Learning and Efficiency in the Market for Physician Referrals

PRELIMINARY

Ian M. McCarthy
Emory University & NBER

Seth Richards-Shubik
Johns Hopkins University & NBER

March 2026

Abstract

In many areas of care, primary care physicians (PCPs) greatly influence a patient's choice of specialist, but how well do PCPs learn about specialist quality, and how does this affect market allocations and patient outcomes? In this paper, we study PCP referrals to specialists, using the population of orthopedic procedures for Medicare fee-for-service beneficiaries from 2008 through 2018. We document substantial heterogeneity in specialist quality and cost within geographic markets, and we present an event study showing that PCPs adjust their referrals specifically based on the outcomes of their own patients. We then employ a structural learning model to measure learning and efficiency in this market. In the model, PCPs update their beliefs about specialist quality based on the outcomes of their patients. The model also accounts for habit persistence and capacity constraints, which are important limitations on potential reallocations when there is better information about quality. In both analyses, we find that PCPs respond modestly to the outcomes of their referred patients. The effects of these informational frictions are substantial: one-third of the patients would be referred to a different specialist under perfect information, and there would be small but meaningful improvements in patient outcomes.

1 Introduction

The lack of complete information about quality differences across products and providers is a major barrier to market efficiency in health care. Although patients rely on highly trained experts such as primary care physicians to recommend specific treatments (e.g., medications) or individual specialists (e.g., surgeons), large quality differences are observed across the commonly chosen alternatives within a given market, even in settings with minimal variation in prices.[1]

A substantial literature has developed using structural learning models to measure the effects of incomplete information about quality on consumer choices and market efficiency in various industries (see Ching, Erdem and Keane, 2013, 2017, for surveys). In health care, these models are often applied to prescribing decisions, where physicians learn about the quality of new pharmaceutical drugs (e.g., Coscelli and Shum, 2004; Ching, 2010; Ferreyra and Kosenok, 2011; Chan, Narasimhan and Xie, 2013) or patient-specific matches with different drugs (e.g., Crawford and Shum, 2005; Dickstein, 2018).[2] Broadly, these papers find that uncertainty about quality substantially delays the adoption of new treatment options and limits market efficiency.[3]

In this paper, we use a structural learning model to study referrals from primary care physicians (PCPs) to orthopedic surgeons for major joint replacements. Referrals are a major determinant of market allocations in health care and consequently affect a large amount of economic activity. Over $2.5 trillion is spent on clinical services and hospital care in the US each year, and the vast majority of this is in specialty care rather than primary care.[4] Existing research has found that PCP referrals greatly influence patients' choices of specialized providers (Freedman, Kouri, West and Keating, 2015; Gaynor, Propper and Seiler, 2016; Barkowski, 2018; Chernew, Cooper, Hallock and Scott Morton, 2021). Therefore, PCPs acting on incomplete information about specialist quality could have profound consequences for health outcomes, health care expenditures, and health care equity. In addition, prior

[1] See, e.g., Kolstad and Chernew (2009) and Dranove and Ody (2019) for surveys.

[2] In addition, Chernew, Gowrisankaran and Scanlon (2008) and Gong (2018) have applied structural learning models to choices about about health insurance plans and surgical procedures, respectively.

[3] Crawford and Shum (2005), however, find that learning about patient-specific matches is quite fast in their application.

[4] National Health Expenditures: `www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data`.

work has found that referring physicians consider past patient experiences in future referral decisions (Schneider and Epstein, 1996; Barnett, Keating, Christakis, O'Malley and Landon, 2012), which motivates our use of a learning model based on patient outcomes.

We estimate the model with data on over 4.5 million joint replacement surgeries in the United States, using the universe of Medicare fee-for-service (FFS) claims from 2008 to 2018. In the model, a specialist's quality is defined as the probability of a positive or negative outcome from each surgery. PCPs refer their patients among the set of available surgeons in their geographic market, and they learn about the qualities of those specialists from the outcomes experienced by their patients. Given the FFS environment, the PCP's referral is unaffected by insurance network restrictions, allowing a referral to any relevant specialist in the patient's geographic market.[5] Differences in the histories of patient outcomes across PCPs referring to the same specialist provide the variation needed to identify the learning process, *net of specialist fixed effects*, which are included in the model to absorb other demand factors not related to quality. We solve both a myopic and a forward-looking version of the model, the latter using a highly tractable solution method developed by Gittins (1979) and Gittins and Jones (1979). The model also accounts for habit persistence in referral choices, and for capacity constraints using an approximation proposed by Richards-Shubik, Roberts and Donohue (2022). This is crucial for counterfactuals where informational frictions are reduced, because the capacity constraints of individual surgeons may limit the market's ability to reallocate patients toward better specialists.

Before applying the model, a simple descriptive analysis of the heterogeneity in quality within geographic markets suggests that meaningful gains would be possible if patients could be reallocated. Comparing the rates of all negative health outcomes—primarily surgical complications and hospital readmissions, but even death in some cases—between surgeons at the 75th and 25th percentiles within each market, we find these differences are between 4 and 6 percentage points in most markets. These are large differences relative to the national average failure rate of 9.5 percentage points. We also provide suggestive evidence that surgeons in the top quarter of the quality distribution appear to have capacity available to take on more patients. Accurately quantifying the potential gains (or conversely, the current

[5]With other types of health insurance, a referral could be an administrative requirement as well as an informative recommendation.

losses) due to information about quality, while allowing for capacity constraints and habit persistence, is the ultimate goal of this study and the reason a structural model is needed.

The results from the model indicate that there are substantial losses due to informational frictions. With perfect information about specialist quality, our counterfactual simulations show that about one-third of patients would be referred to different specialists. The rate of negative outcomes would decrease by 5%, in relative terms, and episode spending would fall similarly. Because the simulations account for specialist capacity constraints and habit persistence in referrals, these changes could reasonably be attributed to informational frictions.

In addition to the structural analysis, we provide design-based evidence that PCPs reduce their referrals to a surgeon after one of their patients experiences a negative outcome with that surgeon. We find that, among the PCPs sending patients to a given surgeon, the number of referrals per quarter drops by about 5 percent for the PCP whose patient has a negative outcome, relative to the other PCPs referring to that surgeon. This is strong evidence of a learning process based on the outcomes of a PCP's own patients. The event-study design controls for surgeon fixed effects, and comparing the "treated" PCPs who observe a negative outcome from each surgeon with the "control" PCPs who do not, we see negligible differences in pre-trends before the event, and no convergence after the event, using up to nine quarters of leads and lags.

In terms of the modeling framework in a health application, our work is most related to two working papers by Gong (2018) and Dickstein (2018), which use the same class of learning models (Bayesian models with beta distribution beliefs and binary outcomes) and study a particular treatment decision. Gong (2018) examines learning about a new surgical procedure for brain aneurysms, where individual surgeons learn about their own skill with the new procedure compared to an established alternative. Patient health outcomes are observed in that application, as in ours, while the choice set is much smaller (three treatment options in Gong (2018) vs. an average of over 80 orthopedic surgeons for a given patient in our analysis). Dickstein (2018) considers antidepressant prescribing, where there are relatively large choice sets (nineteen drugs), but health outcomes are not observed, so the estimation procedure integrates out the possible sequences of outcomes and relies entirely on the observed sequences of choices.

Other papers applying structural learning models to treatment decisions use a different

class of models, typically assuming normally distributed beliefs and outcomes, without directly observing patient outcomes. The majority of this literature focuses on learning about the quality of prescription drugs. For example, Coscelli and Shum (2004), Crawford and Shum (2005), and Ferreyra and Kosenok (2011) study physician learning in the anti-ulcer drug market. Ching (2010), Chan et al. (2013), and Dickstein (2018) similarly consider learning with regard to side effects and treatment effectiveness of other pharmaceuticals. In addition to this work on treatment decisions, Chernew et al. (2008) study consumer learning about health insurance plan quality, in a two-period setting where signals are observed, using both classes of models (beta-binomial and normal-normal).

Our work also relates to the rapidly growing literature on physician networks and referrals. Starting with Barnett et al. (2011), researchers have inferred professional relationships among physicians from shared patients and have examined aggregate associations between features of the inferred networks and variations in expenditures and utilization across geographic markets. Kaur, Perloff, Tompkins and Bishop (2017) and Agha, Ericson, Geissler and Rebitzer (2022) study referrals at the provider level for specific areas of care and find that greater concentration of referrals to fewer specialists is associated with somewhat lower expenditures and utilization. Also, Zeltzer (2020) and Sarsons (2023) consider the role of gender homophily in referrals and the subsequent effect on earnings differences among physicians. Our descriptive analysis of the referral networks for major joint replacements offers new information on heterogeneities in network structure and quality outcomes in this important area of care where referrals can be reasonably well inferred.

Last, in addition to our empirical analysis, we contribute to the understanding of identification in structural learning models. While most existing papers in this literature provide thoughtful, verbal discussions of identification, we develop a formal argument for the identification of all the parameters in our model, including the weight placed on the patient's health outcome in the physician's utility.[6] Notably, some previous studies omit this parameter, which implicitly fixes the marginal rate of substitution between quality and all the factors in the idiosyncratic shocks at an assumed value. While this is innocuous in some cases (e.g., where the outcomes are unobserved and the scale of the outcomes is estimated),

[6]Ching et al. (2013) provide a clear verbal argument for a similar parameter. We formalize this with an "identification at infinity" type of argument. Also, although not applicable to our model, Bunting et al. (2024) have recently developed general results on identification in normal-normal learning models.

it may be restrictive in others. We also provide a small simulation study on the identification of a parameter that governs the precision of the prior beliefs, which is imprecisely estimated in our data and seems to be problematic in some prior studies as well.

The remainder of the paper is organized as follows. We first explain the data and present a descriptive analysis of the referral networks for major joint replacements in Section 2, which includes a simple exercise to illustrate the potential gains from more efficient referrals within geographic markets. In Section 3, we present design-based evidence on how PCPs change their referral patterns in response to negative patient outcomes, using an event study. Section 4 turns to our learning model, where we specify the Bayesian learning process and other key features, briefly explain the solution method, and discuss identification. Finally, Section 5 describes the details of our structural estimation, provides our estimates of the myopic and forward-looking models, and presents the results from our counterfactual simulations. Section 6 concludes.

2 Data and Descriptive Statistics

In order to focus on a predictable care pathway where patients are likely to rely on PCP referrals, our analysis considers planned and elective major joint replacements performed in an inpatient setting. We have data covering all Medicare FFS beneficiaries aged 65 and above, from 2008 to 2018. The relevant procedures, outcomes, and referrals are identified using 100% Medicare FFS claims files, containing all Part A and Part B claims. Data on patient characteristics come from Medicare Beneficiary Summary Files (MBSF), and data on physician practice characteristics come from Medicare Data on Provider Practice and Specialty (MD-PPAS). We also obtain data on hospital characteristics from the American Hospital Association (AHA) Annual Surveys as well as the Provider of Services (POS) files.

The relevant inpatient procedures are identified from the diagnosis related group (DRG) codes on the claims,[7] and the admission source codes on the claims indicate planned and elective procedures. In total there are just over 4.53 million inpatient stays for planned and elective major joint replacements from 2008 to 2018. Among this set of inpatient stays,

[7]We use DRG codes 461, 462, 469, 470, 480-489, and 507-51. For a comprehensive list of DRGs and descriptions, see http://data.nber.org/drg/drgdesc11.pdf.

we limit our analysis to surgeries performed by an orthopedic surgeon, as indicated by the specialty codes and descriptions in the MD-PPAS data.

We infer the referring PCP based on the greatest frequency of "evaluation and management" (E&M) visits to PCPs over the prior 12-month period before a given surgery, following Pham et al. (2009) and Agha et al. (2017).[8] Eligible PCPs are indicated by the specialty codes in the MD-PPAS data, and we require that patients visited the referring PCP at least 2 times in the prior 12-month period. Restricting to procedures performed by orthopedic surgeons and where a referring PCP can be identified reduces the sample to about 2.97 million inpatient stays.

Finally, we impose two additional restrictions on the sample in order to focus on physicians with sufficient experience in diagnosing and treating orthopedic patients. First, we restrict our sample to PCPs with at least 20 total referrals over our time period and with at least three consecutive years with one or more referrals. Second, we restrict our sample to surgeons with at least 20 operations per year. Our final analytic sample consists of just over 2 million planned and elective major joint replacements, in which the referring physician is a PCP with regular referrals for these procedures and in which the operating physician is an orthopedic surgeon with sufficient yearly volume.

As our measure of quality, we follow the CMS Comprehensive Care for Joint Replacement Model and the National Quality Forum. Intuitively, our quality measure derives from mortality, readmissions, and complications within 90 days of discharge.[9] All quality outcomes are observed from claims during and following the focal inpatient stay, which are identified in the data using the patient's unique identification number. We also define measures of total spending and health care utilization around this episode of care by collecting all claims over the 90-day period starting with the admission date for the surgery.

Table 1 presents the basic descriptive statistics of our final sample. The unit of observation is one patient and procedure—i.e., an individual referral. For each referral, we observe

[8]Zeltzer (2020) and Sarsons (2023) instead identify referrals using the "Claim Referring Physician NPI Number" field in the carrier claims data. However in our data we find that this field is missing in 34% of the relevant claims, and it contains the same identification number as the operating physician in 64% of the claims where it is populated.

[9]As per the Comprehensive Care for Joint Replacement Model, some complications are only tracked during the inpatient stay or within 30 days of discharge. Additional details of these measures are presented in the supplemental appendix.

the operating surgeon, the referring PCP, and the admitting hospital. As discussed in more detail in Section 5.3, our estimation process first divides the data into two periods: 1) the baseline period from 2008-2012, which provides a common time frame from which to form observed histories for PCP/specialist pairs; and 2) the estimation period from 2013-2018, for which the baseline period provides the initial values of the histories. From Table 1, we see fairly similar patients, physicians, practices, and outcomes between the baseline and estimation periods. One apparent difference is that PCP and specialist practice sizes are larger in our estimation period, consistent with broader national trends in practice consolidation. We also see a slight reduction in readmissions, mortality, and complications in the estimation period compared to the baseline period.

2.1 Description of PCP Referral Networks

Using the physician's national provider identification (NPI) number, we identify 51,067 unique PCPs and 9,254 unique specialists over our estimation period (2013-2018). As summarized in Table 2, each PCP in our sample refers an average of 4.4 patients for elective orthopedic surgery per year to an average of 3.1 unique orthopedic specialists per year. PCPs therefore refer an average of roughly 1.4 patients to a given specialist in a year, conditional on sending any patients to that specialist that year.

We refer to the number of unique specialists referred to over some time period as the "network size" of a PCP. The distribution of PCP referral network sizes combining the six years from 2013 to 2018 is presented in Figure 1(a). This shows the total number of unique specialists to which a PCP refers patients over the estimation period (for ease of presentation, the top 1% of PCPs with network sizes above 50 are not included). Consistent with the yearly average network sizes in Table 2, the typical PCP sends patients to relatively few specialists (around 10) over the entire estimation period. The frequency distribution also reveals a long tail, indicating the presence of a small set of PCPs with very large network sizes. However for most PCPs, these network sizes are substantially smaller than the total number of orthopedic surgeons in the market, which is around 69 per year on average.

To assess the concentration of PCP referral networks, Figure 1(b) presents the distribution of the proportion of a PCP's referrals going to their "highest-share" specialist. Here,

we calculate each specialist's share of a given PCP's total referrals from 2013 to 2018 and take the maximum of these shares for each PCP. As is evident from Figure 1(b), many PCPs send 20% to 40% of their referrals to a single specialist, suggesting that the modal PCP's referral network is quite concentrated among a small number of specialists.

Table 2 also presents the average number of negative outcomes, or "failures," associated with a PCP's patients and the failure rates per referral. Failures are infrequent but not necessarily rare. On average, there are nearly 0.4 failures among a PCP's patients each year. Per referral the failure rate is about 0.1; in other words, around 9% of a PCP's referrals for major joint replacements result in some form of failure, defined based on the National Quality Forum's measure of quality in joint replacement surgery.

2.2 Within-market Variation in Specialist Quality and Efficiency

As reflected in Table 2 and Figure 1(a), most PCPs have experience referring patients to between 3 and 5 unique specialists per year and around 10 unique specialists over our entire 2013-2018 estimation period. For context, we observe an average of nearly 69 orthopedic surgeons performing a major joint replacement per market per year. So the average yearly network size of about 3.1 specialists implies that an average PCP remains unfamiliar with over 95% of available specialists in their market in a given year. This suggests a lack of experimentation with other specialists and therefore a potential opportunity for quality improvements and cost reductions if PCPs could better allocate their referrals to higher quality and more efficient specialists in their market.

To illustrate the potential gains from such a reallocation, we consider the variation in failure rates and episode spending across specialists *within* each HRR. Specifically, Figure 2(a) plots the differences between the 75th and 25th percentiles of surgeon-specific failure rates within an HRR. This illustrates the hypothetical reduction in failure rates if PCPs could replace referrals to surgeons at the 75th percentile of failure rates (within their market) with referrals to surgeons at the 25th percentile. The implication of Figure 2(a) is that failure rates could be reduced by between 4 and 6 percentage points in most markets, for those patients who are reallocated from surgeons at the highest quartile of failure rates (with a mean of 13.0%) to surgeons at the lowest quartile (with a mean of 6.5%). We present a

similar analysis of 90-day episode spending in Figure 2(b), which suggests possible savings of nearly $8,000 per episode when moving from specialists initiating relatively high-spending episodes (75th percentile of the market-specific spending distribution) to specialists initiating relatively low-spending episodes (25th percentile).

A simple descriptive analysis of specialist capacity further suggests that such a reallocation of patients (from lowest-performing to highest-performing quartile of specialist quality) may be feasible, as illustrated in Figure 3. Figure 3(a) presents a histogram of "potential excess capacity" among the relevant (highest-performing) specialists per HRR and year. We calculate excess capacity as the difference between a specialist's yearly operations and that same specialist's 75th percentile of yearly operations across all years. We then limit to the best 25 percent of specialists with the lowest failure-rates, and we sum their specialist-specific excess capacities for each HRR and year. Figure 3(b) presents the total count of operations by (lowest-performing) specialists in the highest quartile of the failure-rate distribution per HRR and year, therefore summarizing the number of possible patients to be reallocated. Note that these totals are not restricted to planned and elective procedures, and instead are based on a broader set of all procedures in which a given specialist is listed as the operating physician. From this descriptive analysis, over 39% of markets appear to have sufficient capacity among top-quartile specialists to accommodate the potential reallocation from bottom-quartile specialists. While there are many other demands on specialists' time outside of surgical procedures, Figure 3 is at least consistent with the presence of excess surgical capacity to accommodate some within-market reallocations of patients.

3 Do PCPs Respond to Bad Outcomes?

Before presenting our structural learning model, we first consider design-based evidence of PCP responses to negative surgical outcomes of their patients. We estimate this response by exploiting variation in the patient outcomes across PCPs referring to the same specialist, using event studies that compare referral rates between PCPs whose patients do and do not experience failures following procedures by the same specialist.

The sample and data for this analysis are constructed as follows:

1. Create a quarterly panel of all PCP-specialist pairs with at least one referral between

them. Quarters in which there are no referrals are treated as zeros. We also weaken the restriction on specialist volume from requiring 20 operations performed in each year to requiring 20 operations performed in at least one year over our panel.

2. Find all specialists, indexed with j, with at least one failure event during the estimation period from 2013 through 2018. For each failure event, $f = 1, ..., F_j$, denote the quarter of each event by $q_j(f)$, so that $q_j(1)$ denotes the quarter of first failure for specialist j, $q_j(2)$ denotes the quarter of the second failure, etc. Further denote by $\underline{q}$ and $\overline{q}$ the first and last quarter of the analysis, respectively.

3. Find all PCPs who ever refer to specialist j between $\underline{q}$ and $\min(q_j(2) - 1, \overline{q})$, which is either the quarter before the second failure or the last quarter in the data. Drop all observations outside of these time windows, as defined for each specialist. We construct similar cohorts based on the second, third, and fourth failure events.

4. Identify the PCP(s) whose patient(s) experienced a failure following a referral to specialist j in $q_j(1)$.[10] Denote this PCP (or set of PCPs) as type $k = 1$, which constitutes the treatment group for this analysis. Similarly, identify the other PCPs whose patients do not experience a failure with specialist j over the relevant time period, and denote this set of PCPs as type $k = 0$, which is the control group.

5. Calculate the means by PCP type of the number of referrals per PCP per quarter to specialist j, removing the referral(s) that resulted in failure. Denote the mean quarterly patients referred to specialist j from PCP type k by $\bar{r}_{jkt}$.[11]

With this notation and sample construction, we then estimate by OLS the following event study specification:

$$\bar{r}_{jkt} = \gamma_{j,t} + \delta I(k = 1) + \sum_{\substack{\tau=-9 \\ \tau \neq -1}}^{9} \lambda_\tau I(k = 1, t = \tau) + \varepsilon_{jkt}, \tag{1}$$

[10] Typically only one patient experiences a failure from a given specialist in a given quarter, but there are cases where specialists are associated with multiple failures in the same quarter.

[11] We collapse the data by "PCP type" primarily to ease the computational burden. Results are unchanged if we instead estimate Equation 1 at the PCP-specialist-quarter level, allowing for both specialist and PCP fixed effects.

where $\gamma_{j,t}$ denotes specialist-by-quarter fixed effects, $I(k = 1)$ is an indicator set to one for the treated PCPs (i.e., those PCPs whose referrals to specialist j experienced a failure), and $I(k = 1, t = \tau)$ is an indicator set to one if the quarter is in period τ relative to the failure quarter, $q_j(1)$.[12] We estimate Equation (1) separately for four treatment cohorts: the first cohort is defined using the first failure event, as specified above; the subsequent cohorts are defined in an analogous way using the subsequent failure events.[13] We then estimate a "stacked" event study that combines all four cohorts into a single analysis. This follows Cengiz et al. (2019), and we accordingly cluster standard errors by specialist and cohort.[14]

We present the main results from the stacked event study in Figure 4, where we find a statistically significant reduction of approximately 0.056 referrals per quarter per specialist from treated PCPs (i.e., those whose patients experienced a bad outcome) relative to other PCPs (i.e., those who refer to the same specialists but whose patients did not experience a bad outcome).[15] This is compelling evidence of learning specifically based on the outcomes of referred patients, because any changes in supply from the specialists with a bad outcome, or changes in demand from other PCPs, are absorbed by the specialist-by-quarter fixed effects that capture the control observations. To interpret the magnitude of the effect, the average quarterly referral rate, among treated PCP-specialist pairs with at least one referral in the pre-failure quarters, is 1.061. Our estimates therefore imply that a failure event leads to an approximately 5% reduction in referrals from the PCP whose patient experienced the failure to the specialist with whom the failure occurred.[16]

[12] Results are similar if we include specialist and quarter fixed effects separately, or specialist-by-type fixed effects and quarter fixed effects.

[13] For example, the second cohort considers the second failure event, with the sample constructed using $q_j(2)$ as the reference period for each specialist and containing all quarters from $\max(q_j(1) + 1, q_j(2) - 9)$ to $\min(q_j(3) - 1, \overline{q})$.

[14] In the stacked estimation, the same control observations may repeat across failure cohorts (e.g., first failure, second failure, etc.). We therefore include a richer set of specialist-cohort-type fixed effects, in addition to specialist-quarter fixed effects.

[15] The results for the separate failure cohorts are consistent with these stacked estimates and are presented in the supplemental appendix.

[16] The unconditional average quarterly referral rate is 0.194, which includes many quarters with zero referrals for a given PCP-specialist pair. Using this as the denominator would yield a 29% relative reduction, but this overstates the behavioral response since PCPs cannot reduce referrals below zero.

4 Bayesian Learning Model

Our structural model describes a PCP who refers patients to orthopedic surgeons for major joint replacements. There are multiple surgeons available in the market (a geographic area). When a patient arrives who needs this treatment, the PCP selects one surgeon to recommend to the patient, and the patient follows the recommendation and receives treatment from that surgeon.[17] The outcome of the surgery is binary (success or failure), and is observed by the PCP. The probability of success, which is the definition of "quality" in this model, varies across surgeons. The PCP does not know these probabilities, but they have beliefs about specialist quality that are updated based on the outcomes experienced by their patients.

We solve and estimate the model for both myopic and forward-looking behavior. Also, in addition to the learning process, the model includes two other important mechanisms that could limit the allocation of patients to higher quality surgeons: habit persistence and capacity constraints. The former is captured with a *familiarity effect* relating to the number of patients referred to a surgeon in the past, and the latter is approximated with a *congestion effect* arising from the total number of patients being treated by the surgeon at the time (i.e., referrals from other PCPs).

4.1 Model Specification

The PCP, i, refers each patient to some surgeon, j, from a set of available surgeons, J_i. Patients arrive sequentially, so patients and time can both be denoted by t. The choice of specialist is given by a set of indicators, $D_{ijt}, j \in J_i$, where $D_{ijt} = 1$ if patient t is referred to specialist j, otherwise $D_{ijt} = 0$. The binary health outcome is Y_{ijt}, with $Y_{ijt} = 1$ for success and $Y_{ijt} = 0$ for failure (i.e., complication, readmission, or death).

The probability of success with surgeon j is q_j. This is the same for all patients, such that our model features vertical but not horizontal differentiation.[18] The PCP values the patient's outcome, which could reflect intrinsic motivations such as altruism as well as ex-

[17] Freedman et al. (2015) find that the PCP's recommendation is the most commonly cited reason from a patient in their selection of an oncologist, with similar results documented in Barkowski (2018).

[18] The models in Coscelli and Shum (2004), Ching (2010), Ferreyra and Kosenok (2011), and Chan et al. (2013) similarly emphasize vertical differentiation among drugs, while those in Crawford and Shum (2005) and Dickstein (2018) consider horizontal differentiation in terms of the match quality between drugs and individual patients.

trinsic motivations such as malpractice liability. There are other factors, denoted by x_{ijt}, affecting the patient's net benefit from a particular surgeon, such as their distance to the surgeon's facility, and these factors are also considered and valued by the PCP.

Beyond the patient's outcome, the PCP may value working with specialists with whom they have some prior experience. Defining $e_{ijt} \equiv \sum_{s=1}^{t-1} D_{ijs}$ as the number of past patients referred to specialist j, this *familiarity effect* is given by $f(e_{ijt})$, where f is increasing and concave. In addition, the capacity constraints of individual surgeons generate a negative congestion effect among patients referred to the same surgeon. Let n_{jt} denote the total number of patients being treated by surgeon j around time t, which includes referrals from *other* PCPs around that time (e.g., $n_{jt} \equiv \sum_k D_{kjs},\ s \in [t-\tau, t+\tau]$ for some τ), and let z_j denote fixed attributes that affect a surgeon's capacity, such as non-clinical work (e.g., administrative responsibilities or academic research). The *congestion effect* is then given by $c(n_{jt}, z_j)$, where c is decreasing in n.

The PCP's realized utility from referring patient t to specialist j is

$$U_{ijt} \equiv \alpha Y_{ijt} + u(x_{ijt}) + f(e_{ijt}) + c(n_{jt}, z_j) + \xi_j + \epsilon_{ijt}, \tag{2}$$

where α is the weight on the patient's outcome, $u(x_{ijt})$ is the value placed on the other patient-specific factors, and $f(e_{ijt})$ and $c(n_{jt}, z_j)$ are the familiarity and congestion effects described above. The term ξ_j is a specialist fixed effect, which captures time-invariant demand factors that need not be related to surgical outcomes (e.g., office amenities, health system branding, other advertising). Finally, ϵ_{ijt} is an idiosyncratic shock that captures other choice-specific factors, which has an assumed Type I extreme value distribution.

The presence of $c(n_{jt}, z_j)$ in U_{ijt} can be interpreted as the PCP internalizing the effects of congestion on the patient's net benefit from a surgeon. This can be given a structural interpretation based on waiting times, or it can be considered an approximation for other mechanisms. Either way, including this term is important to prevent bias in estimates of the responsiveness to patient outcomes, given here by α and the learning parameters, and counterfactuals based on those estimates (see Richards-Shubik et al., 2022).

4.2 Learning Process

The PCP does not know the quality of each specialist, but rather has beliefs about the possible values of q_j, for each j. These beliefs are specified with the beta distribution, denoted Beta(a, b), which is a natural and tractable modeling choice when outcomes are binary or binomial. All PCPs have the same initial beliefs (i.e., priors) about all specialists in their market, with the parameters equal to (a_0, b_0).

PCPs use Bayesian inference to update their beliefs about q_j and thereby learn about the quality of each specialist based on the outcomes experienced by their patients. Specifically, the parameters (a, b) are updated based on the numbers of successes and failures among the patients referred to specialist j in the past, as follows:

$$a_{ijt} = a_0 + \sum_{s=1}^{t-1} Y_{ijs} \quad \text{and} \quad b_{ijt} = b_0 + \sum_{s=1}^{t-1} (D_{ijs} - Y_{ijs}).$$

Then, from the beta distribution, the mean and variance of the beliefs about q_j are given by

$$m_{ijt} \equiv \frac{a_{ijt}}{a_{ijt} + b_{ijt}} \text{ and } v_{ijt} \equiv \frac{a_{ijt} b_{ijt}}{(a_{ijt} + b_{ijt})^2 (a_{ijt} + b_{ijt} + 1)}. \tag{3}$$

Thus, m_{ijt} denotes the Bayesian expectation of the probability of success for referral t to specialist j, which is a function of the priors (a_0, b_0) and the observed outcomes among patients previously sent to that specialist. More successes will increase a_{ijt}, thus increasing m_{ijt}; however, the magnitude of this increase is affected by (a_0, b_0). Larger values of a_0 and b_0 decrease the responsiveness of m_{ijt} to patient outcomes, and we accordingly refer to the sum of $a_0 + b_0$ as the "strength" of the priors. Intuitively, a_0 and b_0 enter the updated beliefs like fictitious numbers of successes and failures, so the sum $a_0 + b_0$ is like a fictitious number of patients referred to specialist j before the first actual patient. Finally, because the variance, v_{ijt} in Equation (3), decreases in $a + b$, this sum is sometimes referred to as the "precision" of the beliefs (analogously, $a_0 + b_0$ is the precision of the priors).

4.3 Myopic and Forward-Looking Behavior

A myopic PCP simply chooses the specialist with the highest expected payoff for the current patient:

$$\max_{j \in J_i} \mathrm{E}\left[U_{ijt} | \dots\right] = \max_{j} \left\{\alpha m_{ijt} + f(e_{ijt}) + u(x_{ijt}) + c(n_{jt}, z_j) + \xi_j + \epsilon_{ijt}\right\}. \tag{4}$$

Thus if PCPs are myopic, they unambiguously tend to refer to specialists with whom they have had more successes in the past, all else equal.

If PCPs are forward-looking, they also value experimenting with relatively unknown specialists. The choice of specialist for referral t involves both the utility for the current patient and the value of learning more about the quality of the specialists in the market, which could benefit future patients. The solution to this dynamic problem simplifies with the use of a *Gittins index* (Gittins, 1979; Gittins and Jones, 1979), which expresses the value of learning about specialist j as a function of the mean and variance of the current beliefs about that specialist. We denote the index abstractly as $g(m_{ijt}, v_{ijt})$, but this function is well approximated with a fairly simple and tractable closed-form expression developed by Brezzi and Lai (2002).

The Gittins index requires four key assumptions: (i) one option is chosen at a time; (ii) the unchosen options do not affect the current outcome; (iii) the expected returns do not change for the unchosen options; and (iv) beliefs are independent across options. These assumptions are not uncontroversial, but they are relatively unproblematic in our application. Typically, (i) one surgeon treats the patient, and (ii) other surgeons do not assist or interfere with the operation. For (iii) and (iv), we assume that PCPs perceive surgeons to have fixed quality that is independent across surgeons. There would be no depreciation of skills among unchosen surgeons (i.e., unchosen surgeons continue to receive patients from other PCPs) and no important spillovers in quality among surgeons. To this last point, we may rely on the surgeon fixed effects to absorb time-invariant correlations among surgeons, for example among those who operate in the same hospitals.

With forward-looking behavior, the PCP's present discounted utility is given by value

functions, which can be written abstractly as follows:

$$V_{it}(\dots) = \max_{j \in J_i} \left\{ \mathrm{E}\left[U_{ijt} | \dots \right] + \beta \mathrm{E} V_{i,t+1}(\dots) \right\},$$

where β is a known discount factor, and the ellipses $(\dots)$ represent the relevant state variables. These value functions simplify because there are no dynamics in x, n, z, ξ, and ϵ. Patients are assumed to arrive randomly, which makes x exogenous, z and ξ are fixed over time while ϵ is independent, and we assume there are enough patients and PCPs such that no individual PCP can affect n. Hence the components of the future value function related to these terms do not vary across choices in the current period. Because they are also additively separable in the utility function, they drop out of the differences across choices, and therefore they can be ignored in estimation. This leaves the value of learning about the quality of a specialist, captured by the Gittins index, and the present discounted value of increasing familiarity with a specialist. The former has a closed-form approximation as noted above, and the latter is computed as the present discounted value of sending every future patient to the same specialist, denoted by $\overline{\overline{f}}(e_{ijt})$.[19]

Thus, the forward-looking choice problem simplifies as

$$\max_{j \in J_i} \left\{ \alpha g(m_{ijt}, v_{ijt}) + \overline{\overline{f}}(e_{ijt}) + u(x_{ijt}) + c(n_{jt}, z_j) + \xi_j + \epsilon_{ijt} \right\}. \tag{5}$$

Unlike the myopic model, the forward-looking model assigns some value to experimenting with specialists with whom the PCP has less past experience. That is because the Gittins index is increasing in the variance, v_{ijt}, and the variance is decreasing in the number of past patients (as seen in Equation (3), v_{ijt} is decreasing in $a_{ijt} + b_{ijt}$). It is therefore possible for the Gittins index to assign a higher value to a specialist with a lower current expectation for the probability of success but a larger variance: e.g., $m_{ijt} < m_{ikt}$, $v_{ijt} > v_{ikt}$, and $g(m_{ijt}, v_{ijt}) > g(m_{ikt}, v_{ikt})$.

[19] The Gittins index is defined as a lump sum payment that makes an individual indifferent between using one option and retiring it. Because f is increasing in e, the value from f is maximized by sending every future patient to the currently chosen specialist.

4.4 Identification

Here we develop a formal argument for the identification of the parameters in the myopic model, and we offer a simple conjecture for the identification of the forward-looking model. We also discuss simulation results that indicate the difficulty of estimating the strength of the priors $(a_0 + b_0)$ in practice.

First, given a distribution for the shocks ϵ, the conditional choice probabilities identify differences in utility between the alternatives. The effects of the pairwise characteristics, $u(x)$, and the specialist fixed effects, ξ, are identified subject to standard normalizations (e.g., one fixed effect is set equal to zero). Because the number of patients at a specialist, n_{jt}, is endogenous, the congestion effect, $c(n_{jt}, z_j)$, is identified with an instrument for the probability that patients *other than patient* t see specialist j. A natural instrument for this is the distances between *other* patients and specialist j (see the description of the estimation procedure in Section 5.3, and Richards-Shubik et al., 2022, for further details).

This leaves the terms αm and $f(e)$ in Equation (4). The parameter α is identified in the limit as the number of past patients sent to a specialist grows large (i.e., so-called "identification at infinity").[20] To show this, recall $e_{ijt} \equiv \sum_{s=1}^{t-1} D_{ijs}$ and let $\bar{y}_{ijt} \equiv \sum_{s=1}^{t-1} Y_{ijs}/e_{ijt}$. Then αm_{ijt} from Equation (3) can be rearranged as

$$\alpha m_{ijt} = \alpha \frac{a_0}{(a_0 + b_0) + e_{ijt}} + \alpha \frac{e_{ijt}}{(a_0 + b_0) + e_{ijt}} \times \bar{y}_{ijt}. \tag{6}$$

In the limit as $e_{ijt} \to \infty$, we have $\alpha m_{ijt} = \alpha \cdot 0 + \alpha \cdot 1 \cdot \bar{y}_{ijt}$. Thus, intuitively, α is identified by variation in the success rate among specialists to whom a PCP has sent many patients in the past.

For the next step in the argument, it is useful to re-parameterize a_0 and b_0 as $\eta \equiv (a_0 + b_0)$ and $\rho \equiv \frac{a_0}{a_0 + b_0}$. The strength of the priors, η, is identified by the interaction between e_{ijt} and $\bar{y}_{ijt}$ when e_{ijt} is finite. This interaction appears only in the second term in (6), where η is the remaining unknown parameter (since α is identified in the limit). Because η is positive, the marginal effect of the success rate $(\bar{y}_{ijt})$ increases with the number of past patients (e_{ijt});[21] η governs this interaction—i.e., how quickly the marginal effect of the success rate increases

[20] A related argument is given verbally in Ching et al. (2013), for a normal-normal learning model.

[21] Note that $e/(\eta + e)$ goes from 0 to 1 as e goes from 0 to ∞.

with the number of past patients referred to a specialist.

To identify the prior mean, ρ, we use the observed average success rate in the market. This assumes that PCPs have a general awareness of the overall success rate of the procedure in their local area, and have rational expectations.[22] It might instead be possible to identify ρ (equivalently a_0) from the choice data, but this would require a strong functional form assumption to separate the first term in (6), $\alpha\frac{a_0}{(a_0+b_0)+e}$, from the familiarity effect, $f(e)$.

Finally, then, the familiarity effect is identified by the remainder of the marginal effect of experience with a specialist that is not absorbed by αm (i.e., what is left over from $\alpha\frac{a_0}{(a_0+b_0)+e}$). A standard normalization is required, such as $f(0) = 0$. Also f must be bounded as $e \to \infty$; otherwise, familiarity would dominate the effects of all other attributes as $e \to \infty$. This does not seem desirable from a modeling perspective, and it would nullify the argument for the identification α in the limit.

Most of the related papers in the literature provide thoughtful, verbal discussions of identification. The formal argument above offers additional insights because it makes precise the variation that is crucial for each parameter. In particular, it shows how the parameter α is identified. This parameter, which gives the utility weight on the outcome (or "use experience") is omitted from the models in Crawford and Shum (2005), Ferreyra and Kosenok (2011), Dickstein (2018), and Gong (2018) but included in those in Ching (2010) and Ching et al. (2013), for example. In some cases, this omission may be innocuous because outcomes are not observed, so the scale of the outcome is estimated and naturally incorporates the utility weight (i.e., the scale is in terms of utility). However, it is not clear that this is unrestrictive when outcomes are binary, as in Dickstein (2018) and Gong (2018), which fixes the scale of the experienced outcome even if it is unobserved. This is potentially important because the shocks in these models have a parametric distribution with an assumed scale. Consequently, fixing the utility weight on the outcome would imply that it has the identical impact on utility as all the factors captured in the shocks—or equivalently, the marginal rate of substitution between the outcome and the shocks is assumed to be known.

Separately, simulation exercises suggest that the strength of the priors, $\eta \equiv (a_0 + b_0)$, may be poorly identified in practice (see the supplemental appendix for details). While α

[22] Coscelli and Shum (2004), Crawford and Shum (2005), Chan et al. (2013), and Dickstein (2018) similarly use a rational expectations assumption to identify the prior beliefs in their models.

is recovered reasonably well in simulated data, η is noisy, and the likelihood function with respect to η is relatively flat. For some intuition into the problem, consider the second term in Equation (6), which shows how η is identified from the interaction between the number of past patients (e_{ijt}) and the success rate among those patients ($\bar{y}_{ijt}$), as discussed above. From that term (omitting α), the cross partial of e and $\bar{y}$ would be $1/(\eta + e) - e/(\eta + e)^2$, which exhibits little variation over a range of values for η. For example, using the mean number of past patients (conditional on any past referrals from i to j), shown in Table 2 (see "running referrals"), this cross partial ranges from 0.033 to 0.056 when η ranges from 1 to 20, which is little variation relative to the scale of the error terms in the model. Consequently, for the estimation of the model, we fix η at one of a set of alternative values.

Last, we turn to the identification of the forward-looking model. From Brezzi and Lai (2002), the Gittins index is closely approximated by

$$\tilde{g}(m, v) \equiv m + \sqrt{v} \cdot \psi \left(\frac{v}{-\ln(\beta)\, m(1-m)} \right), \tag{7}$$

where ψ is a known function derived by Brezzi and Lai (2002), and β is a known discount factor.[23] Substituting $\tilde{g}$ for g, the forward-looking model shown in Equation (5) has no additional free parameters beyond those of the myopic model in Equation (4). The only differences between the myopic model (4) and the dynamic model (5) are that $\tilde{g}(m, v)$ and $\overline{\overline{f}}(e)$ appear in place of m and $f(e)$, respectively. The Gittins index approximation, $\tilde{g}(m, v)$, involves the same free parameters as the mean of the current beliefs, m, which are a_0 and b_0 (or equivalently, ρ and η), and the same data (the numbers of past referrals and successes). As for the present discounted value of familiarity, $\overline{\overline{f}}(e)$, it has a known relationship with the current value of familiarity (see Gong, 2018, for an equivalent result applied to an effect of learning by doing). Thus, while this remains a conjecture, it seems reasonable to suppose that the forward-looking model is identified under the same conditions as the myopic model.

[23] Ching et al. (2013) discuss the conditions and challenges for identifying the discount factor in structural learning models. We treat it as a known parameter.

5 Estimation of the Learning Model

Below we provide details of the empirical specification, variable definitions, and estimation procedure. We then present the parameter estimates and counterfactual results.

5.1 Empirical Specification

Recall the maximization problem for a myopic PCP from Equation (4),

$$\max_{j \in J_i} \mathrm{E}\left[U_{ijt} | \dots\right] = \max_{j} \left\{\alpha m_{ijt} + f(e_{ijt}) + u(x_{ijt}) + c(n_{jt}, z_j) + \xi_j + \epsilon_{ijt}\right\}.$$

In estimating this equation, we assume that the the idiosyncratic shocks, ϵ, have a Type I extreme value distribution, which yields a multinomial logit model. The fixed effects, ξ, are free parameters, which are separately identified due to the large number of patients relative to the number of surgeons per market. Congestion is assumed to be constant within pre-specified time periods (e.g., a calendar year), so that $c(n_{jt}, z_j) \equiv c(n_{j,\tau(t)}, z_j)$, where $\tau(t)$ is the period containing t. This is for computational feasibility because, in the first step of estimation (described below), the congestion effect is absorbed into a composite fixed effect, $\delta_{j,\tau(t)} \equiv c(n_{j,\tau(t)}, z_j) + \xi_j$, and one of these must be recovered for each surgeon in each relevant time period. The capacity shifters, z_j, are omitted due to data limitations, and we specify the congestion effect to be linear so that $c(n_{j,\tau(t)}, z_j) \equiv \gamma n_{j,\tau(t)}$. There is one patient-specific factor, distance to the surgeon's facility, specified with a linear effect on utility: $u(x_{ijt}) \equiv \pi x_{ijt}$. The effect of familiarity, $f(e_{ijt})$, is specified non-parametrically using a series of indicators for different levels of e_{ijt}. Specifically, we separate e_{ijt} into bins [0], [1], [2], [3], [4], [5], [6, 7], [8, 10], [11, 15], [16, ∞) and include indicator variables κ_b for each bin, b, excluding κ_0.

For the prior beliefs, we use the re-parameterization introduced in Section 4.4:

$$\rho \equiv \frac{a_0}{a_0 + b_0} \quad \text{and} \quad \eta \equiv a_0 + b_0, \tag{8}$$

which are respectively the mean and strength of the priors. Hence, the mean of the current

beliefs about the quality of specialist j in period t is specified as

$$m_{ijt} = \frac{\rho\eta + y_{ijt}}{\eta + e_{ijt}}, \tag{9}$$

where $y_{ijt} \equiv \sum_{s=1}^{t-1} Y_{ijs}$ is the number of success among patients referred to j in the past, and $e_{ijt} \equiv \sum_{s=1}^{t-1} D_{ijs}$ is the number of patients referred to j, as defined previously. The expected utility for a myopic PCP is thus:

$$\mathrm{E}\left[U_{ijt}|\dots\right] = \alpha\frac{\rho\eta + y_{ijt}}{\eta + e_{ijt}} + \pi x_{ijt} + \sum_{b\in B}\kappa_b I(e_{ijt} \in b) + \underbrace{\gamma n_{j,\tau(t)} + \xi_j}_{\delta_{j,\tau(t)}} + \epsilon_{ijt}. \tag{10}$$

As discussed in Section 4.2, the analogous expected utility for a forward-looking PCP is

$$\mathrm{E}\left[U_{ijt}|\dots\right] = \alpha\tilde{g}(m_{ijt}, v_{ijt}) + \pi x_{ijt} + \sum_{b\in B}\kappa_b I(e_{ijt} \in b) + \underbrace{\gamma n_{j,\tau(t)} + \xi_j}_{\delta_{j,\tau(t)}} + \epsilon_{ijt}, \tag{11}$$

where $\tilde{g}(m_{ijt}, v_{ijt})$ is approximated using Equation (7). In both the myopic and forward-looking models, the parameters $\alpha, \pi, \kappa_1, ..., \kappa_B$, and the composite fixed effects $\delta_{j,\tau(t)}$ are estimated from the referral decisions, while ρ is set equal to the mean success rate in the market and η is given one of a set of fixed values, as discussed in Section 4.4.

5.2 Variable Definitions

For each PCP-specialist pair, the running variables e_{ijt} and y_{ijt} give the numbers of referrals and successes prior to patient and time t. However, because our sample begins in 2008, we do not observe the complete history for every pair. Experiences from many years in the past may also be less salient for current referral decisions. To address these issues and employ comparable observation periods to construct histories for all PCPs and specialists, we limit the running variables to a five-year look-back period.[24] Referrals and outcomes occurring more than five years in the past are thereby assumed not to affect the current referral decision. While this five-year period is somewhat arbitrary, some cutoff is necessary

[24]For example, consider a referral from PCP i to specialist j for an inpatient admission on January 1, 2017. Our running variables for this referral are constructed using admissions from January 1, 2012 through December 31, 2016.

in order to ensure a common length of history for all PCP-specialist pairs. We selected five years as it seems reasonably long, and it splits our available data into two similar periods: 2008-2012 is the baseline period from which the initial histories are constructed, and 2013-2018 is the estimation period. The specialist-level patient volumes, $n_{j,\tau(t)}$, which are used for the congestion effect, are defined using three-year periods (2013-2015 and 2016-2018).

We construct choice sets based on the patient's HRR. For concreteness, denote a patient i's HRR as h_i. We then take the set of specialists operating on any patient in h_i, limiting to specialists with at least 20 surgeries in a calendar year. This initial choice set might include specialists located outside of a given HRR but that nonetheless operate on patients residing in h_i. Next, we exclude any specialists whose nearest practice location is farther than the 90th percentile of distances to other specialists in that patient's choice set, or beyond 75 miles. The distance between the patient and the specialist is computed using the zip codes of the patient's residence and the specialist's nearest practice location. The resulting choice sets contain about 40 orthopedic surgeons, on average, compared to 69 specialists per HRR per year on average, so our distance restriction does little to limit the effective choice sets.

5.3 Estimation Procedure

The specifications in Equations (10) and (11) yield relatively standard multinomial logit models, with the exception of the congestion effect. Richards-Shubik et al. (2022) provides a two-step approach to estimate these models with a congestion effect, which is based on methods from Bayer and Timmins (2007) and closely related to the demand estimation from Berry, Levinsohn and Pakes (1995).

In the first step, the congestion effect, $\gamma n_{j,\tau(t)}$, is absorbed into the composite fixed effect, $\delta_{j,\tau(t)}$ in Equations (10) and (11), and the resulting multinomial logit is estimated via maximum likelihood. In the second step, γ is estimated with a linear regression defined by the identity $\delta_{j,\tau(t)} \equiv \gamma n_{j,\tau(t)} + \xi_j$. Because $n_{j,\tau(t)}$ is endogenous to ξ_j, and since ξ_j represents part of the error term in this second step regression, we estimate γ via two-stage least squares. We employ as an instrument the predicted number of patients, $\hat{n}_{j,\tau(t)}$, generated using a simple multinomial logit with only distance as the independent variable. In relation to the individual referral decision in the first step, the exclusion restriction for this instrument

comes from the distances between *all other patients* and specialist j. Specifically, $\hat{n}_{j,\tau(t)}$ is generated using the distances, x_{kjs}, from all PCPs k and patients s in the market during the relevant time period, while only the current patient's distance, x_{ijt}, affects the given referral decision.

The estimation of the multinomial logit in the first step is further simplified because ρ is set equal to the mean specialist quality in that HRR and η is fixed at predetermined values of 1 and 5.[25] The utility specification is therefore linear in the remaining parameters. Finally, we constrain α to be weakly positive so that successful patient outcomes ($Y_{ijt} = 1$) are weakly preferred by PCPs. We impose this constraint using an exponential transformation (i.e., $\alpha = \mathrm{e}^{\tilde{\alpha}}$), implemented using the algorithm proposed in Fader et al. (1992). In our setup, that algorithm amounts to an iterative procedure in which we propose an initial value, $\tilde{\alpha}_0 = \ln(\alpha_0)$, construct the variable $\tilde{m}_{ijt,0} = \mathrm{e}^{\tilde{\alpha}_0} \times m_{ijt}$, estimate the coefficient on this transformed variable, $\widehat{\Delta}_0$, and then update to $\tilde{\alpha}_1 = \tilde{\alpha}_0 + (\widehat{\Delta}_0 - 1)$. This process continues until $\tilde{\alpha}$ converges, which we denote as $\tilde{\alpha}^*$. Our final estimate for α then follows from the transformation of $\tilde{\alpha}^*$, so that $\hat{\alpha} = \mathrm{e}^{\tilde{\alpha}^*}$.

Four alternative models are estimated: assuming myopic or forward-looking behavior, and using $\eta = 1$ or $\eta = 5$. We estimate these models separately in each of the 306 HRRs in our data, both to examine variation across markets and for computational feasibility, avoiding the simultaneous estimation of many thousands of specialist FEs. The distributions of estimates across HRRs are summarized by taking a weighted average of the estimates from each HRR, weighted by the number of observations per HRR. The congestion effect, γ, is estimated separately in each HRR along with the other parameters, and it is also estimated nationwide by pooling the composite fixed effects ($\hat{\hat{\delta}}_{j,\tau(t)}$) and predicted numbers of patients ($\hat{n}_{j,\tau(t)}$) across all HRRs, with HRR fixed effects included in the pooled 2SLS specification.[26]

[25] From Table 2, PCPs refer an average of 4.4 patients per year in total, or about 4.4 cumulative patients to a given specialist during the estimation period, so these values for η span a reasonable range for the strength of the priors in relation to the information learned from actual patients.

[26] The HRR fixed effects in the second-stage regression account for the fact that specialist fixed effects in the first step are estimated separately within each HRR, with one specialist's fixed effect omitted per market as the reference category.

5.4 Parameter Estimates

Out of 306 HRRs, our constrained maximum likelihood estimation converged in over 90% of markets, with slight variation across the myopic versus forward-looking models and the values of η.[27] Parameter estimates from the myopic model are presented in Table 3, with analogous results from the forward-looking model in Table 4. The tables present the coefficient estimates of α, π, γ, and κ using different values for η, as indicated, along with the estimates of the prior mean ρ that come from the observed success rates (and therefore do not vary with η). We summarize the estimates for familiarity focusing on the intervals up to 10 patients for brevity. As noted above, α, π, and κ are estimated separately in each HRR while γ is estimated by pooling across HRRs. For α, π, and κ, we present the mean across HRRs along with certain percentiles.

In both the myopic and forward-looking models, the estimates of π are in line with other estimates of the effect of distance found in the literature, and are fairly similar across the values of η. The estimates for familiarity are also similar across η and are generally large and increasing, suggesting that familiarity plays a potentially large role in driving referrals. Finally, the estimates of γ are negative and significant, indicating the presence of congestion effects, meaning that specialists with larger patient volumes are somewhat less likely to receive a given referral. The estimates of about -0.1 (per 100 patients) imply that specialists with 100 more patients (over a 3-year period) are 7-9% less likely to receive a particular patient, depending on the baseline referral probability.[28]

The estimates of α indicate that PCPs place small to moderate weights on patient outcomes depending on the market. Larger values of η yield larger estimates of α, which is natural because α is divided by $\eta + e$ to form a composite "coefficient" on y in Equation (10), with a similar scaling effect in the forward-looking case. We estimate slightly larger values for α in the forward-looking models, consistent with the additional informational value from patient outcomes in a forward-looking versus myopic decision. Focusing on the smallest estimate of the mean value of α, in the myopic model with $\eta = 1$ (in Table 3,

[27] Specifically, in the myopic (forward-looking) case, the estimation converged for 280 (277) markets for $\eta = 1$ and 281 (279) markets for $\eta = 5$.

[28] In a multinomial logit, the relative marginal effect on a choice probability p_j is $(1 - p_j)\gamma$. The average referral probability is about 0.3 for a PCP's most-chosen specialist and is under 0.1 for any specialist with a past referral.

column "Mean"), its ratio to the corresponding mean estimate of π implies that on average PCPs view the difference between success and failure as worth about 2.9 miles of extra travel distance for the patient. Such a small marginal rate of substitution is consistent with other estimates in the literature finding small quality elasticities of demand. Also, the national mean estimate of α is reasonably precise. A simple estimate of its standard error can be computed by dividing the standard deviation across HRRs (in column "SD/SE") by the square root of the number of HRRs, which gives a naive standard error of 0.0107.

To further interpret the role of patient outcomes in PCP referrals, we present the partial effect of one additional failure in Figure 5 among PCP-specialist pairs with non-zero referral histories. The figure plots the distribution of average partial effects across HRRs. It illustrates two key findings. First, there are some markets in which PCPs essentially do not respond meaningfully to specialist failures (i.e., $\hat{\alpha} = 0$). Second, in markets where there is a response, one additional failure reduces the referral probability by 0.003, on average, which is a nontrivial relative reduction for pairs with non-zero referral histories.[29]

5.5 Counterfactuals

Using our model of PCP learning and referrals, and the estimates summarized in Tables 3 and 4, we consider two primary counterfactuals. Both sets of counterfactuals simulate the reallocation of patients and the effects on patient health under full information, in which PCPs have perfectly accurate beliefs about specialist quality.[30] The counterfactuals differ in regard to our assumptions on the role of familiarity. In one case, we initialize familiarity to 0, and in the other, we initialize familiarity as observed in the data as of the time of the first referral. We refer to the first set of counterfactuals as "resetting familiarity" and the second set as "existing familiarity."[31] In both cases, the number of patients referred to each specialist can change going forward, which naturally allows the familiarity effect

[29]Note that our reduced-form evidence in Section 3 focuses only on cases where more than one PCP refers patients to the same specialist and where at least one (but not all) of the specialist's patients experienced a failure in a given quarter. This yields a relatively unique sample of PCP-specialist pairs, such that effect sizes across the different analyses are not comparable.

[30]This is implemented by replacing m in the myopic version and g in the forward-looking version of the model with the true probability of success for each specialist, as observed in the data. For the Gittins index, note that when the variance of the beliefs, v, is zero, the expression in Equation (7) becomes just m.

[31]Under "existing familiarity", the counterfactuals therefore simulate a shock where the PCPs unexpectedly receive perfect information about specialist quality on January 1, 2013.

to change along with quality information. Familiarity therefore compounds the effect of improved quality information.[32]

The results under "resetting familiarity" are summarized in Figure 6, where we find that 39.6% (32.1%) of existing referrals would be reallocated in the myopic (forward-looking) model. This represents a substantial change in demand for many specialists. The changes in health outcomes are more modest, which is not surprising given the relatively small average failure rate. Under full information, the probability of a negative outcome falls by about 0.5 percentage points relative to the average failure rate of around 0.09, making the relative reduction in the failure rate about 5%. These results include adjustments for changes in congestion, as discussed in Section 5.3, and are relatively similar across different values of η. The results under "existing familiarity" are similar and appear in Figure 7.

Collectively, our counterfactuals indicate that informational frictions have substantial effects on the allocation of referrals. With perfect information, about one-third of patients would be referred to a different specialist. In the existing informational environment, PCPs learn slowly about quality differences among specialists—partly because opportunities for learning (e.g., negative outcomes) are somewhat rare. Familiarity also plays a role, maintaining referral patterns regardless of patient outcomes. The congestion effect among patients referred to the same specialists further reduces the ability of the market to allocate patients to better performing specialists. Consequently, in the existing environment, the referring physician's response to patient outcomes is relatively muted.

6 Discussion

This paper examines how primary care physicians (PCPs) learn about specialist quality and make referral decisions in the context of orthopedic surgery. Using 100% Medicare FFS claims data spanning 11 years, we estimate a structural learning model to understand how PCPs respond to patient outcomes, and we use counterfactual simulations to quantify the losses due to informational frictions in this market.

Employing both a design-based event study and a structural model of PCP referrals, we

[32]In the supplemental appendix, we also consider a simulation in which familiarity is shut off entirely. These results are qualitatively similar to those in the main text, but the effects of full information are more muted since the feedback loop between quality information and familiarity no longer plays a role.

find that PCPs modestly adjust their referrals based on negative patient outcomes; however, this responsiveness remains relatively limited. PCP referral patterns are also heavily influenced by familiarity with specialists, often overpowering the role of observable quality differences. The slow pace of learning and the weight of prior relationships highlight a significant inefficiency in referral markets. Additionally, capacity constraints among specialists present a further challenge to reallocating patients to higher-quality providers, even when quality information is improved. Nonetheless, counterfactual analyses indicate that substantial changes in referral allocations and meaningful gains in patient outcomes could be achievable with more complete information about quality.

Our results suggest that while PCPs demonstrate some responsiveness to negative outcomes, their overall adaptability is constrained. For example, referral probabilities change by less than 5% on average following adverse events, underscoring the inertia inherent in referral behaviors. This modest responsiveness is exacerbated by the infrequency of negative outcomes, which limits opportunities for learning. Market characteristics also appear to play a meaningful role in shaping PCP responses. In concentrated markets with fewer specialists, PCPs exhibit stronger responses to patient outcomes, likely due to more salient quality differences. Conversely, in less concentrated markets, the abundance of options may dilute attention to individual outcomes.

From a policy perspective, our results suggest a muted effect of efforts toward quality transparency, such as centralized platforms for sharing specialist performance metrics or public reporting of specialist quality. Tailored interventions in concentrated markets could amplify the impact of quality-focused policies by leveraging the stronger PCP responses observed in these settings. Moreover, given the strength of PCP-specialist familiarity in driving future referrals, information campaigns or decision-support tools may be most effective when targeted to PCPs entering a new geographic market.

References

Agha, Leila, Brigham Frandsen, and Jambes B. Rebitzer, "Causes and Consequences of Fragmented Care Delivery: Theory, Evidence, and Public Policy," Working Paper 23078, National Bureau of Economic Research April 2017.

_ , Keith Marzilli Ericson, Kimberley H. Geissler, and James B. Rebitzer, "Team Relationships and Performance: Evidence from Healthcare Referral Networks," *Management Science*, May 2022, *68* (5), 3735–3754.

Barkowski, Scott, "Physician Response to Prices of Other Physicians: Evidence from a Field Experiment," Working Paper, Clemson University 2018.

Barnett, Michael L, Bruce E Landon, A James O'malley, Nancy L Keating, and Nicholas A Christakis, "Mapping physician networks with self-reported and administrative data," *Health services research*, 2011, *46* (5), 1592–1609.

_ , Nancy L Keating, Nicholas A Christakis, A James O'Malley, and Bruce E Landon, "Reasons for choice of referral physician among primary care and specialist physicians," *Journal of general internal medicine*, 2012, *27* (5), 506–512.

Bayer, Patrick and Christopher Timmins, "Estimating Equilibrium Models of Sorting Across Locations," *The Economic Journal*, March 2007, *117* (518), 353–374.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, 1995, pp. 841–890.

Brezzi, Monica and Tze Leung Lai, "Optimal learning and experimentation in bandit problems," *Journal of Economic Dynamics and Control*, 2002, *27* (1), 87–108.

Bunting, Jackson, Paul Diegert, and Arnaud Maurel, "Heterogeneity, Uncertainty and Learning: Semiparametric Identification and Estimation," February 2024.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, August 2019, *134* (3), 1405–1454.

Chan, Tat, Chakravarthi Narasimhan, and Ying Xie, "Treatment effectiveness and side effects: A model of physician learning," *Management Science*, 2013, *59* (6), 1309–1325.

Chernew, M., G. Gowrisankaran, and D.P. Scanlon, "Learning and the value of information: Evidence from health plan report cards," *Journal of Econometrics*, 2008, *144* (1), 156–174.

Chernew, Michael, Zack Cooper, Eugene Larsen Hallock, and Fiona Scott Morton, "Physician agency, consumerism, and the consumption of lower-limb MRI scans," *Journal of Health Economics*, March 2021, *76*, 102427.

Ching, Andrew T., "Consumer learning and heterogeneity: Dynamics of demand for prescription drugs after patent expiration," *International Journal of Industrial Organization*, November 2010, *28* (6), 619–638.

_ , Tülin Erdem, and Michael P. Keane, "Learning Models: An Assessment of Progress, Challenges, and New Developments," *Marketing Science*, November 2013, *32* (6), 913–938.

_ , _ , and _ , "Empirical Models of Learning Dynamics: A Survey of Recent Developments," in Berend Wierenga and Ralf van der Lans, eds., *Handbook of Marketing Decision Models*, International Series in Operations Research & Management Science, Cham: Springer International Publishing, 2017, pp. 223–257.

Coscelli, Andrea and Matthew Shum, "An empirical model of learning and patient spillovers in new drug entry," *Journal of Econometrics*, October 2004, *122* (2), 213–246.

Crawford, Gregory S and Matthew Shum, "Uncertainty and learning in pharmaceutical demand," *Econometrica*, 2005, *73* (4), 1137–1173.

Dickstein, Michael, "Efficient provision of experience goods: Evidence from antidepressant choice," Working Paper 18-17, New York University, Leonard N. Stern School of Business 2018.

Dranove, David and Christopher Ody, "Employed for higher pay? How Medicare facility fees affect hospital employment of physicians," *American Economics Journal: Economic Policy*, 2019, *11* (4), 249–271.

Fader, Peter S., James M. Lattin, and John D. C. Little, "Estimating Nonlinear Parameters in the Multinomial Logit Model," *Marketing Science*, November 1992, *11* (4), 372–385.

Ferreyra, Maria Marta and Grigory Kosenok, "Learning About New Products: An Empirical Study of Physicians' Behavior," *Economic Inquiry*, 2011, *49* (3), 876–898.

Freedman, Rachel A, Elena M Kouri, Dee W West, and Nancy L Keating, "Racial/ethnic differences in patients' selection of surgeons and hospitals for breast cancer surgery," *JAMA oncology*, 2015, *1* (2), 222–230.

Gaynor, Martin, Carol Propper, and Stephan Seiler, "Free to Choose? Reform, Choice, and Consideration Sets in the English National Health Service," *American Economic Review*, November 2016, *106* (11), 3521–57.

Gittins, J. C. and D. M. Jones, "A dynamic allocation index for the discounted multiarmed bandit problem," *Biometrika*, December 1979, *66* (3), 561–565.

Gittins, John C, "Bandit processes and dynamic allocation indices," *Journal of the Royal Statistical Society: Series B (Methodological)*, 1979, *41* (2), 148–164.

Gong, Qing, "Physician Learning and Treatment Choices Evidence from Brain Aneurysms," Working Paper, University of North Carolina at Chapel Hill 2018.

Kaur, Ramandeep, Jennifer N. Perloff, Christopher Tompkins, and Christine E. Bishop, "Hospital Postacute Care Referral Networks: Is Referral Concentration Associated with Medicare-Style Bundled Payments?," *Health Services Research*, 2017, *52* (6), 2079–2098.

Kolstad, Jonathan T. and Michael E. Chernew, "Quality and Consumer Decision Making in the Market for Health Insurance and Health Care Services," *Medical Care Research and Review*, February 2009, *66* (1_suppl), 28S–52S.

Pham, Hoangmai H, Ann S O'Malley, Peter B Bach, Cynthia Saiontz-Martinez, and Deborah Schrag, "Primary care physicians' links to other physicians through Medi-

care patients: the scope of care coordination," *Annals of internal medicine*, 2009, *150* (4), 236–242.

Richards-Shubik, Seth, Mark S. Roberts, and Julie M. Donohue, "Measuring quality effects in equilibrium," *Journal of Health Economics*, May 2022, *83*, 102616.

Sarsons, Heather, "Interpreting Signals in the Labor Market: Evidence from Medical Referrals," Working Paper, Harvard University 2023.

Schneider, Eric C. and Arnold M. Epstein, "Influence of Cardiac-Surgery Performance Reports on Referral Practices and Access to Care — A Survey of Cardiovascular Specialists," *New England Journal of Medicine*, July 1996, *335* (4), 251–256. _eprint: https://www.nejm.org/doi/pdf/10.1056/NEJM199607253350406.

Zeltzer, Dan, "Gender Homophily in Referral Networks: Consequences for the Medicare Physician Earnings Gap," *American Economic Journal: Applied Economics*, April 2020, *12* (2), 169–97.

Tables and Figures

Table 1: Descriptive Statistics for Patients and Referrals[a]

	2008-2012	2013-2018	Overall
Patient Characteristics			
Age	74.668	74.138	74.360
	(6.351)	(6.214)	(6.277)
White	0.941	0.930	0.934
	(0.236)	(0.256)	(0.247)
Black	0.038	0.038	0.038
	(0.192)	(0.192)	(0.192)
Male	0.351	0.372	0.364
	(0.477)	(0.483)	(0.481)
PCP Characteristics			
Practice Size	12.633	18.462	16.017
	(20.859)	(27.553)	(25.130)
Male	0.805	0.771	0.786
	(0.396)	(0.420)	(0.410)
Specialist Characteristics			
Practice Size	6.741	9.815	8.525
	(8.290)	(12.614)	(11.113)
Male	0.988	0.986	0.987
	(0.107)	(0.117)	(0.113)
Distance			
to PCP	30.663	28.215	29.237
	(159.829)	(158.802)	(159.236)
to Hospital	39.791	39.915	39.865
	(167.214)	(166.554)	(166.823)
90-day Readmission	0.084	0.068	0.075
	(0.277)	(0.252)	(0.263)
90-day Mortality	0.008	0.005	0.006
	(0.087)	(0.073)	(0.079)
90-day Complication	0.112	0.093	0.101
	(0.315)	(0.291)	(0.301)
Failure	0.114	0.095	0.103
	(0.318)	(0.293)	(0.304)
Observations	939,849	1,301,005	2,240,854

[a]Mean values with standard deviations in parenthesis. As discussed in the main text, the period from 2008-2012 is our baseline period used to form histories of PCP/specialist pairs over a common time period, and the period from 2013-2018 is our estimation period. We present summary statistics separately by those time periods for consistency with our final estimation.

Table 2: Descriptive Statistics for PCPs and PCP/Specialist Pairs[a]

	2008-2012	2013-2018	Overall
Total Referrals	4.159	4.740	4.477
	(3.089)	(3.611)	(3.398)
Network Size	3.002	3.435	3.239
	(1.756)	(1.990)	(1.900)
Total Failures	0.476	0.450	0.462
	(0.747)	(0.727)	(0.736)
Referrals per Specialist	1.385	1.380	1.382
	(0.953)	(0.975)	(0.966)
Failure Rate	0.118	0.098	0.106
	(0.302)	(0.279)	(0.289)
Running Referrals		4.101	
		9.996	
Running Failure Rate		0.102	
		0.206	
Observations	678,535	942,812	1,621,347

[a]Mean values calculated per year, with standard deviations in parenthesis. As discussed in the main text, the period from 2008-2012 is our baseline period used to form histories of PCP/specialist pairs over a common time period, and the period from 2013-2018 is our estimation period. We present summary statistics separately by those time periods for consistency with our final estimation.

Table 3: Structural Parameter Estimates, Myopic Model[a]

			Percentile				
Parameter	Mean	(SD/SE)	10th	25th	50th	75th	90th
α (utility weight on outcome)							
($\eta = 1$)	0.2020	(0.1791)	0.0000	0.0000	0.1166	0.2830	0.5372
($\eta = 5$)	0.3972	(0.3723)	0.0000	0.0000	0.2121	0.5605	1.0891
π (utility weight on distance)							
($\eta = 1$)	-0.0694	(0.0060)	-0.0990	-0.0824	-0.0678	-0.0525	-0.0384
($\eta = 5$)	-0.0695	(0.0061)	-0.1002	-0.0826	-0.0679	-0.0528	-0.0390
ρ (prior mean)							
(all η)	0.8996	(0.0117)	0.8853	0.8920	0.8995	0.9068	0.9139
γ (congestion effect, per 100 patients)							
($\eta = 1$)	-0.0913	(0.0153)					
($\eta = 5$)	-0.1158	(0.0169)					

κ_b (familiarity)	Range of e_{ijt}						
	1	2	3	4	5	[6,7]	[8,10]
($\eta = 1$)	1.1075	1.4307	1.6991	1.8802	2.0294	2.2027	2.3996
	(0.0982)	(0.1141)	(0.1278)	(0.1418)	(0.1574)	(0.1497)	(0.1712)
($\eta = 5$)	1.1141	1.4372	1.7081	1.8892	2.0377	2.2116	2.4058
	(0.0981)	(0.1141)	(0.1278)	(0.1419)	(0.1576)	(0.1500)	(0.1713)

[a]Estimates of specification in Equation (10), using the procedure described in Section 5.3, with different values of η as indicated. Parameters α, π, ρ are estimated separately by HRR; the table reports their distribution, where SD is standard deviation across HRRs. For γ, the "mean" is the point estimate and SE is the standard error from a single nationwide estimation. All familiarity coefficients are presented as the mean across HRRs with standard errors in parentheses. Results limited to HRRs in which our constrained maximization routine converged, 280 HRRs for $\eta = 1$ and 281 for $\eta = 5$.

Table 4: Structural Parameter Estimates, Forward-looking Model[a]

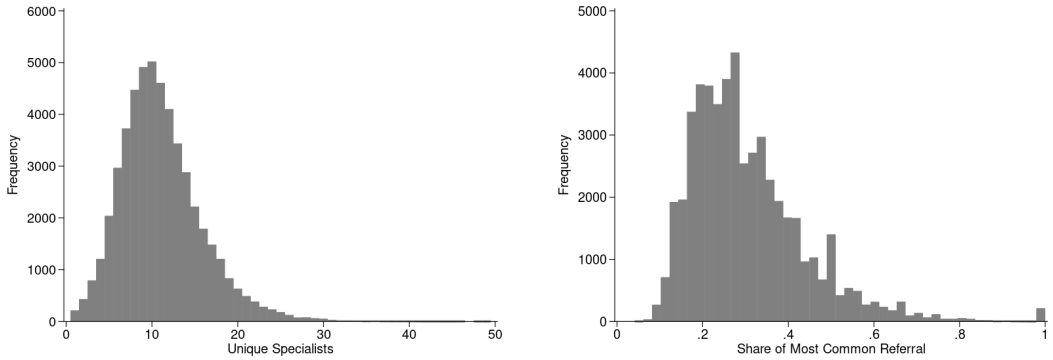
			Percentile				
Parameter	Mean	(SD/SE)	10th	25th	50th	75th	90th
α (utility weight on outcome)							
($\eta = 1$)	0.2631	(0.2308)	0.0000	0.0000	0.1607	0.3652	0.7071
($\eta = 5$)	0.5058	(0.4937)	0.0000	0.0000	0.2524	0.7287	1.3829
π (utility weight on outcome)							
($\eta = 1$)	-0.0692	(0.0060)	-0.0985	-0.0819	-0.0678	-0.0523	-0.0390
($\eta = 5$)	-0.0698	(0.0061)	-0.1002	-0.0830	-0.0684	-0.0528	-0.0379
ρ (prior mean)							
(all η)	0.8996	(0.0117)	0.8853	0.8920	0.8995	0.9068	0.9139
γ (congestion effect, per 100 patients)							
($\eta = 1$)	[b]						
($\eta = 5$)	[b]						

κ_b (familiarity)	Range of e_{ijt}						
	1	2	3	4	5	[6,7]	[8,10]
($\eta = 1$)	1.1559	1.4855	1.7583	1.9407	2.0900	2.2651	2.4631
	(0.1147)	(0.1329)	(0.1464)	(0.1604)	(0.1747)	(0.1684)	(0.1890)
($\eta = 5$)	1.2020	1.5406	1.8212	2.0066	2.1607	2.3346	2.5339
	(0.1451)	(0.1701)	(0.1863)	(0.2011)	(0.2152)	(0.2114)	(0.2318)

[a] Estimates of specification in Equation (11), using the procedure described in Section 5.3, with different values of η as indicated. Parameters α, π, ρ are estimated separately by HRR; the table reports their distribution, where SD is standard deviation across HRRs. For γ, the "mean" is the point estimate and SE is the standard error from a single nationwide estimation. All familiarity coefficients are presented as the mean across HRRs with standard errors in parentheses. Results limited to HRRs in which our constrained maximization routine converged, 277 HRRs for $\eta = 1$ and 279 for $\eta = 5$.

[b] Estimation currently underway. Values have been similar to the myopic model in past runs.

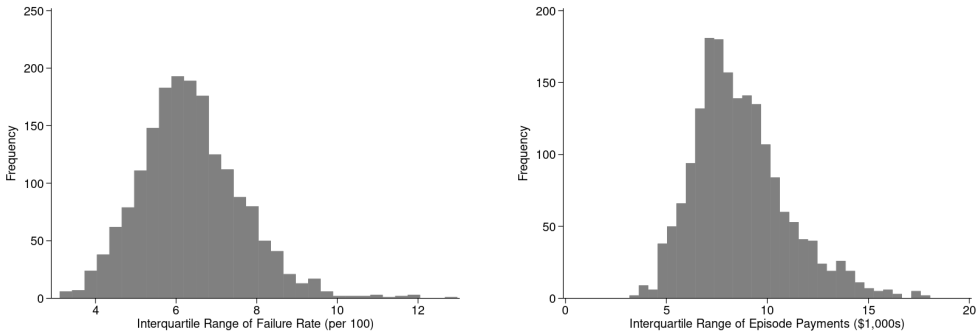
Figure 1: PCP Referral Networks over 2013-2018[a]



(a) Network Size (b) Highest Referral Share

[a]Panel (a) depicts the frequency histogram of PCP network size (number of unique specialists referred to) over the 2013-2018 estimation period. Panel (b) depicts the frequency histogram of the "highest-share" specialist per PCP over the same period.

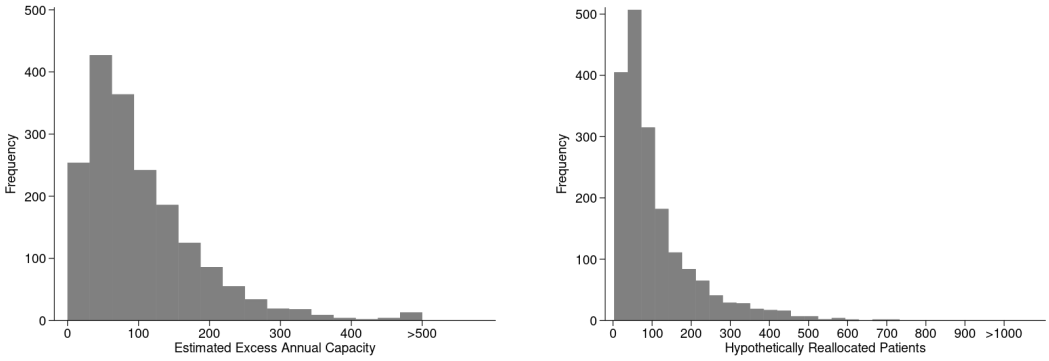
Figure 2: Within-Market Variation in Specialist Quality and Spending[a]



(a) Failure Rate (b) Episode Spending

[a]Panel (a) depicts the interquartile range of failure rates per 100 surgeries across specialists within each HRR over the 2013-2018 estimation period. Panel (b) depicts the interquartile range of mean episode spending per specialist per HRR over the same period. Episode spending is measured as the total Medicare payments over a 90-day period beginning from the date of admission.

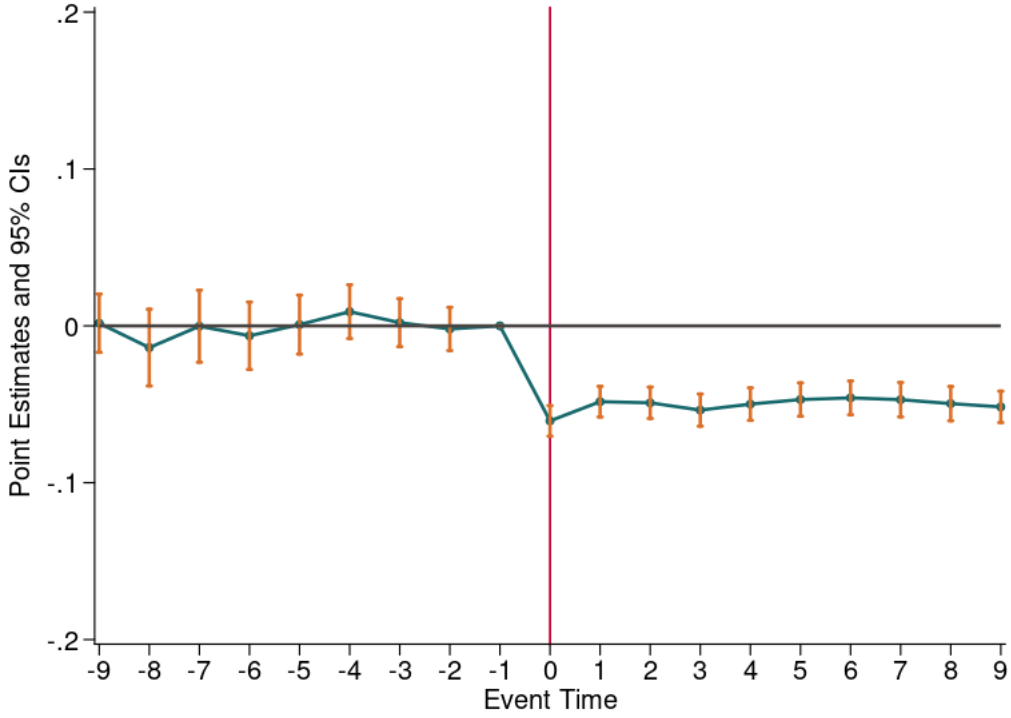
Figure 3: Specialist Capacity and Hypothetical Reallocation[a]



(a) Potential Excess Capacity (b) Hypothetical Reallocation

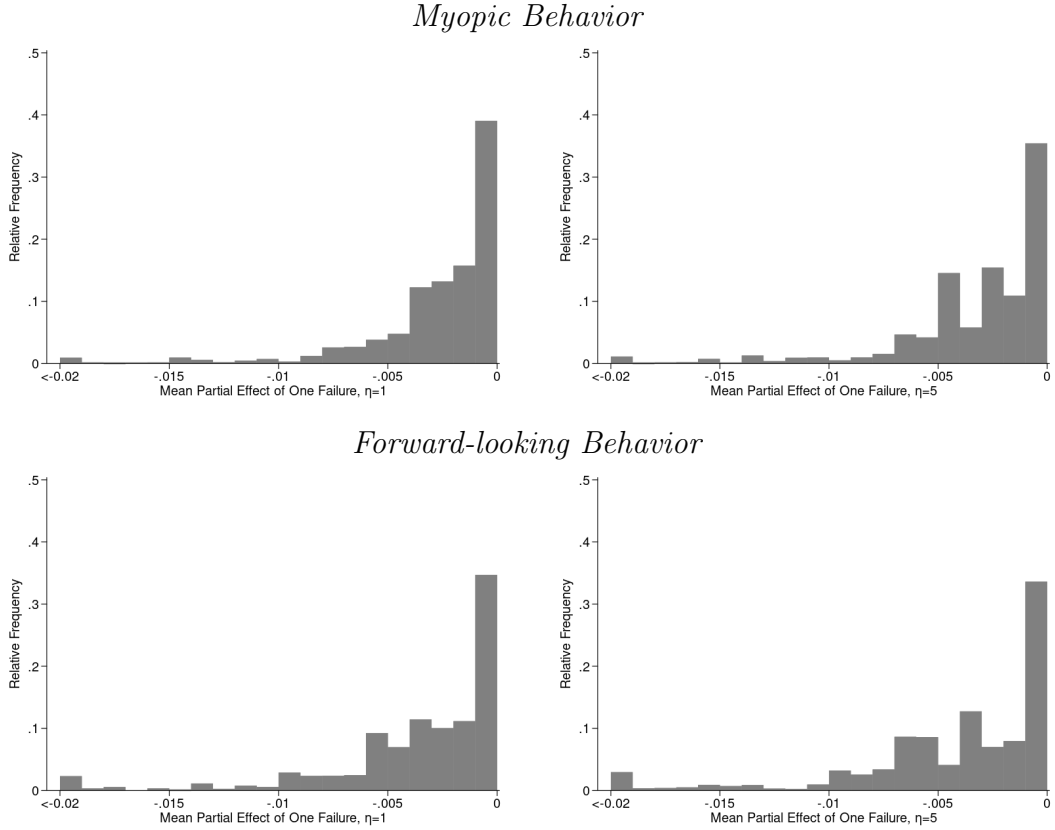
[a]Panel (a) depicts the estimated number of additional annual procedures that could be performed collectively among the specialists in the top 25 percent of the quality distribution. Excess capacity for a single specialist is measured as the difference between the 75th percentile of that specialist's count of yearly operations less the observed count of operations in that year. Panel (b) depicts the count of procedures among the lowest 25th percentile of specialists, aggregated across all specialists in a given HRR.

Figure 4: Event Study of Specialist Failures[a]



[a]Figure depicts stacked event study estimates following Cengiz et al. (2019), based on Equation (1). Our stacked specification includes specialist X quarter fixed effects as well as specialist X cohort X PCP 'type' fixed effects. The event study coefficient at time $\tau = -1$ is normalized to 0.

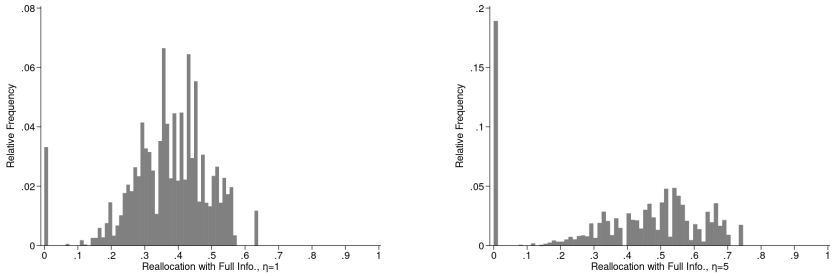
Figure 5: Partial Effect of Failures on Referral Probabilities[a]



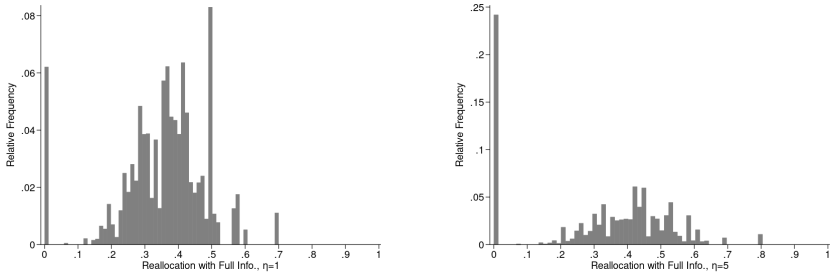
[a]Figure depicts the mean change in referral probabilities among PCP-specialist pairs with some referral history, averaged within an HRR and presented as a histogram across HRRs. Results limited to HRRs in which our constrained maximization routine converged, 280 (277) HRRs in the myopic (forward-looking) case for $\eta = 1$ and 281 (279) in the myopic (forward-looking) case for $\eta = 5$.

Figure 6: Counterfactual Changes with Full Information, Resetting Familiarity[a]

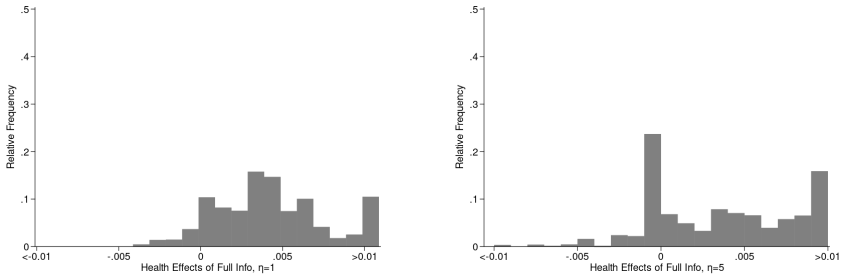
Reallocation of referrals under myopic behavior



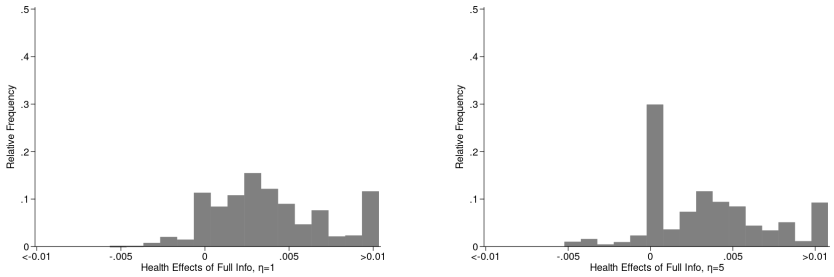
Reallocation of referrals under forward-looking behavior



Change in probability of success under myopic behavior



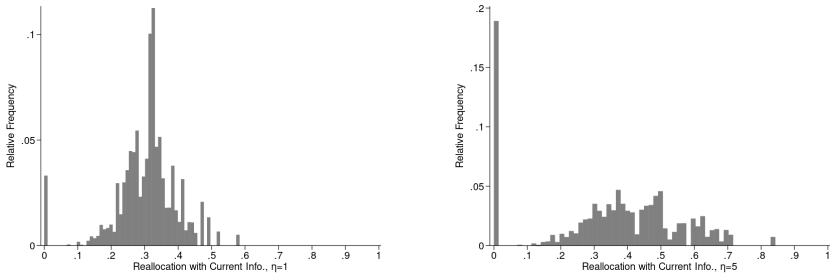
Change in probability of success under forward-looking behavior



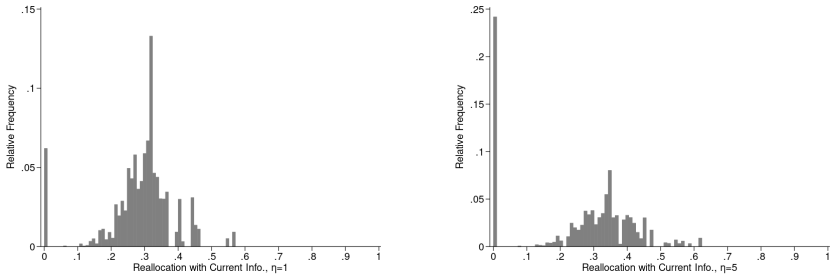
[a]Figure depicts the share of referrals reallocated and the resulting effects on patient health (i.e., the change in the probability of a successful surgery) under a counterfactual in which PCPs have full information about specialist quality and in which familiarity is initialized to 0. Estimates are averaged across individuals within an HRR and presented as a histogram across HRRs. Results limited to HRRs in which our constrained maximization routine converged, 280 (277) HRRs in the myopic (forward-looking) case for $\eta = 1$ and 281 (279) in the myopic (forward-looking) case for $\eta = 5$.

Figure 7: Counterfactual Changes with Full Information, Existing Familiarity[a]

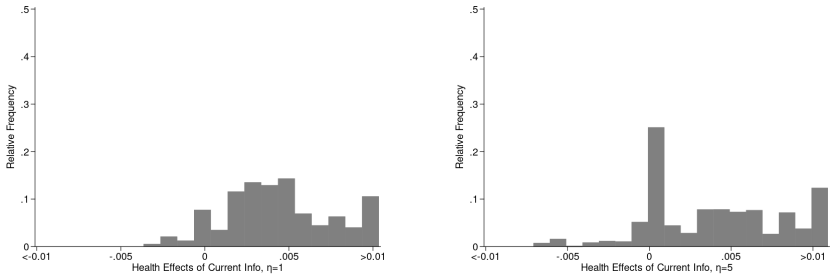
Reallocation of referrals under myopic behavior



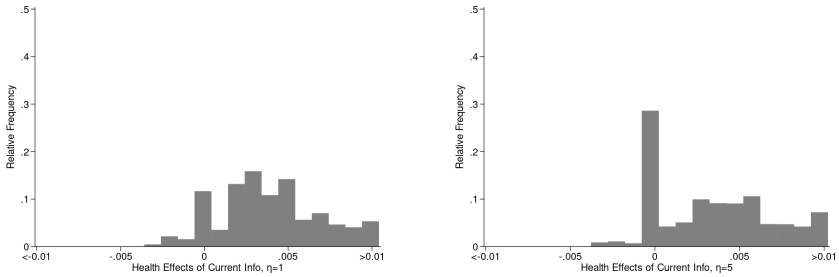
Reallocation of referrals under forward-looking behavior



Change in probability of success under myopic behavior



Change in probability of success under forward-looking behavior



[a]Figure depicts the share of referrals reallocated and the resulting effects on patient health (i.e., the change in the probability of a successful surgery) under a counterfactual in which PCPs have full information about specialist quality and in which familiarity is initialized to observed values as of January 2013. Estimates are averaged across individuals within an HRR and presented as a histogram across HRRs. Results limited to HRRs in which our constrained maximization routine converged, 280 (277) HRRs in the myopic (forward-looking) case for $\eta = 1$ and 281 (279) in the myopic (forward-looking) case for $\eta = 5$.