C-50

Center for Research on Economic and Social Theory
Research Seminar in Quantitative Economics

Discussion Paper

DEPARTMENT OF ECONOMICS

University of Michigan

Ann Arbor, Michigan 48109

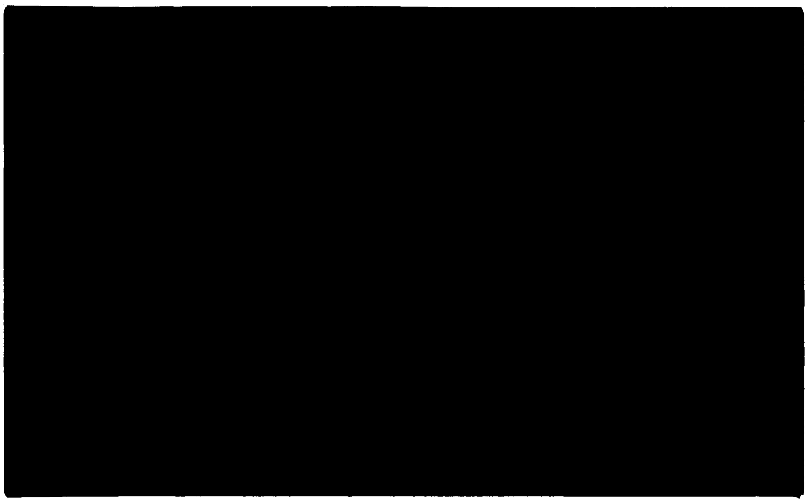

On Settling Many Suits

Lawrence Blume

C-50

25 April 1983

Department of Economics, University of Michigan

This research was begun while the author was a Visiting Fellow at Harvard University's Kennedy School of Government, where David Lax (HBS) and James Sebenius supplied needed advice and encouragement

Preliminary Draft

1. Introduction

Bargaining situations frequently arise where at least one party must take account not only of the immediate negotiation confronting it, but also of its impact on the outcome of subsequent negotiations with other parties. Labor unions may negotiate contracts with several different firms. Contracts to provide certain services may be negotiated sequentially with several different purchasers. These bargaining situations arise frequently in the law, where the tortious actions of a single tortfeasor may have caused injury to several parties. The reckless driving of an individual may have injured not one but several different parties. A particular corporate strategy in violation of the antitrust laws may have injured not one but several competitors, all of whom to decide to litigate.

Two problems arise in repeated bargaining situations that are not present in "one-time" bargaining problems. First, in repeated bargaining problems parties may earn reputations for being easy or tough bargainers. This kind of situation gives rise to a version of the chain store paradox, and can be modeled as a game of incomplete information where there is uncertainty about the payoffs of various strategies to one player. These models have been extensively discussed in recent years.[1] In this context we can think of the reputation issue as one player's attempt to manipulate the information other players receive about him to his own advantage. The second problem is similar. One player may wish to manipulate the information both he and other players receive about exogenous payoff-relevant events to his own advantage. The problem again is one of manipulating information which is valuable to players of the game, but here, and unlike the reputation problem, the information is not asymmetrically observed. As in the incomplete information analysis of reputation problems, here I wish to study how the opportunity to control exogenous payoff-relevant information affects play.

I like to think of this problem in the legal context. A firm engages in an activity which is potentially anticompetitive and in violation of the antitrust laws. The firm's competitors bring a series of suits against the firm. The firm, when considering the first suit against it, has two alternatives. The firm can settle the suit by paying a mutually agreed upon sum of money to the plaintiff or it can litigate. If it litigates, it either wins or loses, and this outcome is not known in advance to either side. Other plaintiffs will revise their beliefs about the merits of their case--the possibility of success for their suit--in light of this first outcome. If bringing suit is costly for the plaintiffs, then an adverse decision for the first plaintiff may convince subsequent plaintiffs not to sue. Similarly, a favorable decision for the first plaintiff may convince other injured parties to litigate, and it will increase the minimum settlement they would agree to. In both cases the random and unknown outcome of litigation will affect the subsequent losses of the defendant, and so he must take this into account in deciding whether to litigate or settle the first suit.[2] Each plaintiff brings only one suit, but the defendant must respond to all suits brought by all plaintiffs.

In this setting, information about the probability of successful litigation is revealed to each litigant every time there is a trial. At any point in time, the defendant can control the flow of information to subsequent plaintiffs by settling suits so that no subsequent observations on the judicial process are observed. The effect of the defendant's ability to control information revelation can be measured by comparing the incentives to settle in the "one-stage" game (the game with only one plaintiff) with the incentives to settle in games with many plaintiffs.

The conclusions for the simple model studied here are strong. The incentives for the defendant to settle are less in the many plaintiff game than in the single plaintiff game. In other words, there are suits that the defendant would settle against only only one plaintiff that he would not settle were he contemplating actions by subsequent plaintiffs. This result clearly has implications for litigation strategy by plaintiffs. For example, class action suits allow the bundling of many suits by individual plaintiffs into one suit brought on behalf of all plaintiffs. The advantages of this and other bundling devices should be examined in light of the effect on defendants' expected return from bringing suit of changing favorably the defendant's incentive to settle.

This and other topics will be taken up in the concluding section of this paper. The next section lists all notation to be used in the sequel, and the formal model is presented in section 3. Results for the one plaintiff model are presented in section 4, and for the many plaintiff model in section 5.

2. Notation

The following notation will be used in the description of the model and subsequent analysis in sections 3 and 4.

a	amount of settlement
g	gain to defendant in the event of no suit
h	plaintiff's loss
l_s	plaintiff's litigation costs in the event of a settlement
l_t	plaintiff's litigation costs in the event of a trial
l	$l_t - l_s$
m_s	defendant's litigation costs in the event of a settlement
m_t	defendant's litigation costs in the event of a trial
m	$m_t - m_s$
p	plaintiff's predicted probability of an award in favor of the plaintiff
q	defendant's predicted probability of an award in favor of the plaintiff
r	plaintiff's prior probability of θ_L
s	defendant's prior probability of θ_L
θ_L	frequency of awards in favor of the plaintiff in the 'low' state
θ_H	frequency of awards in favor of the plaintiff in the 'high' state
u_s	expected value to the plaintiff of a settlement

u_t expected value to the plaintiff of a trial

v_s expected value to the defendant of a settlement

v_t expected value to the defendant of a trial

3. The Model

A risk neutral firm, called player 0, undertakes some action which injures another risk neutral firm, called player 1. This action has generated a gain for player 0 of g dollars, and caused player 1 h dollars of harm. Player 1 contemplates bringing suit against player 0 for the full amount of the damages. If a suit is brought against player 0, three outcomes are possible: a settlement, litigation with a victory for the plaintiff (player 1), and litigation with a victory for the defendant (player 0). First suppose the suit is settled for dollar amount a. Reaching this agreement is costly for both parties. The cost is l_s to the defendant and m_s to the plaintiff. Thus the net benefits of settling for amount a are

$$u_s = g - l_s - a$$

for the defendant and

$$v_s = a - h - m_s$$

for the plaintiff. Suppose next that the suit is litigated.

There are only two possible outcomes of the litigation process. There can be a finding in favor of the plaintiff, in which case the defendant must pay the plaintiff the entire amount h of the claimed damages, or there can be a finding in favor of the defendant, in which case the defendant pays nothing and the plaintiff receives nothing. Litigation is costly for both sides in the suit, with costs l_t for the defendant and m_t for the plaintiff. In the event of a victory for the plaintiff, the net gain to the defendant is $g-I_t-h$ and the net gain to the plaintiff is $-m_t$. In the event of a victory for the defendant, the defendant's net gain is $g-I_t$ and the plaintiff's net gain is $-h-m_t$. These payoff possibilities are summarized in table 1.

Table 1

Payoffs to Litigation Strategies

	defendant	Plaintiff
No Suit:	g	$-h$
Suit:		
Settle:	$g-I_s-a$	$a-h-m_s$
Litigate:		
Pl. Win:	$g-I_t-h$	$-m_t$
Pl. Lose:	$g-I_t$	$-h-m_t$

The outcome of the judicial process is random. If the suit is actually tried, the outcome--win or lose--is random. Furthermore, neither plaintiff nor defendant knows the true probability of plaintiff's success. (The words "win" and "success" will always refer to a victory by the plaintiff, and "lose" will refer to a loss for the plaintiff.) There are two possibilities: success probability θ_L and θ_H, with the latter greater than the former. The plaintiff places prior probability r on θ_L being the true process, and the defendant assigns it prior probability s. The resulting predicted probabilities of success for the plaintiff and defendant are, respectively:

$$p = r\theta_L + (1-r)\theta_H$$

and

$$q = s\theta_L + (1-s)\theta_H$$

I interpret the judicial process in the following way. The

probability of success depends upon the merits of the case. If the plaintiff has a "good" case, the success probability in court is θ_H. If he has a "bad" case, the probability of success is θ_L. Both plaintiff and defendant are uncertain about the merits of the case, and thus about whether the correct success probability is θ_H or θ_L.

The game tree for this game is pictured in figure 1. The opening random move captures the players' uncertainty about the parameter describing the judicial process--the merits of the case. The settlement subgame is not modeled explicitly. Only its outcomes need be noted. The possible outcomes are a(r,s) or litigation. Note that the settlement rule is assumed to depend only upon the two players' predicted probabilities of success.

Of course this is a very stylized model of the legal process. There is no model of the litigation technology, and a connection between litigation costs and the probability of a

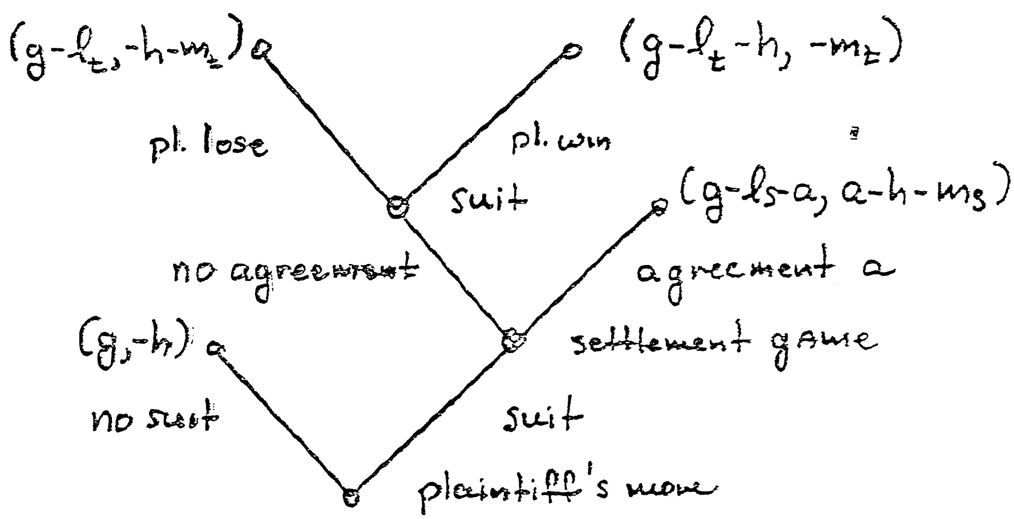

$(g-l_t, -h-m_t)$
$(g-l_t-h, -m_t)$
pl. lose
pl. win
suit
$(g-l_s-a, a-h-m_s)$
no agreement
agreement a
$(g, -h)$
settlement game
no suit
suit
plaintiff's move

fig. 1.

successful suit. The distribution of awards in the litigation process is oversimplified. Most important, there is no model of how the settlement value a gets determined. This is itself presumably the outcome of some non-cooperative game.[3] Presumably, the outcome of the settlement process must be a dollar amount that is individually rational. This is to say, the amount a must exceed the plaintiff's expected gain of litigation and not exceed the defendant's expected loss from litigation. This range is the object of study, and more will be said about it in the next section.

So far only the game with one plaintiff, player 1, has been described. I want to compare the one-plaintiff game with a game containing a sequence of plaintiffs. In this game plaintiffs are players 1 through T, identical in all respects save the information which they have at the time of their opportunity to sue. The plaintiffs are considered sequentially by the defendant. When each plaintiff's turn comes up, the game is exactly that which was just described. At stage t it is plaintiff t's turn to decide whether or not to sue, and if he sues the possible outcomes are litigation or settlement. However, at stage t both the defendant and plaintiff t have had the opportunity to learn about the judicial process from the experience of previous litigation. I assume that the relevant case law is well established and known to all parties, so that there is no room for precedent. Given the underlying merits of the cases, the only uncertainties are due to the random

T Plaintiff Model

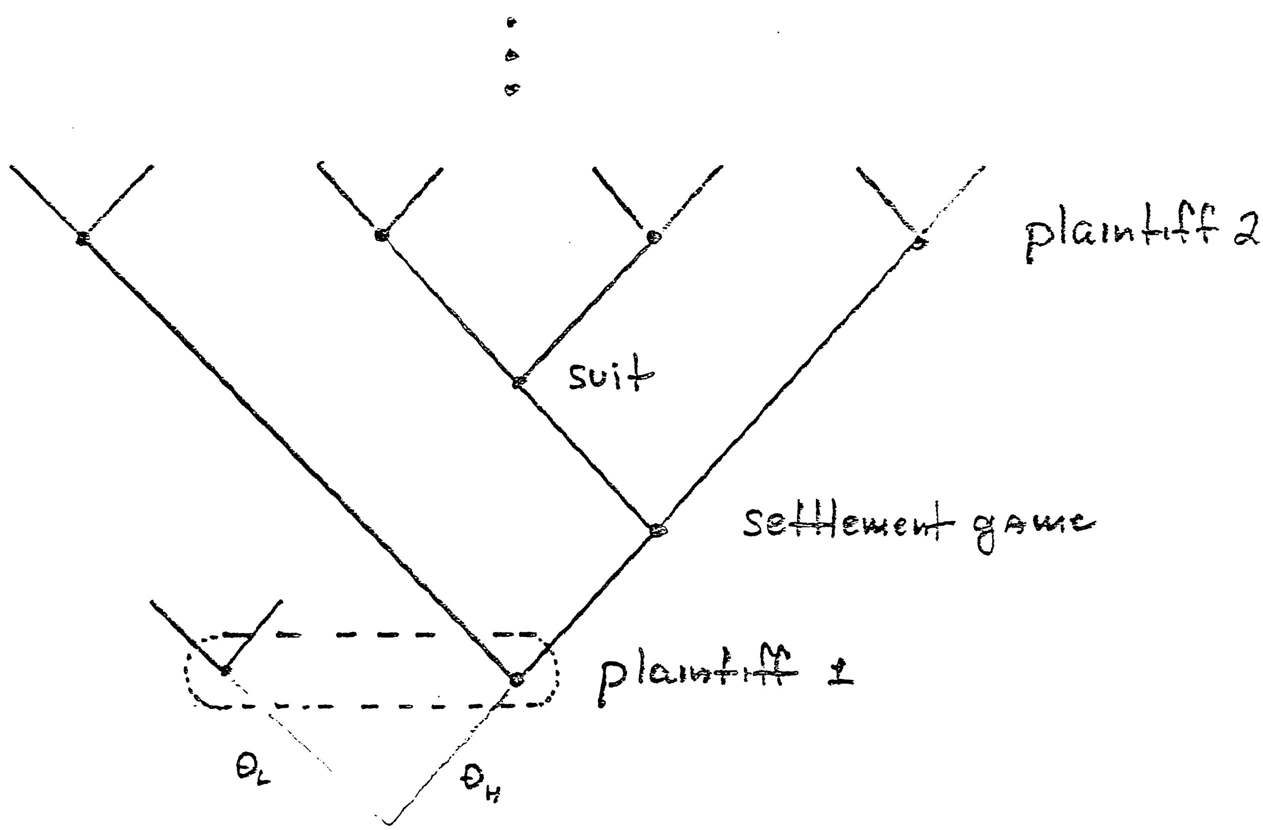

fig. 2

behavior of juries, or the qualities of the presentation of the cases, etc. Furthermore, we explicitly ignore (for the moment) the effect of procedural devices, such as collateral estoppel, which create a causal link between the outcome of a current case and the disposition of previous cases. Thus the outcomes of the judicial process are, in this case, assumed to be i.i.d. Now the defendant must be concerned not only with the expected outcome of the current suit, but its effects upon the incentives of subsequent players to sue. Let r_t, p_t, s_t, and q_t denote the respective posterior and predicted probabilities of the defendant and plaintiff t at the beginning of round t. This is to say, r_t and s_t are the posterior probabilities of θ_L after t-1 rounds of play (which may include less than t-1 observations of the judicial process if some of the previous suits have been settled or not been brought. The game tree for this game (rather, a representative part of it) is pictured in figure 2.

4. Analysis of the One Plaintiff Model

Consider first the one-plaintiff model. In this case there is only one stage, and the defendant need not consider the effects of information revelation. Define the expected values of litigation

$$u_t = q(g-l_t-h)+(1-q)(g-l_t)$$
$$= g-l_t-qh$$

for the defendant and

$$v_t = -pm_t-(1-p)(m_t+h)$$
$$= -m_t-(1-p)h$$

for the plaintiff.

The outcome of the settlement subgame is an amount $a=a(r,s)$. This settlement is feasible only if both the defendant and the plaintiff prefer the settlement to litigation--either player can otherwise refuse to accept an offered settlement. For the plaintiff, this is only possible if

$$v_s \geq v_t$$
$$a-h-m_s \geq -m_t-(1-p)h$$
$$a \geq ph+m_s-m_t.$$

Similarly for the defendant,

$$u_s \geq u_t$$
$$l_t-l_s+qh \geq a.$$

Thus the range of possible settlements is

$$l_t - l_s + qh \geq \quad a \geq ph - m_t + m_s$$

A settlement is only possible, then, if

(1) $\quad l + m \geq (p-q)h \quad .$

Suppose, then, that (1) is satisfied. Then the suit will be settled, and the value to the plaintiff of bringing the suit is v_s. Thus the plaintiff will bring the suit only if

$$v_s \geq -h$$

and this occurs if and only if

$$a \geq m_s .$$

This is to say, the amount of the settlement must exceed the cost of reaching it.

Consider first the extreme case where the settlement value is as low as possible. This corresponds to a settlement mechanism where after the suit is filed the defendant makes a "take it or leave it" offer, if he makes any offer at all. The cost minimizing acceptable offer for the defendant is the lowest acceptable offer to the plaintiff. In this case

$$a = ph - m .$$

This settlement offer gives the plaintiff a return equal to the expected return from litigation. With an offer this low, the plaintiff is indifferent between litigating and settling, and his return from doing either is:

$$-(1-p)h-m_t.$$

The plaintiff will not bring suit unless this exceeds -h, and so in this case, no suit will be brought by the plaintiff unless

$$ph \geq m_t.$$

Note that if there is no cost to bringing suit, the plaintiff will always sue because the expected court award, ph, is positive.

The plaintiff's decision whether or not to bring suit clearly depends upon his subjective beliefs about the probability of success, and several cases are possible. Recall that the judicial process is assumed to decide for the plaintiff with probability either θ_L or θ_H, with the latter greater than the former. The worst case for the plaintiff is when he knows with certainty that the true process is described by θ_L. This is to say, his prior r is 1. The best case is when r=0 and he knows for sure that the judicial process is described by parameter value θ_H. It is possible to imagine configurations of h, m_t, θ_L and θ_H such that the plaintiff will never sue, or always sue. These cases occur, respectively, when

$$\theta_H < m_t/h$$

and

$$\theta_L > m_t/h.$$

For the remainder of the paper the most interesting case will be assumed:

$$\theta_L < m_t/h < \theta_H.$$

In this case, we see that if r is sufficiently high, no suit will be brought. The plaintiff will sue if and only if

(2) $$r < (\theta_H h - m_t)/(\theta_H - \theta_L),$$

and our assumption implies that the right hand side is between zero and one.

Now consider the other extreme, where the settlement is the highest amount that the defendant is willing to pay. This corresponds to a settlement mechanism where the plaintiff makes a "take it or leave it" offer. In this case the only rational offer is for the highest amount that the plaintiff is willing to pay. If this offer is less than the plaintiff's expected return from litigation, then the suit will be settled rather than litigated. The most that the defendant will be willing to pay is the amount that leaves him indifferent between settling and litigation. A calculation shows that this amount is

$$a = qh+l.$$

The suit will be settled only so long as

$$l+m > (p-q)h.$$

The plaintiff will sue only if

$$\max\{l+qh, -m+ph\} > m_s.$$

This condition can be reduced to a condition on the prior probability assessments of the plaintiff and defendant. Clearly the plaintiff is no less likely to settle in this case then at the other extreme, and the defendant is no more likely to settle in this case then as before.

5. Analysis of the Many Plaintiff Model

In this section the decisions to settle and to litigate are considered in the context of not one but many plaintiffs. Suppose that there are T plaintiffs, ordered 1 through T. The plaintiffs are considered sequentially. First plaintiff 1 decides whether or not to sue. If he brings suit, the the suit is either settled or litigated, and the outcome is observed by the defendant and every subsequent plaintiff. Then it is plaintiff 2's turn, and so forth. Analysis of non-zero-sum repeated games of incomplete information such as this can be very complicated because it is frequently the case that there are either no equilibria or very many equilibria, and in the latter case many of the equilibria are implausible. Perfection and properness represent two attempts to refine the set of Nash equilibria of any game to those equilibria satisfying certain plausibility criteria. Here I use an equilibrium concept called sequential equilibrium.[4] This concept is analogous to Selten's concept of perfect equilibrium, but is more natural for games of incomplete information. The concept of subgame perfection requires of equilibrium strategies not only that they be an equilibrium for the whole game but that they also be an equilibrium for any subgame of the original game. Consider the branching of the game tree from any node of the extensive game. This constitutes a subgame, and the strategies--rules which assign to each player a play at each node where it is his turn to play--must also describe an equilibrium for this game as well. The notion of sequential equilibrium goes one step further by considering not just nodes of the game tree but information sets as well. Whenever a player has the move, he also has subjective beliefs about where within his information set the game is actually being played. Sequential equilibrium requires optimal play from each information set (and so it includes the subgame perfection concept), and it also requires that the probability assessments each player assigns to the nodes within his current information set be consistent with previous play.

Equilibrium strategies for the plaintiff are easy to describe. When it is plaintiff t's turn to sue, he can look at the previous history of judicial outcomes and revise his beliefs about the merits of his case using Bayes rule. He will then compute his expectations for suit and settlement, and then decide whether or not to bring suit. Given that he brings suit, he will accept the settlement award a(r,s) only if it exceeds his expected value of litigation. The nature of the equilibrium for the T plaintiff game can be seen from the following proposition.

<u>Proposition 1.</u> If it is optimal for the defendant to settle with plaintiff t, then it is also optimal for him to settle with all plaintiffs t' > t.

This proposition is true because first, the judicial process is stationary--which is to say that the outcome of litigation is independent of the name of the plaintiff--and second, because settling gives no information to either plaintiffs or defendant.

Proof of proposition 1: It suffices to show that, in the T plaintiff game, the defendant will litigate plaintiff T-1 if he litigates plaintiff T. The extension back to previous plaintiffs then follows from the standard backward induction arguments familiar from dynamic programming theory. Suppose that after the completion of the T-2'nd round the defendant has posterior s and the two remaining plaintiffs have posterior r. Suppose that it was optimal for the defendant to settle with plaintiff T-1 and to litigate with plaintiff T. First note that if the case with plaintiff T-1 is settled, no new information is introduced, and so beliefs about the judicial process remain the same. Since the settlement award depends only on beliefs of the two parties involved in the case, the settlement with plaintiff T-1 and the settlement with plaintiff T are identical. The hypothesis is that the expected value of litigation with plaintiff T given the information available after plaintiff T-2 exceeds the value to the defendant of settling with plaintiff T. Therefore it exceeds the expected value of settling with plaintiff T-1.

$$u_s(T-1) = u_s(T),$$

$$E\{u_t(T):T-2\} > u_s(T),$$

and so

$$E\{u_t(T):T-2\} > u_s(T-1),$$

where $u_s(k)$ is the immediate value to the defendant from settling with plaintiff k assuming both T and T-1 are settled, and $E\{u_t(T):T-2\}$ is the expected value to the defendant from settling with plaintiff T given the information available after T-2. Since the judicial process is a martingale, it is easy to see that

$$E\{u_t(T):T-2\} = E\{u_t(T-1):T-2\},$$

and so

$$E\{u_t(T-1):T-2\} > u_s(T-1).$$

Hence the immediate reward to the defendant from litigating the case with plaintiff T-1 exceeds the immediate reward from settling the case. Hence, if it was optimal for the defendant to settle with T-1 and litigate T, it must be the case that the expected value to the defendant of the optimal action against plaintiff T conditioned on the information available after T-2 and assuming that T-1 is litigated is less than $E\{u_t(T):T-2\}$. But this is not the case because

$$\begin{aligned} &E\{E\{\text{return from optimal action}(T):T-1\}:T-2\} \\ &\geq E\{E\{u_t(T):T-1\}:T-2\} \\ &= E\{u_t(T):T-2\}. \end{aligned}$$

Thus the hypothesis that the strategy of settling T-1 and litigating T is contradicted.

Q.E.D.

Proposition 1 describes the nature of equilibrium strategies for the defendant. The defendant's equilibrium strategy is a stopping rule which tells him when to stop litigating. Once he first settles, he will continue to settle for the remainder of the game. The existence of equilibrium, then, depends upon the existence of an optimal stopping rule for a particular sequence of random variables.

The judicial process gives a sequence of 0-1 valued random variables, corresponding to victories for the plaintiff and defendant, respectively. The probability of observing a success at any given time is independent of previous outcomes, and is given by the parameter (which is not known to any litigant). A stopping rule for this process is a random variable with values in the set of positive integers {1,2,...} and such that the event {t = n} depends solely on the outcomes observed by plaintiffs and the defendant up to time n. In the T plaintiff game stopping rules are rules which tell the defendant when to stop litigating.

The stochastic process being stopped is easy to describe. I view the defendant as stopping the process of settlements (described below). The cost of sampling from this process is the random cost of litigation. Suppose there are T plaintiffs, and suppose the first k plaintiffs have sued the defendant and that he has litigated rather than settled. What should he do in subsequent suits? After litigation of suits 1 through k the defendant's posterior is s_{k+1} and the plaintiffs' common prior is r_{k+1}. Should the defendant settle all subsequent suits, his payoff is

$$x_k = (T-k)(g-\max\{a(r_t, s_t)-m_s), 0\}-l_t),$$

the total return from settling all subsequent suits (and taking account of the prospect that no subsequent plaintiff will choose to bring suit). Should he choose to sample again--to observe another outcome of the judicial process--the cost is random. The return to the defendant from sampling is:

$$c_k = \begin{cases} g & \text{if the plaintiff does not sue} \\ g-l_t & \text{if the plaintiff loses} \\ g-l_t-h & \text{if the plaintiff wins.} \end{cases}$$

The three states of c_k can be defined in terms of inequalities for r_{k+1} and s_{k+1}, and their occurances can be described in terms of the success probability of the judicial process.

Define:

$$y_k = x_k - c_1 - \ldots - c_k.$$

The T plaintiff game has a sequential (Nash) equilibrium if there is an optimal stopping rule for this sequence of random variables.

Theorem 1. All T plaintiff games with bounded (measurable) settlement rule $a(r,s)$ have a sequential (Nash) equilibrium.

Proof of Theorem 1: We have already computed each plaintiff's optimal action given the information available when it is his turn to bring suit, and we have seen that the defendant's optimal response to these strategies--the optimal way to manipulate information available to the plaintiffs--is a stopping rule detailing when litigation should stop. The payoffs from stopping after plaintiff k are described by the stochastic process $\{y_k\}$. An equilibrium for the T plaintiff game with settlement rule $a(r,s)$ exists if and only if there exists an optimal stopping rule for the payoff process $\{y_k\}$. An optimal stopping rule for this process does indeed exist. See Chow, Robbins and Siegmund Theorem 3.2.

Q.E.D.

Comparisons between the T plaintiff game and the one plaintiff game are facilitated by the characterization of optimal strategies found in proposition 1. The main result in this direction is the observation that the incentives for the plaintiff to settle are reduced when there are many plaintiffs. Any suit settled in the T plaintiff game will also be settled in the one plaintiff game. At first this result seems surprising. One can imagine two possible stories for this model. In the first, the plaintiff would be more likely to settle in contemplation of the enormous costs involved should he lose and should many plaintiffs then be induced to bring suit. In the second, the plaintiff is eager to litigate because if he wins many plaintiffs may decide not to bring suit. It is this second story that is correct, and it is easily seen to be correct even when the defendant is allowed to exhibit aversion to risk.

This result depends upon two assumptions about the judicial process and the behavior of the plaintiffs. First, the judicial process is stationary so that the average expected outcome for many trials equals the expected outcome of one trial. Second, it is costly to bring suit, and so if the possibility of success for the plaintiffs is sufficiently remote, they will choose not to bring suit. Suppose, for example, that our model were specified so that in the one plaintiff game the plaintiff would sue were he to know that the parameter value describing the merits of the case, θ, was θ_H, but not if it was θ_L. Suppose the plaintiff's prior was such that he chooses to bring suit. Let's compare the expected return to the defendant from litigating in the one plaintiff game with the return from continual litigation in the T plaintiff game, assuming that the parameter value governing the judicial process is in fact θ_H. If litigation was to proceed against every single plaintiff, the expected return per plaintiff would equal the expected return from litigation in the one plaintiff game. But if the defendant were to litigate every single suit, it is possible that a string of losses (for the plaintiffs) would occur that is sufficiently long to make the plaintiffs think that in all likelihood the merits of their case is described not by θ_H but by θ_L. Were this to occur they would cease to bring suit, and the defendant will be even better off than if he had won the suit, since he need not pay litigation costs l_t. This possibility of driving plaintiffs out of court increases the value of litigation to the defendant in the T plaintiff game, and so the zone of feasible settlement decreases. It will be clear that the probability of driving out, say, the last half of the potential plaintiffs, increases with T, and so larger values of T make settlement less likely.

Theorem 2.

For each specification of parameters regarding damages, litigation and settlement costs, prior beliefs and specification of the judicial process, and for each bounded (measurable) settlement rule a(r,s), any suit settled in the T plaintiff game will also be settled in the one plaintiff game.

Proof of theorem 2. Let ${}_ku$ denote the expected return to the defendant from litigating the first suit and optimal play in subsequent suits. Let ${}_ku_t$ denote the expected return to the defendant from litigating every suit in the k plaintiff game. Let ${}_ku_s$ denote the return to the defendant from settling every suit in the k plaintiff game. Let S denote the set of all sample paths of the judicial process (of length T) and for k < T let S_k denote those sample paths that lead the plaintiffs to cease litigation after the k'th trial. Let $S_T = S/\bigcup_{k=1}^{t=1} S_k$. Finally, let a denote the value of settling the first suit. Denote by $\{v_k\}$ the stochastic process of awards to plaintiffs. This is to say, $v_k = h$ if the k'th trial is won by a plaintiff, and 0 otherwise. Computing average returns per plaintiff:

$$
\begin{aligned}
T^{-1}{}_T u &\geq T^{-1}{}_T u_t \\
&= T^{-1}\Sigma_{k=1}^{t=1} \Pr\{S_k\} E\{Tg - kl_t - \Sigma_{j=1}^{k} v_j\} \\
&\geq T_{-1}\Sigma_{k=1}^{t=1} \Pr\{S_k\} E\{Tg - Tl_t - \Sigma_{j=1}^{k} v_j\} \\
&= {}_1u_t.
\end{aligned}
$$

Proposition 1 states that if it is optimal to settle the first suit in the T plaintiff game, it is optimal to settle every suit in the k plaintiff game, and the average return per plaintiff is the same as the return from settling in the one plaintiff game. Suppose, then, that settlement is the optimal strategy in the T plaintiff game. Then

$$
\begin{aligned}
{}_1u_s &= T^{-1}{}_T u_s \\
&\geq T^{-1}{}_T u \\
&\geq {}_1u_t,
\end{aligned}
$$

and so it is optimal to settle in the one plaintiff game.

Q.E.D.

If T is sufficiently large, one of the S_t, $t < T$, will occur with positive probability and so settling will strictly dominate litigation in the one plaintiff game. Alternatively, for any settlement rule there are parameters θ_L, θ_H, prior beliefs r and s, and settlement and litigation costs such that for T sufficiently large, suits will be litigated in the T plaintiff game that would be settled in the one plaintiff game. For an extreme example, consider the settlement rule discussed in section 4 where the plaintiff makes a "take it or leave it" offer to the defendant. This is the case where

$$a = qh+l.$$

In the one plaintiff game the defendant is indifferent

between settling and litigating any suit. However, if T is sufficiently large, then so long as $r < 1$, the predicted probability of some S_t is strictly positive, and so no suit will be settled. This example is particularly interesting because it answers a question about the modelling strategy adopted for the treatment of settlements. Consider the admissable range of settlement offers. The amount agreed upon (net of litigation costs), if it is to be feasible, must exceed the expected value of litigation for the plaintiff and be bounded above the expected value of litigation for the defendant. This last concept is easy to define in the one plaintiff game but it is hard to get a handle on when there are many plaintiffs. In this case we have to consider the value to the defendants of the information revealed by the litigation process. This really involves studying recursion relationships defining the value function of the dynamic programming problem faced by the defendant. Theorem 2 states that the expected value to the defendant of the information revealed by litigation is positive, and the example shows that the one-plaintiff "take it or leave it" offer bounds from above the set of feasible offers.

6. Conclusion

In a simple model of repeated litigation it appears that the presence of many potential plaintiffs for a given cause of action decreases the incentives for the defendant to settle a case. This result depends upon the characterization of the defendant's optimal strategy as a stopping rule. The stopping rule property is very robust. For example, it continues to hold under a variety of rules for distributing the expense of litigation between the two parties involved in any trial, and under more general judicial processes where the damages awarded by the court are random and possibly do not equal h, the damages actually incurred by the plaintiffs.

It is interesting to use the analysis of the previous sections to compare the incentives to settle under different procedural rules. A natural question to ask in light of the results on sequential litigation is what happens if the plaintiffs are allowed to consolidate their individual suits into one big suit. Mechanisms for consolidating suits do exist--perhaps the most well-known device for achieving this goal is the class action suit, but other mechanisms also can reach this end.

To be specific, suppose that the settlement rule is linear in damages, and suppose that the technologies for litigating and settling exhibit constant returns to scale. This is to say, in considering any suit, whether it be an individual suit or the consolidated suit, the settlement rule is $a(r,s) = b(r,s)D$, where D represents the reward at stake--h for an individual suit and Th for the consolidated suit, and settlement and litigation costs in the consolidated suit are T times those in the individual suits. In this case it is straightforward to show that if the consolidated suit will be litigated, then so will the first of the sequential suits. The reasoning is similar to that employed in the proof of theorem 2. The return from litigating all of the sequential suits exceeds the return from litigating the consolidated suit because of the possibility of driving plaintiffs out of the courts due to a run of adverse decisions. This truncates the distribution of rewards to the plaintiffs and also economizes on legal costs. The return from settling all of the sequential suits equals the return from settling the consolidated suit. If it is ever optimal to litigate the sequential suits, it is optimal to do so at the outset. These statements can be put together exactly as they were in the proof of theorem 2 to prove the result.

The conclusion here is that class action suits (and similar procedural devices) are more likely to be settled than are a sequence of individual suits. This is not to say that the plaintiffs are better off with a class action suit, because at least some of them are foregoing payoff relevant information that would otherwise be obtainable from observed from previous suits. So far I have been unable to discern which situation is better for the plaintiffs.

Other procedural rules can be evaluated as well. In a subsequent paper I will use a variant of the model presented here to discuss collateral estoppel.

Proposition 1 and theorems 1 and 2 make use of the fact that the defendant is risk neutral. It is interesting to see how far we can go when we allow for risk aversion. How far we can go depends on how the reward process is modeled. If we assume that utility is additively separable over time (over plaintiffs) and strictly concave, it is easy to see that all the results go through just as before. If we assume that the defendant assigns utility to sums of all outcomes then immediatly things are harder because proposition 1 need no longer be true. A defendant may prefer to settle initial cases in contemplation of horrendous losses from subsequent plaintiffs, but he may be willing to risk litigation once the pool of potential plaintiffs becomes sufficiently small. However there is still a result of sorts about consolidating cases. It is the case that settling a consolidated case gives the same return as settling all of the sequential cases. It is also the case that the return from litigating all of the sequential cases exceeds the return from litigating the consolidated case because it the distribution of net payouts by the defendant (including for litigation costs) can be shown to have a lower expected value and be less risky in the sequential game. Thus, so long as the defendant's utility function does not exhibit increasing absolute risk aversion, it will be the case that if he litigates the consolidated case, he will also litigate at least one of the sequential cases. However, since the defendant's optimal strategy is no longer a stopping rule, he may not litigate the first case.

The models discussed in this paper are applicable not only to bargaining in the shadow of repeated litigation but also to other bargaining situations. Labor arbitration presents another example of a similar process. The general question addressed in this paper also arises whenever state contingent contracts must be written to cover some long period of time, but contract writing is costly. In this setting each party must decide whether to bear the cost of specifing the contract for a particular state of nature or to rely on some external process to resolve any difficulties should this state arise. A model for this question will be more complicated than the model discussed here because of the bilateral nature of the contracting process, but it would be rich in applications. The issue of information revelation raised in this paper will be important in the contract setting for determining the optimal coverage and run of contracts--both important problems in contemporary contracts theory.

REFERENCES

Chow, Y.S., Herbert Robbins and David Siegmund, Great Expectations: The Theory of Optimal Stopping. Boston, 1971.

Cooter, Robert, Stephen Marks and Robert Mnookin, "Bargaining in the Shadow of the Law: A Testable Model of Strategic Behavior," Journal of Legal Studies, 11 (1982), 225-252.

Kreps, David and Robert Wilson, "Sequential Equilibria," Econometrica, 50 (1982), 863-94.

Kreps, David and Robert Wilson, "Reputation and Imperfect Information," Journal of Economic Theory, 27 (1982), 253-279.

Milgrom, Paul and John Roberts, "Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis," Econometrica, 50 (1982), 443-460.

McClennen, Andrew, "A Noncooperative Definition of Two Person Bargaining," University of Toronto Working Paper No. 8303, August 1982.

Rubenstein, Ariel, "Perfect Equilibrium in a Bargaining Model," Econometrica, 50 (1982), 97-109.

----

1. See Kreps and Wilson, and Milgrom and Roberts.

2. Through some mechanism of judicial precedent a previous decision may increase the probability of a subsequent favorable decision. Also, the parameters of the judicial process may be unknown, and so learning behavior must be considered.

3. For examples of formal models of non-cooperative bargaining games see, among others, papers by McLennen and Rubinstein. In an explicitly legal setting, see Cooter et. al.

4. See Kreps and Wilson, Econometrica