

Learning and Efficiency in the Market for Physician Referrals

Ian McCarthy

Emory University & NBER

Seth Richards-Shubik

Johns Hopkins University & NBER

AHEW 2026


Experts who guide buyers must learn about quality themselves

- Buyers of complex services often rely on an expert intermediary
 - Financial advisors, staffing agencies, real estate agents
- The intermediary is imperfectly informed too, and learns about quality from the outcomes it observes
- PCP referrals are a leading case
 - A PCP repeatedly sends patients to surgeons whose quality is hard to observe
 - Quality is revealed only gradually, through the outcomes of the PCP's own patients

How well do PCPs learn about specialist quality, and how much does it matter for where patients go and how they fare?


Surgeon quality varies within markets, and PCPs use only a few surgeons

- Medicare FFS joint replacements, 2008–2018: about 2 million PCP referrals
- The mean failure rate is 9 per 100, and the 25th and 75th percentile surgeons differ by 4 to 6.4 per 100




Surgeon quality varies within markets, and PCPs use only a few surgeons

- Medicare FFS joint replacements, 2008–2018: about 2 million PCP referrals
- The mean failure rate is 9 per 100, and the 25th and 75th percentile surgeons differ by 4 to 6.4 per 100
- A PCP uses about 3 of the roughly 69 surgeons in the market each year, and many send a large share to one surgeon




Surgeon quality varies within markets, and PCPs use only a few surgeons

- Medicare FFS joint replacements, 2008–2018: about 2 million PCP referrals
- The mean failure rate is 9 per 100, and the 25th and 75th percentile surgeons differ by 4 to 6.4 per 100
- A PCP uses about 3 of the roughly 69 surgeons in the market each year, and many send a large share to one surgeon
- In 39% of market-years, top-quartile surgeons appear to have room for the bottom quartile's patients



PCPs refer less to a surgeon after their own patient has a bad outcome

- Compare PCPs who refer to the same surgeon, when only one of them has a patient with a failure
- Specialist-by-quarter fixed effects absorb anything that changes at the surgeon
- The PCP with the failure sends about 0.056 fewer patients per quarter to that surgeon, about 16% fewer than the controls, and the drop persists



Stacked event study, quarters relative to the failure

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \alpha m_{ijt} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \boxed{\alpha m_{ijt}} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

belief about quality, with weight α

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \alpha m_{ijt} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

Diagram illustrating the components of the utility function U_{ijt} :

- αm_{ijt} : belief about quality, with weight α
- πx_{ijt} : distance to the surgeon

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \alpha m_{ijt} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

Diagram illustrating the components of the utility function U_{ijt} :

- αm_{ijt} : belief about quality, with weight α
- πx_{ijt} : distance to the surgeon
- $f(e_{ijt})$: familiarity: past referrals to j

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \alpha m_{ijt} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

The diagram illustrates the components of the utility function U_{ijt} . Each term in the equation is linked to a descriptive box below it:

- αm_{ijt} is linked to "belief about quality, with weight α " (blue box).
- πx_{ijt} is linked to "distance to the surgeon" (yellow box).
- $f(e_{ijt})$ is linked to "familiarity: past referrals to j " (yellow box).
- γn_{jt} is linked to "congestion: the surgeon's other patients" (blue box).

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \alpha m_{ijt} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

The diagram illustrates the components of the utility function U_{ijt} . Each term in the equation is linked to a descriptive box below it:

- αm_{ijt} : belief about quality, with weight α
- πx_{ijt} : distance to the surgeon
- $f(e_{ijt})$: familiarity: past referrals to j
- γn_{jt} : congestion: the surgeon's other patients
- ξ_j : surgeon fixed effect
- ϵ_{ijt} : error term

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \alpha m_{ijt} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

belief about quality, with weight α

distance to the surgeon

familiarity: past referrals to j

congestion: the surgeon's other patients

surgeon fixed effect

Learning

Beliefs start from a common beta prior and update with each patient's outcome

$$m_{ijt} = \frac{\rho\eta + \text{past successes}}{\eta + \text{past referrals}}$$

ρ = prior mean, η = prior strength

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \alpha m_{ijt} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

belief about quality, with weight α

distance to the surgeon

familiarity: past referrals to j

congestion: the surgeon's other patients

surgeon fixed effect

Learning

Beliefs start from a common beta prior and update with each patient's outcome

$$m_{ijt} = \frac{\rho\eta + \text{past successes}}{\eta + \text{past referrals}}$$

ρ = prior mean, η = prior strength

Behavior

Myopic PCPs refer on current beliefs

Forward-looking PCPs also value what a referral teaches them, through a Gittins index

PCPs learn about surgeon quality from their own patients' outcomes

$$U_{ijt} = \alpha m_{ijt} + \pi x_{ijt} + f(e_{ijt}) + \gamma n_{jt} + \xi_j + \epsilon_{ijt}$$

belief about quality, with weight α

distance to the surgeon

familiarity: past referrals to j

congestion: the surgeon's other patients

surgeon fixed effect

Learning

Beliefs start from a common beta prior and update with each patient's outcome

$$m_{ijt} = \frac{\rho\eta + \text{past successes}}{\eta + \text{past referrals}}$$

ρ = prior mean, η = prior strength

Behavior

Myopic PCPs refer on current beliefs

Forward-looking PCPs also value what a referral teaches them, through a Gittins index

Estimation

Logit in each of 306 HRRs

Congestion by 2SLS, distance-based instrument

ρ = market success rate, $\eta = 1$ or 5


What we find

PCPs weigh quality lightly relative to other factors

0 to 2 pp

drop in referral probability after one more failure, across markets. Where PCPs respond, the average is 0.3 pp.

Roughly equivalent to 2.9 miles of additional travel (surgeons average 29 miles from the PCP).



What we find


0 to 2 pp PCPs weigh quality lightly relative to other factors

Full information

23%

of patients go to a different surgeon when
PCPs know true quality

Failures fall by 0.07 per 100, with gains in
most markets



What we find

0 to 2 pp PCPs weigh quality lightly relative to other factors


23% reallocated under full information

Quality weighted like distance

20%

of failures avoided when one SD of quality counts as much as one SD of distance

150 to 200 more successes per 10,000,
still subject to distance and surgeon
capacity



What we find

0 to 2 pp

PCPs weigh quality lightly relative to other factors

23%

reallocated under full information

20%

of failures avoided when quality counts like distance

Takeaway

- The gains from quality transparency alone are bounded by the small weight PCPs place on quality relative to other referral factors
- Tools that raise that weight could deliver substantially larger health improvements

Discussant comments

Comment 1: Identification argument

Comment

“I disagree with the authors’ characterization that Ching, Erdem, and Keane (2013) offer only a ‘verbal’ identification argument while their own presentation is formal.”

Comment 1: Identification argument

Comment. “...Ching, Erdem, and Keane (2013) offer only a ‘verbal’ identification argument...”

- Agreed that the belief variance going to zero with experience is a straightforward limit
- Our argument concerns the marginal effect of the belief on the choice, which separates α from familiarity and from the prior strength η
- Identification is usually argued verbally (e.g., Crawford and Shum 2005, Dickstein 2018, Gong 2018). Formalizing it for a given model has some value, though perhaps not this much emphasis
- Separately, a model without α fixes the rate of substitution between the outcome and the logit error, which matters with binary outcomes (e.g., Gong 2018)

Comment 2: Risk adjustment

Comment

“[T]his dispersion could also reflect unobserved variation in patient severity (e.g., sicker patients selecting higher-quality surgeons).”

“The remaining $\sim 14\%$ of non-DRG 470 cases may account for a disproportionate share of the observed dispersion in failure rates across markets.”

Comment 2: Risk adjustment

Comment. "...this dispersion could also reflect unobserved variation in patient severity..."

- Planned, elective, referred surgeries: 86% are DRG 470, the rest spread over 17 codes
- Risk-adjusted quality uses age, sex, race, DRG, and the Charlson index, and correlates 0.84 with raw quality
- The model needs a binary outcome, so it keeps raw failures. We re-estimate the referral logit with each alternative

	Mean $\hat{\alpha}$
Baseline	0.192
DRG 470 only	0.216
Risk-adjusted quality	0.192

Comment 3: Risk aversion

Comment

“The authors assume risk neutrality in the experience part, while incorporating a term to capture familiarity. But the estimated familiarity effect may just reflect risk aversion.”

Comment 3: Risk aversion

Comment. "...the estimated familiarity effect may just reflect risk aversion."

Surgeon	Past outcomes	Success rate	Outcome variance
A	1, 1, 0, 0	0.5	0.25
B	1, 1, 1, 1, 0, 0, 0, 0	0.5	0.25

- Same success probability, same outcome variance, so risk aversion over the outcome cannot separate A and B. α already is the utility gap between success and failure
- Uncertainty in beliefs about q_j does differ, larger for A with fewer patients
- It depends only on the success rate and the count, so it would be identified only by functional form. As is, aversion to it likely loads onto familiarity

Comment 4: Counterfactuals

Comment

“If the familiarity term simply represents the posterior variance & risk aversion, a counterfactual expt of setting everyone to have perfect info should also set the posterior variance to zero.”

“That would set the familiarity term to the limit.”

Comment 4: Counterfactuals

Comment. "...should also set the posterior variance to zero."

- If familiarity reflects uncertainty about quality, full information should remove it as well
- With familiarity the same for every surgeon, it no longer affects the choice. The appendix reports this case: 35% of referrals reallocated, and success rates 0.27 percentage points lower on average
- Without familiarity, the reallocation follows distance and the specialist fixed effects more than quality

Comment 5: Event study vs. structural model

Comment

“The event study finds a 16% relative reduction in referrals following a bad outcome; the structural model estimates a mean partial effect of one additional failure of -0.32 percentage points among ‘responsive’ HRRs.”

“[T]he sample used in estimating the structural model could have many physicians who hardly refer patients to a specialist and with very low baseline referral probability (close to zero), creating a ‘floor’ that dampens the negative outcome effect on choice behavior.”

Comment 5: Event study vs. structural model

Comment. “The event study finds a 16% relative reduction in referrals following a bad outcome...”

- The two estimates are not directly comparable
- Part of the gap may be compositional: the event study uses PCPs who share a surgeon and refer to it often
- Part may come from how the event study is built: the effect is measured against other PCPs referring to the same surgeon in the same quarter, a different comparison from the change in referral probability in the structural model

Topics

SETTING AND EVIDENCE

- 1 Why joint replacement**
elective, predictable pathway; other settings
- 2 Assigning the PCP**
E&M visits vs. referring-physician field
- 3 Patients follow PCPs**
patient choice, FFS, no networks
- 4 Quality measure**
mortality, readmissions, complications
- 5 Risk adjustment**
case mix, patient selection
- 6 Quality and spending variation**
within-market gaps in quality and spending
- 7 Excess capacity**
room at top surgeons
- 8 Referral networks**
network size and concentration
- 9 Event study**
reduced-form response to a failure

MODEL AND RESULTS

- 10 Beta-binomial model**
binary outcomes, vertical quality
- 11 Learning vs. familiarity**
habit, separating the two
- 12 Priors (ρ, η)**
prior mean and strength
- 13 Practice and peer learning**
integration, colleagues' outcomes
- 14 Specialist fixed effects**
reputation, amenities, unmeasured quality
- 15 Congestion**
capacity, wait times, instrument
- 16 Myopic vs. forward-looking**
experimentation, Gittins index
- 17 Index vs. full solution**
why a bandit does not need a dynamic program
- 18 Size of α**
weight on quality, variation across markets
- 19 Full information**
report cards, transparency
- 20 Quality-weighted referrals**
decision support, EHR tools

- Planned and elective: a predictable care pathway where patients are likely to rely on a PCP referral
- High volume: about 4.5 million inpatient joint replacements in Medicare FFS, 2008–2018
- Performed almost entirely as inpatient procedures during 2008–2018
- Well-defined outcomes: CMS's Comprehensive Care for Joint Replacement model uses the same failure measures
- A setting where referrals can be reasonably well inferred from claims

The model applies wherever an intermediary routes clients to providers whose quality it learns from outcomes.

- The referring PCP is the PCP with the most E&M visits in the 12 months before surgery, with at least 2 visits (following Pham et al. 2009; Agha et al. 2017)
- This identifies a PCP for about 80% of procedures
- The referring-physician field on the surgeon's claim is less useful
 - Blank in about 32% of cases
 - When filled, it most often lists the surgeon (about 43%) or another specialist (about 7%)
 - Only about 18% of cases list a PCP
- Where both methods identify a PCP, they agree in more than 82% of cases in every year

- The PCP's recommendation is the most commonly cited reason patients give for their choice of specialist (Freedman et al. 2015; Barkowski 2018)
- PCP referrals strongly influence which specialists patients see (Gaynor, Proper and Seiler 2016; Chernew et al. 2021)
- Traditional Medicare has no network restrictions
 - A referral is a recommendation, not an administrative requirement
 - Any surgeon in the patient's market is available
- In the model, the PCP values the patient's outcome and other patient factors such as distance

- Follows the CMS Comprehensive Care for Joint Replacement model and the National Quality Forum
- A failure is death or a readmission within 90 days of discharge, or a complication
- Surgeon quality is the probability of success, the same for every patient

2013–2018	Rate
90-day readmission	0.066
90-day mortality	0.005
Complication	0.089
Any failure	0.091

Mean $\hat{\alpha}$

0.192

- Baseline

Mean $\hat{\alpha}$

- Baseline
- Restrict to DRG 470
 - Joint replacement without major complications, 86% of cases

0.192

0.216

Mean $\hat{\alpha}$

- Baseline **0.192**
- Restrict to DRG 470 **0.216**
 - Joint replacement without major complications, 86% of cases
- Use risk-adjusted surgeon quality **0.192**
 - Observed minus expected failures, with expected failures from age, sex, race, DRG, and the Charlson index
 - Correlation with raw quality is 0.84


Mean $\hat{\alpha}$

- Baseline **0.192**
- Restrict to DRG 470 **0.216**
 - Joint replacement without major complications, 86% of cases
- Use risk-adjusted surgeon quality **0.192**
 - Observed minus expected failures, with expected failures from age, sex, race, DRG, and the Charlson index
 - Correlation with raw quality is 0.84


Takeaway

The estimated weight on quality is stable.

Quality and spending vary widely within markets




Failure rates: 4 to 6.4 per 100 between the 25th and 75th percentile surgeon




90-day episode spending: nearly \$8,000 between the 25th and 75th percentile surgeon

Top surgeons appear to have room for more patients





Top-quartile surgeons



Bottom-quartile surgeons

In 39% of market-years, top-quartile capacity covers the bottom-quartile volume. The model handles capacity through the congestion term.

PCP referral networks are small and concentrated




Each year a PCP refers about 4.4 patients to 3.1 surgeons, out of about 69 surgeons in the market.

- Quarterly panel of PCP–surgeon pairs with any referral
- For each surgeon's first failure, the treated PCP is the one whose patient had the failure; control PCPs refer to the same surgeon and have no failure in the window
- The failed referral itself is removed from the outcome
- Repeat for the second, third, and fourth failures and stack the cohorts (Cengiz et al. 2019)

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

$\gamma_{j,t}$: surgeon-by-quarter fixed effects, so changes at the surgeon or in other PCPs' demand drop out. Standard errors clustered by surgeon and cohort.

- About 0.056 fewer referrals per quarter from the treated PCP to that surgeon, about 16% relative to the controls
- Flat pre-trends, and no recovery within nine quarters
- Each failure cohort shows the same pattern, somewhat larger for the first and second failures



- Each surgery succeeds or fails; surgeon j succeeds with probability q_j , the same for all patients (vertical quality)
- PCPs start from a common prior, $\text{Beta}(a_0, b_0)$, and update with their own patients' outcomes

$$a_{ijt} = a_0 + \sum_{s < t} Y_{ijs}, \quad b_{ijt} = b_0 + \sum_{s < t} (D_{ijs} - Y_{ijs}), \quad m_{ijt} = \frac{a_{ijt}}{a_{ijt} + b_{ijt}}$$

- The same class of model as Gong (2018) and Dickstein (2018), with outcomes observed and about 57 surgeons in the typical choice set
- Crawford and Shum (2005) and Dickstein (2018) instead model horizontal match quality between drugs and patients

The belief term splits into the prior and the past success rate, $\bar{y}$:

$$\alpha m_{ijt} = \alpha \frac{a_0}{\eta + e_{ijt}} + \alpha \frac{e_{ijt}}{\eta + e_{ijt}} \bar{y}_{ijt}$$

- α : a PCP with many past patients at two surgeons but different success rates. As $e \rightarrow \infty$ the belief term becomes $\alpha \bar{y}$
- η : how fast the effect of the success rate grows with the number of past patients
- Familiarity $f(e)$: whatever effect of past referrals remains once the belief term is accounted for, with $f(0) = 0$ and f bounded. It cannot absorb the response to outcomes, so its interpretation does not affect $\hat{\alpha}$
- The event study isolates this variation: same surgeon, different outcomes

Past patients	1	2	3	4	5	6–7	8–10
κ_b , myopic ($\eta = 1$)	1.11	1.43	1.70	1.88	2.03	2.20	2.40
κ_b , forward-looking ($\eta = 1$)	1.16	1.49	1.76	1.94	2.09	2.27	2.46

- Familiarity enters as indicators for bins of past referrals, relative to none
- Estimates are nearly identical for $\eta = 5$

- Prior mean $\rho = a_0 / (a_0 + b_0)$ equals the market's average success rate (about 0.90)
 - Assumes PCPs have a general awareness of the local success rate (rational expectations), as in Coscelli and Shum (2004), Crawford and Shum (2005), and Dickstein (2018)
- Prior strength $\eta = a_0 + b_0$ acts like a number of patients seen before the first real one
 - The likelihood is flat in η in simulations, so it is fixed at 1 or 5. The cross-partial of e and $\bar{y}$ that identifies η moves only from 0.033 to 0.056 as η goes from 1 to 20
 - PCPs refer about 4.4 patients a year, so these values bracket a sensible range
 - Other estimates and the counterfactuals are similar at both values; α scales with η

Practice ties and peer outcomes leave $\hat{\alpha}$ unchanged

Mean $\hat{\alpha}$

0.192

- Baseline
 - Choice model re-estimated in each HRR, $\eta = 1$

Practice ties and peer outcomes leave $\hat{\alpha}$ unchanged

Mean $\hat{\alpha}$

0.192

- Baseline
 - Choice model re-estimated in each HRR, $\eta = 1$
- Integrated
 - Adds an indicator for the PCP and surgeon sharing a practice

0.192

Practice ties and peer outcomes leave $\hat{\alpha}$ unchanged

Mean $\hat{\alpha}$

- **Baseline**
 - Choice model re-estimated in each HRR, $\eta = 1$
- **Integrated**
 - Adds an indicator for the PCP and surgeon sharing a practice
- **Peer outcomes**
 - Also adds surgeon quality from practice colleagues' patients

0.192

0.192

0.203

Mean $\hat{\alpha}$

- Baseline
 - Choice model re-estimated in each HRR, $\eta = 1$
- Integrated
 - Adds an indicator for the PCP and surgeon sharing a practice
- Peer outcomes
 - Also adds surgeon quality from practice colleagues' patients

0.192

0.192

0.203

Takeaway

The estimated weight on quality is stable, and the counterfactuals scale with it.

- ξ_j absorbs time-invariant reasons to choose a surgeon: office amenities, health-system affiliation, advertising
- Learning is identified from differences in outcome histories across PCPs referring to the same surgeon, net of ξ_j
- Within HRRs, the correlation between $\hat{\xi}_j$ and observed quality is near zero in most markets
- Predicted referral shares track actual shares closely across the range



- γn_{jt} : a surgeon treating more patients is less attractive for the next one (for example, wait times)
- Two steps (Richards-Shubik, Roberts and Donohue 2022, following Bayer and Timmins 2007)
 - Estimate the logit with a surgeon-by-period fixed effect $\delta_{j\tau} = \gamma n_{j\tau} + \xi_j$
 - Regress $\hat{\delta}_{j\tau}$ on $n_{j\tau}$ by 2SLS
- Instrument: the surgeon's patient count predicted from distances alone
 - Other patients' distances shift the surgeon's volume but not this patient's referral
 - First-stage coefficient about 1.4, F above 1,000
- $\hat{\gamma} \approx -0.1$ per 100 patients: 100 more patients over three years makes a referral 7–9% less likely

- Myopic: refer to the surgeon with the highest expected payoff for this patient
- Forward-looking: also value what the referral teaches about the surgeon, for future patients
 - Gittins index $g(m, v)$, rising in the variance of beliefs, with the closed-form approximation of Brezzi and Lai (2002)
 - Familiarity enters as the discounted value of sending future patients to the same surgeon
 - No parameters beyond the myopic model
- The forward-looking α is somewhat larger (0.26 vs. 0.20 at $\eta = 1$)
- Counterfactual results are similar in the two versions

- Each referral sends one patient to one surgeon, and the outcome informs beliefs about that surgeon only. The problem is therefore a multi-armed bandit, which the Gittins index solves
- The index has no estimated parameters ($\beta = 0.95$, ρ , and η are fixed, and α multiplies it), so the forward-looking model is a standard logit with a precomputed regressor
- A full dynamic solution would be needed with any of the following, none of which is in the model:
 - Beliefs correlated across surgeons, so that one outcome informs several surgeons
 - Prospects that change while a surgeon is not chosen, such as drifting quality or depreciating familiarity
 - Learning without referring, from other PCPs' outcomes or a surgeon's volume



Myopic model	Mean	Percentile across HRRs				
		10th	25th	50th	75th	90th
$\alpha (\eta = 1)$	0.202	0.000	0.000	0.117	0.283	0.537
$\alpha (\eta = 5)$	0.397	0.000	0.000	0.212	0.561	1.089
π (distance)	-0.069	-0.099	-0.082	-0.068	-0.053	-0.038

- α/π : roughly equivalent to 2.9 miles of additional travel (surgeons average 29 miles from the PCP), in line with small quality elasticities elsewhere
- Where PCPs respond, one more failure lowers the referral probability by about 0.003
- Estimated in each HRR; the naive standard error of the national mean is about 0.011



PCPs learn true quality as of January 2013 and keep the familiarity they have, with referrals and congestion re-solved. Surgeon quality is held at its observed rate. Myopic model, $\eta = 5$.


Resetting familiarity moves more patients with no health gain





With familiarity reset, distance and the fixed effects drive the reallocation. The within-market spread of fixed effects is more than 7 times the quality signal.

Longer referral relationships are with better surgeons

- Mean surgeon success rate rises with the number of patients in the relationship, from .891 to .903
- Consistent with familiarity gradually sorting PCPs toward better surgeons
- This is why keeping familiarity helps in the full-information counterfactual



Weighting quality like distance avoids about 20% of failures



Beliefs are set to true quality, and fixed effects and familiarity are removed. α is set so one SD of quality moves utility as much as one SD of distance. Distance and congestion still apply. Myopic model, $\eta = 1$.