

Optimal Contracting and the Organization of Knowledge*

William Fuchs Luis Garicano Luis Rayo
Berkeley-Haas and IZA LSE and CEPR Utah and CEPR

August 11, 2014

Abstract

We study contractual arrangements that support an efficient use of time in a knowledge-intensive economy in which agents endogenously specialize in either production or consulting. The resulting market for advice is plagued by informational problems, since both the difficulty of the questions posed to consultants and the knowledge of those consultants are hard to assess. We show that spot contracting is not efficient since lemons (in this case, self-employed producers with intermediate knowledge) cannot be appropriately excluded from the market. However, an ex-ante, firm-like contractual arrangement uniquely delivers the first best. This arrangement involves hierarchies in which consultants are full residual claimants of output and compensate producers via incentive contracts. This simple characterization of the optimal ex-ante arrangement suggests a rationale for the organization of firms and the structure of compensation in knowledge-intensive sectors. Our findings correspond empirically to observed arrangements inside professional service firms and between venture capitalists and entrepreneurs.

*We are grateful to Marina Halac, David Rahman, Esteban Rossi-Hansberg, Ilya Segal, John Sutton, our editor Marco Ottaviani, and three anonymous referees for their valuable suggestions. We also thank seminar participants at the AEA meetings, Berkeley, Boston University, Carlos III, CEMFI, Chicago Booth, Naples, Frankfurt, University of Washington, Washington University, the Utah WBEC, and the NBER for their comments. Miguel Espinosa provided valuable research assistance. Garicano would like to thank the Toulouse Network in Information Technology for financial support. A previous version of this paper was circulated under the title “Trading Know-How.”

1 Introduction

A key role of markets and other organizations is ensuring that society uses the knowledge available to economic agents optimally (Hayek, 1945). Doing so requires conserving the superior knowledge of experts for the right questions; that is, efficiently matching questions and expertise. A solution to this problem is a “knowledge-based hierarchy” in which experts leverage their knowledge by having less expensive workers deal with routine tasks. Firms and other organizations often rely on such hierarchies by placing agents in different positions according to their expertise.¹

The allocation of knowledge is complicated by the presence of two-sided information asymmetries: both the difficulty of the problems posed to experts and the knowledge of those experts may be hard to assess. As a result the market may be unable to ensure an efficient use of expert time. In this paper, we study whether or not there exist contractual arrangements that support an efficient allocation despite these informational asymmetries, and, if so, the form that such arrangements take.

We carry out our analysis in a general-equilibrium setting, in which agents, each endowed with a privately-observed knowledge level and a unit of time, choose their occupation (production or consulting) and whether or not they wish to participate in the market for knowledge. Problems are encountered in the course of production, and vary in their difficulty. If a producer faces a problem that is too hard relative to her knowledge level, she may seek help from a consultant.

We start by deriving, as a benchmark, the first-best allocation under full information. This allocation has the following features. First, agents are split into three sets: *matched producers* (who receive advice from consultants upon failing to solve their problems); *self-employed producers* (who never receive advice); and *consultants*. Second, consultants are more knowledgeable than both types of producers, and self-employed producers, in turn, are more knowledgeable than matched producers. Accordingly, self-employed producers have intermediate knowledge levels. Third, matching between producers and consultants is positive assortative: more knowledgeable consultants devote their time to problems that are expected to be more difficult, namely, those problems that the most knowledgeable producers failed to solve on their own. Fourth, whenever communication between consultants

¹Garicano (2000) studies knowledge-based hierarchies in which ex-ante homogeneous workers acquire knowledge and choose their positions in the hierarchy. Garicano and Rossi-Hansberg (2006) analyze the problem of knowledge acquisition and optimal organizational structure with heterogeneous agents. Both these papers focus on economies with full information.

and producers is sufficiently inexpensive in terms of the time it consumes, the set of self-employed producers is empty. In this case, even the hardest problems are worth a second try and all agents participate in the market.

Intuitively, team production is most effective when the most knowledgeable agents (who specialize in consulting on difficult problems) are matched with the least knowledgeable ones (who specialize in production and get help from consultants only when needed). Agents with intermediate knowledge, in contrast, contribute to team production the least. The existence of self-employed producers, as we shall see, is the key obstacle to an efficient spot market. It is also one of the features that distinguishes the present setting from a standard asymmetric-information setting: in our model, the ‘lemons’ who threaten market efficiency are precisely the agents with intermediate knowledge levels.

Our main focus is on contracting with asymmetric information. We consider two forms of contracts: ex-post (or spot) and ex-ante. Under ex-post contracting, producers seek help from consultants only after facing a problem they cannot solve. Such contracts rely on two pieces of information to condition payments between agents: (i) participation in the contract; and (ii) whether or not the producer solves her problem upon receiving help from the consultant. Under ex-ante contracting, in contrast, consultants and producers match before producers attempt their problems (in a similar way, for example, as contracting inside a firm). Crucially, ex-ante contracts use an extra piece of information to condition payments: whether or not a producer solves her problem on her own.

We find that spot markets (i.e. those with ex-post contracting) are, in general, not efficient. The reason is that self-employed producers may want to overstate their knowledge and become consultants and/or understate their knowledge and seek advice. Excluding them from the market would require that asking for help is costly (to prevent self-employed producers from pretending to have easy unsolved problems) and, at the same time, would require that consultants pay producers up front for the right to offer advice (to prevent self-employed producers from becoming consultants). We show that spot contracts cannot achieve both goals simultaneously. Accordingly, when the first best calls for self-employed producers it cannot be supported by such contracts.

Our main result is that, in contrast to spot markets, ex-ante contracting always supports the first-best allocation, despite agents contracting under two-sided information asymmetries. Moreover, we show that there is a unique ex-ante arrangement that guarantees an efficient allocation, regardless of the underlying parameters. This arrangement has the following

intuitive structure:

1. Each consultant matches, up front, with a number of producers and enters into a bilateral contract with each one. Consultants with higher knowledge match with more knowledgeable producers (and with a greater number of producers).
2. Consultants pay each producer in their team a positive fixed fee (e.g. a wage) plus a performance bonus in the event that the producer succeeds at solving her problem on her own. The fixed fee is decreasing, and the bonus is increasing, in the producer's knowledge level. In other words, as producer knowledge grows, incentives become higher-powered. It is this feature of the arrangement that screens producers.
3. Net of the above payments to producers, each consultant is the full residual claimant of all output produced by the team. The total expected payment from the consultant to each producer is increasing in the producer's knowledge at a rate equal to the marginal contribution of producer knowledge to expected team output. It is this feature that screens consultants: they simply "hire" producers at their marginal value.

Crucially, these contracts are able to simultaneously screen agents on both sides of the market. When serving as residual claimants, consultants care only about their total expected payment to producers. Producers, in contrast, care also about the way in which the total expected payment is split between fixed and bonus payments.

These contracts also guarantee that the most knowledgeable (self-employed) producers stay away from the market. On the one hand, those producers are willing to remain self-employed because being matched with a consultant (under the pretense of having lower knowledge) commits them to share their output with the consultant when succeeding on their own. Since such output sharing is, in expectation, more costly for producers with higher knowledge (who have a higher likelihood of success), it precisely induces producers with the highest knowledge levels to remain unmatched. On the other hand, self-employed producers do not wish to become consultants because, as such, they would become the full residual claimants of the unsolved problems they consult on, a feature they dislike because they are less capable, compared to existing consultants, of solving such problems. Crucially, unlike ex-ante contracts, spot contracts do not allow for the above output-sharing commitments that prevent self-employed producers from entering the market.

We show, moreover, that whenever the first best calls for self-employed producers (i.e. consultant time is sufficiently scarce), the above ex-ante contracts are the *only* ones that

implement the first best. The reason is that granting full residual claim to consultants maximally deters self-employed producers from wanting to become consultants, as it would force them to internalize the full marginal value of their (relatively-low) knowledge. Such maximal deterrence is in turn required to protect the scarce talent of the most knowledgeable consultants.

Our results are reminiscent of Alchian and Demsetz's (1972) seminal work. While they motivate the appearance of a firm on moral-hazard grounds by arguing that the "specialization in monitoring plus reliance on residual claimant status will reduce shirking" (p. 782), our paper instead shows that specialization in advising and residual-claimant status for consultants solves a double-sided adverse selection problem. Such contracts ensure that experts seek to "employ" producers of the right level (as more knowledgeable producers receive, in expectation, a higher compensation) and, simultaneously, ensure that producers wish to "work" for the right experts (as the power of their incentive pay increases with their reported knowledge level).

This simple characterization of the optimal ex-ante contracts suggests a rationale for the organization of firms and the structure of compensation in the knowledge-intensive sector (e.g. legal, consulting, and advertising). A salient stylized fact of this sector is that agents with the highest expertise tend to be the residual claimants of income.² In the manufacturing sector, in contrast, residual claimants tend to be the suppliers of capital, as opposed to the agents with the highest substantive knowledge about production. Our model suggests that this structure of professional service firms is the optimal arrangement in the presence of double-sided adverse selection.

Beyond professional services firms, an industry whose contractual and matching practices fit those in our model is the venture capital industry (with investors taking the place of our consultants and entrepreneurs taking the place of our producers). As Kaplan and Stromberg (2004) find empirically, investors give entrepreneurs cash up front (akin to the fixed fee in our model) and receive equity if the entrepreneur succeeds (akin to paying the entrepreneur a bonus lower than the value of output). Investors also hold convertible debt, which in case the entrepreneur fails is converted into equity, so that investors receive all the equity in the firm (like in our optimal contract). In addition, as our model predicts, higher-quality

²Evidence that law firms display positive assortative matching, strict stratification, and a scale-skill effect (more skilled partners run larger firms) as predicted by the model, is presented in Garicano and Hubbard (2013) on the basis of US Census of Services data.

entrepreneurs tend to give up less equity (akin to receiving a higher bonus in our model).³

Lastly, we consider two extensions. First, we consider the case in which all matching between producers and consultants is controlled by a monopolistic intermediary. We show that the intermediary finds it in her interest to implement the first-best allocation while keeping all rents for herself. This result illustrates an alternative resolution to the failure of spot markets. Second, to expose the critical role of contingent payments, we consider the case in which consultants cannot be paid success-contingent fees, but we otherwise allow for general random allocations. We show that, in such case, the market for knowledge breaks down completely, regardless of the potential gains from trade.

No previous study has, to our knowledge, examined the problem of matching consultants and producers under double-sided asymmetric information. For a complete literature review, see Section 6.

The remainder of the paper is organized as follows. Section 2 presents the model setup. Section 3 derives the first-best benchmark. Section 4, which contains our main results, addresses contracting under asymmetric information. Section 5 presents our two extensions. Section 6 discusses related literature. Finally, Section 7 concludes.

2 Model setup

We model an economy in which, in the spirit of Garicano (2000) and Garicano and Rossi-Hansberg (2006), agents use time and knowledge to solve problems.

Agents, knowledge, and problems. The economy is formed by a continuum of risk-neutral agents, each endowed with one unit of time and an exogenous, privately-known knowledge level $z \in [0, 1]$. Knowledge levels are distributed according to a c.d.f. F with associated density f . We assume $f(z) > 0$ for all z . Each agent may or may not engage in production, which requires time and knowledge as inputs. When an agent engages in production, she devotes her full unit of time, and applies her knowledge, to a single problem of difficulty x . The difficulty of each problem is unknown and is distributed uniformly on $[0, 1]$, independently across problems.⁴ The agent manages to solve the problem on her own if her knowledge exceeds the problem's difficulty level, i.e. $z > x$.

³As the model also predicts, Sorensen (2008) finds evidence of positive sorting between entrepreneurs and VCs.

⁴Since the distribution F is arbitrary, the assumption that x is distributed uniformly is without loss (i.e. is merely a normalization) under suitable regularity conditions (namely, that both z and x are distributed on $[0, 1]$ according to a density function that is everywhere positive).

Occupations. Agents sort into one of two occupations: production and consulting. Producers initially attempt to solve their respective problems on their own. Those who fail may opt to seek advice from consultants. Consultants do not spend any time on production. Instead, they specialize on transferring knowledge to producers who seek advice. Transferring knowledge consumes h (“help”) units of consultant time per problem, with $h \in (0, 1)$, and, for simplicity, consumes no additional producer time. A producer who seeks advice from a consultant manages to solve her problem if the knowledge of the consultant exceeds the problem’s difficulty level. Whether or not a producer receives advice, she dedicates her full unit of time to working on a single problem.

Teams and output. Producers may either never seek advice or form teams with consultants. Teams, in turn, may be formed either ex-ante (before producers attempt their problems) or ex-post (after producers have failed their problems).⁵

Throughout, we describe teams as pointwise matches between a number of producers of a given type z and a consultant of a given type m , while allowing the number of producers in the team to be non integer. The literal interpretation of such pointwise matching is that a small interval of producers $(z - \varepsilon_1, z + \varepsilon_1)$ is matched with a small interval of consultants $(m - \varepsilon_2, m + \varepsilon_2)$ with the requirement that the combined mass of consultants in the team, namely, $F(m + \varepsilon_2) - F(m - \varepsilon_2)$, is h times as large as the mass of problems left unsolved by the producers in the team, which from Sun’s (2006) exact Law of Large Numbers, equals $\int_{z-\varepsilon_1}^{z+\varepsilon_1} (1-t)dF(t)$ with probability one.⁶ This requirement ensures that the combined measure of consultant time is fully occupied and all producers with unsolved problems receive advice. Pointwise matching (potentially involving a non-integer number of producers) simply serves as a convenient approximation for the case in which ε_1 and ε_2 are taken to be arbitrarily small.⁷

Note that we restrict attention from the outset to matching arrangements in which all

⁵We restrict attention to teams with a single layer of consultants. See, for example, Garicano and Rossi-Hansberg (2006) for an analysis of multi-layered hierarchies in a full-information setting.

⁶Such instance of the exact Law of Large Numbers follows from Sun’s (2006) Corollary 2.9 (which implies that, with probability one, the sample distribution of problem difficulties is uniform) and his Corollary 2.10 (which implies that, with probability one, the mass of problems solved by a positive-measure interval of producers equals the integral over this interval of the expected success level z for each producer). We are grateful to Yeneng Sun for pointing out this fact.

⁷Note that this interpretation of team formation is meaningful despite the facts that: (1) the number of problems left unsolved by a single producer is not deterministic; (2) the number of problems that a single consultant can advise on (namely, $\frac{1}{h}$) need not be an integer; and (3) the density of problems left unsolved by producers of type z (namely, $(1-z)f(z)$) is generically not equal to the density of problems that consultants of type m can advise on (namely, $\frac{1}{h}f(m)$).

producers matched with a given consultant have the same type. This restriction is without loss because we seek efficient contractual arrangements, all of which, as we show in Section 3, exhibit (strict) positive-assortative matching.

A solved problem generates a value of \$1 and an unsolved problem generates a value of \$0. Accordingly, the output generated by the various possible arrangements above is as follows:

- A producer who never seeks advice generates an expected output equal to z .
- From an ex-post standpoint, if a consultant of type m devotes her unit of time to advise $\frac{1}{h}$ producers of type $z < m$ with unsolved problems, the expected output of such team is

$$\pi(m, z) \equiv \frac{1}{h} \text{Prob}[x < m \mid x \geq z] = \frac{1}{h} \cdot \frac{m - z}{1 - z},$$

where $\text{Prob}[x < m \mid x \geq z]$ is the conditional probability that the problem is solved once the producer receives advice given that the producer initially failed. We refer to $\pi(m, z)$ as the “ex-post production function.” Note that this function is supermodular (i.e. $\pi_{mz} > 0$). Intuitively, higher consultant knowledge is more valuable, in the margin, when applied to a harder problem. Harder problems, in turn, are precisely those left unsolved by more knowledgeable producers.

- From an ex-ante standpoint, each producer with knowledge z will need help with probability $(1 - z)$. Thus, a consultant of type m is able to form a team with $n(z) \equiv \frac{1}{h(1-z)}$ producers of type z , as, in expectation, these producers need help on $\frac{1}{h}$ problems. (While the number of problems left unsolved by those $n(z)$ producers is random, the fact that matching actually occurs between small intervals of types on either side allows us to treat such number of unsolved problems as deterministic.) The expected output of such team is

$$\Pi(m, z) \equiv n(z)z + n(z)(1 - z)\frac{m - z}{1 - z} = n(z)m.$$

We refer to $\Pi(m, z)$ as the “ex-ante production function.” Note that this function is also supermodular (i.e. $\Pi_{mz} > 0$). Intuitively, since the consultant’s type only matters when advice is given, Π inherits the supermodularity of π .

Note a key difference between the ex-ante and ex-post production functions. A producer with a higher type is more valuable from an ex-ante standpoint ($\Pi_z > 0$), as she is less likely

to require advice. In contrast, a producer with a higher type is less valuable from an ex-post standpoint ($\pi_z < 0$), as a problem left unsolved by a more knowledgeable producer is, in expectation, more difficult.

Contracts and timing. Contracts are pairwise agreements between a consultant and each producer in her team. We consider two types of contracts, which differ in their timing and thus in the information that can be used to condition payments: ex-post (or spot) and ex-ante.

Under ex-post contracting, timing is as follows. First, agents choose their occupations and producers attempt their problems. Second, each producer who failed to solve her problem decides whether or not to enter into a spot contract with a consultant. If she does not, the producer's problem remains unsolved. If she does, the producer receives advice and a fixed transfer $w \in \mathbb{R}$ from the consultant. In addition, the proceeds from the problem are split between the consultant and the producer according to shares α and $1 - \alpha$, with $\alpha \in [0, 1]$. Since there are only two possible outcomes (the problem is either solved or not) an ex-post contract of the form $\langle w, \alpha \rangle$ is without loss.⁸

Under ex-ante contracting, timing is as follows. First, agents choose their occupations. Second, each producer may or may not enter into an ex-ante contract with a consultant. If she does not, she remains in autarchy. If she does, the producer receives a fixed transfer $\omega \in \mathbb{R}$ up front from the consultant. Third, each producer in a team attempts her problem and seeks advice if needed. If the producer solves the problem on her own, the producer obtains a share β (for bonus) of output and the consultant obtains a share $1 - \beta$, with $\beta \in [0, 1]$. If instead the producer receives help from the consultant, output is split between the consultant and the producer according to shares α and $1 - \alpha$. The key distinction vis-a-vis ex-post contracts is that ex-ante contracts may condition payments on an additional event: whether or not the producer solves the problem on her own. Since there are only three possible outcomes, an ex-ante contract of the form $\langle \omega, \alpha, \beta \rangle$ is without loss.

Note that ex-post contracting is a special case of ex-ante contracting. Indeed, any ex-post contract $\langle w, \alpha \rangle$ can be replicated by an ex-ante contract that sets $\langle \omega, \alpha, \beta \rangle = \langle w, \alpha, 1 - w \rangle$. Equivalently, ex-post contracting is formally identical to ex-ante contracting with the additional restriction that $\beta = 1 - \omega$. It is worth pointing out that since ex-ante contracts are more general, teams are always weakly better off, by assumption, when they use such

⁸We restrict α to lie in $[0, 1]$ to prevent budget-breaking schemes in which a player acquires a negative share of output. We impose the same restriction on the share β defined below. Such restriction prevents incentives to sabotage output and is standard in the literature (see, for example, Eswaran and Kotwal, 1984).

contracts. Therefore, our model is not suitable for providing a complete theory of when ex-ante, rather than ex-post, contracts are used in practice. What our model will deliver is an important advantage of ex-ante contracts as well as a characterization of their optimal shape.⁹

Although the main focus of our paper is the case in which the agents' knowledge z is unobservable, it is useful to first analyze the full-information case. All omitted proofs are in the Appendix.

3 Benchmark: the first-best allocation

Consider a fully informed planner who wishes to maximize social surplus. The planner selects the occupation of each agent and a match between producers and consultants. Without loss, we focus in this section on ex-ante matching in which the planner selects a match before producers have attempted their problems, with the understanding that consultants will offer advice on the unsolved problems of the producers they are matched with. Let $\mathcal{Z}, \mathcal{I} \subset [0, 1]$ denote the sets of matched producers and unmatched (i.e. "self-employed") producers, respectively, and let $\mathcal{M} \subset [0, 1]$ denote the set of consultants. (Note that the planner wishes to match all consultants because an unmatched consultant would deliver zero output.)¹⁰

We begin with some basic properties of any optimal allocation:

Lemma 1 *A surplus-maximizing allocation has the following features (except at most over a zero-measure subset of the type space):*

1. **Positive-assortative matching.** *Matching between producers and consultants is (strictly) positive assortative in the sense that if a producer of type z is allocated to a consultant of type m it cannot be that a producer of type $z' > z$ is allocated to a consultant of type $m' \leq m$.*
2. **Producer/consultant stratification.** *The set $\mathcal{Z} \cup \mathcal{I}$ of producers lies below the set \mathcal{M} of consultants.*

⁹In practice, ex-ante contracting may entail costs relative to spot contracting, including the need for consultants to contract up front with a number of suitable producers and to cope with a potentially-random number of questions asked by those producers. Throughout the paper we abstract from such costs.

¹⁰The results in this section generalize the results of Garicano and Rossi-Hansberg (2004), who focus on a uniform type distribution F , and extend the results in Garicano and Rossi-Hansberg (2006) to the case in which skills are exogenous.

3. **Within-producer stratification.** *The set \mathcal{Z} of matched producers lies below the set \mathcal{I} of self-employed producers.*


Intuition for each part is as follows. Positive-assortative matching is optimal since the output function $\Pi(m, z) = n(z)m$ is supermodular and so higher consultant knowledge is best exploited when applied to a harder problem. Producer/consultant stratification follows from the fact that the marginal value of knowledge is higher for a consultant than it is for a producer, which in turn follows from the fact that a consultant leverages his knowledge by applying it to more than one problem.¹¹ Consequently, the planner optimally allocates only the most knowledgeable agents to consulting. Within-producer stratification is optimal because, provided the planner wishes consultants to attempt only a fraction of the producers' unsolved problems, it is best that consultants attempt the problems that are, on average, easiest to solve. These are precisely the problems that come from producers with the lowest types.

From this Lemma we immediately obtain:

Corollary 1 *A surplus-maximizing allocation consists of a partition of the type space into three intervals, $[0, z_1]$, (z_1, z_2) , $[z_2, 1]$, such that:*

- A. *Agents in $[0, z_1]$ are matched producers who seek advice if needed.*
- B. *Agents in (z_1, z_2) are self-employed producers who never seek advice.*
- C. *Agents in $[z_2, 1]$ are consultants.*
- D. *There is (strict) positive-assortative matching between producers in $[0, z_1]$ and consultants in $[z_2, 1]$.*

The diagram below summarizes parts A-C.



¹¹Indeed, the marginal value of knowledge in the hands of a consultant is equal to the number of producers he is matched with (i.e. $\frac{\partial}{\partial m} \Pi(m, z) = n(z) > 1$); the marginal value of knowledge in the hands of a self-employed producer is 1 (since the expected output of such a producer is z); and the marginal value of knowledge in the hands of a matched producer is no greater than 1 (the reason is that when a matched producer succeeds with dz higher probability, $h \cdot dz$ units of consultant time are freed, which the consultant in turn uses to solve at most dz alternative problems).

In order to match producers and consultants, the planner selects a function $M : \mathcal{Z} \rightarrow \mathcal{M}$ that assigns to each matched producer z a consultant $m = M(z)$. Positive-assortative matching means that M is an increasing function. When selecting M , the planner must satisfy a simple resource constraint. Fix the producer cutoff type z_1 . For any type $z \in \mathcal{Z}$, the total time required to consult on the problems left unsolved by producers in the interval $[z, z_1]$ must equal the total time available to consultants in the interval $[M(z), 1]$:

$$h \int_z^{z_1} (1-t) dF(t) = \int_{M(z)}^1 dF(m) \text{ for all } z \leq z_1. \quad (1)$$

For any given z_1 , there is a unique increasing function $M(z)$ that satisfies the above constraint, namely, the function satisfying $M(z_1) = 1$ and

$$M'(z) = h(1-z) \frac{f(z)}{f(M(z))} \text{ for all } z \leq z_1, \quad (2)$$

which follows from differentiating (1) with respect to z . Denote such function $M(z; z_1)$. Note that the consultant cutoff type z_2 is also uniquely determined by z_1 , namely, $z_2 = M(0; z_1)$.

Let z_1^* denote a surplus-maximizing value for z_1 (a function of both F and h) and let $z_2^* = M(0; z_1^*)$ denote the associated optimal value of z_2 . For notational simplicity, let $M^*(z) \equiv M(z; z_1^*)$ denote the corresponding optimal matching function and let $Z^*(m)$ denote the inverse of $M^*(z)$ (namely, consultant m is optimally matched with producer $z = Z^*(m)$).

Lemma 2 *There is a unique surplus-maximizing allocation. Depending on the value of the helping cost h , this allocation takes one of two forms, depending on whether or not having self-employed producers is efficient:*

1. An “interior solution” with $z_1^* < z_2^*$ (i.e. $\mathcal{I} \neq \emptyset$) and

$$\frac{1}{h} - z_2^* = \int_{z_2^*}^1 n(Z^*(t)) dt. \quad (3)$$

2. A “corner solution” with $z_1^* = z_2^*$ (i.e. $\mathcal{I} = \emptyset$) and

$$\frac{1}{h} - z_2^* \geq \int_{z_2^*}^1 n(Z^*(t)) dt. \quad (4)$$

Moreover, for any given F : (a) there exists a unique number $h_0(F) \in (0, 1)$ such that $z_1^* < z_2^*$ if and only if $h > h_0(F)$; and (b) z_1^* is decreasing in h .

Intuition is as follows. Part 1 corresponds to an interior solution for the cutoff z_1^* and equation (3) describes the corresponding first-order condition. This first-order condition has a simple interpretation. Consider a marginal increase in z_1^* that has the effect, owing to the planner's resource constraint, of transforming one more agent (namely, agent z_2^*) into a consultant. The L.H.S. measures the net gain along the extensive margin: on the one hand, the highest-type consultant is now free to solve $\frac{1}{h}$ unsolved problems of types marginally above z_1^* (who were original self-employed); on the other hand, agent z_2^* no longer produces her original level of surplus z_2^* . The R.H.S. measures the combined loss along the intensive margin: every single producer z below z_1^* is now matched with a marginally worse consultant (since every consultant is matched with a marginally better producer), which in turn lowers the output of producer z 's team by $\frac{\partial}{\partial m}\Pi(M^*(z), z) = n(z)$. Part 2 corresponds to a corner solution for z_1^* and inequality (4) states that, when increasing z_1 the net gain along the extensive margin is no lower than the combined loss along the intensive margin.

Finally, the last part of the Lemma implies that the efficient set of matched producers grows as the helping cost h falls. The reason is that, when offering advice is less time consuming, production in teams becomes more efficient relative to self-employment.

The First Welfare Theorem implies that when types are full information, a decentralized competitive equilibrium delivers the first best. Indeed, in Section 1 of the Supplement (available online) we show that, under full information, the first best is achieved under a variety of contractual arrangements, all of which lead to the same payoffs.¹² Once double-sided asymmetric information is introduced, in contrast, the set of contractual arrangements supporting the first best is significantly reduced.

For future reference, we describe some properties of the full-information competitive-equilibrium payoffs for the interior case with self-employment (see Figure 1 for an illustration). First, the marginal producer z_1^* and the marginal consultant z_2^* earn their autarky payoffs. Second, the payoffs of consultants grow with their type at rate $n(z)$, which is the marginal contribution of consultant skill to team output $\Pi(m, z)$. Third, the agents who

¹²The Supplement also shows that, by means of suitable contracts, the first best is achieved in a competitive equilibrium even when there is one-sided asymmetric information. For instance, in a "referral market" in which consultants are the residual claimants of output, the first best is implemented even when consultant types are unknown.




Figure 1: Full-information payoffs for F uniform and $h = 0.8$. (The equilibrium occupation thresholds are $z_1^* = 0.58$ and $z_2^* = 0.67$.)

obtain the highest gains from trade are those at the extremes of the type space.¹³

In what follows, we restrict attention to the interior case in which the first-best allocation involves self-employed producers – i.e. the case in which consultant time is sufficiently scarce. In the Supplement we extend our results to the corner case with no self-employed producers.

4 Two-sided adverse selection: optimal contracts

We now turn to the case in which the agents' knowledge levels are private information. In order to obtain more favorable contractual terms, agents may wish to misrepresent their knowledge. Consultants, on the one hand, may want to overstate their knowledge in order to exaggerate their contribution to team output. Producers, on the other hand, may want to either understate or overstate their knowledge depending on whether they contract ex-post (in which case higher knowledge is undesirable) or ex-ante (in which case higher knowledge is desirable).

¹³Indeed, the agents with the highest knowledge levels benefit from being able to leverage their knowledge by serving as consultants, whereas the agents with the lowest knowledge levels benefit from being able to pose problems that are relatively easy to solve. In contrast, agents with intermediate knowledge levels provide limited value as both producers (since they pose problems that are relatively hard to solve) and as consultants (since their knowledge is low relative to that of other consultants). For this reason, agents with intermediate knowledge obtain particularly low (even zero) gains from trade.

We are interested throughout in characterizing the family, or families, of contracts

$$\langle \omega(z), \alpha(z), \beta(z) \rangle_{z \in [0, z_1^*]}$$

that implement the first best. In what follows, to avoid notational clutter, let $M(\cdot)$ denote the first-best matching function $M^*(\cdot)$ and let $Z(\cdot)$ denote its inverse $Z^*(\cdot)$. In addition, let $\mathcal{Z} = [0, z_1^*]$, $\mathcal{I} = (z_1^*, z_2^*)$, and $\mathcal{M} = [z_2^*, 1]$.

We are interested in truth-telling mechanisms of the form $\langle \omega(z), \alpha(z), \beta(z), M(z) \rangle_{z \in \mathcal{Z}}$. A mechanism works as follows. After agents sort into occupations, they are invited to report their types to an abstract mechanism designer, with producers invited to report a type in $\mathcal{Z} \cup \mathcal{I}$ and consultants invited to report a type in \mathcal{M} . First, if a producer reports a type $z' \in \mathcal{Z}$ she is matched with consultant $M(z')$ under the bilateral contract $\langle \omega(z'), \alpha(z'), \beta(z') \rangle$. Second, if a producer reports a type $z' \in \mathcal{I}$ she is left unmatched (i.e. becomes self employed). Third, if a consultant reports type $m' \in \mathcal{M}$ she is matched with $n(z')$ producers of type $z' = Z(m')$ under bilateral contracts $\langle \omega(z'), \alpha(z'), \beta(z') \rangle$. Finally, any agent who does not participate in the mechanism is left unmatched.

We say that a family of contracts $\langle \omega(z), \alpha(z), \beta(z) \rangle_{z \in \mathcal{Z}}$ implements the first best if the mechanism $\langle \omega(z), \alpha(z), \beta(z), M(z) \rangle_{z \in \mathcal{Z}}$ induces truth telling once agents have sorted into their efficient occupations and, in addition, the mechanism induces agents to select their efficient occupations up-front.

The expected payoffs for agents who participate in the mechanism are as follows:¹⁴

1. A producer of type $z \in [0, 1]$ who reports type $z' \in \mathcal{Z}$ receives

$$R(z, z') = \omega(z') + z\beta(z') + (M(z') - z)(1 - \alpha(z')),$$

where z is the probability that the producer solves her problem on her own, and $M(z') - z = (1 - z) \frac{M(z') - z}{1 - z}$ is the probability that she seeks advice from consultant $M(z')$ and manages to solve her problem upon seeking such advice.

2. A producer of type $z \in [0, 1]$ who reports type $z' \in \mathcal{I}$ receives her self-employment payoff z .

¹⁴A non-participating producer of type z receives her self-employment payoff z and a non-participating consultant receives zero.

3. A consultant of type $m \in [0, 1]$ who reports type $m' \in \mathcal{M}$ receives

$$S(m, m') = n(z') [-\omega(z') + z' (1 - \beta(z')) + (m - z') \alpha(z')] \text{ for } z' = Z(m'),$$

where z' is the probability that each producer z' she is matched with solves her problem on her own, and $m - z' = (1 - z') \frac{m - z'}{1 - z'}$ is the probability that producer z' requires advice from consultant m and solves her problem conditional on receiving such advice.

For future reference, note that both $R(z, z')$ and $S(m, m')$ are linear in true type, with

$$R_z(z, z') = \alpha(z') + \beta(z') - 1 \text{ and } S_m(m, m') = \alpha(Z(m'))n(Z(m')). \quad (5)$$

Accordingly, R has increasing differences in z and $(\alpha(z') + \beta(z') - 1)$ and S has increasing differences in m and $\alpha(Z(m'))n(Z(m'))$.

With slight abuse of notation, let $R(z) = R(z, z)$ and $S(m) = S(m, m)$ denote, respectively, the payoffs of a producer in \mathcal{Z} and a consultant in \mathcal{M} who report their types truthfully. Note, finally, that R and S satisfy the condition that team surplus ($\Pi(M(z), z) = n(z)M(z)$) is fully allocated within each team:

$$S(M(z)) + n(z)R(z) = n(z)M(z) \text{ for all } z \in \mathcal{Z}. \quad (TS)$$

4.1 Incentive and occupational constraints

The following are necessary and sufficient conditions for a family of ex-ante contracts to implement the first best in the presence of self-employment:

First, the relevant incentive constraints *within* the efficient set of producers are met:

$$R(z) \geq R(z, z') \text{ for all } z, z' \in \mathcal{Z}. \quad (IC_1)$$

Second, the relevant incentive constraints *within* the efficient set of consultants are met:

$$S(m) \geq S(m, m') \text{ for all } m, m' \in \mathcal{M}. \quad (IC_2)$$

Third, and finally, every agent finds it optimal to select her first-best occupation:

$$R(z) \geq \max \left\{ z, \max_{m' \in \mathcal{M}} S(z, m') \right\} \text{ for all } z \in \mathcal{Z}, \quad (i)$$

$$S(m) \geq \max \left\{ m, \max_{z' \in \mathcal{Z}} R(m, z') \right\} \text{ for all } m \in \mathcal{M}, \quad (ii)$$

$$z \geq \max \left\{ \max_{z' \in \mathcal{Z}} R(z, z'), \max_{m' \in \mathcal{M}} S(z, m') \right\} \text{ for all } z \in \mathcal{I}. \quad (iii)$$

We refer to these three inequalities collectively as the *occupational-sorting constraints*.¹⁵

We now simplify the above five constraints following standard steps. We begin with the incentive constraints (IC_1) and (IC_2). Making use of the Envelope Theorem (e.g. Milgrom and Segal, 2002, Theorem 2), together with the derivatives of R and S in (5), these constraints are each reduced to two simple conditions:

On the one hand, the producer incentive constraints (IC_1) are met if and only if

$$R(z) = R(0) + \int_0^z \underbrace{[\alpha(t) + \beta(t) - 1]}_{R'(t)} dt \text{ for all } z \in \mathcal{Z}, \quad (E_R)$$

$$R'(z) = \alpha(z) + \beta(z) - 1 \text{ is nondecreasing in } z, \quad (M_R)$$

where (E_R) is the usual envelope condition that captures local incentive constraints and (M_R) is the relevant monotonicity condition that ensures that all incentive constraints are met globally (across producers in \mathcal{Z}).

On the other hand, the consultant incentive constraints (IC_2) are met if and only if

$$S(m) = S(z_2^*) + \int_{z_2^*}^m \underbrace{\alpha(Z(t))n(Z(t))}_{S'(t)} dt \text{ for all } m \in \mathcal{M}, \quad (E_S)$$

$$S'(m) = \alpha(Z(m))n(Z(m)) \text{ is nondecreasing in } m, \quad (M_S)$$

where, as before, (E_S) is the usual envelope condition that captures local incentive constraints and (M_S) is the relevant monotonicity condition that ensures that all incentive constraints

¹⁵The first inequality indicates that no matched producer wishes to become either a self-employed producer (and receive her autarky payoff z) or a consultant (and receive payoff $\max_{m'} S(z, m')$). The second inequality indicates that no consultant wishes to become either a self-employed producer (and receive her autarky payoff m) or a matched producer (and receive payoff $\max_{z'} R(m, z')$). The third inequality indicates that no self-employed producer wishes to become either a matched producer (and receive payoff $\max_{z'} R(z, z')$) or a consultant (and receive payoff $\max_{m'} S(z, m')$).

are met globally (across consultants in \mathcal{M}).

We now turn to the occupational-sorting constraints (i)-(iii). The following Remark, which relies on standard results, indicates that these constraints can be reduced to two equalities (those in (PC)) and two inequalities (those in (DD_R) and (DD_S)).

Remark 1 *Suppose the first best prescribes self-employment ($\mathcal{I} \neq \emptyset$) and suppose the family of contracts $\langle \omega(z), \alpha(z), \beta(z) \rangle_{z \in \mathcal{Z}}$ satisfies (IC₁) and (IC₂). Then, that family of contracts satisfies the occupational-sorting constraints (i)-(iii) if and only if:*

1. *The agents' equilibrium payoffs are continuous at z_1^* and z_2^* :*

$$R(z_1^*) = z_1^*, \quad S(z_2^*) = z_2^*. \quad (PC)$$

2. *The directional derivatives of R and S satisfy:*

$$R'(z_1^*-) \leq 1, \quad (DD_R)$$

$$1 \leq S'(z_2^*+), \quad (DD_S)$$

(where $R'(z_1^* -)$ denotes the left-hand derivative of R at z_1^* and $S'(z_2^* +)$ denotes the right-hand derivative of S at z_2^*).

Intuition is as follows. The agents' equilibrium payoffs are given by the function

$$V(z) = \begin{cases} R(z) & \text{for all } z \in \mathcal{Z}, \\ z & \text{for all } z \in \mathcal{I}, \\ S(z) & \text{for all } z \in \mathcal{M}. \end{cases}$$

Note that V is the upper envelope of the payoffs that agents can obtain under the family of available contracts, and, under any such contract, payoffs are continuous in type. It follows that V must be continuous in type, which delivers (PC). Moreover, it follows that at every point z the directional derivatives of V must satisfy $V'(z-) \leq V'(z+)$, which delivers (DD_R) and (DD_S).¹⁶

¹⁶See Milgrom and Segal (2002), Theorem 1. Indeed, if instead the payoff $R(z)$ of matched producers was steeper than the payoff z of self-employed producers at the cutoff type z_1^* , then agents marginally to the left of z_1^* would prefer to be self-employed. Similarly, if the payoff z of self-employed agents was steeper than the payoff of consultants at the cutoff type z_2^* , then agents marginally to the right of z_2^* would prefer to be self employed.

Conversely, (DD_R) and (DD_S) together with the monotonicity conditions ensure that $R(z)$ is (weakly) flatter than z and that $S(z)$ is (weakly) steeper than z . Accordingly, it suffices that the occupational sorting constraints are met at the occupational boundaries z_1^* and z_2^* , which is guaranteed by (PC) .

In summary, the relevant constraints are the within-occupation incentive constraints (captured by (E_R) , (M_R) , (E_s) , and (M_s)) and the occupational-sorting constraints (captured by (PC) , (DD_R) , and (DD_S)).

4.2 A necessary condition for optimality: consultants as residual claimants

We now prove a result that will become a central building block for the results that follow:

Lemma 3 *Suppose the first best prescribes self-employment ($\mathcal{I} \neq \emptyset$). Then, the family of contracts $\langle \omega(z), \alpha(z), \beta(z) \rangle_{z \in \mathcal{Z}}$ implements the first best **only if** $\alpha(z) = 1$ for every $z \in \mathcal{Z}$.*

Intuitively, granting full residual claim to consultants is the only way to prevent agents with low knowledge from entering the consulting occupation. To see why, note that the highest consultant $m = 1$ must receive *all* gains from trade when contracting with her producers z_1^* , as such producers are held to their reservation payoffs z_1^* . Accordingly, the payoff of $m = 1$ is pinned down at the highly-attractive level $\frac{1}{h}$ (which is the highest possible payoff that any agent in the economy can hope to get). Granting $m = 1$ such a high payoff causes a problem for the mechanism designer: since consultant payoffs are necessarily continuous in type, many agents will in principle wish to become consultants and earn payoffs similar to that of $m = 1$. As a result, the designer's problem becomes, de facto, one of maximally deterring self-employed producers from entering the consulting occupation. In order to achieve maximal deterrence, the designer must ensure that as a consultant's type m falls, $S(m)$ falls at the highest possible rate consistent with incentive compatibility. This goal is achieved by forcing consultants to internalize the full marginal contribution of their knowledge, namely, setting $\alpha(Z(m)) \equiv 1$.

Note that granting consultants the full residual claim of the problems in their teams implies that $S'(m) \equiv n(Z(m))$ (which is equal to the full marginal contribution of consultant knowledge to team output). This observation, combined with the fact that $S(z_2^*)$ is pinned down at z_2^* , implies that the only way to achieve the first best in the presence of self-

employment is to grant every consultant, and therefore every agent, an equilibrium payoff exactly equal to her full-information competitive payoff.

4.3 Ex-post (spot) contracting

We begin with the case in which contracting takes place in a spot market, after producers have attempted their problems. Only producers who have failed participate in this market. Recall that, formally, ex-post (spot) contracting is a special case of ex-ante contracting in which $\beta(z)$ is restricted to equal $1 - \omega(z)$ for all $z \in \mathcal{Z}$.

The following result shows that spot contracting can never achieve the first best in the presence of self-employment:

Proposition 1 *Suppose the first best prescribes self-employment ($\mathcal{I} \neq \emptyset$). Then, the first best cannot be supported by spot contracting.*

Intuitively, once consultants are granted the full residual claim of output, the “law of one price” implies that all unsolved problems must trade for the same (positive) price. As result, “lemons” (high-knowledge producers with hard unsolved problems) cannot be excluded from the market.

To see why, recall that the marginal consultant z_2^* is matched with $n(0) = \frac{1}{h}$ producers of type $z = 0$, and those producers always fail to solve their $\frac{1}{h}$ problems. Note, moreover, that consultant z_2^* must become the full residual claimant of all such unsolved problems (per Lemma 3), yielding a total expected output $\frac{1}{h}z_2^*$, and yet must earn a low expected payoff equal to her autarky payoff z_2^* (per Remark 1). Accordingly, under ex-post contracting, such arrangement would require that the consultant pay a positive fixed price $w(0)$ to her producers. In that case, however, every self-employed producer with an unsolved problem would enter the market by pretending to have knowledge $z = 0$, namely, by selling the full residual claim of her problem, at a positive price, to consultant z_2^* .¹⁷

As we show below, ex-ante contracting prevents such deviation from self-employed producers by forcing them to commit up front to share their output when succeeding on their

¹⁷This result opens the possibility that, in spot markets, lemons in the middle of the type distribution are likely to benefit from asymmetric information (at the cost of agents in the extremes). Indeed, it is plausible that such a result would hold in a broad class of models where agents with unknown skill levels self-select into occupations. While deriving second-best equilibria is beyond the scope of the present paper, exploring such possibility is a promising avenue for future work. We are grateful to an anonymous referee for this observation.




Figure 2: Optimal ex-ante contracts for F uniform and $h = 0.8$. (The equilibrium occupation thresholds are $z_1^* = 0.58$ and $z_2^* = 0.67$.)

own. Since such output sharing is more costly for producers with higher types, it successfully keeps the most talented producers away from the market, as required in the efficient allocation.

4.4 Ex-ante contracting

This section presents our main result: Theorem 1. This Theorem shows that, through an appropriate use of fixed wages and ex-ante output sharing, the first best can always be achieved. Moreover, the family of contracts that supports the first best is unique.

Theorem 1 *Suppose the first best prescribes self-employment ($\mathcal{I} \neq \emptyset$). A family of contracts implements the first best if and only if*

$$\alpha(z) = 1; \quad \beta(z) = hS(M(z)); \quad \text{and } \omega(z) = M(z) - hS(M(z)) \text{ for all } z \in \mathcal{Z},$$

where $S(M(z)) = z_2^* + \int_{z_2^*}^{M(z)} n(Z(t)) dt$.

These contracts are illustrated in Figure 2. Producers, on the one hand, have the option of either remaining self-employed or opting for a simple contract $\langle \omega(z), 1, \beta(z) \rangle$. Under such contract, a producer receives a guaranteed payment $\omega(z)$, which is decreasing in z , together

with a bonus $\beta(z)$ upon succeeding on her own, which is increasing in z .¹⁸ These contracts screen producers precisely because they become higher-powered as z grows. Indeed, as a producer's knowledge grows she is more willing to accept a lower wage in exchange for a higher contingent reward.¹⁹

Consultants, on the other hand, acquire the residual claim, net of bonus payments, of $n(z)$ problems, where z is the type of producers they opt to match with. In exchange, consultants offer each of the corresponding $n(z)$ producers an expected payment $\omega(z) + z\beta(z)$. These contracts screen consultants (who would prefer, other things equal, to be matched with more knowledgeable producers) by ensuring that the combined expected payment $n(z)[\omega(z) + z\beta(z)]$ to producers increases in z at a rate $n'(z)M(z)$ equal to the marginal contribution of producer knowledge to ex-ante team output, of which consultants are the residual claimants.

Note that consultants care only about the total expected payment $n(z)[\omega(z) + z\beta(z)]$ to producers, not about the way in which this payment is split between fixed and bonus payments. This property is what allows ex-ante contracts, with consultants as residual claimants, to simultaneously screen agents on both sides of the market.

The above contracts also guarantee that self-employed agents stay away from the market. That self-employed producers do not wish to become matched producers follows from the fact that being matched (under the pretense of having a type $z' \in \mathcal{Z}$) forces them to pay a share $1 - \beta(z')$ of output when succeeding on their own. Since such output sharing is, in expectation, more costly for producers with higher knowledge, it induces the producers with the highest knowledge levels (namely, self-employed producers) to remain unmatched. That self-employed producers do not wish to become consultants follows from the fact that, as consultants, they would become the full residual claimants of the unsolved problems they consult on, which they dislike because they are less capable, compared to existing consultants, of solving such problems.²⁰

¹⁸Notice that the lowest-type producers receive a positive wage despite the fact that they always need advice to solve their problems. The reason is that these producers are still the ones who carry out the production activities in their teams.

¹⁹Note that in the limit as z converges upward to z_1^* , producers receive zero fixed wages and the full residual value of their self-produced output, which is precisely the form of compensation received by self-employed producers.

²⁰Under a standard assumption on beliefs, the arrangement in Theorem 1 is robust to any given agent rejecting every ex-ante contract and instead seeking to trade ex-post. In particular, assume that if a producer rejects all ex-ante contracts the consultants in the economy infer that his type is between z_1^* and z_2^* (namely, he is considered to be among those types who chose to remain in autarky in equilibrium). On the one hand, a matched producer will not deviate because, ex-post, he cannot receive a positive price for his problem if

Finally, that the above contracts are necessary to implement the first best follows from Lemma 3 (namely, $\alpha(z) \equiv 1$ is necessary to deter self-employed agents from becoming consultants) and the occupational-sorting constraints (namely, the marginal consultant z_2^* must earn her autarky payoff). These two requirements pin down the payoffs of all consultants, and, consequently, the payoffs of all matched producers. Moreover, screening matched producers requires that their equilibrium payoff $R(z)$ grows at rate $R'(z) = \beta(z)$, which uniquely pins down the bonus schedule.

As we show in the Supplement (Section 2), when the first best calls for no self-employed producers (i.e. when helping costs h are sufficiently small), a contract identical to that in Theorem 1 (except for the payoff of the marginal producer/consultant) is guaranteed to implement the first best. The key difference vis-a-vis the present case is that the first best could in principle be implemented using other contractual arrangements as well.

Income inequality and variable pay

A testable prediction of the model is that, in the presence of double-sided asymmetric information, the level of income inequality $R'(z)$ across producers is directly reflected in the level of variable pay $\beta(z)$, with more inequality corresponding to higher variable pay.²¹ One could test this implication, for example, by using data on ‘non partners’ in professional service firms.

The fact that $R'(z) = \beta(z)$ also allows us to discuss some comparative statics. Suppose that the distribution F is uniform. On the one hand, if h falls, income inequality across all producers falls, in turn reducing the level of variable pay for each producer.²² This result is caused by a matching effect: when h drops, any two producers are matched with more similar consultants (since each consultant is now matched with more producers). As a result, the contribution of these two producers to surplus becomes more similar. On the other hand, if the distribution of knowledge in the economy shifts through a counter-clockwise rotation of $f(z)$ around $z = 1/2$ (thus leading to a larger mass of highly-knowledgeable agents),

unsolved. The reason is that any consultant can opt to contract ex-ante with producers of type z_1^* , which, given that $\omega(z_1^*) = 0$ and $\beta(z_1^*) = 1$, means that the consultant obtains $\frac{1}{h}$ unsolved problems for free, and these problems are weakly easier in expectation than those of the deviating producer. On the other hand, a consultant will not deviate because the best he can hope to achieve ex-post is to purchase $\frac{1}{h}$ unsolved problems for free from self-employed producers with types in (z_1^*, z_2^*) , and therefore contracting ex-ante with producers of type z_1^* (which means purchasing $\frac{1}{h}$ easier unsolved problems for free) is a better option.

²¹In contrast, if information was symmetric, producers could be compensated using an arbitrary mix of fixed and variable pay, as we illustrate in Section 1 of the online Supplement.

²²Indeed, when F is uniform, both $\beta(0)$ (which equals hz_2^*) and $\beta'(z)$ (which equals h) are increasing in h .

inequality across producers grows. The reason is that producer knowledge is more scarce and therefore becomes more expensive in the margin. (Numerical simulations for this result are available upon request.)

4.5 Empirical significance of the optimal contract

The optimal contracts derived above shed light on the ownership structure of professional service firms and on the contractual arrangements used in the venture capital industry. Here we discuss these empirical implications.

Consultants as endogenous principals

Under the optimal contracts, each team can be interpreted as being formed by a principal (the consultant) and $n(z)$ employees (the producers). The principal pays each employee a fixed wage plus a performance bonus, and, net of these payments, she is the residual claimant of the team's output. Thus, the optimal contracts have a simple firm-like flavor that is reminiscent of the contracts proposed by Alchian and Demsetz (1972). In Alchian and Demsetz, specialized monitors endogenously become principals. In contrast, in the present setting, consultants who specialize in providing advice, and are not directly involved in production, endogenously become principals. Note that these principals are the most knowledgeable individuals in the economy. This arrangement corresponds closely with what we observe in professional service firms, where indeed the most knowledgeable agents typically become principals/partners and reward the rest of the employees using a combination of fixed wages and bonuses.²³

Venture Capital and private equity

Our model captures various salient aspects of contracts used in venture capital settings. Venture capitalists (VCs) and other private equity investors (the consultants of our model) match with entrepreneurs (the producers of our model) and provide them with a mix of financing and knowledge. As in our setting, transactions are plagued by informational asymmetries on both sides. VCs ignore, a priori, the quality of the entrepreneur, and the entrepreneur ignores crucial characteristics of VCs, such as which other projects they are involved with and how useful their experience is for the specific project at stake.²⁴

²³To the best of our knowledge this feature has not been formally explained in previous work.

²⁴We are grateful to an anonymous referee for suggesting this application.

The most salient empirical facts about these relations seem in accordance with our model:

First, matching between VCs and entrepreneurs is positive assortative. In a comprehensive analysis of 1,666 investments made by 75 VCs between 1975 and 1995 in California and in Massachusetts, Sorensen (2007) finds strong evidence of such positive sorting: more experienced investors have access to higher-quality projects. He finds that an entrepreneur being of higher quality (a higher z in our model) explains 60% of the variance in a project's success, whereas the value added by venture capital investors (arising from a higher $M(z)$ in our model) explains another 40%.

Second, contracts are effectively structured like those predicted by Theorem 1. As Kaplan and Stromberg (2003) show in their analysis of 200 investment contracts used by VCs, private equity investors generally finance start-ups with a combination of equity and convertible debt (or other financial arrangements with equivalent properties). Investors give entrepreneurs cash up-front (ω in our model) and hold a fraction of equity if the entrepreneur succeeds ($1 - \beta$ in our model). Moreover, investors hold convertible debt, which in case the entrepreneur fails is converted into equity, and so investors receive all equity of the firm ($\alpha = 1$ in our model). A key implication of our model in this context is that higher-quality entrepreneurs should give up less equity $1 - \beta$. This feature has indeed been a systematic finding in the empirical literature, e.g. Kaplan and Stromberg (2004).

5 Extensions

In this section we consider two extensions: (1) the case of a monopolistic intermediary who matches producers and consultants; and (2) the case of contracting absent contingent pay for consultants. In both extensions, we maintain our assumption of double-sided asymmetric information.

5.1 A monopolistic intermediary

Consider a simple variation of the model in which all matching between producers and consultants is controlled by a monopolistic intermediary. For concreteness, we study an ex-post contracting environment in which, after agents have selected their occupations and producers have attempted their problems, the monopolist has the ability to purchase the residual claim of unsolved problems from producers and hire consultants to offer advice on

those problems. We allow the monopolist to offer contracts consisting of both fixed and output-contingent payments.

Proposition 2 *Suppose a monopolist controls all matching between producers and consultants. Then, the monopolist's profits are uniquely maximized at the first-best allocation.*

This result illustrates how an intermediary who hires consultants and finds unsolved problems (in this case a stylized monopoly) can solve the failure of spot markets in the presence of double-sided asymmetric information. Intuition is as follows. By purchasing unsolved problems for free and offering consultants output-contingent contracts that hold them to their outside options, the monopolist is able to keep all gains from trade for herself. As a result, she finds it in her interest to maximize such gains.

Note that this outcome differs from the standard nonlinear-pricing monopoly outcome in which quality is distorted in order to reduce consumer information rents (e.g. Mussa and Rosen, 1978, Maskin and Riley, 1984). The reason is that, in the present setting with output-contingent contracts, the expected transfers to consultants depend on both the consultants' reported types, as usual, as well as on their *true* types (since a consultant with a higher type is paid a given output-contingent bonus with higher probability). By exploiting this fact, the monopolist is able to extract all information rents from its clients without distorting the allocation of resources.²⁵

5.2 Market breakdown absent contingent pay

Here we illustrate the key role played by success-contingent consultant pay (through either output-contingent contracts or straight transfers of problem ownership). Specifically, we assume that consultants can only be paid a fixed fee (independent of success), but we otherwise allow for general random allocations, including rationing on either side of the market. Note that this is the first time we consider random matching. The reason is that, up to this point, our goal was to support the first best, which (per Lemma 1) calls for deterministic matching.

Proposition 3 shows that, absent contingent pay, there is a full market breakdown regardless of the potential gains from trade:

Proposition 3 *Suppose consultants cannot receive success-contingent pay. Then, there is no competitive equilibrium with a positive level of trade.*

²⁵We are grateful to an anonymous referee for suggesting this extension.

The intuition for this result is that, on the one hand, the expected earnings of a producer are increasing in her knowledge and, on the other hand, the expected earnings of a consultant are independent of her knowledge (owing precisely to the lack of contingent pay). As a result, the occupational-sorting incentives are fully misaligned: if an agent with a given knowledge level wishes to become a consultant, then all agents with lesser knowledge follow the same path.

Unlike in a standard lemons market (e.g. Akerlof, 1970), where unravelling from the top leads to various degrees of market failure (depending on the magnitude of the potential gains from trade), in the present setting occupational deviations from the lowest types guarantee a full market breakdown.

To illustrate, consider the market for brain surgeons. That very good surgeons may not be differentially compensated from merely good surgeons appears to be a second-order consideration when supporting trade. The first-order challenge arises from the average Joe donning a white robe and offering to crack your head open. Accordingly, certification can be particularly valuable in environments in which residual ownership cannot be transferred and contingent contracts cannot be written. Fuchs and Garicano (2010) take the present result as their point of departure and show how a simple certification process partially restores efficiency.²⁶

6 Related literature

The previous literature on consulting markets emphasizes moral hazard in the provision of consultant services. Demski and Sappington (1987) examine the trade-off between productive effort and information-gathering incentives faced by consultants. Wolinsky (1993) considers the type of incentives that must be provided for consultants to recommend a proper treatment, i.e. a minor treatment for minor problems and a major treatment for major problems. Wolinsky shows that specialization is optimal in such a setting. Similarly, Pesendorfer and Wolinsky (2003) study the provision of adequate diagnostic effort by consultants.²⁷ We depart from this literature in that we abstract from moral hazard altogether and focus instead on double-sided adverse selection.²⁸

²⁶See also early work by Leland (1979) and Shapiro (1986) on licensing.

²⁷Taylor (1995) studies how insurance can solve informational asymmetries in a context where only the consultant can determine the necessary treatment.

²⁸See also Ottaviani and Sorensen (2006), who study cheap talk by an expert who cares about his reputation, and Inderst and Ottaviani (2012), who study incentives in the form of commissions and kickbacks for

Garicano and Santos (2004) study a referral problem with one-sided adverse selection, where each production opportunity is always dealt with by a single individual. Agents’ skills are known, occupational choices are immaterial, and the agent who first receives an opportunity automatically learns its value. The efficient allocation involves low-skill agents tackling low-value opportunities and high-skill agents tackling high-value ones. Their model also features moral hazard and, indeed, their focus is on the trade-off between moral hazard and efficient problem allocation. In contrast, we study knowledge hierarchies, with two-sided adverse selection, where more than one agent could be involved in solving a particular problem (“management by exception”). Agents’ skills are unknown and they can choose their occupation. Moreover, the hierarchy structure exacerbates adverse selection as only the agents with problems that turned out to be hardest are in need of advice from high-skill agents. Finally, in Garicano and Santos (2004) the first best need not be achieved, whereas we show that a simple, firm-like contract always attains the first best.

Our paper is also related to the literature on trade in markets with bilateral asymmetric information. Most of this literature stems from Myerson and Satterthwaite’s (1983) seminal analysis of trade between buyers and sellers with private valuations (see also, for example, Lu and Robert, 2001). In this literature, in contrast to our setting, buyers and sellers do not care about each other’s types directly, players rely only on non-contingent payments, and matching considerations are not studied. The only prior paper we are aware of that studies equilibria in matching markets with two-sided adverse selection is Gale (2001). There are several important differences between our models. First, Gale takes as exogenous the side of the market agents belong to, while in our model agents select into different occupations. Second, in Gale’s paper all agents have equal outside options, while in our setting higher-quality agents have more attractive outside options, which further exacerbates the adverse selection problem. Third, Gale focuses on obtaining refinements that guarantee a fully-separating equilibrium. In contrast, we focus on characterizing contracts supporting the first best.

The paper is also related to the literature on management-worker sorting under full information, and, in particular, to Garicano and Rossi-Hansberg (2004, 2006). These models have been generalized by Eeckhout and Kircher (2012), who study general conditions under which interactions in production between worker and managerial skills generate skill-scale effects and positive sorting under full information. As in these papers, the economy-wide

advisors who act as intermediaries between consumers and firms.

problem we study here is one of matching talent with problems. In contrast to these papers, introducing informational asymmetries and optimal contracting allows us to study the contractual structure underpinning hierarchical team formation and, indeed, leads us to the finding that “firm-like” contracts are the unique ones that allow for such teams to be (optimally) formed under double-sided informational asymmetries. Our analysis is thus novel in its ability to characterize optimal contracting, which was immaterial in the settings above. We also offer a more general characterization of the first best with exogenous knowledge than previous work.

In independent and complementary work, Acemoglu, Mostagir, and Ozdaglar (2014) consider a problem of crowdsourcing. In their model, a principal who owns a set of unsolved problems contracts with agents (either his employees or external contractors) who attempt to solve these problems. They show that the first-best allocation of talent, which involves an endogenous hierarchy of problem solvers, is implementable regardless of informational asymmetries. Their approach differs from our own in two key respects. First, our technologies differ. In Acemoglu et al., in accordance to their focus on crowdsourcing, there is a limited supply of problems and agents cannot advise each other (instead they pass unsolved problems to one another). In our setting, in accordance to our focus on knowledge hierarchies with specialized advisors, the supply of problems is unlimited (i.e. all agents can work on their own problems if they so desire) and agents can communicate with each other, allowing the most knowledgeable among them to withdraw from direct production and instead specialize on advising others. Second, while in Acemoglu et al. the principal who owns all problems is exogenous, we obtain an endogenous principal who acts as an advisor and is optimally assigned the residual claim of output.

7 Conclusion

Over the last fifty years, the share of the service sector in the GDP of Western economies has grown substantially, e.g. in the USA by about 20%. Approximately half of this increase is due to growth in the share of professional and business services, such as computing, consulting, and legal services (Herrendorf, Rogerson, and Valentinyi, 2009). While production in the simplest manufacturing and agricultural sectors requires mostly labor and capital, production in the service sector is knowledge-intensive: capital equipment is largely irrelevant (as seen, for instance, in the law) and firms’ success comes down strictly to their ability to deploy

human capital appropriately.

In this paper, we study contractual arrangements in such knowledge-intensive industries. Leveraging expertise requires the use of knowledge-based hierarchies, where experts specialize in providing help to a number of non-experts. A crucial obstacle to such vertical specialization is asymmetric information – it is difficult to assess the knowledge levels of agents asking for, and those offering to, help. Can contracting resolve this problem?

We give a sharp answer to this question. In our model, there is a unique and simple arrangement that is guaranteed to organize knowledge efficiently. This arrangement is consistent with what we observe in practice: experts endogenously become residual claimants of output and pay a fixed wage and a performance bonus to the producers they are matched with. Moreover, producer incentives become higher-powered as their knowledge level grows. This arrangement simultaneously screens agents on both sides of the market because, on the one hand, consultants care only about their total payment to producers and, on the other hand, producers care about the way in which this total payment is split between fixed and bonus payments. Notably, these contracts are the only ones that deter self-employed agents from overstating their knowledge and becoming consultants, as they force them to internalize the full marginal value of their relatively-low knowledge.

This characterization illuminates contracting practices in the knowledge-intensive sectors of the economy, particularly inside professional service firms and between venture capital firms and entrepreneurs, as we have discussed above. Our analysis suggests that both sectors solve a matching problem under double-sided adverse selection by giving a full residual claim to experts on the one hand, and an appropriate incentive bonus (or equity share) to non-experts on the other hand.

We have focused throughout on contracting as a device to solve a problem of double-sided adverse-selection. In practice, firms and markets rely on various additional remedies including, notably, certification and reputation. A possible avenue for future work is studying how these different mechanisms interact. Finally, we have studied the value of ex-ante contracting while abstracting from the costs that such contracting is likely to entail relative to spot contracting, including the need for experts to contract up front with a number of suitable producers and to cope with a potentially-random number of questions asked by those producers. An additional possible avenue for future work is building a more comprehensive theory of the emergence of ex-ante contracting in the light of such costs.

8 Appendix: omitted proofs

This section contains omitted proofs.

Proof of Lemma 1. Part 1, which follows from the (strict) supermodularity of Π , is proven in a more general framework by Eeckhout and Kircher (2012) (see their Proposition 1 and Section 3.3.4).

For part 2 we must show that there is a cutoff z_2 such that (almost) all types $m > z_2$ are consultants and (almost) all types $z < z_2$ are producers. We proceed by contradiction. Suppose the desired claim does not hold – namely, there is a positive-measure subset of consultants \mathcal{M}_0 that lies below a positive-measure subset of producers \mathcal{Z}_0 . Select \mathcal{M}_0 and \mathcal{Z}_0 so that they have equal measures and they each fit in an interval of length $\varepsilon > 0$. Denote the measure of each of these sets Q . Note that no unsolved problem of a type z in \mathcal{Z}_0 is allocated to a consultant with a type m in \mathcal{M}_0 because that consultant’s time would be wasted.

Now (1) transform all types in \mathcal{Z}_0 into consultants and allocate to them the unsolved problems initially allocated to the original consultants in \mathcal{M}_0 and (2) transform all types in \mathcal{M}_0 into producers and allocate their unsolved problems to the consultants that initially received the unsolved problems of the original producers in \mathcal{Z}_0 (provided such consultants exist and only up to the time constraint of those consultants). Note that change (1) increases social surplus by at least $\frac{1}{h}(\inf \mathcal{Z}_0 - \sup \mathcal{M}_0)Q$ (which is a lower bound on how many more problems are solved by the new consultants in \mathcal{Z}_0 relative to the old consultants in \mathcal{M}_0) and change (2) reduces social surplus by no more than $(\sup \mathcal{Z}_0 - \inf \mathcal{M}_0)Q$ (which is an upper bound on how many fewer problems are solved by the new producers in \mathcal{M}_0 relative to the old producers in \mathcal{Z}_0 , which, given that some of those unsolved problems may be attempted by a consultant, constitutes in turn an upper bound on how many fewer problems attempted by the new producers in \mathcal{M}_0 are ultimately solved relative to the problems originally attempted by \mathcal{Z}_0). Since $\frac{1}{h} > 1$, the gain from (1) exceeds the loss from (2) for all small ε .

For part 3, let z_2 be such that almost all $z < z_2$ are producers. We must show that there is a cutoff $z_1 \leq z_2$ such that (almost) all producers $z \leq z_1$ seek advice and (almost) all $z \in (z_1, z_2)$ are self-employed producers. We proceed by contradiction. Suppose the desired claim does not hold – namely, there is a positive-measure subset of self-employed producers \mathcal{Z}_0 that lies below a positive-measure subset of matched producers \mathcal{Z}_1 . Select \mathcal{Z}_0 and \mathcal{Z}_1 so that both subsets of producers leave the same mass of unsolved problems. Now swap the roles of \mathcal{Z}_0 and \mathcal{Z}_1 . Since the unsolved problems of producers in \mathcal{Z}_0 are on average easier than

those of producers in \mathcal{Z}_1 , the associated consultants solve them with a higher probability while consuming the same amount of time, which increases social surplus. ■

Proof of Lemma 2. We begin by stating the planner's problem as a function of z_1 alone:

$$\begin{aligned} \max_{z_1} \int_0^{z_1} M(z; z_1) dF(z) + \int_{z_1}^{M(0; z_1)} z dF(z) \\ \text{s.t.} \\ z_1 \leq M(0; z_1). \end{aligned} \quad (6)$$

The objective is equal to social surplus. The first term is the total mass of solved problems among those problems initially attempted by the matched producers in $\mathcal{Z} = [0, z_1]$. The second term is the total mass of problems solved by the self-employed producers in $\mathcal{I} = [z_1, M(0; z_1)]$. Constraint 6 indicates that, per Corollary 1, z_1 cannot exceed $z_2 = M(0; z_1)$.

Now let \hat{z} denote the largest value of z_1 such that constraint (6) is met and replace that constraint with the equivalent constraint $z_1 \leq \hat{z}$. The planner's first-order condition is

$$\int_0^{\hat{z}} \frac{\partial}{\partial z_1} M(z; z_1^*) dF(z) + M(0; z_1^*) f(M(0; z_1^*)) \frac{\partial}{\partial z_1} M(0; z_1^*) + (1 - z_1^*) f(z_1^*) = \lambda,$$

where λ denotes the Lagrange multiplier for constraint $z_1 \leq \hat{z}$, and the planner's complementary slackness condition is $\lambda[\hat{z} - z_1] = 0$. Note, in addition, that

$$\frac{\partial}{\partial z_1} M(z; z_1) = -h(1 - z_1) \frac{f(z_1)}{f(M(z; z_1))} \text{ for all } z_1, \quad (7)$$

which follows from differentiating constraint (1) with respect to z_1 . The first-order condition therefore simplifies to²⁹

$$-\int_{M(0; z_1^*)}^1 n(Z^*(m)) dm - z_2^* + \frac{1}{h} = \lambda \frac{1}{h(1 - z_1^*) f(z_1^*)}.$$

²⁹To see why, note that from eq. (2) and integrating by substitution, $\int_{M(0; z_1^*)}^{M(x; z_1^*)} n(Z^*(m)) dm = \int_0^x \frac{f(z)}{f(M(z; z_1^*))} dz$ for all $x \leq z_1^*$, which in turn implies from eq. (7) that $\int_{M(0; z_1^*)}^{M(x; z_1^*)} n(Z^*(m)) dm = -\frac{1}{h(1 - z_1^*) f(z_1^*)} \int_0^x \frac{\partial}{\partial z_1} M(z; z_1^*) dF(z)$ for all $x \leq z_1^*$, and therefore $\int_{M(0; z_1^*)}^1 n(Z^*(m)) dm = -\frac{1}{h(1 - z_1^*) f(z_1^*)} \int_0^{z_1^*} \frac{\partial}{\partial z_1} M(z; z_1^*) dF(z)$.

Consequently, at the optimum, either: (1) $z_1^* < \widehat{z}$, $\lambda = 0$, and

$$z_2^* + \int_{z_2^*}^1 n(Z^*(m))dm = \frac{1}{h}; \quad (8)$$

or (2) $z_1^* = \widehat{z}$, $\lambda \geq 0$, and

$$z_2^* + \int_{z_2^*}^1 n(Z^*(m))dm \leq \frac{1}{h}.$$

As claimed in parts 1 and 2 of the Lemma.

That z_1^* is unique follows from equation (8) having a unique solution, which in turn follows from its L.H.S. being decreasing in z_1^* .³⁰ That z_1^* is a maximum, not a minimum, and is strictly positive, follows from noting that, in addition, the derivative of the planner's objective evaluated at $z_1 = 0$ is positive, namely,

$$\begin{aligned} M(0;0)f(M(0;0)) \frac{\partial}{\partial z_1} M(0;0) + f(0) = \\ (1-h)f(0) > 0, \end{aligned}$$

where the equality follows from (7) and the fact that $M(0;0) = 1$.

Finally, we turn to the last part of the Lemma. Note that, making use of (7), and rearranging terms, the marginal contribution of z_1 to social surplus is

$$(1-z_1)f(z_1) \left\{ 1 - h \left[\int_{M(0;z_1)}^1 n(Z(m; z_1))dm + M(0; z_1) \right] \right\} \quad (9)$$

(where $Z(\cdot; z_1)$ is the inverse of $M(\cdot; z_1)$). Note that the expression in square brackets is greater than 1 for all $z_1 \in (0, \widehat{z}]$ (since $n(Z(m; z_1)) > 1$ for all $m \in [M(0; z_1), 1]$). It follows that, for any $z_1 > 0$, expression (9) is decreasing in h and can be either positive or negative depending on the value of $h \in (0, 1)$. It follows that, whenever constraint (6) is slack, z_1^* is decreasing in h and therefore the desired cutoff $h_0(F) \in (0, 1)$ exists. Moreover, when

³⁰To see why note that

$$\frac{d}{dz_1^*} \left[z_2^* + \int_{z_2^*}^1 n(Z^*(m))dm \right] = \frac{d}{dz_1^*} M(0; z_1^*) (1 - n(Z^*(z_2^*))) + \int_{z_2^*}^1 \frac{d}{dz_1^*} n(Z^*(m))dm.$$

The first term on the R.H.S. is positive because $\frac{d}{dz_1^*} M(0; z_1^*) < 0$ and $n(Z^*(z_2^*)) = n(0) = \frac{1}{h} > 1$. The second term on the R.H.S. is also positive because $\frac{d}{dz_1^*} n(Z^*(m)) > 0$ for all m , as increasing z_1^* implies that each consultant is matched with producers of a higher type and so can be matched with a greater number of such producers.

constraint (6) binds, z_1^* is decreasing in h owing to the time constraint (1) (i.e. when h falls a given set of consultants can help a larger set of producers). ■

Proof of Remark 1. Necessity (\Rightarrow). (PC) follows from the fact that the agents' equilibrium payoffs must be continuous in their type (a straightforward implication of local incentive compatibility). (DD_R) and (DD_S) follow from Theorem 1 in Milgrom and Segal (2002).

Sufficiency (\Leftarrow). Conditions (M_R), (M_S), (DD_R), and (DD_S) ensure that $R'(z) \leq 1$ and $S(z) \geq 1$. It therefore suffices that the occupational-sorting constraints are met at the occupational boundaries z_1^* and z_2^* , which in turn is guaranteed by (PC). ■

Proof of Lemma 3. On the one hand, from Lemma 2, the planner's first-order condition implies that

$$\frac{1}{h} - z_2^* = \int_{z_2^*}^1 n(Z(t))dt. \quad (a)$$

On the other hand, the consultant envelope condition (E_S), combined with the fact that $S(z_2^*) = z_2^*$, implies that

$$S(1) - z_2^* = \int_{z_2^*}^1 \alpha(Z(t))n(Z(t))dt. \quad (b)$$

From (TS), and the fact that $R(z_1^*) = z_1^*$, it follows that $S(1) = \frac{1}{h}$. Consequently, given that $\alpha(z) \in [0, 1]$ for all z , equations (a) and (b) can only be simultaneously satisfied if $\alpha(z) = 1$ for almost every z . That $\alpha(z) = 1$ for every z then follows from the consultant monotonicity constraint (M_S). ■

Proof of Proposition 1. Suppose toward a contradiction that the family of ex-post contracts $\langle w(z), \alpha(z) \rangle_{z \in \mathcal{Z}}$ implements the first best. First, note that consultant z_2^* is matched with $n(0) = \frac{1}{h}$ producers of type $z = 0$ and that team generates output $\Pi(z_2^*, 0) = \frac{1}{h}z_2^*$. It follows from (TS) that

$$S(z_2^*) + \frac{1}{h}R(0) = \frac{1}{h}z_2^*.$$

Second, recall that Remark 1 requires that $S(z_2^*) = z_2^*$ and Lemma 3 requires that $\alpha(0) = 1$. It follows that $R(0) = w(0)$ and so the above equality simplifies to

$$z_2^* + \frac{1}{h}w(0) = \frac{1}{h}z_2^*,$$

which implies that $w(0) > 0$.

Finally, given the existence of contract $\langle w(0), \alpha(0) \rangle = \langle w(0), 1 \rangle$, every self-employed producer $z \in \mathcal{I}$, upon failing to solve her problem, would gain by pretending to be of type 0

and thus selling her unsolved problem to a consultant of type z_2^* , a contradiction. ■

Proof of Theorem 1. Necessity (\Rightarrow). Suppose the family of contracts $\langle \omega(z), \alpha(z), \beta(z) \rangle_{z \in \mathcal{Z}}$ implements the first best. First, from Lemma 3, we obtain $\alpha(z) = 1$ for all $z \in \mathcal{Z}$. Second, by differentiating (TS) , namely, $R(z) = M(z) - \frac{1}{n(z)}S(M(z))$, with respect to z we obtain

$$R'(z) = M'(z)[1 - \alpha(z)] + hS(M(z)) \text{ for all } z \in \mathcal{Z}.$$

Since $R'(z) = \alpha(z) + \beta(z) - 1$ and $\alpha(z) = 1$, the above equality delivers $\beta(z) = hS(M(z))$ for all $z \in \mathcal{Z}$. Third, and finally, (TS) evaluated at the above values for $\alpha(z)$ and $\beta(z)$ implies that

$$\omega(z) + zhS(M(z)) = M(z) - h(1 - z)S(M(z)) \text{ for all } z \in \mathcal{Z},$$

which after rearranging terms delivers $\omega(z) = M(z) - hS(M(z))$ for all $z \in \mathcal{Z}$.

Sufficiency (\Leftarrow). We need to show that the family of contracts in the Theorem meets the seven relevant constraints, namely, (E_R) , (M_R) , (E_S) , (M_S) , (PC) , (DD_R) , and (DD_S) . Note that $R(z) = M(z) - (1 - z)hS(M(z))$ (from (TS)) and so $R'(z) = hS(M(z))$.

(E_R) is met because $R'(z) = hS(M(z)) = \beta(z)$. (M_R) is met because $S(M(z))$ is increasing in z . (E_S) is met from the definition of $S(M(z))$. (M_S) is met because $S'(m) = n(Z(m))$ is increasing in m . (PC) is met because, on the one hand, from (TS) ,

$$R(z_1^*) = M(z_1^*) - h(1 - z_1^*)S(M(z_1^*)) = z_1^*,$$

where the second equality follows from the fact that $M(z_1^*) = 1$ and $S(1) = \frac{1}{h}$ (owing to the definition of S and Lemma 2, part 1), and, on the other hand, $S(z_2^*) = z_2^*$ by definition. Finally, (DD_R) and (DD_S) are met because, on the one hand, $R'(z_1^*) = hS(1) = 1$, and, on the other hand, $S'(z_2^*) = n(Z(z_2^*)) > 1$. ■

Proof of Proposition 2. It suffices to show that the monopolist can implement the first best and, simultaneously, obtain all gains from trade (namely, hold all agents to their self-employment payoffs). Assume that all agents sort into their first-best occupations (we will show that, given the contracts that follow, such choices are indeed optimal). First, let the monopolist offer to purchase the residual claim of all unsolved problems from producers $z \in [0, z_1^*]$ for free while at the same time asking those producers to truthfully reveal their types. As producers have nothing to lose, they agree to both. Second, let the monopolist offer consultants a menu of contracts $\langle \tau_0(m), \tau_1(m), Z^*(m) \rangle_{m \in [z_2^*, 1]}$, where $\tau_0(m)$ is a fixed

payment, $\tau_1(m)$ is a bonus per problem solved, and $Z^*(m)$ is the producer match enjoyed by a consultant who reports to be of type m . Now set

$$\tau_0(m) = Z^*(m) \text{ and } \tau_1(m) = h(1 - Z^*(m)) \text{ for all } m \in [z_2^*, 1].$$

Note that every consultant $m \in [z_2^*, 1]$ earns her self-employment payoff m regardless of her report. As a result, all consultants agree to report their types truthfully. Since both producers and consultants report their types truthfully, the monopolist can indeed implement the first-best match $Z^*(\cdot)$. Finally, note that under such contracts, all agents earn their self-employment payoffs regardless of their career choices and reports, and therefore they find it optimal to select their first-best occupations from the outset. ■

Proof of Proposition 3. Assume toward a contradiction that there exists an equilibrium with a positive mass of consultants. Let \mathcal{M}^* and \mathcal{Z}^* denote, respectively, the equilibrium set of consultants and the equilibrium set of producers matched with those consultants, let $\mu(\mathcal{Z}^*)$ denote the measure of \mathcal{Z}^* , and let $m^* = \sup \mathcal{M}^* > 0$ denote the highest consultant type. We proceed in five steps. First, since consultant types are unobserved and consultant earnings are not contingent on their success, all consultants must earn the same expected payoff. Denote such payoff S^* . Second, note that $S^* \geq m$ for all consultants $m \in \mathcal{M}^*$ (as the outside option of a type m consultant is at least m) and therefore $S^* \geq m^*$. Third, as any agent may become a consultant, the equilibrium payoff of all producers is bounded below by S^* . Fourth, the combined (ex-ante) output produced by all matched agents ($\mathcal{M}^* \cup \mathcal{Z}^*$) is no larger than $\mu(\mathcal{Z}^*) \cdot m^*$ (as $\mu(\mathcal{Z}^*)$ is the mass of attempted problems and m^* is an upper bound for the probability that a given problem is solved). Fifth, since each producer must earn, in expectation, at least $S^* \geq m^*$, the combined earnings of all producers \mathcal{Z}^* must be at least $\mu(\mathcal{Z}^*) \cdot m^*$. It follows that the combined earnings of all consultants \mathcal{M}^* must be no greater than zero, a contradiction. ■

REFERENCES

- ACEMOGLU D., MOSTAGIR, M. and OZDAGLAR A. (2014), “Managing Innovation in a Crowd”, (Mimeo).
- AKERLOF, G. A. (1970), “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism”, *Quarterly Journal of Economics*, 84 (3), 488-500.

- ALCHIAN, A. and DEMSETZ, H. (1972), "Production, information costs, and economic organization", *American Economic Review*, 62 (5), 777-795.
- DEMSKI, J. S. and SAPPINGTON, D. E. (1987), "Delegated Expertise", *Journal of Accounting Research*, 25 (1), 68-89.
- EECKHOUT, J. and KIRCHER, P. (2012), "Assortative Matching with Large Firms: Span of Control over More versus Better Workers", (Mimeo).
- ESWARAN, M. and KOTWAL, A. (1984), "The Moral Hazard of Budget-Breaking", *Rand Journal of Economics*, 15 (4), 578-581.
- FUCHS, W. and GARICANO, L. (2010), "Matching Problems with Expertise in Firms and Markets", *Journal of the European Economic Association*, 8 (2-3), 354-364.
- GALE, D. (2001), "Signalling in Markets with two-sided adverse selection", *Economic Theory*, 18 (2), 391-414.
- GARICANO, L. (2000), "Hierarchies and the Organization of Knowledge in Production", *Journal of Political Economy*, 108 (5), 874-904.
- GARICANO, L. and ROSSI-HANSBERG, E. (2004), "Inequality and the Organization of Knowledge", *American Economic Review*, 94 (2), 197-202.
- GARICANO, L. and ROSSI-HANSBERG, E. (2006), "Organization and Inequality in a Knowledge Economy", *Quarterly Journal of Economics*, 121 (4), 1383-1436.
- GARICANO, L. and SANTOS, T. (2004), "Referrals", *American Economic Review*, 94 (3), 499-525.
- HAYEK, F. A. (1945), "The Use of Knowledge in Society", *American Economic Review*, 35 (4), 519-530.
- HERRENDORF, B., ROGERSON, R. and VALENTINYI, Á. (2009), "Two Perspectives on Preferences and Structural Transformation", NBER Working Papers 15416.
- INDERST, R. and OTTAVIANI, M. (2012), "Competition through commissions and kick Backs", *American Economic Review*, 102 (2), 780-809.

- KAPLAN, S. N. and STROMBERG, P. (2003), "Financial contracting theory meets the real world: An empirical analysis of venture capital contracts", *Review of Economic Studies*, 70 (2), 281-315.
- KAPLAN, S. N. and STROMBERG, P. (2004), "Characteristics, contracts, and actions: Evidence from venture capitalist analyses", *Journal of Finance*, 59 (5), 2177-2210.
- KLEINER, M. M and KUDRLE, R. T. (2000), "Does Regulation Affect Economic Outcomes? The Case of Dentistry", *Journal of Law and Economics*, 43 (2), 547-582.
- LELAND, H. E. (1979), "Quacks, Lemons, and Licensing: A Theory of Minimum Quality Standards", *Journal of Political Economy*, 87 (6), 1328-46.
- LU, H. and ROBERT, J. (2001), "Optimal Trading Mechanisms with Ex Ante Unidentified Traders", *Journal of Economic Theory*, 97 (1), 50-80.
- MASKIN, E. and RILEY, J. (1984), "Monopoly with incomplete information", *Rand Journal of Economics*, 15 (2), 171-196.
- MILGROM, P. and SEGAL, I. (2002), "Envelope Theorems for Arbitrary Choice Sets", *Econometrica*, 70 (2), 583-601.
- MUSSA, M. and ROSEN, S. (1978), "Monopoly and product quality", *Journal of Economic Theory*, 18 (2), 301-317.
- MYERSON, R. and SATTERTHWAITTE, M. (1983), "Efficient mechanisms for bilateral trading", *Journal of Economic Theory*, 28 (2), 265-281.
- OTTAVIANI, M. and SORENSEN, P. N. (2006), "Professional Advice", *Journal of Economic Theory*, 126 (1), 120-142.
- PENSENDORFER, W. and WOLINSKY, A. (2003), "Second Opinions and Price Competition: Inefficiency in the Market for consultant Advice", *Review of Economic Studies* 70 (2), 417-437.
- SHAPIRO, C. (1986), "Investment, Moral Hazard and Occupational Licensing", *Review of Economic Studies*, 53 (5), 843-862.
- SORENSEN, M. (2007), "How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital", *Journal of Finance*, 62 (6), 2725-2762.

SUN, Y. (2006), “The exact law of large numbers via Fubini extension and characterization of insurable risks”, *Journal of Economic Theory*, 126 (1), 31-69.

TAYLOR, C. (1995), “The Economics of Breakdowns, Checkups and Cures”, *Journal of Political Economy*, 103 (1), 53-74.

WOLINSKY, A. (1993), “Competition in a Market for Informed consultants Services”, *Rand Journal of Economics*, 24 (3), 380-398.