Sabety

By 2025 the United States is projected to have a shortage of 90,400 physicians. This is because demand for physicians will increase, due to a rapidly aging population, while the supply of physicians will simultaneously shrink, as providers themselves retire.[i] More than 20 percent of primary care physicians (PCPs) are over 60 years of age, compared to about 18 percent in the total workforce.[ii] As a result, understanding the impact of provider departures and, in particular, retirement will be particularly important over the next decade.

In order to better understand how worker turnover affects firm productivity, I first ask how patients with a departing provider are directly affected. For instance, if older PCPs are less well versed in the latest pharmacological agents, replacing a retiring PCP with a younger PCP may improve patient outcomes. I then focus on indirect effects, quantifying disruption costs for patients who did not lose their provider, but seek care at a physician group that experienced a departure. For example, incumbent providers may take on more work when a provider departs, affecting patients who were not directly affected by the transition. In order to shed light on these effects, I use Medicare data to match patients in physician groups with a departure to patients in groups without a departure. I then quantify patient outcomes in an event study framework.

1. Direct effects: how do transition costs vary by provider age, medical training, and tenure as well as by the length of the patient-provider relationship? I ask how departure costs and benefits differ based on the characteristics of the leaving versus replacement provider. This allows me to separately estimate the returns to experience, training, tenure, and longstanding patient-provider relationships.

I first ask if retiring providers provide better or worse care than younger replacement providers along margins such as missed conditions and their patients' rate of adherence to medications.[1] Figure 1 suggests that these kinds of replacements may be common with replacement providers being on average younger than their departing peers. Age is an important determinant of general human capital and, thus, potentially important for patient outcomes.[iii] On one hand, patients' health could improve if the replacement provider is better versed in new drugs and/or medical technologies. On the other hand, health could be negatively affected if there are large returns to experience.

I examine how provider labor force patterns depend on age and then quantify the resultant effect on patient outcomes. For instance, Figure 2 plots the average number of patients seen per month and the average patient risk score by PCP age. The graph shows that the average PCP's case load peaks at age 60 and then slowly declines. Comparatively, the average patient risk score reaches its minimum around 60, increasing thereafter. Although more effective PCPs may deliver better care and, thus, have lower risk scores, these descriptive patterns may suggest that when PCPs age they decrease their patient load, letting go of their healthier patients first. If patient outcomes are positively related to experience throughout the age distribution, older PCPs could be acting optimally by keeping their sickest patients. However, if this relationship does not hold throughout the age distribution, this behavior could be detrimental to high-risk patients' outcomes.

Returns to training may also be important, especially for more complex patients. For example, if the departing provider is a PCP and the replacement provider is a physician assistant (PA), the training gap between providers may affect patient outcomes. This is a potentially important channel because preliminary work shows that PAs and nurse practitioners (NPs) replace a departing PCP half of the time.[2]

To quantify how provider tenure affects patient outcomes, I estimate how outcomes are affected when a leaving provider's patients are absorbed by an incumbent worker opposed to a

[1] I categorize retirements as departures where a provider of retirement age disappears from the data

[2] PAs practice nearly independently in most settings, with the overseeing physician sometimes even located in another state, NPs have independent authority in many states, and both are allowed to independently prescribe in almost all states.

new hire. For example, a new provider may have to get up to speed on the practice's electronic health records (EHR) system and may not know how to recognize flags relevant to patient care in the chart. This, in turn, could adversely affect patient outcomes.

I also quantify how the loss of a longstanding patient-provider relationship, something I term "patient-specific" knowledge, affects outcomes. For instance, a provider who has known the patient for a long period of time may be better able to pick up that a patient is deteriorating, even if the patient is unable to verbalize it. I capture this by controlling for the differences between departing and replacement providers, ascribing the residual effect of turnover on outcomes to the discontinuity of care.

2. **Indirect effects: how do physician groups compensate for the loss of the provider?** In order to estimate how long it takes for the firm to replace the worker externally through hiring, I directly observe when the replacement hire is made. I then focus on how a worker's exit affects the firm's demand for its remaining workers. If firms cannot replace skilled workers from the external market in the short run, remaining providers may absorb the departing provider's patients until a replacement can be hired and trained.

I also assess the effect of turnover on patients at the same practice who do not directly interact with the departing provider. This is a potentially important channel because patient care throughout the practice could be influenced by the loss of a provider. For instance, if staying providers temporarily take on some of the departing provider's patients, it could leave less time for their own patient population. I estimate these effects by comparing patients in practices that did not experience a disruption to patients in practices that experienced a disruption, but did not directly lose a provider.

Data: As my main data set, I use a 20 percent sample of Medicare patients from 2008 to 2014. Table 1 summarizes the working sample, which contains over 5 million patients and 200,000 providers per year. I find that about 14 percent of providers depart per year.[3] My sample also contains about 100,000 TIN-ZIPs per year. The TIN-ZIP is the finest practice unit available in the Medicare data, where the TIN is the tax identifier and the ZIP is the zip code. As a result, the TIN-ZIP counts practices owned by the same entity in the same zip code once, underestimating the number of practices. In addition to this weakness, although the Medicare data contains a unique provider identifier, it does not record information on provider characteristics.

To overcome these weaknesses, I merge three different data sets into the Medicare data (this merged-in data set underlies Figures 1 and 2). My first data set, SK&A, is proprietary data that obtains a census of physician groups for marketing. SK&A allows me to assign providers to a specific practice unit, improving on the lack of precision in the Medicare TIN-ZIP metric. Second, I use the National Plan and Provider Enumeration System (NPPES), a registry of national provider identifiers (NPI) to identify PCPs, PAs, and NPs. It also contains information on when a NPI is deactivated, allowing me to identify provider deaths. Third, I use MD-PPAS, created by Medicare, to identify provider birth date.

Because Medicare is a 20 percent sample of patients, not physician groups, it may be underpowered for physician group outcomes. As a result, I supplement the Medicare data with proprietary athenahealth data, a national provider of EHR, medical billing, and practice management services. This gives me a more comprehensive picture of physician group level outcomes because it contains all patients and providers in the physician group. The data covers 44,038 unique primary care providers and about 24 million unique patients from 2010 to 2014 (Table 1). I use the athenahealth data as my secondary data set because it is a short panel, it only captures the physician setting, and providers are non-identifiable.

[3] I define a departure as moving practices and not seeing a patient for the subsequent 12 months or more. I also restrict departures to departing three months before the clinic departs from the data set to avoid capturing mergers and/or acquisitions.

Empirical Strategy: I intend to causally identify the effect of turnover on patient outcomes using an event study framework, matching patients in treatment physician groups to observably similar patients in control physician groups that did not lose a provider.[4] The identification strategy assumes that patients in treatment and control practices would have parallel trends in the absence of a transition. In order to partially test this assumption, I compare pre-trends by quantifying patient outcomes in treatment and control practices before the disruption occurs.

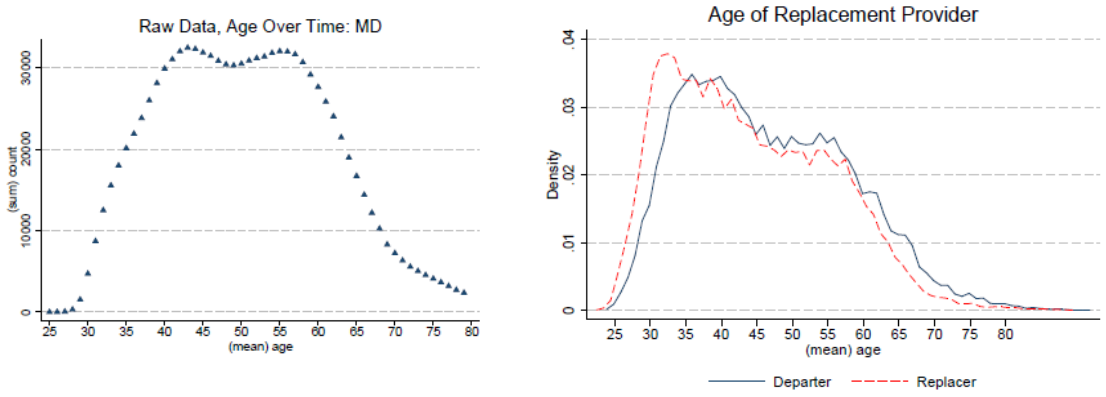
To give a general sense of my regression framework, I model the direct effect of a departure in an event study framework as follows:

$$Y_{ik} = \rho_i + \sum_k \beta_k \mathbf{1}_{Tx\ practice,ik} * \boldsymbol{D}_{post,ik} + \varepsilon_{ik}$$

I create a variable *k* to reflect the relative *k* months from the transition (at *k*=0). I define the regression such that *i* represents a patient, ρ_j are patient match fixed effects, $\mathbf{1}_{Tx\ practice,ik}$ is a dummy variable indicating that the patient seeks care in a treatment practice, and $\boldsymbol{D}_{post,ik}$ is a dummy variable indicating that the provider left *k* periods ago. The direct effect of losing a provider, or the effect of interest, is captured by β_k. I assign patients to providers if the patient sees a provider for over half of their pre-period visits. Outcomes (Y_{jt}), outlined in Table 2, are measured at the patient level and were chosen to capture the effectiveness of care. Regressions are clustered at the patient-pair level.

Conclusion: I quantify the substitutability of workers through a healthcare lens. This not only allows me to quantify productivity in terms of patient outcomes, but also enables me to think about tasks and, thus, task-shifting in measureable patient units. As a result, when faced with a rapidly aging population, lessons from this work will help policymakers think creatively about how to lessen transition costs in sectors with highly skilled workers who may have a large amount of general, firm, and patient specific human capital that is being lost in the transition.

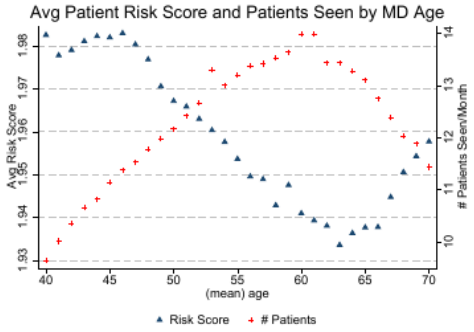
Figure 1:



Notes: The first plot shows the count of each PCP by age. The Age of Replacement graph overlays two histograms: one for departing providers (blue) and one for replacement providers (red).

[4] I use the term "treatment practice" to describe practices with a transition of care and the term "control practice" to describe practices where no transitions occurred.

Figure 2:



Notes: The plot contains the average PCPs' age on the x-axis. The blue triangles plot the average patient risk score by PCP age and the red crosses plot the average number of patients seen per month seen by PCP age.

Table 1:

Medicare

Year	TIN-ZIP	APs	Patients	Dep. Rate
2008	109,509	207,773	5.7m	13%
2009	99,397	213,610	5.6m	15%
2010	98,603	219,355	5.3m	15%
2011	94,947	225,950	5.3m	16%
2012	91,861	235,183	5.4m	16%
2013	91,137	285,264	5.5m	17%
2014	90,519	259,620	5.5m	9%

athenahealth

Year	Clinics	APs-Clinc	Patients	Dep. Rate
2010	4,457	27,705	10.4m	16%
2011	6,151	36,112	13.0m	16%
2012	7,598	50,398	16.8m	16%
2013	9,899	69,085	22.0m	17%
2014	12,932	84,841	29.0m	15%

Notes: The TIN-ZIP is the finest practice unit in the Medicare data, where the TIN is the tax identifier and the ZIP is the zip code. APs represent the number of unique advanced providers (PCPs, PAs, and NPs). Dep. rate = the # of departures over total # of providers by year. I restrict departures to departing three months before the clinic departs from the data set to avoid capturing mergers and/or acquisitions.

Table 2:

Outcomes (Y_{it})	**Specific Examples Pre/Post**
Care coordination	• Rate of duplicate testing and referrals • How quickly patients are able to make appointments with providers
Patient satisfaction	• Rate of patients leaving with departing providers (proxy v. precisely observe)
Adherence to medications	• Rate of refills after a provider departs
Rate of ACSCs	• Rate of ambulatory care sensitive conditions presenting in the physician setting
ED/IP utilization	• Re-admissions • Rate of emergency department (ED) / inpatient (IP) visits
Rate of missed conditions	• Rate of patients presenting with heart attacks in the ED/IP several days after presenting in the doctors' office reporting risk factors • Rate of patients presenting with pulmonary embolisms in the ED/IP several days after presenting in the doctors' office with 7 risk factors

Notes: Blue outcomes are specific to the athenahealth data and green outcomes are specific to Medicare. **Black** outcomes are quantified in both data sets.

[i] "The Complexities of Physician Supply and Demand: Projections from 2013 to 2025". AAMC. March, 2015
[ii] Authors tabulations of Medicare data and general workforce data from the Administration on Aging, 2014
[iii] Mincer, J. A. (1974). "The Human Capital Earnings Function" in *Schooling, Experience, and Earnings*. Cambridge, MA: National Bureau of Economic Research, pp. 83-96.; Becker, G.S. (1975). "Human Capital and the Personal Distribution of Income: an Analytical Approach" in *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, 2nd edition. New York: National Bureau of Economic Research.

Adrienne H. Sabety

Harvard University
Ph.D. Program in Health Policy
14 Story St.
Cambridge, MA 02138

Phone: (805) 441-6311
E-mail: asabety@g.harvard.edu
Website: http://scholar.harvard.edu/asabety/

Graduate Studies

Fields: Health Economics, Labor Economics, Industrial Organization

Committee: David Cutler (chair), Claudia Goldin, Jonathan Gruber, Anupam Jena

Education

A.S. in Mathematics, Cuesta Community College, 2008 – 2010

B.A. in Economics, *honors*, University of California, Berkeley, 2010 – 2012

Ph.D. Candidate in Health Policy & Economics, Harvard University, 2014 – present

Research Experience

Research Assistant for Professor David Card, National Bureau of Economic Research, 2011 – 2012

Research Assistant for Professor Jonathan Gruber, National Bureau of Economic Research, 2012 – 2014

Teaching Experience

Lead Math Tutor, University of California, Berkeley, 2010 – 2012

Grader, Harvard Business School, Technology and Operations Management Unit, 2016 – 2017

Fellowships, Honors, and Awards

Graduate Research Fellowship, National Science Foundation, 2014 – 2017

Inequality and Social Policy Fellowship, Weiner Center for Social Policy, 2015 – 2017

Research Grant, Mossavar-Rahmani Center for Business & Government, 2017

Works in Progress

"Does the Patient-Provider Relationship Matter for Health Care Quality: Evidence from Worker Turnover"

Abstract: There is significant discussion in health care about the importance of continuity of care, but very little analysis about how health provider turnover affects patient care. Using a difference-in-differences design and both Medicare and athenahealth, Inc. data, this project has three main aims. First, it will quantify how patient outcomes are affected when replacement providers have a different level of skill, tenure at the firm, and/or a disruption to a patient's continuity of care. Second, it will quantify spillover effects on providers and patients who remain at the firm after a provider departs. Third, it will explore how providing care as a solo practitioner versus in teams may lessen turnover costs. Better understanding the costs of losing primary care providers will aid policymakers in effectively designing policies to mitigate turnover costs and improve the quality of patient care.

"Adverse Selection and Imperfect Competition in Health Insurance Markets" (with Victoria Marone)

Abstract: Policies aimed at reducing market power in health insurance markets may in fact harm consumers when adverse selection is also present. We begin by showing that when selection can occur along an intensive margin, theoretically there are circumstances in which equilibrium price distortions arising from insurer market power (namely, markups) can increase both social surplus and consumer surplus, where the necessary, but not sufficient, condition is that low coverage plans have higher markups than high coverage plans. This result runs contrary to what is true of the case where selection can only occur along an extensive margin. We then use proprietary data on the health insurance market of a large, fully-insured employer, in which both adverse selection and imperfect competition exist, to show that we in fact observe low coverage contracts to be marked up more than higher coverage contracts, and thus empirically demonstrate imperfect competition working against the forces of adverse selection. Using a structural model of demand for and supply of insurance contracts, we provide evidence that fixed employer subsidies are the primary cause of this pattern of markups. We perform various counterfactuals to explore the full nature of this result and characterize its generalizability beyond the single setting we consider.

"Short-Term High-Interest Borrowing and Low Income Workers" (with Ariella Kahn-Lang)

Abstract: There is an ongoing debate between policymakers who view short-term, high-interest loans as a crucial source of credit for consumers managing budget shortfalls versus those who believe that these loans offer predatory interest rates and are ultimately harmful to borrowers. This is partially due to the dearth of evidence on why borrowers turn to these loans, which offer seemingly exorbitant interest rates, rather than saving or cutting costs. In order to better answer this question we quantify the circumstances surrounding demand for high-interest, short-term loans as well as the impacts of short-term credit on economic well-being and mobility.

"The ActionHealthNYC Experiment: Coordinating Unauthorized Immigrants Health Care" (with Jonathan Gruber)

Abstract: The ActionHealthNYC program is a year-long randomized control trial that provides health care to New York City unauthorized immigrants with incomes at or below 200% of the federal poverty level. It provides the first opportunity to rigorously investigate the effect of providing health care to unauthorized immigrants, a group that is largely uninsured. Using both survey and administrative data, we examine how the intervention affects the populations' utilization of health care, health status, financial strain, and trust in medical institutions.

"How the Onset of Medicare Coverage Affects Opioid Use for Pain" (with Nicole Maestas and Tisamarie Sherry)

Abstract: Medical use of prescription opioids for pain has risen dramatically over the past two decades. This increase has coincided with a period of insurance coverage expansions, notably the establishment of the Medicare Part D prescription drug benefit in 2006. Access to insurance that reduces cost-sharing for prescription drugs has been shown to increase drug utilization (Ghosh et

al. 2017), including in the specific case of Medicare Part D (Yin et al. 2008). We therefore examine the impact of Medicare Part D implementation on the treatment of pain among adults age 65 and older, and in particular the use of prescription opioids in this age group. Using a combined regression discontinuity and differences-in-differences design, we use data from the National Ambulatory Medical Care Survey (NAMCS) to compare the incidence of pain diagnoses and prescription opioid use for non-disabled adults just below and above age 65, both before and after the roll-out of Medicare Part D in 2006. This allows us to characterize how pain management in older adults changes around the time of Medicare enrollment, and whether Part D had any further incremental impact on prescription opioid use. We find that the probability of being diagnosed with a pain condition increases when individuals turn 65 (and hence become eligible for Medicare) after the introduction of Part D in 2006, but there is no significant increase in the probability of being prescribed an opioid. We then provide descriptive evidence that the growth in opioid utilization among the Medicare-eligible is due to the stock of opioid users aging into Medicare, as opposed to the onset of Medicare eligibility at age 65.

"Opioids without Pain? Medical Indications for Opioid Prescriptions in Ambulatory Care in the United States, 2006-2011" (with Nicole Maestas and Tisamarie Sherry), *submitted*

Abstract: Medical use of opioids has grown dramatically over the past two decades, far exceeding increases in the prevalence of pain. This discrepancy has raised concerns about the appropriateness of physicians' prescribing practices, and specifically whether patients' medical indications justify opioid therapy. We therefore examine the medical indications associated with opioid prescriptions in ambulatory care between 2006 and 2014, to determine the share of prescriptions written for conditions that cause pain.

Professional Activities

Referee Service: Journal of Health Economics, The American Journal of Managed Care, The Journal of Delivery Science and Innovation
Presentations: AcademyHealth Annual Research Meeting (2017)
Restricted Data: Special Sworn Status at the U.S. Census Bureau