Demand Fluctuations, Capacity Constraints and Repeated Interaction: An Empirical Analysis of Service Quality Adjustments*

Olivier Chatain†
Strategy and Business Policy Department
HEC Paris

Alon Eizenberg‡
Department of Economics
The Hebrew University of Jerusalem

March 2015

Abstract

We utilize a unique database from a large legal services provider to examine how service quality responds to the firm's available capacity and workload, and to the nature of the firm-client relationship. We develop empirical measures of both the (internal) level of resources available to the firm at different points in time, and of the (external) value creation for customers. Our results indicate that a service provider can use quality adjustment as a substitute for price adjustment in order to tackle demand fluctuations in the presence of capacity constraints, and to foster long-term relationships with customers. Specifically, we show that service quality increases with the amount of the firm's available resources, decreases with the firm's workload, and increases with the number of previous successful interactions with the client. By documenting these relationships, we wish to shed light on the limitations of current estimates of consumer surplus in service industries, as well as on potential inefficiencies in such industries.

Keywords: Resource Allocation, Endogenous quality, Capacity constraints, Service industries, Customer loyalty, Resource allocation

---

*A previous draft was circulated under the title "Timing is Everything? An Empirical Analysis of the Determinants of Service Quality Provision." We thank the managers of the legal service provider for giving us access to the data and for their time and patience during the development of the paper. We thank Victor Aguirregabiria, Heski Bar-Isaac, Dan Ben Moshe, Greg Crawford, Zvi Eckstein, Oded Galor, David Genesove, Saul Lach, Yanping Liu, Michelle Sovinsky, and seminar participants at Bar Ilan University, the Hebrew Univeristy of Jerusalem, IDC Herzlia, University of Toronto, Tel Aviv University, and the University of Zurich, and conference participants at the Israeli I.O. day (Tel Aviv) EARIE (Milan) and MaCCI (Romrod) for comments and suggestions. Kostia Kofman provided excellent research assistance.

†chatain@hec.fr

‡alon.eizenberg@mail.huji.ac.il

1 Introduction

The adjustment of prices is a fundamental aspect of economic activity, and it allows firms to achieve a variety of goals. One such goal is the optimal allocation of scarce capacity in the presence of demand fluctuations. Capacity choices involve a tradeoff: choosing a low capacity restricts the firm's ability to meet high demand states, while setting a high capacity level leaves much of that capacity idle in low demand states. Price adjustment offers a natural solution to this issue, often practiced by airlines: setting high (low) prices when demand is high (low) allows the firm to meet various demand levels given a fixed capacity. Price then serves as a signal of the underlying cost of serving the marginal client, suggesting that this mechanism can be viewed as economically efficient. Another natural example of how firms use price adjustment is *price discrimination*: the practice of charging different prices to different consumers. Many firms, for example, run loyalty programs that reward repeated customers with discounts in order to provide incentives for repeated purchases.

In light these important roles played by price adjustments, it is interesting to note that firms operating in professional service industries (e.g., consulting, legal services, investment banking advisory) often refrain from this practice to a large extent, keeping their pricing highly stable and largely uniform. Attorneys, for instance, often charge an hourly rate for their services, with this rate remaining stable for substantial periods of time. Expert physicians, accountants and consultants also tend to have fixed rates. While some negotiation over the fixed rate is likely, resulting in some customers getting discounts, the price still remains largely fixed, serving as a stable signal of the service provider's level of expertise. Frequent price adjustments, or substantial price differentials across customers, would likely hamper the firm's ability to use the price as a signal, providing one reason why professional service providers may be reluctant to adjust prices flexibly and frequently.

This reluctance to adjust prices raises a natural question: is there an alternative channel that allows professional service providers to achieve the goals described above, i.e., managing demand fluctuations given a fixed capacity level, or providing incentives for customer loyalty? In this paper we suggest that such a mechanism exists: service providers adjust the *quality of service* as a substitute to adjusting prices. Indeed, service providers have a lot of leeway in adjusting service quality, as a consequence of the non-contractible dimensions of the client-firm interaction. A service provider may, at her discretion, decide how much effort and resources to spend on a given client. The provider may, for example, choose whether or not to be available for follow-up consultations with the client over the phone, beyond the standard meeting time. The provider may also choose which professionals to assign to the client's case, reserving the better professionals for favored customers, or restricting their availability when demand for them is high.

Our paper contributes to the literature by documenting such patterns using internal records from a large legal service provider. Our analysis documents two such mechanisms:

first, we show that the firm provides better service quality at times when its capacity constraints are less binding, and interpret this as evidence that quality adjustments can be used in lieu of price adjustments for the purpose of tackling demand fluctuations given a fixed capacity. Second, we show that service quality also increases with the number of previous interactions with the same client in the past. We view this mechanism as an analog of the familiar loyalty card mechanism. Unlike the loyalty card, however, the mechanism we document is an informal one, and is therefore less transparent and more subtle. This finding is also interesting since theory alone would not necessarily predict this pattern: the firm also has a conflicting incentive to divert resources toward *new potential customers* in order to lure them in.

Our paper, therefore, documents the practice of adjusting service quality as an alternative to the adjustment of prices. This practice allows the firm to keep prices relatively fixed and has clear advantages in an environment in which price adjustment is costly. Quality adjustment can be made in a highly-flexible fashion, allowing the service provider to respond immediately to an unexpected surge in demand, for instance, by dedicating lower-than-normal resources to each client being currently served. Price adjustments, in many setups, cannot be made with the same degree of flexibility.

Empirical strategy. We utilize a data set of internal records from a single firm (hereafter "the Firm") that provides legal staff to customers in need of legal services. We observe very rich information on each of the Firm's clients and the projects associated with them, and on each of the attorneys employed by the Firm.

Our analysis follows two distinct steps. The first step defines measures of service quality, as well as measures of the firm's capacity and workload over time. We also define a measure of the extent to which the interaction with any specific customer is repeated. In the second step, we use the measures of service quality as dependent variables, and the measures of capacity, workload and repeated interaction as explanatory variables in regression analyses. This enables us to document the patterns described above: we show that service quality increases with available capacity and decreases with the firm's workload, and, in addition, increases with the extent of previous interaction with the customer.

Defining and measuring service quality—our dependent variable— is a nontrivial task. Such measurement is motivated by the nature of the service rendered by the studied firm. This firm provides legal services to business customers. Its business model differentiates it from traditional law firms: it provides each potential customer with a "shortlist" of attorneys-employees who are deemed to be adequate candidates to work on the client's case. The client can then interview these attorneys, and either choose to work with one of them, or choose none of them, effectively turning the Firm's services down in favor of some outside alternative. Our study focuses on this aspect of the Firm's activity, and defines the relevant service as *providing customers with an opportunity to match with suitable attorneys* (as opposed to the quality of the legal service itself). The quality of this service is, therefore, reflected in various aspects of the aforementioned shortlist: how much choice is granted to specific customers, and to what extent do these choices provide a good match to customers' tastes and needs?

We therefore define two measures of service quality, each capturing a different aspect of the shortlist. The first measure focuses on the number of options, i.e., on the length of the shortlist. Providing a client with three options rather than two increases the probability of a good match, implying that the length of the shortlist can serve as a measure of service quality. One may, of course, wonder whether providing a customer with more choice is always beneficial—under certain circumstances, it may actually be viewed as reflecting lower service quality. As we discuss below in detail, however, the 25th, 50th and 75th percentiles of the shortlist length distribution are 1, 2 and 3, respectively. The firm, therefore, does not overwhelm customers with huge choice sets, supporting the notion that a longer list does reflect more choice and better service.

Our second measure of service quality goes beyond the question of how many choices are provided, and examines the options themselves. This second measure is defined as the client's estimated *expected utility* given the provided shortlist. This measure is computed by noting that the client's problem has the familiar econometric structure of a discrete choice problem: we observe a decision maker (the client) choosing from among several alternatives (the attorneys on the shortlist), with key characteristics of both the decision maker and of each alternative being observed by the econometrician. We can therefore estimate a simple model of client preferences: the conditional logit model (McFadden 1974). Having estimated this model, we can compute the client's expected utility given the list. This measure allows us to collapse the shortlist into a single statistic that captures the extent to which the characteristics of the offered attorneys provide a good match to the client's preferences and needs.

The detailed nature of the dataset, derived directly from the Company's IT system, allows us to additionally develop measures that serve as explanatory variables. We measure the available capacity at the time in which the client needs to be served with a shortlist by counting the number of attorneys who, at that point in time, are employed by the firm and have the relevant expertise and seniority level. We measure the company's relevant workload by counting the number of projects that are "similar" to the relevant client's project and are being contemporaneously processed by the firm. Finally, we also measure the extent of previous interaction with the client by counting the number of previous instances in which the client was assigned an attorney by the firm.

Implications and some caveats. A clear caveat to our analysis is its reliance on data from a single firm, which, as described in detail below, has certain unique features. It is therefore not easy to evaluate the extent to which our findings generalize to other firms or service industries. While this is a limitation, we believe that the focus on a specific firm allows us to capture patterns that are very difficult to observe using data from multiple firms or industries. In particular, the specificity of the setup allows us to tackle the difficult task of measuring service quality in a way that is guided by institutional details. By doing this, we are able to document patterns that may be present much more broadly, beyond the particular firm studied here. Future research of additional case studies could help complement the approach taken in this paper.

Another caveat is that we do not formally address the issue of learning over time, by both the Firm and its clients. One interesting possibility is that both parties improve their understanding of the interaction over time: the client may have a better sense of the Firm's abilities and resources, and the Firm may develop a better sense of a specific client's preferences and needs. While this is a clear possibility, our analysis of repeated interactions abstracts from this dynamic, complicated aspect of repeated interactions. In particular, our discrete choice model does not allow the customer's preferences to evolve across interactions with the Firm. This static approach could be justified on the grounds that clients tend to return to the Firm in different contexts over time (e.g., with legal matters that require very different expertise). We view this static approach as a first step to modeling the issues at hand, leaving the modeling of learning to future work.

Notwithstanding the above caveats, we believe that the patterns documented in this paper have important implications for our understanding of service industries. In particular, they illuminate important limitations of typical studies that measure output and welfare in service industries. Consider, for example, two clients who get an annual checkup for their automobiles at a service provider, each of them paying 50 US$. Imagine that one of them arrived on a "good" day when the provider was not too busy, and had a better and more careful job done on her car than the other customer. Typical economic data from such an industry would typically not reflect such differences in the quality of service across different customers *conditional on provider and price*, and instead would treat the two transactions symmetrically.

While a considerable body of research in economics (notably Griliches 1961, Pakes 2003) emphasizes the importance of computing quality-adjusted indices of *physical output*, the literature remains largely silent on the issue of how to measure service quality and account for possible differences across customers. Our paper, while having the limitation of providing evidence from one specific firm, reveals that within-provider quality differences are indeed important.

Another potential implication of quality differences across customers concerns the efficiency of the market mechanism. Note, in this context, that adjusting prices and adjusting quality may have very different economic implications. The former practice may be viewed as an efficient mechanism for reasons described above. The latter practice may lead to inefficiencies in cases where customers fail to fully observe the quality of service they receive, while being aware of the fact that, in general, different customers enjoy different quality of service. Since customers often lack the expertise that would enable them to determine the quality of the professional service they received, such scenarios are likely. Consistent with the classic work of Akerlof (1970), such information asymmetries naturally lead to suboptimal market outcomes. While our study does not address information asymmetries, our findings highlight a potential for inefficiency in service industries. Such inefficiencies may, under some circumstances, justify policy interventions: for instance, regulations that mandate a minimum standard of service.

Another interesting aspect of our findings is that little is known about the exact manner with which firms use their resources in this context. Documenting the practice of adjusting service quality according to the demand level informs us about the firm's ability to identify the level of demand it faces at any point in time, and act upon this information in a sophisticated fashion. The extent to which managers are able to employ such sophisticated strategies has clear managerial implications: for instance, the greater is the ability to use this strategy, the lower is the long-run level of capacity that the firm needs to maintain.

The finding that service quality increases with the extent of previous firm-client interaction also has important implications. This mechanism induces switching costs, much like any typical loyalty program (e.g., frequent flyer cards). The possibility that customers may be "locked into" continued use of service providers (e.g., accountants or attorneys) via this mechanism suggests an anticompetitive effect of this practice. In general, the impact of such practices on the degree to which the market is competitive is ambiguous.[1] Even less is known about the welfare implications of the mechanism we document here, given the fact that it is an informal and nontransparent mechanism. This motivates additional work on the welfare implications of history-based quality adjustment in service industries.

Relationship to previous literature. Our paper treats product quality as an endogenous variable, taking into account both the internal constraints faced by the firm when delivering a given quality level to its customers, and the nature of its long-term relationship with customers. In this regard, our study relates to several recent empirical papers. Macchiavello and Morjaria (2014) test a model of relational contracts using data from the Kenyan flower market and, in particular, show that negative supply shocks affect delivery in a way that depends on the age of the relationship with clients. The research of Batt and Terwiesch (2012), set in a hospital context, shows how service time and quality in emergency rooms responds to variation in the demand level. Our paper is similarly concerned with the utilization of available resources but takes a complementary view point: we go beyond measuring quantitative service measures (such as the time spent with the patient, or mortality rates) and measure a client's expected utility given the resources spent by the firm.

Our work also relates to the "insider econometrics" (Ichniowski and Shaw, 2009) approach to studying organizational performance. Similar to work in that stream, we use fine-grained employee level data to understand firm behavior. However, while that line of work tends to focus on employee productivity (with a recent example being Hendel and Spiegel 2013), our main object of analysis is rather at the level of the firm's decision to allocate employees to potential clients. We therefore combine *internal* information on resource allocation with *external* information on client choices and preferences to provide a more complete picture of the interaction between such elements than that provided by the extant literature.

[1] See for example Von Weizsäcker (1984), Klemperer (1987).

Our study also complements work in empirical industrial organization. While work in this area traditionally treated price as the endogenous variable and other product characteristics as fixed, a more recent literature treats product quality and characteristics as endogenous.[2] This literature, however, focuses on physical product quality and treats it as a strategic choice made once and for all with respect to all units sold. In this paper, in contrast, we treat quality as a short run variable that can be adjusted on a case-by-case basis. This is akin to product customization, but is driven by the firm's choice given its internal functions and constraints.

Finally, our work relates to the business strategy literature. This literature has consistently asserted that creating value for customers is key to the competitive performance of firms (Brandenburger and Stuart, 1996) and that firms should organize to best deliver this value to their target customers (Porter, 1985). While work in this tradition has recently attempted to quantify the link between value creation and performance (e.g., Chatain, 2011) there is no empirical study measuring the value a firm creates and relating it explicitly to the tradeoffs it faces internally regarding the allocation of its scarce resources. Our paper thus empirically brings together both the external (value creation for customers) and internal (tradeoffs) sides of business strategy.

The paper is organized as follows. Section 2 describes the data and important facts regarding the Firm's operations. Section 3 presents results that use the length of the shortlist as the dependent variable, while Section 4 estimates a client preferences model, uses it to construct a measure of client utility (or "value"), and then uses this estimated utility as a dependent variable. Section 5 concludes.

2 Setting and Data

2.1 Basic setup of the legal service provider

This study uses internal data from a firm that matches highly skilled legal professionals to corporate clients in a major metropolitan area for short to long term projects. The Firm's business model is to offer access to attorneys who are as skilled as those in the best law firms but at a significantly lower cost, as the clients are not charged hourly rates. Clients are granted flexible access to these attorneys in order to complete specific projects, without having to hire the attorneys. The Firm attracts and selects attorneys who used to work at top law firms or at "in house" legal departments and offers them the opportunity to continue being engaged in sophisticated legal work while enjoying much more personal flexibility than is possible in traditional career tracks. The client pays the Firm a daily rate that varies with the seniority of the attorney and the Firm pays the attorney a salary, negotiated *ex ante* with the Firm, that is proportional to the time spent on the project. The Firm targets large corporate clients who are concerned about the rising costs of using

[2]Examples include Crawford, Scherbakov and Shum (2011), Fan (2012), and Eizenberg (2014). See Crawford (2012) for a recent survey.

traditional law firms and who are sophisticated enough to know which legal work can be farmed out to attorneys on a short term contract.

The data we use have been generated by the Firm's IT system and cover the Firm's activity over a time period of 2.5 years. It provides an accurate description of the "life cycle" of a potential project. This process begins with a potential lead for a project, and the emergence of such a lead is recorded in the system, so that we observe the date in which the Firm begins the task of trying to staff such projects. A project lead could stem from a potential client's indication of a need for legal services in some designated area (for instance, Intellectual Property litigation, or Employment law). Alternatively, the firm itself may decide to approach a potential client. The process ends in one of the following two possibilities: either an attorney was assigned to work on the project, or not. The project is "closed" in the system once an assigned attorney completed working on the project, or when the project is designated as closed by the Firm since no attorney was assigned to it.

Once a lead emerges, the firm considers the pool of attorneys that are available for this task. We shall refer to this general pool as the "long list." Attorneys who belong in this long list are those who are under contract with the Firm, and are potentially interested in being assigned to a project.[3] Given the long list, the Firm proceeds to assign a "short list" that consists of a few attorneys. The shortlist contains, on average, 2.71 attorneys with the median shortlist length being two.[4] This short list is presented to the client, who is then able to interview each attorney.

Descriptively, the relationship between shortlist length and the probability of "landing" the project appears weak and even slightly negative. Specifically, the simple correlation between shortlist length and a dummy variable taking the value 1 if the project was landed is -0.09. Regressing the "landed" dummy variable on shortlist length, accounting for client fixed effect, yields a negative and statistically significant coefficient of about -0.02.[5] This finding should not be surprising. On the one hand, we should expect more choices to be correlated with higher client takeup rates. on the other hand, however, the shortlist length is endogenous, and the firm may offer more options to legal projects that are *a-priori* difficult to staff. It is possible that these effects largely cancel each other out, yielding the small observed correlation. The fact that this continues to hold when controlling for client fixed effects suggests a substantial scope for heterogeneity at the level of the individual project, over and above heterogeneity at the client level.

In many cases, the short list consists of a single attorney. Interviews with managers at the Firm suggests that these cases may happen when the client trusts the Firm to pick

[3]Attorneys under contract with the firm are not obligated to accept projects. However, they have an incentive to be assigned since their pay is determined by the extent to which they are employed in projects.

[4]A handful of projects have rather long "shortlists," the longest list being 70 with the next-longest being 24. The final set of projects that we use in estimation excludes the projects with these extreme shortlist lengths on account of missing data issues, and, in any event, robustness checks indicate that our results are not sensitive to excluding projects with shortlists consisting of more than ten attorneys.

[5]This coefficient remains stable when controlling for the relevant (macro) practice areas, described below. These descriptive findings were derived from the full sample of 787 projects used in estimation, see below.

the attorney by itself. Moreover, sometimes the client makes its choice without actually interviewing the candidates, an issue to which we return below. Ultimately, the client may choose one of the suggested attorneys, or choose none of them to work on the project. The latter outcome is viewed as a project that is "lost" from the point of view of the Firm. This assignment process is captured in Figure 1.

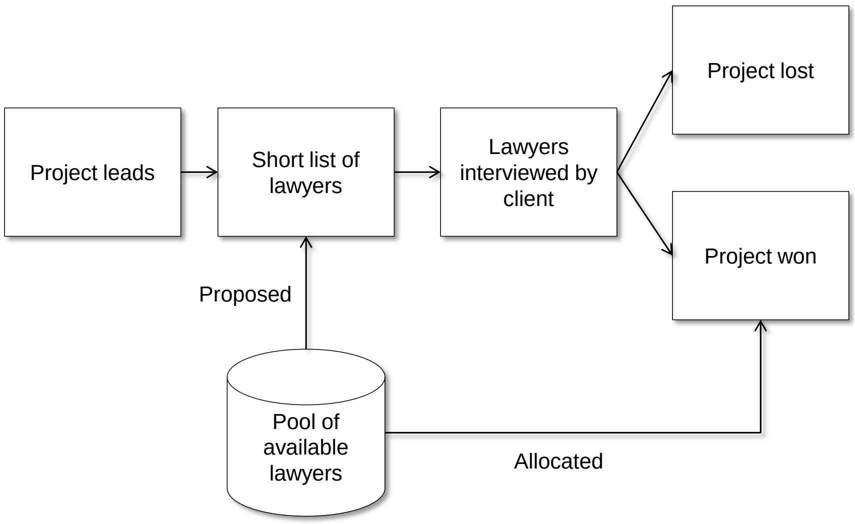

Figure 1: The assignment process

If the project was "landed," i.e., one of the lawyers was selected, that lawyer is assigned to the project. The client pays the Firm an amount that is determined by the project's length, and by the Seniority level of the selected attorney (more on Seniority below). The attorney is compensated by the Firm in proportion to a designated annual salary pre-determined in her contract. The attorney is paid a fraction of this salary, determined by the fraction of days in which the attorney was assigned to projects by the Firm.

Our goal is to use this setup in order to create empirical measures of the resources at the disposal of the firm, and of the resources it decides to allocate to specific projects. The next subsection explains the structure of the data, and the variables that we define in order to capture these aspects.

2.2 Data description and variable definitions

The data provide a very rich description of the attorneys that belong to the Firm's pool of available attorneys (the "long list"), of the potential projects the Firm is trying to staff, and of the clients themselves. A total of 955 attorneys are observed during the sample period. We observe several variables characterizing each attorney. The first is the attorney's *practice area*. A total of 33 such practice areas are defined, allowing us to capture an attorney's relevant area with relative precision. Attorneys are often associated with multiple such areas, but typically no more than two or three. For some purposes, we aggregate these 33 areas up to 5 "macro" areas: (i) Employment, Benefits and Erisa,

(ii) Litigation, (iii) Corporate & Securities, (iv) IP and Commercial Transactions, and (v) Real Estate.

We also observe the attorney's designated *seniority level.* Attorneys are classified into four such levels based on the Firm's assessment of their credentials. While clients may prefer a more senior attorney, choosing a more senior attorney results in higher client cost, as the payment to the Firm is a function of this seniority as determined by a list of daily rates.[6] Finally, we also observe additional attorney characteristics: the attorney's designated annual salary at the Firm, and information about the law school from which the attorney graduated, which we use to create a dummy variable that takes the value 1 if the school is one of the top 15 in the U.S. News & World Report ranking, and 0 otherwise.[7]

Several observed variables characterize potential projects, of which we observe 1,535 (noting that we drop some of these due to several issues as described below). First, we observe the relevant practice area, enabling us to determine the extent to which lawyers on and off of the short list match the practice area that is needed for the specific project. We also observe the project's designated seniority level. This is the level deemed by the Firm to be adequate for the project at hand, and it is likely to be determined using information provided by the client. Observing both the seniority and the practice area that is relevant for the project allows us to ascertain the quality of the match between each attorney and the project, and, ultimately, to empirically estimate the importance of the match on these dimensions.

For some projects, however, information about the relevant practice area and / or seniority is missing, and we drop these projects from our sample. We also drop a handful of projects (3.7% of the total number of projects) in which more than a single attorney was selected by the client, since our reliance on a discrete-choice model requires that only a single attorney, at most, would be picked. We also drop matters which shortlist contains attorneys with missing seniority, and matters with multiple practice areas.[8] This leaves us with a total of 787 project leads.

These leads involve 334 unique potential clients, implying that a given client, on average, is associated with 2.35 project leads. Repeated interaction with the same client is therefore an important aspect of the Firm's activity, as our empirical analysis emphasizes. Conversely, 323 clients appear in the sample only once. The maximum number of repeated client appearances is 44. Robustness checks indicate that our results are not driven by a small set of such "dominant" clients. The data also reveal some scope for repeated interaction that involves the same client and the same attorney. In total, 45% of clients with multiple project leads ever had the same attorney included in shortlists for

[6]This practice is consistent with the way professional services are priced in legal services, accounting and consulting.

[7]We have age and gender information for some but not all attorneys. We therefore do not use these variables in our analysis.

[8]The latter is done in order for some of our explanatory variables to be well-defined: namely, the variables that capture the number of attorneys and the number of "live" projects that the Firm has in practice area-seniority cells.

different projects. In addition, 14.6% of clients who worked with the company more than once (in the sense of actually selecting an offered attorney for at least one project) had the same attorney work on their projects more than once.

We found that within the sample of attorney-project paired observations (excluding projects relating to clients who had a single interaction with the company), the simple correlation between the number of times an attorney was included in the clients' shortlists and an indicator for whether the attorney was selected by the client is -0.03. This suggests that being offered to the client multiple times is not associated with a higher probability of being selected by the client. We also note that clients may often return to the Firm with potential projects that are in different practice areas, limiting the scope for the desirability of working with the same attorney multiple times. For the above reasons, our analysis does not focus on a repeated interaction with the same attorney but rather on the repeated interaction between the client and the Firm.

We observe a number of additional variables that characterize the client associated with the project. These include the client's revenue and industry (for some clients).[9] Finally, a crucial aspect of the data is that we observe, for each project, the short list of attorneys designated to it, and the client's choice. We also observe an indicator for whether each attorney on the short list was interviewed. In some cases, no interviews take place. In these cases, we still interpret the short list as an accurate description of the client's choice set. As we discuss in the sections below, however, we pay special attention to this issue when estimating the client preferences model.

Descriptive evidence of capacity and demand. Figure 2 displays the number of project leads that emerged in each week during the sample period. The figure reveals substantial variation in this flow of new project leads, with certain weeks generating a much bigger number of such leads compared to other weeks. As discussed above, such demand fluctuations create challenges for firms that compete by providing a scarce resource—qualified professionals—which supply is limited and fixed in the short run.

Figure 3 displays the number of attorneys included in the pool of available attorneys over the sample period, both in total and within each of the five macro practice areas. Substantial expansion of the pool is clearly observed in "Corporate & Securities" and in "IP and Commercial Transactions," while the other areas demonstrated more stable attorney counts. Figure 4 demonstrates that this expanded pool has been put to work on an increasing number of projects over the studied period.

Finally, Figure 5 examines the ratio of available attorneys to active project leads, presented by week and (macro) practice area. For this purpose we define "available attorneys" by counting, in each week and within each macro area, how many attorneys are in the pool but are not currently assigned to a project. We refer to such attorneys as being "on the beach." We also define "live projects" as projects that, at the relevant point in time, were initiated and were not yet either staffed or lost. The date in which a project is no

[9] We also observe the size of the client's internal legal department, but do not use this information in our regressions.

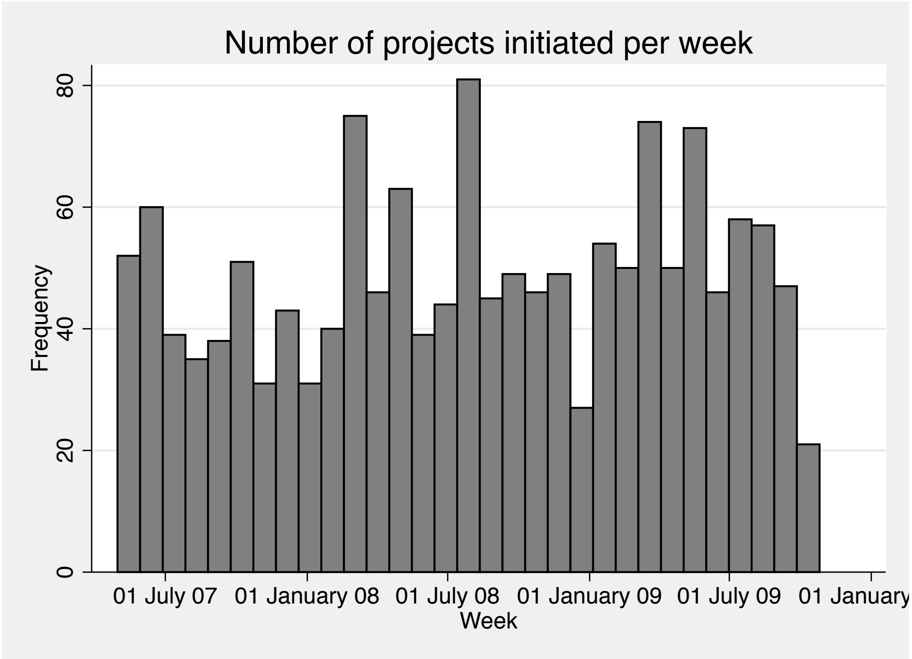

Figure 2: Demand fluctuations reflected in the weakly flow of new project leads

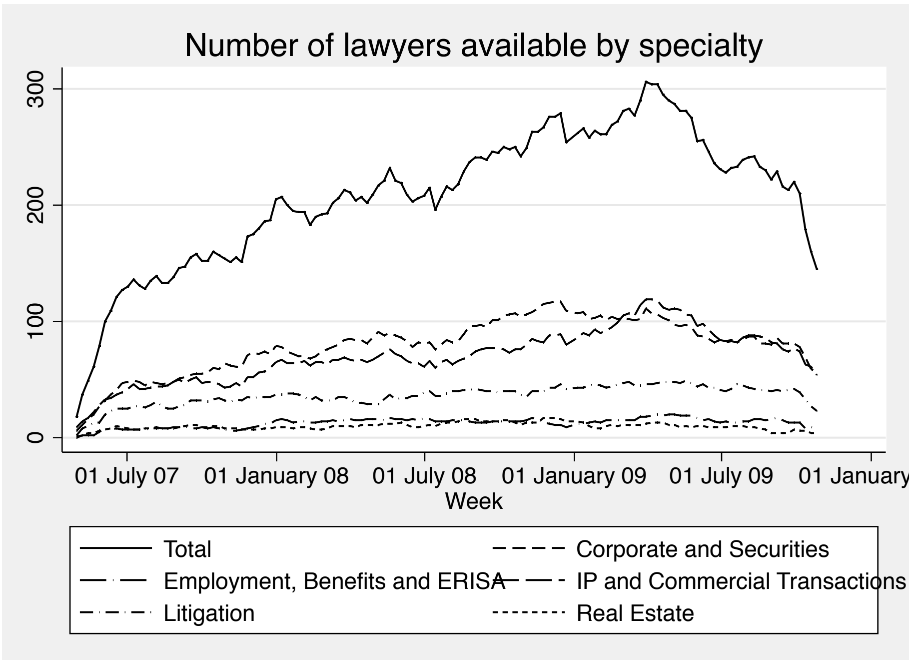

Figure 3: Available capacity: the number of attorneys in the Firm's "pool" in each macro practice area

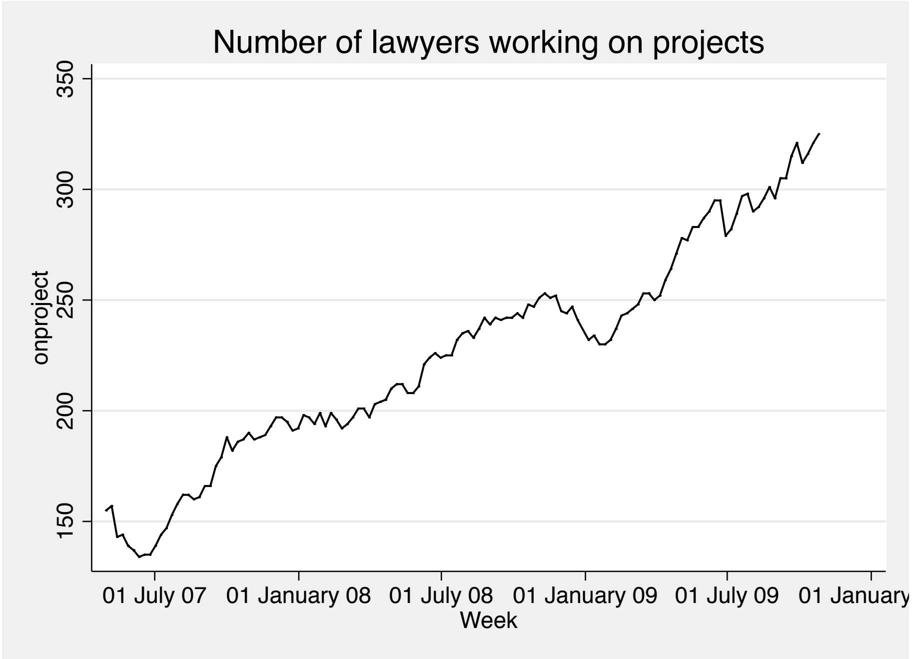

Figure 4: The number of attorneys assigned to projects, by week

longer a "live" project is therefore either the first day an attorney actually works for the client, or the date in which the project is declared closed in the internal IT system.[10]

The figure displays a clear picture for most areas, including the influential Corporate & Securities and IP and Commercial Transactions: the ratio of "on the beach attorneys" to the number of active leads is increasing at first, concomitantly with heightened recruiting of new attorneys. As many of these attorneys get assigned to projects, the ratio declines and stabilizes from the middle of the sample period through its end. Importantly, all our subsequent analyses yield qualitatively similar results with respect to signs and significance levels of our key variables when we restrict the sample to projects originated after July 1st 2008, i.e., following the stabilization of this ratio. This indicates that our conclusions are not driven by the initial expansionary period.

3 Determinants of shortlist length

Framework. Having described the variables we observe, we now motivate an empirical analysis that treats the length of the shortlist as a dependent variable. Consistent with several aspects of the setting, we view this length as a quantitative indication of the amount of resources that the Firm spends on a project lead. First, the Firm must exert managerial efforts in order to add attorneys to this list: it must actively screen the large pool of available attorneys in search of attorneys that match the profile of the potential

[10]For many "lost" projects, but not all, the data document the reason why the project was lost. Examples for such indications are: the client used a competitor, used internal resources, did not follow up, etc. We do not use this information in our formal analysis, and treat all "lost" projects symmetrically, regardless of the reasons that led to the loss of the project.

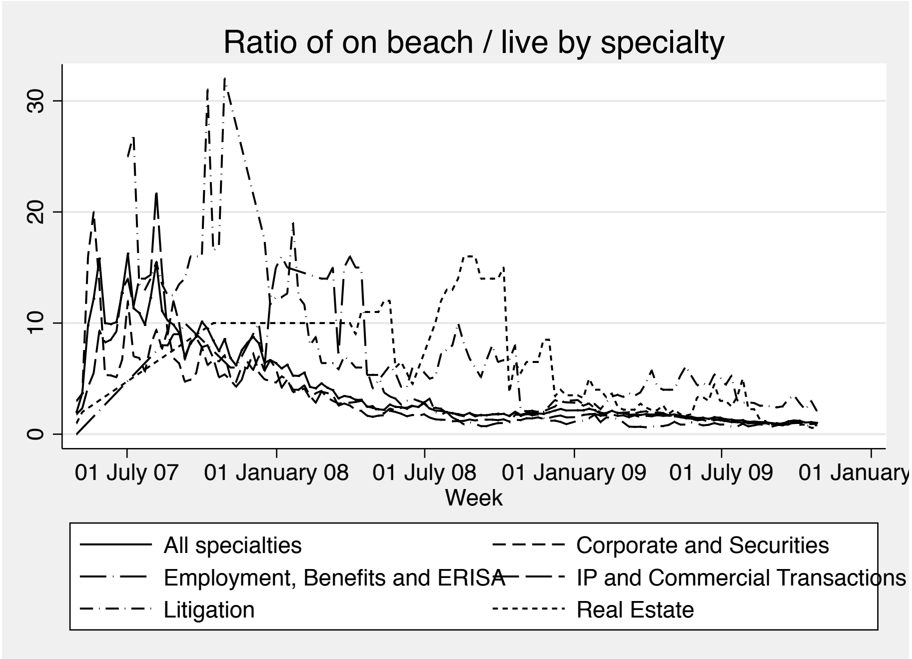

Figure 5: Spare capacity vs. live projects

project. It must also check with such attorneys whether they are disposed to meet the client and ultimately work on the project.

Second, assigning an attorney to the short list for the current project implies an opportunity cost: this limits the Firm's ability to make the same attorney available for assignment in other projects—both current project leads, and the future flow of new project leads—some of which may be more lucrative than the current project. In determining how many options to provide to the current project, the Firm is implicitly solving a complicated dynamic problem: it must consider the tradeoff between attending to the needs of the current project vs. other current and future projects.

While we do not formally state or solve this complicated dynamic problem in this paper, our regressions can be interpreted as estimating the Firm's *policy function* that describes the solution to this implicit problem. Given state variables that determine the firm's available capacity and workload, and the current project's characteristics, we may imagine that the Firm determines the length of the current short list with the goal of maximizing some long-run stream of discounted payoffs. Our estimates of this policy function reveal how the Firm trades off the costs and benefits associated with increasing the amount of resources provided to the current project.

Bajari, Benkard and Levin (2007) present a method for estimating dynamic models of firm behavior following a two-step approach. In the first step, a policy function is estimated by regressing the firm's observed choices on the state variables. In the second step, the structure of the dynamic problem is used to estimate the remaining parameters of the model by matching the estimated policy choices to the ones predicted by the model. Our approach could be viewed as a much more modest treatment of a complex dynamic problem. We do not formally state the dynamic problem, but we do specify the control variable (the length of the short list in this section, and the customer's expected utility in the next) and the state variables (current capacity in the relevant area, current workload in that area, and previous interactions with client). We then estimate the policy function by regressing the "control variable" on these "state variables." We do not, however, state the dynamic problem formally, or perform the involved structural estimation of the "second step." Instead, we focus on estimating the policy function and on interpreting its findings in the context of using quality adjustments as a substitute to price adjustments, consistent with the discussion in the introduction section.

Theoretical predictions. While we do not state a formal theoretical model, our framework allows us to arrive at some theoretical predictions. We should expect the Firm's capacity (measured by the number of "on the beach" attorneys in the relevant practice area-seniority cell) to have a *positive* effect on the shortlist length. The reason is that higher spare capacity reduces both the opportunity cost, and the managerial cost of populating the list with additional names. With respect to the opportunity cost, one can intuitively view this mechanism as a reduction in the shadow price of spending more resources on the current client. We note that our interpretation of "capacity constraints" here is not necessarily that of a binding constraint (i.e., the situation where no attorneys are available to serve the current client), but a more subtle mechanism that depends on a buffer—how much spare capacity does the Firm have at any point in time in each practice area-seniority data cell.

Similar arguments suggest that we expect the Firm's workload (measured by the number of "live" projects in the relevant practice area - needed seniority cell) to have a *negative* effect on the dependent variable. A higher workload implies higher managerial costs associated with populating the shortlist associated with the specific project, stemming from the need to simultaneously staff a large number of projects with similar characteristics. The high workload also reflects higher opportunity costs: there are more alternative assignments for each attorney who may be assigned to the project at hand.

Finally, we consider the predicted sign of the variable capturing the number of previous interactions with the same client in the past. In this case, no clear theoretical prediction arises due to two conflicting forces. The effect of a past relationship may go in either direction: if an existing client is viewed as "captive," the firm may choose to divert resources away from existing clients and toward new potential clients with which it may wish to establish a relationship. If, on the other hand, the firm wishes to strengthen its long-term relationships with existing clients, it would do the exact opposite. Theoretically, the Firm should decide which of these countervailing effects dominates by taking into account its expectations regarding the future flow of lucrative project from both the current client, and from other, "new" clients. We do not formally model these expectations, and instead allow the data to inform us about which effect dominates in the balance.

Results. To test the hypotheses presented above, and learn about the effects of the various explanatory variables on the shortlist's length, we use the 787 projects for which full data is available (see above). Our dependent variable is the number of attorneys in the short list. We define the following independent variables. The first is the number of attorneys who are currently available for staffing (i.e., "on the beach") and whose "macro" practice area *and* seniority correspond to those required for the project. This gives a measure of how tight the internal supply of this type of lawyer is within the firm. The second is the number of projects that are being concurrently considered by the firm (i.e., "live" projects) for the same practice area-seniority combination. This gives us a measure of the perceived demand for attorneys of this seniority and specialization at the time when the short list is designed. A third explanatory variable of interest is the number of projects the client started with the Firm in the past.

We control for five dummy variables for the "macro" practice areas in order to estimate differences in baseline short list length and for certain client characteristics such as revenues. In some regressions, these variables are interacted with the above "attorneys on the beach" and "live projects" variables in order to allow the effects of these supply and demand variables to differ across these macro areas. Importantly, we also include client fixed effects. As the dependent variable is a count variable, we use a Poisson regression. We report robust standard errors clustered by client.

Tables 1 and 2 provide the results of these analyses. Focusing on Table 1 first, we see that across a variety of specifications, a projects' shortlist length is positively affected by the number of attorneys who are "on the beach" and belong in the area-seniority cell that is relevant to the project. This validates our theoretical prediction regarding the likely effect of this variable. These effects are statistically significant (at the 5% level in the first three specifications, and at the 10% level in the fourth column, when client fixed effects are included). As the different columns show, this finding is robust to the inclusion of various project and client characteristics such as the client's revenues, and the project's "macro" area.

Using the fixed effects specification from the fourth column of Table 1, the average partial effect of the number of "on the beach attorneys" on the dependent variable is about 0.046, suggesting that slightly more than 20 additional on-the-beach attorneys in the relevant macro area-seniority cell would add a single attorney to the shortlist. This number is by no means small: given that the 25th, 50th and 75th percentiles of the distribution of the shortlist length are 1,2 and 3, respectively, and that the "long list" includes hundreds of attorneys, it appears that a non-radical shift in the number of on-the-beach lawyers can substantially expand the shortlist.

Table 1 also validates our second prediction: the number of "live" projects who belong in the relevant area-seniority cell has a negative, strongly significant effect on the length of the shortlist. The average partial effect is -0.05, suggesting that about 20 additional "live" projects in the relevant data cell reduce the shortlist length by 1. As explained above, this finding is consistent with higher demand being associated higher managerial costs and opportunity costs, such that the firm optimally reduces the amount of resources it spends on the specific project at hand. Our findings are strongly supportive, therefore, of the notion that the amount of resources spent on a client is significantly affected by the timing in which the client needs to be served: clients who approach the Firm at times in which spare capacity is abundant and few "similar" projects need to be processed have more options to choose from.

We next turn to the effect of the client's past relationship with the firm, measured by the number of previous projects. In the first three columns of Table 1, this variable's effect is found to be insignificant. On the fourth column, however, we control for fixed effects at the client level, and obtain a significant result: the shortlist length is increasing in the number of such previous interactions. As discussed above, one may *a priori* envision two conflicting mechanisms: if returning customers have switching costs that "lock them in" as customers, the Firm may need to invest less resources in serving them. On the other hand the Firm may wish to reward returning customers in order to strengthen its long-term relationship with them, either because no significant lock-in exists, or because the Firm is still investing to create this lock-in. The finding we obtain here is consistent with the second scenario, rather than with the first.

In Table 2, we repeat the analysis, this time including not only fixed effects as in the fourth column of Table 1, but also interactions of macro area dummies with the "attorneys on the beach" and the "live projects" variables, respectively, in order to allow the capacity and demand effects to vary across legal areas. The results reveal that the "beach" effect is positive in all five "macro" areas, but is significant only within the "Benefits and ERISA" and "Real Estate" areas. Revisiting Figure 3 and Figure 5 above, it seems that both of these areas maintained rather stable attorney counts over time (Figure 3), and that "Real Estate" experienced substantial periods with a high ratio of "on the beach" attorneys to "live" projects. This latter area, therefore, had more spare capacity over the sample period relative to most other areas. It is, however, difficult to draw a causal relationship between this fact and the relatively-strong role that appears to be played by the capacity issue within this area. Table 2 continues to report that the "live" effect is present and significant in almost all areas.

Overall, the results in Tables 1 and 2 demonstrate that the quality of service rendered to a client, measured by the amount of choices provided, is significantly affected by the timing in which the client is served, as well as by the past relationship. As stressed in the introduction, the extant empirical literature offers very scarce evidence on such mechanisms, which we interpret as a substitute firm strategy to price adjustments. The length of the short list, however, is just one dimension of the service provided by the Firm, and, in particular, it ignores other aspects of the shortlist, namely, quality and fit of the attorneys on the list. The rest of the paper addresses these additional dimensions by considering an alternative dependent variable: clients' expected utility given the shortlist, estimated via an empirical model of client preferences.

4 Determinants of client utility

In this section we repeat the analysis of the previous section, but this time with a different dependent variable—one that captures not only the number of options provided to the client, but also the quality of the match between these attorneys and the relevant project. This results in an index we refer to as the "Value" provided to the project, essentially capturing the expected utility of the client given the short list provided by the Firm. The analysis follows in two steps: the first step estimates a discrete choice model of client preferences, enabling the computation of the index. The second step uses this index as a dependent variable in regressions of the same nature as those presented in the section above. These tasks are carried out in the following two subsections.

4.1 A model of client preferences

We observe a set of projects $\mathcal{I}$, where, as explained above, a given client may appear in this set several times, each time with a different project. Define $\mathcal{J}_i$ as the set of lawyers considered by project i (the "short list"). Each project-lawyer observation $j \in \mathcal{J}_i$ is characterized by a $1 \times p$ vector of observed characteristics, x_{ij}. This notation indicates that the vector contains characteristics of the project, characteristics of the attorney, and variables that capture the interaction (or "match") between client and attorney characteristics.

In our application, these variables include the seniority of the attorney, a measure of distance between the requested seniority and the attorney's seniority, the number of "micro" practice areas that characterize both the attorney and the project (presumably capturing the quality of the match between the attorney and the project), and additional variables: the client's revenue, the attorney's annual salary, and a dummy variable taking the value 1 for attorneys that graduated from a top-15 law school. We provide more discussion of these variables when discussing the estimation results below.

Our client utility model builds on the conditional logit model (McFadden 1974). Project's i's utility from choosing lawyer j is given by:

$$u_{ij} = x_{ij}\beta + \epsilon_{ij} \tag{1}$$

where ϵ_{ij} are extreme-value deviates that are distributed IID across lawyers and clients, and β is our parameter of interest. The project also has the option of choosing none of the lawyers, to which we also refer as the "outside option," providing a utility of $u_{i0} = \epsilon_{i0}$. This effectively normalizes the mean (across clients) utility from the outside option to zero. Given these assumptions, the probability that lawyer j is chosen is given by the familiar logit model formula as $exp(x_{ij}\beta)/\big(1 + \sum_{j\in\mathcal{J}_i} exp(x_{ij}\beta)\big)$, while the probability of the outside option being chosen, reflecting a "loss" of the project from the point of view of the Firm, is $1/\big(1+\sum_{j\in\mathcal{J}_i} exp(x_{ij}\beta)\big)$. As usual, these probabilities motivate estimation of the model via Maximum Likelihood. This model does not account for the possibility that the choice set is endogenous, in the sense that the shortlist is selected by the Firm in a way that possibly depends on client characteristics that are unobserved to the econometrician. The fact that we control for detailed observed client characteristics hopefully mitigates such concerns to a large extent.[11]

To estimate the model, we do not use all 787 projects, but rather a subsample of 413 projects that have the following feature: the client interviewed at least one of the attorneys on the shortlist. In other words, for the purposes of estimating client preferences (and only for that purpose), we drop projects in which the client made a choice without actually interviewing any of the attorneys on the shortlist. As long as at least one attorney was interviewed by project i, we do use the project in our estimation, and consider all attorneys $j \in \mathcal{J}_i$ as the relevant choice set—including non-interviewed ones. The reason for this choice is that we believe that this subsample may reflect choices that more accurately capture clients' underlying preferences. In unreported robustness checks, we include all 787 projects in estimating the client preferences model, and find qualitatively similar findings.[12]

Client preferences: estimation results. Estimation results are presented in Table 3. Across the columns, several robust conclusions emerge. First, the attorney's seniority is found to have an insignificant effect on client utility. This could be explained by two conflicting effects of attorney seniority: a more senior attorney may be more attractive from the client's point of view, but this is offset by higher client costs.

We do find, however, an important role for *the match* between the seniority requested by the project, and the attorney's seniority: the "Seniority distance" variable has a negative and significant effect on utility. This variable is defined by taking the absolute difference between the two seniorities. Since seniority takes the integer values from 1 to 4, this variable ranges between zero (perfect match) and three.[13] Using the specification in the fourth column of Table 3, increasing the seniority distance variable by 1 reduces the probability of selecting the attorney by about 0.126 (averaged across projects).[14]

Another dimension of the project-attorney match that emerges as important from the estimation results is the practice area. The number of practice areas on which the attorney and the project match has a positive and significant effect on client utility, indicating that clients attribute particular importance to getting an attorney with capabilities and experience in the relevant legal area. An additional "matching" practice area raises the probability of selecting the attorney by seven percentage points (again, this is the average

[11]Note that choice set endogeneity in discrete choice models is rarely addressed in the literature due to its complexity. In ongoing separate work, the authors provide a methodology that would allow one to estimate preferences while taking the issue into account, but it is left outside the scope of the current paper.

[12]Specifically, we retain the results reported below, that client utility significantly decreases in "seniority distance." The effect of the number of matching practice areas remains positive, but loses its significance.

[13]The explanatory variables that enter this estimation procedure are normalized: the seniority distance variable is divided by 4, the number of matching practice areas is divided by 2, and additional normalizations were applied to some of the other variables.

[14]For robustness, we also ran specifications where the seniority match was defined to be binary: a dummy variable took the value 1 if the seniorities associated with the attorney and the project were the same, and zero otherwise. The results remain very similar.

partial effect across projects).

Also included in the model are controls for attorney quality, such as being a top-15 law school graduate, or the attorney's annual salary. With respect to salary, it is important to note that clients do not observe it: recall that clients pay the Firm for the legal service, and the Firm pays the attorney a salary. One may view this variable as a proxy for quality that is unobserved to us, but may be observed by the client. For instance, attorneys with strong verbal abilities may be paid more by the Firm (as their outside options are better), and this is a feature that clients may observe when interviewing the attorney. In this sense, this variable can help us control for unobserved attorney quality. Interestingly, however, both salary and top-15 law school fail to raise client utility significantly, and this pattern is consistent throughout the various specifications.

The client's revenue has a negative and significant effect on utility. The interpretation of this finding is that corporate clients with higher revenues have better outside options, reflecting a systematically smaller utility from the "inside options" (namely, the Firm's attorneys). We use a "missing revenues" dummy variable to account for the fact that revenue data is not available for all clients. Another client characteristic for which we control in some specifications (columns 1,3,4 and 6) is the overall number of projects (won and lost) with the Firm. While this variable is an endogenous outcome, we experiment with including it since it may capture some stable, unobserved factors that affect the client's utility from using the Firm's services. This serves as a suboptimal substitute for a client fixed effect, which is computationally difficult to incorporate into a nonlinear model (client fixed effects are, however, included in the Poisson regressions discussed above, and in the Value regressions discussed below). We do not, however, find a significant effect for this variable, when included.

To sum, estimation results for the client preferences model reveal a clear picture: the only variables that come out significant are those that capture the quality of the match between the attorney and the project at hand. "Vertical" client characteristics such as top-15 law school status or salary fail to explain client utility. While these findings may seem surprising, they are actually consistent with institutional details. The Firm caters to corporate clients in need for specific, targeted legal services. It is likely that clients' first-order consideration is to choose an attorney that has the right expertise for the project, and that other aspects—that may be quite important in other legal setups—are of a lesser importance.[15]

Measuring the value provided to projects. The client preferences model helps us quantify an important aspect of our framework: the expected utility that the Firm provides to potential projects. A familiar result in the logit model is that the expected utility enjoyed by project i, which we denote by V_i, is given by:

[15] One can imagine, for example, that additional attorney characteristics would become increasingly important in scenarios in which a firm is hiring a client as an in-house employee.

$$V_i = \ln\left[1 + \sum_{j\in\mathcal{J}_i} exp(x_{ij}\beta)\right] \tag{2}$$

The index V_i is easily calculated for each project using the estimate $\hat{\beta}$ obtained from the logit estimation described above (specifically, we use the estimates on the fourth column of Table 3). Recalling that the probability of losing project i is given by $1/\big(1 + \sum_{j\in\mathcal{J}_i} exp(x_{ij}\beta)\big)$, it is easy to see that this probability is decreasing in V_i. In order to increase the probability of landing the project, therefore, more Value must be provided. Providing more Value to the current project, however, subjects the Firm to important costs, some of which being opportunity costs that are associated with leaving fewer resources available for future projects. Other costs involve the managerial costs of compiling a shortlist with attorneys that provide a good match with the client's needs.

To see this, note from equation (2) that the firm may increase V_i in several ways. First, it may increase the length of the short list, making more attorneys available to the client's choice. As discussed in Section 2, providing a longer list requires higher levels of managerial effort, and is therefore costly. Another strategy to increase V_i is to improve the match between the characteristics of the attorneys on the short list, and those of the project. Our estimates of the client preferences model reveal that clients strongly value a match between the project's designated practice area and seniority with the attorney's practice area and seniority. In order to increase V_i, therefore, the short list must be populated with attorneys in the relevant seniority-practice area cell.

If these attorneys are currently in short supply (i.e., not many attorneys who are "on the beach" belong in this cell), or the Firm has a high workload (i.e., it is simultaneously trying to populate additional shortlists within the same cell), both managerial and opportunity costs should be greater. For this reason, much like in the analysis in Section 3, we should expect the number of attorneys "on the beach" within the relevant cell to have a positive effect on the Value conferred upon the client, and the workload, measured by the number of "live" projects within the cell, to have a negative effect on this Value. In addition, for analogous reasons as discussed in Section 3, the theoretical prediction for the effect of interactions with the client in the past is ambiguous. The next subsection estimates these effects using a regression framework.

4.2 Value regressions

The amount of Value that the Firm allocates to a project lead, as a function of state variables, is described by the Firm's *policy function* (or, its *assignment function*). This function is estimated by regressing V_i on the state variables associated with project i (as explained above, while we borrow terminology from the literature on dynamic models, we do not actually state or estimate a dynamic model explicitly). Similarly as in the analysis of the length of the shortlist, the main state variables are the current number of "on the beach" attorneys in the macro area and seniority requested by the project, the number of "live" projects in this data cell, and the number of the Firm's previous "landed" projects involving this client. The number of observations in this regression is the number of projects observed with full information, 787. Recall that we estimated the client preferences model with a subsample of 413 projects in which at least one attorney was interviewed. We then use the estimated β coefficients in order to compute the Value V_i from equation (2) for all 787 projects.

Considering the relationship between the dependent variable utilized here (i.e., the "Value") and the dependent variable used in the previous section (i.e., the shortlist length), it is easy to see from equation (2) that the two should be correlated: increasing the length of the list also increases the "Value" or expected utility. This reflects the fact that additional opportunities for a match are available (noting the role of the project-attorney error ϵ_{ij} in the utility specification, often viewed as a "taste for variety"). We find that the correlation between V_i and the length of the shortlist is .75, suggesting that, while the two dependent variables are correlated, the Value V_i does contribute information on top of that conveyed by the shortlist length. We therefore believe that the two analyses complement each other.

Tables 4 and 5 show the results of these "Value regressions." The general pattern of the findings is very similar to that found in section 3, where we used the shortlist length as the dependent variable. Consistent with our predictions, the number of "on the beach" attorneys in the relevant data cell has a positive and significant (at the 5% level) effect on the client's Value, while the number of "live" projects in this cell has a negative and significant (at the 1% level) effect. The number of previously-landed projects with the same client has a positive and significant (at the 1% level) on the Value, including in column 3 that controls for client fixed effects. Once again, therefore, the results indicate that the Firm diverts resources toward existing clients, consistent with a desire to retain them as clients into the future, rather than with viewing the clients as captive. Macro area dummy variables reveal that no systematic differences are observed across such areas in terms of Value provision.

Table 5 continues to follow the spirit of our analysis of the length of the shortlist above: it now regresses the Value conferred upon the project on state variables in a way that allows the "beach" and "live" effects to differ across the five macro areas. The "beach" effect is positive in all areas, but is only significant (at the 5% level) in the "Corporate and Securities" area and (at the 10% level) in the "Real Estate" area (recalling that the Real Estate finding is consistent with our findings from the analysis that treated the length of the shortlist as the dependent variable). The "live" effect is negative in all five areas, but is only significant (at the 1% level) in "Corporate and Securities" and (at the 5% level) in the "Litigation" macro area.

Our predictions regarding the "beach" and "live" effects therefore continue to be supported in this more detailed analysis, with the strongest effect being associated with the "Corporate and Securities" area. Discussions with the Firm confirm the notion that capacity and project-load considerations may indeed be more important in this area since good professionals in this area are in high demand. Importantly, our finding regarding the effect of past interactions continues to be strongly supported in the analysis of Table 5: the number of previous “landed” projects has a positive and significant (at the 1% level) effect on Value.

To sum, the analysis that treats the Value embedded in the shortlist as the dependent variable largely confirms and reinforces the conclusions from the analysis that treated the length of the shortlist as the dependent variable. We find that the Firm invests more resources in clients at times in which its capacity is more abundant, or at times in which its workload of similar projects is lower. The Firm also systematically rewards returning customers relative to new ones. Our concluding section offers some final reflections upon these findings.

5 Concluding remarks

We study the manner with which a service provider allocates its resources across different customers given internal supply and perceived external demand. The analysis highlights that in a professional service industry, firms may have substantial degrees of freedom to adjust the quality of the service provided to customers, for instance, by regulating the extent of choice given to the client, and the quality of the match between the client's needs, on the one hand, and the resources that are made available, on the other hand. The key insight from our analysis is that quality adjustment is used as a substitute for price adjustment by professional service providers.

Our focus on a specific case study—the internal records of a single firm—naturally restricts the scope of our findings. At the same time, this focused lens allows us to discern patterns that are likely to be missed when studying more aggregate (e.g. industry-level) data. While we cannot check the extent to which such patterns hold outside of the particular firm we study, one may speculate based on anecdotal evidence that these patterns may, indeed, be prevalent.

Our main finding suggests an important difference between service industries and industries that produce physical goods. While the quality of a physical good may sometimes be difficult to ascertain from the point of view of the customer, it is less likely to vary substantially across customers who purchase the same good and pay the same price. In a professional service industry, in contrast, the quality of the service may very well vary across such customers, as we demonstrate in this paper. In particular, *timing* seems to play a very important role: customers enjoy different levels of service depending on when they arrive. Customers who arrive at “better” times — when more personnel are available and fewer similar customers are being served — may enjoy better quality services, since the firm optimally chooses to adjust the amount of resources spent on its customers according to its capacity and workload.

Such subtle mechanisms can have several important consequences for service industries. First, the ability to deploy the firm's resources in such flexible fashions implies that smaller long-run capacity levels can be maintained than those that would have been required absent such practices. Second, if customers understand these mechanisms, they are likely to respond to them by strategically timing their purchases of professional services. Some customers may try, for example, to have their accountant handle their tax return well before the deadline, so as to secure more attentive treatment of their case. Some car owners may try to arrive at the service provider in the early morning hours, when it is more likely that the mechanic would devote adequate attention to their automobile.

In many other cases, however, customers may have very little predictive abilities about their service providers workloads. Such information asymmetries may lead to inefficient market outcomes. Similar information asymmetries may arise in connection with the manner in which the firm allocates resources across *customer types*. A company may identify some customers as more lucrative than others, and actively divert resources toward this type of customers at the expense of others. In our context, we find that returning customers are rewarded by being granted a larger (and more relevant) menu of options to choose from. Some other firms, however, may consider returning customers as being locked in, and divert resources away from them and toward new potential customers. It may take customers quite some time to understand which (if any) of such policies is being implemented by a service provider, creating another informational wedge. While firms may be able to use such wedges to their advantage, their equilibrium properties are unclear. In particular, it is possible that firms have to compensate customers for the uncertainty generated by such intricate policies.

The distribution of resources across customers in service industries may, therefore, have important consequences for the efficiency of the relevant markets. As discussed above, they may also imply important limitations of the standard measurement of output and welfare in such industries. While our paper does not offer general methods to improve the economic analysis of service industries, it does use a specific example in order to shed light on these issues. Of note, the service sector has grown substantially in developed economies over the past few decades, motivating additional research into the issues alluded to in this work.

Our paper stops short of addressing some very interesting issues that are left for future work. An important limitation is that we do not formally model information asymmetries which are likely to play a very important role in the mechanisms we study. Our client preferences model does not incorporate client uncertainty, and our analysis of the firm's behavior also does not formally address the information asymmetry between the Firm and its clients. One particularly intriguing aspect is the potential role of *learning mechanisms*. Our work identifies implications of long-run, evolving relationships between the Firm and its customers. Clearly, such relationships could involve mutual learning: customers learn to appreciate what the Firm has to offer (and can perhaps better identify its policies with respect to the choice set it grants them), while the Firm may learn how to better address a specific client's needs. We hope to pursue this important issue in future work.[16]

[16] Israel (2005) offers an example of using data from a single firm—an insurer—to empirically investigate learning

Finally, and again related to informational issues, our analysis emphasizes that some customers (specifically, returning customers) enjoy higher quality on average. But the optimal strategy for the Firm may actually involve distinguishing customers by the extent of *variance* in the Value they receive. This could be optimal if customers have very few observations with which to form an expectation about the quality of service that the Firm offers (Spiegler 2006). An interesting question for future work is what can be learned from the second moment of the distribution of "Value" regarding a service provider's optimal behavior.

effects. Israel's strategy is aided by the existence of exogenous "learning events," namely claims, which absence from our framework forbids us from following the approach developed there.

References

Akerlof, G. (1970). 'The market for lemons', *The Quarterly Journal of Economics*, vol. 84, pp. 488–500.

Batt, R.J. and Terwiesch, C. (2012). 'Doctors under load: An empirical study of state-dependent service times in emergency care', .

Brandenburger, A. and Stuart, H. (1996). 'Value-based business strategy', *Journal of Economics and Management Strategy*, vol. 5(1), pp. 5–24.

Chatain, O. (2011). 'Value creation, competition, and performance in buyer-supplier relationships', *Strategic Management Journal*, vol. 32(1), p. 76102, ISSN 1097-0266.

Crawford, G.S. (2012). 'Endogenous product choice: A progress report', *International Journal of Industrial Organization*, vol. 30(3), pp. 315–320, ISSN 0167-7187.

Crawford, G.S., Shcherbakov, O. and Shum, M. (2011). 'The welfare effects of monopoly quality choice: Evidence from cable television markets', University of Warwick.

Eizenberg, A. (2014). 'Upstream innovation and product variety in the U.S. home PC market', *Review of Economic Studies*, vol. 81, pp. 1003–1045.

Fan, Y. (2013). 'Ownership consolidation and product characteristics: A study of the us daily newspaper market', *American Economic Review*, vol. 103(5), pp. 1598–1628.

Griliches, Z. (1961). 'Hedonic price indexes for automobiles: An econometric analysis of quality change', *NBER The Price Statistics of the Federal Government*, vol. Staff Report No. 3, General Series No. 73, pp. 173–196.

Hendel, I. and Spiegel, Y. (2013). 'Tweaking and the horndal effect', Social Science Research Network, Rochester, NY.

Ichniowski, C. and Shaw, K.L. (2009). 'Insider econometrics: Empirical studies of how management matters', National Bureau of Economic Research.

Israel, M. (2005). 'Services as experience goods: An empirical examination of consumer learning in automobile insurance', *The American Economic Review*, vol. 95(5), pp. 1444–1463.

Klemperer, P. (1987). 'The competitiveness of markets with switching costs', *The RAND Journal of Economics*, vol. 18(1), pp. 138–150.

Macchiavello, R. and Morjaria, A. (2014). 'The value of relationships: Evidence from a supply shock to kenyan rose exports', *The American Economic Review, forthcoming.*

McFadden, D. (1974). 'The measurement of urban travel demand', *Journal of Public Economics*, vol. 3(4), pp. 303–328, ISSN 0047-2727.

Pakes, A. (2003). 'A reconsideration of hedonic price indices with an application to pc's', *The American Economic Review*, vol. 93, pp. 1576–1593.

Porter, M.E. (1985). *Competitive Advantage: Creating and Sustaining Superior Performance*, Simon and Schuster, ISBN 9781416595847.

Spiegler, R. (2006). 'Competition over agents with boundedly rational expectations', *Theoretical Economics*, vol. 1, pp. 207–231.

Von Weizsäcker, C. (1984). 'The costs of substitution', *Econometrica*, vol. 52, pp. 1085–1116.

Tables

Table 1: Poisson regression with the shortlist length as the dependent variable

	(1)	(2)	(3)	(4)
Variables				
# Lawyers on Beach in Area-Seniority Cell	0.0339**	0.0368**	0.0346**	0.0496*
	(0.0165)	(0.0163)	(0.0158)	(0.0257)
# Live Projects in Area-Seniority Cell	-0.0381**	-0.0371**	-0.0349**	-0.0558***
	(0.0175)	(0.0170)	(0.0164)	(0.0204)
# Previous Projects with Client	0.00204	0.00162	0.000248	0.0166**
	(0.00140)	(0.00160)	(0.00171)	(0.00716)
Client's Revenues		-9.17e-07	-8.62e-07	
		(6.09e-07)	(5.38e-07)	
Revenues Missing		-0.188***	-0.207***	
		(0.0598)	(0.0554)	
Corporate and Securities			0.0945	
			(0.0978)	
IP and Commercial Transactions			-0.0529	
			(0.0952)	
Litigation			0.165	
			(0.132)	
Real Estate			-0.148	
			(0.232)	
Constant	0.894***	0.966***	0.937***	Absorbed
	(0.125)	(0.124)	(0.153)	
Quarterly Fixed Effects	Yes	Yes	Yes	Yes
Firm Fixed Effects	No	No	No	Yes
Observations	787	787	787	577
Number of clusters				124

Robust standard errors in parentheses, clustered by client
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2: Poisson regression with area-specific effects

Nbr Lawyers on the Beach in Corporate and Securities	0.0466
	(0.0293)
Nbr Lawyers on the Beach in Benefits and ERISA	0.110**
	(0.0535)
Nbr Lawyers on the Beach in IP and Commercial Transactions	0.0344
	(0.0257)
Nbr Lawyers on the Beach in Litigation	0.0411
	(0.0310)
Nbr Lawyers on the Beach in Real Estate	0.125***
	(0.0466)
Nbr Projects Being Processed (live) in Corporate and Securities	-0.0617***
	(0.0236)
Nbr Projects Being Processed (live) in Benefits and ERISA	-0.131***
	(0.0480)
Nbr Projects Being Processed (live) in IP and Commercial Transactions	-0.0395*
	(0.0214)
Nbr Projects Being Processed (live) in Litigation	-0.0890***
	(0.0292)
Nbr Projects Being Processed (live) in Real Estate	-0.0195
	(0.0521)
Nbr Projects Won with Client up to Current Opened Project	0.0201***
	(0.00603)
Corporate and Securities	1.651**
	(0.774)
IP and Commercial Transactions	1.622**
	(0.742)
Litigation	2.139***
	(0.670)
Employment, Benefits and ERISA	1.338
	(0.874)
Real Estate	NA
Constant	Absorbed
Quarterly Fixed Effects	Yes
Firm Fixed Effects	Yes
Observations	577
Number of clusters	124
Robust standard errors in parentheses, clustered by client	
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$	
Real estate used as omitted category	

Table 3: Client preference model: estimation results

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES						
Constant	0.433	0.357	0.316	0.229	0.352	0.361
	(0.309)	(0.298)	(0.333)	(0.588)	(0.579)	(0.561)
Seniority Attorney	-0.0342	-0.0302	-0.0301	-0.0676	-0.0410	-0.0677
	(0.121)	(0.121)	(0.121)	(0.173)	(0.171)	(0.171)
Seniority Difference	-2.672***	-2.696***	-2.723***	-2.756***	-2.696***	-2.703***
	(0.551)	(0.551)	(0.553)	(0.557)	(0.553)	(0.555)
No. of Practice Areas with Match	0.766**	0.794**	0.815**	0.795**	0.780**	0.746**
	(0.361)	(0.359)	(0.360)	(0.362)	(0.360)	(0.362)
Total projects	-0.821		1.027	1.063		-0.801
	(0.865)		(1.136)	(1.141)		(0.870)
Revenues			-1.417**	-1.431**	-1.032*	
			(0.714)	(0.716)	(0.580)	
Revenues Missing			0.321	0.321	0.299	
			(0.259)	(0.260)	(0.259)	
Top 15 Law Schools				0.152		0.148
				(0.155)		(0.154)
Salary				0.0577	0.0187	0.0470
				(0.344)	(0.344)	(0.339)
Salary Missing				-0.00346	-0.0906	-0.0157
				(1.032)	(1.032)	(1.014)
Observations	1,533	1,533	1,533	1,533	1,533	1,533
Number of cases	413	413	413	413	413	413
Likelihood ratio	-483.6	-484.0	-479.9	-479.4	-479.4	-483.1

Standard errors in parentheses
See notes on variable normalizations in the text
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 4: Value regressions

	(1)	(2)	(3)
Variables			
Nbr Lawyers on the Beach	0.0228**	0.0229**	0.0365**
	(0.0100)	(0.0100)	(0.0157)
Nbr Projects Being Processed (live)	-0.0248***	-0.0244***	-0.0361***
	(0.00956)	(0.00941)	(0.0138)
Nbr Projects Won with Client up to Current Opened Project	0.00473***	0.00416***	0.0114***
	(0.000815)	(0.000882)	(0.00425)
Corporate and Securities		0.128	
		(0.113)	
Employment, Benefits and ERISA		-0.0346	
		(0.126)	
IP and Commercial Transactions		0.0867	
		(0.111)	
Litigation		0.187	
		(0.123)	
Client's Revenues	-4.46e-06***	-4.41e-06***	
	(3.02e-07)	(2.86e-07)	
Revenues Missing	0.134***	0.129***	
	(0.0402)	(0.0389)	
Constant	1.169***	1.063***	1.095***
	(0.0821)	(0.134)	(0.0898)
Quarterly dummies	Yes	Yes	Yes
Client Fixed Effects	No	No	Yes
Observations	787	787	577
Number of clients	124	124	124
R-squared	0.186	0.195	0.469

Robust standard errors in parentheses, clustered by client
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 5: Value regression with area-specific effects

Variables	
Nbr Lawyers on the Beach in Corporate and Securities	0.0374**
	(0.0180)
Nbr Lawyers on the Beach in Benefits and ERISA	0.0264
	(0.0448)
Nbr Lawyers on the Beach in IP and Commercial Transactions	0.0317
	(0.0209)
Nbr Lawyers on the Beach in Litigation	0.0332
	(0.0299)
Nbr Lawyers on the Beach in Real Estate	0.0625*
	(0.0330)
Nbr Projects Being Processed (live) in Corporate and Securities	-0.0456***
	(0.0147)
Nbr Projects Being Processed (live) in Benefits and ERISA	-0.0446
	(0.0429)
Nbr Projects Being Processed (live) in IP and Commercial Transactions	-0.0201
	(0.0163)
Nbr Projects Being Processed (live) in Litigation	-0.0650**
	(0.0286)
Nbr Projects Being Processed (live) in Real Estate	-0.00945
	(0.0321)
Nbr Projects Won with Client up to Current Opened Project	0.0148***
	(0.00412)
Corporate and Securities	0.737*
	(0.410)
Employment, Benefits and ERISA	0.733
	(0.539)
IP and Commercial Transactions	0.619
	(0.411)
Litigation	1.056**
	(0.503)
Constant	0.379
	(0.395)
Quarterly dummies	Yes
Client Fixed Effects	Yes
Observations	577
Number of clients	124
R-squared	0.484

Robust standard errors in parentheses, clustered by client
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$
Real Estate is the omitted category