and hospital (or nearest hospital in system). Minimum and median distances between physicians in same TIN and hospital (or nearest hospital in system)
- TIN-level SK&A reported ownership measures: number and share of NPIs in TIN that report being owned by hospital/system, plus one-year lags and leads of these variables

Once restricted to physicians, the MD-PPAS data contains 791,649 NPIs and 260,609 TINs over 2008-2016, for a total data set consisting of 5,393,622 physician-years. For the remainder of our analyses, we make the following, additional sample restrictions:

- We exclude TIN-years for which all physicians have missing geographic coordinates (no CBSA code or practice address from SK&A, Physician Compare, or NPPES). This removes 923 physician-years from the sample.
- We exclude TIN-years for which all physicians are not in any HRR. This removes 46,320 physician-years (from Puerto Rico and US territories) from the sample.

After these restrictions, we are left with 785,414 NPIs, 253,318 TINs, and 5,346,252 physician-years.

3 Training Data and Performance of Algorithms

The composition of the training sample we use is shown in Table 1. The training data was built from a random sample of 572 TINs, a random sample of 214 TINs associated with mergers between hospitals and physician practices identified by Irving Levin Associates (a market intelligence firm that tracks the health care sector), and a random sample of 131 TINs with at least one cardiologist in the Philadelphia and Miami HRRs.[3] Integration status for these TINs was verified using a wide variety of publicly-available sources (including IRS 990 filings, SEC 10-K filings, physician practice websites, government financial reports, disclosures for tax exempt bonds, and trade websites like Becker's Hospital Review). Wherever possible, we relied on IRS and SEC filings of acquiring hospitals.[4]

The first triplet of columns describe our sample TINs; the second triplet of columns describe our sample NPIs. Over 2008-2016, we observe 916 unique TINs and 124,725 unique NPIs in our training sample. The last triplet of columns describes a subset of 5,152 unique NPIs who were involved in group practice

[2] A single SK&A code might refer to multiple similarly-named entities hundreds of miles apart.

[3] We originally pursued cardiologists as a particularly interesting specialty, and the Philadelphia and Miami markets as ones with which we had previous familiarity. This last sample is therefore a sample of convenience.

[4] Of 6,498 total TIN-years (543,183 unique NPI-years), 5,006 TIN-years (408,257 NPI-years) were validated in IRS/SEC filings. Verification in IRS/SEC filings was less often possible in small TINs.

mergers—these NPIs were present in a group practice TIN before and after it changed integration status. Our validation is performed at the TIN-year level, and we follow all sample TINs for the full time horizon 2008-2016. The small variation in the number of TINs across years is driven by TIN entry/exit from the MD-PPAS data. The dramatic growth we observe in the number of NPIs in our sample over time is driven by additional physician practices joining our sample TINs over time. When this occurred, we did not track down the integration status of those NPIs' previous TINs and did not include any of their previous TIN-years in our training sample.

Table 1: Random TIN, Levin, & Cardiology Sample Size by Year

	TINs in Sample			All NPIs in Sample			NPIs Involved in Group Practice Mergers		
Year	Integrated	Non-Integrated	Total	Integrated	Non-Integrated	Total	Integrated	Non-Integrated	Total
2008	289	398	687	31,391	13,004	44,395	11	4,095	4,106
2009	317	405	722	34,903	13,304	48,207	322	3,859	4,181
2010	334	405	739	39,041	12,696	51,737	1,368	2,745	4,113
2011	363	373	736	43,798	12,448	56,246	2,146	1,816	3,962
2012	378	337	715	48,297	12,062	60,359	2,903	793	3,696
2013	386	339	725	51,566	12,008	63,574	3,216	230	3,446
2014	394	335	729	55,970	12,709	68,679	3,031	180	3,211
2015	386	340	726	59,711	12,991	72,702	2,859	115	2,974
2016	384	335	719	63,585	13,699	77,284	2,765	1	2,766

3.1 Comparison Algorithms

For one comparison algorithm, we follow Neprash et al. (2015) and use MD-PPAS to compute the percent of outpatient billing provided at a hospital outpatient department. We then flag doctors as integrated if the percent of hospital outpatient billing is greater than 25%. In Table 2, we display the extensive margin error rate (misclassification of integrated vs. non-integrated status) for this approach.

The "Overall" error columns display the percent of physicians misclassified by this algorithm for all physician-years. The "Integrated" columns display the percent of integrated physicians that are erroneously predicted to be non-integrated (the "false negative" rate for our context). Similarly, the "Non-Integrated" columns display the percent of non-integrated physicians that are erroneously predicted to be integrated (the "false positive" rate). For context, note that a simple coin flip would (in expectation) result in a 50% error rate in each column.

Across all training sample physicians, 45% of physician-years were misclassified using this approach, and misclassification rates were similar in the full sample (left panel) and in the subsample of NPIs involved in group practice mergers (right panel). Misclassification rates were higher for integrated physician-years, indicating that the algorithm misses many physicians who are truly integrated with hospitals, but do not do much outpatient billing. As in Neprash et al. (2015), we also tried using cutoffs at 50, 75, 95, and 99.9% outpatient billing; results were worse (higher overall error rate) with these stricter thresholds.

For another comparison algorithm, we follow Baker et al. (2016) and measure vertical integration by relying directly on self-reporting of hospital/system ownership in SK&A survey data. We assign each physician to her reported system owner (or, as appropriate, to the system of her reported hospital owner) in SK&A. If a doctor reports ownership in more than one system, we use Jaro-Winkler distance to assign her to the system whose name is most similar to the physician's TIN legal name. For hospitals that are not a part of a system, we compare hospital and TIN legal names. Any ties are broken at random.

Table 3 below is structured similarly to Table 2 above, but with a focus on intensive margin error as well as extensive margin error. That is, the hospital outpatient billing algorithm only classifies physicians as integrated or non-integrated, while the SK&A algorithm also assigns integrated physicians to a specific hospital/system. Even though it is therefore held to a stricter standard, SK&A self-report performs better than hospital outpatient billing, with an error rate of 41% in the full sample and 23% in the subsample of physicians involved in group practice mergers. A significant driver of the still-high error rate is survey nonresponse, which disproportionately affects integrated physicians, 50% of whom are erroneously classified as non-integrated.

Table 2: Individual Physician Hospital Outpatient Department Billing: Prediction Error (%)

	All NPIs			NPIs Involved in Group Practice Mergers		
Year	Overall	Integrated	Non-Integrated	Overall	Integrated	Non-Integrated
2008	45.89	53.19	28.28	24.74	63.64	24.64
2009	44.59	50.34	29.49	29.8	77.02	25.86
2010	44.68	49.69	29.26	38.9	66.67	25.06
2011	45.92	50.5	29.8	50.61	70.5	27.09
2012	45.33	49.37	29.17	51.73	61.76	15.01
2013	43.84	47.31	28.91	57.23	60.73	8.26
2014	44.24	47.46	30.07	54.69	57.77	2.78
2015	44.02	47.19	29.48	55.28	57.36	3.48
2016	44.25	47.29	30.14	58.53	58.55	0
All Years	44.67	48.76	29.41	45.51	61.41	24.1

Table 3: SK&A Ownership Self-Report: Prediction Error (%)

	All NPIs			NPIs Involved in Group Practice Mergers		
Year	Overall	Integrated	Non-Integrated	Overall	Integrated	Non-Integrated
2008	54.75	76.54	2.14	3.04	72.73	2.86
2009	48.69	65.81	3.77	12.87	75.47	7.64
2010	45.99	59.98	2.98	26.04	66.52	5.87
2011	43.12	53.36	7.12	38.62	55.59	18.56
2012	37.9	45.99	5.52	26.52	31.76	7.31
2013	37.1	43.81	8.31	25.25	26.71	4.78
2014	36.83	43.7	6.57	24.63	25.8	5
2015	36.4	42.15	10.01	24.34	24.48	20.87
2016	36.89	43.07	8.23	25.05	25.06	0
All Years	40.98	50.34	6.07	22.56	33.89	7.32

An alternative way of using the SK&A data is to pool survey responses within each TIN-year. Table 4 below shows error rates from a procedure that pools responses by fitting a decision tree to predict integration status, where the only input is the share of physicians in each TIN-year reporting ownership by a given hospital/system. Cutoffs used for share of physicians are determined by a decision tree with one split; the split in the classification tree is determined by the largest decrease in Gini impurity. Results are qualitatively similar (but a bit worse) when we instead use a logistic regression with the same single regressor. Since the training sample is used to train the model and to assess performance, error rates below are out-of-sample error rates from repeated (three repeats) five-fold cross-validation. Pooling data across NPIs within a TIN improves performance substantially: full sample error rates drop from 41% to 6% and the "Group Practice Merger" subsample error rates drop from 23% to 14%.

Table 4: SK&A Reported Ownership Grouped by TIN: Prediction Error (%)

	All NPIs			NPIs Involved in Group Practice Mergers		
Year	Overall	Integrated	Non-Integrated	Overall	Integrated	Non-Integrated
2008	13.55	13.54	13.58	18.45	9.09	18.48
2009	7.92	8.3	6.93	18.96	37.89	17.38
2010	6.79	6.22	8.54	23.02	24.63	22.21
2011	9.37	7.4	16.31	27.18	26.19	28.36
2012	2.83	2.42	4.44	4.11	2.62	9.58
2013	5.02	3.99	9.44	2.79	2.83	2.17
2014	5.34	5.46	4.81	5.4	5.59	2.22
2015	5.52	4.52	10.11	10.17	9.82	18.84
2016	4.42	3.77	7.41	11.18	11.19	0
All Years	6.37	5.65	9.06	14.2	10.46	19.22

Note: The out-of-sample error in this table is determined by repeated 5-fold cross validation.

3.2 Random Forest Model

In our preferred approach to assigning integration, we use all the variables defined in Section 2 above to fit a random forest model predicting each NPI's integration status with each candidate hospital/system. We tune the random forest parameters (number of trees and number of variables tried at each node of each tree) to reduce overall misclassification, in repeated five-fold cross-validation. Ties are broken by the vote share (over trees) in the random forest; remaining ties are broken at random using a fixed seed.

Out-of-sample performance of the model is shown in Table 5 below. We fit the model separately in each year; tuning parameters employed for each year are reported in the first two columns. The remaining columns are as in the above, reflecting prediction error at the NPI-year level. This algorithm improves significantly upon the SK&A-by-TIN model above, reducing the full sample error rate to 3% and the "Group Practice Merger" subsample error rate to 10%. Performance is significantly better in 2010-2016 than in 2008-2009 because of improvements in SK&A reporting in later years.

Several features of our analysis bear noting. First, we use a number of both geographic distance and string distance variables to fit the model. The latter is more difficult to interpret than the former, so we also fit models using no string distance variables. This causes out-of-sample error rates to increase from 2.7% to 3.4% for the full sample. Second, we use our full training sample to fit the model, which involves combining a pure random sample of TINs with the Levin sample, which explicitly upsamples TINs changing integration status, and the convenience cardiology sample, which explicitly overweights two particular MSAs. However, we have also trained and fit each model (outpatient billing, SK&A individual, SK&A-by-TIN, and random forest) on the random sample of TINs only and found the performance of each algorithm to be nearly identical. Thus, in the remainder of this document, we use the full training sample to generate all metrics.

Table 5: Random Forest Model Full Specification: Prediction Error (%)

			All NPIs			NPIs with an Ownership Transition		
Year	γ	λ	Overall	Integrated	Non-Integrated	Overall	Integrated	Non-Integrated
2008	1	750	7.696	9.977	2.187	4.738	9.091	4.726
2009	1	750	5.143	5.726	3.612	15.809	61.491	11.993
2010	10	500	3.267	3.451	2.702	18.626	32.919	11.492
2011	5	500	2.446	2.477	2.336	15.282	15.750	14.727
2012	3	750	1.581	1.868	0.431	3.952	5.029	0
2013	5	1,500	2.188	2.561	0.586	6.310	6.758	0
2014	20	500	1.473	1.499	1.356	4.412	4.674	0
2015	15	1,000	1.582	1.497	1.973	7.039	7.322	0
2016	5	750	1.766	1.726	1.949	9.038	9.042	0
All Years			2.731	2.946	1.929	9.833	10.481	8.959

Note: The out-of-sample error in this table is determined by repeated 5-fold cross validation. γ is the number of variables considered at each node within a decision tree. λ is the number of trees.

3.2.1 Characterizing Prediction Error in Random Forest Models

As noted above, the random forest model's prediction error is very low overall (2.7%), but higher for a subsample of NPIs involved in TIN transitions (9.8%). In the analysis below, we provide a taxonomy of errors resulting from the random forest model, with the following definitions:

- For TIN j with billing activity in T_j periods, we correctly predict some periods $t_{cj} \subseteq T_j$. Define $\underline{t}_{cj} = \min t_{cj}$, that is the first year that we predict correctly. Define $\bar{t}_{cj} = \max t_{cj}$, the last year we predict correctly.

- For each physician i in pre-merger TIN j, define $t_{j(i)m}$ as the year each physician becomes integrated as a part of (possibly multiple) post-merger TINs j'. Note that j may equal j' when the health system keeps all of its acquired physicians in the original pre-merger TIN.

- Incorrect Prediction in First Years of TIN Existence: this captures physicians in TINs with incorrect predictions in periods $t_j < \underline{t}_{cj}$. Since the underlying quality of SK&A generally improves from 2008-2016, a large share of overall error typically falls in this category. This means that most of the error here involves a lag in identifying integration.

- Incorrect Prediction Between First and Last Correct Year (a.k.a. "flip-floppers"): this includes physicians in TINs with incorrect predictions in t_j where $t_j \in [\underline{t}_{cj}, \bar{t}_{cj}]$.

- Incorrect Prediction in Last Years of TIN Existence: these are physicians in TINs with incorrect predictions in periods $t_j > \bar{t}_{cj}$

- No Correctly Predicted Year: these are physicians in TINs with incorrect predictions for all t_j such that $t_{cj} = \emptyset$.

- System ID Change/Hospital Merger: occasionally the hospital with which a TIN is integrated will switch systems and physicians will remain integrated with this hospital. This can cause error when lags or leads of ownership overlap at the system level used in prediction do not reflect this system change.

- Early Prediction of TIN Acquisition: this category includes physicians who are eventually acquired by hospitals, but we predict the acquisition too early. That is, non-integrated physicians in pre-merger TIN j who are predicted to be integrated prior to the merger year, $t_{j(i)m}$.

- Late Prediction of TIN Acquisition: this category captures physicians who are acquired by hospitals, but for whom we do not predict the any change in ownership until after the true date of acquisition.

These are integrated physicians that 1) are in pre-merger TIN j prior to $t_{j(i)m}$, 2) are in post-merger TINs j', and 3) that we predict are non-integrated at $t \geq t_{j(i)m}$

The breakdown of errors for the whole sample of physician-years ("Total") and within the subsample of incorrect physician-years ("Incorrect") is shown in Table 6 below. For the overall sample, 42% of errors result from incorrect assignment in the first year of TIN existence. For the subsample of NPIs with ownership changes, where the error rate was higher at 9.8% across all NPI-years, 94% of errors are driven by the algorithm making a mistake in determining the timing of TIN mergers. Specifically, 56% of errors are driven by late assignment of integration status to truly integrated NPI-years, and 39% of errors are driven by early assignment of integration status to truly non-integrated NPI-years.

Table 6: Taxonomy & Frequency of Random Forest Prediction Errors: Full Specification

	Integrated Physicians		Non-Integrated Physicians	
	% of Total	% of Incorrect	% of Total	% of Incorrect
Overall Sample				
Incorrect Prediction in First Years of TIN Existence	1.138	41.667	0.057	2.086
Incorrect Prediction Between First and Last Correct Year	0.230	8.424	0.208	7.602
Incorrect Prediction in Last Years of TIN Existence	0.357	13.067	0.143	5.254
No Correctly Predicted Year	0.460	16.828	0	0
System ID Change/Hospital Merger	0.138	5.071	0	0
NPIs with Ownership Changes				
TIN Hospital Merger-Early Prediction	0	0	3.812	38.791
TIN Hospital Merger-Late Prediction	5.458	55.511	0	0
System ID Change/Hospital Merger	0.559	5.656	0	0

Note: NPI-Years are averaged across repeated cross-validation folds.

3.2.2 Comparison across Algorithms

Table 7 below displays the extensive margin concordance between different algorithms we have implemented, for our full sample of NPI-years. Each column compares an alternative approach to the baseline, full sample random forest model. For example, the value of 0.980 under "No String" indicates that 98% of NPI-years had the same integration vs. non-integration status predicted in our full sample random forest model, whether or not string distance variables were employed in the model. Similarly, the value of 0.986 under "Random Only" indicates that over 98% of NPI-years had the same integration status predicted in the random forest model, whether or not we restrict the training sample to the random sample of TINs. As expected, concordance is high for all random forest models, and lower for SK&A and hospital outpatient ("HOPD") models. There is no clear pattern in which specific specialties have higher concordance rates for most comparisons, with the exception of "HOPD Billing." In the last column, we observe that concordance rates are lowest for hospital-based specialties like anesthesiology and radiology, where high outpatient billing rates would lead the HOPD algorithm to erroneously flag integration where none exists.

3.2.3 Comparison with AHRQ 2016 Data

In 2015, AHRQ created an initiative to study health systems, in collaboration with researchers at Dartmouth College, the National Bureau of Economic Research, the RAND Corporation, and Mathematica Policy Research. AHRQ has since published lists of U.S. health systems for 2016 and 2018, including indicators for system ownership and provider affiliations with systems.[5] By AHRQ's definition, "a health system includes at least one hospital and at least one group of physicians that provides comprehensive care (including primary and specialty care) who are connected with each other and with the hospital through common ownership or joint management." AHRQ's list explicitly excludes candidate systems without at least one general acute care

[5] See https://www.ahrq.gov/sites/default/files/wysiwyg/chsp/compendium/techdocrpt_0.pdf.

Table 7: Comparing Integration Algorithms: Share of Physician-Years with Matching Integration Status

	No String	Random Only	Grouped SK&A	Individual SK&A	HOPD Billing
Overall	0.980	0.986	0.908	0.759	0.665
Anesthesia	0.976	0.987	0.912	0.787	0.394
Cardiology	0.982	0.985	0.903	0.759	0.695
Nephrology	0.986	0.990	0.905	0.819	0.805
Neurosurgery	0.981	0.986	0.896	0.736	0.686
Ophthalmology	0.992	0.994	0.927	0.877	0.884
Orthopedics	0.983	0.988	0.903	0.826	0.776
Otolaryngology	0.985	0.989	0.910	0.817	0.800
Plastic Surgery	0.989	0.993	0.916	0.841	0.620
Primary Care	0.984	0.988	0.925	0.749	0.670
Radiology	0.960	0.975	0.880	0.721	0.478
Urology	0.985	0.991	0.901	0.815	0.795

Notes: Each column compares a possible approach to determining integration status with the baseline random forest specification. Specifically, each column displays the share of physicians whose integration status matches that in the baseline random forest. "Cardiology" includes cardiologists and cardiac surgeons. "Anesthesia" includes anesthesiologists and pain management physicians.

hospital, 50 total physicians, or 10 primary care physicians. These and other exclusions imply that AHRQ's list is not fully comparable with our database, which is meant to detect all instances of hospital/health system ownership of physician practices.

For the 96.6% of NPIs in TINs that we are able to match uniquely to AHRQ in our 2016 training data, we can evaluate how AHRQ's list deviates from manually validated integration status. For these NPIs, the error rate in extensive margin concordance was 6.4% overall, 6.6% among integrated TINs, and 5.8% among non-integrated TINs. Further investigation of these TINs revealed two key patterns of disagreement. First, we found a number of instances where we were able to confirm ownership using SEC/IRS data, but no "health system" was flagged by AHRQ. This may be driven by those instances not meeting AHRQ's standard of health system due to type of hospital or number of physicians. Second, we observe that AHRQ occasionally flags integration for certain large TINs (e.g., hospital staffing companies) with tight affiliations to hospitals, but with physician or private equity ownership, rather than hospital/health system ownership.

4 National Integration Trends

Focusing on physicians in the continental United States, Figure 1 shows the national trends in integration predicted by each of the above algorithms, with the AHRQ prediction indicated for 2016 only.[6] Each algorithm predicts steady growth in integration between 2008 and 2016. Growth is flattest for the outpatient billing algorithm. By 2016, predicted integration is similar for the "Random Forest" and "Individual SK&A Survey Respondents" algorithms, a bit higher for "Grouped SK&A," and a bit lower for "AHRQ Compendium."

Given the superior performance of the random forest algorithm in fitting vertical integration in our training sample, we focus on that algorithm in the remainder of this document. Figure 2 shows national trends in integration, overall and by specialty. While vertical integration has trended upward in each of the top specialties shown, growth has been most dramatic for Cardiology and Neurosurgery, each of which was more than 50% integrated by 2016.

Lastly, Figure 3 shows a bubble plot of integration in each of the approximately 300 U.S. hospital referral regions (HRRs), across all specialties, comparing 2016 on the vertical axis to 2008 on the horizontal axis. Bubble sizes indicate the number of NPIs in each HRR. Essentially all of the mass is northeast of the 45-degree line, indicating growth in vertical integration. However, there is a vertically dispersed pattern of bubbles at each value of "Share Integrated in 2008"; some healthcare markets experienced much more growth in integration than others.

[6] The Figure has two separate lines for "Individual SK&A Survey Respondents" (self-reported ownership among NPIs responding to the SK&A survey) and "Individual SK&A" (defaults to non-integration for all non-respondents). Unsurprisingly, the former is shifted upward relative to the latter.

Figure 1: Comparing National Physician Integration Trends Across Methods

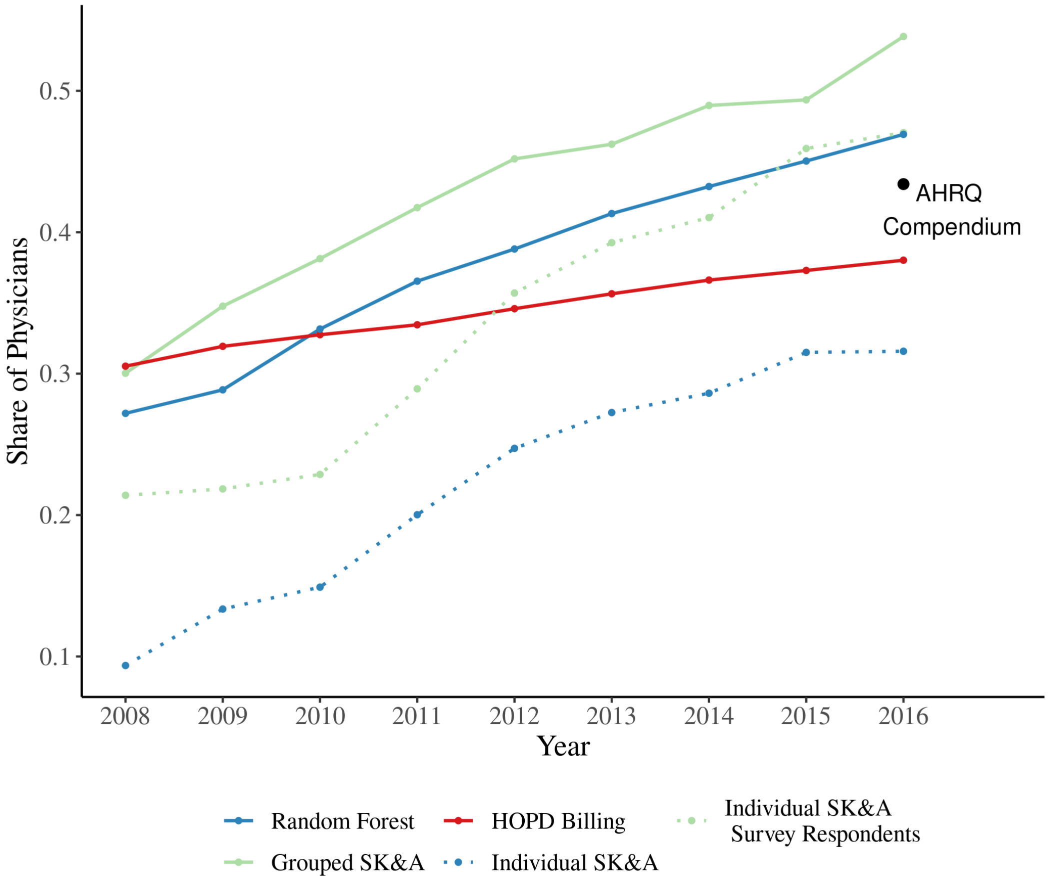

Figure 2: National Trends in Physician Integration (2008-2016)–Full Specification

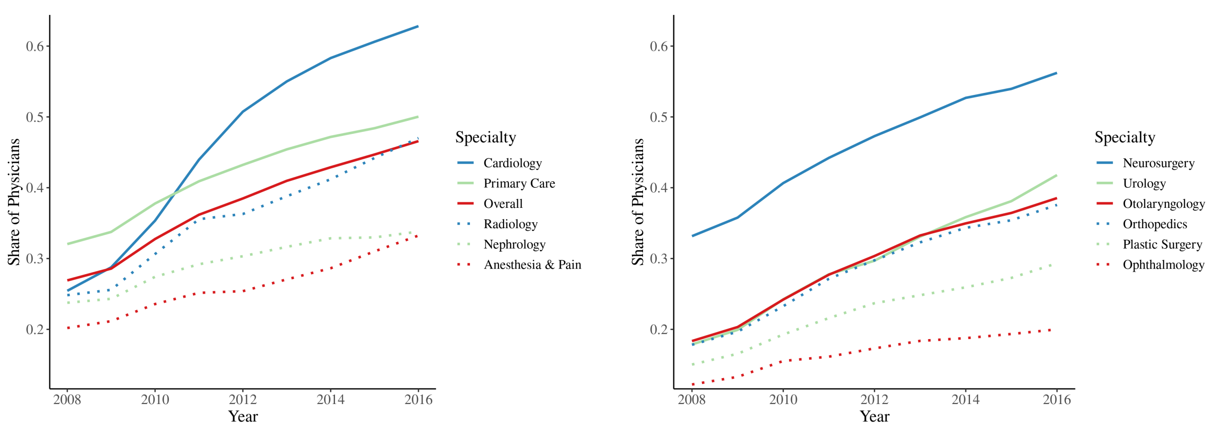

Figure 3: Growth in Physician Integration by HRR

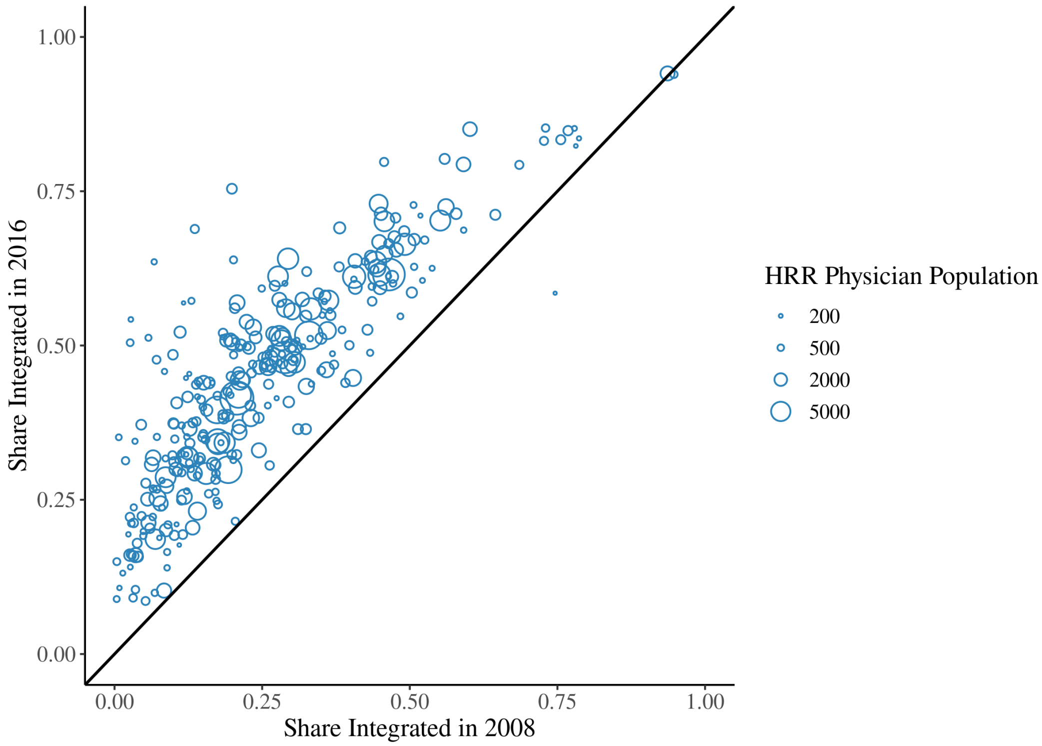

References

Baker, L. C., Bundorf, M. K., and Kessler, D. P. (2016). The effect of hospital/physician integration on hospital choice. *Journal of Health Economics*, 50:1–8.

Cooper, Z., Craig, S. V., Gaynor, M., and Van Reenen, J. (2018). The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured*. *The Quarterly Journal of Economics*, 134(1):51–107.

Neprash, H. T., Chernew, M. E., Hicks, A. L., Gibson, T., and McWilliams, J. M. (2015). Association of Financial Integration Between Physicians and Hospitals With Commercial Health Care Prices. *JAMA Internal Medicine*, 175(12):1932–1939.