

Imperfect Monitoring and Fixed Spreads in the Market for IPOs*

Richard Lowery

Carnegie Mellon University

Tepper School of Business

PRELIMINARY AND INCOMPLETE

Draft 4: December 23, 2008

Abstract

Characteristics of the investment banking industry, particularly the extreme concentration of spreads at exactly 7%, seem consistent with some form of collusion through which underwriters can extract surplus from the IPO. I present a model of investment banking that, under the assumption of optimal collusion, generates a distribution of spreads qualitatively similar to that observed. The model is extended to show that underpricing and spread rigidity may arise together, each one reinforcing incentives to engage in the other.

*This paper has benefited from comments from Juan Dubra, Rick Green, Burton Hollifield, Bryan Routledge, Bilge Yilmaz, and participants at the UNC/Duke Corporate Finance Conference. I have also received very useful guidance from Spencer Martin and Sevin Yeltekin. All errors are my own.

1 Introduction

In the U.S., over 95% of privately held firms of intermediate value pay a spread of exactly 7% to an underwriter when they decide to go public (Chen and Ritter [2000]¹). Very large firms pay lower spreads and very small firms higher. Since the total payment for investment banking services is determined by the spread and the proceeds of the offering, this rigidity is surprising. Costs to an underwriter almost certainly contain a fixed component. Either competition or efficient collusion between underwriters should thus lead to spreads that fall with the size of the offering.

A debate has developed in the corporate finance literature as to whether this extreme spread rigidity is evidence of collusive behavior by underwriters. Chen and Ritter [2000], who discovered the pervasiveness of the seven percent spread, argue informally that it is evidence of collusion, while Hansen [2001] argues that coordination on 7% naturally arises in an efficient contract. The question of whether underwriters collude on spread offers is important. Standard concerns about inefficient provision of goods and services in monopolistic or collusive industries apply, but, more importantly, if underwriters can collude to extract profits that would otherwise have accrued to entrepreneurs and venture capitalists, the incentives to engage in risky but positive expected value projects are diminished. If, on the other hand, underwriters are competitive in pricing underwriting services, then incentives for entrepreneurial activity will be appropriately aligned.

I present a formal model of the IPO process as an infinitely repeated game of imperfect information with public monitoring. In this context, I show that optimal collusion by underwriters will lead to spreads qualitatively similar to those observed, while competition or monopoly provision will imply that spreads depend on the size of the firm over the entire distribution of firm values.

I assume firms have two incentives to go public. Going public increases the present value of the firms expected future earnings. This “common value” element of the IPO process likely follows from improved access to capital markets.² Going public also provides private benefits (or

¹See also Jay Ritter’s website for more current data.

²Chemmanur and Fulghieri [1999] treat this as the primary motivation for going public, and Hale and Santos [2006] present direct evidence that having held an IPO for *debt* reduces the interest rate charged on bank loans

implies private costs) to the managers of the firms, who are ultimately responsible for choosing whether or not to go public. This private element would follow primarily from preferences for personal liquidity,³ which will vary across the owners or managers of different firms.

In the model, underwriters collude to extract the most profits possible from issuing firms. Underwriters exploit their repeated interactions in the IPO market to maintain a pricing strategy that provides greater profits than one-shot competition. They cannot, however, collude on the spread schedule that maximizes total profits, even conditional on the information that they receive about the value of the firm and the preferences of the managers. Because underwriters receive different signals about the value of the firm, colluding on a fully efficient spread schedule becomes a game of imperfect monitoring and consequently bears costs associated with on-equilibrium-path punishments. I show that underwriters can, under certain circumstances, improve profits by using a partially rigid spread that calls for high spreads for small firms and a lower, uniform spread for intermediate and large sized firms. I also argue that the results on cyclical pricing in Rotemberg and Saloner [1986] explain why very large firms are charged lower spreads.

In addition to demonstrating that collusion is a likely explanation for spread rigidity, the model also suggests links between unobserved variables in the investment banking industry, such as the costs of providing underwriting services and the effect of IPOs on the value of firms, and the observed characteristics of the distribution of spreads. While estimating the parameters of the model is beyond the scope of the present work, such an exercise could prove useful in quantifying the costs of collusion and investigating the welfare effects of anti-collusive policies.

The idea that rigidity in pricing may make collusion easier has been understood informally in the industrial organization literature for some time. More recently, Athey et al. [2004] developed a model of collusion in Bertrand competition and show that under asymmetric and imperfect signals about firm costs, rigid pricing schemes may be optimal. The techniques applied rely crucially on the assumption that players receive private information about their own costs, a *private* value component of payoffs. Their approach does not readily generalize to cases where

and private bond issues for a firm. Pagano et al. [1998] and Bharath and Dittmar [2007] also provide useful discussions of the motivations for a firm to go public from the “common value” perspective.

³See Ritter [2003].

participants in the market are differentially informed but symmetric in all other respects. The model of Hanazono and Yang [2007] introduces private information over a common value aspect of payoffs, specifically the state of demand, and shows that in a simple two-state environment price rigidity can also arise under optimal symmetric collusion. While the tradeoffs between price rigidity and price flexibility in this model are similar to those that arise in the auction setting that I consider, the restriction to two states prevents their model from addressing the extent of rigidity, and in particular it cannot be used to investigate the partially rigid schedule observed in the cross-section of IPO spreads.

The paper proceeds as follows. I first present a model of the underwriting process. I then explore the likelihood that spreads will exhibit rigidity when underwriters behave competitively, as a competitive oligopoly, or as a monopoly. Under each of these assumptions, price rigidity will only be observed in very special cases that I argue are unlikely to characterize the underwriting industry.

I then show that partially rigid spreads similar to those observed in the data can arise when underwriters receive different but informative signals about firm value and play symmetric strategies in each period. Partially rigid spreads are *strictly* preferred to any fully rigid or fully flexible (i.e. everywhere downward sloping) spread and so the conclusion that partially rigid spreads can arise under collusion with imperfect information is robust to small perturbations of any element of the environment considered.

Next, I discuss an extension of the model to permit underpricing and argue that spread rigidity and underpricing will reinforce each other. When underwriters collude on a rigid spread to reduce the costs of monitoring collusive agreements they will be unable to extract all possible rents from those firms that choose to go public. By introducing an “offer price” stage *after* the firm has committed to one underwriter, I show that underpricing has a natural role as a means to capture more of these rents.

I also address the marked difference between IPO spreads, where rigidity is pervasive, and SEO spreads, where rigidity is nonexistent. In the SEO market the restriction to symmetric equilibria is shown to be inappropriate, which in turn permits efficient, and therefore non-rigid, spread pricing.

The paper concludes with a brief discussion of international differences in the market for underwriting services and the implications for regulatory policy that follow from the analysis here.

2 Model

This section presents a model of IPO underwriting. The model admits heterogeneity in firm value and manager preferences for going public. Underwriters receive (potentially noisy and private) signals about these quantities before bidding for the right to take a firm public. The model therefore allows for the possibility of a relationship between firm value and IPO spreads and can thus provide insight into the source of spread rigidity over a wide range of firm values.

Firms have two incentives to go public. First, there is a “common value” element of going public; the net present value of a firm increases when it goes public. Second, managers of firms have idiosyncratic private incentives to take their firms public (or to keep their firms private). A detailed examination of the source of the “common value” or “private value” benefits of taking a firm public is beyond the scope of this paper. The common value benefits most likely relate to improved access to capital markets, while private benefits can be thought of as representing a tradeoff between preferences for control and preferences for personal liquidity.

Institutional constraints prevent a firm from going public without the assistance of one of the underwriters. In taking a firm public, an underwriter incurs a fixed cost, which implies that returns to scale are increasing. Certain costs associated with underwriting, particularly those involving accounting tasks, might increase with the value of the firm going public. Many costs, however, will have a fixed component; for example, the costs of certain legal and regulatory tasks would not vary much with firm value. Other costs, such as those associated with advertising the issue, might actually fall in firm value as larger firms would tend to be more well known and therefore require less effort to elicit interest during the book-building process. Post-issue stabilization, a potentially significant cost to underwriters, also has an ambiguous relationship to firm value since the expected magnitude of such activity most likely depends both on the size of the issue and on the riskiness of the price from the perspective of the underwriter (Prabhala

and Puri [1998]), which may be decreasing in the value of the firm.⁴

Underwriters compete for a sequence of opportunities to take firms public by making simultaneous spread offers after they receive their signal about the firms, where the spread determines the proportion of the common value of the firm that the underwriter receives as fees for facilitating the IPO; the IPO market is thus modeled as a repeated procurement auction with security bids (see DeMarzo et al. [2005] for a general treatment of auctions with security bids). I assume that the offer price of the issue is exactly the true value; that is, book building perfectly reveals the value of the firm and underwriters have no incentive to underprice offerings. Section 5 addresses the role of underpricing in the model considered here. The repeated setting captures the presence of the established investment banks which dominate the market for IPO services. The entire structure of the auction is exogenous. This abstracts away from considerations about optimal security design for IPOs, which has been explored, for example, in Chakraborty et al. [2008]. Optimal design in the presence of collusion is left for future research.

Following the bidding for each firm, underwriters learn the true common value of the firm and the winning bid is made public. This additional information, which becomes common knowledge among firms, captures a crucial element of the IPO process; before going public, information about a firm is disaggregate and held privately, whereas after the IPO there is a perfectly observable summary statistic aggregating such information, the trading price. This public signal allows the repeated auction to be modeled as a game of imperfect public monitoring. By confining attention to strategies conditioned on current private information and the history of the public signal, the repeated interaction can be recast in a recursive structure and analyzed using standard techniques developed in Abreu et al. [1990].

Formally, there is a sequence of short-lived firms and two long lived investment banks who discount the future at rate δ . Each firm has type $\{x_t, \varepsilon_t\}$, where x_t is value and ε_t summarizes the idiosyncratic preferences of the owners and managers of the firm. Each underwriter, indexed by i , observes a potentially two dimensional signal $\{\xi_t^i, \eta_t^i\}$, where ξ_t^i is informative about x_t and η_t^i is informative about ε_t . Underwriters simultaneously submit spread offers $\alpha_t^i \in [0, \infty]$.

⁴The direction of the change in the costs of stabilization is an empirical question that, to my knowledge, has not been addressed. Aggarwal [2000] provides a method for determining the extent of after-market activities by underwriters, and could likely be extended to address this question.

⁵⁶ The firm then chooses between the two underwriters, or decides to remain private.

The payoffs in the stage game can then be summarized as follows:

- If the firm goes public using underwriter i :

firm: $(1 - \alpha_t^i)\beta x_t + \varepsilon_t$

underwriter i : $\alpha_t^i\beta x_t - \kappa$

underwriter j : 0

- If the firm does not go public:

firm: x_t

underwriter i : 0

underwriter j : 0

Here, β summarizes the common value benefits of going public. Since the IPO provides a proportional increase in value, going public is more valuable to large firms from an absolute perspective.

2.1 Public History and Equilibrium

After each stage of the game, the true value x_t and a public signal $a_t = \{\min_{i \in \{1,2\}} \alpha_t^i\}$ are revealed publicly. For simplicity, I assume that the lowest spread offer is always reported publicly, regardless of whether the firm decides to go public. See Figure 1 for a summary of the timing of the stage game.

[Figure 1 about here.]

⁵For expositional clarity, throughout the paper I treat the action space as a continuum. Because I will use the bang-bang result from Abreu et al. [1986] and Abreu et al. [1990], formally we must consider arbitrarily fine approximations to the continuous stage game as bang-bang has not, to my knowledge, been shown to apply to games with continuous action spaces. All of the arguments showing that the results for the continuous case also apply to the limit of an appropriate discretization of the game are collected in appendix A.

⁶It is obviously unlikely to observe a bid greater than 1 as that would require payments in excess of the total value of the firm. Since some of the benefits of going public are not directly related to the increase in the value of the firm, however, it is possible that managers of very small firms would pay more than the entire value of the firm in order to realize the private benefits of going public. For simplicity, I do not rule out even such obviously pathological examples.

I will analyze the game by restricting attention to “almost pure strategy symmetric perfect public equilibria” where this set of strategies is defined as follows:

Definition A *perfect public equilibrium* is a profile of public strategies that, for any public history, specifies a Nash equilibrium for the repeated game starting at that history. A perfect public equilibrium is *symmetric* if all players use the same strategy profile following every history.

The restriction to symmetric equilibria is not innocuous and provides significant restrictions on the strategies available to the underwriters. Fully efficient collusion in Bertrand competition (Athey and Bagwell [2001] and Kandori and Matsushima [1998]) and in auctions (Aoyagi [2003] and Aoyagi [2007]) generally relies on exploiting asymmetric bid rotation schemes and communication. The restriction to strongly symmetric strategies renders the techniques exploited in these papers out of bounds in my application. In fact, permitting even simple bid rotation schemes with pre-play communication will immediately permit underwriters to collude efficiently in a pure common values setting, which will imply spreads that do depend on firm value. I interpret the restriction to strongly symmetric equilibria as a restriction to collusive strategies that are not easily detectable by regulators. Since firms are effectively anonymous to underwriters, asymmetric collusion would require bid rotation. Anti-trust regulators are experienced in identifying bidding rings of this type, and maintaining such coordination in a more complex setting than described here would likely require either an extremely sophisticated initial agreement or regular cartel discussions to resolve questions and assure continued adherence. Such discussions are the primary means of identifying and prosecuting collusion, and thus the restriction to strongly symmetric equilibria makes sense within the legal framework faced by underwriters. The difference between symmetric and asymmetric collusion also provides a useful way to understand the differences between spreads charged on IPOs and spreads charged on SEOs, which do not cluster to any meaningful extent. This comparison is explored in section 6.

Definition An *almost pure strategy symmetric perfect public equilibrium* is a symmetric perfect public equilibrium in which in each period underwriters choose pure actions *or* choose mixed

actions that are consistent with an equilibrium of the one-shot version of the game.

That is, I consider only strategies that generally call for both long-lived players to play the same pure strategy action profile following a given public history, with the action chosen by any underwriter being measurable with respect to the public history and the current signal received by that underwriter. Only when long-lived players are playing as if the game were not repeated can mixed strategies be included.⁷ And, the continuation strategy following each history must be a Nash equilibrium of the repeated game. To avoid cumbersome repetition of this phrase, I will simply refer to the perfect public equilibria of the game with the understanding that I am referring to this particular class of perfect public equilibria.

Note that this equilibrium concept implies that all short-lived players must play their static best response to the long-lived players' equilibrium actions. For a full treatment of the concept of symmetric perfect public equilibrium see Mailath and Samuelson [2006], and for related applications to price rigidity see Hanazono and Yang [2007] and Athey et al. [2004]. Finally, I introduce a public correlation device to simplify optimal punishments. This assumption is without loss of substantial generality as such devices are readily available, particularly in financial markets.

3 Spreads under Competition and Perfect Monitoring

The primary goal of this study is to show that the concentration of spreads at exactly one number and the behavior of spreads for the largest and the smallest firms are best explained by a combination of collusion and imperfect information among underwriters. To reach this conclusion, it is necessary to consider the likelihood that spreads would have the documented characteristics in the absence of collusive arrangements of the type considered. The model presented here will not generate rigid spreads under either perfect competition or oligopolistic competition, except in certain pathological cases. As long as signals about firm value are

⁷Permitting mixed strategies when long-lived players are simply playing stage-game equilibria permits a larger class of punishment strategies. Optimal collusion will generally take the form of pure strategies, while stage game play may involve mixed strategies, particularly if the distribution of signals is not atomless. These mixed strategies are admitted to avoid artificially and unnecessarily reducing the set of punishments available to long-lived players.

informative, the bid that leads to zero profit in expectation and the bid that maximizes profits in the one-shot auction will both depend on the signal. Since for a given spread a more valuable firm will generate a higher payoff to the underwriter, and since returns to scale are increasing, larger firms must receive lower spread offers, on average, than smaller firms. This will be true locally at every point, so contrary to the claims in Hansen [2001] and Torstila [2003] a competitive underwriting industry will not use a contract with a partially rigid spread.

On the other hand, if underwriters can effectively act as a monopoly, signals about value will be informative about the reservation spread of the firm as long as manager preference is not identically zero. Since signals are informative about firm value and firm value is informative about the distribution of the reservation spread, underwriters will charge slightly different spreads in response to slightly different signals. If manager preference were identically zero for all firms and the common value benefit of going public took exactly the proportional form described, then monopoly spreads would be rigid, as in Chen [2001]. Any slight perturbation of either of these assumptions would undo rigidity, and such assumptions cannot account for the increase in spreads charged on the smallest issues.

The impossibility of observing spreads that are rigid even over a small region is not sensitive to the parametric assumptions on costs and benefits assumed in the main text of the paper. Appendix B presents a more general model and argues that almost all cost and benefit functions will necessarily imply spreads that are not rigid over any interval. Here, however, I present a parametric example that develops the intuition for why spreads will not be rigid under competition, oligopolistic competition, or collusion with symmetric information.

In the absence of potentially imperfect monitoring (that is, when both underwriters receive the same signal before each auction), the costs and benefit functions above generate a pattern of spreads that bears some resemblance to the observed spread distribution but does not imply rigid spreads. Figures 2 and 3 show examples of this, first in the case where underwriters both learn the true value of the firm and the true value of the preferences of the manager and then in the case where they only learn the value of the firm. Figure 4 shows the spreads that would arise under perfect information and competitive pricing. Details of the straightforward derivation of

these spread schedules can be found in appendix C and D.⁸

[Figure 2 about here.]

[Figure 3 about here.]

[Figure 4 about here.]

In both of these cases, the fact that the signals are identical makes collusion easy since any deviation from the pure strategy spread schedule prescribed in a collusive equilibrium is detectable and can therefore be deterred without costs; the perfect public equilibrium then coincides with collusion at the monopoly spread schedule. Under such a collusive outcome, information about firm value and manager preference will prove useful for choosing the spread offer. When both value and preferences are observed, underwriters will push firms to their participation constraints; that is, firms will be exactly indifferent between going public at the collusive spread offer and remaining private. This participation constraint will depend both on the preferences of the manager and the offer price. When preferences are not observed, signals about value effectively become signals about the elasticity of demand for holding an IPO.

Other things equal, a firm with a manager with strong preferences for going public will be willing to pay a higher spread and will consequently be charged the higher spread. The value of the firm will determine how sensitive the spread offer is to individual preferences. More valuable firms will in effect care more about the spread relative to individual preferences. Spreads above the “common value” benefit of the IPO are costly in the sense that the post-IPO firm is less valuable than the pre-IPO firm since the payments to underwriters exceed the common value gains. For small firms, the individual manager benefits of going public can justify the costs, but for very valuable firms, even ones controlled by a manager with a strong preference for taking his firm public, the loss associated with holding an IPO at a very high spread will overwhelm the idiosyncratic private benefits. Symmetrically, firms with managers who prefer to keep their firms private will require spreads below the common value component of the benefit of going public. Small firms may demand very low spreads or, in extreme cases, refuse to go public at

⁸There is no closed form expression for the spread function for unobservable ε ; in the appendix, I solve for the closed form expression when ε is unobserved but distributed uniformly rather than normally.

any non-negative spread. But, since underwriters must pay a fixed cost to take a firm public, they will refuse to offer a spread low enough to induce a small firm with low preference for going public to hold an IPO. The distribution of spreads for small firms is thus truncated from below, and consequently the average spread charged to small firms will be high relative to the average spread charged to larger firms.

Additionally, in the case of observable manager preference, underwriters are assumed to be patient but not arbitrarily patient. Since the value of going public increases with the value of the firm while the optimal spread remains approximately the same, larger firms will, on average, prove more profitable for underwriters. Thus, to sustain an equilibrium, spreads must decrease for such firms so that a deviation will not be too profitable relative to expected future profits from maintaining collusion. This is, of course, an application of the result in Rotemberg and Saloner [1986].

Note that the spreads charged when underwriters collude and only value is observable is not necessarily monotone decreasing. For very small firms, the underwriters will charge a high spread to capture rents from those firms with strong preferences for going public. In intermediate ranges, charging relatively low spreads may be valuable because it can induce firms with idiosyncratic preferences for staying private to agree to go public. For the very largest firms spreads will be higher again since the relative importance of idiosyncratic preferences decreases. The participation constraints of very large firms will cluster very tightly around one particular spread, although the spread function will never be exactly flat. The fact that optimal collusive spreads asymptote to one particular spread strictly greater than zero will play an important role in explaining why spreads can be rigid over such a wide range of firm values while still responding to value for small and large firms.

4 Imperfect Monitoring

We have now seen that when underwriters receive identical signals about firm value and manager preference they optimally collude on spread offers that call on average for high spreads for small firms. Spreads for moderate sized firms decline sharply from these high levels, and the spread offers become more “flat” in the sense that firms of intermediate size are charged similar,

but not identical, spreads.

This section shows that when signals are informative but asymmetric such that underwriters have private information about the characteristics of the firm optimal collusion may call for a partially rigid spread. For simplicity, we focus on a model where underwriters receive conditionally independent signals about the value of the firm and no signal about the preferences of the manager, but all conclusions will be robust to admitting minimally informative private signals about manager preference.

Colluding on a rigid or partially rigid spread implies that underwriters ignore information that would be useful for setting a spread that extracts the most possible surplus from issuing firms. The benefit of ignoring such information comes from the fact that deviations from a rigid spread schedule can be perfectly detected and prevented with punishments that do not occur along the equilibrium path. A spread schedule that uses all available information about firms would require a different spread offer for each signal.⁹ Since underwriters cannot observe the signal of the other underwriter, deviations from the prescribed spread schedule cannot be directly observed. The punishments necessary to enforce such a spread schedule are then triggered by *apparent* deviations, and as such will occur along the equilibrium path. This inefficiency will in certain cases prove more severe than that associated with ignoring private information.

Finally, in the non-identical-signals setting the model still can predict that spreads will be high for small firms even when rigidity is better than fully separating spreads. Such a partially rigid spread function can arise without requiring on-path punishments. An underwriter with a very low signal about firm value will not necessarily have an incentive to imitate a higher signal (and thus increase his chances of winning the IPO). If such an action requires him to bid far below the first-best spread offer, the efficiency loss may exceed the gain to himself of capturing more of the market. Thus, a spread schedule that calls for most firms to be charged the same spread but for very small firms to be charge a higher spread can provide greater profits than either a fully rigid or fully separating spread function, all without requiring on-on path punishments.

⁹In certain cases not considered in this section, the spread schedule may be non-monotonic and thus a measure zero set of signals could call for the same spread offer.

For simplicity, I consider a particular class of problems, summarized in the following assumptions:

Assumption 1.

1. $x \sim U[0, \bar{x}]$
2. $\xi_i = \begin{cases} x & \text{with probability } p \\ U[0, \bar{x}] & \text{with probability } 1 - p \end{cases}$
3. $\varepsilon \sim \exp(\lambda)$
4. $\delta \rightarrow 1$

It follows immediately that $\xi_i \sim U[0, \bar{x}]$ and that

$$x|\xi_i \sim \begin{cases} \xi_i & \text{w.p. } p \\ U[0, \bar{x}] & \text{w.p. } 1 - p \end{cases}$$

Assumption 3 indicates that we are in a special circumstance in which the manager-specific value of an IPO is always positive. This assumption is made only for analytical tractability, as it helps guarantee that the optimal spread schedule will be weakly decreasing everywhere.

I now solve for the optimal rigid spread, the optimal two-step self-enforcing spread, and an approximately optimal fully separating spread. Calculating these spreads makes it possible to compare the value of each type of equilibrium and to conclude that, in certain cases, the two-step spread schedule can provide profits higher than any fully rigid spread or any fully separating spread. From this finding, it is possible to demonstrate that the optimal collusive spread schedule when a two-step self-enforcing spread is preferred to either a rigid spread or a fully separating spread will exhibit characteristics similar to those observed in data. Spreads for intermediate and large sized firms will almost always be the same, while small firms will be charged spreads that will decrease on average as the firm gets larger (but remains “small”), but firms of the exact same value will be charged different spreads with positive probability. These characteristics are exactly those documented in Chen and Ritter [2000]. Finally, we demonstrate that if underwriters are less than completely patient and the upper bound on

firm value is high enough, underwriters must demand lower spreads following the very highest realizations of their signals about firm value. These elements of the optimal spread function combine to demonstrate that the “seven percent solution” almost certainly results from attempts by underwriters to collude to extract as much surplus as possible from issuing firms.

4.1 Optimal Rigid Spread

The optimal rigid spread is the solution to a relatively straightforward univariate maximization problem. The only complications are that certain parameter values imply that spreads should be set either so low as to guarantee that all firms hold an IPO or so high as to guarantee that no firms hold an IPO. Specifically, when both the costs of holding the IPO and the mean of the idiosyncratic manager preferences are sufficiently low, the underwriters will set a rigid spread of $\alpha^r = 1 - \frac{1}{\beta}$, which guarantees that even the firm with $\varepsilon = 0$ holds an IPO. When costs are sufficiently high and the mean of the manager preference is sufficiently low, no spread leads to positive profit. We will ignore such cases by imposing a technical restriction on the relationship between the costs of holding an IPO, κ , and the mean of the distribution of manager preference, λ :¹⁰

Assumption 2.

$$\kappa\lambda < 1$$

The following proposition shows how to derive the optimal rigid spread. The derived expression is an immediate consequence of the necessary condition for an interior optimum for profit maximization and the proof is thus omitted.¹¹

Proposition 4.1. *The optimal rigid spread is given by $\alpha^{r*} = 1 - \frac{1}{\beta} + \gamma^*$, where*

$$\gamma^* = \begin{cases} \gamma' & \text{if } \gamma' > 0 \\ 0 & \text{otherwise} \end{cases},$$

¹⁰Note that this condition is not tight, in the sense that there are problems with $\kappa\lambda > 1$ where there is an interior optimum for α ; we ignore these cases for simplicity. In general, characterizing the set of parameter values where conditions hold is tedious and uninteresting; further details on these sets are available from the author.

¹¹The objective function in this problem fails to be quasiconcave for certain parameter values, but it is possible to show that there are no more than two local maxima and that identifying the global maximum is not difficult. Details are omitted to save space.

and γ' satisfies

$$\gamma' = \frac{1}{\beta\lambda\bar{x}} \log \left[1 + B + (\beta - 1 + \beta\gamma') \frac{B^2}{A} \right]$$

with

I. $A = 2(1 - \beta) - \gamma'\beta(1 - \kappa\lambda)$

II. $B = \gamma'\beta\lambda\bar{x}$.

Charging the optimal rigid spread would imply that firms are ignoring all information contained in their signals. This eliminates the difficulties associated with monitoring deviations from equilibrium since any deviation can be identified perfectly. It is, however, possible to increase the profits accruing to the underwriters without requiring on-path punishments. The following subsection demonstrates the procedure for finding just such an equilibrium.

4.2 Partially Rigid Spread

The benefits of a fully rigid spread come from the elimination of the need for punishments that occur along the equilibrium path. The costs are that spreads cannot be chosen optimally to extract as much surplus as possible from the firms given the information available to the underwriters. However, it is possible to choose a spread function that both uses information contained in the signals received by the intermediaries and does not require punishments along the equilibrium path. This section presents the form of such a “self-enforcing” spread function, proves the existence of such spread functions for the case where an optimal rigid spread exists, and describes the procedure for finding the optimal spread function within a restricted class of such self-enforcing equilibria. This two-step spread function will call for higher spreads for the smallest firms and a single, fixed spread for all other firms.

Definition A self-enforcing spread function is a function $\alpha^{se} : [0, \bar{x}] \rightarrow \mathbb{R}^+$ such that if players 1 and 2 play $\alpha(\xi_k), k \in \{1, 2\}$ in the pricing stage following any history where no deviation has been detected with probability one (that is, in period t , $a_s \Rightarrow P(\exists \xi \in [0, \bar{x}] \text{ s.t. } \alpha_s = \alpha^{se}(\xi)) > 0$ for all $s < t$) and play the stage-game equilibrium otherwise, then no arbitrarily patient underwriter has an incentive to deviate.

Such spread functions must take on a very specific form over the range of signals for which an underwriter would anticipate positive profits. To slightly simplify the discussion, we restrict attention to the class of spread functions that are weakly decreasing in signals about firm value:

Proposition 4.2. *Let $X^p(\alpha)$ be the set of signals for which an underwriter expects positive profits in the stage game under spread function α . Then, if α is part of a decreasing self-enforcing equilibrium, α restricted to $X^p(\alpha)$ is a step function.*

Proof. Assume the contrary. This implies that, in some region where expected profits are positive, the spread function is continuous but not flat. Then, there is some signal ξ' such that an arbitrarily small deviation from $\alpha(\xi')$ is not detectable with probability one. But, such an arbitrarily small deviation will increase expected profits by $\frac{p^2}{2}(\alpha(\xi')\beta\xi - \kappa)e^{-\lambda x(\alpha\beta+1-\beta)}$ since the implied change in α is infinitesimal and the deviator now captures the entire market when both agents receive the same, correct signal. Thus, a profitable deviation exists for underwriter type ξ' . \square

Given this result, I can restrict attention to decreasing step functions without loss of substantial generality. Such step functions effectively generate partially separating equilibria. That is, all underwriter types effectively pool with those other types within their step. Such a self-enforcing step function is guaranteed to exist under general conditions:

Proposition 4.3. *For any configuration of parameter values for which there is a rigid spread that implies positive profits, there is some self-enforcing two-step spread.*

See the appendix for a proof.

Calculating an optimal two-step self-enforcing spread is a relatively simple multivariate optimization subject to a single incentive compatibility constraint since indifference at the threshold implies strict preference away from the threshold, a form of a single crossing property. The optimization problem can be written as follows:

Let

$$\begin{aligned}\pi(\alpha, x) &= \alpha\beta x - \kappa \\ h(\alpha, x) &= e^{-\lambda x(\alpha\beta+1-\beta)}.\end{aligned}$$

$$\begin{aligned}
& \max_{\alpha^h, \alpha^l, x^*} \int_0^{\bar{x}} \left(\mathbf{1}_{x < x^*} \left(p^2 \left(1 - \frac{x^*}{\bar{x}} \right)^2 + \frac{2px^* \left(1 - \frac{x^*}{\bar{x}} \right)}{\bar{x}} \right. \right. \\
& \left. \left. + \mathbf{1}_{x > x^*} (1-p)^2 \left(\frac{x^*}{\bar{x}} \right)^2 \right) \pi(\alpha^h, x) h(\alpha^h, x) \right. \\
& \left. + \left(\mathbf{1}_{x < x^*} \left(1 - \left(1 - \frac{x^*}{\bar{x}} \right) \left(p^2 \left(1 - \frac{x^*}{\bar{x}} \right) + \frac{2px^*}{\bar{x}} \right) - \frac{x^*}{\bar{x}} \right) \right. \right. \\
& \left. \left. + \mathbf{1}_{x > x^*} \left(1 - (1-p)^2 \left(\frac{x^*}{\bar{x}} \right)^2 \right) \right) \pi(\alpha^l, x) h(\alpha^l, x) \right) \frac{1}{\bar{x}} dx
\end{aligned}$$

subject to

$$\begin{aligned}
& \frac{1}{2} \left(\left(p^2 + \frac{x^* p (1-p)}{\bar{x}} \right) \pi(\alpha^h, x^*) h(\alpha^h, x^*) \right) \\
& + p(1-p) \int_0^{x^*} \pi(\alpha^h, x) h(\alpha^h, x) \frac{1}{\bar{x}} dx + (1-p)^2 \frac{x^*}{\bar{x}} \int_0^{\bar{x}} \pi(\alpha^h, x) h(\alpha^h, x) \frac{1}{\bar{x}} dx \\
& \geq \\
& \frac{1}{2} \left(p^2 \left(1 - \frac{x^*}{\bar{x}} \right) + p \left(1 + \frac{x^*}{\bar{x}} \right) \right) \pi(\alpha^l, x^*) h(\alpha^l, x^*) \\
& + p(1-p) \left(\int_0^{x^*} \pi(\alpha^l, x) h(\alpha^l, x) \frac{1}{\bar{x}} dx + \frac{1}{2} \int_{x^*}^{\bar{x}} \pi(\alpha^l, x) h(\alpha^l, x) \frac{1}{\bar{x}} dx \right) \\
& + \frac{1}{2} (1-p)^2 \left(1 + \frac{x^*}{\bar{x}} \right) \int_0^{\bar{x}} \pi(\alpha^l, x) h(\alpha^l, x) \frac{1}{\bar{x}} dx.
\end{aligned}$$

The existence of such a two-step self-enforcing spread arises because the optimal spread rises sharply as firms become small. A two-step spread with a relatively high threshold could be more profitable as it would allow more efficient pricing of the most valuable IPOs, but it would be impossible to enforce such a spread since the temptation of the firms near the threshold to bid as if they had a higher signal would be too strong. Self-enforcement requires that the step be large and that the higher spread offer be relatively close to the optimal spread offer for the threshold firm. In this case, when deviating the benefits of capturing more of the market are offset by the costs of charging a less efficient spread. The costs to the underwriter of holding an IPO play an important though not absolutely essential role in the nature of the self-enforcing spread. When underwriting costs (i.e. κ) are relatively high the costs of misreporting one's signal following a low signal can be substantial as the market share increase comes largely from winning IPOs that

are unprofitable at the lower spread. With very small or even zero underwriting costs, however, it is still possible to construct two-step self-enforcing spreads; these will generally call for very high spreads for the very smallest firms. In this case the self-enforcement is driven exclusively by the fact that signals about firm value are informative about the expected reservation price of the firm; very small firms should be charged very high spreads since there is a small chance that the firm has a high idiosyncratic preference for going public and the firm's small size effectively magnifies this preference. Introducing underwriting costs implies two-step spreads much closer to what is observed in data.

A two-step self-enforcing spread will not, in general, prove optimal, but we will not seek to precisely characterize the optimal self-enforcing spread. Instead, it is sufficient to show that the two-step spread dominates the optimal fully flexible spread in order to demonstrate that the optimal spread function does not respond to firm size everywhere. Since relying on the self-enforcing characteristics of step functions is the only way to decrease the costs associated with monitoring for a fully separating spread and since such step functions must have thresholds calling for increased spreads for the smallest firms, we can conclude that if the optimal self-enforcing two-step spread provides more profit than the optimal fully separating spread then the optimal spread function must call for rigid spreads over a large region of intermediate valued firms and higher spreads for smaller firms. These smaller firms may face a more complicated spread schedule than a simple two-step function, but we will show that when a two-step spread is preferred to either a rigid spread or a fully flexible spread the optimal spread schedule will exhibit both rigidity over a large region and behavior generally consistent with observed data for small firms.

4.3 Flexible Spreads

We have so far confined attention to spread functions that do not require punishments along the equilibrium path. These are exactly the partially rigid spreads that appear to match the empirical data. We now consider instead those spread functions that serve to fully separate types and thus exploit the information about firm value contained in the signals to the underwriters. For the class of problems under consideration, this reduces to searching for the optimal strictly

decreasing spread function.

In this section, we discuss a method for deriving an upper bound on the equilibrium payoffs of a strictly decreasing spread function. We follow a simplified version of the dynamic programming approach developed in Abreu et al. [1990]. For details on implementing the full procedure, see Mailath and Samuelson [2006] and Judd et al. [2003]. We assume throughout that underwriters have access to a public randomization device.

The problem of finding the optimal fully separating symmetric perfect public equilibrium reduces to finding the everywhere downward sloping spread function that satisfies all incentive compatibility constraints, conditional on promised continuation values falling in the interval $[\underline{v}, v^*]$, where \underline{v} is the payoff associated with the worst symmetric perfect public equilibrium and v^* is the payoff associated with the best such equilibrium. Incentive constraints take the basic form of requiring that, for every possible misreport, the increased likelihood of reverting to the “punishment” continuation value following any deviation from the proposed spread schedule is sufficiently high to overcome any short-term improvement associated with capturing more market share. Continuation values must, in turn, be drawn from the set of feasible equilibrium payoffs, $[\underline{v}, v^*]$. Furthermore, we know from the bang-bang result of Abreu et al. [1986] and Abreu et al. [1990] that any equilibrium payoff can be achieved with strategies that call for continuation values drawn only from the extreme points of the set of equilibrium values.¹² In the symmetric case considered here, this reduces the set of strategies under consideration to those that call for play of the “optimal” spread schedule until punishment is triggered, followed by perpetual play of the stage game equilibrium following transition to the “punishment” phase.¹³ Note that the optimal spread schedule will *not* be the spread schedule that a monopolist observing the first order statistic of the signals would offer. The need to satisfy incentive compatibility constraints will introduce a tradeoff between choosing an efficient spread schedule and choosing

¹²Readers concerned about the application of the bang-bang theorem to a game with a continuous action space may consult appendix A for an explanation of the applicability of the results to games with a finite but arbitrarily fine action space.

¹³It is straightforward to show that the worst sustainable symmetric perfect public equilibrium payoff is the worst payoff associated with repeated play of a stage-game symmetric Nash equilibrium for arbitrarily patient underwriters. When underwriters are not arbitrarily patient, it may be possible to sustain a lower value as a SPPE payoff for impatient underwriters. However, this value will clearly be above 0 (a lower bound on the set of individually rational payoffs), and abusing notation slightly we will continue to refer to this worst value as the stage-game payoff, Π^{SG} .

a spread schedule that does not tempt underwriters to deviate. This will, in general, lead to a spread function that is distorted downward for larger values.

The problem then becomes to find the optimal spread function and reversion probabilities to maximize the value of the game. Reversion probabilities, of course, cannot depend directly on whether an underwriter deviated from his assigned bid. Reversion probabilities can only depend on the public history, specifically the history of the true value of the firms and the lowest spread offers. The bang-bang restriction, in turn, allows us to consider only those strategies where reversion probabilities depend only the most recent realization of the public signal. The value of the game, then, will depend on the probability of reversion given that everyone adheres to the equilibrium strategy. The dynamic programming approach to solving for this value consists of finding a superset of the set of payoffs sustainable in a perfect public equilibrium, maximizing the value of the game over all downward sloping spread schedules, reversion probabilities and continuation values that are, taken together, incentive compatible for every signal and every possible deviation and where the continuation values fall in the posited superset. The maximum value that can be achieved by solving this problem is also an upper bound on the set of sustainable perfect public equilibrium payoffs. The same procedure with minimization replacing maximization gives and new lower bound. This procedure can be repeated indefinitely to obtain smaller and smaller supersets of the equilibrium value set, and will iterate to convergence in sufficiently regular programs. We will be concerned only with calculating a sufficiently restrictive upper bound and will therefore rely only on the monotonicity of this operator.

To summarize the approach used, note first that it is straightforward to find a superset of the set of symmetric perfect public equilibria. Underwriters would certainly be able to guarantee themselves at least $-\varepsilon$, where ε can be made arbitrarily small, by demanding an extremely high “proportion” of the firm in exchange for underwriting services. If underwriters are allowed to bid $\alpha = \infty$, it is clear that no perfect public equilibrium can call for payoffs below 0. On the other hand, we know that it is impossible for underwriters to do better in equilibrium than would a monopolist who observes the first-order statistic from the two signals about firm size. Calculating the optimal spread function for such a monopolist is a relatively simple matter,

and the payoffs associated with playing such a strategy in every period represent an upper bound of what underwriters can achieve in a collusive equilibrium. Designate this value \bar{v}^0 . We can now maximize the value of the game treating continuation values as choices from the set $[0, \bar{v}^0]$. The continuation values, spread schedule, and reversion probabilities must satisfy the incentive compatibility conditions that prevent any underwriter from effectively misreporting his signal about the value. The maximum value that can be achieved by solving this program is an upper bound for the value that can be achieved in a symmetric perfect public equilibrium. This procedure can now be repeated indefinitely to obtain a decreasing upper bound on the equilibrium value set. This upper bound can then be compared directly to the maximum value that can be achieved with a rigid spread or a self-enforcing two-step spread function, both of which can be calculated directly without relying on a dynamic programming approach since neither require punishments along the equilibrium path.

Details of the procedure can be found in the appendix.

4.4 Comparing Values: An Example

We have now established a procedure to find an upper bound on the value of a fully separating spread schedule under optimal collusion. We can also find the optimal value that can be achieved with a fully rigid spread or a two-step self-enforcing spread. If either the fully rigid spread or the two-step self-enforcing spread imply greater value than the fully separating spread, we can conclude that the optimal symmetric perfect public equilibrium strategy implies partial price rigidity.

It is clear that, for sufficiently small values of p , a rigid spread will dominate a flexible spread since the optimal flexible spread without incentive constraints converges pointwise to the optimal rigid spread. That is, when the private signals contain almost no information about the value of the firm, even a monopolist would not alter his spread offer much in light of the information received from the private signals. Maintaining incentive compatibility, however, remains costly.

It is not as clear, however, that a two-step self-enforcing spread schedule will be optimal for intermediate values of p . As $p \rightarrow 1$, a fully flexible spread will dominate the two-step

spread since the value from the fully flexible spread converges to perfect information monopoly profits.¹⁴ For $p \rightarrow 0$, the rigid spread dominates the two-step spread since the two-step spread requires a large gap between the high spread and the low spread in order to maintain self-enforcement. A two-step self-enforcing spread will be optimal over some intermediate range of p if the costs of enforcement of the fully flexible spread rise sufficiently fast as p decreases relative to the decline in value of segmenting the market based on spreads. It is difficult to characterize the region of p for which the two-step spread will dominate. The technique described above, however, enables us to find examples where exactly this occurs.

Specifically, consider the case where $\bar{x} = 20$, $\beta = 1.01$, $\lambda = 3$, $\kappa = .1$, and $p = 0.6$. Then, an upper bound on the value¹⁵ that can be achieved with a fully flexible spread is 0.0416, while the rigid spread gives 0.0432 and the two-step self-enforcing spread gives 0.0465. This self-enforcing rigid spread calls for a spread offer of 0.22 following signals below 2.87 and a spread offer of 0.035 for those above this threshold. See figure 6 for a plot of the optimal two-step self-enforcing spread and an approximation to the optimal fully separating spread schedule resulting from the iterative procedure above. Also, see Figure 5 for an approximation of the punishment function ρ necessary to sustain the optimal flexible spread. Note the shape; small deviations are often more profitable than large deviations, so the probability of reverting to the punishment phase are lower following signals that indicate the possibility of a large deviation. For larger signals, however, payoffs will increase in the size of the deviation, at least initially, as more and more valuable market share is captured. Thus, the punishment function is not in general monotonic.

[Figure 5 about here.]

[Figure 6 about here.]

Since the two-step spread is strictly preferred to either the fully rigid or the fully flexible spread, it is clear that minimal perturbations of the environment will not undermine the incentives to use a partially rigid spread. That is, small changes in the distribution of ε, x , or ξ would generate small changes in the value of the two-step spread and the value of the fully flexible

¹⁴This follows from the fact that, as $p \rightarrow 1$, the underwriters will have almost perfect information. Furthermore, the probability that reversion will be triggered is bounded above by $1 - p^2$.

¹⁵All values are expressed as the expected normalized discounted sum of payoffs.

spread, so the two-step spread would remain more profitable. Minimally informative private signals about manager preference would also not change the ordering of profitability; the optimal fully separating spread would now be a function both of the signal about firm value and the signal about manager preference, but the value of using such a schedule would converge to the value of ignoring manager preference as the informativeness of the signal decreases. So, in this sense, spread rigidity can arise for generic parameter values and functional forms for the relevant structural elements of the IPO industry when spread rigidity is motivated by the need to collude in an environment with private information. This contrasts with the scenarios where spreads can be partially rigid under competition, competitive oligopoly, or monopoly provision, which rely on knife-edge cases to generate spread rigidity.

The example considered in this section is particularly revealing as it represents a case where, arguably, colluding on a fully separating spread schedule should be relatively easy. With positive probability, both underwriters receive the same signal and therefore should even in a fully separating equilibrium frequently bid exactly the same when adhering to a collusive spread schedule. If signals were instead absolutely continuous, every single stage would result in different bids with probability 1. While the relative costs of monitoring will depend on the signal structure in a complex manner, the example presented here suggest that partially rigid spreads would be optimal for many signaling environments.

4.5 Optimal Spreads

The above discussion shows that optimal spreads may exhibit partial rigidity when underwriters collude but have private information about firms. Deriving the two-step spread function that provides higher payoffs than the optimal feasible fully separating spread function is sufficient to show that the optimal spread function must exhibit some degree of rigidity, but such a two-step spread is not necessarily the optimal feasible spread schedule. Indeed, data indicate that underwriters do not collude on a two-step spread schedule. Instead, for small firms, several different spreads are charged, with the spread generally, but not strictly, declining in firm value.

In this section, I demonstrate that in fact the optimal self-enforcing spread schedule in the example above must involve more than the single step up. I also argue that all of these steps

must occur over the region of low value firms, where “low value” refers to those firms whose value is small enough that manager preference plays a very important role and, as a result, first-best spread offers decline steeply in the signal about firm value. The empirical implications of a spread function of this type provide a very close qualitative match to the distribution of spreads documented in Chen and Ritter [2000].

To see that the optimal self-enforcing spread function must call for more than two steps in the numerical example above, we rely on the following proposition:

Definition An n -step self-enforcing spread function is an $n + (n - 1)$ -tuple $\left\{ \left\{ \alpha^i \right\}_{i=1}^n, \left\{ x^i \right\}_{i=1}^{n-1} \right\}$ where $\left\{ \alpha^i \right\}$ is a decreasing sequence of spreads and $\left\{ x^i \right\}$ is an increasing sequence of thresholds, and $x^0 = 0$ and $x^n = \bar{x}$, where underwriter j is assigned to demand spread α_i if it has a signal $\xi_j \in (x^{i-1}, x^i]$. Furthermore, the schedule can be enforced without recourse to on-path punishments.

This is just a generalization of the two-step self-enforcing spread function.

Proposition 4.4. Consider an n -step self enforcing spread function $\left\{ \left\{ \alpha^i \right\}_{i=1}^n, \left\{ x^i \right\}_{i=1}^{n-1} \right\}$. If

$$\begin{aligned} \frac{p^2}{2}(-\kappa) &+ \frac{p(1-p)}{2} \frac{1}{\bar{x}} \int_0^{x^1} \pi(\alpha^1, x) h(\alpha^1, x) dx \\ &+ \frac{(1-p)^2}{2} \frac{x^1}{\bar{x}} \int_0^{\bar{x}} \pi(\alpha^1, x) h(\alpha^1, x) dx \\ &< 0 \end{aligned}$$

then there exists a self-enforcing spread schedule with $n + 1$ steps that provides higher payoffs by calling for higher spreads for a subset of signals $[0, x^1]$.

Proof. The expression in the proposition is simply the limit of the expected payoff to an underwriter for adhering to the proposed equilibrium as the signal goes to zero. If this expression is negative, that implies that underwriters with the smallest signals expect negative profit. A schedule calling for those underwriters with very small signals to charge higher spreads would thus be more profitable overall. By choosing the region for this higher spread to be arbitrarily small but with positive measure, the necessary changes to the existing steps in order to make

the new schedule an equilibrium would be arbitrarily small. A standard continuity argument then shows that the $n + 1$ step spread is an equilibrium. \square

Applying this proposition to the numerical example shows that there must be at least 3 steps in the optimal self-enforcing step spread schedule. Note that it is not immediate that any self-enforcing step spread schedule can be improved upon with a spread schedule with one additional step. An underwriter with an arbitrarily small signal may still make positive profits if the profits accruing when he is wrong about the firm value exceed the losses when he is right. We can, however, conclude that all of the steps that do arise in an optimal or approximately optimal self-enforcing spread schedule will have the steps concentrated in the low value region.

To see this, note first that large steps are easier to enforce without on-path punishments. A small step cannot be easily enforced since total value accruing will not change much following a deviation from one step to another while the probability of winning the market will increase appreciably. Therefore, if underwriters want to enforce a step with a threshold at an intermediate value firm, the large step required to maintain self-enforcement will generate a large efficiency loss since the first-best spread schedule becomes relatively flat for intermediate and large sized firms. On the other hand, first-best spreads rise sharply as firm value becomes small, so imposing a large enough step to maintain self enforcement in this region does not involve such a severe efficiency loss.

Exactly what constitutes the region of “small firms” is difficult to succinctly define, but the basic intuition is that spreads will be high where the decision to go public is driven primarily by the idiosyncratic private benefits to managers rather than by the opportunity to increase the common value of the firm.

The empirical implication of a spread schedule that calls for offering a relatively low, fixed spread following most signals and offering higher spreads following signals that the firm value is low, with the exact spread offer for low value firms depending on the signal in a coarse but non-trivial manner, will prove quite similar to the distribution of spreads in data. Since signals about firm value are imperfect, the relationship between true firm value (as measured by the closing price on the first day of trading or, in the model, the issue price) and realized spread will have the steps overlap. With positive probability, both the underwriters will receive and

incorrect signal. If both of these incorrect signals are outside the (endogenously determined) step containing the true value, the realized spread offer will differ from the spread that would be charged if the signal were correct. If only one signal is incorrect but that signal leads to a spread offer below that of the correct signal, then again the realized spread will differ from the spread implied by the schedule. For example, if the true value of the firm were such that a correct signal would lead to the highest spread offer, but one firm receives an incorrect signal that the firm is actually of intermediate value, then that firm would be charged the lower spread.

Note that this argument is not symmetric. A firm of intermediate or high value will only be charged a high spread if *both* signals are wrong since the firm can choose the lowest spread offer. Furthermore, when both signals are wrong and both signals imply that the firm is small, an already unlikely event, the firm is very likely to decline to go public since its participation constraint is likely to be violated; the increasingly tight participation constraint is the primary reason why intermediate value firms are charged the lower spread in the first place, so clearly very few will agree to go public if underwriters demand the relatively high spread charged to small firms. Thus, while small firms will occasionally go public at lower spreads than most firms of their size (as observed in the data), few if any large or intermediate size firms will be observed going public at a high spread. This is consistent with the key finding of Chen and Ritter [2000] that price rigidity is nearly absolute over a significant range of firms while also explaining the richer distribution of spreads for smaller firms.¹⁶

4.6 Impatient Firms

Up to this point, we have demonstrated that a model with private information and collusion by underwriters generates a relationship between firm value and spreads that closely resembles the observed data for small and intermediate value firms. In this section, we show that if underwriters are not arbitrarily patient, spreads may decline for the very largest firms while continuing to exhibit the same pattern described above for low and intermediate value firms.

Underwriter impatience can thus account for the decline in spreads observed for the largest

¹⁶This point undermines the argument that the observation of 7% spreads in certain small offerings shows that collusion is not responsible for the price clustering (see Hansen [2001]). Once imperfect information about value is introduced, observing exactly 7% for certain small firms will not be surprising since only one underwriter need believe that the firm has intermediate, rather than low, value in order for 7% to arise.

firms to hold IPOs.

The intuition for this decline is the same as described for the case where firm value and idiosyncratic preferences are perfectly observable. After receiving a signal that the value of the firm is very high, an underwriter expects profits from an optimally chosen spread to also be very high. If the underwriter is impatient and the signal indicates that the firm is sufficiently large, the underwriter will not adhere to the optimal *ex ante* collusive equilibrium but will instead make an (observable) deviation to undercut the other underwriter. This will trigger punishments, but the impatient underwriter will prefer the short-term benefit of the deviation to remaining in the collusive equilibrium.

Given that underwriters are too impatient to enforce optimal collusion, a collusive scheme that takes this constraint into consideration can provide greater profits than a naive attempt to collude on the optimal collusive equilibrium that arbitrarily patient underwriters would choose. The optimal spread schedule for impatient firms requires lower spread offers following high signals and on equilibrium path punishments, but, if impatience is only relevant for the very highest signals, the spread schedule will be effectively unchanged from the arbitrarily patient case except following the very highest signals.

Thus, collusion with private information and some degree of impatience can produce spread schedules that respond to firm value at the extremes of the distribution of firm value but also cause almost all firms of intermediate value to face the exact same spread. This model then explains the key qualitative characteristics of the “seven percent solution.”

To demonstrate the above point formally, we will consider the effect of introducing impatience to those games with optimal collusive equilibria in the class described for the numerical example. The following assumption describes the set of problems under consideration and introduces minimal impatience:

Assumption 3.

1. *Parameter values are such that, if underwriters were arbitrarily patient ($\delta \rightarrow 1$), they would collude optimally on a spread schedule calling for a rigid spread for intermediate and large firms and higher spreads for the smallest firms, and would strictly prefer such*

a schedule to a fully separating or fully rigid schedule. This schedule will henceforth be referred to as the arbitrarily patient equilibrium

2. Underwriters are impatient such that, where δ^* is the minimal patience needed to sustain the equilibrium where underwriters are arbitrarily patient, $\delta < \delta^*$.

3. $\delta^* - \delta \approx 0$

These assumptions guarantee that there is some $\hat{x} \in [0, \bar{x})$ such that underwriters have an incentive to deviate from the arbitrarily patient equilibrium if and only if the signal ξ is in $(\hat{x}, \bar{x}]$, and that this interval is arbitrarily small. With this δ , underwriters who attempt to enforce the arbitrarily patient equilibrium will eventually observe an “off path” deviation and revert to perpetual one-shot equilibrium play. A better schedule, and one that would remain an equilibrium, would call for underwriters to offer a lower spread following the signals above \hat{x} . This alternative equilibrium would not be self-enforcing when spreads are chosen optimally since underwriters receiving signals close to but below \hat{x} would have an incentive to claim a higher signal. Consequently, a public signal that the winning bid implied a signal $\xi \in (\hat{x}, \bar{x}]$ but that the true value x was less than \hat{x} must trigger reversion to stage game play with positive probability. Assumption 3 above guarantees that the probability that such an event will occur along the equilibrium path is arbitrarily small. Thus, by distorting the spreads down for the very highest signals, underwriters can avoid triggering the off-schedule deviations that would occur when attempting to collude on the arbitrarily patient equilibrium. Since the punishments required to enforce this nearly optimal equilibrium occur infrequently and the very large realizations of the signal that require downward distortion also occur infrequently (since \hat{x} is arbitrarily close to \bar{x} and the probability of one of these large signals is $\frac{\bar{x}-\hat{x}}{\bar{x}}$), there is no incentive for underwriters to employ a schedule over signals in the range $[0, \hat{x}]$ that differs qualitatively from that called for in the arbitrarily patient equilibrium. That is, as $\delta \uparrow \delta^*$ the equilibrium that is identical to the arbitrarily patient equilibrium over $[0, \hat{x}]$ and exploits the most efficient combination of downward distortions for $\xi \in (\hat{x}, \bar{x}]$ and on-equilibrium-path punishments will provide payoffs converging to the payoffs to the arbitrarily patient equilibrium. Since a partially rigid spread schedule is strictly preferred to either a fully separating or fully rigid schedule in the case of arbitrarily patient underwriters, moving to a rigid or fully separating spread schedule cannot

improve payoffs in the case of minimal impatience.

We will not attempt here to trace out the exact empirical implications of impatience. Since the distribution of large firms is unlikely to be well approximated by a uniform distribution, and a distribution with a long, thin tail would produce substantially different predictions about spreads charged by impatient underwriters, such an exercise would provide little in the way of validating or rejecting the model. We have, however, shown that in the context of the model with private information, impatience by underwriters will lead to a decrease in the spreads charged to the most valuable firms without undermining the incentives to apply the partially rigid spread described in the preceding subsection over the range of signals not near the maximum of the distribution.

5 Underpricing

While the great extent of spread rigidity in the market for IPOs has received significant attention in financial economics, the tendency for IPOs to be significantly underpriced has been the subject of far more research. Prominent examples include Loughran and Ritter [2004], Booth and Chua [1999], and Lowery and Shu [2002]. In this section, I show that spread rigidity and underpricing may be closely related.

Under the rigid or partially rigid equilibrium described above, underpricing has a natural role. Underwriters collude to extract as much surplus from firms as possible. However, given the need to collude on a rigid or partially rigid spread and the imperfect information about firm value and preferences, those firms that choose to hold an IPO will still benefit. That is, firms are not pushed to their participation constraint, and indeed underwriters are not even extracting all surplus that they could given their information. This section shows that underwriters may have an incentive to underprice issues in order to extract this additional surplus and to exploit any additional information received during the book-building process. The opportunity to underprice, in turn, reinforces the incentives to collude on a rigid or partially rigid spread; part of the loss from failing to charge the most efficient, flexible spread is recouped through the underpricing stage without requiring costly punishments.

Admitting underpricing into the analysis requires explicitly considering the choice of the offer price. Above, I have implicitly assumed that the offer price for the shares perfectly reflects the true value of the firm. That is, while underwriters have imperfect information about the value of the firm when bidding to hold the IPO, the book building process functions perfectly to reveal the true value and there is no incentive to underprice the offer.

Given the high and highly variable underpricing observed in the industry, these simplifying assumptions clearly do not capture all important elements of the industry. Either underwriters remain uninformed about the true value of the firm relative to information available to the general investing public or they intentionally underprice issues. It is difficult to believe that the underwriter would be so poorly informed relative to the first pair of investors to engage in a transaction following the IPO,¹⁷ so we proceed under the assumption that by the end of the book building process the underwriter has access to all information necessary to accurately price the security. Additionally, the underwriter is likely to learn more about the preferences of the managers of the firm and will thus update the distribution of ε .

To introduce underpricing into the model, let the underwriter also have access to a costly technology for turning underpricing into profits. Specifically, assume that an underwriter is in a long term relationship with a class of institutional investors who will “kick back” part of their gains from receiving an underpriced issue. This sort of behavior is well documented, from the brazen (SEC [2004]) to more subtle *quid pro quo*. Underpricing will also have other less direct costs; a large degree of underpricing may attract scrutiny from regulators or lawsuits from issuers, and too much underpricing may hurt the reputation of the underwriter and make other firms less likely to choose that underwriter in the future.

We assume that underwriters are unable to compete on underpricing. Since the issue price is set very late in the IPO process, underwriters cannot contract on a particular issue price. Thus, if underwriters wished to compete on underpricing, they would have to send credible signals to the firms that they will not underprice even when it is profitable to them. Any such signal would be just as observable to the other long-lived underwriter as it would be to the sequence of short-lived firms, so underwriters would be able to punish any attempt to use more

¹⁷Kirgman et al. [1999] show that an overwhelming proportion of the first day return, the standard measure of underpricing, is realized on the first trade of that day.

favorable issue pricing to capture more of the market.

Formally, we consider the following environment:

- Once the underwriter is chosen to hold the IPO, it observes x perfectly
- With probability μ , the underwriter observes ε
- The underwriter chooses an offer price $x_o \leq x$, which implies a degree of underpricing $u = x - x_o$.
- The firm has an opportunity to withdraw from the issue and receive its reservation value x (in which case the underwriter receives $-\kappa$)
- If the issue proceeds, the payoffs are:

underwriter: $\alpha\beta x_o + \theta_t\beta u - \zeta_t\beta^2 u^2$, where $\theta_t \sim [0, 1]$ and $\zeta_t \sim [0, \infty)$

firm: $(1 - \alpha)\beta x_o + \varepsilon$

This environment guarantees that a firm’s participation constraint in the underwriter-selection stage is unchanged; with positive probability the underwriter will have no incentive to underprice. That is, if $\theta < \alpha$ or ζ sufficiently large, a firm will prefer to choose $x_o \approx x$; when this occurs with positive probability, any firm that would choose to go public at a “fair” offer price will continue to agree to hold an IPO in the selection stage. Those firms who are almost indifferent between going public at the equilibrium α and remaining private will often be forced to their participation constraint when ε is revealed and will often cancel the IPO when ε is not revealed. They, however, still expect positive unconditional surplus, with the expected surplus decreasing in α .¹⁸

Underwriters will clearly benefit from the introduction of the underpricing stage. They have more information and another means to extract surplus. Furthermore, the value of coordinating

¹⁸It is not immediate that, in the presence of underpricing, firms would continue to choose the lowest spread. If underpricing is a costless way to generate additional profits, firms will be indifferent between a high spread with the expectation of low underpricing and a low spread with the expectation of high underpricing. The convex costs associated with underpricing guarantee that most firms will still prefer a low spread to a high spread. It is possible to construct examples even with convex costs where some small firms prefer larger spreads (since larger spreads serve, in part, to align the interest of the firm and the underwriter), but these cases are not of particular interest.

on a fully flexible spread relative to a self-enforcing step function is reduced, thus preserving the spread-rigidity result above. That is, the offer-price phase will allow underwriters to extract some of the surplus lost by coordinating on a self-enforcing partially rigid spread, still without relying on on-equilibrium path punishments. To see this, consider the case where the mass of the distribution of θ becomes concentrated around 1 while the mass of the distribution of ζ becomes concentrated around 0, and $\mu \rightarrow 1$. Then, the underwriter extracts almost the full surplus from each firm, regardless of the spread offer. Taking a very small firm public at a relatively low α , however, will still imply an expected loss for the underwriter unless ε is very high. This situation implies that underwriters would still have an incentive to coordinate on a step-spread function rather than simply charging a flat, low spread regardless of the signal.

Under the assumptions presented above, it is straightforward to solve for optimal underpricing. When the underwriter does not observe ε , it will choose the optimal offer price taking into account the risk that too much underpricing will cause the firm to cancel the offer. The underwriter, of course, has effectively received a signal that the firm has $\varepsilon \geq x(\alpha\beta + 1 - \beta)$, and will therefore maximize conditional on an updated posterior distribution of ε . The optimal level of underpricing is then given by:

$$u^* = \frac{\left(\frac{2\zeta\beta}{\lambda(1-\frac{1}{x})} + \theta - \alpha\right) - \sqrt{\left(\frac{2\zeta\beta}{\lambda(1-\frac{1}{x})} + \theta - \alpha\right)^2 + 4\zeta\beta\frac{\theta-\alpha-\alpha x}{\lambda}}}{2\zeta\beta}$$

if said quantity falls in the interval $[0, x]$; a negative offer price or an offer price above the true value is infeasible.

In the cases where the underwriter does observe ε , the issue will never be withdrawn as the underwriter can now perfectly observe the participation constraint. In cases where the participation constraint does not bind, underpricing is given by

$$u^* = \frac{\theta - \alpha}{2\zeta\beta},$$

again with the *caveat* that underpricing must fall in the feasible interval. When the participation

constraint does bind, optimal underpricing is instead given by

$$u^* = \frac{\varepsilon}{(1-\alpha)\beta} - \left(\frac{\alpha\beta + 1 - \beta}{(1-\alpha)\beta} \right) x.$$

It is immediate that money left on the table will not depend on x when the participation constraint does not bind. It is less obvious, but nonetheless true, that the expected value of money left on the table conditional on firm value does not depend on the value of the firm when the participation constraint does bind. That is, the expected value of ε conditional on a firm choosing to go public at a given spread increases in x at a rate that exactly offsets the decrease in x of the degree of underpricing conditional on ε . This exact balancing is in part an artifact of the modeling structure, in particular the assumption that the costs of underpricing are convex in the absolute rather than relative degree of underpricing, but may help explain why the relationship between money left on the table and firm value is not obviously monotonic.

6 Seasoned Equity Offerings

While the spreads on IPOs exhibit remarkable price rigidity, spreads charged for seasoned equity offerings demonstrate dependence on issue size and also exhibit variance even conditional on issue size (See Chen and Ritter [2000], figures 4 and 5.) This suggests that the logic underlying the collusive solution for IPOs does not apply to SEOs. Since IPOs and SEOs are fundamentally different in many ways, this difference should not be taken as evidence that the market for SEOs is more competitive than that for IPOs, even though spreads on SEOs are on average much lower than those on IPOs. There are two crucial differences between the market for IPO services and the market for SEO services. First, there is no asymmetric information about firm value when a firm is considering holding an SEO since the stock price is already a publicly observable variable. And, firms holding SEOs, almost by definition, have a preexisting relationship with one investment bank. Krigman et al. [2001] document that 70% of firms completing an SEO within three years of their IPO use the same underwriter as they did for the IPO. This suggests that underwriters may use their preexisting relationships with various firms to coordinate on an asymmetric bid-rotation scheme, which would permit flexible spreads

on SEOs that make the most use of all information available to the underwriter who held the firm's IPO. The lower spreads observed on SEOs would then arise not from greater competition amongst underwriters, which is hard to reconcile with the 70% figure cited above, but from the presence of alternative sources of capital available to firms that are already public and the relative absence of idiosyncratic preferences over capital structure among managers of already public firms.

7 International Comparisons

The distribution of IPO spreads varies significantly across countries. In particular, both the average spread level and the degree of clustering at the modal spread differs markedly in cross country comparisons. Torstila [2003] presents evidence on spread rigidity in many countries. In Singapore and Germany the shape of the spread distribution is roughly similar to that in the United States, although spreads are on average much lower. The Netherlands, Italy, and the United Kingdom all exhibit very little rigidity, while Hong Kong, India, and Malaysia exhibit rigidity across a wide range of values without any apparent steps. This observation is broadly consistent with the model in this paper as collusion can produce any of these outcomes depending on, among other things, the quality of information available to the underwriting industry in a country. The exact implications for relative competitiveness in each country are ambiguous. The presence of rigidity is convincing evidence of collusion, but the absence of rigidity can imply either competitive pricing or efficient collusion exploiting bid rotation and communication. However, in each of the countries mentioned above as not exhibiting rigidity, there is significant variance in the spreads charged for a *given* level of offering proceeds. This is some evidence that underwriters are exploiting information about firm preferences to extract collusive rents. Conditional variance will arise under either myopic oligopolistic competition or perfect competition, largely as a result of imperfect signals about firm value. The wide variance relative to the slope of the average spread, however, is weak evidence for the hypothesis that underwriters are engaged in efficient collusion.

The relatively low spreads charged outside of the United States also has ambiguous implications for collusion. Lower average spreads in countries where collusion is apparent imply that

the value of access to public capital markets is lower, which will drive the average reservation spread of private firms lower. Given the relative importance of public capital markets in the United States, it is not surprising that firms in other countries are generally charged spreads lower than 7%. In particular, in countries where bank financing is an important source of capital for firms, it is unlikely that many firms would be willing to go public at such high spreads. If, on the other hand, spreads were determined competitively, the cross-country difference in spreads would primarily imply differences in costs to underwriters for running IPOs. It is not at all clear that this activity should be more expensive in the U.S., a country with relatively efficient and well functioning capital markets and regulatory infrastructure.

8 Policy Implications

Collusion in the pricing of IPO services is costly from a social perspective. Inefficiently high spreads result in fewer IPOs than would be optimal, and the incentives to start positive expected value firms will be suppressed if much of the value of taking the company public, which seems to represent a significant fraction of the total value of the firm (Ritter [1987]), is captured by the underwriters.

Regulatory intervention to reduce or eliminate collusion would therefore seem desirable. A natural if blunt approach to such regulation would involve imposing a binding cap on the spread charged for IPOs. Such a regulation might prove unwise if, as conjectured in Chen and Ritter [2000], small firms are charged high spreads in order to cover the fixed costs associated with taking a firm public. In such a circumstance, small firms would be effectively excluded from holding IPOs, even when going public is quite valuable to the firm. The model presented here indicates, however, that the higher spreads charged to firms with low value are driven in part, and perhaps for the most part, by “demand” effects rather than “supply” effects. Underwriters charge higher spreads to small firms in order to capture relatively large fees from those firms whose idiosyncratic preferences for going public are large. The small value of the firm effectively magnifies the importance of idiosyncratic preferences to the manager, relative to his concern about minimizing the fee. As long as underwriters charge higher spreads to small firms for this reason, a price cap will not drive small firms out of the market and will instead prove even more

beneficial to those firms.

9 Conclusion

We have now seen that the empirically observed distribution of IPO spreads can best be explained by assuming that underwriters collude on price but receive noisy signals about firm value and preferences. Attempts to reach optimal collusion in a symmetric perfect public equilibrium are doomed to failure by the requirement that punishments occur along the equilibrium path, but underwriters will still capture significant rents. Collusion, price rigidity, and underpricing are closely linked, as underwriters exploit underpricing to extract additional surplus from firms. The model suggests links between unobservable variables relevant for underwriting and observed data, providing a blueprint for a structural analysis of the industry. Such analysis could guide regulatory policy designed to address the social costs associated with collusion in investment banking.

References

- D Abreu, D Pearce, and E Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, September 1990.
- D. Abreu, P. Milgrom, and D. Pearce. Information and timing in repeated partnerships. *Econometrica*, 59:1713–33, 1991.
- Dilip Abreu, David Pearce, and Ennio Stacchetti. Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory*, 39(1):251 – 269, 1986.
- Reena Aggarwal. Stabilization activities by underwriters after initial public offerings. *The Journal of Finance*, 55(3):1075–1103, 2000.
- Masaki Aoyagi. Bid rotation and collusion in repeated auctions. *Journal of Economic Theory*, 112(1):79–105, September 2003.
- Masaki Aoyagi. Efficient collusion in repeated auctions with communication. *Journal of Economic Theory*, 134(1):61 – 92, 2007.

- S Athey, K Bagwell, and C Sanchirico. Collusion and price rigidity. *Review of Economic Studies*, 71(2):317–349, 2004.
- Susan Athey and Kyle Bagwell. Optimal collusion with private information. *The RAND Journal of Economics*, 32(3):428–465, 2001.
- Sreedhar T. Bharath and Amy K. Dittmar. Why do firms use private equity to opt out of public markets. Working paper, March 2007.
- James Booth and Lena Chua. Ownership dispersion, costly information, and IPO underpricing. *Journal of Financial Economics*, 41(2):291–310, 1999.
- Archishman Chakraborty, Simon Gervais, and Bilge Yilmaz. Security design in initial public offerings. October 2008.
- Thomas Chemmanur and Paolo Fulghieri. A theory of the going-public decision. *Review of Financial Studies*, 12(2):249–279, Summer 1999.
- Hsuan-Chi Chen. Competition and collusion in the IPO market. September 2001.
- Hsuan-Chi Chen and J.R. Ritter. The seven percent solution. *The Journal of Finance*, 55(3): 1105–1131, June 2000.
- Peter M. DeMarzo, Ilan Kremer, and Andrzej Skrzypacz. Bidding with securities: Auctions and security designs. *American Economic Review*, 95(4):936–959, September 2005.
- Galina Hale and Joao Santos. Evidence on the costs and benefits of bond ipos. 2006.
- Makoto Hanazono and Huanxing Yang. Collusion, fluctuating demand, and price rigidity. *International Economic Review*, 48(2):483–515, 2007.
- Robert S. Hansen. Do investment banks compete in IPOs?: The advent of the ‘7% plus contract’. *Journal of Financial Economics*, 59(3):313–346, March 2001.
- K.L. Judd, S. Yeltekin, and J. Conklin. Computing supergame equilibria. *Econometrica*, 71(4): 1239–1254, 2003.

- Michihiro Kandori and Hitoshi Matsushima. Private observation, communication and collusion. *Econometrica*, 66(3):627–652, May 1998.
- Laurie Kirgman, Wayne H. Shaw, and Kent L. Womack. The persistence of IPO mispricing and the predictive power of flipping. *The Journal of Finance*, 54(3):1015–1044, 1999.
- Laurie Krigman, Wayne H. Shaw, and Kent L. Womack. Why do firms switch underwriters? *Journal of Financial Economics*, 60:245–284, 2001.
- Tim Loughran and Jay Ritter. Why has IPO underpricing changed over time? *Financial Management*, 33(3):5–37, 2004.
- Michelle Lowery and Susan Shu. Litigation risk and IPO underpricing. *Journal of Financial Economics*, 65(3):309–355, 2002.
- George J. Mailath and Larry Samuelson. *Repeated Games and Reputations: Long-Run Relationships*. Oxford UP, 2006.
- Marco Pagano, Fabio Panetta, and Luigi Zingales. Why do companies go public? an empirical analysis. *The Journal of Finance*, 53(1):27–64, 1998.
- Nagpurnanand R. Prabhala and Manju Puri. December 1998.
- Jay Ritter. Investment banking and securities issuance. In G. Constantinides, M. Harris, and R. Stultz, editors, *Handbook of the Economics of Finance*. Elsevier, 2003.
- Jay Ritter. The costs of going public. *Journal of Financial Economics*, 19(2):269–281, December 1987.
- J Rotemberg and G Saloner. A supergame-theoretical model of price wars during booms. *American Economic Review*, 76:390–407, 1986.
- SEC. SEC charges CSFB with abusive IPO allocation practices; CSFB will pay 100 million to settle SEC and NASD actions: Millions in IPO profits extracted from customers in exchange for allocations in “hot” deals. press release, January 2004. Available from: <http://www.sec.gov/news/headlines/csfbipo.htm>.

Sami Torstila. The clustering of ipo gross spreads: International evidence. *Journal of Financial and Quantitative Analysis*, 38(3):673–694, September 2003.

A Finite Approximation

In this appendix, I consider finite approximations to the continuous game described in the main text. Such finite games will exhibit the bang-bang property, whereas it is possible that there is some equilibrium of the continuous version of the game that provides better payoffs than the best bang-bang equilibrium. The arguments below presuppose familiarity with sections 3 and 4.

There are two complications that must be considered when treating the continuous stage game discussed in the main text as an approximation to a stage game with a finite but arbitrarily large action space. First, it is necessary to show that the payoffs to the optimal rigid and partially rigid spread schedules are approximated by the payoffs in the continuous version, and that the optimal strictly decreasing spread function is an appropriate approximation to the payoffs to a separating equilibrium in the finite game. The first two of these requirements are trivial. The third is less so because it is necessary to define the analog in the finite case of a strictly decreasing spread schedule. To address this, I consider a discretization of the following form. Spreads can be chosen from a finite subset of $[0, M]$ with cardinality N , and the spread schedule must be a decreasing step function. Furthermore, no step can have measure greater than ι . To approach the continuous case, let $M \rightarrow \infty$, $N \rightarrow \infty$, and $\iota \rightarrow 0$. As long as these parameters change at an appropriate rate relative to each other, this approximation converges to the continuous game with a strictly decreasing spread schedule in the following sense. As $N \rightarrow \infty$ (at a rate fast relative to the increase in M), the underwriters cease to be constrained by the finiteness of the action space, while $\iota \rightarrow 0$ implies that they cannot exploit spread rigidity to reduce the probability of triggering punishment. Thus, in the limit, the underwriters will choose a spread that effectively uses all information available in the signals without exploiting rigidity to decrease the probability of on equilibrium path punishments. The maximum payoff to such an equilibrium employing bang-bang strategies converges to the maximum payoff

to the optimal equilibrium supported with bang-bang strategies in the continuous game, and thus the upper bound on payoffs to an appropriate discretization of the game converges to the upper bound on the payoffs in bang-bang equilibria in the continuous case. Specifically, if there is some equilibrium of the continuous game in which payoffs are higher than the best bang-bang equilibrium of the continuous game, in an arbitrarily fine discretization of the game there is some profitable deviation from such an equilibrium since such an equilibrium cannot be supported with bang-bang strategies. This is *not* a contradiction of the claim that the discretization converges to the continuous game in the appropriate sense. Assuming there is some better-than-bang-bang equilibrium in the continuous game, there may be some deviation from the equilibrium that provides the same payoff to the deviator as adhering. In this case, the payoffs to the discrete analog of this deviation may give a higher payoff than adhering to the discrete analog of the better-than-bang-bang equilibrium. As the discretization becomes increasingly fine, the payoffs from the deviation must converge to the payoffs from adhering, but they are only equal in the limit and *not* for any actual finite game.

The second concern when evaluating a discrete version of the game is that the results from section 3 must be interpreted more carefully. These results depended effectively on showing that spreads will not be locally flat. In any finite approximation spreads will clearly have flat regions by construction; in fact, the spread schedule will be a step function. The appropriate approximation here is that as the cardinality of the finite action space increases, the number of steps will increase, and as the cardinality goes to infinity so will the number of steps. Furthermore, the measure of the set of signals generating any single spread will go to zero. This is the appropriate sense in which spreads will not be rigid in finite approximations to the continuous game under competition, competitive oligopoly, or monopoly.

That this statement is true follows from the fact that, for any finite strategy with underwriters attempting to reach the imposed criteria (zero profits, one-shot equilibrium, or optimal monopoly), firms would always “deviate” from prescribed actions with finite support *if* they were not restricted to a finite action space. For example, if firms chose actions such that they would be as “competitive” as possible, then a restriction to a finite action space would imply that expected profit would be small but positive for some signals and small but negative for

others.¹⁹ Thus, when the action space expands by a sufficiently rich set of additional spreads, the most competitive spread will have more steps than when the action space was smaller.

Further details on these approximations are available from the author on request.

B Competitive and Monopoly Spreads: General Model

This section explores in a more general framework the spread schedules that would arise in the absence of collusion of the form posited. Specifically, we assume that the common value benefits of holding the IPO take the form of a potentially nonlinear function $\tilde{\beta}$, where $\tilde{\beta}'(x) > 0$ for all x . I also consider a general, weakly increasing cost function $C(x)$. Spreads are considered under three alternative assumptions about the process for bidding for an IPO. First, underwriters are assumed to bid in a “competitive” market for underwriting services such that they expect zero profit following every realization of the signal. Second, underwriters bid strategically as in an oligopoly without collusion. Finally, underwriters act effectively as a monopoly.²⁰ I consider the conditions necessary to generate intervals over which a fixed spread is charged. In each of these cases the conditions for flat spreads over some interval of firm value are shown to rely on very specific relationships among the costs to the underwriter of holding an IPO, the common value benefits, and the distribution of idiosyncratic preferences. As such, a slight perturbation of any one of these functions will eliminate the region of flat spreads. This requirement that the forms of the functions representing the structural elements of the underwriting industry exhibit remarkable coincidences²¹ in order to generate even small regions of fixed spreads will later be seen to contrast with the simple and robust setting in which collusion with imperfect information leads to partially rigid spreads that are qualitatively consistent with observed spreads.

Throughout this section we confine attention to symmetric, pure strategies that call for weakly decreasing spreads. While these restrictions are in part for simplicity, the distribution of spreads generated by asymmetric, increasing, or mixed strategies would not be compatible

¹⁹Here we ignore the irrelevant fact that firms could guarantee themselves zero profit by charging ∞ , which is clearly not descriptive of the data and not in the spirit of the competitive assumption.

²⁰Note that of these three solution concepts, only competitive oligopoly represents an equilibrium of the game as defined above.

²¹The argument is that spreads will respond to information generically. Defining the appropriate sense of “generic” and formally proving the statement are tedious, uninteresting, and add little beyond the intuition presented here. Details are available from the author.

with observed data.

With these restrictions, I can consider a general, weakly decreasing spread schedule $\alpha(\cdot)$ where spreads are rigid over some interval $[x^d, x^u] = \mathcal{I} \subset X$. That is, for all $\xi \in \mathcal{I}$, $\alpha(\xi) = \alpha \in \mathbb{R}^+$. The payoff to one underwriter following a signal in $\xi \in \mathcal{I}$ conditional on being among the lowest bidders can then be given in each case by

$$\mathbf{\Pi}(\xi, \alpha) = \mathcal{A}(\xi)E[\mathcal{H}(\alpha, x)(\alpha\tilde{\beta}(x)-C(x))|\xi, \xi' \in \mathcal{I}] + (1-\mathcal{A}(\xi))E[\mathcal{H}(\alpha, x)(\alpha\tilde{\beta}(x)-C(x))|\xi, \xi' < x^d] \quad (\text{B.1})$$

where ξ' is the signal observed by the other underwriter, $\mathcal{A}(\xi) = \frac{P(\xi' \in X|\xi)}{P(\xi' \in X|\xi) + P(\xi' < x^d|\xi)}$, and $\mathcal{H}(\alpha, x) = P(\varepsilon > x - (1-\alpha)\tilde{\beta}(x))$. These expressions are, respectively, the probability of making the same spread offer as the other underwriter and being chosen to hold the IPO conditional on either making the lowest spread offer or making the same spread offer and being selected at random,²² and the probability that the firm will agree to go public given a spread α and firm value x . This expression for expected revenue can be decomposed into $\mathbf{\Pi}^T$, the part of the expected profit conditional on winning that accrues when the auction is a ‘‘tie,’’ and $\mathbf{\Pi}^S$, the part that accrues when the underwriter wins outright. The values $\mathbf{\Pi}$ and $\mathbf{\Pi}^T$ will provide the conditions necessary for spreads to be rigid under competition and competitive oligopoly; a related function will provide the conditions for monopoly.

B.1 Competitive Spreads

Here, we assume that following every signal about firm value the underwriter sets his spread such that he will earn, in expectation, zero profit. That is, $\mathbf{\Pi}(\xi, \alpha) = 0$ for all $\xi \in X$. Under the assumption that spreads are the same for every $\xi \in \mathcal{I}$, there is an additional restriction that $\frac{\partial \mathbf{\Pi}}{\partial \xi} = 0$ for all $\xi \in \text{int}(\mathcal{I})$. In fact, $\frac{\partial^n \mathbf{\Pi}}{\partial^n \xi} = 0$ for all n and all $\xi \in \text{int}(\mathcal{I})$. It is then clear that if signals are informative about value and value is informative about the probability that a firm will agree to the IPO at a given spread, flat regions can only occur when there is a remarkable coincidence among the costs of holding an IPO, the common value benefits from an IPO, and distribution of manager preference.

²²Here, we use the convention that if both underwriters offer the same spread, the firm first chooses between the two underwriters and *then* consults its participation constraint. This is just a convenience.

The following proposition notes one such relationship that would generate flat spreads; the proof is an immediate consequence of equation B.1 and is therefore omitted:

Proposition B.1. *Spread offers are invariant in the signal about firm value over the interval $\mathcal{I} \subset X$ if*

$$C(x) = \gamma\tilde{\beta}(x).$$

for all x that occur with positive probability following signals in \mathcal{I} .

The proposition states that if the costs to the underwriter of arranging the IPO are exactly proportional to the common value benefits of the IPO, then flat spreads will be observed. This strict proportionality imposes strong restrictions on the relationship between the cost function of an underwriter, which will depend on the costs of providing analyst coverage and engaging in roadshows, and the benefits of being a public firm, which are related to improved access to capital markets. Relaxing this strict proportionality will then effectively require that $\frac{1-A}{A}$, which is determined entirely by the joint distribution of firm value and signals, and $\mathcal{H}(x, \alpha)$, which is determined by the distribution of the preference of the manager, to exactly adjust for the lack of proportionality in costs and benefits, an extremely unlikely event. This simple case thus demonstrates why it is so unlikely to observe flat spreads in a competitive market for IPO services.

B.2 Oligopolistic Competition

When the two firms bid for the right to hold the IPO and seek to maximize their current payoff (i.e. they do not attempt to collude), spreads are again unlikely to contain regions with flat areas. In this case, it is possible to derive some simple necessary conditions for the existence of an interval \mathcal{I} over which the spread schedule is flat. These conditions, however, are not sufficient for a the spread schedule to be flat in the interval.

Proposition B.2. *A necessary condition for the spread schedule to be flat but greater than zero over the interval X is that $\mathbf{\Pi}^T(\xi, \alpha) = 0$.*

Proof. If spreads are flat over some interval in equilibrium, a deviation to an infinitesimally higher spread guarantees that the deviator never wins when both signals are in said interval,

and an infinitesimal deviation to a lower spread guarantees that the deviator always wins when both signals are in said interval. For both of these deviations to be unprofitable, it must be the case that the expected profit conditional on both signals being in the interval be zero at α (since the change in profit accruing in all events other than $\xi, \xi_i \in \mathcal{I}$ changes infinitesimally following an infinitesimal deviation). \square

The above condition can be satisfied in two ways; either costs are exactly proportional to benefits such that $C(x) = \gamma\tilde{\beta}(x)$ for all x with positive support following $\xi, \xi' \in X$, in which case $\alpha = \gamma$ satisfies this condition, or $\mathcal{H}, \tilde{\beta}$, and C have a particular dependence on α and x such that this condition is met for some α . Neither of these events are likely in the sense that small perturbations of any of the functions would undo the necessary relationship. Furthermore, this condition is effectively independent of the differential equation that will determine the symmetric equilibrium in the auction. Assuming that equilibrium spreads are in fact flat over \mathcal{I} , and thus the expected profit accruing conditional on both signals falling in \mathcal{I} is zero, the optimal bid for an underwriter with a signal in \mathcal{I} will be determined by maximizing expected profit conditional on the other signal not being in \mathcal{I} and on the equilibrium spread schedule outside of the interval. Not only must this happen to give the α that satisfies the zero profit condition for one particular signal in \mathcal{I} , it must do so for *all* signals in \mathcal{I} , an extremely unlikely event when signals are at all informative.

B.3 Monopoly

Several authors (for example Chen [2001]) have argued that fixed spreads can arise when underwriters collude on the optimal monopoly spread. In this section I argue that the optimal monopoly spread is unlikely to be characterized by fixed spreads over a significant region. It is important to draw a distinction between the argument that spreads are rigid because underwriters collude on a monopoly spread that happens to be rigid and the argument that rigid spreads arise in a second-best equilibrium where collusion at monopoly spreads is impossible. First, as I will show below, monopoly spreads will only be rigid for very specific environments that are unlikely to describe the IPO process. As such, explanations of spread rigidity based on monopoly pricing do not really provide theoretical support for the claim that spread rigidity is

evidence of collusion rather than competition; both monopoly pricing and competition generate rigid spreads in effectively non-generic environments. Furthermore, optimal policy responses to collusion may differ depending on whether the market is effectively a monopoly or only imperfectly collusive. Finally, the empirical implications for costs and benefits in the IPO industry of the data on spreads will depend crucially on the exact nature of the collusion in the industry. Attempting to recover empirical facts relating to the costs to underwriters and the benefits to firms of holding IPOs, while a daunting task, is worthwhile both because it would inform policy decisions and because few alternatives for measuring these quantities are available to financial economists.

In this section, I consider the case of a disaggregated monopoly, where the two underwriters each receive signals and set spreads to maximize joint profits but where underwriters cannot communicate their signals to each other before bidding.²³ This schedule must by construction satisfy the condition that, following a given signal, the underwriter receiving the signal cannot unilaterally “deviate” to an infinitesimally different spread and, by doing so, increase the expected *total* profits of the two underwriters. That is, defining $\tilde{\mathcal{A}}(\xi)$ as the probability of the event $\xi' \in \mathcal{I}$ conditional on ξ' either in \mathcal{I} or less than x^d , it must be the case that for all²⁴ $\xi \in \mathcal{I}$:

$$\begin{aligned} \frac{d}{d\alpha} \{ & (1 - \mathcal{A}(\xi))E[\mathcal{H}(\alpha, x)(\alpha\tilde{\beta}(x) - C(x))|\xi, \xi' < x^d] \\ & + \mathcal{A}(\xi)E[\mathcal{H}(\alpha, x)(\alpha\tilde{\beta}(x) - C(x))|\xi, \xi' \in \mathcal{I}] \} \\ & \geq 0 \end{aligned}$$

That is, an infinitesimal decrease (to below α , the rigid spread over \mathcal{I}) in the spread offer following a signal $\xi \in \mathcal{I}$ must decrease the total profit to both underwriters conditional on $\{\xi' \in \mathcal{I}\} \cup \{\xi' < x^d\}$. (The marginal decrease is irrelevant if the other underwriter would have won the IPO at a lower spread than α before the deviation.) The above expression follows from two facts: after a marginal decrease in the spread on the rigid interval, the total increase in

²³Permitting communication will not change the conclusions of this section in any meaningful way, although the exact environments where fixed spreads will arise will differ.

²⁴Technically, except on a set that arises with probability zero; we ignore such concerns.

profits conditional on winning must be negative, and a deviator to a marginally smaller spread following a signal that calls for the rigid spread now captures all of IPO's where the other signal also calls for the rigid spread.²⁵

If the above inequality does not bind for almost every $\xi \in \mathcal{I}$, then there is some interval over which a marginal decrease in spreads can increase total profits, contradicting the optimality of $\alpha(\cdot)$. But, when the inequality does bind almost everywhere, we again have a condition that cannot hold in the generic sense described in the discussion of competition and non-collusive oligopoly.

Spreads will, however, be rigid in the case where firm value is not informative about the probability that a firm will accept a given spread offer. Specifically, if $\varepsilon = 0$ for all firms and $\tilde{\beta}(x) = \beta x$, a monopolist would charge exactly

$$\alpha = 1 - \frac{1}{\beta}$$

regardless of his signal, since this would guarantee that every firm exactly met its participation constraint, regardless of its true value.

Rigid spreads over some interval could also arise when manager preference is bounded from below. It is then possible to find cases where, for a large enough signal, underwriters will set spreads to guarantee that all firms go public, regardless of manager preference or firm value. That is, for signals that indicate that the distribution of firms is skewed toward high values, the monopolist underwriter might prefer to set a spread that guarantees that the largest possible firm goes public. The costs of failing to serve such firms may be so high as to make marginal differences in earnings elsewhere irrelevant, and the spread may thus be rigid exactly at the level that guarantees that the largest possible firm will agree to the IPO.

These two cases both involve a discontinuity in the distribution of ε so that the condition described above does not eliminate them. The second case also presents the most compelling

²⁵The analogous condition to guarantee that the underwriter will not increase his spread is just

$$\frac{d}{d\alpha} \left\{ (1 - \tilde{A}(\xi)) E[\mathcal{H}(\alpha, x)(\alpha\tilde{\beta}(x) - C(x)) | \xi, \xi' < x^d] \right\}$$

since, by raising the spread, the underwriter affects profits only when the other underwriter would have charged a spread higher than the rigid spread.

alternative explanation for spread rigidity as it will involve decreasing spreads for small firms and flat spreads for large firms. Without extreme assumptions on the distribution of value and signals about value, however, it cannot generate the specific pattern observed in data, where spreads appear to step up for small firms rather than smoothly decreasing toward the flat spread. Furthermore, this explanation relies on the existence of a largest possible firm value and on significant mass being concentrated around that value.

C Perfectly Observable Value and Preference

To analyze the game under perfect information, I will proceed under the following additional assumptions:

Assumption 4.

1. x and ε are common knowledge at the beginning of every period.
2. $E[\varepsilon] = 0$
3. Underwriters are sufficiently patient to sustain any degree of collusion.

The first assumption is just a restatement of the condition that both agents receive perfect signals at the beginning of the period. The second assumption is just for simplicity.

The spread that will be charged if the two underwriters can collude perfectly and set the monopoly price is then given by the following proposition:

Proposition C.1. *When underwriters act as a monopoly, if*

$$\varepsilon \geq \kappa - (\beta - 1)x$$

the spread function is given by

$$\alpha^* = 1 - \frac{1}{\beta} + \frac{\varepsilon}{\beta x}$$

and the firm chooses to hold the IPO. Otherwise,

$$\alpha \geq 1 - \frac{1}{\beta} + \frac{\varepsilon}{\beta x}$$

and the firm does not hold the IPO.

Proof. Since the firm is short-lived, he will accept the lowest spread offer as long as his residual claim on the profits to the public firm plus his idiosyncratic private value for going public exceed the private value of the firm. Underwriters will force the firm to its participation constraint:

$$\begin{aligned}(1 - \alpha)\beta x - x + \varepsilon &= 0 \\ \alpha &= 1 - \frac{1}{\beta} + \frac{\varepsilon}{\beta x}.\end{aligned}$$

But, underwriters will only make such an offer if it is profitable to them. This condition is given by

$$\begin{aligned}\left(1 - \frac{1}{\beta} + \frac{\varepsilon}{\beta x}\right)\beta x - \kappa &> 0 \\ \varepsilon &> \kappa - (\beta - 1)x.\end{aligned}$$

□

I now consider how spreads behave for very small and very large firms. On average, very small firms will face arbitrarily high spreads, while larger firms face spreads that converge to the spread that would be charged to a firm with no idiosyncratic preferences for going public.

Proposition C.2. *The expected spread as firms become infinitely valuable and as firms cease to have any value are given by:*

$$\lim_{x \rightarrow \infty} E[\alpha | x, \varepsilon > \kappa - (\beta - 1)x] = 1 - \frac{1}{\beta} \tag{C.1}$$

$$\lim_{x \rightarrow 0} E[\alpha | x, \varepsilon > \kappa - (\beta - 1)x] = \infty. \tag{C.2}$$

Proof.

$$\begin{aligned}E[\alpha | x, \varepsilon > \kappa - (\beta - 1)x] &= E\left[1 - \frac{1}{\beta} + \frac{\varepsilon}{\beta x} \mid x, \varepsilon > \kappa - (\beta - 1)x\right] \\ &= 1 - \frac{1}{\beta} + \frac{1}{\beta x} E[\varepsilon | \varepsilon > \kappa - (\beta - 1)x].\end{aligned}$$

Since $\kappa - (\beta - 1)x$ is decreasing in x and covers the real line, we can conclude that $\lim_{x \rightarrow \infty} E[\varepsilon | \varepsilon > \kappa - (\beta - 1)x] = 0$, while $\lim_{x \rightarrow 0} \frac{1}{x} E[\varepsilon | \varepsilon > \kappa - (\beta - 1)x] = \infty$ since $E[\varepsilon | \varepsilon > \kappa]$ is positive, and, more generally, $\frac{d}{dx} E[\varepsilon | \varepsilon > \kappa - (\beta - 1)x] > 0$ since $\beta > 1$. Thus,

$$\begin{aligned} \lim_{x \rightarrow \infty} E[\alpha | x, \varepsilon > \kappa - (\beta - 1)x] &= 1 - \frac{1}{\beta} \\ \lim_{x \rightarrow 0} E[\alpha | x, \varepsilon > \kappa - (\beta - 1)x] &= 1 - \frac{1}{\beta} + \infty \\ &= \infty, \end{aligned}$$

establishing the proposition. □

Corollary C.3. *The variance of the spread conditional on firm value disappears for large firm values.*

Proof. Direct calculation of the variance in spreads conditional on x gives

$$\frac{1}{\beta^2 x^2} E[\varepsilon^2].$$

□

The implications of the above results are as follows. First, without idiosyncratic preferences all firms would be charged exactly $1 - \frac{1}{\beta}$. This is exactly the spread that captures all of the “common value” of the IPO process, regardless of the size of the firm. However, this is not a complete explanation for the concentration of spreads at seven percent; rigidity would not be robust to the introduction of some small degree of idiosyncratic manager preference, and spreads would not rise at all for small firms, as they do in the data. With the introduction of manager preferences, spreads do depend on the value of the firm, but in such a way that the dependence disappears as firms grow large but matters a great deal for the smallest firms.

I will now address why spreads decrease for the most valuable firms.

C.1 Partial Collusion with Impatient Firms

When underwriters are impatient, it will not in general be possible to sustain optimal collusion for all firm values. When $\beta > 1$, more valuable firms provide, on average, more profitable opportunities for collusion. Consequently, when a very valuable firm enters and underwriters are insufficiently patient, they will have an incentive to deviate from an equilibrium that calls for optimal collusion. Thus, to sustain an equilibrium, spreads must decrease for such firms so that a deviation will not be too profitable relative to expected future profits from maintaining collusion. This is, of course, an application of the result in Rotemberg and Saloner [1986].

When ε has distribution F and associated density f , and x has distribution G with density g over support $[\underline{x}, \bar{x}]$,²⁶ the optimal spread can now be expressed as follows:

Proposition C.4. *When underwriters are impatient and the upper bound on the value of firms is sufficiently large, the optimal collusive duopoly spread is given by*

$$\alpha = \begin{cases} \alpha^* & \text{if } \Pi^m(x, \varepsilon) \leq \bar{\Pi} \\ \frac{\bar{\Pi} + \kappa}{\beta x} & \text{otherwise,} \end{cases}$$

where α^* is again the collusive spread schedule when firms are perfectly patient, $\Pi(x, \varepsilon)$ is the equilibrium profit accruing to the underwriter in a period with a firm of type (x, ε) , $\Pi^m(x, \varepsilon)$ is the profit that would accrue to an underwriter if he forces the firm to its participation constraint, and $\bar{\Pi}$ is the largest value satisfying the condition

$$\bar{\Pi} = \sum_{t=1}^{\infty} \delta^t \left(\int_{\underline{x}}^{\bar{x}} \int_{-\infty}^{\infty} \Pi(x, \varepsilon) f(\varepsilon) g(x) d\varepsilon dx \right).$$

The proof of this proposition is standard and is omitted.

The following parametric example highlights the implications of the results in this section. Let $x \sim U[0, 1000]$ and $\varepsilon \sim N(0, 1)$ and set $\beta = 1.075$ and $\kappa = 1$. Note that this β implies $\lim_{x \rightarrow \infty} \alpha = 1 - \frac{1}{\beta} = 0.07$, the seven percent spread pervasive for IPOs. Finally, assume that δ is such that $\bar{\Pi}$, the highest level of profits that can be sustained through collusion, is 50. We

²⁶The existence of a density function for either x or ε is not necessary but is assumed to simplify notation.

can now, for specific draws of x and ε , calculate the equilibrium spread. Note that, despite the uniform distribution of x , fewer IPO's will be observed at low values of x since small firms will not find IPO's valuable unless they get an improbably high draw of ε .

This pattern contrasts with the spreads that would be observed if the underwriting industry were competitive. In this case, patience or impatience would be irrelevant, and spreads would be given by:

Proposition C.5. *Competitive underwriters charge spreads*

$$\alpha = \frac{\kappa}{\beta} \frac{1}{x}.$$

The above is an immediate consequence of the zero profit condition for competition, since it must be the case that $\alpha\beta x = \kappa$. Note that the presence of idiosyncratic manager preference is irrelevant and spreads smoothly decline toward zero for the most valuable firms. The conditional variance of the spread charged will, of course, be zero over the entire distribution of values.

D Perfectly Observable Value and Unobservable Preference

The specific environment considered in this section is described formally in the following assumption:

Assumption 5.

1. x is common knowledge at the beginning of every period.
2. Both underwriters observe F , the unconditional distribution of ε , but receive no other signal about ε .
3. $\varepsilon \sim U[-\eta, \eta]$
4. Underwriters are sufficiently patient to sustain any degree of collusion.
5. Both underwriters are risk-neutral.
6. Underwriters cannot demand a spread greater than 1.

Since there is no information about ε contained in the signals to the underwriters, a stationary, symmetric, public pure strategy will be a function mapping signals to spread offers. And, for any strategy, the public history is sufficient to identify deviations. Therefore, the underwriters can collude on the monopolist spread. I now derive the spread function implied by the above assumptions and then discuss the implications.

Proposition D.1. *The optimal collusive equilibrium spread with symmetric imperfect information is $\alpha^* = 1$ if $x \leq \frac{\eta+\kappa}{1+\beta}$. Otherwise, the optimal spread is given by*

$$\alpha^*(x) = \begin{cases} 1 - \frac{1}{\beta} + \frac{\eta}{\beta x} & \text{if } x \leq \frac{\kappa-\eta}{\beta-1} \\ \frac{1}{2} \left(1 - \frac{1}{\beta} + \frac{\eta+\kappa}{\beta x} \right) & \text{if } x \in \left[\frac{\kappa-\eta}{\beta-1}, \frac{\kappa+3\eta}{\beta-1} \right] \\ 1 - \frac{1}{\beta} - \frac{\eta}{\beta x} & \text{if } x \geq \frac{\kappa+3\eta}{\beta-1} \end{cases}$$

The proof is by standard optimization and is omitted.

The first thing to observe is that, unsurprisingly, $\alpha^*(x) \rightarrow 1 - \frac{1}{\beta}$ as $x \rightarrow \infty$ since the constraint set collapses to $1 - \frac{1}{\beta}$. This occurs because, as x grows large, the idiosyncratic element of preferences for IPO's becomes relatively unimportant; a firm will accept an offer of a spread slightly below $1 - \frac{1}{\beta}$ with probability increasing toward 1. While the assumption of bounded support for ε makes this effect particularly dramatic, the intuition will hold for virtually any distribution of ε independent of x . Even if mass is concentrated in the negative tails, indicating that most entrepreneurs prefer to keep their firms private, this mass will eventually constrict into a tight region around 0 as x grows and financial benefits become the overwhelming concern.

Furthermore, when costs are relatively small, α is declining in value for small firms. This result holds even if the costs of holding an IPO are zero. That is, the upward pressure on price as firms get very small does not result entirely, or necessarily at all, from the need to cover costs.

E Proof of Proposition 4.3

Proof. The proof is by construction. Choose some $x^* \in (0, \bar{x})$, where there is some α such that the profits accruing to the underwriter of type x^* for pooling with all types $x < x^*$ are positive

at α (such and x^*, α pair must exists by the assumption that there exists some profitable rigid spread). Now, choose α^h to maximize profits of type x^* conditional on pooling with the lower types. The profit accruing to type x^* for pooling with the types higher than x^* is clearly greater than the profit for pooling with low types at α^h , continuous in α , and reaches a minimum that is less than zero as $\alpha \rightarrow 0$. Thus, there is some α^l such that type x^* is indifferent between pooling with the low types at α^h and pooling with the high types at α^l .

Since α^h is chosen as the maximum for x^* , we know that the difference between the value of adhering and the value of deviating increases as ξ decreases away from x^* . So, types below x^* do not have an incentive to deviate. The same argument shows that types above x^* do not have an incentive to imitate a type below x^* .

□

F Flexible Spread Upper Bound

The procedure used to find an upper bound on the set of symmetric perfect public equilibrium payoffs is described in greater detail here.

Define

$$\begin{aligned} R(\rho, \alpha) &= \int_0^{\bar{x}} R(\rho, x, \alpha) \frac{1}{x} dx \\ R(\rho, x, \alpha) &= \int_0^{\bar{x}} \rho(x, y) dF^{FOS}(y) \\ \rho(x, y) &= P(\text{reversion} | x, y = \max \alpha^{-1}(\alpha_i)). \end{aligned}$$

That is, $R(\rho, \alpha)$ is the unconditional reversion probability given than the game reverts to repeated play of the worst one-shot equilibrium strategy with probability $\rho(x, y)$ when the public signal reveals that the true value of the firm is x and the private signal implied by the lowest bid under the spread schedule α is y , assuming that both players use spread schedule α .

Let \mathfrak{J} be the set of closed intervals on the real line. And, let \mathfrak{C} be the family of incentive

constraints:

$$\begin{aligned} & (1 - \delta)\pi(w, w) + \delta E_{z,y}[\rho(z, y)\Pi^{SG} + (1 - \rho(z, y))v | \xi_i = w, \alpha, \rho, m_i = \xi_i] \\ \geq & (1 - \delta)\pi(w, w') + \delta E_{z,y}[\rho(z, y)\Pi^{SG} + (1 - \rho(z, y))v | \xi_i = w, \alpha, \rho, m_i = w'] \end{aligned}$$

for all w and $w' > w$, where $m_i = \alpha^{-1}(\alpha_i)$; that is, m_i is the implicit report of the signal received by underwriter i when underwriters are both using the spread schedule α . Here,

$$\pi(w, w') = E_q[\mathbf{1}_{\{\xi_j < w'\}^*}(\alpha(w')\beta q - \kappa)e^{-\lambda q(\alpha(w')^{\beta+1-\beta})} | \xi_i = w, \alpha]$$

where

$$\mathbf{1}_{\{a < b\}^*} = \begin{cases} 1 & \text{if } a < b \\ \frac{1}{2} & \text{if } a = b \\ 0 & \text{if } a > b \end{cases}.$$

These incentive constraints are simply the requirement that an underwriter finds it optimal to truthfully reveal his signal through his bid rather than attempting to capture additional market share by decreasing the spread he demands.²⁷

Define $\mathfrak{B} : \mathfrak{J} \rightarrow \mathfrak{J}$ such that

$$\mathfrak{B}([a, b]) = [a, b'],$$

with

$$b' = \max_{\alpha, \rho, v^h, v^l} \pi(\alpha) - \frac{1 - \delta}{\delta} R(\rho, \alpha)[v^h - v^l]$$

subject to \mathfrak{C} and $v^h, v^l \in [a, b]$. The expression for b' is the maximized present value of the game and is the direct analogy of the expression derived in Abreu et al. [1991]. Let Π^{SG} be the per-period unconditional expected payoff from repeated play of the symmetric stage game equilibrium with the lowest payoff, and let v^* be the maximum of the set of symmetric perfect public equilibrium payoffs. We will apply the following proposition to derive an upper bound on the set of symmetric perfect public equilibrium payoffs:

Proposition F.1. \mathfrak{B} is a contraction over the domain $\{[a, b] | a \leq \Pi^{SG}, b \geq v^*\}$.

²⁷Constraints preventing the underwriter from reporting a signal higher than that which he received will not bind in the problem of finding the maximum SPPE payoff and so are ignored.

Proof. Define

$$\begin{aligned} R(\rho) &= \int_0^{\bar{x}} R(\rho, x) \frac{1}{x} dx \\ R(x, \rho) &= \int_0^{\bar{x}} \rho(x, y) dF^{FOS}(y) \\ \rho(x, y) &= P(\text{reversion} | x, y = \max \alpha^{-1}(\alpha_i)). \end{aligned}$$

Let $R^*(\alpha)$ be the minimal unconditional reversion probability required to maintain incentive compatibility for all signals when both underwriters use spread schedule α .

Let α^* represent the spread schedule used initially in the optimal symmetric perfect public equilibrium.

By lemma F.2 below, we know

$$v^* = (1 - \delta)\pi(\alpha^*) + \delta[R^*(\alpha^*)\Pi^{SG} + (1 - R^*(\alpha^*))v^*]$$

so

$$v^* = \pi(\alpha) - \frac{\delta}{1 - \delta} R^\alpha(\alpha^*)[v^* - \Pi^{SG}].$$

Thus if $a = \Pi^{SG}$ and $b = v^*$, a solution to

$$\max_{\alpha, \rho, v^h, v^l} \pi(\alpha) - \frac{\delta}{1 - \delta} R(\alpha)[v^h - v^l]$$

subject to \mathfrak{C} and $v^h, v^l \in [a, b]$, is a^*, ρ^*, v^* , and Π^{SG} .

If $a < \Pi^{SG}$ and/or $b > v^*$, the problem is exactly as above with one or two constraints relaxed.

Thus, $[a, v^*] \subseteq \mathfrak{B}([a, b])$. □

Lemma F.2. *Any symmetric perfect public equilibrium payoff can be supported with a strategy that uses only Π^{SG} and v^* as promised continuation values.*

Proof. The existence of a public correlating device permits the construction of an absolutely

continuous public signal for any discretized version of the game.²⁸ Therefore, any discrete approximation to the game will satisfy the requirements to apply the bang-bang result of Abreu et al. [1990]. Since we are considering arbitrarily fine discrete approximations to the continuous game, this is sufficient to apply the bang-bang result. \square

Having established that \mathfrak{B} is a monotone operator, it is straightforward to find an upper bound on the SPPE payoffs for a strictly decreasing spread function. By choosing a lower bound below Π^{SG} , in this case 0, and applying \mathfrak{B} to $[0, \bar{v}^0]$, we find a new upper bound \bar{v}^1 . This procedure can be repeated to find a (weakly) decreasing sequence of upper bounds. The maximization procedure is somewhat difficult because it is necessary to find α and ρ simultaneously. Since the game must be solved as the limit of a sequence of discrete approximations, this requires maximization over $3 * K + 2$ variables, where K is the number of grid points for the given discretization.²⁹ Fortunately, in the implementation studied a relatively coarse grid appears to give a good approximation for the continuous game, at least for the purposes of finding the upper bound. See figure 7 for a plot of the calculated upper bounds as a function of the fineness of the grid; the value appears to reach a maximum around 30 grid points, indicating that further increases in the number of grid points used would not lead to a significantly higher upper bound on the value of a fully separating spread.

[Figure 7 about here.]

²⁸The procedure is similar to the procedure when signals are drawn from a finite space. See, for example, Mailath and Samuelson [2006], remark 2.5.1.

²⁹The procedure used here takes into consideration that ties can occur when one or both signals are incorrect by chance. The fact that the spread schedule in the case where spreads are restricted to the rationals is a function mapping a continuum of signals into a countable set is then not a cause for concern.

List of Figures

1	Timing of the Game	59
2	Perfect Information Spreads with Collusion: $x \sim U[0, 1000]$, $\varepsilon \sim N(0, 1)$, $\beta = 1.075$, and $\kappa = 1$	60
3	Collusive Spreads with Unobservable ε	61
4	Competitive Spreads with Perfect Information	62
5	Approximate Reversion Function: $K = 30$	63
6	Optimal Two-Step Self-Enforcing Spread and Approximate Optimal Fully Separating Spread	64
7	Calculated Upper Bound as a Function of Grid Fineness	65

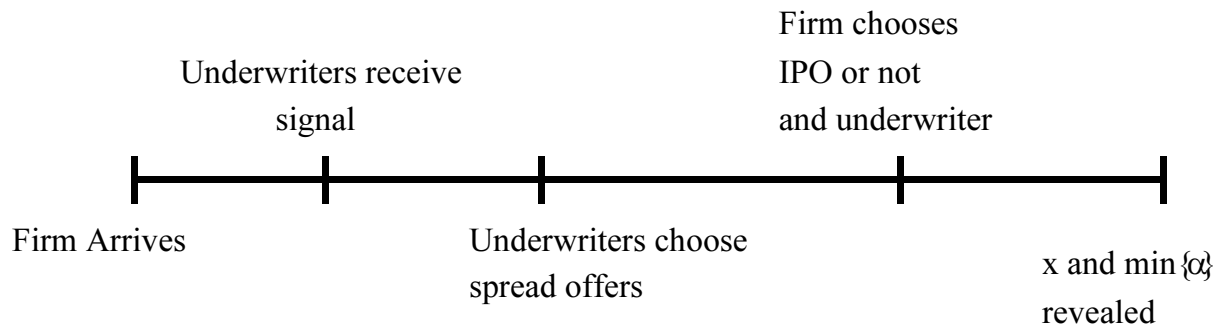


Figure 1: Timing of the Game

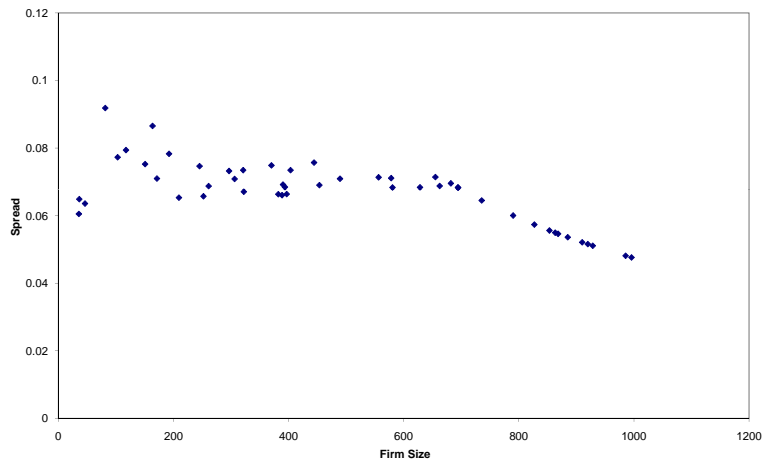


Figure 2: Perfect Information Spreads with Collusion: $x \sim U[0, 1000]$, $\varepsilon \sim N(0, 1)$, $\beta = 1.075$, and $\kappa = 1$

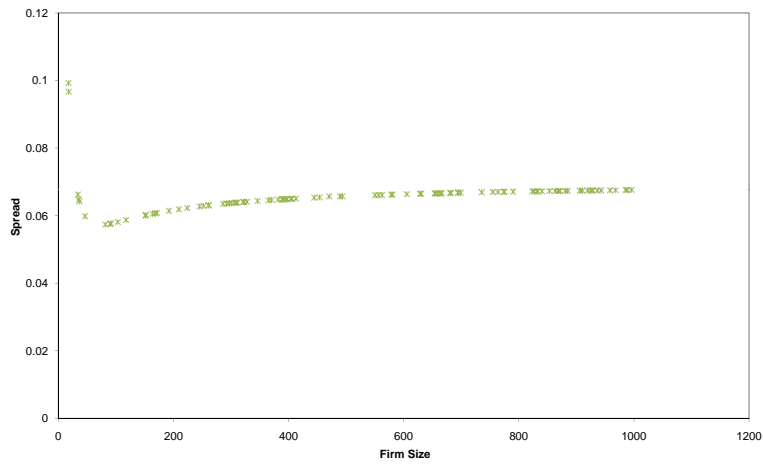


Figure 3: Collusive Spreads with Unobservable ε

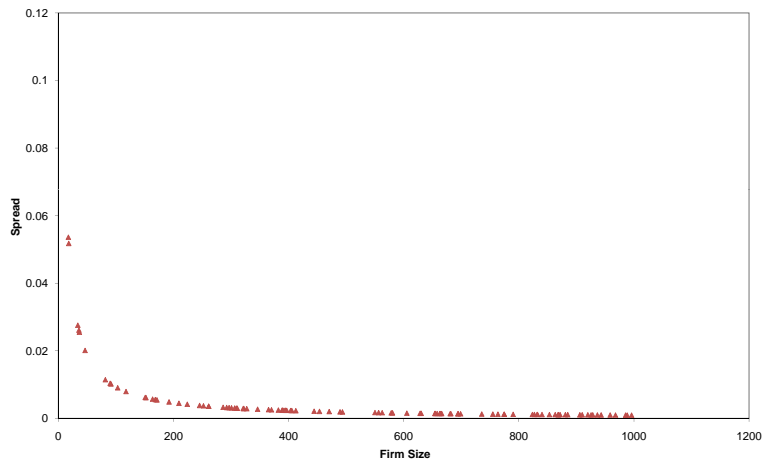


Figure 4: Competitive Spreads with Perfect Information

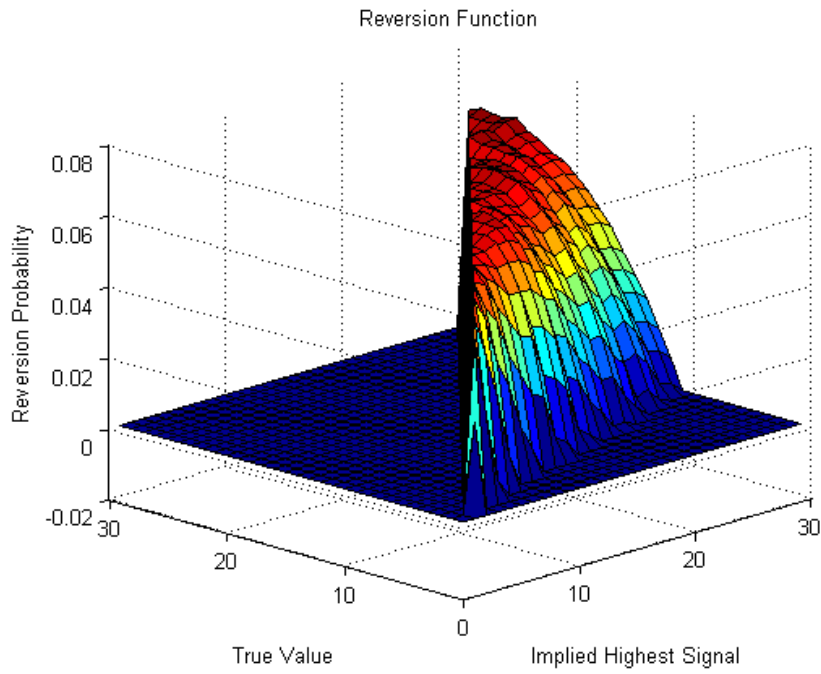


Figure 5: Approximate Reversion Function: $K = 30$

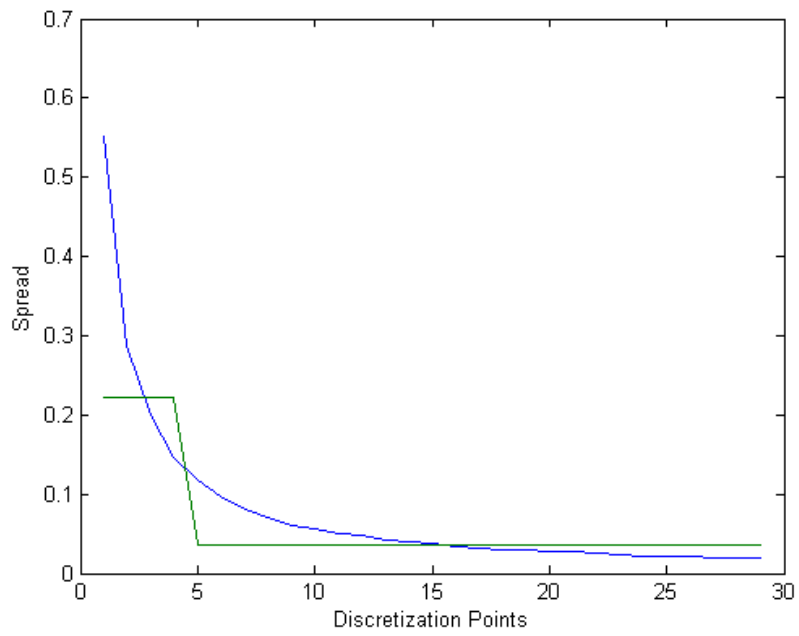


Figure 6: Optimal Two-Step Self-Enforcing Spread and Approximate Optimal Fully Separating Spread

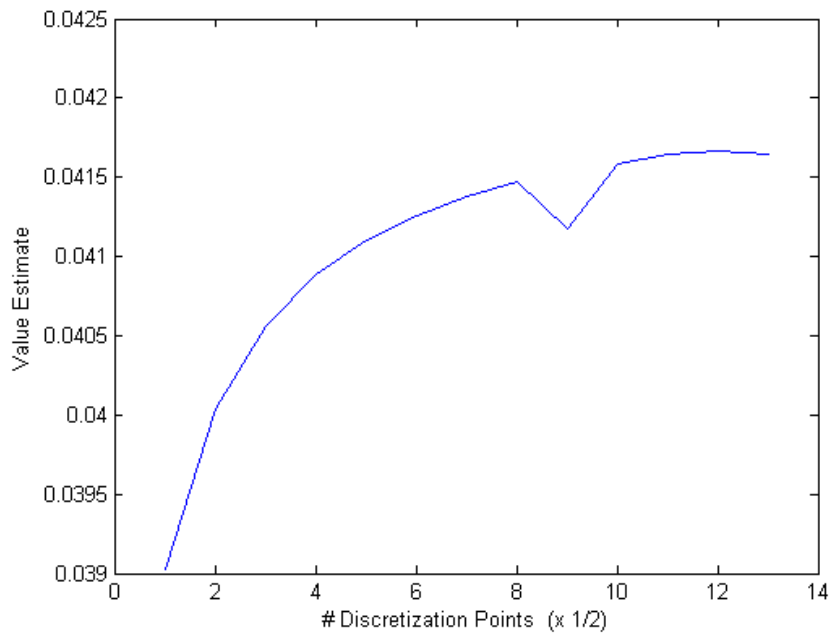


Figure 7: Calculated Upper Bound as a Function of Grid Fineness