The Journal of Industrial Economics

WILEY

ORIGINAL ARTICLE

Asset Specificity and Inefficient Bargaining: Theory and Evidence From Television Shows

Luís Cabral[1] | Gabriel Natividad[2]

[1]New York University, New York, USA | [2]Department of Economics, Universidad de Piura, Lima, Peru

Correspondence: Luís Cabral (luis.cabral@nyu.edu)

Accepted: 9 May 2025

ABSTRACT

Evidence from TV shows suggests that failed contract negotiations may lead to inefficient show cancellation. We propose a theoretical model of bargaining with asymmetric information that allows for this possibility. We derive various testable implications, all of which are borne by the data. We show that an increase in asset specificity of the actor-show match implies an *increase* in the probability of an agreement under efficient bargaining but a *decrease* under asymmetric information. We use this result as an identification strategy to place a 2% lower bound on the probability that a TV show is cancelled even though it would be efficient for it to continue.

1 | Introduction

The majority of multi-episode television shows—from sitcoms to courtroom dramas—last for multiple seasons. Although many contracts include long-term provisions, it is common for contract negotiations to take place at the end of a season. In this context, extension decisions (will the show continue next season?) are frequently fodder for media hype, at times reaching higher levels of drama than the show itself. These extension decisions are also interesting from an economics point of view, as they combine key concepts from game theory and contract theory: Outside options and bargaining power, asymmetric information and efficiency, asset specificity and hold-up, to name a few.

In this paper, we focus on the extension decision of television shows.[1] By means of a motivating example, consider the show *On My Block*, the Netflix teen comedy-drama series centered on the lives of a group of teenagers in a tough LA neighborhood. The show's top stars each earned $200,000 per season during the first two seasons. *On My Block* was Netflix's most-binged show in 2018, the year it was first released. This unexpected success prompted the cast to negotiate a better deal for season 3, initially pushing for $250,000 per episode per actor. After Netflix reportedly countered with a $45,000 offer, the two sides eventually agreed to $650,000 for the season, which came to $81,250 per episode (Netflix decided to produce eight shows only during season 3). The pre-season 3 agreement also included the provision for an increase to $850,000 for a potential fourth season and $1.05 million for a potential fifth [1].

Netflix's Cindy Holland acknowledged that the actors "did pretty well in its first season" and that "people have fallen in love with the characters and that cast" (ibid). *On My Block*'s popular appeal remained high throughout its life. For example, IMDb episode ratings averaged 7.25, 7.44, 7.43, and 7.35 over the show's four seasons. Why was it then not renewed after season 4? One explanation is that Netflix had different plans all along. A tantalizing alternative explanation, however, is that "it [had] little to do with popularity" but rather "most likely [had] to do with money," that is, with the prospect of paying an excessive wage bill [2].

Days before the last episode of season 4 was released, Netflix ordered the spin-off show *Freeridge* (the fictional neighborhood where *On My Block* was set). *Freeridge*, which features a different cast from *On My Block*, was released in February 2023 (eight episodes), but in April, Netflix announced it was cancelling the show after the first season.

Anecdotal evidence suggests that the pattern of events found in *On My Block* is also present in other shows. Following early success, the show's key actors become more popular and, possibly, more critical to the show's success (i.e., their talent is a match-specific asset). Emboldened by this sudden increase in power, the cast asks for more. The result is that actors do get more, but also that there is a chance the show is not renewed, possibly inefficiently so (to the extent that the actors' outside option is worth less than the future wage they bargained for).

Our goal is to analyze the dynamics of show renewal from an economic perspective. From a theoretical point of view, we have many modeling paradigms to choose from. However, not all are equally appropriate to address the patterns described above. For example, recent structural empirical IO work (e.g., Crawford and Yurukoglu [3]) models negotiated prices as arising from a Nash-bargaining surplus-splitting rule, that is, in a complete information framework. In this context, an agreement is always reached when such an agreement is efficient, thus precluding the possibility of inefficient show terminations. Coalitional equilibrium concepts such as the Core or Shapley value suffer from the same limitation.

Myerson and Satterthwaite [4] show that it is generally impossible to find efficient bargaining mechanisms in the presence of two-sided asymmetric information.[2] This suggests that asymmetric information provides a natural path to model the possibility of inefficient bargaining breakdown. However, the literature on non-cooperative bargaining with two-sided asymmetric information suggests that the outcome may depend delicately on details of the game's extensive form. Moreover, there may exist multiple equilibria.

Anecdotal evidence from TV show negotiations suggests that there is no clear negotiation protocol and that both sides have private information. To address this problem, we propose a simple extensive form that may be viewed as the reduced form of a potentially complicated (and possibly protocol-free) negotiation process. Specifically, building on Bolton and Whinston [5] and Rey and Tirole [6], we assume that Nature selects one of the parties to make a take-it-or-leave-it (TIOLI) offer to the other party. With respect to the previous literature, we add the feature that the probability that each party gets to make such a TIOLI offer is endogenous, and we allow for asymmetric information. Our model has two desired features. First, each party's payoff from a negotiated agreement is increasing in their bargaining power. Second, and more importantly, the model allows for the possibility of inefficient bargaining failure (in the present context, a show that is inefficiently cancelled).

The model suggests a series of testable empirical implications. We show that the probability of a show's extension is: (a) increasing in the show's value (which we measure by user ratings); (b) decreasing in the quality of the cast (which we measure by the actor's historical success); and (c) decreasing in the degree of talent match specificity (which we measure by how regular the top cast is over the show's history). The latter prediction requires that there is substantial private information on the producer's side, which we argue is the case.[3]

Our reduced-form analysis estimates are consistent with these predictions. Our estimates imply that: (a) a one-standard deviation increase in IMDb rating is associated with a 3.6 percentage points increase in the probability of extension (from a baseline of 69%, so an increase form 69% to 72.6%); (b) a one-standard deviation increase in cast talent is associated with a 2.6 percentage points decrease in the probability of extension; and (c) a one-standard deviation increase in the degree of match specificity is associated with a 1.7 percentage points decrease in the probability of extension.

We then use the theoretical model as a lens with which to interpret the empirical estimates. We show theoretically that, under efficient bargaining, the probability of an agreement is *increasing* in the degree of match specificity. Intuitively, the more match-specific the value of actors, the lower the combined outside option, while the value of an agreement remains constant. However, in equilibrium with asymmetric information, the probability of an agreement is *decreasing* in the degree of match specificity.

The contrast between the efficient and the equilibrium solutions, together with our empirical estimates, provides strong evidence of bargaining inefficiency. It also suggests a strategy for placing a lower bound on the extent of bargaining inefficiency, that is, a lower bound on the probability that a show is canceled even though it would be efficient for it to be extended. In fact, any increase in the probability of a breakdown resulting from an increase in talent specificity can be ascribed to bargaining inefficiency.

Following this process, a conservative quantification of the probability of inefficient breakdown is given by the effect of a one-standard-deviation increase in talent specificity. Different econometric specifications lead to somewhat different values for this lower bound, from 1.7 to 3.2 percentage points. We believe this is a sufficiently high value to warrant a careful examination of the importance of inefficient bargaining.

1.1 | Related Literature

Recent empirical IO research typically models negotiated prices as arising from a Nash-bargaining surplus-splitting rule, that is, in a complete information framework. See, for example, Crawford and Yurukoglu [3], Grennan [10], Gowrısankaran, Nevo, and Town [11], Ho and Lee [12], Backus et al. [13]. In this context, there is no negotiation breakdown by assumption. Grennan and Swanson [14] show that information is valuable and suggest that asymmetric information is an important component of negotiations, but they do not explicitly model such a negotiation process.

Some papers, however, explicitly allow for asymmetric information and the possibility of bargaining breakdown. In these papers, the structure of bargaining follows closely the particular features of the setting. Silveira [15] develops a model of pre-trial negotiation in criminal cases, allowing for the possibility of bargaining failures due to asymmetric information.[4] Like Silveira [15], we consider a take-it-or-leave-it bargaining protocol (in his case, the prosecutor offers the defendant to settle for a sentence). One important difference is that we allow for the roles of maker and receiver of a TIOLI offer to vary depending on each player's bargaining power.

Backus et al. [13] study patterns of behavior in bilateral bargaining situations using a rich, new dataset describing back-and-forth sequential bargaining occurring in over 25 million listings from eBay's Best Offer platform. They derive and test implications from various theoretical models. They find that the majority of sequences play out differently, ending in disagreement or delayed agreement, which can be rationalized by incomplete information models.

Larsen [18] develops a model of bargaining based on the actual institutional details of the car auction industry. When the highest bid is lower than the seller's (secret) reserve price, a period of alternating offers ensues. Larsen [18] assumes a cap on the number of offers and counteroffers as well as a cost per offer, thus ensuring the process ends after a finite number of steps. About a third of the time there is no agreement, some of which Larsen [18] suggests is due to bargaining inefficiencies.[5] Our model differs in that, in our case, there is no set formal protocol for negotiations (or available data on offers and counteroffers). In this sense, our paper is closer to the reduced-form empirical literature on bargaining mentioned earlier, with the difference that our main focus is on the possibility of failure to reach an agreement.

We believe our paper contributes to the literature in three ways. First, we propose a tractable model of bargaining that allows for meaningful comparative statics. Second, we propose an identification strategy that allows us to place a lower bound on the probability of inefficient bargaining failure. And third, we provide evidence regarding the negotiation process in a specific industry.

2 | Theory

Consider a game played between a producer and an actor.[6] The producer's (net) value of an agreement for the actor to star in the show at a wage w is given by

$$v_1 = k + b - w \quad (1)$$

where $k \in \mathbb{R}^+$ stands for actor talent (human capital), $b \in \mathbb{R}^+$ reflects other sources of value (e.g., script), and w is the wage paid to the actor. If no agreement is reached, then the producer's outside option is worth

$$v_0 = (1 - s)\, b + \epsilon \quad (2)$$

where $s \in [0, 1]$ measures the producer-actor match degree of asset specificity and $\epsilon \in \mathbb{R}$ measures a residual value which is the producer's private information. Implicitly, the above formulation assumes that, if negotiations break down, then the actor leaves the show and is replaced by a new actor, with the actor change costing the producer a drop in value to $(1 - s)\, b$. An example of this might be Charlie Sheen leaving *Two and a Half Men* and being replaced by Ashton Kushner. Our data suggests that this is the exception, not the rule: If a given show-actor pair is not continued, then the show is not continued either, in which case the producer's outside option would be some value ϵ. That said, we would still expect the resulting value for the producer/network to be declining in s. For example, viewers may switch to a different network because their favorite actor has left the show. And the critical feature in the model is that the producer's outside option be declining in s.

We assume that ϵ is distributed according to $F(x)$, that F is continuously differentiable, and that $H(x) \equiv F(x)/f(x)$ is strictly increasing.[7]

The above formulation implicitly assumes that there is a competitive market for actor talent k at a cost $w = k$; moreover, that actors have no private information regarding their outside option.[8] Below, we discuss this and other key assumptions. For now, we note that if $s = 0$, then the producer can hire any actor with talent k at a wage $w = k$ and get a net value b. At the opposite extreme, if $s = 1$, then the value b is destroyed when the current actor departs, leaving the producer with an outside option worth ϵ.

Regarding the actor, we assume that his value from an agreement is given by

$$u_1 = w \quad (3)$$

whereas his outside option is given by

$$u_0 = k \quad (4)$$

In principle, any bargaining procedure under incomplete information can be recast as a direct revelation game (Myerson [19]). In this paper, following Bolton and Whinston [5] and Rey and Tirole [6], we take the approach of modeling bargaining as a random-proposer game. Specifically, with probability $\alpha(k, s)$, the actor makes a take-it-or-leave-it (TIOLI) offer to the producer, a salary w that the actor requires to continue with the show. With probability $1 - \alpha(k, s)$, the opposite happens: The producer makes a TIOLI offer to the actor (also a salary w). We assume that $\alpha(k, s)$ is strictly increasing in k and s. The idea is that the greater the actor's talent or the uniqueness of their contribution to the show, the greater the actor's bargaining power, which is reflected in their ability to make a TIOLI offer (i.e., the probability that the actor is able to commit to a TIOLI offer).

2.1 | Key Assumptions

Our simple model is based on several assumptions. First, we assume that actors have no private information and that their outside option is to earn a salary w commensurate with their talent k. Ravid [7], Zuckerman et al. [8], and Elberse [9] provide evidence consistent with this assumption. For example, evidence shows that a higher k is associated with higher movie revenues ($b + k$ in our model) but not higher movie profits ($b + k - w$ in our model). We do not require the extreme assumption of no private information on the actor's side. However, the simple way in

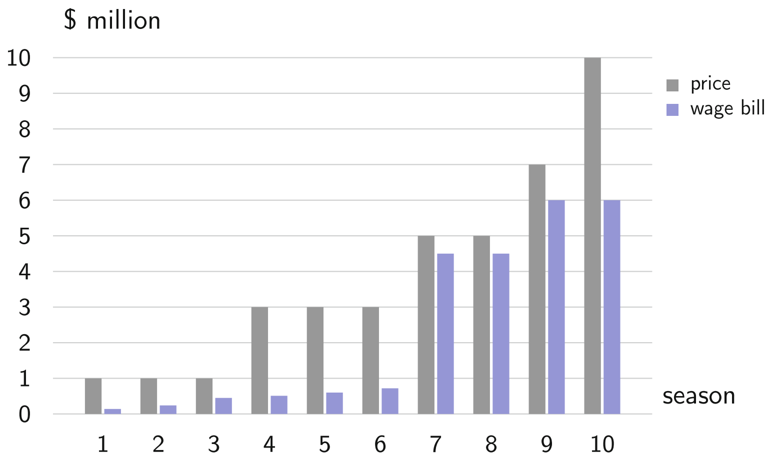

FIGURE 1 | The evolution of *Friends*: 1994–2002. Cabral [20]. Depicted value of the wage bill for seasons 8–10 does not include revenue sharing payments. [Colour figure can be viewed at wileyonlinelibrary.com]

which we model asymmetric information greatly simplifies the analysis.

The assumptions regarding $\alpha(k, s)$ are quite central to our argument. They thus require some justification. As a motivating example, consider the contrast between two leading NBC shows, *Friends* and *Law & Order*. *Friends* premiered on NBC on September 22, 1994. By the end of the second season, it was clear that *Friends* was a great success. Not surprisingly, the actors thought they could get a bigger slice of the pie: Without them, there would be no show. By season 9, each actor was paid $1 million per episode plus a share in the show's revenues. Not only was this substantially more than the $22,500 they were paid during the first season, but the ratio between the wage bill paid by the producer, WBTV, and the price WBTV received from NBC increased from 13.5% to 85.7% (see Figure 1). Despite very solid ratings, the show was not continued beyond its 10th season.

Whereas *Friends* was very centered on a cast of characters, so that actor talent was a very match-specific asset, one of the distinctive features of *Law & Order*'s business model is that it is centered on the plot rather than on the characters. Virtually nothing is known about the main characters, which implies that actor talent is not match-specific. As a result, cast changes are easy to implement—and indeed take place quite frequently: Since 1992, every one of the main characters has been replaced at least once, sometimes three or four times. For example, in 2004, Annie Parisse, one of the leading actresses, announced she wanted to pursue a movie career. Producer Dick Wolf's reaction was typical: "It was: 'Oh, thank you for coming in early. You don't mind if we kill you, do you?'" In fact, in the season's last episode, Ms. Parisse's character ends up "dead in the trunk of a car, a casualty of a drug-and-murder investigation left unresolved in anticipation of next season." Appropriately, Ms. Parisse was replaced by Alana De La Garza, whose character on *CSI: Miami* had been killed the previous season [20]. *Law & Order* premiered on September 13, 1990, ran for 20 seasons, and, after an 11-year hiatus, continued its run for more seasons.

The above anecdotal evidence seems broadly consistent with our assumption that $\alpha(k, s)$ is increasing in k and in s. To conclude this section, we should also mention that the assumption that $\alpha(k, s)$ is increasing in s is broadly consistent with the result in Fudenberg, Levine, and Tirole [21] that a seller can credibly play a take-it-or-leave-it strategy when bargaining with many buyers between whom the seller can costlessly switch. A higher value of s means that the producer has fewer alternative options to the current cast, which is equivalent to there being fewer buyers in the Fudenberg, Levine, and Tirole [21] context.

2.2 | Equilibrium Wage Offer

Consider the subgame when the actor makes a TIOLI offer. The offer is accepted by the producer if and only if $v_1 > v_0$, given by (1) and (2), or simply

$$\epsilon < s\ b + k - w \tag{5}$$

which results in a probability of acceptance given by $F(s\ b + k - w)$. The actor's expected payoff is then given by

$$\bar{u} = w\ F(s\ b + k - w) + k\ (1 - F(s\ b + k - w)) \tag{6}$$

The actor's first-order condition is given by

$$F(s\ b + k - w) - w\ f(s\ b + k - w) + k\ f(s\ b + k - w) = 0$$

or simply

$$w = k + \frac{F(s\ b + k - w)}{f(s\ b + k - w)} \tag{7}$$

We cannot, in general, obtain an explicit expression for w (it appears on both sides of the equation). However, we can state that the solution is unique.

Lemma 1. *There exists a unique solution to (7). Moreover, $w > k$.*

(The proof of this and the following results may be found in the Appendix.) Consider now the case when the producer makes a TIOLI offer. Since the actor has no private information regarding their outside option, the producer is able to extract all of the actor's surplus by offering $w = k$ if and only if the producer benefits from an agreement at that wage level, that is,

$$b + k - k > (1 - s)\ b + \epsilon$$

which happens with probability $F(s\ b)$.

2.3 | Complete Information

As mentioned earlier, Myerson and Satterthwaite [4] show that it is generally impossible to find efficient bargaining mechanisms in the presence of two-sided asymmetric information. Our purpose is to evaluate the extent to which asymmetric information leads to inefficient bargaining breakdown. For this purpose, it helps to consider the equilibrium solution under complete information, that is, when the value of ϵ is known by both players (including, in particular, the actor). If the TIOLI is made by the producer, then we get the same solution as before, since we assume the producer always observes ϵ. By contrast, if the TIOLI offer is made by the actor, then we get a different outcome than before. Since now the producer has no private information regarding their outside option, the actor is able to extract all of the producer's surplus by offering a w such that the producer is indifferent between accepting and taking the outside option. From Equations (1) and (2), this implies

$$k + b - w = (1 - s)\ b + \epsilon \tag{8}$$

The actor makes this offer if and only if the resulting w exceeds the actor's outside option, that is, $w - k > 0$. (8) may be rewritten as

$$w - k = s\ b - \epsilon$$

It follows that an agreement takes place with probability $F(s\ b)$, just as in the case when the producer makes an offer.

To summarize, under complete information, the outcome is efficient, in the sense that an agreement takes place if and only if the joint payoff of actor and producer is greater with an agreement than under the outside option. Moreover, the probability of an agreement is independent of the identity of the party making a TIOLI offer. In the particular case at hand, such a probability is given by

$$r^\star = F(s\ b) \tag{9}$$

2.4 | Comparative Statics

Let $\hat{r}$ be the probability that an agreement is reached (under asymmetric information). Let $\hat{r}_i$ $(i = a, p)$ be the probability that an agreement is reached given that i is making a TIOLI offer. From the analysis above,

$$r_a = F(s\ b + k - w)$$
$$r_p = F(s\ b)$$

where w is determined by (7). It follows that

$$\hat{r} = \alpha(k, s)\ r_a + (1 - \alpha(k, s))\ r_p$$

or simply

$$\hat{r} = F(s\ b) - \alpha(k, s)\ (F(s\ b) - F(s\ b + k - w)) \tag{10}$$

Our first result compares the complete and the incomplete information cases.

Proposition 1. $\hat{r} < r^\star$: *The probability of an agreement is greater under complete information.*

Intuitively, under complete information, an agreement is reached when it is efficient for an agreement to be reached, which happens with probability $F(s\ b)$. Under asymmetric information, and as shown by (10), the "baseline" probability of an agreement is also given by $F(s\ b)$. However, with probability α, the actor makes a TIOLI offer under asymmetric information. This implies the possibility that the offer be rejected even though it would be efficient for an agreement to be reached (at a different w). The difference $F(s\ b) - F(s\ b + k - w)$ measures the probability of such an inefficient bargaining breakdown.

Proposition 1 is consistent with the idea that asymmetric information implies, with some probability, an inefficient bargaining breakdown. In the particular setting we consider, the inefficiency we capture is isomorphic to the inefficiency of monopoly pricing with market power: With some probability, consumers fail to make a purchase even though it would have been efficient for a purchase to take place, that is, willingness to pay is greater than seller cost.[9]

Our next result concerns the comparative statics of r^*.

Proposition 2. *$r^\star$ is strictly increasing in b and s (and invariant with respect to k).*

Consider first an increase in b. Just as an expansion of the demand curve increases the probability of a sale (all else equal), so a better show is more likely to be extended. As to an increase in s, notice that the value of an agreement remains the same, but the value of the outside option strictly decreases. Therefore, the higher s is, the more likely it is that it is efficient for the show to continue.

Next, we turn to the comparative statics regarding $\hat{r}$.

Proposition 3. *$\hat{r}$ is strictly increasing in b and strictly decreasing in k. If b is sufficiently small and $\partial\alpha/\partial s$ sufficiently high, then $\hat{r}$ is decreasing in s.*

The intuition for the effect of an increase in b is similar to Proposition 2. Consider now the effect of an increase in k. A higher k implies higher actor power (by assumption). This increase in actor power is reflected in a higher probability that the TIOLI offer during negotiations is made by the actor. Finally, to the extent that the producer has private information regarding their valuation, this decreases the probability of extension. Note that an increase in k has no effect on $r^\star$. In fact, an increase in k is reflected, dollar for dollar, in both an increase in the value of an agreement and the value of the outside option. As such, the probability of an efficient agreement does not change.

With regard to s, a similar argument applies: A higher s leads to a higher α, which in turn increases the probability of inefficient bargaining breakdown. However, we must also take into consideration a second effect. Unlike k, which has no effect on the efficiency of an agreement, an increase in s increases the relative value of an agreement vis-a-vis the outside option. Intuitively, the greater s is, the more the parties have to lose from failing to come to an agreement. This is reflected in the fact that $r^\star$ is strictly increasing in s. Under asymmetric information, however, the effect of an increase in s on the probability of an agreement is only positive if it outweighs the negative bargaining-power effect. In other words, under asymmetric information, and to the extent that we have two effects of opposite sign, the effect of s is largely an empirical question.

The combination of Propositions 2 and 3 provides a strategy for both testing our theory of negotiations and estimating the extent of bargaining inefficiencies. Proposition 2 states that, everything else constant, the greater the degree of asset specificity, the more likely it is for efficiency to dictate that an agreement should be reached. However, Proposition 3 states that, in the presence of asymmetric information, negotiations may break down, and that the probability that this happens is increasing in the degree of asset specificity (due to the shift in bargaining power). For these reasons, we argue that a negative relation between r and s provides evidence of bargaining inefficiency.

We are unable to tell whether a particular negotiation should have led to an agreement or not (in efficiency terms). However, when we observe an increase in the probability of a negotiation's breakdown resulting from an increase in s, we can attribute such an increase to bargaining inefficiency. In fact, efficiency considerations dictate that the probability of an agreement should have increased, not decreased. This allows us to place a lower bound on the extent of bargaining inefficiency.

In Section 4, we test Proposition 3, whereas in Section 6 we estimate a lower bound on the extent of bargaining inefficiencies. Before that, in the next section, we describe the data we use in Sections 4 and 6.

3 | Data

We draw data from the raw IMDb files publicly available in March 2016, which allow us to assemble a dataset at the show-season level for all shows that were released in the US from 1970 until 2014. The IMDb files identify TV programming episodes that belong to different seasons of different shows, as well as actors participating in those episodes and user ratings at the episode level. We have data for 3243 shows, some of which last for one season, a few for more than 20 seasons. Each season, a show typically comprises several episodes. Shows can be of different kinds: Comedy, documentary, drama, and so on.

To test our theory, we need to measure the variables r, b, k, and s at the show-season level. We next turn to each of these.

3.1 | Show Extension (r)

In our theoretical model, r is the probability that a given show is extended into an extra season. To estimate this probability directly, we create a dummy variable, at the show-season level, that takes the value 1 if the show is continued for at least one more season.

3.2 | Show Value (b)

As a proxy for the value of a given show, we use the average IMDb user ratings of all episodes in a given season. We note that user ratings information is time-invariant in our source. In other words, we only observe the total rating as of 2016. Since ratings are at the episode level, for a given show, we obtain a series b that varies over episodes.

We believe IMDb user ratings provide a good proxy for show value. Figure 2 illustrates the relation between IMDb ratings and DVD sales during the 2000–2009 period, a period when DVD sales represented a significant fraction of revenues. While there is considerable variation, we observe a clear positive correlation between the two measures. A simple log-linear regression yields a coefficient of 0.6397 (estimated with a p value lower than 2E-16),

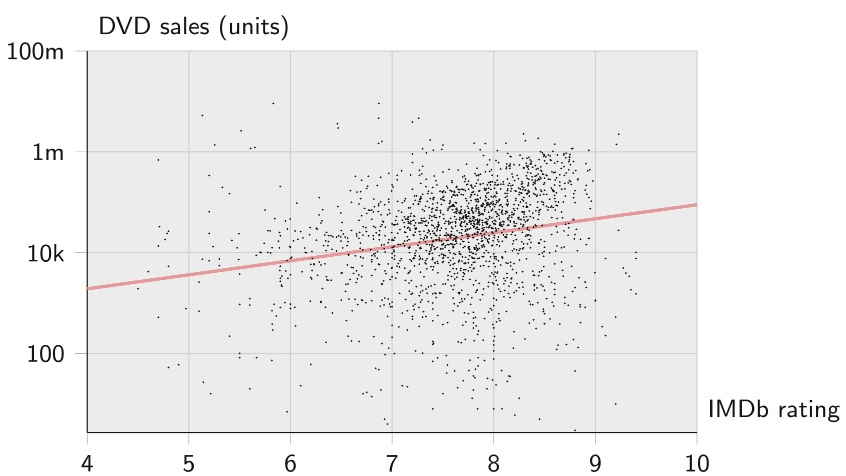

FIGURE 2 | IMDb rating and DVD sales. [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 1 | Correlation matrix of talent specificity depending on number of top talent.

	Top-1	Top-2	Top-3
Top-2	0.74		
Top-3	0.65	0.88	
Top-5	0.53	0.72	0.83

TABLE 2 | Correlation matrix of independent variables.

	b	k
k	0.12	
s	0.07	0.06

implying that a 1 point increase in IMDb rating is associated with a 64 percent increase in DVD sales.

3.3 | Talent (k)

Next, we turn to a measure of the actor's human capital, k. IMDb data includes a registry of acting credits, specifically the billing order of the talent participating in each episode. From this registry, we take the top three billed actors. We first count all distinct TV shows, feature films, TV films, or video films in which each given actor of a show's top three billed actors has participated in the past up to the current year. We then average this count (weighed by the rating of each show in each person's record to account for varying degrees of quality) over the top 3 actors and express it in natural logarithms. As a robustness measure, we do the same calculations, but considering the top 5 rather than the top 3.

3.4 | Talent Specificity (s)

We measure asset specificity, s, by the persistence of actors throughout a show's history. For each top-3 actor, we measure the percentage of all episodes (in all seasons, up to the current season) in which the actor was billed in the top 3. Then, for the current season, we take the average of this variable across all actors billed in the top 3 during the current season. As a robustness measure, we do the same calculations, but considering the top 5 rather than the top 3.

Given the importance of the talent-specificity variable, we consider a number of variations, some of which we report in the robustness section. For now, we note that measures that select the top 1, 2, 3, or 5 cast members are highly correlated, as shown in Table 1. We also note that the values of b, k, and s are positively correlated, as shown in Table 2 (all correlations are significant at the 5% level).

3.5 | Other Variables

The genre of each show is available for 98.8% of shows; we assign those without a genre to an undefined genre category. When more than two genre classifications are available on each show, we only assign these shows the two most common genre categories in the sample, thus reaching 99 blends of genres (for instance, action-drama).

3.6 | A First Look at the Data

Figure 3 shows the distribution of our 3243 shows in terms of the number of seasons. While the average count of seasons per show is 2.95, there is considerable variation across shows. Almost 60% of all shows last for one season only, and the distribution is highly skewed.

As an illustration of the variety of shows in our sample, Table 3 provides a few examples. One variable that plays a central role in our theory of show extension is actor talent specificity. Among the particular examples considered in Table 3, it takes the lowest value with *Are You Afraid of the Dark?*, a show where "a group of teenagers meet in the woods, and tell scary stories." Ross Hull, the actor with the most episodes, appeared 68 times from 1990 to 2000. However, most actors appeared in a much lower number of episodes during the show's seven seasons. This explains the low value of the talent specificity variable: Actors on *Are You Afraid of the Dark?* come and go, and the show is not particularly dependent on a particular actor. By contrast, shows like *Friends* or *Veep* score 1 on actor specificity. In fact, the leading three actors appeared in every episode.

A quick scan through the data on Table 3 fails to show a clear relation between talent specificity and the number of seasons. However, we must recall that the number of seasons results from a variety of factors, including the show's popularity. The purpose of our analysis is precisely to examine the separate contribution of each variable to the likelihood that a show is extended into the next season, and ultimately, for how many seasons it lasts.

Table 4 reports summary statistics of the main variables, at the show-season level, that we will use in our regressions. Variable s plays a central role in our analysis, and as such, warrants special attention. Figure 4 plots the distribution of s in bins of 0.05 width. As can be seen from the top panel, the distribution is very skewed, with considerable mass in the top bin. In fact, we note that 35% of our observations have s exactly equal to one: The top 3 cast members are present in each of the show's episodes. If we exclude these observations with $s = 1$, then we obtain a rather more even distribution, as can be seen in the bottom panel of Figure 4. In the regressions we present in Section 4, as well as in the numerical exercise we present in Section 6, we will take this asymmetry into account.

4 | Empirical Results

Proposition 3 suggests a series of testable predictions. Specifically, we expect r to be increasing in b and decreasing in k and s. Table 5 presents the results from a set of regressions, beginning with our base regression. In all regressions, the dependent variable is a dummy that takes the value 1 when a show is extended into the next season. The regressions differ in terms of functional form. All regressions include fixed effects for the number of

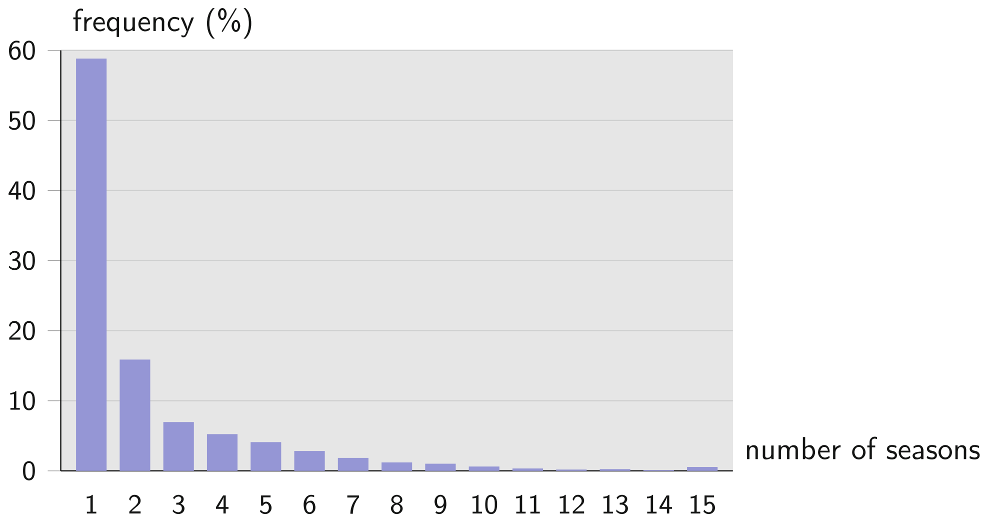

FIGURE 3 | Histogram of total number of seasons by show. The rightmost bin includes all shows with 15 or more seasons. [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 3 | Some examples of show-level observations in the data.

Show title	Num. seasons	Start year	End year	Talent	Talent specificity	IMDb rating
30 Rock	7	2006	2012	7.4	0.91	8.12
Are You Afraid of the Dark?	7	1990	2000	1.21	0.03	8.18
Arrested Development	5	2003	2013	7.93	0.99	8.52
Breaking Bad	5	2008	2012	8.86	0.98	8.98
Friends	10	1994	2003	6.85	1.00	8.52
Gold Rush: Alaska	5	2010	2014	0.64	0.14	6.81
Law & Order	20	1990	2009	6.06	0.56	7.73
Lost	6	2004	2010	5.28	0.53	8.68
Mad Men	7	2007	2014	8.8	0.97	8.53
Paranormal Witness	3	2011	2013	1.12	0.08	7.75
Seinfeld	9	1989	1997	3.36	0.93	8.48
The Sopranos	6	1999	2006	7.8	0.96	8.68
The Wire	5	2002	2008	3.62	0.68	8.84
The X Files	9	1993	2001	5.17	0.19	8.10
Veep	3	2012	2014	9.94	1.00	8.19

TABLE 4 | Summary statistics at the show-season level.

Var	Desc	Count	Mean	Std	Min	25%	50%	75%	Max
r	Show renewal dummy	7474	0.69	0.46	0.00	0.00	1.00	1.00	1.00
b	IMDb rating	7474	7.53	1.00	1.10	7.10	7.68	8.16	9.90
s	Talent specificity	7474	0.72	0.32	0.01	0.50	0.86	1.00	1.00
k	Talent (in logs)	7474	1.55	0.67	0.00	1.02	1.54	2.04	4.00

episodes, the season number, and the combination of genre and year.

The effect of "IMDb rating," our measure of b, is estimated with statistical precision at 0.036. The sign is consistent with the theory. We estimate that a one-standard-deviation increase in IMDb rating (which, per Table 4, is equal to one) is associated with a 3.6 percentage point increase in the probability of extension, from a baseline of 69%. In other words, an increase from 69% to 72.6%.

Next, we consider the comparative statics with respect to actor talent. The effect of "Talent," our measure of k, is estimated with

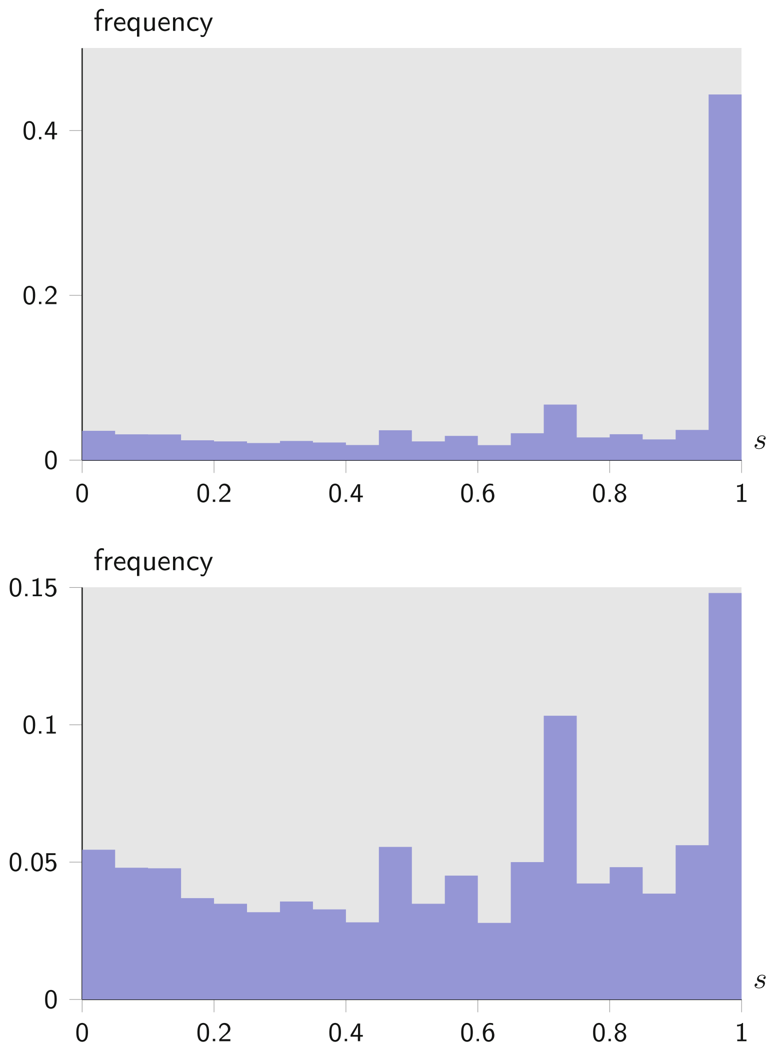

FIGURE 4 | Distribution of s. Top panel: All observations (7474 observations). Bottom panel: Excluding $s = 1$ (4879 observations). Note that the vertical scale is different in the bottom panel. [Colour figure can be viewed at wileyonlinelibrary.com]

statistical precision at -0.034. The sign is consistent with the theory. We estimate that a one-standard-deviation increase in talent quality is associated with a decrease in the probability of extension by about 2.5 percentage points (0.034×0.74).

Finally, the effect of "Talent specificity," our measure of s, is estimated with statistical precision at -0.053. We estimate that a one-standard-deviation increase in Top-3 talent specificity is associated with a decrease in the probability of extension by about 1.7 percentage points (specifically, 0.052×0.32).

Our second regression introduces a quadratic term in the variable s. The linear and quadratic terms are statistically significant and have values -0.554 and 0.415, respectively. This implies that the relation between s and r (the probability of continuation) is convex. In fact, since $0.554/(2 \times 0.415) = 0.67$ lies between zero and 1 (the domain of s), we estimate a total effect that is first negative but positive for high values of s.

At this point, one suspects that the non-monotonic relation between s and r may be related to the high density of observations with $s = 1$ (as indicated by Figure 4). To check this possibility, we run a third regression where we add a dummy that takes the value 1 if and only if $s = 1$. We can either add this dummy as a separate regressor or interact it with s: The resulting equation is the same. The coefficients are estimated with precision. They imply that the contribution of s to r increases linearly in s from zero for $s = 0$ to -0.098 as s tends to 1. At $s = 1$, the contribution of s to r is estimated at $-0.099 + 0.047 = -0.052$. Therefore, this second regression confirms that the effect of s on the likelihood of an agreement is negative for all values of s. We also tried a fourth regression where we combine a quadratic form *and* a dummy for $s = 1$. However, in this case, the dummy coefficient is not statistically significant, which suggests that a dummy for $s = 1$ or a quadratic functional form are alternative ways of addressing the nonlinearity of the relation between s and r.

4.1 | Model Fit

One way to evaluate the model's fit is to compare the observed duration of each show (number of seasons) with the model prediction based on estimated probabilities of show renewal into the

TABLE 5 | Season renewal, r, as a function of b, k, and s.

Dependent variable	Renewal after current season (1/0)		
IMDb rating	0.036***	0.035***	0.036***
	(0.01)	(0.01)	(0.01)
Talent	−0.034***	−0.034***	−0.034***
	(0.01)	(0.01)	(0.01)
Talent specificity	−0.053**	−0.554***	−0.099***
	(0.02)	(0.18)	(0.03)
Talent specificity squared		0.415***	
		(0.14)	
Talent specificity = 1 dummy			0.047*
			(0.02)
Count of episodes f.e.	Yes	Yes	Yes
Season number f.e.	Yes	Yes	Yes
Genre × year f.e.	Yes	Yes	Yes
Adjusted R^2	0.16	0.17	0.16
N. observations	7474	7474	7474
N. clusters	99	99	99

next season. Let T_i be show i's number of seasons. We then compute the model's estimated value of T_i, which we denote by $\widehat{T}_i$, for a given show (conditional on the show having been produced at least one season) as

$$\widehat{T}_i = \sum_{t=1}^{\infty} t\,(1 - r_{it})\,\Pi_{\tau=1}^{t} r_{i\tau}$$

where r_{it} is the probability that show i is renewed after season t. Note that, in some cases, we do not observe some of the show's variables for a given season. When that is the case, we estimate the value of r_{it} by using the nearest estimate of those variables. Having all of the values of r_{it}, we compute $\widehat{T}_i$ for each show i. Finally, for each *observed* show duration, we compute $\mathbb{E}(\widehat{T}|T)$, the average value of all $\widehat{T}$ of shows that lasted for T seasons.

Table 6 shows the result of this exercise. The model does a good job of predicting when shows are expected to last for longer seasons. That said, the relation between T and $\mathbb{E}(\widehat{T}|T)$ is flatter than the diagonal.

4.2 | Additional Predictions

Although the main focus of our empirical analysis is on the comparative statics with respect to r, our model also has testable implications regarding actor performance. Conditional on a show not being extended, an actor's subsequent performance should be better than average. By "better than average" we mean better than predicted by other controls, in particular actor and year fixed effects. Intuitively, based on our model of failed negotiations, an actor with a better outside option is more likely to reach an agreement with the producer of an outside TV show, feature film, TV film, or video film. In other words, the sample of actors who have just had their last season with a given show is biased towards actors who, given their favorable outside option, are more likely to reject an offer from producers.

TABLE 6 | Average expected number of seasons according to estimated model, $\mathbb{E}(\widehat{T}|T)$, conditional on actual number of seasons, T.

T	$\mathbb{E}(\widehat{T}\|T)$	N
1	2.38	1908
2	3.46	515
3	3.77	226
4	3.95	170
5	4.51	133
6	4.63	92
7	4.96	60
8	4.60	39
9	4.77	33
10	6.64	20

TABLE 7 | Event of no renewal for top-3 actors and subsequent personal career outcomes.

Dependent variable	Number of titles		Average billing order	
	Year $t+1$	Year $t+2$	Year $t+1$	Year $t+2$
No renewal$_{i,t}$	0.175***	0.244***	0.067	−0.343***
	(0.01)	(0.01)	(0.10)	(0.10)
Person f.e.	Yes	Yes	Yes	Yes
Year f.e.	Yes	Yes	Yes	Yes
Adjusted R^2	0.30	0.30	0.14	0.14
N. obs.	532,968	509,536	274,975	266,493
N. clusters	22,921	22,410	21,196	20,784

Table 7 displays four regressions that test this prediction. We consider two measures of performance. First, the number of titles (where titles are distinct TV shows, feature films, TV films, and video films) in which the actor is engaged. Second, his or her (average) billing order in those titles. Note that a *lower* billing order number is a measure of better performance (1 being best), so our theory's prediction corresponds to a positive coefficient on the first performance measure and a negative coefficient on the second measure. We also consider two periods, the year after a show is terminated and the year after that. This results in a total of four regressions.

As can be seen, the results in Table 7 are consistent with theory. The coefficients both on the number of shows and billing order are statistically significant for $t+2$. In that year, an actor present in a recently-terminated show is expected to be present at 0.244 more distinct titles than predicted by actor and year fixed effects, and to be placed 0.343 points higher in the billing order.

5 | Extensions and Robustness

We consider a series of robustness checks that essentially extend the set of regressions shown in Table 5. In particular, since the main variable of focus in our analysis is s, our robustness checks are primarily centered on s.

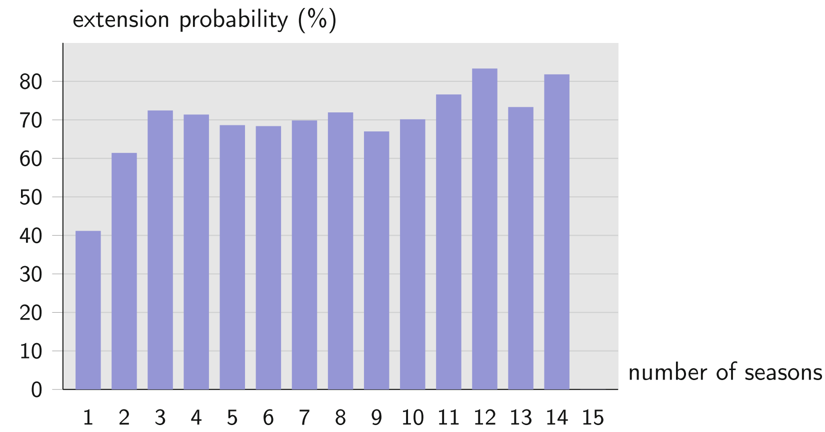

FIGURE 5 | Probability of a show extension. Lower probabilities (higher hazard rates) during the first two periods suggest that some learning takes place during the first two periods. [Colour figure can be viewed at wileyonlinelibrary.com]

First, we consider two alternative definitions of s. Recall that our basic definition is based on the fraction of episodes where a given actor was billed in the top 3. An alternative definition is based on the number of *seasons* in which an actor was billed top 3 in *some* episode during that season. A second alternative definition of s is based on the number of top 3 episode appearances during the *current* season. In both cases, we find our estimates are statistically significant, have the same sign as the estimates in Table 5, and imply effects of the same order of magnitude.

A second set of robustness tests consists of using the top 5 billed actors in lieu of the top 3. We try this alternative approach both for our measure of talent and for our measure of talent specificity. As in the previous robustness tests, we find significant estimates with the same sign and size.

Our basic set of regressions is based on a linear probability model. We also considered Poisson pseudo-likelihood estimation. Again, the results are fairly similar.

5.1 | Learning

Figure 5 shows the sample probability of an extension in a given season. The first-period probability is considerably lower than in subsequent periods. This suggests that learning takes place during the first season (or the first two seasons). Our regressions include season fixed effects. However, to allow for the possibility of the effects of k and s having a different nature during the first season, we considered alternative regressions where a $t = 1$ dummy interacts with the value of s. We conclude that the coefficient is not statistically significant.

5.2 | Alternative Narrative

Our theoretical model assumes that the actor's utility is limited to their compensation. One possible alternative narrative for our results is that an actor's utility from featuring in a given show declines over time. If this is the case, then what we deem to be an inefficient show cancellation may actually be an efficient outcome: The value of an agreement is lower than the alternative because the actor's utility from an agreement is lower.

To consider this possibility, we extend the regressions in Table 7 to include the value of s in the show the actor has left. We find the coefficient not to be statistically significant.

6 | Bargaining Inefficiency

Proposition 2 states that, *under efficient bargaining*, an increase in s leads to a higher probability that an agreement is reached. In fact, the joint value of not extending a show declines, whereas the value of extending the show remains constant. By contrast, Proposition 3 provides sufficient conditions such that, *in equilibrium* with asymmetric information, an inefficient breakdown in negotiations takes place with positive probability and with a probability that is increasing in s. Similarly, our results indicate that, under efficient bargaining, the probability of an agreement is constant with respect to k but decreasing in k in equilibrium. These relations can be derived from the expressions of efficient and equilibrium probabilities of continuation, which are given by

$$r^\star = F(s\ b)$$
$$\hat{r} = F(s\ b) - \alpha(k,s)\ (F(s\ b) - F(s\ b + k - w))$$

Our empirical analysis shows that the probability of a show's continuation is decreasing in s, a relation that is economically sizable, statistically significant, and robust to a variety of changes in regression specification. Altogether, we suggest that our theoretical and empirical analyses add up to strong evidence that the efficient bargaining model fails to capture an important feature of reality.

As mentioned earlier, we cannot determine whether it would be efficient for a specific show to be extended. As such, we are unable to distinguish an efficient show termination from an inefficient one. We can, however, provide a lower bound on the extent of inefficient bargaining. Suppose that, when $s = 0$, there is no inefficient bargaining breakdown. This does not follow from the theoretical model, that is, even when $s = 0$, inefficient breakdown takes place with positive probability. But suppose that $\alpha(k, s) = 0$, which stacks the cards against inefficient bargaining. This allows us to assign to inefficient bargaining any *increase* in the probability of breakdown resulting from *higher* values of s. In fact,

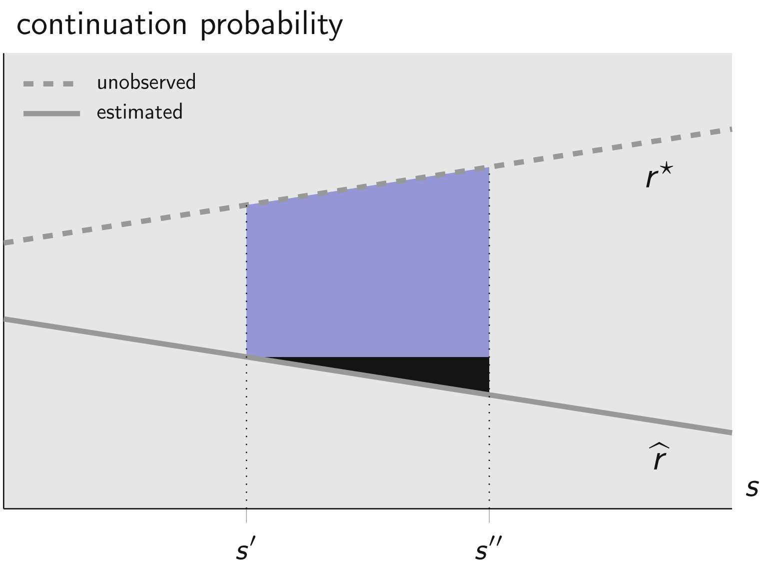

FIGURE 6 | Inefficient bargaining breakdown. Theory shows that $r^\star$ is increasing, $r^\star > \widehat{r}$, and (under certain conditions) $\widehat{r}$ is decreasing. Therefore, the height of the black (darkest) area, weighted by the density of s, provides a lower bound on the extent of inefficient bargaining breakdown. [Colour figure can be viewed at wileyonlinelibrary.com]

per Proposition 2, under efficient bargaining, higher values of s should be associated with a higher probability of a negotiated agreement. A similar argument applies to increases in k.

The idea is illustrated by Figure 6, which considers variations in s (measured on the horizontal axis). The vertical axis measures the probability that a show is continued. Our theoretical results imply that $r^\star$, the efficient continuation probability, is strictly increasing in s. We plot $r^\star$ as a dashed line to highlight the fact that we do *not* observe its value. Our theoretical results also imply that $\widehat{r}$, the equilibrium continuation probability, is strictly decreasing in s and strictly lower than $r^\star$.

Since we do not observe $r^\star$ (dashed line), we cannot measure the efficiency-equilibrium gap (the gap between the dashed line and the solid line). However, since $r^\star$ is strictly increasing and $\widehat{r}$ strictly decreasing, as s moves from s' to s'', the area in black provides a *lower bound* on the total area between $r^\star$ and $\widehat{r}$.

Specifically, we estimate a lower bound on the probability of inefficient bargaining breakdown as follows. Consider first variation in s. For each show-season observation, we compute the increase in the probability of a negotiation breakdown in excess of what it would be if $s = 0$. We then average this over all show-season observations. We repeat this process by considering the three regressions in Table 5: Linear, quadratic, and linear with a dummy at $s = 1$. A similar process can be followed for variation in k, with the simplifying feature that the coefficient on k is invariant across regression models.

The first numeric column in Table 8 displays the results from this exercise. Considering variation in k, we estimate a lower bound of 5.68% on the probability of inefficient bargaining breakdown. When it comes to variation in s, our results depend on the model we consider. Our estimate based on a quadratic model is significantly higher than those based on linear models. We believe the reason for this is that our exercise amounts to a significant degree of out-of-sample prediction. This is particularly dangerous when we use a quadratic equation. In other words, our quadratic model implies that, as s tends to zero, the probability of an agreement increases at an increasing rate. However, we have very few observations for low values of s, and so such an estimate has to be taken with a grain a salt.

TABLE 8 | Lower-bound estimate of the probability of inefficient bargaining (in percentage points) based on variation in k or s.

	Variation interval	
Source of variation	$[0, 1]$	$[\mu - \sigma, \mu]$
k	5.68	3.55
s (linear model)	3.82	1.70
s (quadratic model)	13.99	7.05
s (linear with a dummy at $s = 1$)	5.34	3.17

Excluding the estimate based on the quadratic model, we obtain a lower bound of the order of 4 percentage points (considering variation in s) and 5 percentage points (considering variation in k). We believe these values are sufficiently high to make us take seriously the possibility of inefficient bargaining. Note that the values for k and s are alternative paths for a lower bound; in other words, they are not additive.

As mentioned in the previous paragraphs, the lower-bound estimates in Table 8 may be criticized for being unrealistically based on "out of sample" extrapolation. For example, we compare the actual values of s with a counterfactual where $s = 0$. However, as Table 4 suggests, $s = 0$ is several standard deviations away from the mean. A more conservative lower bound might be obtained by simply considering a one-standard-deviation shift in the value of s around its mean and assuming that, at the lower end of that shift (i.e., at a point between the mean of s and zero), the probability of inefficient bargaining breakdown is zero.

The rightmost column of Table 8 displays the results from this exercise. Not surprisingly, the values are lower, and the estimates of the four different estimation paths are closer in magnitude, consistent with our out-of-sample extrapolation conjecture. That said, we still get a lower bound of about 2 or 3 percent.

The existence and importance of bargaining inefficiency likely depend on the particular context under consideration. Larsen [18], arguably the closest paper to ours, estimates that "over one-half of failed negotiations are cases where gains from trade exist." This is one order of magnitude greater than our result, but we derive a lower bound, not a point estimate.

7 | Conclusion

Much of the recent structural empirical IO literature assumes efficient bargaining. Evidence from several industries, including TV shows, suggests that the negotiation outcome is not always efficient. We propose a theoretical model and an identification strategy that allow us to place a 2%–5% lower bound on the probability that a TV show is cancelled even though it would be efficient for it to continue (the actual lower bound depending on how conservative one wants to be in estimating it).

One possible next step along this line of research would be to devise and estimate a structural model of show renewal. One advantage of such a model is that it allows us to obtain an estimate of the probability of bargaining breakdown as well as to perform a welfare analysis.

Acknowledgements

We are grateful to two anonymous referees and seminar audiences at NYU, Universidad de Piura, Duke, NC State, IIOC 2024, and Carlos III (Madrid) for helpful comments. Gabriel Natividad acknowledges the hospitality of the Economics department at NYU Stern during this project's development.

Endnotes

[1] For simplicity, we use "TV shows" as a generic concept that includes both linear television and streaming.

[2] More specifically, when the supports of buyer and seller types overlap, there does not exist any incentive-compatible, individually rational bargaining mechanism that is ex-post efficient and that also satisfies an ex-ante balanced budget.

[3] For example, the evidence presented in Ravid [7], Zuckerman et al. [8], and Elberse [9] is consistent with the assumption that actors have no private information and that their outside option is to earn a salary w commensurate with their talent k. By contrast, the business of television is complex as it requires combining multiple pieces of content. In this context, it may be difficult to estimate the buyer's opportunity cost and/or outside option.

[4] See also Sieg [16], which in turn is based on Nalebuff [17]. Like ours, these papers assume an extensive form with a TIOLI offer.

[5] The data on Table 3, Panel A of Larsen [18], show that, conditional on trade not occurring through the auction price exceeding the reserve price (i.e., conditional on beginning the bargaining stage), trade occurs with probability 0.646. Thus, trade fails with probability $1 - 0.646$.

[6] While for simplicity we cast our analysis in terms of a producer-actor game, the qualitative results in this section should be of more general interest.

[7] The assumption that $H \equiv F/f$ is strictly increasing is common in various settings (e.g., auction theory) and is satisfied by various distributions. In a pricing context where the demand curve is given by $1 - F$, the assumption that H is strictly increasing is equivalent to the assumption that the marginal revenue curve is strictly decreasing.

[8] A general formulation of our model has private information on both sides. Our central results depend on private information on the producer's side being sufficiently more significant than private information on the actor's side. By assuming no private information on the actor's side, we obtain that feature by construction. The more general framework requires that we qualify the results with somewhat awkward parameter assumptions, while the gain in terms of intuition is limited.

[9] We should note that the classic Myerson and Satterthwaite [4] inefficiency theorem is about the case where there is uncertainty about whether gains from trade exist (two-sided incomplete information with overlapping support of the buyer and seller type distributions). That is not the case in our model. In our setting, a first-best mechanism always exists, namely to let the informed party—the producer—make the TIOLIO.

References

1. L. Goldberg, "Netflix's 'On My Block' Stars Score Sizable Pay Raises," 2019 Hollywoodreporter.com.

2. M. Douglas, "Why Netflix is Canceling 'On My Block' After Season Four," 2021 Yahoo.com.

3. G. S. Crawford and A. Yurukoglu, "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review* 102 (2012): 643–685.

4. R. Myerson and M. A. Satterthwaite, "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory* 29 (1983): 265–281.

5. P. Bolton and M. D. Whinston, "The 'Foreclosure' Effects of Vertical Mergers," *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für Die Gesamte Staatswissenschaft* 147 (1991): 207–226.

6. P. Rey and J. Tirole, "A Primer on Foreclosure," in *Handbook of Industrial Organization*, vol. 3, ed. M. Armstrong and R. Porter (Elsevier, 2007), 2145–2220.

7. S. A. Ravid, "Information, Blockbusters, and Stars: A Study of the Film Industry," *Journal of Business* 72 (1999): 463–492.

8. E. Zuckerman, T.-Y. Kim, K. Ukanwa, and J. Von Rittmann, "Robust Identities or Nonentities? Typecasting in the Feature-Film Labor Market," *American Journal of Sociology* 108 (2003): 1018–1074.

9. A. Elberse, "The Power of Stars: Do Star Actors Drive the Success of Movies?," *Journal of Marketing* 71 (2007): 102–120.

10. M. Grennan, "Price Discrimination and Bargaining: Empirical Evidence From Medical Devices," *American Economic Review* 103 (2013): 145–177.

11. G. Gowrisankaran, A. Nevo, and R. Town, "Mergers When Prices Are Negotiated: Evidence From the Hospital Industry," *American Economic Review* 105 (2015): 172–203.

12. K. Ho and R. S. Lee, "Insurer Competition in Health Care Markets," *Econometrica* 85 (2017): 379–417.

13. M. Backus, T. Blakee, B. Larsen, and S. Tadelis, "Sequential Bargaining in the Field: Evidence From Millions of Online Bargaining Interactions," *Quarterly Journal of Economics* 135 (2020): 1319–1361.

14. M. Grennan and A. Swanson, "Transparency and Negotiated Prices: The Value of Information in Hospital-Supplier Bargaining," *Journal of Political Economy* 128 (2020): 1234–1268.

15. B. S. Silveira, "Bargaining With Asymmetric Information: An Empirical Study of Plea Negotiations," *Econometrica* 85 (2017): 419–452.

16. H. Sieg, "Estimating a Bargaining Model With Asymmetric Information: Evidence From Medical Malpractice Disputes," *Journal of Political Economy* 108 (2000): 1006–1021.

17. B. Nalebuff, "Credible Pretrial Negotiation," *RAND Journal of Economics* 18 (1987): 198–210.

18. B. J. Larsen, "The Efficiency of Real-World Bargaining: Evidence From Wholesale Used-Auto Auctions," *Review of Economic Studies* 88 (2021): 851–882.

19. R. Myerson, "Incentive Compatibility and the Bargaining Problem," *Econometrica* 47 (1979): 61–73.

20. L. Cabral, *TV Power Games: Friends and Law & Order* (NYU Stern Case Study, 2010).

21. D. Fudenberg, D. Levine, and J. Tirole, "Incomplete Information Bargaining With Outside Opportunities," *Quarterly Journal of Economics* 102 (1987): 37–50.

Appendix A

Proof of Lemma 1. We can re-write (7) as

$$x = \frac{F(s\ b - x)}{f(s\ b - x)} \tag{A1}$$

where $x \equiv w - k$. Since $F(x)/f(x)$ is strictly increasing, the left-hand side of (A1) is strictly increasing in x. The right-hand side of (A1) is strictly decreasing in x. At $x = 0$, the left-hand side is equal to zero, whereas the right-hand side is positive, since $f(x) > 0$ for $x \in \mathbb{R}$. Finally, by the intermediate value theorem, there exists a unique $x > 0$ satisfying (A1). □

Proof of Proposition 1. Since $w > k$ (which follows from Lemma 1), and since F is strictly increasing, $F(s\ b) - F(s\ b + k - w) > 0$. Since $\alpha > 0$, the result then follows from Equations (9) and (10). □

Proof of Proposition 2. Since $r^* = F(s\ b)$ and F is strictly increasing, it follows that r^* is increasing in s and in b. □

The following lemma is used in the proof of Proposition 3.

Lemma 2. $dw/dk = 1$

Proof of Lemma 2. The assumption that the actor's outside option is given by k implies that, when the producer makes a TIOLI offer, we have $w = k$, in which case the result is trivial. The result is less immediate in the case when the actor makes an offer. When that is the case, from Equation (7) we have that

$$z = \frac{F(s\ b - z)}{f(s\ b - z)}$$

where $z \equiv w - k$. This implies that, for given s and b, z is uniquely determined, which in turn implies that w changes with k dollar for dollar. □

Proof of Proposition 3. The equilibrium probability of continuation is given by (10), that is,

$$\hat{r} = F(s\ b) - \alpha(k, s)\ (F(s\ b) - F(s\ b + k - w)) \tag{A2}$$

Consider first a variation in b. Taking the derivative of (A2) with respect to b, we get

$$\begin{aligned} \frac{d\hat{r}}{db} &= s\ f(s\ b) - \alpha(k, s)\ \frac{d}{db}\ (F(s\ b) - F(s\ b + k - w)) \\ &= s\ f(s\ b) > 0 \end{aligned} \tag{A3}$$

where the second equality follows from Lemma 2. Specifically, since $dw/dk = 1$,

$$\frac{d}{db}\ F(s\ b + k - w) = \frac{d}{db}\ F(s\ b)$$

which in turn implies that the second term on the right-hand side of (A3) is zero.

Consider now a variation in k. Taking the derivative of (A2) with respect to k, we get

$$\begin{aligned} \frac{d\hat{r}}{dk} &= -\frac{d\alpha}{dk}(F(s\ b) - F(s\ b + k - w)) \\ &\quad - \alpha(k, s)\ \frac{\partial}{\partial k}\ (F(s\ b) - F(s\ b + k - w)) \\ &= -\frac{d\alpha}{dk}(F(s\ b) - F(s\ b + k - w)) < 0 \end{aligned}$$

where the second equality follows from Lemma 2 and the inequality from the assumption that $d\alpha/dk > 0$. The case of a variation in s is a bit trickier, given that we have two effects with opposite signs. Taking the derivative of (A2) with respect to s, we get

$$\begin{aligned} \frac{d\hat{r}}{ds} &= b\ f(s\ b) - \alpha\ \frac{\partial}{\partial s}\ (F(s\ b) - F(s\ b + k - w)) \\ &\quad - \frac{d\alpha}{ds}(F(s\ b) - F(s\ b + k - w)) \\ &= b\ f(s\ b) - \frac{d\alpha}{ds}(F(s\ b) - F(s\ b + k - w)) \end{aligned}$$

The second equality follows from Lemma 2. Since F is strictly increasing and Lemma 1 implies that $w > k$, it follows that $F(s\ b) > F(s\ b + k - w)$. Finally, if b is sufficiently small and $d\alpha/ds$ is sufficiently high, then $d\hat{r}/ds < 0$. □